
Morpheus Documentation

Morpheus

Dec 11, 2020

MORPHEUS UI

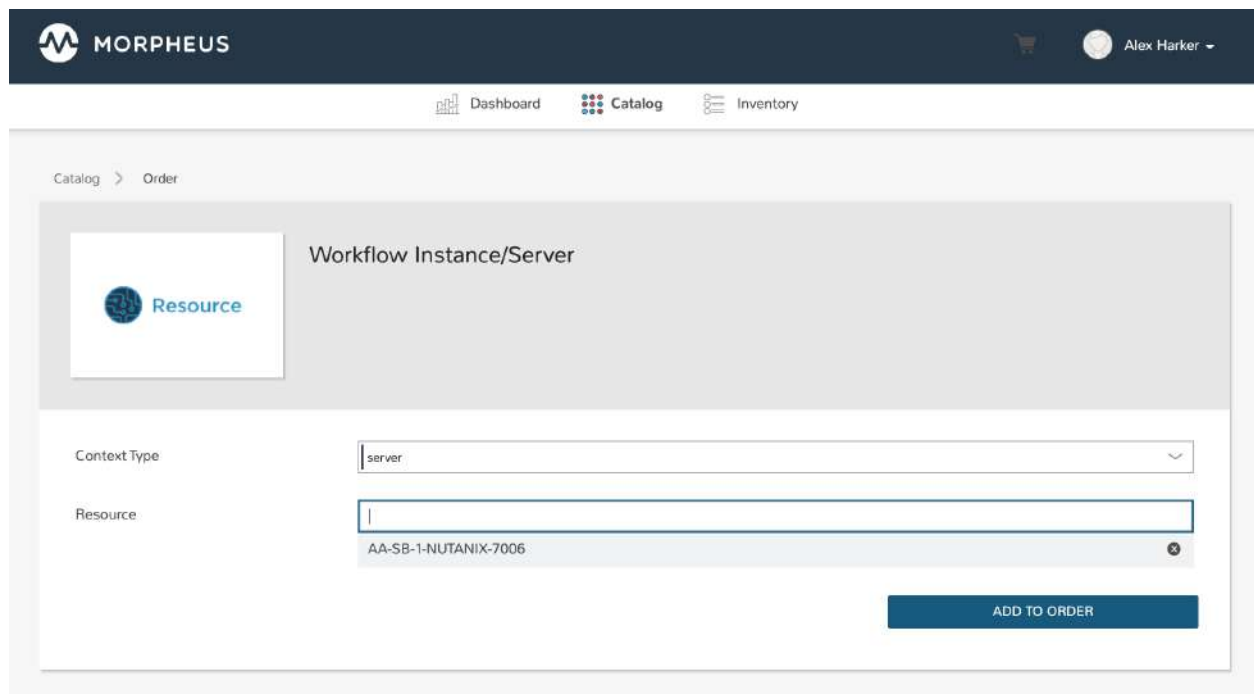
1	v5.2.0 Highlights	1
1.1	Service Catalog Persona Improvements	1
1.2	ServiceNow Integration Improvements	2
1.3	Hide Blueprint fields	2

V5.2.0 HIGHLIGHTS

1.1 Service Catalog Persona Improvements

The Morpheus version 5.0.0 beta introduced Personas, which are a new approach for optimizing and simplifying self-service for targeted audiences. The first Persona to ship is [Service Catalog](#), which sees additional improvements in the 5.2.0 LTS release.

- Make Morpheus Operational Workflows available for order from the Service Catalog and run them against selected targets
- With added API/CLI support, work with Personas, create and manage Catalog Items, and make selections from the catalog through Morpheus API and CLI tools
- Inventory list view now includes much greater detail about each inventory item
- Categorize items under selected headers for enhanced discoverability as the catalog grows
- With the added quantity selector, order additional copies of items in your cart without creating duplicate selections



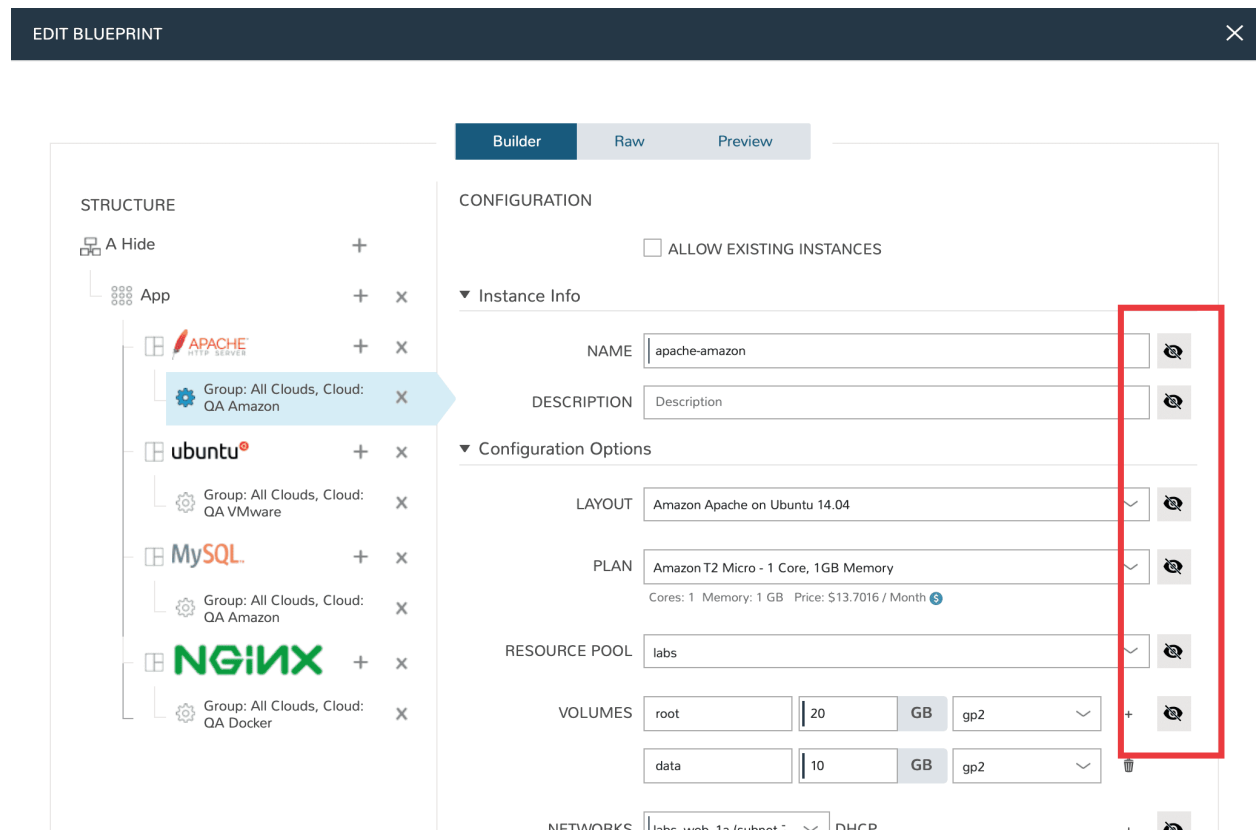
1.2 ServiceNow Integration Improvements

Morpheus 5.2.0 brings improvements to the ServiceNow integration, including upgrades to incident surfacing, integration with service catalog items, and more.

- “Morpheus Incident” alerts are now more insightful including links to the related Morpheus incident or check, severity information, and other details about the failing check
- Provision Service Catalog Items through the Morpheus ServiceNow plugin
- Added the capability to identify a MID server once on the properties page rather than setting it individually for each call
- Pricing data is now available to the ServiceNow plugin when ordering Service Catalog items. This is made available on the XML as a monthly price, users would have to modify the form UI to surface this information

1.3 Hide Blueprint fields

Administrators have been able to lock Blueprint fields to restrict editing during App provisioning for a long time. Fields can now be hidden from view completely by toggling the lock/unlock icon to the hidden setting. When users provision Apps based on this Blueprint, they will not see hidden fields at all. This gives administrators additional flexibility to mask unneeded complexity from users.



1.3.1 Getting Started

Requirements

Morpheus is a software based appliance installation capable of orchestrating many clouds and hypervisors. Before an installation is started it is important to understand some of the base requirements.

In the simplest configuration Morpheus needs one Appliance Server. The Appliance Server, by default, contains all the components necessary to orchestrate both VMs and containers. To get started some base requirements are recommended:

Base Requirements

Table 1: Supported Appliance Operating Systems

OS	Version(s)	Notes
Amazon Linux	2	
CentOS	7.x, 8.x	
Debian	8, 9, 10	FreeRDP 2.0 is not compatible with Debian 9. Guacd will remain at 1.0.0 for Appliances running on 9.
RHEL	7.x, 8.x	
SUSE SLES	12, 15	
Ubuntu	16.04, 18.04	14.04 is no longer supported for Appliance OS. Existing Appliances on 14.04 must upgrade to 16.04 or 18.04 PRIOR to upgrading to v4.2.1. Note: 14.04 is still supported by the Morpheus Agent.

- **Memory:** 16 GB recommended for default installations. 8 GB minimum required with 4 GB+ available storage swap space
- **Storage:** 200 GB storage minimum (see Storage Considerations below)
- **CPU:** 4-core, 1.4 GHz (or better), 64-bit CPU recommended for all-in-one systems. For a distributed-tier installation, it's recommended each tier have 2-core, 1.4 GHz (or better), 64-bit CPU
- Network connectivity from your users to the appliance over TCP 443 (HTTPS)
- Superuser privileges via the `sudo` command for the user installing the Morpheus appliance package
- Access to base `yum` or `apt` repos. Access to Optional RPMs repo required for RHEL 7.x
- An appliance license is required for any operations involving provisioning
- **Internet Connectivity (optional)**
 - To download from Morpheus' public docker repositories and system Virtual Image catalog
 - Offline installation require installing the supplemental package in addition to the regular installation package. Local `yum/apt` repo access still required for offline installations.

Note: Access to `yum` and `apt` repos is still required for offline installations.

- **VM and Host Agent Install (optional)**

- Inbound connectivity access from provisioned vm's and container hosts on ports 443 (Agent install and communication) and 80 (Linux Agent installs via yum and apt)
- An Appliance URL that is accessible/resolvable to all managed hosts. It is necessary for all hosts that are managed by Morpheus to be able to communicate with the appliance server ip on port 443. This URL is configured under Admin->Settings.

Note: Ubuntu 16.10 is not currently supported.

Storage Considerations

Upon initial installation Morpheus takes up less than 10 GB of space, however Morpheus Services, Virtual Images, Backups, Logs and stats and user uploaded and imported data require adequate space on the Morpheus Appliance(s) per Appliance Configuration and activity.

Important: It is the customers responsibility to ensure adequate storage space per configuration and use case.

Default Paths

/opt/morpheus Morpheus Application and Services Files

/var/opt/morpheus User, Application and Services Data, including default config Elasticsearch, RabbitMQ and Database data, and default Virtual Image path.

/var/log Morpheus Service logs

/tmp/morpheus Working directory for Backups

Images

Virtual Images can be uploaded to Morpheus Storage Providers for use across Clouds. By default when no Storage Provider has been added, images will write to `/var/opt/morpheus/morpheus-ui/vms`. Please ensure adequate space when uploading Images using local file paths.

Backups

Morpheus can offload snapshots when performing backups to local or other Storage Providers. By default when no Storage Provider has been added, backups will write to `/tmp/morpheus/backups/`. When using none NFS Storage providers, the backup file(s) must be written to `/tmp/morpheus/working/` before they can be zipped, sent to the destination Storage provider such as S3, and removed from `/tmp/morpheus/working/`. Please ensure adequate space in `/tmp/morpheus/` when offloading Backups.

Migrations

When performing a Hypervisor to Hypervisor migration, such as VMware to AWS, Virtual Images are written to local storage before conversion and/or upload to the target hypervisor. Please ensure adequate space in `/var/opt/morpheus/morpheus-ui/vms` or other configured local Storage Provider paths when performing Migrations.

VM Logs and Stats

When using a Morpheus configuration with locally installed Elasticsearch, VM, Container, Host and Appliance logs and stats are stored in Elasticsearch. Please ensure adequate space in `/var`, specifically `/var/opt/morpheus/elasticsearch` in relation to the number of Instances reporting logs, log frequency, and log retention count.

Morpheus Services Logs

Logs for services local to the Morpheus Appliance, such as the Morpheus UI, elasticsearch, rabbitmq, mysql, nginx and guacd are written to `/var/log/morpheus/`. Current logs are rotated nightly, zipped, and files older than 30 days are automatically removed. Misconfigured services, ports and permissions can cause excessive log file sizes.

Network Connectivity

Morpheus primarily operates via communication with its agent that is installed on all managed vm's or docker hosts. This is a lightweight agent responsible for aggregating logs and stats and sending them back to the client with minimal network traffic overhead. It also is capable of processing instructions related to provisioning and deployments instigated by the appliance server.

The Morpheus Agent exists for both linux and windows based platforms and opens NO ports on the guest operating system. Instead it makes an outbound SSL (https/wss) connection to the appliance server. This is what is known as the `appliance url` during configuration (in Admin->Settings). When the agent is started it automatically makes this connection and securely authenticates. Therefore, it is necessary for all vm's and docker based hosts that are managed by morpheus to be able to reach the appliance server ip on port 443.

Morpheus has numerous methods to execute agent installation, including zero open port methods.

Components

The Appliance Server automatically installs several components for the operation of Morpheus. This includes:

- RabbitMQ (Messaging)
- MySQL (Logistical Data store)
- Elasticsearch (Logs / Metrics store)
- Tomcat (Morpheus Application)
- Nginx (Web frontend)
- Guacamole (Remote console service for clientless remote console)
- Check Server (Monitoring Agent for custom checks added via UI)

All of these are installed in an isolated way using chef zero to `/opt/morpheus`. It is also important to note these services can be offloaded to separate servers or clusters as desired. For details check the installation section and high availability.

Common Ports & Requirements

The following chart is useful for troubleshooting Agent install, Static IP assignment, Remote Console connectivity, and Image transfers.

Table 2: Common Ports & Requirements

Feature	Method	OS	Source	Destination	Port	Requirement
Agent Communication	All	All	Node	Appliance	443	DNS Resolution from node to appliance url
Agent Install	All	Linux	Node	Appliance	80	Used for appliance yum and apt repos
	SSH	Linux	Appliance	Node	22	DNS Resolution from node to appliance url Virtual Images configured SSH Enabled on Virtual Image
	WinRM	Windows	Appliance	Node	5985	Not required for agent installation in VMware vCenter and vCloud Director type clouds. Otherwise, access from Morpheus App Nodes to Instance Node on 5985 Virtual Images configured WinRM Enabled on Virtual Image(<i>winrm quickconfig</i>)
	Cloud-init	Linux				Cloud-init installed on template/image Cloud-init settings populated in User Settings or in <i>Admin -> Provisioning</i> Agent install mode set to Cloud-Init in Cloud Settings
	Cloudbase-init	Windows				Cloudbase-init installed on template/image Cloud-init settings populated in User Settings or in <i>Admin -> Provisioning</i> Agent install mode set to Cloud-Init in Cloud Settings
	VMtools	All				VMtools installed on template Cloud-init settings populated in Morpheus user settings or in <i>Administration -> Provisioning</i> when using Static IP's Existing User credentials entered on Virtual Image when using DHCP RPC mode set to VMtools in VMware cloud settings.
Static IP Assignment & IP Pools	Cloud-Init	All				Network configured in Morpheus (Gateway, Primary and Secondary DNS, CIDR populated, DHCP disabled) Cloud-init/Cloudbase-init installed on template/image
1.3. Hide Blueprint fields						Cloud-init settings populated in Morpheus user settings or in <i>Administration -> Provisioning</i>
	VMware Tools	All				

Communication Data

The following table contains communication information, including frequency and configurability between Morpheus and its supported technology integrations.

Source	Push/Pull	Destination
Cloud Foundry App Check	Server Pull	Cloud Foundry Applications that exist within Morpheus
Docker Container Check	Server Pull	Docker containers that exist within Morpheus
Elastic Search Check	Server Pull	Elastic Search application
Microsoft SQL Server Check	Server Pull	Microsoft SQL application
Mongo Check	Server Pull	Mongo DB application
MySQL Check	Server Pull	MySQL application
Postgres Check	Server Pull	Postgres application
Push API Check	Client Push	Morpheus API
Rabbit MQ Check	Server Pull	Rabbit MQ application
Redis Check	Server Pull	Redis application
Riak Check	Server Pull	Riak application
SNMP Check	Server Pull	SNMP
Socket Check	Server Pull	Web Socket
Virtual Machine Check	Server Pull	Virtual Machine that exists within Morpheus
Web Check	Server Pull (GET) or Server Push (POST)	Web application
Public Cloud Integration	Server Pull	Alibaba Cloud
Public Cloud Integration	Server Pull	Amazon AWS
Public Cloud Integration	Server Pull	Amazon AWS GovCloud
Public Cloud Integration	Server Pull	DigitalOcean
Public Cloud Integration	Server Pull	Google Cloud Platform
Public Cloud Integration	Server Pull	Huawei Cloud
Public Cloud Integration	Server Pull	IBM Cloud
Public Cloud Integration	Server Pull	Microsoft Azure
Public Cloud Integration	Server Pull	Open Telekom Cloud
Public Cloud Integration	Server Pull	Oracle Public Cloud
Public Cloud Integration	Server Pull	UpCloud
Public Cloud Integration	Server Pull	VMware on AWS
Private Cloud Integration	Server Pull	Cisco UCS Manager
Private Cloud Integration	Server Pull	Dell EMC
Private Cloud Integration	Server Pull	HPE
Private Cloud Integration	Server Pull	HPE OneView
Private Cloud Integration	Server Pull	KVM
Private Cloud Integration	Server Pull	MacStadium
Private Cloud Integration	Server Pull	Microsoft Azure Stack
Private Cloud Integration	Server Pull	Microsoft Hyper-V
Private Cloud Integration	Server Pull	Microsoft SCVMM
Private Cloud Integration	Server Pull	Nutanix Acropolis
Private Cloud Integration	Server Pull	Openstack
Private Cloud Integration	Server Pull	Oracle VM
Private Cloud Integration	Server Pull	Pivotal Cloud Foundry
Private Cloud Integration	Server Pull	Supermicro
Private Cloud Integration	Server Pull	Vmware vCloud Director
Private Cloud Integration	Server Pull	Vmware ESXi
Private Cloud Integration	Server Pull	VMware Fusion

Source	Push/Pull	Destination
Private Cloud Integration	Server Pull	VMware vCenter
Private Cloud Integration	Server Pull	Xen Server
Automation Integration		Ansible
Automation Integration	Server Pull	Ansible Tower
Automation Integration	Server Pull	Chef
Automation Integration	Server Pull	Puppet
Automation Integration	Server Pull	Salt
Automation Integration		Terraform
Automation Integration	Server Pull	vRealize Orchestrator
Backup Integration	Server Pull	Commvault
Backup Integration	Server Pull	Veeam
Backup Integration	Server Pull	Rubrik
Backup Integration	Server Pull	Zerto
Backup Integration	Server Pull	Avamar
Build Integration	Server Pull	Jenkins
Container Integration	Server Pull	Docker
Container Integration		Docker Registry
Container Integration	Server Pull	Kubernetes
Deployment Integration	Server Pull	Git/Github
DNS Integration	Server Pull	AWS Route53
DNS Integration	Server Pull	Microsoft DNS
DNS Integration	Server Pull	PowerDNS
Identity Management Integration	Server Pull	Microsoft AD
Identity Management Integration	Server Pull	OneLogin
Identity Management Integration	Server Pull	Okta
Identity Management Integration	Server Pull	Jump Cloud
Identity Management Integration	Server Pull	LDAP
Identity Management Integration	Server Pull	SAML
IPAM Integration	Server Pull	Infoblox
IPAM Integration	Server Pull	phpIPAM
IPAM Integration	Server Pull	Bluecat
IPAM Integration	Server Pull	SolarWinds
ITSM Integration	Server Pull	ServiceNow
ITSM Integration	Server Pull	Cherwell
ITSM Integration	Server Pull	Remedy
Key & Certificate Integration	Server Pull	Venafi
Load Balancer Integration	Server Pull	AzureLB
Load Balancer Integration	Server Pull	F5 BigIP
Load Balancer Integration	Server Pull	Citrix NetScaler
Logging Integration		LogRhythm
Logging Integration		Splunk
Logging Integration		Syslog
Monitoring Integration	Server Pull	ServiceNow
Monitoring Integration		AppDynamics
Monitoring Integration		NewRelic
Network Integration	Server Pull	NSX-T
Network Integration	Server Pull	NSX-V
Network Integration	Server Pull	Cisco ACI
Network Integration	Server Pull	Unisys Stealth

Source	Push/Pull	Destination
Service Discovery Integration		Consul
Storage Integration	Server Pull	3Par
Storage Integration	Server Pull	Azure Storage
Storage Integration	Server Pull	Dell ECS
Storage Integration	Server Pull	Isilon
Morpheus Agent	Agent Pull	Application Tier

SELinux

If not required by organizational policy, we recommend setting SELinux to “Permissive” or “Disabled” modes to prevent any unnecessary security-related issues. Morpheus versions 3.6.0 and higher do support “Enforcing” mode if it is required by your organization due to IT policies. Set the mode appropriately prior to running the Morpheus installer and it will make the required changes based on your chosen SELinux context.

Important: Setting SELinux to “Enforcing” mode requires policies to be configured correctly in order for the Morpheus appliance to function correctly.

Supported Languages

Morpheus supports a number of different UI languages, including:

- English
- German
- Spanish
- Chinese (Simplified)
- Portuguese (Brazil)

Currently, UI language is not configurable from within Morpheus itself. Changing the language within the application will involve some combination of operating system and web browser language setting changes. Morpheus must also have a translation set for your chosen language to see a change. Depending on the browser and the operating system, you may need to fully close and reopen the web browser or restart the machine completely.

Note: Many of Morpheus’ language packs are generated by our clients. For that reason, we cannot guarantee accuracy and completeness of the translation. As new UI elements are added, existing language sets may not have immediate updates to keep pace with application changes. If you would like to contribute to a new or existing language pack, contact your account team or Morpheus support. Contributed content would be included with the next application update.

Installation

Installation Overview

Important: Morpheus v4.2.0 enhanced security configuration restricts incoming appliance connections to TLS v1.2, potentially impacting front-end load balancer monitoring/health checks that support only TLS v1.1 or lower, as well as Morpheus Agent installations for Windows nodes using .net versions that do not support TLS v1.2. Refer to TLS

Morpheus comes packaged as a `deb`ian or `yum` based package. The default configuration installs all required services on a single vm or bare metal Host. Morpheus can be configured in a distributed architecture to use one or multiple external services, and multiple application Hosts can be configured for High Availability configurations.

All components required for Morpheus are installed and configured by default during the Morpheus `reconfigure` command. The Morpheus config file, `morpheus.rb`, can optionally be configured to point the Morpheus App to external services (distributed configuration).

Morpheus can optionally be configured to use external Database, Messaging, and/or Search Tiers. This means instead of installing, for example, MySQL on the same host as the Morpheus App, the Morpheus configuration file (`morpheus.rb`) is setup to point to an external MySQL host, cluster or service, and MySQL will not be installed or configured on the Appliance Host.

Install Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

Configuration Options

- **Single Host (All-In-One/default)** All tiers a single host. The reconfigure process installs all required services. This is the default configuration.
- **Single Hosts with Distributed Service(s)** Transactional Database, Non-Transactional Database, and/or Message tiers are externalized, with the remaining services on a single host. The reconfigure process installs all services not set to false in `/etc/morpheus.morpheus.rb`
- **Clustered Hosts with Distributed Transactional Database (3-Node HA)** Application, Message and Non-Transactional tiers are installed and clustered on three or more hosts, with all three hosts pointing to externalized database tier. The reconfigure process installs all services except `mysql`.
- **App Host(s) with Distributed Services (Full HA)** Application tier is installed on one or more hosts. All UI hosts point to externalized Transactional Database, Non-Transactional Database, and Message Tiers. The reconfigure process installs only Application services.

Distributed Configurations

Morpheus provides a wide array of options when it comes to deployment architectures. It can start as a simple one machine instance where all services run on the same machine, or it can be split off into individual services per machine and configured in a high availability configuration, either in the same region or cross-region. Naturally, high availability can grow more complicated, depending on the configuration you want to do and this article will cover the basic concepts of the Morpheus HA architecture that can be used in a wide array of configurations.

There are four primary tiers of services represented within the Morpheus appliance. They are the Application Tier, Transactional Database Tier, Non-Transactional Database Tier, and Message Tier. Each of these tiers have their own recommendations for High availability deployments that we need to cover.

Application Tier

The application tier is easily installed with the same Debian or yum repository package that Morpheus is normally distributed with. Advanced configuration allows for the additional tiers to be skipped and leave only the “stateless” services that need run. These stateless services include Nginx and Tomcat. These machines should also have at least 8gb of Memory. They can be configured across all regions and placed behind a central load-balancer or Geo based load-balancer. They typically connect to all other tiers as none of the other tiers talk to each other besides through the central application tier. One final piece when it comes to setting up the Application tier is a shared storage means is necessary when it comes to maintaining things like deployment archives, virtual image catalogs, backups, etc. These can be externalized to an object storage service such as amazon S3 or Openstack Swiftstack as well. If not using those options a simple NFS cluster can also be used to handle the shared storage structure.

Transactional Database Tier

The Transactional database tier usually consists of a MySQL compatible database. It is recommended that a lockable clustered configuration be used (Currently Percona XtraDB Cluster is the most recommended in Permissive Mode). There are several documents online related to configuring and setting up an XtraDB Cluster but it most simply can be laid out in a many master configuration. There can be some nodes setup with replication delay as well as some with no replication delay. It is common practice to have no replication delay within the same region and allow some replication delay cross region. This does increase the risk of job run overlap between the 2 regions however, the concurrent operations typically self-correct and this is a non-issue.

Non-Transactional Database Tier

The Non-Transactional tier consists of an ElasticSearch (version 7.6.0) cluster. Elastic Search is used for log aggregation data and temporal aggregation data (essentially stats, metrics, and logs). This enables for a high write throughput at scale. ElasticSearch is a Clustered database meaning all nodes no matter the region need to be connected to each other over what they call a “Transport” protocol. It is fairly simple to get setup as all nodes are identical. It is also a java based system and does require a sizable chunk of memory for larger data sets. (8gb) is recommended and more nodes can be added to scale either horizontally or vertically.

Messaging Tier

The Messaging tier is an AMQP based tier along with STOMP Protocol (used for agent communication). The primary model recommended is to use RabbitMQ for queue services. RabbitMQ is also a clustered based queuing system and needs at least 3 instances for HA configurations. This is due to elections in the failover scenarios RabbitMQ can manage. If doing a cross-region HA RabbitMQ cluster it is recommended to have at least 3 rabbit queue clusters per region. Typically to handle HA a RabbitMQ cluster should be placed between a load balancer and the front-end application server to handle cross host connections. The ports necessary to forward in a Rabbit MQ cluster are (5672, and 61613). A RabbitMQ cluster can run on smaller memory machines depending on how frequent large requests bursts occur. 4–8gb of Memory is recommended to start.

Pros/Cons

Single Host

Advantages

- Simple Installation - Morpheus Installs all required services
- Simple Configuration - Morpheus configures all required services
- Simple Maintenance - All service connections and credential are local - All logs are local - All Data is local (by default)
- Not dependent on network connections for vital services - Facilitates speed and reliability

Disadvantages

- Single point of failure
- Individual services cannot be scaled
- Upgrades require (minimal) downtime
- Single region

Single Hosts with Distributed Service(s)

Advantages

- Individual services can be scaled
- Managed Services such as RDS can be utilized

Disadvantages

- Single region
- External services require additional configuration and maintenance
- Morpheus is subject to network performance, configuration and availability
- Increased Installation time possible

Clustered Hosts with Distributed Transactional Database

Advantages

- Database can be scaled vertically and/or horizontally
- Managed Services such as RDS can be utilized
- Zero down time upgrades
- No single point of failure
- RabbitMQ and Elasticsearch Clusters

Disadvantages

- External Database services requires additional configuration and maintenance
- App Host Clustering requires additional configuration and maintenance
- Extended Installation time

- Increased Infrastructure requirements
- Load Balancer required to front App Hosts
- Shared Storage configuration required

App Host(s) with Distributed Services

Advantages

- Individual services can be scaled vertically and/or horizontally
- Managed Services such as RDS can be utilized
- Zero down time upgrades
- No single point of failure
- Multi region support


Disadvantages

- External services require additional configuration and maintenance
- Extended Installation time
- Increased Infrastructure Requirements
- Increased Networking requirements
- Load Balancer required to front App Hosts
- Shared Storage configuration required
- Rabbit Load balancer required

Single Node Installation

Single Node Install Overview

In a Single Host/All-in-one configuration, all components required for Morpheus are automatically installed and configured during the Morpheus `reconfigure` command.



Appliance Host

- **Application**
 - Morpheus App
- **Web Server/Proxy**
 - Nginix
- **Database**
 - MySQL
- **Messaging**
 - RabbitMQ
- **Search**
 - Elasticsearch
- **Console**
 - Guacamole
- **Monitoring**
 - Check Server

Default Paths

Morpheus follows several install location conventions. Below is a list of the system paths.

Important: Altering the default system paths is not supported and may break functionality.

- Installation Location: `/opt/morpheus`
- Log Location: `/var/log/morpheus`
 - Morpheus-UI: `/var/log/morpheus/morpheus-ui`
 - MySQL: `/var/log/morpheus/mysql`
 - NginX: `/var/log/morpheus/nginx`
 - Check Server: `/var/log/morpheus/check-server`
 - Elastic Search: `/var/log/morpheus/elasticsearch`
 - RabbitMQ: `/var/log/morpheus/rabbitmq`
- User-defined install/config: `/etc/morpheus/morpheus.rb`

Single Node Install on CentOS

To get started installing Morpheus on CentOS a few preparatory items should be addressed first.

1. Configure firewalld to allow access from users on port 443 (Or remove firewall if not required).
2. Make sure the machine is self resolvable to its own hostname.

Important: If the machine is unable to resolve its own hostname `nslookup hostname` some installation commands will be unable to verify service health during installation and fail.

3. Next simply download the relevant `.rpm` package for installation. This package can be acquired from <https://morpheushub.com> downloads section.

Tip: Use the `wget` command to directly download the package to your appliance server. i.e. `wget https://downloads.morpheusdata.com/path/to/package.rpm`

4. Next we must install the package onto the machine and configure the morpheus services:

```
sudo rpm -i morpheus-appliance-x.x.x-1.x86_64.rpm
sudo morpheus-ctl reconfigure
```

5. Once the installation is complete the web interface will automatically start up. By default it will be resolvable at `https://your_machine_name` and in many cases this may not be resolvable from your browser. The url can be changed by editing `/etc/morpheus/morpheus.rb` and changing the value of `appliance_url`. After this has been changed simply run :

```
sudo morpheus-ctl reconfigure
sudo morpheus-ctl stop morpheus-ui
sudo morpheus-ctl start morpheus-ui
```

Note: The morpheus-ui can take 2-3 minutes to startup before it becomes available.

There are additional install settings that can be viewed in the [Additional Configuration Options](#) section.

Once the browser is pointed to the appliance a first time setup wizard will be presented. Please follow the on screen instructions by creating the master account. From there you will be presented with the license settings page where a license can be applied for use (if a license is required you may request one or purchase one by contacting your sales representative).

More details on setting up infrastructure can be found throughout this guide.

Tip: If any issues occur it may be prudent to check the morpheus log for details at `/var/log/morpheus/morpheus-ui/current`.

Single Node Install on Debian/Ubuntu

To get started installing Morpheus on Ubuntu or Debian a few preparatory items should be addressed first.

1. First make sure the apt repository is up to date by running `sudo apt-get update`. It is advisable to verify the assigned hostname of the machine is self-resolvable.

Important: If the machine is unable to resolve its own hostname `nslookup hostname` some installation commands will be unable to verify service health during installation and fail.

1. Next simply download the relevant `.deb` package for installation. This package can be acquired from <https://morpheushub.com> downloads section.

Tip: Use the `wget` command to directly download the package to your appliance server. i.e. `wget https://downloads.morpheusdata.com/path/to/package/morpheus-appliance_x.x.x-1_amd64.deb`

1. Next we must install the package onto the machine and configure the morpheus services:

```
sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
sudo morpheus-ctl reconfigure
```

2. Once the installation is complete the web interface will automatically start up. By default it will be resolvable at `https://your_machine_name` and in many cases this may not be resolvable from your browser. The url can be changed by editing `/etc/morpheus/morpheus.rb` and changing the value of `appliance_url`. After this has been changed simply run:

```
sudo morpheus-ctl reconfigure
sudo morpheus-ctl stop morpheus-ui
sudo morpheus-ctl start morpheus-ui
```

Note: The *morpheus-ui* can take 2-3 minutes to startup before it becomes available.

There are additional install settings that can be viewed in the [Additional Configuration Options](#) section.

Once the browser is pointed to the appliance a first time setup wizard will be presented. Please follow the on screen instructions by creating the master account. From there you will be presented with the license settings page where a license can be applied for use (if a license is required you may request one or purchase one by contacting your sales representative).

More details on setting up infrastructure can be found throughout this guide.

Tip: If any issues occur it may be prudent to check the morpheus log for details at `/var/log/morpheus/morpheus-ui/current`.

Single Node Install on RHEL

To get started installing Morpheus on RHEL/RedHat a few prerequisite items are required.

1. Configure firewalld to allow access from users on 443 (Or remove firewall if not required).
2. Make sure the machine is self resolvable to its own hostname.
3. For RHEL 7.x, the Optional RPMS repo needs to be added for Reconfigure to succeed. Its does not need to be added For RHEL 8.x, as the Optional RPMs repo is now part of the appstream repo that is enabled by default in RHEL 8.x.
 - **RHEL 7.x Amazon:** `yum-config-manager --enable rhel-7-server-rhui-optional-rpms`
 - **RHEL 7.x:** `yum-config-manager --enable rhel-7-server-optional-rpms`

Note: For Amazon users a Redhat subscription is not required if the appropriate yum REGION repository is added instead as demonstrated above.

The RedHat Enterprise Linux server needs to be registered and activated with Redhat subscription. The server optional rpms repo needs to be enabled as well.

To check if the server has been activated please run the subscription-manager version. Subscription manager will return the version plus the python dependency version.

If the server has not been registered and activated then the subscription manager version will return the below message.

```
sudo subscription-manager version
server type: This system is currently not registered
subscription management server: 0.9.51.24.-1
subscription-manager: 1.10.14-7.el7 python-rhsm: 1.10.12-2.el7
```

When a server has been registered and activated with Redhat the subscription manager will return the below message.

```
sudo subscription-manager version
server type: Red Hat Subscription Management
subscription management server: 0.9.51.24-1
subscription-manager: 1.10.14-7.el7 python-rhsm: 1.10.12-2.el7
```

If the subscription manager re-turns the message `This system is currently not registered` please follow the below steps to register the server.

Tip: To register the server you will need to have sudo permissions [Member of the Wheel group] or root access to the server. You will also need your Redhat registered email address and password.

subscription-manager register

```
sudo subscription-manager register
Username: redhat@example.com
Password: . subscription-manager auto --attach
```

Note: This can take a minute to complete

```
sudo subscription-manager attach --auto

Installed Product Current Status: Product Name: Red Hat Enterprise Linux
Server Status: Subscribed
```

To check to see if the RHEL server has the Red Hat Enterprise Linux 7 Server - Optional (RPMs) repo enabled please run the following command to return the repo status.

Tip: To check the server repos you will need to have sudo permissions [Member of the Wheel group] or root access to the server.

```
sudo yum repolist all | grep "rhel-7-server-optional-rpms" rhel-7-server-optional-
↳rpms/7Server/x86_64 disabled
```

If the repo status was returned as disabled then you will need to enable the repo using the subscription manager like below.

```
sudo subscription-manager repos --enable rhel-7-server-optional-rpms
Repository 'rhel-7-server-optional-rpms' is enabled for this system.
```

The message Repo 'rhel-7-server-optional-rpms' is enabled for this system. will appear after enabling the repo. This will confirm that the repo has been enabled.

Next simply download the relevant .rpm package for installation. This package can be acquired from morpheushub.com.

Tip: Use the wget command to directly download the package to your appliance server. i.e. `wget https://downloads.morpheusdata.com/path/to/package.rpm`

Next we must install the package onto the machine and configure the morpheus services:

```
sudo rpm -i morpheus-appliance_x.x.x-1_amd64.rpm
sudo morpheus-ctl reconfigure
```

Once the installation is complete the web interface will automatically start up. By default it will be resolvable at `https://your_machine_name` and in many cases this may not be resolvable from your browser. The url can be changed by editing `/etc/morpheus/morpheus.rb` and changing the value of `appliance_url`. After this has been changed simply run:

```
sudo morpheus-ctl reconfigure
sudo morpheus-ctl stop morpheus-ui
sudo morpheus-ctl start morpheus-ui
```

Note: The `morpheus-ui` can take 2-3 minutes to startup before it becomes available.

There are additional install settings that can be viewed in the *Additional Configuration Options* section.

Once the browser is pointed to the appliance a first time setup wizard will be presented. Please follow the on screen instructions by creating the master account. From there you will be presented with the license settings page where a license can be applied for use (if a license is required you may request one or purchase one by contacting your sales representative).

More details on setting up infrastructure can be found throughout this guide.

Tip: If any issues occur it may be prudent to check the morpheus log for details at `/var/log/morpheus/morpheus-ui/current`.

Distributed Installations

Distributed Installation Overview

Morpheus provides a wide array of options when it comes to deployment architectures. It can start as a simple one machine instance where all services run on the same machine, or it can be split off into individual services per machine and configured in a high availability configuration, either in the same region or cross-region. Naturally, high availability can grow more complicated, depending on the configuration you want to do and this article will cover the basic concepts of the Morpheus HA architecture that can be used in a wide array of configurations.

There are four primary tiers of services represented within the Morpheus appliance. They are the App Tier, Transactional Database Tier, Non-Transactional Database Tier, and Message Tier. Each of these tiers have their own recommendations for High availability deployments that we need to cover.

Important: This is a sample configuration only. Customer configurations and requirements will vary.

Application Tier

The application tier is easily installed with the same debian or yum repository package that Morpheus is normally distributed with. Advanced configuration allows for the additional tiers to be skipped and leave only the “stateless” services that need run. These stateless services include Nginx and Tomcat. These machines should also have at least 8gb of Memory. They can be configured across all regions and placed behind a central load-balancer or Geo based load-balancer. They typically connect to all other tiers as none of the other tiers talk to each other besides through the central application tier. One final piece when it comes to setting up the Application tier is a shared storage means is necessary when it comes to maintaining things like deployment archives, virtual image catalogs, backups, etc. These can be externalized to an object storage service such as amazon S3 or Openstack Swiftstack as well. If not using those options a simple NFS cluster can also be used to handle the shared storage structure.

Transactional Database Tier

The Transactional database tier usually consists of a MySQL compatible database. It is recommended that a lockable clustered configuration be used (Currently Percona XtraDB Cluster is the most recommended in Permissive Mode). There are several documents online related to configuring and setting up an XtraDB Cluster but it most simply can be laid out in a many master configuration. There can be some nodes setup with replication delay as well as some with no replication delay. It is common practice to have no replication delay within the same region and allow some replication delay cross region. This does increase the risk of job run overlap between the 2 regions however, the concurrent operations typically self-correct and this is a non-issue.

Non-Transactional Database Tier

The Non-Transactional tier consists of an ElasticSearch (version 7.6.0) cluster. Elastic Search is used for log aggregation data and temporal aggregation data (essentially stats, metrics, and logs). This enables for a high write throughput at scale. ElasticSearch is a Clustered database meaning all nodes no matter the region need to be connected to each other over what they call a “Transport” protocol. It is fairly simple to get setup as all nodes are identical. It is also a java based system and does require a sizable chunk of memory for larger data sets. (8gb) is recommended and more nodes can be added to scale either horizontally or vertically.

Messaging Tier

The Messaging tier is an AMQP based tier along with STOMP Protocol (used for agent communication). The primary model recommended is to use RabbitMQ for queue services. RabbitMQ is also a clustered based queuing system and needs at least 3 instances for HA configurations. This is due to elections in the failover scenarios rabbitmq can manage. If doing a cross-region HA rabbitmq cluster it is recommended to have at least 3 rabbit queue clusters per region. Typically to handle HA a RabbitMQ cluster should be placed between a load balancer and the front-end application server to handle cross host connections. The ports necessary to forward in a Rabbit MQ cluster are (5672, and 61613). A rabbitmq cluster can run on smaller memory machines depending on how frequent large requests bursts occur. 4–8gb of Memory is recommended to start.


3-Node HA Install

Distributed App Nodes with Externalized DB

Assumptions

This guide assumes the following:

- The Baremetal nodes cannot access the public internet
- The base OS is RHEL 7.x
- Shortname versions of hostnames will be resolvable
- All nodes have access to a shared volume for `/var/opt/morpheus/morpheus-ui`. This can be done as a post startup step.
- This configuration will support the complete loss of a single node, but no more. Specifically the Elasticsearch tier requires at least two nodes to always be clustered..



Default Locations

Morpheus follows several install location conventions. Below is a list of system defaults for convenient management:

- Installation Location: `/opt/morpheus`
- Log Location: `/var/log/morpheus`
 - Morpheus-UI: `/var/log/morpheus/morpheus-ui`
 - NginX: `/var/log/morpheus/nginx`
 - Check Server: `/var/log/morpheus/check-server`
 - Elastic Search: `/var/log/morpheus/elasticsearch`
 - RabbitMQ: `/var/log/morpheus/rabbitmq`
- User-defined install/config: `/etc/morpheus/morpheus.rb`

Database Cluster Setup (Percona XtraDB Cluster)

Out of the box Morpheus uses MySQL but Morpheus supports any mySQL-compliant database. There are many ways to set up a highly available, MySQL dialect-based database. One which has found favor with many of our customers is Percona's XtraDB Cluster. Percona's product is based off of Galera's WSREP Clustering, which is also supported.

Important: Additional configuration for Percona Clusters with TLS enabled is required. Refer to Percona TLS Configuration in our full HA docs for details.

Requirements

Note: Morpheus idiomatically connects to database nodes over 3306

Once you have your database installed and configured:

1. Create the Database you will be using with morpheus.

```
mysql> CREATE DATABASE morpheus CHARACTER SET utf8 COLLATE utf8_general_ci;
mysql> show databases;
```

2. Next create your morpheus database user. The user needs to be either at the IP address of the morpheus application server or use '@%' within the user name to allow the user to login from anywhere.

```
mysql> CREATE USER '$morpheus_db_user_name'@$source_ip IDENTIFIED BY '$morpheus_
↳db_user_pw';
```

3. Next Grant your new morpheus user permissions to the database.

```
mysql> GRANT ALL PRIVILEGES ON morpheus_db_name.* TO 'morpheus_db_user'@$source_
↳ip IDENTIFIED BY 'morpheus_db_user_pw' with grant option;

mysql> GRANT SELECT, PROCESS, SHOW DATABASES, SUPER ON *.* TO 'morpheus_db_user'@
↳'$source_ip' IDENTIFIED BY 'morpheus_db_user_pw';

mysql> FLUSH PRIVILEGES;
```

4. Checking Permissions for your user.

```
SHOW GRANTS FOR '$morpheus_db_user_name'@$source_ip';
```

Continued Installation Steps

1. First begin by downloading the requisite Morpheus packages either to the nodes or to your workstation for transfer. These packages need to be made available on the nodes you wish to install Morpheus on.

```
[root@app-server-1 ~]# wget https://example/path/morpheus-appliance-ver-1.el7.x86_
↳64.rpm
[root@app-server-1 ~]# wget https://example/path/morpheus-appliance-offline-ver-1.
↳noarch.rpm
```

2. Once the packages are available on the nodes they can be installed. Make sure that no steps beyond the rpm install are run.

```
[root@app-server-1 ~] rpm -i morpheus-appliance-ver-1.el7.x86_64.rpm
[root@app-server-1 ~] rpm -i morpheus-appliance-offline-ver-1.noarch.rpm
```

3. Next you will need to edit the Morpheus configuration file `/etc/morpheus/morpheus.rb` on each node.

Node 1

```
appliance_url 'https://morpheus1.localdomain'
elasticsearch['es_hosts'] = {'10.100.10.121' => 9200, '10.100.10.122' => 9200,
↪ '10.100.10.123' => 9200}
elasticsearch['node_name'] = '10.100.10.121'
elasticsearch['host'] = '0.0.0.0'
rabbitmq['host'] = '0.0.0.0'
rabbitmq['nodename'] = 'rabbit@node01'
mysql['enable'] = false
mysql['host'] = '10.100.10.111'
mysql['morpheus_db'] = 'morpheusdb'
mysql['morpheus_db_user'] = 'morpheus'
mysql['morpheus_password'] = 'password'
```

Node 2

```
appliance_url 'https://morpheus2.localdomain'
elasticsearch['es_hosts'] = {'10.100.10.121' => 9200, '10.100.10.122' => 9200,
↪ '10.100.10.123' => 9200}
elasticsearch['node_name'] = '10.100.10.122'
elasticsearch['host'] = '0.0.0.0'
rabbitmq['host'] = '0.0.0.0'
rabbitmq['nodename'] = 'rabbit@node02'
mysql['enable'] = false
mysql['host'] = '10.100.10.111'
mysql['morpheus_db'] = 'morpheusdb'
mysql['morpheus_db_user'] = 'morpheus'
mysql['morpheus_password'] = 'password'
```

Node 3

```
appliance_url 'https://morpheus3.localdomain'
elasticsearch['es_hosts'] = {'10.100.10.121' => 9200, '10.100.10.122' => 9200,
↪ '10.100.10.123' => 9200}
elasticsearch['node_name'] = '10.100.10.123'
elasticsearch['host'] = '0.0.0.0'
rabbitmq['host'] = '0.0.0.0'
rabbitmq['nodename'] = 'rabbit@node03'
mysql['enable'] = false
mysql['host'] = '10.100.10.111'
mysql['morpheus_db'] = 'morpheusdb'
mysql['morpheus_db_user'] = 'morpheus'
mysql['morpheus_password'] = 'password'
```

Important: The elasticsearch node names set in `elasticsearch['node_name']` must match the host entries in `elasticsearch['es_hosts']`. `node_name` is used for `node.name` and `es_hosts` is used for `cluster.initial_master_nodes` in the generated `elasticsearch.yml` config. node names that do not match entries in `cluster.initial_master_nodes` will cause clustering issues.

4. Reconfigure on all nodes

```
[root@app-server-1 ~] morpheus-ctl reconfigure
```

Morpheus will come up on all nodes and Elasticsearch will auto-cluster. The only item left is the manual clustering of RabbitMQ.

5. Select one of the nodes to be your Source Of Truth (SOT) for RabbitMQ clustering. We need to copy the secrets

for RabbitMQ, copy the erlang cookie and join the other nodes to the SOT node.

Begin by copying secrets from the SOT node to the other nodes.

```
[root@app-server-1 ~] cat /etc/morpheus/morpheus-secrets.json

"rabbitmq": {
  "morpheus_password": "***REDACTED***",
  "queue_user_password": "***REDACTED***",
  "cookie": "***REDACTED***"
},
```

Then copy the erlang.cookie from the SOT node to the other nodes

```
[root@app-server-1 ~]# cat /opt/morpheus/embedded/rabbitmq/.erlang.cookie

# 754363AD864649RD63D28
```

6. Once this is done run a reconfigure on the two nodes that are NOT the SOT nodes.

```
[root@app-server-2 ~] morpheus-ctl reconfigure
```

Note: This step will fail. This is ok, and expected. If the reconfigure hangs then use Ctrl+C to quit the reconfigure run and force a failure.

7. Subsequently we need to stop and start Rabbit on the NOT SOT nodes.

Important: The commands below must be run at root

Note: If you receive an error unable to connect to epmd (port 4369) on app-server-1: nxdomain (non-existing domain) make sure to add all IPs and hostnames to the etc/hosts file like so:

```
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
127.0.0.1 app-server-1.localdomain app-server-2 localhost
127.0.0.1 container16
10.100.10.113 app-server-1
10.100.10.114 app-server-2
10.100.10.115 app-server-3
```

```
[root@app-server-2 ~]# morpheus-ctl stop rabbitmq
[root@app-server-2 ~]# morpheus-ctl start rabbitmq
[root@app-server-2 ~]# source /opt/morpheus/embedded/rabbitmq/.profile
[root@app-server-2 ~]# rabbitmqctl stop_app

Stopping node 'rabbit@app-server-2' ...

[root@app-server-2 ~]# rabbitmqctl join_cluster rabbit@app-server-1

Clustering node 'rabbit@app-server-2' with 'rabbit@app-server-1' ...
```

(continues on next page)

(continued from previous page)

```
[root@app-server-2 ~]# rabbitmqctl start_app  
Starting node 'rabbit@app-server-2' ...
```

8. Now make sure to reconfigure

```
[root@app-server-2 ~] morpheus-ctl reconfigure
```

9. Once the Rabbit services are up and clustered on all nodes, apply required ha-mode and expires policies to the morpheus vhost:

```
[root@app-server-2 ~] rabbitmqctl set_policy -p morpheus --apply-to queues --  
→priority 2 statCommands "statCommands.*" '{"expires":1800000, "ha-mode":"all"}'  
[root@app-server-2 ~] rabbitmqctl set_policy -p morpheus --apply-to queues --  
→priority 2 morpheusAgentActions "morpheusAgentActions.*" '{"expires":1800000,  
→"ha-mode":"all"}'  
[root@app-server-2 ~] rabbitmqctl set_policy -p morpheus --apply-to queues --  
→priority 2 monitorJobs "monitorJobs.*" '{"expires":1800000, "ha-mode":"all"}'  
[root@app-server-2 ~] rabbitmqctl set_policy -p morpheus --apply-to all --  
→priority 1 ha ".*" '{"ha-mode":"all"}'
```

Important: Failure to set the proper policies will result in degraded RabbitMQ performance, Java Heap issues, and/or refused RabbitMQ connections resulting in degraded Morpheus UI performance, unconsumed messages or UI failure.

10. The last thing to do is restart the Morpheus UI on the two nodes that are NOT the SOT node.

```
[root@app-server-2 ~]# morpheus-ctl restart morpheus-ui
```

If this command times out then run:

```
[root@app-server-2 ~]# morpheus-ctl kill morpheus-ui  
[root@app-server-2 ~]# morpheus-ctl start morpheus-ui
```

11. You will be able to verify that the UI services have restarted properly by inspecting the logfiles. A standard practice after running a restart is to tail the UI log file.

```
root@app-server-2 ~]# morpheus-ctl tail morpheus-ui
```

12. Lastly, we need to ensure that Elasticsearch is configured in such a way as to support a quorum of 2. We need to do this step on EVERY NODE.

```
[root@app-server-2 ~]# echo "discovery.zen.minimum_master_nodes: 2" >> /opt/  
→morpheus/embedded/elasticsearch/config/elasticsearch.yml  
[root@app-server-2 ~]# morpheus-ctl restart elasticsearch
```

Note: For moving /var/opt/morpheus/morpheus-ui files into a shared volume make sure ALL Morpheus services on ALL three nodes are down before you begin.

```
[root@app-server-1 ~]# morpheus-ctl stop
```

13. Permissions are as important as is content, so make sure to preserve directory contents to the shared volume.

- Subsequently you can start all Morpheus services on all three nodes and tail the Morpheus UI log file to inspect errors.

Database Migration

If your new installation is part of a migration then you need to move the data from your original Morpheus database to your new one. This is easily accomplished by using a stateful dump.

- To begin this, stop the Morpheus UI on your original Morpheus server:

```
[root@app-server-old ~]# morpheus-ctl stop morpheus-ui
```

- Once this is done you can safely export. To access the MySQL shell we will need the password for the Morpheus DB user. We can find this in the morpheus-secrets file:

```
[root@app-server-old ~]# cat /etc/morpheus/morpheus-secrets.json
```

```
{
  "mysql": {
    "root_password": "***REDACTED***",
    "morpheus_password": "***REDACTED***",
    "ops_password": "***REDACTED***"
  },
  "rabbitmq": {
    "morpheus_password": "***REDACTED***",
    "queue_user_password": "***REDACTED***",
    "cookie": "***REDACTED***"
  },
  "vm-images": {
    "s3": {
      "aws_access_id": "***REDACTED***",
      "aws_secret_key": "***REDACTED***"
    }
  }
}
```

- Take note of this password as it will be used to invoke a dump. Morpheus provides embedded binaries for this task. Invoke it via the embedded path and specify the host. In this example we are using the Morpheus database on the MySQL listening on localhost. Enter the password copied from the previous step when prompted:

```
[root@app-server-old ~]# /opt/morpheus/embedded/mysql/bin/mysqldump -u morpheus -
→h 127.0.0.1 morpheus -p > /tmp/morpheus_backup.sql

Enter password:
```

This file needs to be pushed to the new Morpheus Installation's backend. Depending on the GRANTS in the new MySQL backend, this will likely require moving this file to one of the new Morpheus frontend servers.

- Once the file is in place it can be imported into the backend. Begin by ensuring the Morpheus UI service is stopped on all of the application servers:

```
[root@app-server-1 ~]# morpheus-ctl stop morpheus-ui
[root@app-server-2 ~]# morpheus-ctl stop morpheus-ui
[root@app-server-3 ~]# morpheus-ctl stop morpheus-ui
```

- Then you can import the MySQL dump into the target database using the embedded MySQL binaries, specifying the database host, and entering the password for the Morpheus user when prompted:

```
[root@app-server-1 ~]# /opt/morpheus/embedded/mysql/bin/mysql -u morpheus -h 10.
↪130.2.38 morpheus -p < /tmp/morpheus_backup.sql
Enter password:
```

Recovery

If a node happens to crash most of the time Morpheus will start upon boot of the server and the services will self-recover. However, there can be cases where RabbitMQ and Elasticsearch are unable to recover in a clean fashion and it requires minor manual intervention. Regardless, it is considered best practice when recovering a restart to perform some manual health checks.

```
[root@app-server-1 ~]# morpheus-ctl status
run: check-server: (pid 17808) 7714s;
run: log: (pid 549) 8401s
run: elasticsearch: (pid 19207) 5326s;
run: log: (pid 565) 8401s
run: guacd: (pid 601) 8401s;
run: log: (pid 573) 8401s
run: morpheus-ui: (pid 17976) 7633s;
run: log: (pid 555) 8401s
run: nginx: (pid 581) 8401s;
run: log: (pid 544) 8401s
run: rabbitmq: (pid 17850) 7708s;
run: log: (pid 542) 8401s
run: log: (pid 548) 8401s
```

But, a status can report false positives if, say, RabbitMQ is in a boot loop or Elasticsearch is up, but not able to join the cluster. It is always advisable to tail the logs of the services to investigate their health.

```
[root@app-server-1 ~]# morpheus-ctl tail rabbitmq
[root@app-server-1 ~]# morpheus-ctl tail elasticsearch
```

To minimize disruption to the user interface, it is advisable to remedy Elasticsearch clustering first. Due to write locking in Elasticsearch it can be required to restart other nodes in the cluster to allow the recovering node to join. Begin by determining which Elasticsearch node became the master during the outage. On one of the two other nodes (not the recovered node):

```
[root@app-server-2 ~]# curl localhost:9200/_cat/nodes
app-server-1 10.100.10.121 7 47 0.21 d * morpheus1
localhost 127.0.0.1 4 30 0.32 d m morpheus2
```

The master is determined by identifying the row with the '*' in it. SSH to this node (if different) and restart Elasticsearch.

```
[root@app-server-1 ~]# morpheus-ctl restart elasticsearch
```

Go to the other of the two 'up' nodes and run the curl command again. If the output contains three nodes then Elasticsearch has been recovered and you can move on to re-clustering RabbitMQ. Otherwise you will see output that contains only the node itself:

```
[root@app-server-2 ~]# curl localhost:9200/_cat/nodes
localhost 127.0.0.1 4 30 0.32 d * morpheus2
```

If this is the case then restart Elasticsearch on this node as well:

```
[root@app-server-2 ~]# morpheus-ctl restart elasticsearch
```

After this you should be able to run the curl command and see all three nodes have rejoined the cluster:

```
[root@app-server-2 ~]# curl localhost:9200/_cat/nodes
app-server-1 10.100.10.121 9 53 0.31 d * morpheus1
localhost 127.0.0.1 7 32 0.22 d m morpheus2
app-server-3 10.100.10.123 3 28 0.02 d m morpheus3
```

The most frequent case of restart errors for RabbitMQ is with epmd failing to restart. Morpheus's recommendation is to ensure the epmd process is running and daemonized by starting it:

```
[root@app-server-1 ~]# /opt/morpheus/embedded/lib/erlang/erts-5.10.4/bin/epmd -daemon
```

And then restarting RabbitMQ:

```
[root@app-server-1 ~]# morpheus-ctl restart rabbitmq
```

And then restarting the Morpheus UI service:

```
[root@app-server-1 ~]# morpheus-ctl restart morpheus-ui
```

Again, it is always advisable to monitor the startup to ensure the Morpheus Application is starting without error:

```
[root@app-server-1 ~]# morpheus-ctl tail morpheus-ui
```

Recovery Thoughts/Further Discussion: If Morpheus UI cannot connect to RabbitMQ, Elasticsearch or the database tier it will fail to start. The Morpheus UI logs can indicate if this is the case.

Aside from RabbitMQ, there can be issues with false positives concerning Elasticsearch's running status. The biggest challenge with Elasticsearch, for instance, is that a restarted node has trouble joining the ES cluster. This is fine in the case of ES, though, because the `minimum_master_nodes` setting will not allow the un-joined singleton to be consumed until it joins. Morpheus will still start if it can reach the other two ES hosts, which are still clustered.

The challenge with RabbitMQ is that it is load balanced behind Morpheus for requests, but each Morpheus application server needs to bootstrap the RabbitMQ tied into it. Thus, if it cannot reach its own RabbitMQ startup for it will fail.

Similarly, if a Morpheus UI service cannot reach the database, startup will fail. However, if the database is externalized and failover is configured for Master/Master, then there should be ample opportunity for Morpheus to connect to the database tier.

Because Morpheus can start even though the Elasticsearch node on the same host fails to join the cluster, it is advisable to investigate the health of ES on the restarted node after the services are up. This can be done by accessing the endpoint with curl and inspecting the output. The status should be "green" and number of nodes should be "3":

```
[root@app-server-1 ~]# curl localhost:9200/_cluster/health?pretty=true
{
  "cluster_name" : "morpheus",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 3,
  "number_of_data_nodes" : 3,
  "active_primary_shards" : 110,
  "active_shards" : 220,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
```

(continues on next page)

(continued from previous page)

```
"number_of_in_flight_fetch" : 0
}
```

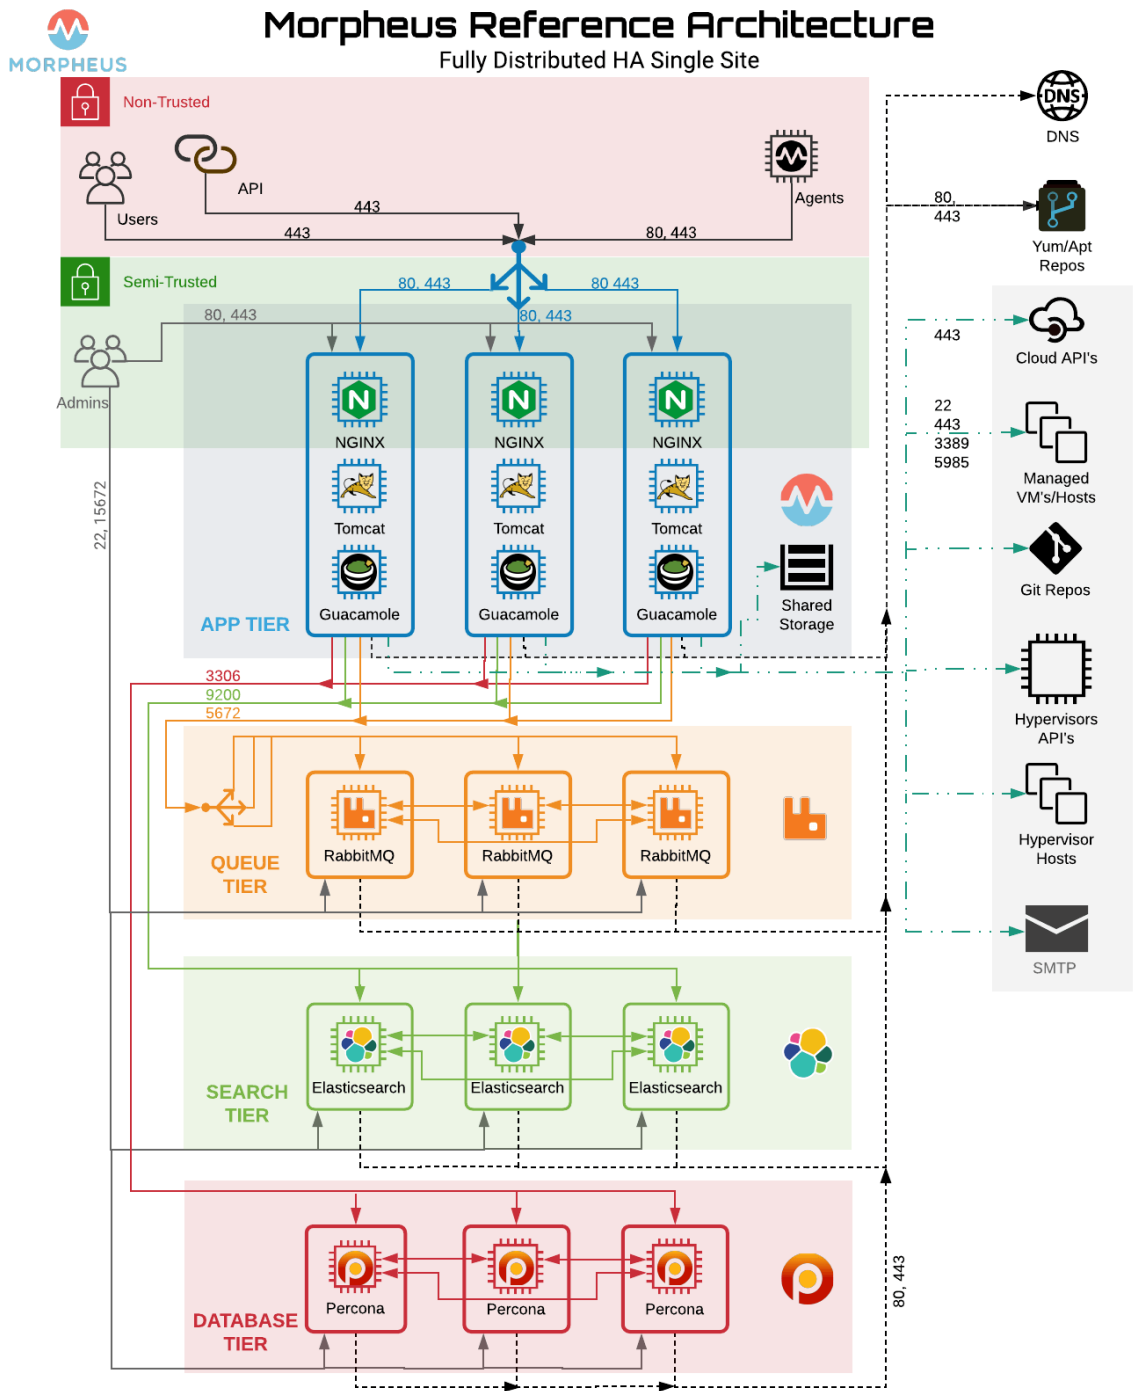
If this is not the case it is worth investigating the Elasticsearch logs to understand why the singleton node is having trouble joining the cluster. These can be found at `/var/log/morpheus/elasticsearch/current`

Outside of these stateful tiers, the “morpheus-ctl status” command will not output a “run” status unless the service is successfully running. If a stateless service reports a failure to run, the logs should be investigated and/or sent to Morpheus for additional support. Logs for all Morpheus embedded services are found in `/var/log/morpheus`.

Full HA Install

Full HA Install Overview

- **App Host(s) with Distributed Services (Full HA)** Application tier is installed on one or more hosts. All UI hosts point to externalized Transactional Database, Non-Transactional Database, and Message Tiers. The reconfigure process installs only Application services.



Minimum Nodes

For Full High-Availability configurations, RabbitMQ, Elasticsearch and mySQL(Galera/Percona) must be configured in minimum 3 Node Clusters, and 2 or more App Nodes are required.

Note: VM requirements assume local services. VM count requirements are not applicable when using hosted services such as AWS RDS mySQL.

Minimum 11 Nodes

- 2+ Application Hosts
- 3 Node RabbitMQ Cluster
- 3 Node Elasticsearch Cluster
- 3 Node Galera/Percona Cluster

Important: Asynchronous Active/Active and Active/Passive Database configurations are not supported for HA configurations. A minimum 3 node mySQL Cluster with synchronous multi-master replication is required for Database Clusters. Morpheus recommends Percona XtraDB Clusters with synchronous multi-master replication. Asynchronous Active/Passive can be used but is not considered an HA configuration.

Important: For Clusters with more than 3 Nodes, always use an odd number of nodes (3,5,7 etc) to ensure Quorum.

Shared Storage

For configurations with 2 or more Applications Nodes, Shared Storage is required between the app nodes for `/var/opt/morpheus/morpheus-ui/*`. Local Storage File Shares will need to be copied to a shared file system so all assets are available on all App nodes.

Shared Assets

- Logos
- Uploaded Virtual Images
- Deployment Uploads
- Ansible
- Terraform
- Morpheus Backups

Note: Backups, deployment and virtual image storage locations can be overridden within the Morpheus-ui.

Port Requirements

Service	Source	Destination	Port(s)
Morpheus	Application Node	mySQL	3306
Morpheus	Application Node	Elasticsearch	9200; 9300
Morpheus	Application Node	RabbitMQ	5672; 61613
Morpheus	Application Node	YUM or APT	443; 80
Elasticsearch	Elasticsearch	Elasticsearch	9200; 9300
mySQL	mySQL	mySQL	3306;4444;4567;4560
RabbitMQ	RabbitMQ	RabbitMQ	5672 or 5671(SSL); 61613 or 61614(SSL)

Default Locations

Morpheus follows several install location conventions. Below is a list of system defaults for convenient management:

- Installation Location: `/opt/morpheus`
- Log Location: `/var/log/morpheus`
 - Morpheus-UI: `/var/log/morpheus/morpheus-ui`
 - NGINX: `/var/log/morpheus/nginx`
 - Check Server: `/var/log/morpheus/check-server`
- User-defined install/config: `/etc/morpheus/morpheus.rb`

Percona XtraDB Cluster

Out of the box Morpheus uses MySQL but Morpheus supports any mySQL compliant database. There are many ways to set up a highly available, MySQL dialect based database. One which has found favor with many of our customers is Percona's XtraDB Cluster. Percona's product is based off of Galera's WSREP Clustering, which is also supported.

Important: Additional configuration for Percona Clusters with TLS enabled is required. Refer to Percona TLS Configuration for details.

Requirements

Note: Morpheus idiomatically connects to database nodes over 3306

Once you have your database installed and configured:

1. Create the Database you will be using with morpheus.

```
mysql> CREATE DATABASE morpheus CHARACTER SET utf8 COLLATE utf8_general_ci;

mysql> show databases;
```

2. Next create your morpheus database user. The user needs to be either at the IP address of the morpheus application server or use `@' % '` within the user name to allow the user to login from anywhere.

```
mysql> CREATE USER '$morpheus_db_user_name'@$source_ip IDENTIFIED BY '$morpheus_
↳db_user_pw';
```

3. Next Grant your new morpheus user permissions to the database.

```
mysql> GRANT ALL PRIVILEGES ON morpheus_db_name.* TO 'morpheus_db_user'@$source_
↳ip' IDENTIFIED BY 'morpheus_db_user_pw' with grant option;

mysql> GRANT SELECT, PROCESS, SHOW DATABASES, SUPER ON *.* TO 'morpheus_db_user'@
↳'$source_ip' IDENTIFIED BY 'morpheus_db_user_pw';

mysql> FLUSH PRIVILEGES;
```

4. Checking Permissions for your user.

```
SHOW GRANTS FOR '$morpheus_db_user_name'@$source_ip';
```

Percona XtraDB Cluster with TLS

Installation and configuration of Percona XtraDB Cluster on CentOS/RHEL 7 with TLS enabled for all comms.

Important: This is a sample configuration only. Customer configurations and requirements will vary.

Requirements

Percona requires the following ports for the cluster nodes. Please create the appropriate firewall rules on your Percona nodes.

- 3306
- 4444
- 4567
- 4568

Configure SELinux

When SELinux is set to `Enforcing`, by default it will block Percona Cluster communication.

To allow Percona XtraDB Cluster functionality when SELinux is `Enforcing`, run the following on each Database Node:

1. Install SELinux utilities

```
[root]# yum install -y policycoreutils-python.x86_64
```

2. Configure Percona ports for SELinux:

```
[root]# semanage port -m -t mysqld_port_t -p tcp 4444
[root]# semanage port -m -t mysqld_port_t -p tcp 4567
[root]# semanage port -a -t mysqld_port_t -p tcp 4568
```

3. Create the policy file PXC.te

```
[root]# vi PXC.te

require {
    type unconfined_t;
    type mysqld_t;
    type unconfined_service_t;
    type tmp_t;
    type sysctl_net_t;
    type kernel_t;
    type mysqld_safe_t;
    class process { getattr setpgid };
    class unix_stream_socket connectto;
    class system module_request;
    class file { getattr open read write };
    class dir search;
}

#===== mysqld_t =====

allow mysqld_t kernel_t:system module_request;
allow mysqld_t self:process { getattr setpgid };
allow mysqld_t self:unix_stream_socket connectto;
allow mysqld_t sysctl_net_t:dir search;
allow mysqld_t sysctl_net_t:file { getattr open read };
allow mysqld_t tmp_t:file write;
```

4. Compile and load the SELinux policy

```
[root]# checkmodule -M -m -o PXC.mod PXC.te
[root]# semodule_package -o PXC.pp -m PXC.mod
[root]# semodule -i PXC.pp
```

Add Percona Repo

1. Add the percona repo to your Linux Distro.

```
[root]# wget https://www.percona.com/downloads/RPM-GPG-KEY-percona && rpm --
→import RPM-GPG-KEY-percona

[root]# yum install -y https://repo.percona.com/yum/percona-release-latest.noarch.
→rpm
```

2. The below commands will clean the repos and update the server.

```
[root]# yum clean all
[root]# yum update -y --skip-broken
```

Installing Percona XtraDB Cluster

1. Install the Percona XtraDB Cluster software and its dependencies.

```
[root]# yum install -y Percona-XtraDB-Cluster-57
```

2. Enable the mysql service so that the service started at boot.

```
[root]# systemctl enable mysql
```

3. Start mysql

```
[root]# systemctl start mysql
```

4. Log into the mysql server and set a new password. To get the temporary root mysql password you will need to run the below command. The command will print the password to the screen. Copy the password.

```
[root]# grep 'temporary password' /var/log/mysqld.log
```

5. Login to mysql

```
[root]# mysql -u root -p
password: `enter password copied above`
```

6. Change the root user password to the mysql db

```
mysql> ALTER USER 'root'@'localhost' IDENTIFIED BY 'rootPassword';
```

7. Create the sstuser user and grant the permissions.

```
mysql> CREATE USER 'sstuser'@'localhost' IDENTIFIED BY 'sstUserPassword';
```

Note: The sstuser and password will be used in the /etc/my.cnf configuration.

```
mysql> GRANT RELOAD, LOCK TABLES, PROCESS, REPLICATION CLIENT ON *.* TO 'sstuser'@
↪ 'localhost';

mysql> FLUSH PRIVILEGES;
```

8. Exit mysql then stop the mysql services:

```
mysql> exit
Bye
[root]# systemctl stop mysql.service
```

9. Install Percona on to the other nodes using the same steps.

Once the service is stopped on all nodes move onto the next step.

Add [mysqld] to my.cnf in /etc/

1. Add the following to /etc/my.cnf. The node_name and node_address needs to be unique on each of the nodes.

Node 01:

```
[root]# vi /etc/my.cnf
```

```
[mysqld]
pxc_encrypt_cluster_traffic=ON
max_connections = 300
wsrep_provider=/usr/lib64/galera3/libgalera_smm.so

wsrep_cluster_name=morpheusdb-cluster
wsrep_cluster_address=gcomm://10.30.20.10,10.30.20.11,10.30.20.12

# for wsrep_cluster_address=gcomm://Enter the IP address of the primary_
↪node first then remaining nodes. Separating the ip addresses with commas

wsrep_node_name=morpheus-node01
wsrep_node_address=10.30.20.10

wsrep_sst_method=xtrabackup-v2
wsrep_sst_auth=sstuser:sstUserPassword
pxc_strict_mode=PERMISSIVE
wsrep_sync_wait=2

skip-log-bin
default_storage_engine=InnoDB
innodb_autoinc_lock_mode=2
```

Node 02

```
$ [root]# vi /etc/my.cnf
```

```
[mysqld]
pxc_encrypt_cluster_traffic=ON
max_connections = 300
wsrep_provider=/usr/lib64/galera3/libgalera_smm.so

wsrep_cluster_name=morpheusdb-cluster
wsrep_cluster_address=gcomm://10.30.20.10,10.30.20.11,10.30.20.12

# for wsrep_cluster_address=gcomm://Enter the IP address of the primary_
↪node first then remaining nodes. Separating the ip addresses with commas

wsrep_node_name=morpheus-db-node02
wsrep_node_address=10.30.20.11

wsrep_sst_method=xtrabackup-v2
wsrep_sst_auth=sstuser:sstUserPassword
pxc_strict_mode=PERMISSIVE
wsrep_sync_wait=2

skip-log-bin
default_storage_engine=InnoDB
```

(continues on next page)

(continued from previous page)

```
innodb_autoinc_lock_mode=2
```

Node 03

```
$ [root]# vi /etc/my.cnf
```

```
[mysqld]
pxc_encrypt_cluster_traffic=ON
max_connections = 300
wsrep_provider=/usr/lib64/galera3/libgalera_smm.so

wsrep_cluster_name=morpheusdb-cluster
wsrep_cluster_address=gcomm://10.30.20.10,10.30.20.11,10.30.20.12

# for wsrep_cluster_address=gcomm://Enter the IP address of the primary_
↪node first then remaining nodes. Separating the ip addresses with commas

wsrep_node_name=morpheus-node03
wsrep_node_address=10.30.20.12

wsrep_sst_method=xtrabackup-v2
wsrep_sst_auth=sstuser:sstUserPassword
pxc_strict_mode=PERMISSIVE
wsrep_sync_wait=2

skip-log-bin
default_storage_engine=InnoDB
innodb_autoinc_lock_mode=2

.. note:: The default setting on [morpheus] app nodes for ``max_active``_
↪database connections is 100. For this example we are setting ``max_
↪connections = 300`` to account for 3 maximum simultaneous morpheus app_
↪node connections. If ``max_active`` is configured higher on the app_
↪nodes, or the number of app nodes is not 3, adjust accordingly for your_
↪configuration.
```

2. Save /etc/my.cnf

Bootstrap Node 01

Important: Ensure mysql.service is stopped prior to bootstrap.

1. To bootstrap the first node in the cluster run the below command.

```
systemctl start mysql@bootstrap.service
```

Note: The mysql service will start during the bootstrap.

Note: Startup failures are commonly caused by misconfigured /etc/my.cnf files. Also verify safe_to_bootstrap is set to 1 on Node 01 in /var/lib/mysql/grastate.dat.

Configure Morpheus Database and User

1. Create the Database you will be using with morpheus.

Login to mysql on Node 01:

```
mysql -u root -p
password:

mysql> CREATE DATABASE morpheus CHARACTER SET utf8 COLLATE utf8_general_ci;

mysql> show databases;
```

1. Next create your morpheus database user. This is the user the morpheus app nodes will auth with mysql.

```
mysql> CREATE USER 'morpheusDbUser'@'%' IDENTIFIED BY 'morpheusDbUserPassword';
```

2. Next Grant your new morpheus user permissions.

```
mysql> GRANT ALL PRIVILEGES ON *.* TO 'morpheusDbUser'@'%' IDENTIFIED BY
↳ 'morpheusDbUserPassword';

mysql> FLUSH PRIVILEGES;

.. important:: If you grant privileges to the morpheusDbUser to only the
↳ morpheusdb database, you will also need to GRANT SELECT, PROCESS, SHOW
↳ DATABASES, SUPER ON PRIVILEGES to the morpheusDbUser on *.* for the Appliance
↳ Health service.

mysql> exit
```

Copy SSL Files to other nodes

During initialization of Node 01 the required *pem* files will be generated in `/var/lib/mysql`. The `ca.pem`, `server-cert.pem` and `server-key.pem` files need to match on all nodes in the cluster.

1. Copy the following files from Node 01 to the same path (default is `/var/lib/mysql`) on Node 02 and Node 03:

```
/var/lib/mysql/ca.pem
/var/lib/mysql/server-cert.pem
/var/lib/mysql/server-key.pem

.. important:: Ensure all 3 files match on all 3 nodes, including path, owner and
↳ permissions.

.. note:: The generated certificate is self signed. Consult Percona documentation
↳ for [mysqld] and SSL file configuration when providing your own.
```

Start the Remaining Nodes

1. Start mysql on Node 02 and Node 03

```
[root]# systemctl start mysql.service
```

The services will automatically join the cluster using the sstuser we created earlier.

Note: Startup failures are commonly caused by misconfigured `/etc/my.cnf` files.

Verify Configuration

1. Verify SELinux is not rejecting any db cluster communication by running the below on all db nodes:

```
[root@allDbNodes]# grep -i denied /var/log/audit/audit.log | grep mysqld_t
```

If there are any results, address the source or update the SELinux Policy to resolve.

1. To verify all nodes joined the cluster, on any db node login to mysql and run `show status like 'wsrep%';`

```
[root@anyDbNode]# mysql -u root -p
mysql> show status like 'wsrep%';
```

2. Verify `wsrep_cluster_size` is 3 and `wsrep_incoming_addresses` lists all 3 node ip addresses.
3. From all Morpheus app nodes, verify that you can login to all 3 database nodes

```
[root@allAppNodes] cd
[root@appNode01]# ./mysql -u morpheusDbUser -p -h 10.30.20.10
[root@appNode02]# ./mysql -u morpheusDbUser -p -h 10.30.20.11
[root@appNode03]# ./mysql -u morpheusDbUser -p -h 10.30.20.12
```

If you are unable to login to mysql from an app node, ensure credentials are correct, privileges have been granted, and mysql is running.

To validate network accessibility, use telnet to verify app node can reach db nodes on 3306: `telnet 10.30.20.10 3306`

RabbitMQ Cluster

An HA deployment will also include a Highly Available RabbitMQ. This can be achieved through RabbitMQ's HA-Mirrored Queues on at least 3, independent nodes. To accomplish this we recommend following Pivotal's documentation on RabbitMQ here: <https://www.rabbitmq.com/ha.html> and <https://www.rabbitmq.com/clustering.html>

Install RabbitMQ on the 3 nodes and create a cluster.

Note: For the most up to date RPM package we recommend using this link: :link: <https://www.rabbitmq.com/install-rpm.html#downloads>

Important: Morpheus connects to AMQP over 5672 or 5671(SSL) and 61613 or 61614(SSL)

RabbitMQ Installation and Configuration

Important: This is a sample configuration only. Customer configurations and requirements will vary.

1. Install epel-release and erlang

```
yum install epel-release
yum install erlang
```

2. Install RabbitMQ on all 3 Nodes

```
wget https://dl.bintray.com/rabbitmq/rabbitmq-server-rpm/rabbitmq-server-3.6.12-1.
↪el7.noarch.rpm

rpm --import https://www.rabbitmq.com/rabbitmq-release-signing-key.asc

yum -y install rabbitmq-server-3.6.12-1.el7.noarch.rpm

chkconfig rabbitmq-server on

rabbitmq-server -detached
```

3. Copy the erlang.cookie from Node 1

```
cat /var/lib/rabbitmq/.erlang.cookie
```

Copy the .erlang.cookie value

4. Overwrite /var/lib/rabbitmq/.erlang.cookie on Nodes 2 & 3 with value from Node 1 and change its permissions using the follow commands:

```
chown rabbitmq:rabbitmq /var/lib/rabbitmq/*
chmod 400 /var/lib/rabbitmq/.erlang.cookie
```

5. Edit /etc/hosts file on all 3 nodes to refer to shortnames of the other nodes

Example for node 1 (adjust for nodes 2 and 3):

```
vi /etc/hosts

10.30.20.101 rabbit-2
10.30.20.102 rabbit-3
```

6. Run the following commands on Node 2 and on Node 3 to join them to the Cluster:

```
rabbitmqctl stop
rabbitmq-server -detached
rabbitmqctl stop_app
rabbitmqctl join_cluster rabbit@<<node 1 shortname>>
rabbitmqctl start_app
```

7. On Node 1, create vhost and add Admin user for Morpheus

```
rabbitmqctl add_vhost morpheus
rabbitmqctl add_user <<admin username>> <<password>>
rabbitmqctl set_permissions -p morpheus <<admin username>> ".*" ".*" ".*"
rabbitmqctl set_user_tags <<admin username>> administrator
```

8. On All Nodes, enable stomp and management plugins:

```
rabbitmq-plugins enable rabbitmq_stomp
rabbitmq-plugins enable rabbitmq_management
```

9. On Node 1, add the required Rabbitmq Policies. The policies will propagate to all nodes.

```
rabbitmqctl set_policy -p morpheus --apply-to queues --priority 2 statCommands
↪ "statCommands.*" '{"expires":1800000, "ha-mode":"all"}'
rabbitmqctl set_policy -p morpheus --apply-to queues --priority 2 _
↪ morpheusAgentActions "morpheusAgentActions.*" '{"expires":1800000, "ha-mode":
↪ "all"}'
rabbitmqctl set_policy -p morpheus --apply-to queues --priority 2 monitorJobs
↪ "monitorJobs.*" '{"expires":1800000, "ha-mode":"all"}'
rabbitmqctl set_policy -p morpheus --apply-to all --priority 1 ha ".*" '{"ha-mode
↪ ":"all"}'
```

Elasticsearch

Sample Install of 3 node Elasticsearch Cluster on CentOS 7

Important: This is a sample configuration only. Customer configurations and requirements will vary.

Important: Morpheus v4.1.2+ requires Elasticsearch v7.x.

Requirements

1. Three Existing CentOS 7+ nodes accessible to the Morpheus Appliance
2. Install Java on each node

You can install the latest OpenJDK with the command:

```
sudo yum install java-1.8.0-openjdk.x86_64
```

To verify your JRE is installed and can be used, run the command:

```
java -version
```

The result should look like this:

```
Output of java -version
openjdk version "1.8.0_65"
OpenJDK Runtime Environment (build 1.8.0_65-b17)
OpenJDK 64-Bit Server VM (build 25.65-b01, mixed mode)
```

Install Elasticsearch 7.x

Important: This is an example Elasticsearch Upgrade for reference only, and is not indicative of the upgrade procedure for every environment/user/customer/configuration.

1. On each ES node run the following to install Elasticsearch.

```
rpm --import https://artifacts.elastic.co/GPG-KEY-elasticsearch
```

```
wget https://artifacts.elastic.co/downloads/elasticsearch/elasticsearch-7.6.2-x86_64.rpm
```

```
sudo rpm -Uvh elasticsearch-7.6.2-x86_64.rpm
```

2. If necessary, update permissions for the specified log and data paths

```
sudo chown -R elasticsearch:elasticsearch /var/log/elasticsearch/
sudo chown -R elasticsearch:elasticsearch /usr/share/elasticsearch/
```

3. Edit `/etc/elasticsearch/elasticsearch.yml` and update each nodes configurations accordingly. Please note several attributes differ in 7.x from 5.x.

Node 1 Example (customer configurations will vary)

```
sudo vi /etc/elasticsearch/elasticsearch.yml

#Sample elasticsearch.yml config. Adjusting values in elasticsearch.yml for
each node in the cluster.
#Note: Sample only, user configurations and requirements will vary.

node.name: "es-node-01" ##unique name of this node
network.host: 10.30.22.152 ##ip of this node
http.port: 9200
discovery.seed_hosts: ["10.30.22.152","10.30.22.153","10.30.22.154"] ## add
all cluster node ip's
cluster.initial_master_nodes: ["10.30.22.152","10.30.22.153","10.30.22.154"] ## add all cluster node ip's
path.logs: /var/log/elasticsearch ## Or your preferred location.
path.data: /usr/share/elasticsearch/ ## Or your preferred location.
discovery.zen.minimum_master_nodes: 2
```

Node 2 Example (customer configurations will vary)

```
sudo vi /etc/elasticsearch/elasticsearch.yml

#Sample elasticsearch.yml config. Adjusting values in elasticsearch.yml for
each node in the cluster.
#Note: Sample only, user configurations and requirements will vary.

node.name: "es-node-02" ##unique name of this node
network.host: 10.30.22.153 ##ip of this node
http.port: 9200
discovery.seed_hosts: ["10.30.22.152","10.30.22.153","10.30.22.154"] ## add
all cluster node ip's
cluster.initial_master_nodes: ["10.30.22.152","10.30.22.153","10.30.22.154"] ## add all cluster node ip's
```

(continues on next page)

(continued from previous page)

```
path.logs: /var/log/elasticsearch ## Or your preferred location.
path.data: /usr/share/elasticsearch/ ## Or your preferred location.
discovery.zen.minimum_master_nodes: 2
```

Node 3 Example (customer configurations will vary)

```
sudo vi /etc/elasticsearch/elasticsearch.yml

#Sample elasticsearch.yml config. Adjusting values in elasticsearch.yml for
↪each node in the cluster.
#Note: Sample only, user configurations and requirements will vary.

node.name: "es-node-03" ##unique name of this node
network.host: 10.30.22.154 ##ip of this node
http.port: 9200
discovery.seed_hosts: ["10.30.22.152","10.30.22.153","10.30.22.154"] ## add
↪all cluster node ip's
cluster.initial_master_nodes: ["10.30.22.152","10.30.22.153","10.30.22.154
↪"] ## add all cluster node ip's
path.logs: /var/log/elasticsearch ## Or your preferred location.
path.data: /usr/share/elasticsearch/ ## Or your preferred location.
discovery.zen.minimum_master_nodes: 2
```

4. Save elasticsearch.yml**5. Start Elasticsearch on each node.**

```
sudo service elasticsearch start
```

6. Verify cluster health

```
curl http://localhost:9200/_cluster/health

or

curl http://node_ip:9200/_cluster/health
```

Application Tier

Morpheus configuration is controlled by a configuration file located at `/etc/morpheus/morpheus.rb`. This file is read when you run `morpheus-ctl reconfigure` after installing the appliance package. Each section is tied to a deployment tier: database is mysql, message queue is rabbitmq, search index is elasticsearch. There are no entries for the web and application tiers since those are part of the core application server where the configuration file resides.

1. Download and install the Morpheus Appliance Package
2. Next we must install the package onto the machine and configure the morpheus services:

```
sudo rpm -i morpheus-appliance-x.x.x-1.x86_64.rpm
```

3. After installing and prior to reconfiguring, edit the `morpheus.rb` file

```
sudo vi /etc/morpheus/morpheus.rb
```

Change the values to match your configured services:

Note: The values below are examples. Update hosts, ports, usernames and password with your specifications. Only include entries for services you wish to externalize.

```
mysql['enable'] = false
mysql['host'] = {'10.30.20.139' => 3306, '10.30.20.153' => 3306, '10.30.20.196' =>
↳3306}
mysql['morpheus_db'] = 'morpheusdb'
mysql['morpheus_db_user'] = 'dbuser'
mysql['morpheus_password'] = 'dbuserpassword'
rabbitmq['enable'] = false
rabbitmq['vhost'] = 'morpheus'
rabbitmq['queue_user'] = 'lbuser'
rabbitmq['queue_user_password'] = 'lbuserpassword'
rabbitmq['host'] = 'rabbitvip'
rabbitmq['port'] = '5672'
rabbitmq['heartbeat'] = 50
elasticsearch['enable'] = false
elasticsearch['cluster'] = 'esclustername'
elasticsearch['es_hosts'] = {'10.30.20.91' => 9200, '10.30.20.149' => 9200, '10.30.20.
↳165' => 9200}
elasticsearch['use_tls'] = true
elasticsearch['auth_user'] = 'morpheus-user'
elasticsearch['auth_password'] = 'xxxxxxxxxxxxxxxxxx'
```

4. Reconfigure Morpheus

```
sudo morpheus-ctl reconfigure
```

Shared Storage

For configurations with 2 or more Applications Nodes, Shared Storage is required between the app nodes. Local Storage File Shares will need to be copied to a shared file system so all assets are available on all App nodes.

Assets

- White label images
- Uploaded virtual images
- Deploy uploads
- Ansible Plays
- Terraform
- Morpheus backups

Tip: Backups, deployments and virtual image storage locations can be overridden within the Morpheus-ui. You can find more information on storage here: [Storage](#)

To copy the `morpheus-ui` directory to the shared storage follow the below steps:

1. SSH into the Appliance

2. `sudo su` (or login as root)
3. `cd` into `/var/opt/morpheus/`
4. Backup morpheus-ui directory by running the command below. This will create a new directory in `/var/opt/morpheus/` called `morpheus-ui-bkp` and copy the contents of `morpheus-ui` into the new directory

```
cp -r morpheus-ui morpheus-ui-bkp
```

5. Move morpheus-ui to your shared storage. Example below:

```
mv morpheus-ui /nfs/appliance-files/
```

6. Mount your shared storage volume to `/var/opt/morpheus/morpheus-ui`. How you mount it is dependent on what kind of storage it is. If you mount the volume after the package install, but before the reconfigure, then you don't need to copy anything to a backup.
7. SSH into the second Appliance and then Backup morpheus-ui directory by running

```
cp -r morpheus-ui morpheus-ui-bkp
```

Tip: when adding additional nodes you will only need to run step 6 and 7

Important: NFS mounts require `sync` option when using Ansible integration with Morpheus Agent command bus execution enabled.

Upgrades & Maintenance

Upgrading

Upgrading Overview

Important: Morpheus v4.1.2+ requires Elasticsearch 7.x. Earlier versions of Morpheus ran against Elasticsearch v5.x.

- The Elasticsearch version for Appliance configurations with the default local Elasticsearch target will automatically be upgraded and no manual upgrade is required.
 - For Appliance configurations with an existing external Elasticsearch service, an upgrade of Elasticsearch to v7.x is required to upgrade Morpheus to v4.1.2+.
 - Morpheus can also be pointed to a new Elasticsearch 7.x cluster or service as an alternate to upgrading an existing cluster.
 - Elasticsearch data will not be retained during a direct 5.x to 7.x upgrade by default. Please refer to Elasticsearch documentation if backing up and restoring your 5.x Elasticsearch Morpheus data is required.
 - If log and stat data stored in Elasticsearch is not critical, a 5.x Elasticsearch data backup and restoration to 7.x, or a 5.x -> 6.x -> 7.x rolling upgrade is not necessary as Morpheus will rebuild the indices upon connection to the 7.x cluster.
 - Please refer to [Elasticsearch Upgrade Documentation](#) before installing or upgrading to v4.1.2 if your Appliance's Elasticsearch is external.
-

Morpheus Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

Upgrade Requirements

3.6.x to v5.2.0 Upgrade

- Only appliances running Morpheus v3.6.0 or higher can upgrade to 4.x.
- MySQL will be upgraded to 5.7.x on Appliances with MySQL running on the app node (Single Node or “all-in-one” Appliances). Backup your database before running the upgrade.

Important: BACKUP YOUR DATABASE before the upgrade. You can use the appliance backup job in Morpheus, but make sure you download it before you do the upgrade.

- RabbitMQ will be upgraded to v3.7 on Appliances with RabbitMQ running on the app node (Single Node or “all-in-one” Appliances). On 3-Node configurations, the RabbitMQ queues and configuration will be dropped and the cluster will need to be configured and established again.
- Elasticsearch will be upgraded from 5.x to 7.x. Refer to [Elasticsearch Upgrade Documentation](#) for upgrading external ES Clusters.
- Stop all morpheus services, not just the morpheus-ui, before the upgrade. Although the upgrade process will also stop the services, take this step to ensure they are stopped.
- Warnings about missing files during the removal phase are expected and can be ignored.
- For firewall/proxy/acl considerations, the domain for Appliance, Supplemental and Agent packages was changed recently to <https://downloads.morpheusdata.com> from <https://downloads.gomorpheus.com>. Please update ACL's to allow access to <https://downloads.morpheusdata.com> when necessary.

Refer to [v5.2.0 Compatibility & Breaking Changes](#) for externalized MySQL, Elasticsearch and/or RabbitMQ version requirements.

4.0.0, 4.1.0, 4.1.1 to v5.2.0 Upgrade

- Elasticsearch will be upgraded from 5.x to 7.x. Refer to [Elasticsearch Upgrade Documentation](#) for upgrading external ES Clusters.
- For firewall/proxy/acl considerations, the domain for Appliance, Supplemental and Agent packages was changed recently to <https://downloads.morpheusdata.com> from <https://downloads.gomorpheus.com>. Please update ACL's to allow access to <https://downloads.morpheusdata.com> when necessary.

4.1.2+ to v5.2.0 Upgrade

No major service changes

Single Node Appliance Upgrade

The following covers upgrading single node (All-In-One) Appliance configurations.

Important: Only appliances running Morpheus v4.2.0 or higher can upgrade to 4.x. Always backup your database before running any upgrade.

4.2.0 to v5.2.0 Upgrade

Debian / Ubuntu

To upgrade Morpheus running on Ubuntu/Debian, download new deb package, stop the morpheus-ui, install the new deb package, then reconfigure.

```
sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_amd64.deb
sudo morpheus-ctl stop morpheus-ui
sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
sudo morpheus-ctl reconfigure
```

Note: Services will be stopped during package installation and started during the reconfigure process, including the Morpheus-ui service.

All services will automatically start during the reconfigure process. After the reconfigure has succeeded, tail the ui service to watch ui startup logs with `morpheus-ctl tail morpheus-ui`.

After the morpheus-ui service finishes loading, the upgrade is complete.

CentOS / RHEL

To upgrade Morpheus running on CentOS/RHEL, download and install the new rpm package, stop the morpheus-ui, reconfigure and then start the morpheus-ui:

```
sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
sudo morpheus-ctl stop morpheus-ui
sudo rpm -Uhv morpheus-appliance-x.x.x-x.x86_64.rpm
sudo morpheus-ctl reconfigure
```

Note: Services will be stopped during package installation and started during the reconfigure process, including the Morpheus-ui service.

All services will automatically start during the reconfigure process. After the reconfigure has succeeded, tail the ui service to watch ui startup logs with `morpheus-ctl tail morpheus-ui`.

After the morpheus-ui service finishes loading, the upgrade is complete.

3-Node HA Upgrade

3-Node HA Appliance represent 3 App nodes with local RabbitMQ and Elasticsearch services clustered across the app nodes, and an external Galera/Percona MySQL cluster.

Morpheus Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

Refer to *v5.2.0 Compatibility & Breaking Changes* for any 3-node variations using externalized MySQL, Elasticsearch and/or RabbitMQ version requirements.

4.2.x -> v5.2.0 upgrade

Due to Database schema changes in v5.2.0 it is important to stop the morpheus-ui service on all app nodes prior to upgrade. Failure to do so may result in errors or database corruption.

Upgrade Instructions

3-Node HA Debian / Ubuntu Upgrade

The following covers upgrading the Morpheus App nodes in 3 Node HA configurations to v5.2.0.

Warning: As a best practice, always backup your database prior to any upgrade.

Important: The following is only for “3 Node HA” Architecture configurations.

Morpheus Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

4.2.0+ -> v5.2.0 Upgrade

Important: Due to Database schema changes in v5.2.0 it is important to stop the morpheus-ui service on all app nodes prior to upgrade. Failure to do so may result in errors or database corruption.

1. Starting with Node 3, on **All** App Nodes, stop the morpheus-ui services via `morpheus-ctl stop morpheus-ui`. If you receive a timeout, run `morpheus-ctl graceful-kill morpheus-ui`.

```
[root@app-server-3 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-2 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-1 ~]# morpheus-ctl stop morpheus-ui
```

2. Upgrade the deb package on Node 1, then run a Reconfigure on Node 1

```
[root@app-server-1 ~]# sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_
↪amd64.deb
[root@app-server-1 ~]# sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
[root@app-server-1 ~]# sudo morpheus-ctl reconfigure
```

Note: All services will automatically stopped and started during the reconfigure process. After the reconfigure has succeeded, tail the ui service to watch ui startup logs with `morpheus-ctl tail morpheus-ui`.

3. Once Node 1 upgrade has completed and the ui is available, upgrade the deb package on Node 2, then run a Reconfigure on Node 2.

```
[root@app-server-2 ~]# sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_
↪amd64.deb
[root@app-server-2 ~]# sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
[root@app-server-2 ~]# sudo morpheus-ctl reconfigure
```

4. Then upgrade the deb package on Node 3, and run a Reconfigure on Node 3

```
[root@app-server-3 ~]# sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_
↪amd64.deb
[root@app-server-3 ~]# sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
[root@app-server-3 ~]# sudo morpheus-ctl reconfigure
```

5. The upgrade is complete and the Morpheus-ui services should be running with clustered Elasticsearch and RabbitMQ services across the 3 nodes.

Important: If reconfigure after a rpm package upgrade stalls or hangs on starting a service (mysql, rabbitmq, elastic-search ...) it is possible the `morpheus-runsvdir` service did not start or a process it was managing was manually shutdown or killed. To resolve, run `systemctl stop morpheus-runsvdir` then `systemctl start morpheus-runsvdir`, then run `reconfigure` again, `morpheus-ctl reconfigure`.

3-Node HA CentOS / RHEL Upgrade

The following covers upgrading the Morpheus App nodes in 3 Node HA configurations to v5.2.0.

Warning: As a best practice, always backup your database prior to any upgrade.

Important: The following is only for “3 Node HA” Architecture configurations.

Morpheus Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

4.2.0+ -> v5.2.0 Upgrade

Important: Due to Database schema changes in v5.2.0 it is important to stop the morpheus-ui service on all app nodes prior to upgrade. Failure to do so may result in errors or database corruption.

1. Starting with Node 3, on **All App Nodes**, stop the morpheus-ui services via `morpheus-ctl stop morpheus-ui`. If you receive a timeout, run `morpheus-ctl graceful-kill morpheus-ui`.

```
[root@app-server-3 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-2 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-1 ~]# morpheus-ctl stop morpheus-ui
```

2. Upgrade the RPM package on Node 1, then run a Reconfigure on Node 1

```
[root@app-server-1 ~]# sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-1 ~]# sudo rpm -Uvh morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-1 ~]# sudo morpheus-ctl reconfigure
```

Note: All services will automatically stopped and started during the reconfigure process. After the reconfigure has succeeded, tail the ui service to watch ui startup logs with `morpheus-ctl tail morpheus-ui`.

3. Once Node 1 upgrade has completed and the u is available, upgrade the RPM package on Node 2, then run a Reconfigure on Node 2.

```
[root@app-server-2 ~]# sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-2 ~]# sudo rpm -Uvh morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-2 ~]# sudo morpheus-ctl reconfigure
```

4. Then upgrade the RPM package on Node 3, then run a Reconfigure on Node 3

```
[root@app-server-3 ~]# sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-3 ~]# sudo rpm -Uhv morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-3 ~]# sudo morpheus-ctl reconfigure
```

5. The upgrade is complete and the Morpheus-ui services should be running with clustered Elasticsearch and RabbitMQ services across the 3 nodes.

Important: If reconfigure after a rpm package upgrade stalls or hangs on starting a service (mysql, rabbitmq, elasticsearch ...) it is possible the `morpheus-runsvdir` service did not start or a process it was managing was manually shutdown or killed. To resolve, run `systemctl stop morpheus-runsvdir` then `systemctl start morpheus-runsvdir`, then run `reconfigure` again, `morpheus-ctl reconfigure`.

Full HA Upgrade

Full HA configurations represent multiple app nodes with external (non-system) MySQL, RabbitMQ and Elasticsearch Clusters or Services.

Morpheus Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

Overview

When upgrading other Appliance Configurations from 4.2.x to v5.2.0 only services local to the Morpheus App node(s) will be upgraded. For fully distributed configurations (Full HA), where MySQL, RabbitMQ and Elasticsearch are external clusters or services, the upgrade process will not upgrade these services.

Refer to *v5.2.0 Compatibility & Breaking Changes* for externalized MySQL, Elasticsearch and/or RabbitMQ version requirements.

Upgrade Instructions

Full HA Debian / Ubuntu Upgrade

The following covers upgrading the Morpheus App nodes in Full HA Architecture configurations to v5.2.0.

Important: The following is only for Full HA Architecture configurations, where MySQL, Elasticsearch and RabbitMQ services are external to the App nodes. The following steps assume 3 app nodes, adjust accordingly.

Morpheus Packages

Morpheus Release Package urls can be obtained from <https://morpheushub.com>

4.2.0+ -> v5.2.0 Upgrade

Important: Due to Database schema changes in v5.2.0 it is important to stop the morpheus-ui service on all app nodes prior to upgrade. Failure to do so may result in errors or database corruption. As a best practice, always backup your database prior to any upgrade.

1. Starting with App Node 3, on **All App Nodes**, stop the morpheus-ui services via `morpheus-ctl stop morpheus-ui`. If you receive a timeout, run `morpheus-ctl graceful-kill morpheus-ui`.

```
[root@app-server-3 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-2 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-1 ~]# morpheus-ctl stop morpheus-ui
```

2. Upgrade the deb package on App Node 1, then run a Reconfigure on App Node 1

```
[root@app-server-1 ~]# sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_
↪amd64.deb
[root@app-server-1 ~]# sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
[root@app-server-1 ~]# sudo morpheus-ctl reconfigure
```

Note: All services will automatically stopped and started during the reconfigure process. After the reconfigure has succeeded, tail the ui service to watch ui startup logs with `morpheus-ctl tail morpheus-ui`.

3. Once App Node 1 upgrade has completed and the ui is available, upgrade the deb package on App Node 2, then run a Reconfigure on App Node 2.

```
[root@app-server-2 ~]# sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_
↪amd64.deb
[root@app-server-2 ~]# sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
[root@app-server-2 ~]# sudo morpheus-ctl reconfigure
```

4. Then upgrade the deb package on App Node 3, and run a Reconfigure on App Node 3

```
[root@app-server-3 ~]# sudo wget https://packageUrl.morpheus-appliance_x.x.x-x_
↪amd64.deb
[root@app-server-3 ~]# sudo dpkg -i morpheus-appliance_x.x.x-1_amd64.deb
[root@app-server-3 ~]# sudo morpheus-ctl reconfigure
```

5. The upgrade is complete and the Morpheus-ui services should be running on allocation 3 nodes.

Important: If reconfigure after a rpm package upgrade stalls or hangs on starting a local service it is possible the `morpheus-runsvdir` service did not start or a process it was managing was manually shutdown or killed. To resolve, run `systemctl stop morpheus-runsvdir` then `systemctl start morpheus-runsvdir`, then run reconfigure again, `morpheus-ctl reconfigure`.

Full HA CentOS / RHEL Upgrade

The following covers upgrading the Morpheus App nodes in Full HA Architecture configurations to v5.2.0.

Important: The following is only for Full HA Architecture configurations, where MySQL, Elasticsearch and RabbitMQ services are external to the App nodes.

4.2.0+ -> v5.2.0 Upgrade

Important: Due to Database schema changes in v5.2.0 it is important to stop the morpheus-ui service on all app nodes prior to upgrade. Failure to do so may result in errors or database corruption. As a best practice, always backup your database prior to any upgrade.

1. Starting with App Node 3, on **All** App Nodes, stop the morpheus-ui services via `morpheus-ctl stop morpheus-ui`. If you receive a timeout, run `morpheus-ctl graceful-kill morpheus-ui`.

```
[root@app-server-3 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-2 ~]# morpheus-ctl stop morpheus-ui
```

```
[root@app-server-1 ~]# morpheus-ctl stop morpheus-ui
```

2. Upgrade the RPM package on App Node 1, then run a Reconfigure on App Node 1

```
[root@app-server-1 ~]# sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
↪x86_64.rpm
[root@app-server-1 ~]# sudo rpm -Uvh morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-1 ~]# sudo morpheus-ctl reconfigure
```

Note: All services will automatically stopped and started during the reconfigure process. After the reconfigure has succeeded, tail the ui service to watch ui startup logs with `morpheus-ctl tail morpheus-ui`.

3. Once App Node 1 upgrade has completed and the u is available, upgrade the RPM package on App Node 2, then run a Reconfigure on App Node 2.

```
[root@app-server-2 ~]# sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
↪x86_64.rpm
[root@app-server-2 ~]# sudo rpm -Uvh morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-2 ~]# sudo morpheus-ctl reconfigure
```

4. Then upgrade the RPM package on App Node 3, then run a Reconfigure on App Node 3

```
[root@app-server-3 ~]# sudo wget https://packageUrl.morpheus-appliance-x.x.x-x.x86_64.rpm
↪x86_64.rpm
[root@app-server-3 ~]# sudo rpm -Uvh morpheus-appliance-x.x.x-x.x86_64.rpm
[root@app-server-3 ~]# sudo morpheus-ctl reconfigure
```

5. The upgrade is complete and the Morpheus-ui services should be running with clustered Elasticsearch and RabbitMQ services across the 3 nodes.

Important: If reconfigure after a rpm package upgrade stalls or hangs on starting a local service it is possible the `morpheus-runsvdir` service did not start or a process it was managing was manually shutdown or killed. To resolve, run `systemctl stop morpheus-runsvdir` then `systemctl start morpheus-runsvdir`, then run reconfigure again, `morpheus-ctl reconfigure`.

morpheus-ctl tips

`morpheus-ctl` is useful beyond reconfigures and starting the ui, and many commands can be run across all services, or scoped to a single service.

Some common commands include:

morpheus-ctl status This list all the installed services and their current Status

morpheus-ctl start (service) This starts all services if no service is specified, or starts the specified service. For example,

- `morpheus-ctl start/stop/restart/kill` on an all-in-one appliance will start, stop, restart or kill mysql, elasticsearch, rabbitmq, check-server, guacd and the morpheus-ui, one by one.
- `morpheus-ctl start/stop/restart/kill morpheus-ui` will only start, stop, restart or kill the morpheus-ui service, leaving the other service in their current state. Same goes for `morpheus-ctl start/stop/restart/kill mysql`, `morpheus-ctl start/stop/restart/kill elasticsearch` etc.

`morpheus-ctl` commands:

General Commands:

```
cleanse
    Delete *all* morpheus data, and start from scratch.
help
    Print this help message.
reconfigure
    Reconfigure the application.
show-config
    Show the configuration that would be generated by reconfigure.
uninstall
    Kill all processes and uninstall the process supervisor (data will be preserved).
```

Service Management Commands:

```
graceful-kill
    Attempt a graceful stop, then SIGKILL the entire process group.
hup
    Send the services a HUP.
int
    Send the services an INT.
kill
    Send the services a KILL.
once
    Start the services if they are down. Do not restart them if they stop.
restart
    Stop the services if they are running, then start them again.
service-list
    List all the services (enabled services appear with a *.)
```

(continues on next page)

(continued from previous page)

```
start
  Start services if they are down, and restart them if they stop.
status
  Show the status of all the services.
stop
  Stop the services, and do not restart them.
tail
  Watch the service logs of all enabled services.
term
  Send the services a TERM.

Elasticsearch Commands:

elastic-util
  Backup/Restore ElasticSearch data

Firewall Commands:

firewall-enable-blocking
  Enables firewall blocking mode.
```

Morpheus DB Migration

If your new installation is part of a migration or you need to move the data from your original Morpheus database, this is easily accomplished by using a stateful dump.

To begin this, stop the Morpheus UI on your original Morpheus server:

```
[root@app-server-old ~] morpheus-ctl stop morpheus-ui
```

Once this is done you can safely export. To access the MySQL shell we will need the password for the Morpheus DB user. We can find this in the morpheus-secrets file:

```
[root@app-server-old ~] cat /etc/morpheus/morpheus-secrets.json | grep morpheus_
↪password
"morpheus_password": "451e122cr5d122asw3de5e1b", <-----this one
"morpheus_password": "9b5vdj4de5awf87d",
```

Take note of the first morpheus_password as it will be used to invoke a dump. Morpheus provides embedded binaries for this task. Invoke it via the embedded path and specify the host. In this example we are using the morpheus database on the MySQL listening on localhost. Enter the password copied from the previous step when prompted:

```
[root@app-server-old ~] /opt/morpheus/embedded/mysql/bin/mysqldump -u morpheus -h 127.
↪0.0.1 morpheus -p > /tmp/morpheus_backup.sql
Enter password:
```

This file needs to be pushed to the new Morpheus Installation's backend. Depending on the GRANTS in the new MySQL backend, this will likely require moving this file to one of the new Morpheus frontend servers.

Once the file is in place it can be imported into the backend. Begin by ensuring the Morpheus UI service is stopped on all of the application servers:

```
[root@app-server-new ~] morpheus-ctl stop morpheus-ui
```

Then you can import the MySQL dump into the target database using the embedded MySQL binaries, specifying the database host, and entering the password for the morpheus user when prompted:


```
[root@app-server-new ~] /opt/morpheus/embedded/mysql/bin/mysql -u morpheus -h 10.1.2.
↪2 morpheus -p < /tmp/morpheus_backup.sql
Enter password:
```

The data from the old appliance is now replicated on the new appliance. Simply start the UI to complete the process:

```
[root@app-server-new ~] morpheus-ctl start morpheus-ui
```

Scaling Morpheus Nodes

Morpheus App nodes can be scaled to accommodate additional load. Appliance nodes can be scaled vertically in centralized architectures, and both vertically and horizontally in distributed architectures.

Vertical Scaling

In all Appliance Architectures, Application nodes can be vertically scaled at any time, however a reconfigure must be performed for additional resources to be utilized by Morpheus on a node, which will result in the `morpheus-ui` restarting on the reconfiguring node.

Morpheus configures memory/ram utilization for services during the `reconfigure` process. If additional memory/ram is added to a Host or VM running the Morpheus App, the additional memory/ram will not be utilized by the Morpheus Application until a `morpheus-ctl reconfigure` command is ran and the additional memory/ram is recognized.

When the `morpheus-ctl reconfigure` command detects changes on available memory/ram, it will trigger a `morpheus-ui` service restart.

Important: When the `morpheus-ctl reconfigure` command detects changes on available memory/ram, it will restart the `morpheus-ui` service.

The impact on Availability depends on the Morpheus Appliance Architecture.

- **Centralized Appliances** Morpheus will be unavailable while the `morpheus-ui` restarts.
- **Distributed Appliances** Zero-down time can be achieved by Reconfiguring one App Node at a time, with proper Front-End Load Balancer configuration.

Horizontal Scaling

Additional Morpheus App Nodes can be added at any time to Fully Distributed Architectures.

- Configure Shared Storage paths for the new App Node(s)
- Install, but do not run the `morpheus-ctl reconfigure` command on the new App Node(s), using the same Morpheus version as the existing Appliance nodes.
- Copy the `morpheus.rb` from an existing App Node to the new App Node(s)
- Ensure permissions and network configuration for the new App Node(s) to access all MySQL and Elasticsearch nodes, and the RabbitMQ VIP.
- Ensure permissions and network configuration for all required UI services and Integrations, such as network access to ESXi hosts over 443 for Hypervisor console and/or image transfers.

- Add associated SSL files and configuration, if not on shared storage.
- Reconfigure the new App Node(s) via `morpheus-ctl reconfigure`
- Verify UI startup succeeded
- Add New App Node(s) to Front End Morpheus UI Load Balancer pool.

During `morpheus-ctl reconfigure`, the new App Node(s) will validate and be configured to use the existing tiers for the UI service. Upon successful reconfigure, the Morpheus service will be available on the App Node(s) with consistent data and capabilities.

Note: No services, including `morpheus-ui`, are required to be shut down on existing nodes when adding new App Nodes

Morpheus UI war files

Pre-release or patched versions of the Morpheus UI are sometimes provided. To deploy the ui war on a Morpheus Appliance:

1. Download the war file to the target appliance

```
wget https://url/war_file
```

Note: If the war file is provided via a droplr link, ensure a + is added to end of droplr url or the file will not download

2. Backup current war file

```
sudo mv /opt/morpheus/lib/morpheus/morpheus-ui.war /opt/morpheus/lib/morpheus/  
↪morpheus-ui.bak.`date +%m-%d-%Y`
```

3. Move and rename new war file

```
sudo mv <file> /opt/morpheus/lib/morpheus/morpheus-ui.war
```

4. Ensure war is owned by `morpheus-app`

```
sudo chown morpheus-app.morpheus-app /opt/morpheus/lib/morpheus/morpheus-ui.war
```

5. Restart the Morpheus UI

```
sudo morpheus-ctl restart morpheus-ui
```

The new ui war will load on startup!

Additional Configuration Options

Advanced morpheus.rb Settings

Morpheus allows for additional advanced customizations to the morpheus.rb file located in /etc/morpheus/morpheus.rb. Below is a list of the supported items available in the morpheus.rb file.

```

appliance_url 'https://morpheus.appliance-url.com'
  # The appliance_url must not contain a trailing `/.`.
  # Appending alternate port to appliance_url is supported. ie 'https://morpheus.
  ↪appliance-url.com:8443'
  # The appliance_url cannot exceed 64 characters

app['encrypted_key_suffix'] = 'suffix'

ui['http_client_connect_timeout'] = 10000 #in seconds
ui['http_client_connect_timeout'] = 600000 #in seconds
ui['kerberos_config'] = nil
ui['kerberos_login_config'] = nil
ui['log_dir'] = '/var/log/morpheus/morpheus-ui'
ui['max_memory_mb'] = nil
ui['memory_alloc_arena_max'] = 2
ui['memory_map_max'] = 65536
ui['memory_map_threshold'] = 131072
ui['memory_top_pad'] = 131072
ui['memory_trim_threshold'] = 131072
ui['vm_images_cdn_url'] = 'https://morpheus-images.morpheusdata.com'

mysql['enable'] = true
mysql['host'] = {'127.0.0.1' => 3306}
mysql['log_dir'] = '/var/log/morpheus/mysql'
mysql['max_active'] = 100
mysql['morpheus_db_user'] = 'morpheus'
mysql['morpheus_db'] = 'morpheus'
mysql['mysql_url_override'] = 'jdbc:mysql://10.30.20.10:3306,10.30.20.11:3306,10.30.20.
  ↪12:3306/morpheusdb?autoReconnect=true&useUnicode=true&characterEncoding=utf8&
  ↪failOverReadOnly=false&useSSL=false'
mysql['tmp_dir'] = '/tmp/mysql'

logging['svlogd_num'] = 30 # keep 30 rotated log files
logging['svlogd_size'] = 209715200 # 200 MB in bytes
logging['svlogd_timeout'] = 86400 # rotate after 24 hours in seconds

rabbitmq['enable'] = true
rabbitmq['heartbeat'] = nil
rabbitmq['host'] = '127.0.0.1'
rabbitmq['log_dir'] = '/var/log/morpheus/rabbitmq'
rabbitmq['nodename'] = 'rabbit@localhost'
rabbitmq['port'] = '5672'
rabbitmq['queue_user'] = 'queue_user'
rabbitmq['vhost'] = 'morpheus'

elasticsearch['log_dir'] = '/var/log/morpheus/elasticsearch'
elasticsearch['enable'] = true
elasticsearch['es_hosts'] = {'127.0.0.1' => 9200}
elasticsearch['host'] = "127.0.0.1"
elasticsearch['memory_alloc_arena_max'] = 2

```

(continues on next page)

(continued from previous page)

```

elasticsearch['memory_map_max'] = 65536
elasticsearch['memory_map_threshold'] = 131072
elasticsearch['memory_top_pad'] = 131072
elasticsearch['memory_trim_threshold'] = 131072
elasticsearch['open_files'] = 204800
elasticsearch['replica_count'] = 1

nginx['cache_max_size'] = '5000m'
nginx['enable'] = true
nginx['loading_pages']['failure_page_h1'] = 'Morpheus Server Error'
nginx['loading_pages']['failure_page_h2'] = 'Please contact your system administrator,
↳for assistance.'
nginx['loading_pages']['failure_page_title'] = 'Morpheus Server Error'
nginx['loading_pages']['iteration_time'] = 10000 # milliseconds
nginx['loading_pages']['loading_page_h1'] = 'Morpheus is Loading...'
nginx['loading_pages']['loading_page_h2'] = 'please wait'
nginx['loading_pages']['loading_page_title'] = 'Morpheus Loading'
nginx['loading_pages']['max_loops'] = 60 # seconds
nginx['loading_pages']['timeout_page'] = '/timeout.html'
nginx['loading_pages']['timeout_page_h1'] = 'Timeout waiting for Morpheus to load,
↳click below to try again.'
nginx['loading_pages']['timeout_page_title'] = 'Morpheus timeout, please try again...'
nginx['ssl_ciphers'] = "ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-GCM-SHA256:DHE-
↳RSA-AES256-GCM-SHA384:DHE-RSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-SHA384:ECDHE-RSA-
↳AES128-SHA256:ECDHE-RSA-AES256-SHA:ECDHE-RSA-AES128-SHA:DHE-RSA-AES256-SHA256:DHE-
↳RSA-AES128-SHA256:DHE-RSA-AES256-SHA:DHE-RSA-AES128-SHA:ECDHE-RSA-DES-CBC3-SHA:EDH-
↳RSA-DES-CBC3-SHA:AES256-GCM-SHA384:AES128-GCM-SHA256:AES256-SHA256:AES128-
↳SHA256:AES256-SHA:AES128-SHA:DES-CBC3-SHA:HIGH:!aNULL:!eNULL:!EXPORT:!DES:!MD5:!
↳PSK:!RC4"
nginx['ssl_company_name'] = "Morpheus, LLC"
nginx['ssl_country_name'] = "US"
nginx['ssl_email_address'] = "personal@email.com"
nginx['ssl_locality_name'] = "San Mateo"
nginx['ssl_organizational_unit_name'] = "DevOps"
nginx['ssl_protocols'] = "TLSv1 TLSv1.1 TLSv1.2"
nginx['ssl_session_cache'] = "builtin:1000 shared:SSL:10m"
nginx['ssl_session_timeout'] = "5m"
nginx['ssl_state_name'] = "CA"
nginx['worker_connections'] = 10240
nginx['workers'] = integer calculated from number of cpus

repo['repo_host_url'] = 'https://downloads.morpheusdata.com'

```

Note: elasticsearch['replica_count'] settings only apply to local Elasticsearch and not an external cluster. The user must set the replica count in the code for each index. The setting in morpheus.rb is only the cluster default and only applies to the all-in-one appliance. If the cluster is external, the user must set the default on their Elasticsearch config file.

Offline Installations and Upgrades

For customers that have an appliance behind a firewall/proxy that does not allow downloads from our Amazon download site, you can add the supplemental package to add the needed packages the standard Morpheus installer would have downloaded.

Offline Installation Requirements

- NTP should be correctly configured and the server is able to connect to the NTP server in the ntp.conf file
- The OS package repositories should be configured to use local LAN repository servers or the server should be able to receive packages from the configured repositories
- The standard Morpheus and supplemental packages must be downloaded from another system and transferred to the Morpheus Appliance server
- The supplemental package is additive, the full installer is also required

Note: The supplemental package is linked 1-to-1 to the appliance release. For example the supplemental package for 4.2.1-1 should be used with the appliance package 4.2.1-1

Offline Install

Ubuntu/Debian

1. Download both the regular Morpheus Appliance package and the Supplemental packages on to the appliance server:

```
wget http://example_url/morpheus-appliance_version_amd64.deb
wget http://example_url/morpheus-appliance-supplemental_version_all.deb
```

2. Install the both the Appliance package AND the Supplemental package.

```
sudo dpkg -i morpheus-appliance_version_amd64.deb
sudo dpkg -i morpheus-appliance-supplemental_version_all.deb
```

3. Set the Morpheus UI appliance url (if needed, hostname will be automatically set).

```
# edit appliance_url to resolvable url (if not configured correctly by default)
sudo vi /etc/morpheus/morpheus.rb
```

4. Reconfigure the appliance to install required packages

```
sudo morpheus-ctl reconfigure
```

The Chef run should complete successfully. There is a small pause when Chef runs the resource `remote_file[package_name]` action create while Chef verifies the checksum. After the reconfigure is complete, the morpheus-ui will start and be up in a few minutes.

Note: Tail the morpheus log file located at `/var/log/morpheus/morpheus-ui/current` with the command `morpheus-ctl tail morpheus-ui` and look for the Morpheus ascii logo to know when the morpheus-ui is

up.

CentOS/RHEL

1. Download both the regular Morpheus Appliance package and the matching Supplemental package on to the Appliance server:

```
wget http://example_url/morpheus-appliance_package_url.noarch.rpm
wget http://example_url/morpheus-appliance_package_supplemental_url.noarch.rpm
```

2. Install the both the Appliance package AND the Supplemental package.

```
sudo rpm -i morpheus-appliance_package_url.noarch.rpm
sudo rpm -i morpheus-appliance_package_supplemental_url.noarch.rpm
```

3. Set the Morpheus UI appliance url (if needed, hostname will be automatically set).

```
#Edit appliance_url to resolvable url (if not configured correctly by default)

sudo vi /etc/morpheus/morpheus.rb
```

4. Reconfigure the appliance to install required packages

```
sudo morpheus-ctl reconfigure
```

The Chef run should complete successfully. There is a small pause when Chef runs the resource `remote_file[package_name]` action create while Chef verifies the checksum. After the reconfigure is complete, the morpheus-ui will start and be up in a few minutes.

Note: Tail the morpheus-ui log file with `morpheus-ctl tail morpheus-ui` and look for the Morpheus ascii logo to know when the morpheus-ui is up.

Proxies

Overview

In many situations , companies deploy virtual machines in proxy restricted environments for things such as PCI Compliance, or just general security. As a result of this Morpheus provides out of the box support for proxy connectivity. Proxy authentication support is also provided with both Basic Authentication capabilities as well as NTLM for Windows Proxy environments. Morpheus is even able to configure virtual machines it provisions to utilize these proxies by setting up the operating systems proxy settings directly (restricted to cloud-init based Linux platforms for now, but can also be done on windows based platforms in a different manner).

To get started with Proxies, it may first be important to configure the Morpheus appliance itself to have access to proxy communication for downloading service catalog images. To configure this, visit the Admin -> Settings page where a section labeled “Proxy Settings” is located. Fill in the relevant connection info needed to utilize the proxy. It may also be advised to ensure that the Linux environment’s `http_proxy`, `https_proxy`, and `no_proxy` are set appropriately.

Defining Proxies

Proxies can be used in a few different contexts and optionally scoped to specific networks with which one may be provisioning into or on a cloud integration as a whole. To configure a Proxy for use by the provisioning engines within Morpheus we must go to `Infrastructure -> Networks -> Proxies`. Here we can create records representing connection information for various proxies. This includes the host ip address, proxy port, and any credentials (if necessary) needed to utilize the proxy. Now that these proxies are defined we can use them in various contexts.

Cloud Communication

When morpheus needs to connect to various cloud APIs to issue provisioning commands or to sync in existing environments, we need to ensure that those api endpoints are accessible by the appliance. In some cases the appliance may be behind a proxy when it comes to public cloud access like Azure and AWS. To configure the cloud integration to utilize a proxy, when adding or editing a cloud there is a setting called “API Proxy” under “Advanced Options”. This is where the proxy of choice can be selected to instruct the Provisioning engine how to communicate with the public cloud. Simply adjust this setting and the cloud should start being able to receive/issue instructions.

Provisioning with Proxies

Proxy configurations can vary from operating system to operating system and in some cases it is necessary for these to be configured in the blueprints as a prerequisite. In other cases it can also be configured automatically. Mostly with the use of cloud-init (which all of our out of the box service catalog utilizes on all clouds). When editing/creating a cloud there is a setting for “Provisioning Proxy” in “Provisioning Options”. If this proxy is set, Morpheus will automatically apply these proxy settings to the guest operating system.

Overriding proxy settings can also be done on the Network record. Networks (or subnets) can be configured in `Infrastructure -> Networks` or on the Networks tab of the relevant Cloud detail page. Here, a proxy can also be assigned as well as additional options like the `No Proxy` rules for proxy exceptions.

Docker

When provisioning Docker based hosts within a Proxy environment it is up to the user to configure the docker hosts proxy configuration manually. There are workflows that can be configured via the Automation engine to make this automatic when creating docker based hosts. Please see documentation on Docker and proxies for specific information.

Proxy setups can vary widely from company to company, and it may be advised to contact support for help configuring morpheus to work in the proxy environment.

SSL Certificates

By default Morpheus generates a Self-Signed SSL Certificate. The Self-Signed SSL Certificate can be replaced with a Trusted CA Signed SSL Certificate.

Trusted CA Signed SSL Certificate Implementation

1. If you don't already have your certificate, run an OpenSSL command to generate an SSL certificate request (.csr) and private key (.key). If you need help formatting the command, [DigiCert provides a helpful tool](#)
2. Submit your certificate request (.csr) and await approval of the request and return of the certificate (.crt)
3. Copy the private key and certificate to /etc/morpheus/ssl/your_fqdn_name.key and /etc/morpheus/ssl/your_fqdn_name.crt respectively

- **Extracting Certificates in PFX Format**

```
# Extract the private key
openssl pkcs12 -in example.pfx -nocerts -nodes -out priv.key
# Extract the public key
openssl pkcs12 -in example.pfx -clcerts -nokeys -out pub.crt
# Extract the CA cert chain
openssl pkcs12 -in example.pfx -cacerts -nokeys -chain -out ca.crt
```

- **Extracting Certificates in PEM Format**

```
# Extract the private key
openssl x509 -outform der -in your-cert.pem -out your-cert.key
# Extract the public key
openssl x509 -outform der -in your-cert.pem -out your-cert.key
```

```
nginx['ssl_certificate'] = 'path to the certificate file'
nginx['ssl_server_key'] = 'path to the server key file'
```

Note: Both files should be owned by root and only readable by root, also if the server certificate is signed by an intermediate then you should include the signing chain inside the certificate file. The key file needs to be decrypted for Morpheus to install it properly.

4. Next simply reconfigure the appliance and restart nginx:

```
sudo morpheus-ctl reconfigure
sudo morpheus-ctl restart nginx
```

Self-Signed SSL Certificate Regeneration

When Morpheus is deployed it generates a 10 year Self-Signed SSL Certificate. Below details the process to regenerate the Certificate and Key files.

Regenerate both the Certificate and Key

1. Delete the certificate and key files in /etc/morpheus/ssl/.
2. Run Reconfigure `morpheus-ctl reconfigure`.
3. Restart NGINX `morpheus-ctl restart nginx`.

Regenerate only the Certificate

1. Delete the certificate file in `/etc/morpheus/ssl/`.
2. Run `Reconfigure morpheus-ctl reconfigure`.
3. Restart NGINX `morpheus-ctl restart nginx`.

Import Trusted Certificates

Important: The following applies to upgrades after modifying the java keystore.

Steps to import trusted certificates to Morpheus after an upgrade.

1. Obtain the full SSL certificate chain in PEM format.
2. Copy them to each appliance and place them in the `/etc/morpheus/ssl/trusted_certificates` directory.
3. Run `morpheus-ctl reconfigure` on each appliance, note you don't need to stop Morpheus before you run this.
4. Run the following command as root:

```
export PATH=/opt/morpheus/sbin:/opt/morpheus/sbin:/opt/morpheus/embedded/
↪sbin:/opt/morpheus/embedded/bin:$PATH
```

5. Run the following command for each certificate in the chain, adjusting the file and alias name as needed. Answer yes for the root certificate when asked if you want to trust it.

```
/opt/morpheus/embedded/java/jre/bin/keytool -import -keystore /opt/
↪morpheus/embedded/java/jre/lib/security/cacerts -trustcacerts -file /
↪etc/morpheus/ssl/trusted_certs/root_ca.pem -alias some_alias -storepass_
↪changeit
```

6. Verify by running:

```
openssl s_client -connect host:port -showcerts -tls1_2``
```

7. You should get an output similar to:

```
New, TLSv1/SSLv3, Cipher is ECDHE-RSA-AES256-GCM-SHA384
Server public key is 2048 bit
Secure Renegotiation IS supported
No ALPN negotiated
SSL-Session:
Protocol : TLSv1.2
Cipher : ECDHE-RSA-AES256-GCM-SHA384
Session-ID:↪
↪5D9E820E4FF2A73A9977BA663E6029AA5415FEE85F49D8B1E541F5997C8E1FB2
Session-ID-ctx:
Master-Key:↪
↪29EEC2E7750C659AECB9942902D9A87B824E571522812B718420FC08F8D2ACE68CB16EC812A7D90B12A86D1970
Key-Arg : None
PSK identity: None
PSK identity hint: None
SRP username: None
```

(continues on next page)

(continued from previous page)

```
Start Time: 1547219217
Timeout   : 7200 (sec)
Verify return code: 0 (ok) #<-----
```

8. If the certificates are installed correctly you should see `Verify return code: 0 (ok)`. If they were not installed correctly then you will see a return similar to: `Verify return code: 21 (unable to verify the first certificate)`
9. Repeat for all App Nodes

Data Encryption

By default, Morpheus encrypts sensitive data in the Database using AES encryption mode with GCM. Passwords and other strings in Morpheus Appliance configuration files such as `morpheus-secrets.json` and `morpheus.rb` are in plain text as they are only accessible by root.

Passwords and other strings in Morpheus Appliance configuration files can be set to an encrypted string using the Morpheus crypto utility to generate ENC strings and then using `ENC(string)` as the value in the configuration file.

Additionally a custom Encryption Key Suffix can be set in the `morpheus.rb` configuration file. This suffix will be combined with a system string to generate a SHA-256 hash, which is used to generate the AES encryption key.

Generate ENC Strings for morpheus-secrets.json

System generated passwords are set in `/etc/morpheus/morpheus-secrets.json`. These entries can be updated to ENC strings with the following steps:

1. On the Morpheus appliance, run `morpheus-ctl get-crypto-string migrate` which will output `ENC()` strings for the passwords in `morpheus-secrets.json`
2. Update the desired password strings in the `morpheus-secrets.json` config file with the matching `ENC()` string.
3. Save `morpheus-secrets.json`
4. Run `morpheus-ctl reconfigure`

Generate ENC Strings for custom morpheus.rb entries

`ENC()` strings can be generated for sensitive data set in `morpheus.rb`, such as the password to an external service.

To generate `ENC()` strings for `morpheus.rb` entries:

1. On the Morpheus appliance, run `morpheus-ctl get-crypto-string string $clear_text '$suffix'` which will output strings for the passwords in `morpheus-secrets.json`
 - Replace `$clear_text` with the string to be encrypted
 - If a suffix is defined in `morpheus.rb` (as described in the next section), replace `$suffix` with your suffix.

Note: It is advisable to disable bash history logging by running `unset HISTFILE` before running the `morpheus-ctl get-crypto-string` command and then `set HISTFILE=$HOME/.bash_history` to reenale.

2. Update the desired password strings in the `morpheus.rb` config file with the matching string output, using `ENC($output)` format

- Example: `mysql['morpheus_password'] = 'ENC($ZI5Dna00quhxKe$kDFD+U2ZeJUuYiNC$F1+czPNyc`

3. Save `morpheus.rb`

4. Run `morpheus-ctl reconfigure`

Encrypted Key Suffix

A custom Encryption Key Suffix can be set in the `morpheus.rb` configuration file. This suffix will be combined with a system string to generate a SHA-256 Hash, which is used in the generation of the system AES encryption key.

Danger: Setting a custom Encryption Key Suffix affects all data encrypted by Morpheus, including database and cypher data. Encryption Key Suffix is required in the event data needs to be migrated or recovered. Once the Encryption Key Suffix is set, data cannot be recovered without it. Store any Encryption Key Suffix externally where it can be referenced for future scenarios.

Important: The Encryption Key Suffix cannot be changed or removed after being set. Changing or removing an existing Encryption Key Suffix will prevent data access. If an existing suffix is altered in the `morpheus.rb` file, it must be restore to its original value.

1. Add `app['encrypted_key_suffix'] = '$suffix'` to `/etc/morpheus/morpheus.rb`, replacing `$suffix` with your suffix.

Danger: Once an Encryption Key Suffix is set and applied via `reconfigure`, it cannot be altered or removed and data cannot be migrated or recovered without it.

2. Run `morpheus-ctl reconfigure`

- Reconfigure will generate a new encryption key using the suffix and set (ENC) values for service password in `application.yml`

Initial Appliance Setup

Appliance Setup

After installation, log into the appliance at the URL presented upon completion. An initial setup wizard walks through the first account and user creations.

1. Enter Master Account name

- Typically, the Master Account name is your Company name.

2. Create Master User

- First Name
- Last Name
- Username

- Email Address
 - Password * Must be at least 8 characters long and contain one each of the following: Uppercase letter, lowercase letter, Number, Special Character
3. Enter Appliance Name & Appliance URL
 - The Appliance Name is used for white labeling and as a reference for multi-appliance installations.
 - The Appliance URL is the URL all provisioned instances will report back to. Example: <https://example.morpheusdata.com>. The Appliance URL can be changed later, and also set to different url per cloud integration.
 4. Optionally Enable or Disable Backups, Monitoring, or Logs from this screen.

Note: You may adjust these settings from the Administration section.

Note: The Master Account name is the top-level admin account.

Note: The Master User is the system super user and will have full access privileges.

Upon completing of the initial appliance setup, you will be taken to the Admin -> Settings page, where you will add your License Key.

Login Methods

Master Tenant

- Enter username or email. and password

Subtenant

To login, subtenants can either use the master tenant URL with `subtenant\username` formatting:

Example: I have a username `subuser` that belongs to a tenant with the subdomain `subaccount`. When logging in from the main login url, I would now need to enter in: `subaccount\subuser`

Or use the tenant specific URL which can be found and configured under Administration > Tenants > Select Tenant > Identity Sources.

Tenants > Morpheus-Tenant > Identity Sources

Morpheus-Tenant

ACCOUNT LOGIN URL

<https://sandbox.morpheusdata.com/login/account/79>

EDIT

Search

+ ADD IDENTITY SOURCE

TYPE	NAME	DETAILS	ACTIVE
------	------	---------	--------

Important: In 3.4.0+ Subtenant users will no longer be able to login from the main login url without specifying their subdomain.

Configure Cloud-init Global Settings

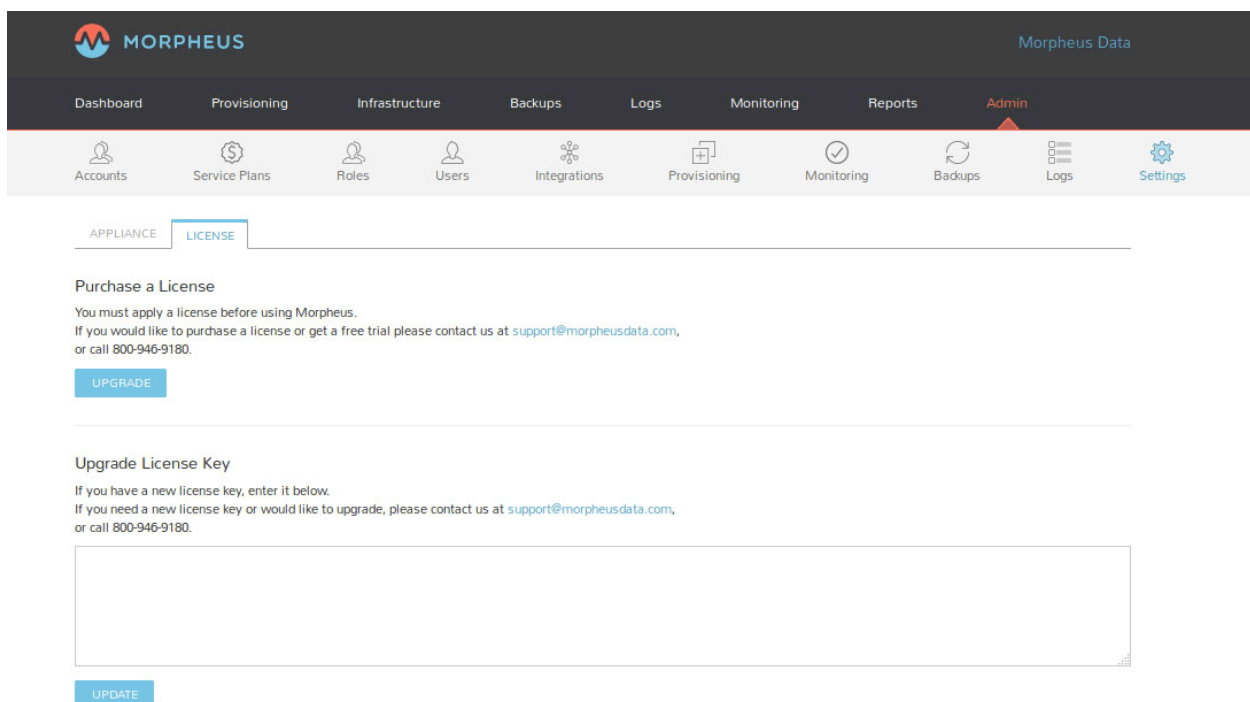
When using cloud-init, cloudbase-init, VMware Tools customizations, or Nutanix Sysprep, Global Linux User and Windows Administrator credentials can be set using the settings in *Administration - Provisioning*. It is recommended to define these settings after installation unless credentials are defined per Virtual Image for Provisioning.

Add a License Key

In order to provision anything in Morpheus, a Morpheus License Key must be applied.

If you do not already have a license key, one may be requested from <https://www.morpheushub.com> or from your Morpheus representative.

In the Administration -> Settings section, select the LICENSE tab, paste your License Key and click *UPDATE*



The screenshot shows the Morpheus Administration interface. At the top, there is a navigation bar with the Morpheus logo and the text "Morpheus Data". Below this is a secondary navigation bar with tabs: Dashboard, Provisioning, Infrastructure, Backups, Logs, Monitoring, Reports, and Admin. The Admin tab is selected and highlighted in orange. Below the navigation bar is a row of icons representing various settings: Accounts, Service Plans, Roles, Users, Integrations, Provisioning, Monitoring, Backups, Logs, and Settings. The Settings icon is highlighted in blue. Below the Settings icon, there are two tabs: APPLIANCE and LICENSE. The LICENSE tab is selected and highlighted in blue. The main content area of the LICENSE tab contains the following text:

Purchase a License
 You must apply a license before using Morpheus.
 If you would like to purchase a license or get a free trial please contact us at support@morpheusdata.com, or call 800-946-9180.

Below this text is a blue button labeled "UPGRADE".

Below the "UPGRADE" button, there is a section titled "Upgrade License Key". It contains the following text:

If you have a new license key, enter it below.
 If you need a new license key or would like to upgrade, please contact us at support@morpheusdata.com, or call 800-946-9180.

Below this text is a large, empty text input field. Below the input field is a blue button labeled "UPDATE".

When the license is accepted, your license details will populate in the Current License section.

If you receive an error message and your license is not accepted, please check it was copied in full and then contact your Morpheus representative. You can also verify the License Key and expiration at <https://www.morpheushub.com>.

Core Functionality

Morpheus Discovery

Morpheus has the ability to ingest existing environments. Existing running workloads will be inventoried into Morpheus and displayed in the UI. In 5-7 days Morpheus will start making recommendations based off of usage and pricing

Note: Work loads that are inventoried do not have to be converted to managed.

Once inventoried, Morpheus can provide valuable data for that instance:

- Morpheus will know about networks
- Start aggregating cost on public clouds
- Start tracking usage
- Some Clouds offer statistical details (Amazon / VMware)
- Power Status

Right away inventorying existing environments will provide you with immediate insight to that environment. Once an existing workload has been discovered it can be converted to managed. Once converted to managed, Morpheus can deliver more capabilities and features.

Note: Workloads do not need the agent installed to be managed

Once a workload is managed:

- Enforce expiration/shutdown policies. This helps reign in environments (sprawl) and reduce cost.
- Can tell what instance type it is
- Can install agent (agent is optional)
- Installing agent provides credentials and allows you to run workflows against it (day 2 operations)

Morpheus Agent

Overview

The Morpheus Agent is an important and powerful facet of Morpheus as an orchestration tool. Though it is not required, which is one unique capability of our platform versus some of our competitors, it is recommended for use as it brings many benefits. Not only does it provide statistics for the guest operating system and resource utilization, it also brings along with it monitoring and log aggregation capabilities. After an initial brownfield discovery, Users can decide to convert unmanaged VMs to managed.

Note: **Agent installation is not required to manage an Instance.** If you don't have the Agent installed, we make every effort to aggregate stats. These will vary based on the Cloud and can be more limited or less accurate without utilizing Morpheus Agent.

The Morpheus Agent is very lightweight and secure. It does not open any inbound network ports but rather only opens an outbound connection back to the Morpheus appliance over port 443 (HTTPS or WSS protocol). This allows for a

bidirectional command bus where instructions can be sent to orchestrate a workload without needing access to things like SSH or WinRM. The tool can even be installed at provision time via technologies like Cloud-Init, such that the Morpheus appliance itself doesn't even need direct network access to the VLAN under which the workload resides. By doing this we address many of the network security concerns that come with the use of an agent while demonstrating its security and analytics benefits. We can even use this statistical data at the guest OS level rather than the hypervisor level to provide extremely precise right-sizing recommendations.

Morpheus Agent Key Features

Key Enhanced Statistics and Benefits

Category	Statistic	With Agent	Without Agent
Memory	Max Memory	Yes	Yes
Memory	Used Memory	Yes	No
Memory	Cache Memory	Yes	No
Storage	Max Storage	Yes	Yes
Storage	Used Memory	Yes	No
Processing	System CPU %	Yes	Yes
Processing	User CPU %	Yes	No
IOPS	Total IOPS	Yes	No
IOPS	IOPS Read	Yes	No
IOPS	IOPS Write	Yes	No
Networking	Net TX Rate	Yes	No
Networking	Net RX Rate	Yes	No
Other	Agent Command Bus	Yes	No
Other	Log Aggregation	Yes	No
Other	Health & Monitoring	Yes	No

Additional benefits:

- Installation is optional for provisioned and managed VMs
- The Linux agent can be shrunk and should be less than 0.2% peak
- Provides a command bus such that Morpheus doesn't need credentials to access a box
- Can still manage workflows if credentials are changed
- SSH agent can be disabled and still get access to the box
- Agent can be installed over Cloud-init, Windows unattend.xml, VMware Tools, SSH, WinRM, Cloudbase-Init, or manually
- Makes a single, persistent connection over HTTPS websocket and runs as a service
- Buffers and compresses logs, then sends them in chunks to minimize packets
- Supports syslog forwarding
- Accepts commands, executes commands, writes files, and manipulates firewalls
- Significantly enhances Guidance recommendations through enhanced statistics

Note: The Morpheus Agent is required for managed Docker, Kubernetes, SCVMM, Hyper-V, KVM, and ESXi Hosts (for ESXi-only Cloud, not vCenter Clouds).

Morpheus Agent OS Support

Name	Bit Count	Category	Code	OS Family	OS Version	Platform	Ver
amazon linux 2 64-bit	64	amazonlinux	amazonlinux.2.64	centos	2	linux	ama
centOS	64	centos	cent	rhel	all	linux	cent
centOS 6	64	centos	cent.6	rhel	6	linux	cent
centOS 6 64-bit	64	centos	cent.6.64	rhel	6	linux	cent
centOS 64-bit	64	centos	cent.64	rhel	all	linux	cent
centOS 7	64	centos	cent.7	rhel	7	linux	cent
centOS 7 64-bit	64	centos	cent.7.64	rhel	7	linux	cent
centOS 8 64-bit	64	centos	cent.8.64	rhel	8	linux	cent
debian	32	debian	debian	debian	all	linux	debi
debian 6	32	debian	debian.6	debian	6	linux	debi
debian 6 64-bit	64	debian	debian.6.64	debian	6	linux	debi
debian 64-bit	64	debian	debian.64	debian	all	linux	debi
debian 7	32	debian	debian.7	debian	7	linux	debi
debian 7 64-bit	64	debian	debian.7.64	debian	7	linux	debi
debian 8	32	debian	debian.8	debian	8	linux	debi
debian 8 64-bit	64	debian	debian.8.64	debian	8	linux	debi
debian 9.4 64-bit	64	debian	debian.9.4.64	debian	9	linux	debi
esxi 4	64	esxi	esxi.4	NULL	4	esxi	vmw
esxi 5	64	esxi	esxi.5	NULL	5	esxi	vmw
esxi 6	64	esxi	esxi.6	NULL	6	esxi	vmw
fedora	32	fedora	fedora	NULL	all	linux	fedo
fedora 64-bit	64	fedora	fedora.64	NULL	all	linux	fedo
linux	64	linux	linux	NULL	all	linux	linu
linux 32-bit	32	linux	linux.32	NULL	all	linux	linu
linux 64-bit	64	linux	linux.64	NULL	all	linux	linu
mac os	64	mac	mac	darwin	all	osx	app
opensuse	32	opensuse	opensuse.32	suse	all	linux	ope
opensuse 15 64-bit	64	opensuse	suse.15.64	suse	15	linux	suse
opensuse 64-bit	64	opensuse	opensuse.64	suse	all	linux	ope
oracle enterprise	32	oracle	oracle.32	NULL	all	linux	orac
oracle enterprise 64-bit	64	oracle	oracle.64	NULL	all	linux	orac
oracle linux 64-bit	64	oracle	oracle.linux.64	NULL	all	linux	orac
oracle vm server	64	oracle	ovs	NULL	all	linux	orac
other 32-bit	32	other	other.32	NULL	all	other	othe
other 64-bit	64	other	other.64	NULL	all	other	othe
redhat	64	redhat	redhat	rhel	all	linux	redh
redhat 6	64	redhat	redhat.6	rhel	6	linux	redh
redhat 6 64-bit	64	redhat	redhat.6.64	rhel	6	linux	redh
redhat 64-bit	64	redhat	redhat.64	rhel	all	linux	redh
redhat 7	64	redhat	redhat.7	rhel	7	linux	redh
redhat 7 64-bit	64	redhat	redhat.7.64	rhel	7	linux	redh
redhat 8 64-bit	64	redhat	redhat.8.64	rhel	8	linux	redh
solaris	32	solaris	solaris.32	NULL	all	solaris	sola
solaris 64-bit	64	solaris	solaris.64	NULL	all	solaris	sola
suse enterprise	32	suse	suse	suse	all	linux	suse
suse enterprise 11	32	suse	suse.11	suse	11	linux	suse
suse enterprise 11 64-bit	64	suse	suse.11.64	suse	11	linux	suse

Table 4 – continued from previous page

Name	Bit Count	Category	Code	OS Family	OS Version	Platform	Ver
suse enterprise 12	32	suse	suse.12	suse	12	linux	suse
suse enterprise 12 64-bit	64	suse	suse.12.64	suse	12	linux	suse
suse enterprise 64-bit	64	suse	suse.64	suse	all	linux	suse
ubuntu	32	ubuntu	ubuntu	debian	all	linux	can
ubuntu 12	32	ubuntu	ubuntu.12.04	debian	12.04	linux	can
ubuntu 12 64-bit	64	ubuntu	ubuntu.12.04.64	debian	12.04	linux	can
ubuntu 13	32	ubuntu	ubuntu.13.10	debian	13.1	linux	can
ubuntu 13 64-bit	64	ubuntu	ubuntu.13.10.64	debian	13.1	linux	can
ubuntu 14	32	ubuntu	ubuntu.14.04	debian	14.04	linux	can
ubuntu 14 64-bit	64	ubuntu	ubuntu.14.04.64	debian	14.04	linux	can
ubuntu 15	32	ubuntu	ubuntu.15.10	debian	15.1	linux	can
ubuntu 15 64-bit	64	ubuntu	ubuntu.15.10.64	debian	15.1	linux	can
ubuntu 16	32	ubuntu	ubuntu.16.04	debian	16.04	linux	can
ubuntu 16 64-bit	64	ubuntu	ubuntu.16.04.64	debian	16.04	linux	can
ubuntu 18.04	32	ubuntu	ubuntu.18.04	debian	18.04	linux	can
ubuntu 18.04 64-bit	64	ubuntu	ubuntu.18.04.64	debian	18.04	linux	can
ubuntu 20.04	32	ubuntu	ubuntu.20.04	debian	20.04	linux	can
ubuntu 20.04 64-bit	64	ubuntu	ubuntu.20.04.64	debian	20.04	linux	can
ubuntu 64-bit	64	ubuntu	ubuntu.64	debian	all	linux	can
unknown	64	other	unknown	NULL	all	unknown	unk
windows	64	windows	windows	windows	all	windows	mic
windows 10	32	windows	windows.10	windows	10	windows	mic
windows 10 64-bit	64	windows	windows.10.64	windows	10	windows	mic
windows 7	32	windows	windows.7	windows	7	windows	mic
windows 7 64-bit	64	windows	windows.7.64	windows	7	windows	mic
windows 8	32	windows	windows.8	windows	8	windows	mic
windows 8 64-bit	64	windows	windows.8.64	windows	8	windows	mic
windows server 2003	64	windows	windows.server.2003	windows	2003	windows	mic
windows server 2008	64	windows	windows.server.2008	windows	2008	windows	mic
windows server 2008 R2	64	windows	windows.server.2008.r2	windows	2008	windows	mic
windows server 2012	64	windows	windows.server.2012	windows	2012	windows	mic
windows server 2016	64	windows	windows.server.2016	windows	2016	windows	mic
windows server 2019	64	windows	windows.server.2019	windows	2019	windows	mic
xen server 6.1	64	xen	xenserver.6.1	xen	6.1	linux	xen
xen server 6.2	64	xen	xenserver.6.2	xen	6.2	linux	xen
xen server 6.5	64	xen	xenserver.6.5	xen	6.5	linux	xen
xen server 7.0	64	xen	xenserver.7.0	xen	7	linux	xen

Note: Other Operating System types may be supported but are not tested.

Agent Installation

There are many methods to install the Morpheus Agent on supported targets. All Agent installation methods are executing a script on the target that calls back to the Morpheus appliance over port 443.

Important: All Agent installation methods require the Target (VM or Host) to resolve and reach the appliance URL over port 443. In addition to the main Appliance URL (in Administration > Settings), additional Appliance URLs can be set per cloud in the Advanced Options section of the Create/Edit Cloud modal. When this field is populated, it will override the main Appliance URL for anything provisioned into that Cloud.

Basic Installation Steps

1. An Agent installation method is used to get the install script onto the target VM or Host
2. The Agent installation script is executed on the target VM or Host, installing the Agent and all dependencies
3. The Agent is started and makes a websocket connection to the configured Appliance URL over port 443 using the target VM or Host API key

It is important to note these are three separate steps with varying requirements. The first requires a path to get the script on the target. The second requires successful execution of the Agent installation script. The third requires a successful websocket connection. Requirements for the successful execution of each step are covered in the sections below.

Tip: The Morpheus UI current log, located at `/var/log/morpheus/morpheus-ui/current`, is very helpful when troubleshooting Agent installations.

Agent Install Modes

Agent Installation Method is determined by:

- The AGENT INSTALL MODE setting on target Cloud: - SSH / WinRM / Guest Execution - Cloud Init / Unattend (when available)
- Platform / OS type on Virtual Image or target (VM or Host)
- Virtual Image configuration
- RPC Mode setting (VMware Clouds only)

Agent Installation Methods

The Morpheus Agent can be installed with a variety of automated methods. It is important to note these methods simply get the Agent install script to the target. Successful execution of the Agent install script is not directly related to the Agent install method.

SSH Morpheus makes an SSH connection to the VM or Host, CURLs, and executes the Agent install script:

```
curl -k -s "https://${applianceUrl}/api/server-script/agentInstall?
apiKey=${apiKey}" | bash
```

WinRM Morpheus makes a WinRM connection to the VM or Host and executes the Agent install script

VMware Tools Morpheus executes agentInstall.sh or agentInstall.ps1 via VMware Tools Guest Execution

Cloud-Init Morpheus executes agentInstall.sh via cloud-init runcmd

Cloudbase-Init Morpheus adds WindowsAgentCloudInitInstallScript to CloudbaseInitUserData

Windows Unattend Morpheus adds getWindowsAgentDownloadScript to unattend.xml (RunSynchronousCommand)

Manual User executes agentInstall.sh or agentInstall.ps1 manually on the VM or Host. These scripts can be obtained on the VM or Host detail page via Actions > Download Agent Script

SSH

Process

Morpheus makes an SSH connection to the VM or Host, CURLs, and executes the Agent install script:

```
curl -k -s "https://${applianceUrl}/api/server-script/agentInstall?apiKey=${apiKey}" | bash
```

Requirements

- Port 22 is open for Linux images, and SSH is enabled
- Credentials have been entered on the image if using a custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section

WinRM

Process

Morpheus makes a WinRM connection to the VM or Host and executes the Agent install script

Requirements

- Port 5985 must be open and winRM enabled for Windows images
- Credentials have been entered on the image if using a custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section
- Administrator User (SID 500) is required for Windows Agent install

VMware Tools

Process

Morpheus executes agentInstall.sh or agentInstall.ps1 via VMware Tools Guest Execution

Requirements

- VMware Tools is installed on the template(s)
- Credentials have been entered on the Image if using an uploaded or synced image when Cloud-init, Guest Customizations, or Sysprep for Windows are not used. Credentials can be entered on Images in the Provisioning > Virtual Images section
- Administrator User (SID 500) is required for Windows Agent install.

Cloud-Init

Process

Morpheus executes agentInstall.sh via Cloud-Init runcmd

Requirements

- Cloud-Init is installed on Virtual Image
- “IS CLOUD INIT ENABLED?” is checked (true) on the Morpheus Virtual Image record
- Cloud-Init User is configured in the Admin > Provisioning section

Cloudbase-init

Process

Morpheus adds WindowsAgentCloudInitInstallScript to CloudbaseInitUserData

Requirements

- Cloudbase-Init is installed on the Virtual Image
- “IS CLOUD INIT ENABLED?” is checked (true) on the Morpheus Virtual Image record
- Windows Administrator password defined in the Administration -> Provisioning section

Windows Unattend

Process

Morpheus adds getWindowsAgentDownloadScript to unattend.xml (RunSynchronousCommand)

Requirements

VMware

- Windows Administrator password defined in the Administration > Provisioning section OR Administrator User (SID 500) and valid Windows password are defined on the Morpheus Virtual Image record
- “FORCE GUEST CUSTOMIZATION?” is checked (true) on the Morpheus Virtual Image record when using DHCP
- “IS CLOUD INIT ENABLED?” is unchecked (false) on the Morpheus Virtual Image record

Nutainx/SCVMM/Openstack

- Windows Administrator password defined in the Administration > Provisioning section OR Administrator User (SID 500) and valid Windows password are defined on the Morpheus Virtual Image record
- “ENABLED SYSPREP?” is checked (true) on the Morpheus Virtual Image record
- “IS CLOUD INIT ENABLED?” is unchecked (false) on the Morpheus Virtual Image record

Manual

Process

- From the VM or Host record page (/infrastructure/servers/\${id}) run *ACTIONS* -> Download Agent Script
- This will generate an Agent install script based off the target OS and platform, Appliance URL, and API key
- Manually execute the downloaded script on the Target VM or Host

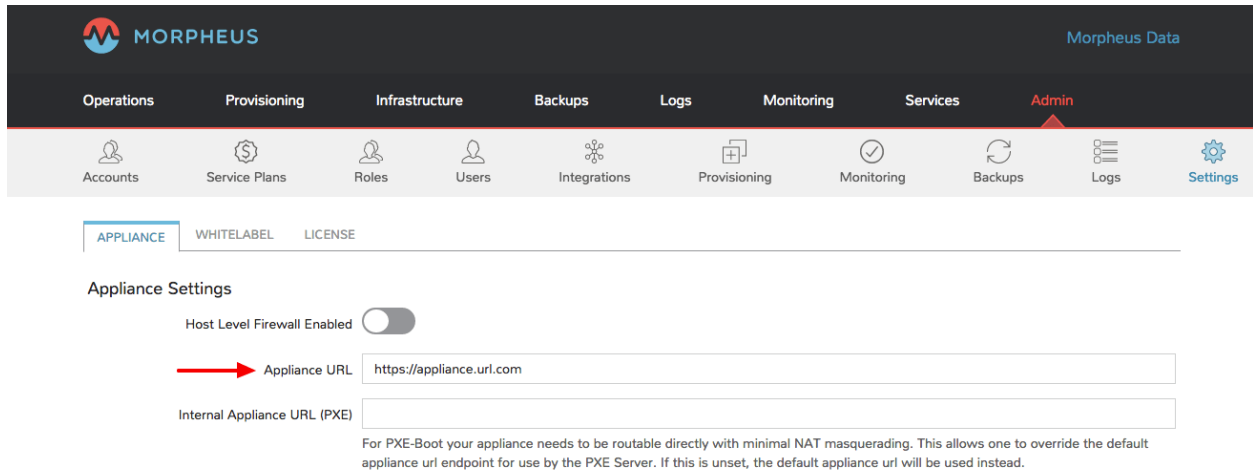
Agent Install Requirements

Agent Installation Requirements							
Requirement	Agent Installation Method						
	SSH	WINRM	VMWARE TOOLS	CLOUD-INIT	CLOUD-INIT	BASE-TEND	BASE-MANUAL
Target (vm/host) to resolve and reach Appliance URL over 443	YES	YES	YES	YES	YES	YES	YES
Target (vm/host) to resolve and reach Appliance URL over 80	Ubuntu 14.04 Only		Ubuntu 14.04 Only	Ubuntu 14.04 Only			Ubuntu 14.04 Only
Websockets enabled when using load balancer	YES	YES	YES	YES	YES	YES	YES
Access to Target VM/Host on port 22	YES	NO	NO	NO	NO	NO	NO
Access to Target VM/Host on port 5985	NO	YES	NO	NO	NO	NO	NO
Vmware Tools installed and flagged on Virtual Image	NO	NO	YES	NO	NO	YES	NO
Syspreped Image and Sysprep Enabled flagged on Virtual Image (Nutanix, Openstack, SCVMM)	NO	NO	NO	NO	NO	YES	NO
Force Guest Customizations flagged on Virtual Image	NO	NO	DHCP	NO	NO	DHCP	NO
Cloud-Init installed and flagged on Virtual Image	NO	NO	NO	YES	YES	NO	NO
Global Cloud-Init user populated in /admin/provisioning	NO	NO	NO	YES	NO	NO	NO
Windows Administrator Password populated in /admin/provisioning	NO	NO	NO	NO	YES	YES	NO
Access to configured YUM or APT repos	NO but will cause delay in Agent Install	N/A	NO but will cause delay in Agent Install	NO but will cause delay in Agent Install	N/A	N/A	NO but will cause delay in Agent Install
.net >=4.5.2 (Windows, Morpheus >= 4.1.2)	N/A	YES	YES	N/A	YES	YES	YES
User with Sudo Access set on Virtual Image (Greenfield)	YES	N/A	YES	NO	N/A	N/A	N/A
Administrator User (SID 500) set on Virtual Image (Greenfield)	N/A	YES	YES	N/A	NO	N/A	N/A
User with Sudo Access set on VM/Host Record (Brownfield)	YES	N/A	YES	N/A	N/A	N/A	N/A
Administrator User (SID 500) set on VM/Host Record (Brownfield)	N/A	YES	YES	N/A	N/A	N/A	N/A

When provisioning an Instance, there are network and configuration requirements to consider in order to successfully

install the Morpheus Agent. Typically, when a VM Instance is still in the provisioning phase long after the VM is up, the Instance is unable to reach Morpheus. Depending on the Agent install mode, it could also mean Morpheus is unable to reach the Instance.

The most common reason an Agent install fails is the provisioned Instance cannot reach the Morpheus Appliance via the Appliance URL set in Administration > Settings over port 443. When an Instance is provisioned from Morpheus, it must be able to reach the Morpheus appliance via the Appliance URL or the Agent will not be installed.



In addition to the main Appliance URL in Administration > Settings, additional Appliance URLs can be set per Cloud in the Advanced Options section of the Cloud configuration modal when creating or editing a Cloud. When this field is populated, it will override the main Appliance URL for anything provisioned into that Cloud.

Tip: The Morpheus UI current log, located at `/var/log/morpheus/morpheus-ui/current`, is very helpful when troubleshooting Agent installations.

Agent Install Methods

Morpheus Agent installation supports multiple install methods.

- SSH/WinRM
- VM Tools
- Cloud-Init & Cloudbase-Init
- Windows Unattended
- Manual

For All Agent Install Methods

When an Instance is provisioned and the Agent does not install, verify the following for any Agent install mode:

- The Morpheus Appliance URL (Administration > Settings) is both reachable and resolvable from the provisioned node
- The Appliance URL begins with <https://>, not <http://>

Note: Be sure to use <https://> even when using an IP address for the appliance.

- Inbound connectivity access to the Morpheus appliance from provisioned VMs and container hosts on port 443 (needed for Agent communication)
- Private (non-Morpheus provided) VM images and templates must have their credentials stored. These can be entered or edited in the Provisioning > Virtual Images section by clicking the Actions dropdown on an imaged detail page and selecting Edit.

Note: Administrator user is required for Windows Agent install.

- The Instance does not have an IP address assigned. For scenarios without a DHCP server, static IP information must be entered by selecting the Network Type: Static in the Advanced Options section during provisioning. IP Pools can also be created in the Infrastructure > Networks > IP Pools section and added to Cloud network sections for IPAM
- DNS is not configured and the node cannot resolve the appliance. If DNS cannot be configured, the IP address of the Morpheus appliance can be used as the main or Cloud appliance

SSH

- Port 22 is open for Linux images, and SSH is enabled
- Credentials set on the image if using a custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section

WinRM

- Port 5985 must be open and WinRM enabled for Windows images
- Credentials have been entered on the image if using a custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section

Note: Administrator user is required for Windows Agent install.

VMware Tools (vmtools)

- VMware Tools is installed on the template(s)
- Credentials have been entered on the image if using custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section
- Sudo privileges required for Linux
- Administrator User required for Windows (SID 500)

Cloud-Init

- Cloud-Init settings configured in Administration > Provisioning section
- Cloud-Init installed on Virtual Image
- Cloud-Init enabled on Virtual Image config

Cloudbase-Init

- Windows Administrator Password defined in Administration > Provisioning section
- Cloudbase-Init installed on Virtual Image
- Cloud-Init enabled on Virtual Image config

Note: Unattend Agent Installation and customizations are recommended over Cloudbase-Init

Windows Unattended

- Windows Administrator Password defined in Administration > Provisioning section
- VMware: Force Guest Customizations set to forced on Virtual Image config when using DHCP (Static Assignment will already force Guest Customizations)
- Nutanix & SCVMM: Virtual Image is sysprepped and shutdown, Sysprep Enabled flagged on Virtual Image config

Manual

- Agent Install scripts can be downloaded from Morpheus and ran manually on the target host when required via Actions -> Download Agent Script on the managed Resource. Please note the script will be unique per managed Resource.

Restarting the Morpheus Agent

In some situations, it may necessary to restart the Morpheus Agent on the host to re-sync communication from the Agent to the Morpheus appliance.

Linux

On the target host, run `sudo morpheus-node-ctl restart morphd` and the Morpheus agent will restart. `morpheus-node-ctl status` will also show the agent status.

Windows

The Morpheus Windows Agent service can be restarted in Administrative Tools -> Services.

Tip: The Morpheus Remote Console is not dependent on Agent communication and can be used to install or restart the Morpheus agent on an Instance.

Uninstall Morpheus Agent

Linux Agents

You can use the following to uninstall the linux agent (contains commands for both rpm and deb agents)

```
sudo rm /etc/apt/sources.list.d/morpheus.list \  
sudo morpheus-node-ctl kill \  
sudo apt-get -y purge morpheus-node \  
sudo apt-get -y purge morpheus-vm-node \  
sudo yum -y remove morpheus-node \  
sudo yum -y remove morpheus-vm-node \  
sudo yum clean all \  
sudo systemctl stop morpheus-node-runsvdir \  
sudo rm -f /etc/systemd/system/morpheus-node-runsvdir.service \  
sudo systemctl daemon-reload \  
sudo rm -rf /var/run/morpheus-node \  
sudo rm -rf /opt/morpheus-node \  
sudo rm -rf /etc/morpheus \  
sudo rm -rf /var/log/morpheus-node \  
sudo pkill runsv \  
sudo pkill runsvdir \  
sudo pkill morphd \  
sudo usermod -l morpheus-old morpheus-node \  

```

Windows Agents

```
$app = Get-WmiObject -Class Win32_Product
                -Filter "Name = 'Morpheus Windows Agent'"
$app.Uninstall()
```

CentOS/RHEL 7 Images

For custom CentOS 7 images we highly recommend setting up Cloud-Init and fixing the network device names. More information for custom CentOS images can be found in the CentOS 7 image guide.

Communication Data

The following page contains communication information between the Morpheus appliance, integrated technologies, managed machines, and services.

Communication Frequency and Configurability

The following table contains communication information, including frequency and configurability between Morpheus and its supported technology integrations.

Source	Push/Pull	Destination
Cloud Foundry App Check	Server Pull	Cloud Foundry Applications that exist within Morpheus
Docker Container Check	Server Pull	Docker containers that exist within Morpheus
Elastic Search Check	Server Pull	Elastic Search application
Microsoft SQL Server Check	Server Pull	Microsoft SQL application
Mongo Check	Server Pull	Mongo DB application
MySQL Check	Server Pull	MySQL application
Postgres Check	Server Pull	Postgres application
Push API Check	Client Push	Morpheus API
Rabbit MQ Check	Server Pull	Rabbit MQ application
Redis Check	Server Pull	Redis application
Riak Check	Server Pull	Riak application
SNMP Check	Server Pull	SNMP
Socket Check	Server Pull	Web Socket
Virtual Machine Check	Server Pull	Virtual Machine that exists within Morpheus
Web Check	Server Pull (GET) or Server Push (POST)	Web application
Public Cloud Integration	Server Pull	Alibaba Cloud
Public Cloud Integration	Server Pull	Amazon AWS
Public Cloud Integration	Server Pull	Amazon AWS GovCloud
Public Cloud Integration	Server Pull	DigitalOcean
Public Cloud Integration	Server Pull	Google Cloud Platform
Public Cloud Integration	Server Pull	Huawei Cloud
Public Cloud Integration	Server Pull	IBM Cloud
Public Cloud Integration	Server Pull	Microsoft Azure
Public Cloud Integration	Server Pull	Open Telekom Cloud
Public Cloud Integration	Server Pull	Oracle Public Cloud

Source	Push/Pull	Destination
Public Cloud Integration	Server Pull	UpCloud
Public Cloud Integration	Server Pull	VMware on AWS
Private Cloud Integration	Server Pull	Cisco UCS Manager
Private Cloud Integration	Server Pull	Dell EMC
Private Cloud Integration	Server Pull	HPE
Private Cloud Integration	Server Pull	HPE OneView
Private Cloud Integration	Server Pull	KVM
Private Cloud Integration	Server Pull	MacStadium
Private Cloud Integration	Server Pull	Microsoft Azure Stack
Private Cloud Integration	Server Pull	Microsoft Hyper-V
Private Cloud Integration	Server Pull	Microsoft SCVMM
Private Cloud Integration	Server Pull	Nutanix Acropolis
Private Cloud Integration	Server Pull	Openstack
Private Cloud Integration	Server Pull	Oracle VM
Private Cloud Integration	Server Pull	Pivotal Cloud Foundry
Private Cloud Integration	Server Pull	Supermicro
Private Cloud Integration	Server Pull	Vmware vCloud Director
Private Cloud Integration	Server Pull	Vmware ESXi
Private Cloud Integration	Server Pull	VMware Fusion
Private Cloud Integration	Server Pull	VMware vCenter
Private Cloud Integration	Server Pull	Xen Server
Automation Integration		Ansible
Automation Integration	Server Pull	Ansible Tower
Automation Integration	Server Pull	Chef
Automation Integration	Server Pull	Puppet
Automation Integration	Server Pull	Salt
Automation Integration		Terraform
Automation Integration	Server Pull	vRealize Orchestrator
Backup Integration	Server Pull	Commvault
Backup Integration	Server Pull	Veeam
Backup Integration	Server Pull	Rubrik
Backup Integration	Server Pull	Zerto
Backup Integration	Server Pull	Avamar
Build Integration	Server Pull	Jenkins
Container Integration	Server Pull	Docker
Container Integration		Docker Registry
Container Integration	Server Pull	Kubernetes
Deployment Integration	Server Pull	Git/Github
DNS Integration	Server Pull	AWS Route53
DNS Integration	Server Pull	Microsoft DNS
DNS Integration	Server Pull	PowerDNS
Identity Management Integration	Server Pull	Microsoft AD
Identity Management Integration	Server Pull	OneLogin
Identity Management Integration	Server Pull	Okta
Identity Management Integration	Server Pull	Jump Cloud
Identity Management Integration	Server Pull	LDAP
Identity Management Integration	Server Pull	SAML
IPAM Integration	Server Pull	Infoblox
IPAM Integration	Server Pull	phpIPAM

Source	Push/Pull	Destination
IPAM Integration	Server Pull	Bluecat
IPAM Integration	Server Pull	SolarWinds
ITSM Integration	Server Pull	ServiceNow
ITSM Integration	Server Pull	Cherwell
ITSM Integration	Server Pull	Remedy
Key & Certificate Integration	Server Pull	Venafi
Load Balancer Integration	Server Pull	AzureLB
Load Balancer Integration	Server Pull	F5 BigIP
Load Balancer Integration	Server Pull	Citrix NetScaler
Logging Integration		LogRhythm
Logging Integration		Splunk
Logging Integration		Syslog
Monitoring Integration	Server Pull	ServiceNow
Monitoring Integration		AppDynamics
Monitoring Integration		NewRelic
Network Integration	Server Pull	NSX-T
Network Integration	Server Pull	NSX-V
Network Integration	Server Pull	Cisco ACI
Network Integration	Server Pull	Unisys Stealth
Service Discovery Integration		Consul
Storage Integration	Server Pull	3Par
Storage Integration	Server Pull	Azure Storage
Storage Integration	Server Pull	Dell ECS
Storage Integration	Server Pull	Isilon
Morpheus Agent	Agent Pull	Application Tier

Ports and Protocols

The following table contains communication port and protocol data between Morpheus appliance tiers, managed machines, and services. All communication to and from Morpheus goes thru the application tier with exception of inter-cluster communications for each of the Morpheus tiers when using a distributed architecture.

Ports used to communicate with integrated technologies are those defined for the integration's API. They are not represented in this table as many of these are configurable and may be different in each customer environment. Additionally, ports used to complete Morpheus checks are customizable and may vary for each check configured. They are also not represented in this table.

Table 6: **Ports and Protocols**

Source	Destination	Port	Protocol	Description
User	Application Tier	443	TCP	User Access
Morpheus Servers	DNS Servers	53	TCP	Domain Name Resolution
Morpheus Servers	Time Source	123	TCP	Time Resolution
Morpheus Servers	Web or Offline Installer	80, 443	TCP	Download repos and Morpheus
Managed Machine	Application Tier	443	TCP	Morpheus Agent Communicat
Managed Machine	Application Tier	80, 443	TCP	Agent Installation. (Requires p
Managed Machine	Application Tier	N/A	N/A	Agent Installation Clout-init (I
Managed Machine	Application Tier	N/A	N/A	Agent Installation Cloudbase-i
Managed Machine	Application Tier	N/A	N/A	Agent Installation VMtools
Managed Machine	Application Tier	N/A	N/A	Static IP Assignment & IP Poc

Table 6 – continued from previous page

Source	Destination	Port	Protocol	Description
Managed Machine	Docker Image Repo	443	TCP	Applicable if using docker
Managed Machine	Application Tier	69	TCP/UDP	PXE Boot (Forwarded to inter)
Application Tier	Managed Machine	5985	TCP	Agent Installation WinRM (W)
Application Tier	Managed Machine	22	TCP	Agent Installation SSH (Linux)
Application Tier	Managed Machine	22, 3389, 443	TCP	Remote Console (SSH, RDP, F)
Application Tier	AWS S3	443	TCP	Morpheus Catalog Image Down
Application Tier	Hypervisor	443	TCP	Hypervisor hostname resolvab
Application Tier	Non- Transactional Database Tier	443	TCP	Applicable if using Amazon E
Application Tier	Docker CE Repo	443	TCP	Applicable only when integrat
Application Tier	Rubygems	443	TCP	
Application Tier	Morpheus Hub	443	TCP	(Optional) Telemetry data (Dis
Application Tier	Mail Server	25 or 465	SMTP	Send email from Morpheus
Application Tier	postmarkapp	2525	TCP	Send email from Morpheus (su
Application Tier	Messaging Tier	5672	TCP	AMQP non-TLS connections
Application Tier	Messaging Tier	5671	TCP	AMQPS TLS enabled connect
Application Tier	Messaging Tier	61613	TCP	STOMP Plugin connections (R
Application Tier	Messaging Tier	61614	TCP	STOMP Plugin TLS enabled c
Messaging Tier	Messaging Tier	25672	TCP	Inter-node and CLI tool comm
Administrator Web Browser	RabbitMQ Server Management	15672	TCP	Management plugin
Administrator Web Browser	RabbitMQ Server Management	15671	TCP	Management plugin SSL
Messaging Tier Cluster Node	Messaging Tier Cluster Node	4369	TCP	erlang (epmd) peer discovery s
Application Tier	Non-Transactional Database Tier	9200	TCP	Elasticsearch requests (Used in
Non-Transactional Database Tier	Non-Transactional Database Tier	9300	TCP	Elasticsearch Cluster
Transactional Database Tier	Transactional Database Tier	4567	TCP/UDP	Write-set replication traffic (ov
Transactional Database Tier	Transactional Database Tier	4568	TCP	Incremental State Transfer (IS
Application Tier	Transactional Database Tier	3306	TCP	MySQL client connections
Backup Solution	Transactional Database Tier	4444	TCP	State Snapshot Transfer (SST)
Application Tier	Integrated Technology	Varies	TCP	Integrations (Uses the port of t

VMware Support Statement

Morpheus Data, LLC will support customers who run Morpheus products on supported operating systems, irrespective of whether they are running in VMware environments or not. Morpheus Data, LLC supports operating systems, not specific hardware configurations. Accordingly, VMware operates as a hardware abstraction layer.

VMware supports a set of certified operating systems and hardware. The customer and VMware will be responsible for any interactions or issues that arise at the hardware or operating system layer as a result of their use of VMware.

Morpheus Data, LLC will not require clients to recreate and troubleshoot every issue in a non-VMware environment; however, Morpheus Data, LLC does reserve the right to request our customers to diagnose certain issues in a native certified operating system environment, operating without the virtual environment. Morpheus Data, LLC will only make this request when there is reason to believe that the virtual environment is a contributing factor to the issue.

Any time spent on investigation of problems that may, in the sole opinion of Morpheus Data, LLC be related to VMware, will be handled in the following fashion:

1. Morpheus Data, LLC will provide standard support to all Morpheus products.
2. If a problem is encountered while Morpheus is/are running in a VMware environment, the client may be required to recreate the problem on a non-VMware server unit, at which time Morpheus Data, LLC will provide regular support.

3. The client can authorize Morpheus Data, LLC to investigate the VMware-related items at normal time and materials rates. If such investigation shows that the problem is VMware related, the client may contract Morpheus Data, LLC to provide a software change to resolve the issue if such a resolution is possible.
4. Regardless of the problem type or source, if the problem is determined to be a non VMware-related issue, time spent on investigation and resolution will be covered as part of regular maintenance and support will be provided as usual.

1.3.2 Provisioning

There are several capabilities in the Morpheus provisioning engine. Things ranging from application / service deployments via containers, virtual machines, and even bare metal. Deployment management and app template construction are also core aspects of the provisioning engine. Take advantage of custom tasks and workflows within any environment by building tasks and workflows from those tasks. There is a lot of information to cover with regards to provisioning but Morpheus makes it intuitive and smooth.

Requirements

Provisioning Instances and Apps typically involves many steps beyond starting a workload. Morpheus is centered around automating everything desired for your application to be fully operational, including networking, storage, hostnames, domains, dns, licenses, scripts/automation, scaling, load balancers, security, accessibility, governance, auditing, monitoring, backups, costs, sizing and on and on. Point being there is a lot that goes on when spinning up an instance or app, and to make the magic happen a few requirements need to be met.

Important: By default, Agent Installation is enabled when provisioning unless deselected on the Virtual Images or *SKIP AGENT INSTALL* is selected when provisioning.

VM Provision Steps

While an infinite number of steps can happen when provisioning an Instance or App using a VM(s) in Morpheus, the basic order is:

- Look for Virtual Image Morpheus will check if the Virtual Image set on the Node Type or selected during provisioning is already available in the source Cloud. If not and it is an Uploaded/Local Image, Morpheus will attempt to upload the Image to the target Cloud.

Upload Image

For Uploaded/Local Images that do not exist in the target cloud, Morpheus will need to upload the Image.

Ensure the Virtual Image is valid for the target Cloud, the Image meets the target cloud upload requirements, and Morpheus has network access and permissions to upload the image.

Note: When uploading an image to a VMware Cloud, the Virtual Image is copied directly to the target ESXi host, NOT through the vCenter server. Ensure the Morpheus Appliance(s) can resolve target ESXi hostnames and connect on port 443 for successful vmdk/ova uploads.

Clone Image Once the Image is confirmed available in the target cloud, Morpheus will clone the Image to the target Datastore.

Note: The target host must have access to the target Datastore of the Image

- Reconfigure Image Once cloned Morpheus will resize the Image based off provisioning parameters
- **Cloud-init (if enabled)**

Attached cloud-init iso When using cloud-init, Morpheus will attach a tiny metadata iso to new VM. Network, Machine, User and any other cloud-init metadata will be sourced from this iso.

VM Tools Morpheus will run Guest Customizations via VMware VM Tools, including network config when assigning static IP's.

- Wait for Power On status and Network info Morpheus will wait to hear back from the target cloud/hypervisor that the VM has successfully started and has an IP address.

Note: If `VM TOOLS INSTALLED?` is NOT checked on the source Virtual Image configuration, Morpheus will skip waiting for network.

- **Finalize** By default this will include Agent Installation and any post-provision scripts or workflows or integration automation steps.

Important: If the VM is stuck in finalize for long periods of time, this typically means the Agent cannot be installed or has not been heard back from. This will result in a ! warning Instance status upon provisioning completion.

If agent installation is not possible or desired, uncheck “Install Agent” on the source Virtual Image configuration or select “Skip Agent Install” during provisioning to speed up provisioning completion.

Virtual Images

While containers are the future, the most common provisioning method involves Virtual Machines, and the most important part of Provisioning a VM is the Virtual Image. When provisioning a VM, Morpheus will need to do a few things depending on the location of the Virtual Image and if agent install, console access, and script execution is desired.

Synced Images need to be properly configured Morpheus gathers as much metadata for synced images as possible, but depending on the cloud, os, image configuration, agent install settings, by default the synced Virtual Images may not be ready to provision until configured. The Virtual Image is already at the target Cloud, but datastore selection, credentials, cloud-init settings, and networks and security settings on the Virtual Image can cause provisioning issues.

Local/Uploaded Virtual Images Images uploaded to Morpheus are configured during the *Add Virtual Image* process, however Morpheus in most scenarios will still need to copy the Virtual Image to the target Hypervisor/Cloud upon the first provision to the target Cloud. In addition to the requirements for provisioning a synced Virtual Image, copying an uploaded Virtual Image to the target Cloud upon is required and network and image configurations can cause upload failures, resulting in provisioning issues.

Marketplace Images AWS and Azure marketplace Images can be provisioned using the generic Amazon or Azure Instance Types, or added as Virtual Images as scoped to Node Types for custom Instance Types. Marketplace items provisioned/added to Morpheus still fall upon the requirements of the target Cloud, such as matching the region with the Image and licensing.

Synced Images

When a Cloud is added to Morpheus, all available Images/Templates records from that Cloud will be synced in regardless of Inventory settings on the Cloud. These Image records will be available in the Virtual Images section and can be provisioned by using the target clouds generic Instance Type, ie VMware, Amazon, Azure, Openstack etc Instance Types, or by creating custom Instance Types and selecting the Image on a Node Type.

Note: Synced Virtual Images are just meta-data records in Morpheus pointing to the Image in the target Cloud. The actual Image files are not copied/imported to Morpheus.

Before provisioning a synced Virtual Images, ensure the image is configured properly:

Name Name of the Virtual Image in Morpheus . This can be changed from the name of the Image, but editing will not change the name of the actual Image.

Operating System Specifies the Platform and OS of the image. All Windows images will need to have Operating System specified on the Virtual Image, as Morpheus will assign Linux as the Platform for all Images without Operating System specified.

Minimum Memory The Minimum Memory setting will filter available Service Plans options during provisioning. Service Plans that do not meet the Minimum Memory value set on the Virtual Image will not be provided as Service Plan choices.

Cloud Init Enabled? On by default, uncheck for any Image that does not have Cloud-Init or Cloudbase-Init installed.

Important: Provisioning a Virtual Images that has *Cloud Init Enabled?* checked on the Virtual Record in Morpheus but does not have cloud-init install will result in immediate provisioning failure.

Install Agent On by default, uncheck to skip Agent install. Note this will result in the loss of utilization statistics, logs, script execution, and monitoring. (Some utilization stats are collected for agent-less hosts and vm's from VMware and AWS clouds).

Username Existing Username on the Image. This is required for authentication, unless Morpheus is able to add user data, Cloud-Init, Cloudbase-Init or Guest Customizations. If Cloud-Init, Cloudbase-Init Guest Customizations or Nutanix Sysprep are used, credentials are defined in *Administration -> Provisioning and User Settings`*. *If credentials are defined on the Image and Cloud-Init is enabled, [morpheus] will add that user during provisioning, so ensure that user does not already exist n the image (aka ``root`)*. For Windows Guest Customizations, Morpheus will set the Administrator password to what is defined on the image if Administrator user is defined. Do not define any other user than Administrator for Windows Images unless using Cloudbase-init. Morpheus recommends running Guest Customizations for all Windows Images, which is required when joining Domains as the SID will change.

Password Password for the Existing User on the image if Username is populated.

Storage Provider Location where the Virtual Image will be stored. Default Virtual Image Storage location is /var/opt/morpheus/morpheus-ui/vms. Additional Storage Providers can be configured in *Infrastructure -> Storage*.

Cloud-Init User Data Accepts what would go in runcmd and can assume bash syntax. Example use: Script to configure satellite registration at provision time.

Permissions

Set Tenant permissions in a multi-tenant Morpheus environment. No impact on single-tenant environments.

Visibility

Private Image is only available in the specified Tenants below.

Public Image is available to all Tenants.

Tenant If Visibility is set to Private, specify Tenants the Image will be available for.

Auto Join Domain? Enable to have instances provisioned with this image auto-join configured domains (Windows only, domain controller must be configured in *Infrastructure -> Network* and the configured domain set on the provisioned to Cloud or Network).

VirtIO Drivers Loaded? Enable if VirtIO Drivers are installed on the image for provisioning to KVM based Hypervisors.

VM Tools Installed? On by default, uncheck if VMware Tools (including OpenVMTools) are not installed on the Virtual Image. Morpheus will skip network wait during provisioning when deselected.

Force Guest Customization? VMware only, forces guest customizations to run during provisioning, typically when provisioning to a DHCP network where guest customizations would not run by default. This is required for host/computer name definitions, domain joining, licenses and user definitions when using DHCP.

Trial Version Enable to automatically re-arm the expiration on Windows Trial Images during provisioning.

Enabled Sysprep? Applicable to Nutanix Only. Enable of the Windows Image has been sys-prepped. If enabled Morpheus will inject Unattend.xml through the Nutanix API (v3+ only)

Important: Provisioning a Virtual Images that has *Cloud Init Enabled?* checked on the Virtual Record in Morpheus but does not have cloud-init install will result in immediate provisioning failure.

Important: For Linux images without CCloud-Init, and existing username and password must be defined on the Virtual Image record for Agent Install, Domain joining, licensing, script execution and other automation, and ssh or RDP Console access.

Local Virtual Images

A Local Virtual Image means it has been uploaded to Morpheus. To provision, Morpheus will need to upload the Image to the target Cloud upon first provision.

- Ensure the Virtual Image is valid for the target Cloud, the Image meets the target cloud upload requirements, and Morpheus has network access and permissions to upload the image.

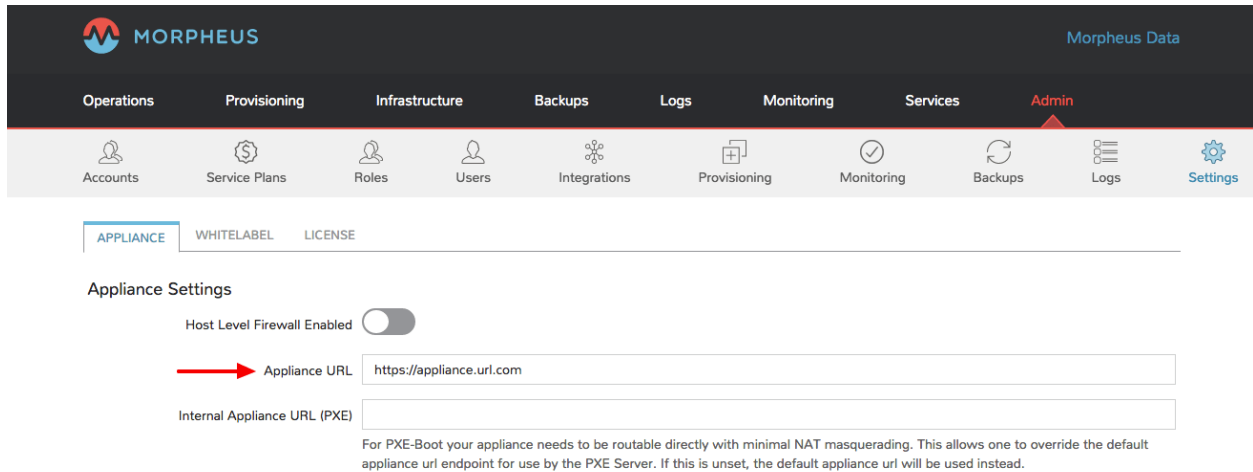
Note: When uploading an image to a VMware Cloud, the Virtual Image is copied directly to the target ESXi host, NOT through the vCenter server. Ensure the Morpheus Appliance(s) can resolve target ESXi hostnames and connect on port 443 for successful vmdk/ova uploads.

Once a Local Virtual Image has been uploaded to a Cloud, subsequent provisions will use the Image local to the cloud instead of uploading again as long as the copied image is still available in the source Cloud.

Agent Install

When provisioning an instance, there are some network and configuration requirements to successfully install the morpheus agent. Typically when a vm instance is still in the provisioning phase long after the vm is up, the instance is unable to reach Morpheus, or depending on agent install mode, Morpheus is unable to reach the instance.

The most common reason an agent install fails is the provisioned instance cannot reach the Morpheus Appliance via the `appliance_url` set in Admin -> Settings over both 443 and 80. When an instance is provisioned from Morpheus, it must be able to reach the Morpheus appliance via the `appliance_url` or the agent will not be installed.



In addition to the main `appliance_url` in Admin -> Settings, additional `appliance_urls` can be set per cloud in the Advanced options of the cloud configuration pane when creating or editing a cloud. When this field is populated, it will override the main appliance url for anything provisioned into that cloud.

Tip: The Morpheus UI current log, located at `/var/log/morpheus/morpheus-ui/current`, is very helpful when troubleshooting agent installations.

Agent Install Modes

There are 3 Agent install modes:

- ssh/winrm
- VMware Tools
- cloud-init

For All Agent Install modes

When an instance is provisioned and the agent does not install, verify the following for any agent install mode:

- The Morpheus `appliance_url` (Admin -> Settings) is both reachable and resolvable from the provisioned node.
- The `appliance_url` begins with `https://`, not `http://`.

Note: Be sure to use `https://` even when using an ip address for the appliance.

- Inbound connectivity access to the Morpheus Appliance from provisioned VM's and container hosts on port 443 (needed for agent communication)
- Private (non-morpheus provided) vm images/templates must have their credentials entered. These can be entered/edited in the Provisioning - Virtual Images section but clicking the Actions dropdown of an image and selecting Edit.

Note: Administrator user is required for Windows agent install.

- The instance does not have an IP address assigned. For scenarios without a dhcp server, static IP information must be entered by selecting the Network Type: Static in the Advanced section during provisioning. IP Pools can also be created in the Infrastructure -> Networks -> IP Pools section and added to clouds network sections for IPAM.
- DNS is not configured and the node cannot resolve the appliance. If dns cannot be configure, the ip address of the Morpheus appliance can be used as the main or cloud appliance.

SSH/Winrm

Linux Agent

- Port 22 is open for Linux images, and ssh is enabled
- Credentials have been entered on the image if using custom or synced image. Credentials can be entered on images in the Provisioning -> Virtual Images section.

Windows Agent

- Port 5985 must be open and winRM enabled for Windows images.
- Credentials have been entered on the image if using custom or synced image. Credentials can be entered on images in the Provisioning -> Virtual Images section.

Note: Administrator user is required for Windows agent install.

VMware tools (vmtools) rpc mode

- VMware tools is installed on the template(s)
- Credentials have been entered on the Image if using uploaded or synced image when Cloud-init or Guest Customizations or Sysprep for Windows are not used. Credentials can be entered on Images in the *Provisioning -> Virtual Images* section.

Cloud-Init agent install mode

- Cloud-Init is configured in Admin -> Provisioning section
- Provisioned image/blueprint has Cloud-Init (linux) or Cloudbase-Init (windows) installed

Provisioning Concepts

Morpheus is a powerful infrastructure agnostic Cloud Application Management Platform. As a result of this there are some differing concepts compared to other CMP platforms in the space. It is here that it is important to notice the qualification difference between Morpheus and other platforms.

Morpheus refers to itself as a CAMP (Cloud Application Management Platform) as opposed to a (Cloud Management Platform). While that may seem minor, it actually is a big deal. Many CMP based applications start at the IaaS layer and work up to the application layer (often needing additional PaaS) architectures to fill out the model. Morpheus was designed from a middle-ground perspective. As such some concepts are a bit different. This provides a more complete platform that allows for greater capabilities out of the box as will be seen when these concepts are covered.

Instances

Morpheus starts with provisioning Instances. In some platforms an Instance is representative of a singular object like a “Virtual Machine” in Amazon. In Morpheus, this concept was rethought. An Instance is more of a representation of a Resource or Service. This service may involve several virtual machines or even several docker containers.

For example, in the morpheus Instance wizard Mongo is an option and contains several “Instance Configurations”. One of these configurations is a full Mongo cluster consisting of either seven virtual machines or seven docker containers. Rather than representing these directly as seven individual “instances”, Morpheus groups them together into a singular instance of a service that contains multiple containers or virtual machines. This even allows for instance actions that can be performed to expand capacity on an instance (either horizontally or vertically). In the past, a database server may have been representative of a singular server, but this model has drastically changed in a big data world. This same concept also can apply to something like a simple Apache web server where there are 10 copies of a web server horizontally scaled out to handle traffic.

When viewing an instance detail page, one is able to look at details/statistics specific to a virtual machine or container. Morpheus simply helps simplify the management model for tracking these services.

Containers / Nodes / Virtual Machines

In relation to *Instances*, an instance can have many nodes. A node is a generic representation of a container or a virtual machine. In most cases, Morpheus will represent a node as a Container or Virtual Machine depending on the provisioning engine used for the instance. Node is just a generic naming representation when referring to these types of items. The public developer API, however, often refers to both virtual machines and docker containers as *Containers*. The UI was since updated to better delineate this concept for easier understanding but in essence the name is valid for both concepts of containerized environments as well as Virtual Machines. In fact, one can even think of a Docker Host as a Hypervisor (which we do).

Hosts / Servers

This concept is mostly tailored to users of morpheus responsible for managing and maintaining the underlying infrastructure integrations. A Host typically refers to a Docker Host in which a container in an instance is running, or a hypervisor virtual machines can be provisioned onto. A Server is the underlying general representation of a physical or virtual server. It could be a Host representation, a Virtual Machine, or even a Bare Metal delineation.

When a user provisions a vm based instance, a corresponding server record is created to represent the link to the actual resource via the underlying provisioning engine. This may seem a bit odd but provides an aspect of Morpheus that is quite powerful. This singular concept is what allows Morpheus to ingest “Brownfield” environments. We do not need to start clean. Morpheus can be integrated into existing environments and manage existing virtual machines. The way Morpheus does this is by periodically syncing existing vms from the added cloud integrations. A server record will be created and periodically updated (5 minutes typically) with realtime information and changes. This, in essence, provides CMDB based capabilities as well. When a server is discovered, the user (given the appropriate access) can convert the virtual machine to a managed instance. When this is done a corresponding Instance is made in the provisioning section of Morpheus and the Morpheus Agent can also optionally be installed to provide more refined guest operating system level statistics and logging.

Apps

On top of all the previous concept, Morpheus provides an Apps layer. An App is a collection of Instances linked together via application tiers. Tiers allow the user to define segregated sections of connectivity between the various elements / instances within an application. Once these instances are all linked together in an application concept, this may affect Instance environments and provide service discovery capabilities for them to cross connect. There are several service discovery aspects within morpheus as well as integrations with services like Consul.

Blueprints

A blueprint is typically referred to as an Application Blueprint. It allows a user to define an application structure for easy reproducibility and deployment into various environments. They can be used to mix and match various instance types to provision an application dependent on multiple layers of services.

Instances

Instances is a great starting point for taking advantage of self service features and spinning up both VM’s and containers. In Morpheus it may be advisable to cover the definition of a few terms used within the application so as to reduce confusion.

Instance A set of containers or virtual machines that can correlate to a single horizontally scalable entity or a service suite like a database. (It is important to note that an instance can contain one or more containers/vms depending on the instance type and configuration).

Container Typically a docker container provisioned via a Morpheus Docker host.

Virtual Machine A virtualized compute server provisioned onto various hypervisor hosts.

The top of the main Instances page shows overall statistic for the listed Instances, including count, status, and resource utilization. You can search for instances by name, or filter by group, instance type, or category.

Note: Instances listed are determined by group access and role permissions.

The Instance list contains important information about each instance, including the instance name, environment tag, instance type icon, ip and port info, instance version, the number of virtual machines or containers in the instance, the group the instance is in, and the cloud or clouds the instance is in.

Creating Instances

The instance catalog is the one stop shop for selecting items to be provisioned and pieced together. It contains not only basic container and vm options but also tailored services for SQL databases, NoSQL databases, cache stores, message busses, web servers, and even full fledged apps. The list contains a lot of items to choose from and they are represented to the user based on what provisioning engines are enabled and integrated in the Morpheus environment.

To get started, simply click the + *Add* button in the upper right of the `Provisioning -> Instances` section. A modal will display allowing the catalog to be searched. Once an item is selected it is just a matter of following the steps through the wizard.

Tip: The instance catalog can be customized via role based access control thereby restricting access to non sanctioned catalog items, as well as added to via the `Provisioning -> Library` section. It is completely customizable.

The next step will ask for a Group and Cloud to be selected. The Group is an abstract representation that can contain multiple cloud integrations. These cloud integrations can also be in multiple groups and is also useful for using role based access control to restrict provisioning access and set retainment policies. If the environment is new and these do not yet exist, It may be advisable to refer to the main section on Getting started by setting up some cloud integrations and infrastructure first. The wizard continues by allowing us to choose a name for the instance as well as an environment.

Note: Currently the Environment option is mostly useful for presenting the user with informative metadata around the instance when coming back to it later.

Moving on, it is now time to configure the Instance. Depending on the option that was chosen and the Instance Configuration that is chosen fields will change. This can include cloud specific fields (i.e. Datastore for VMware or Network). There will also be options like initial username. Some of these fields are optional and will be represented as such.

Configuration options provided in this screen are very powerful. An example is Mysql where a Master/Slave or Master/Master layout can be selected. These configurations will automatically deploy two MySQL VMs or containers and link them together to provide replication. These types of configurations exist for a wide range of instance types and are optimized for high performance and scale. It is even possible to provision entire sharded Mongo clusters.

One last step before the instance can be provisioned is the `Automation` step. This wizard step may or may not appear depending on the capabilities of the instance type or previous configurations in the account. It is here one can easily select a post provisioning workflow to run (see more on Tasks and Workflows), assign a load balancer, or even configure the backup job that gets created.

Now that the steps are completed for provisioning the selected instance type , simply review your selections and complete. The instance will automatically show up in the instances list and its provisioning state will be represented. Depending on what was provisioned this step can range from seconds to minutes (typically a container configuration will be rather quick if the instance type has previously been provisioned before).

Instance Details

The instance detail page is where you can view and fully manage an instance. To get to an instance detail page, navigate to provisioning, instances, and click on an instance. Please note instance details and actions differ between instance types and user permissions.

There are several sections within an Instance page that provide useful capabilities to the user.

Summary Stats and status information

Deploy Track deployment history for instance types that support deployments or manually kick off a deployment (only visible for instance types that support deployments)

Settings Some instance types support custom configuration settings (i.e. mysql presents the my.ini)

Network Useful for configuring security groups and access to the instance.

Monitoring Quick summary of the monitoring system and all checks that were configured to test the state of the instance

Backups Quick backup dashboard. Useful for viewing historical backups as well as kicking off new ones.

Logs View all aggregated logs from the containers or VM's representing the instance.

Environment View the environment variables presented to the instances or exported by the instances via Apps (more on this in the Apps section). Even see Imported environment variables that may be referenced by the running instance.

Scale For instances that support load balancing and auto scaling. Easily configure auto scaling thresholds and load balancer settings that pertain to a particular instance.

Console Access the instance or container via a client-less Console supporting SSH, RDP, VNC, and even hypervisor level remote consoles.

Managing Instances

Instance actions allow you to perform numerous management tasks on instances. The actions available depend on the instance type, hypervisor, roles permissions, and instance state.

Edit Edit the Name, Description, Environment, Group, Metadata, Tags, and Owner for the Instance.

Delete Deletes the Instance.

Important: Deleting an Instance will delete the actual VM's or Containers and cannot be undone, unless a Delayed Removal policy has been applied prior to the Deletion. To delete Instances without deleting associated VM's, delete the Instances VM record(s) from the Infrastructure section with "Remove Infrastructure" deselected and select "Remove Associated Instances" in the VM delete modal options. This will delete the records in Morpheus but leave the infrastructure in place.

Tip: You can change the owner of an instance easily by selecting the edit button and entering a new owner in the corresponding field.

Actions

Available options in the Actions dropdown can include:

Suspend Puts the VM in a suspended state without shutting down the OS.

Stop/Start/Restart Service Stops, Starts or Restarts the service associated with the Instance Type.

Stop/Start/Restart Server Stops, Starts or Restarts the Virtual Machine.

Import as Image Clones and exports VM in its current state to target Storage provider and adds Virtual Image Record with metadata matching the source Instance's configuration.

Clone to Image Clones and converts VM in its current state to image in the source Cloud and adds Virtual Image Record with metadata matching the source Instance's configuration.

Lock/Unlock Instance A locked instance cannot be deleted until it is unlocked.

Reconfigure The Reconfigure action allows service plan, disk, cpu, ram, networks and storage controller changes. Available options depend on the instance type and service plan configuration. Some resize actions require an instance restart.

Clone Creates a new Instance from the Instance at its current state.

Backup Immediately executes a backup of the Instance. Only available for Instances with backups enabled.

Run Workflow Presents workflow options and then immediately runs selected Workflow on the Instance. Workflows can be created in the Provisioning -> Automation section.

Run Script Presents Script options and immediately executes selected Script on the Instance. Scripts can be created in the Provisioning -> Library section.

Apply Template Presents Template options and immediately applies selected Template to the Instance. Templates can be created in the Provisioning -> Library section.

Add Node Adds an additional node to the configuration. Additional options and configurations are required in the add node wizard depending on instance configuration and type.

Eject Disk Ejects attached disk/iso.

Add Slave Adds a database slave in the Instance.

Change Master Changes the database Master node in an Instance.

Clone to Template (VMware) Creates a new VMware Template from the Instance with corresponding Morpheus Virtual Image record.

Tip: Scrolling down in the Actions dropdown may be necessary to see all options.

Performing Instance Actions

1. Select the Provisioning link in the navigation bar.
2. Click the Instance from the list of instances you wish to perform an action on.
3. Click the Actions drop down button and select an Action.

Notes

Every Instance has a Wiki section for adding useful information about the Instance. Wiki can be added by selecting the Wiki tab button on the bottom of Instance Detail pages. Instances with associated VMware VM's will bi-directionally sync Morpheus Instance Wiki content and VMware VM Notes. See the [Wiki](#) Section for more details.

Tip: Markdown Syntax is supported in Wikis.

Remote Console

Morpheus has a built in Remote Console for Instances, Hosts, Virtual Machines and Bare Metal. The following information reviews the Roles Settings, Protocols, and Requirements necessary to configure and troubleshoot Remote Console access.

Role Settings

User Role settings determine if the Console tab or `Open Console` Action appear for a user, and if a login prompt is presented or the user is automatically logged in when using the Console.

- **Remote Console (None, Provisioned, Full)**

None The user will not have access to remote console.

Provisioned The user will only have remote console access for Instances they provisioned.

Full The user will have remote console access for all instances they have access to.

- **Remote Console: Auto Login (No, Yes)**

No A login prompt will be present in the console for Linux platforms, and the main login screen will present for Windows platforms.

Yes Morpheus will automatically login to the remote console using the credentials defined on the VM or Host. For provisioned Instances, the credentials are defined either from the credentials defined on the Virtual Image used, added via cloud-init or VMware Tools using the global cloud-init settings (Administration - Provisioning) or the Linux or Windows settings defined in User Settings. For Instances created when converting a VM or Host to managed, the credentials are entered when converting to managed. These credentials can be changed by editing the underlying VM or Host of the Instance.

Note: If the credentials defined on the VM or Host are not valid, and the `Remote Console: Auto Login` Role setting is set to `Yes`, the console will not be able to connect and no console window or login prompt will be presented. The credentials on the underlying VM or Host must be edited or `Remote Console: Auto Login` Role setting can be set to `No` for a login prompt to present in the console. Credentials cannot be changed from an Instance view, only in the Infrastructure VM or Host view.

Protocols

Platform Type and Cloud Settings determines the protocol and port used for Remote Console connections.

- **SSH** The SSH protocol will be used for Linux and OSX platform types, and 22 is the default port used.
- **RDP** The RDP (Remote Desktop) protocol will be used for Windows platform types over port 3389 by default.
- **VNC** The VNC protocol will be used for all platform types in Clouds with the `Hypervisor Console` option enabled in cloud settings. VNC connection are made directly to the Hypervisor Host over port 443.

Note: Alternative ports can be configured per VM or Host by editing the VM or Host and editing the Port field in the RPC host section.

SSH

For all Linux and OSX platform types, Morpheus will use the SSH protocol via port 22 by default for Remote Console connections, unless the `Hypervisor Console`` option is enabled for VMware type clouds.

Morpheus will SSH using the username, password, RPC Host IP address and Port defined in the VM or Host record.

Default Requirements for SSH Connectivity

- SSH Enabled on the target VM or Host
- Port 22 incoming open on the target VM or Host firewalls and security groups from the Morpheus Appliance (not from the users IP address)
- An IP address defined on the VM or Host record that is routable from the Morpheus Appliance.
- Valid credentials defined on the VM or Host record in the RPC host field.
- *Remote Console* Role Permissions set to *Provisioned* or *Full* if the User provisioned the instance, or *Full* if the user did not provision the instance.

RDP

For all Windows platform types, Morpheus will use the RDP protocol via port 3389 by default for Remote Console connections, unless the `Hypervisor Console`` option is enabled for VMware type clouds.

Morpheus will RDP using the username, password, RPC Host IP address and Port defined in the VM or Host record.

Default Requirements for RDP Connectivity

- Remote Access enabled on the target VM or Host and Remote Desktop enabled in the Windows Firewall settings. If the VM or Host is on a different network than the Morpheus appliance, public access for Remote Desktop must be enabled in the Firewall settings.
- Port 3389 incoming open on the target VM or Host firewalls and security groups from the Morpheus Appliance (not from the users IP address)
- An IP address defined on the VM or Host record that is routable from the Morpheus Appliance.
- Valid credentials defined on the VM or Host record in the RPC host field.
- *Remote Console* Role Permissions set to *Provisioned* or *Full* if the User provisioned the instance, or *Full* if the user did not provision the instance.

Note: If *Remote Console: Auto Login* is set to *No* in a users Role permissions, *Allow connections only from computers running Remote Desktop with Network Level Authentication* in the *Windows System Properties -> Remote* settings must be DISABLED for Remote Console to connect.

VNC (VMware Hypervisor Console)

When the `Hypervisor Console` option is enabled in cloud settings, the VNC protocol will be used for all platform types that Cloud.

When using VNC Hypervisor Console, the Morpheus Appliance connects directly to the host the VM is on, not directly to the VM.

Morpheus features Remote Console support directly to hypervisors. To enable this feature a few prerequisites must be met:

- The Morpheus Appliance must have network access to the host the VM is on over 443.
- The Morpheus Appliance must be able to resolve the hypervisor hostnames.

Note: VNC connections for VMs and Hosts in VMware type clouds are made directly to the ESXi hosts, not vCenter.

Unlike SSH and RDP, valid credentials do not need to be set on the VM or Host records in Morpheus for VNC hypervisor console connections. An IP address is also not required on the VM or Host for VNC hypervisor console connections. Morpheus will be able to connect to the VM or Host as soon as the `Host (Hypervisor)` record is set, which can be viewed in the Info section on the VM or Host detail page.

Note:

- Auto-login is not supported for Hypervisor Console. Auto-login role settings do not apply to console connecting when using Hypervisor Console. Please note Hypervisor Console sessions persist on the ESXi host and once a user manually logs in to the VM they will continue to be logged in, even if the console tab/window in Morpheus is closed, until they manually log out.
 - Copy and Paste and Text selection in Linux terminals is not supported when using VNC (VMware Hypervisor Console).
 - In Morpheus versions 3.2.0 and higher, a newer Guacamole version is installed that is not compatible with MacOS Platform Types over VNC.
-

Copy and Paste

Note: Copy and Paste for Text is supported for SSH and RDP protocols only.

To Copy text from the console:

1. Select text in the Console window.
2. Click the COPY button at the top of the Console window.
3. The selected text is copied to the users clipboard.

To Paste text into console:

1. Copy text on the local computer to you clipboard
2. Right click into the “Paste Text Here” field at the top of the Console window. The field will the display “Text Copied, Use Console to Paste.”
3. Right click into the console window.
4. The text is pasted into the VM.

Guacamole

Overview

Morpheus uses Apache Guacamole, a clientless remote console. Guacamole is installed on the Morpheus Appliance during the initial reconfigure. In Morpheus versions 3.2.0 and higher, Guacamole 0.9.14 is automatically installed. On Morpheus versions older than 3.2.0, 0.9.9 is installed. The 0.9.14 version is required for VNC Hypervisor Console functionality on ESXi v6.5 and later.

The Guacamole proxy daemon, `guacd`, is used for all Remote Console connections and must be running for Remote Console functionality.

Troubleshooting guacd

If all console connections are not functioning, the Guacamole proxy daemon (`guacd`) process may not be running or have a stuck process preventing console connections. This is evident when only the header appears in the console tab/window, and no console window appears below the header and no connection status is show in the console header. The following commands can be used on the Morpheus Appliance to restore console functionality.

morpheus-ctl status Lists all local Morpheus services including `guacd` and their states. If `guacd` is stopped, it will need to be started again for Remote Console to function.

morpheus-ctl start guacd Starts the `guacd` process

morpheus-ctl stop guacd Stops the `guacd` process

morpheus-ctl kill guacd Forcefully kills the `guacd` process

morpheus-ctl restarts guacd Restarts the `guacd` process

morpheus-ctl tail guacd Tails the `guacd` current and state logs, located by default at `/var/log/morpheus/guacd/`. This log is useful when troubleshooting console connections, guacamole service status, and to determine the protocol being used for the Remote Console connection.

If `guacd` continues to stop even after being started, or if `guacd` is running and no properly configured console connections are functioning, there may be a stuck `guacd` or multiple `guacd` processes running, which will need to killed and `guacd` started again.

To kill all `guacd` processes on the Morpheus Appliance and start `guacd` again:

1. Kill the morpheus `gaucd` proccess: `morpheus-ctl kill guacd`
2. Grep for all running `guacd` processes: `sudo ps -aux | grep guacd` and note the `guacd` pid(s) (minus the process from the grep)
3. Kill all running `guacd` processes: `kill -9 pid` replacing *pid* with the pid(s) of the target processes
4. Start `guacd` again: `morpheus-ctl start guacd`

5. Tail the guacd logs to verify guacd is started and listening: `morpheus-ctl tail guacd` The log output will resemble below when guacd is properly running:

```
guacd[16899]: INFO:      Guacamole proxy daemon (guacd) version 0.9.14 started
guacd[16899]: INFO:      Listening on host 127.0.0.1, port 4822
```

6. Additional information in the guacd logs appears when Morpheus is making a console connection. A successful connection will resemble:

```
guacd[24725]: INFO: Creating new client for protocol "ssh"
guacd[24725]: INFO: Connection ID is "$24f67856-f050-4a17-83eb-9101g0cd8869"
guacd[24743]: INFO: Current locale does not use UTF-8. Some characters may not
↳render correctly.
guacd[24743]: INFO: User "@63102f19-eff4-412e-b1f9-718405f55782" joined
↳connection "$24f67856-f050-4a17-83eb-9101g0cd8869" (1 users now present)
guacd[24743]: INFO: Auth key successfully imported.
guacd[24743]: INFO: SSH connection successful.
```

Guacamole Version

In Morpheus versions 3.2.0 and higher, Guacamole version 0.9.14 is automatically installed. On Morpheus versions older than 3.2.0, 0.9.9 is installed. The 0.9.14 version is required for VNC Hypervisor Console functionality on ESXi v6.5 and later.

Note Guacamole version 0.9.14 is not compatible with MacOS Platform Types over VNC on ESXi v6.0 or prior (6.5 is supported). If necessary, the guacamole version can be reverted to 0.9.9.

To revert the guacamole version from 0.9.14 to 0.9.9.

1. Kill guacd: `morpheus-ctl kill guacd`
2. Check if any guacd processes are still running: `ps -aux | grep guacd`
3. If so, kill the processes: `kill -9 pid` with id being the actual process id, like 16101.
4. Go to the guacd 0.9.9 directory: `cd /var/opt/morpheus/guacamole-server-0.9.9`
5. Run: `make install`
6. Start guacd: `morpheus-ctl start guacd`

Apps

Apps allow instances having general relationships to be grouped in a clean and organized manner. App functionality enables full control of which instances belong in an app as well setting Firewall and Access Control List (ACL) rules. Use Apps to structure all necessary components into a single place. Add checks and groups for web servers, database nodes, etc.

Apps can be created from Blueprints, which are made in Provisioning -> Blueprints or from Existing Apps.

Creating Apps from Blueprints

1. Click **+ADD** on the right side of the main Apps section in Provisioning.
2. Select an existing App Blueprint and click **NEXT**.

Note: Blueprints must be created in in Provisioning -> Blueprints. to appear as options when creating an App.

3. Enter a Name for the App and select a Group. Default Cloud and Env can also be selected.
4. Click **NEXT**. Blueprint configurations matching the Group, Cloud and Environment selections will auto-populate the configurations of the Instances in the App. If no Blueprint Configuration matched the Group, Cloud or Env selections, the Instances will have default configurations.
5. Configure your Instances. Depending on the Blueprint Configurations settings, instances may already be fully configured. Fields that are locked in a Blueprint cannot be edited when creating an App.

Note: Once an Instance is fully configured, a green checkmark will appear next to the Instance. Instances that have required fields that need populated will have a red X and must be completed. If your Blueprint is already fully configured you can simply select complete!

6. Select **COMPLETE** and the App will be created and the Instances will begin provisioning.

NEW APP

TEMPLATE

SETUP

CONFIGURE

Builder

Raw

Preview

STRUCTURE

Spud Marketing

Web

Tomcat

Database

MySQL

CONFIGURATION

Instance Info

NAME

DESCRIPTION

Configuration Options

CLOUD

VERSION

INSTANCE CONFIGURATION

PLAN

VOLUMES

HOST

Advanced Options

Automation

Deployment

Load Balancer

Backups

Lifecycle

PREV

COMPLETE

Creating Apps from Existing Instances

1. Click **+ADD** on the right side of the main Apps section in Provisioning.
2. Select **APP FROM EXISTING INSTANCES** from the Blueprints list and click **NEXT**.
3. Enter a Name for the App and select a Group. Default Cloud and Env can also be selected.

Note: Only instances within the selected Group and Cloud will be available to be added to the App.

4. In the **STRUCTURE** section, select **+** to add a Tier

5. Select or enter a Tier Name.
6. Select the Tier to set Boot Order, rename, or once multiple Tiers are added, connect the Tier to other Tiers.
7. In the STRUCTURE section, select + in a Tier to add an Instance
8. Select the Instance Type of the Existing Instance to be added to the App.
9. In the STRUCTURE section, select the Instance.
10. In the CONFIGURATION section, select the Cloud the Existing Instance is in. Existing INSTANCES that match the Group, Cloud and Instance Types set will populate.
11. Select the desired Instance from the INSTANCES list. Selected instance will show in the SELECTED INSTANCE section.

Note: Only one existing Instance can be added per Instance. To add multiple Existing Instances, repeat the step above including adding an Instance for each Existing Instance to be added to the App.

12. Once all Existing Instances have been selected, click *COMPLETE*.
13. A new App will be created out of the Existing Instances.

NEW APP

TEMPLATE

SETUP

CONFIGURE

BuilderRawPreview

STRUCTURE

Morpheus HA01

App

CentOS

CentOS

Database

PERCONA

Messaging

RabbitMQ

ES

elastic

CONFIGURATION

USE EXISTING

INSTANCE

CLOUD

SEARCH

SELECTED INSTANCE

INSTANCES

Nutanix

search for instances

MORPHEUS-HA-AP1

SSH: 10.30.20.181:22 SSH: 10.30.20.62:22

VERSION: 7.3

CONTAINERS:2

STATUS:

JW-CENTOS-AHV

SSH: 10.30.20.165:22 SSH: 10.30.20.91:22 SSH: 10.30.20.149:22

VERSION: 7.3

CONTAINERS:3

STATUS:

JW-CENTOS-AHV-1

SSH: 10.30.20.135:22 SSH: 10.30.20.37:22 SSH: 10.30.20.52:22

VERSION: 7.3

CONTAINERS:3

STATUS:

JW-CENTOS-AHV-2

SSH: 10.30.20.104:22

VERSION: 7.3

CONTAINERS:1

STATUS:

MORPHEUS-HA-DB

SSH: 10.30.20.153:22 SSH: 10.30.20.139:22 SSH: 10.30.20.196:22

VERSION: 7.3

CONTAINERS:3

STATUS:

MORPHEUS-HA-MQ

SSH: 10.30.20.129:22 SSH: 10.30.20.58:22

VERSION: 7.3

CONTAINERS:2

STATUS:

MORPHEUS-HA-MQ-LB-1

SSH: 10.30.20.53:22

VERSION: 7.3

CONTAINERS:1

STATUS:

PREV

COMPLETE

Exporting Configuration JSON

To export a Blueprint as JSON:

1. Navigate to `Provisioning > Apps`
2. Select an App from the list to view the App detail page
3. Click the Actions button and select Export
4. The App export file will be downloaded to your computer as `{app_name}.json`

Provisioning Apps via API

A quick example of how this work: <https://d.pr/i/yxsW7t>

Blueprints

Overview

App Blueprints support a vast array of providers and configurations with programmatic markup or Infrastructure as Code capabilities. Blueprints configs can be manually added or scoped to a git repo. Morpheus blueprints allows for full automation configuration, locked fields, tiered boots, and linked tiers with exported evvars. All blueprints have permission settings for controlling group and tenant access.

Blueprint Types

- Morpheus
- Terraform
- ARM (Azure)
- Cloud Formation (AWS)
- Kubernetes
- Helm

Morpheus Blueprints

Morpheus App Blueprints allow pre-configured full multi-tier application deployments for multiple environments. Blueprints can be provisioned from the `Provisioning -> Apps` section and can be fully configured for one click provisioning. Blueprints can be built within the `Builder` section or by code in the `Raw` section. Blueprints can also be exported as YAML or JSON and created with the Morpheus API and CLI.

A unique capability of the YAML/JSON based Morpheus blueprint structure is the ability to have multiple configurations per instance being provisioned within the app blueprint. This can be a scoped configuration that acts as overrides based on selected cloud, group, and/or environment the app is being provisioned in as a target. For example, maybe the “development” environment doesn’t need as many horizontally scaled nodes as the “production” environment. Another great aspect of this configuration markup is a blueprint can be defined as a hybrid cloud blueprint. This makes the app blueprint structure very powerful and in some ways better than alternative infrastructure as code orchestrators. For Example, ARM is locked into Azure, while Cloud Formation is locked into AWS. Even Terraform does not allow a tf file to expand its bounds beyond a specific provider type.

Basic Blueprint Structure

In a Morpheus App Blueprint there are a few structural concepts to be aware of. Firstly there is a concept of a *Tier*. A *Tier* is a grouping of instances within an app blueprint. Tiers can be used for a variety of things including sequenced booting of instances or even properly creating endpoint groups and security group contexts in network security tools like Cisco ACI. An example of a Tier structure might be a *Web* tier and a *Database* tier. These tiers can also be marked as connected such that network communication rules can appropriately be defined. A basic 2 Tier blueprint skeleton might look something like this:

```
name: Tier Example
type: morpheus
tiers:
  Web:
    linkedTiers:
      - Database
    tier:
      bootOrder: 1
    instances:
  Database:
    tier:
      bootOrder: 0
    instances:
```

This example has defined 2 tiers as yaml properties under the *tiers* object. They are called *Web* and *Database*. A Tier can optionally define its connected tiers which are bi-directional even though only one tier has to define them. This is the *linkedTiers* array and simply lists the connected tiers by tier name. A Boot Order can also optionally be defined under a nested {"tier": {"bootOrder": 1}} object structure.

Configuration Scopes

Another capability of Morpheus App Blueprint structure is its configuration scoping. This allows properties to be overridden based on the apps target environment or even target group and cloud. For example. Maybe we want to use a larger plan size in production vs. development

An example of that can be done using “environments” overrides.

```
name: Simple Nginx
type: morpheus
tiers:
  Web:
    instances:
      - instance:
          type: docker
          name: Sample Nginx
    clouds:
      AWS Cali:
        instance:
          layout:
            code: docker-1.7-single
          config:
            dockerImageVersion: latest
            dockerRegistryId: ''
            dockerImage: nginx
          plan:
            code: container-128
```

(continues on next page)

(continued from previous page)

```

environments:
  Production:
    groups:
      All Clouds Demo:
        clouds:
          AWS Cali:
            plan:
              code: container-256

```

Note the new environments object. The object graph of the morpheus blueprint structure gets merged and flattened at provision time based on the scope of the configurations provided as well as the users target cloud, group, and environment selection. In the Above example, a selective override was done for the *AWS Cali* cloud when using a Production Environment and deploying to the group *All Clouds Demo*. This specific example changes the plan to a larger size. Scoped configurations have various levels of precedence. Cloud is the lowest level of precedence. a cloud configuration in a group is the next level higher and finally an environment configuration in a group in a cloud is the highest level of scoped precedence.

Getting Started

To get started, it may be best to look at a simple App Blueprint configuration. Docker templates are less complex than virtual machine based templates so lets look at a Blueprint that deploys a single Nginx container to a target cloud:

```

name: Simple Nginx
type: morpheus
tiers:
  Web:
    linkedTiers: []
    instances:
      - instance:
          type: docker
          name: Sample Nginx
        clouds:
          AWS Cali:
            instance:
              layout:
                code: docker-1.7-single
                id: 206
            volumes:
              - rootVolume: true
                name: root
                size: 1
            backup:
              createBackup: false
            config:
              dockerImageVersion: latest
              dockerRegistryId: ''
              dockerImage: nginx
            plan:
              id: 68
              code: container-128
            ports:
              - name: HTTP
                port: 80
                lb: HTTP

```

There are some useful things to look at in the above docker example. One is there are different objects based on the different available configuration options for the target provision type. These options are actually data driven and can be extracted from the option types api in the morpheus api doc. That is a useful resource to look at while building morpheus blueprints or by using the *morpheus-cli* which provides prompts for helping build custom morpheus app blueprints.

EDIT APP TEMPLATE

BuilderRawPreview

STRUCTURE

Spud Marketing

Web

Tomcat

Group: All Clouds Demo, Cloud: Appliance KVM & Docker

Group: All Clouds Demo, Cloud: Labs UCS

Database

MySQL

Group: All Clouds Demo, Cloud: Appliance KVM & Docker

Group: All Clouds Demo, Cloud: Labs UCS

CONFIGURATION

Template Summary

NAME

DESCRIPTION

CATEGORY

IMAGE

BROWSE

SAVE

Creating App Blueprints

1. Navigate to Provisioning -> Blueprints
2. Select + ADD
3. Enter a NAME for the Blueprint and select NEXT
4. Optionally add a Description, Category, and Image for the Blueprint.

Add Tiers

1. In the STRUCTURE section, select + to add a Tier
2. Select or enter a Tier Name.
3. Select the Tier to set Boot Order, rename, or once multiple Tiers are added, connect the Tier to other Tiers.

Add Instances to Tiers

1. In the STRUCTURE section, select + in a Tier to add an Instance
2. Select an Instance Type
3. Optionally add a name for the Instance. Instances with blank names will automatically be named based off the App name.

Tip: You can use the variable `${app.name}` in your instance naming convention to reference the name of the application you're deploying.

Add Configurations to Instances

1. In the STRUCTURE section, select + in an Instance to add a Configuration
2. Select at least one option from Group, Cloud or Environment.
3. Select `ADD CONFIG` to create the configuration
4. Populate the Configuration
 - Configurations can be fully partially or populated
 - Fields can be locked or hidden by selecting the Lock icon next to the Field. Locking prevents the field from being editable when provisioning an App using the Blueprint while hidden fields are not revealed to the user at all
 - `ALLOW EXISTING INSTANCE` will allow users to add existing Instances to the App when using the blueprint

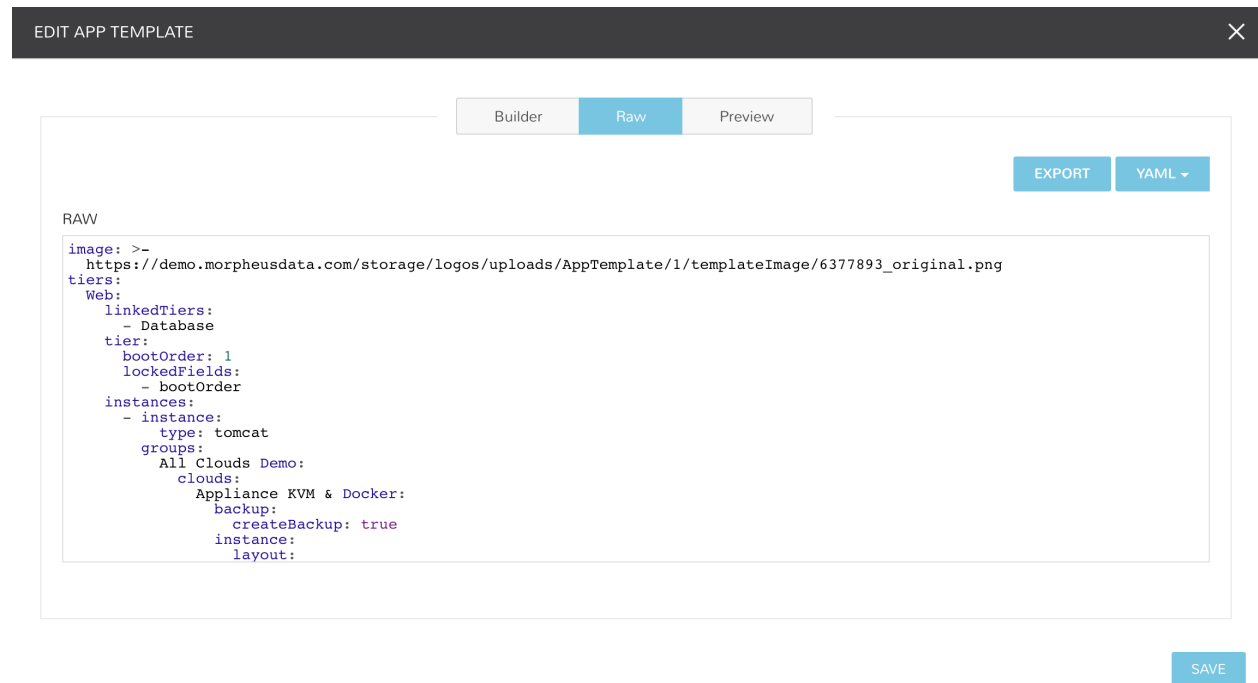
Save

Once all desired Tiers, Instances and Configurations are added, select Save. The Blueprint will be created, can be edited after saving, and will available in the Apps section for provisioning.

Note: Blueprints are not provisioned when created. To provision a Blueprint, use `Provisioning -> Apps`.

RAW

Blueprints can be create, edited or Exported in the RAW section when creating or editing a blueprint.

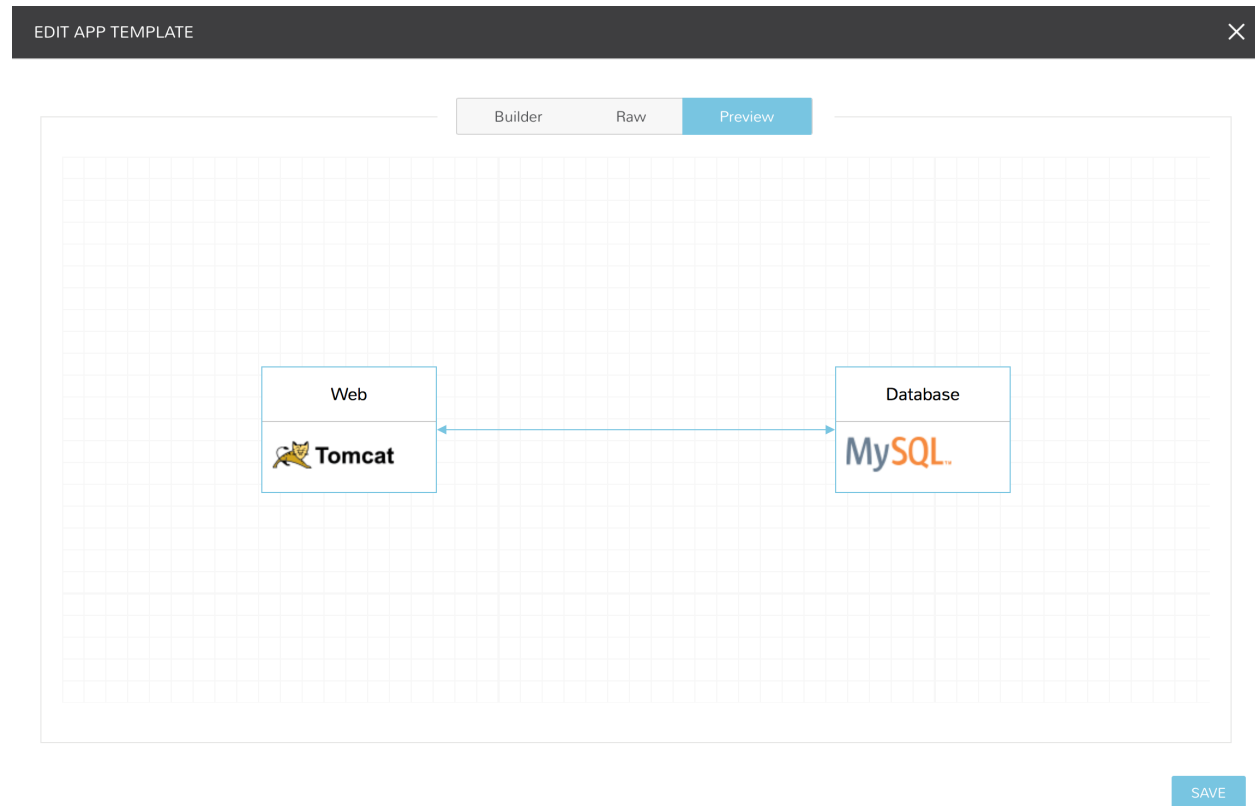


To Export a Blueprint as JSON or YAML:

1. Navigate to Provisioning > Blueprints
2. Edit an existing App by clicking on the pencil icon
3. On the Edit Blueprint modal, select the Raw tab
4. Select YAML or JSON from the dropdown in the top right
5. Click the Export button
6. Select the configurations to include in the export by selecting or deselecting configurations as needed. Selected configurations will be highlighted
7. Click the DOWNLOAD CONFIGURATION button
8. The Blueprint export file will be downloaded to your computer as {app_name}-config.json or {app_name}-config.yaml

Preview

In the APP BLUEPRINT modal, select the Preview section to display a graphical representation of your Blueprint Tiers, Instances and Tier Connections.



Important: When Tiers are connected, the Instances in a Tier will import the envvars from Instances in connected Tiers, and if [morpheus] is managing the Instance Firewalls, communication between the Instances will be facilitated based on the Instances port configurations.

Provisioning

To provision a Blueprint, navigate to Provisioning -> Apps and select the Blueprint when creating an App. See the [App section](#) of Morpheus docs for more on provisioning Apps.

Terraform Blueprints

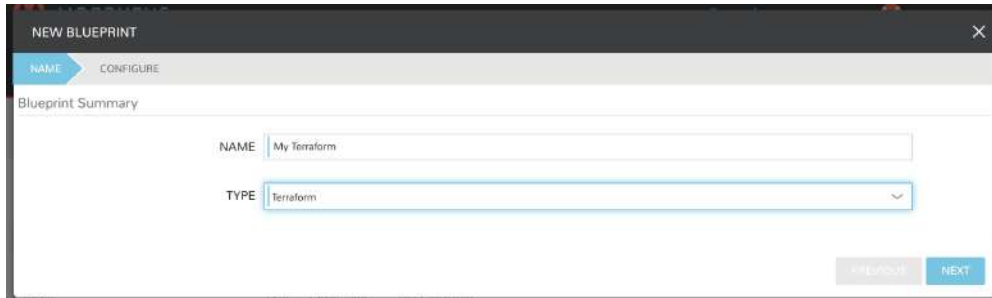
Terraform Blueprints are one way that Terraform can be integrated and leveraged with Morpheus, with the other being the Morpheus Terraform provider which is not discussed in this section. Morpheus and Terraform are complimentary technologies which together can increase efficiency and simplify automation across cloud environments. For more on this relationship, see our [whitepaper](#) on how Morpheus and Terraform are better together.

Currently, Morpheus supports provisioning Apps based on Terraform Blueprints to VMware, Amazon, Azure, and Oracle Clouds with additional Cloud support coming in future releases. On first attempt to provision a Terraform App, Morpheus will automatically install Terraform. It is possible in some operating system configurations for this

automated installation process to fail, requiring you to install Terraform manually. If needed, manual installation instructions and guidance are provided [here](#).

To create a new Terraform Blueprint, navigate to Provisioning > Blueprints. Click + *ADD*.

On the Name tab of the New Blueprint modal, enter a name for your new Blueprint. In the Type dropdown menu, select Terraform. *NEXT*

The image shows a web interface for creating a new blueprint. The modal is titled 'NEW BLUEPRINT' and has two tabs: 'NAME' (active) and 'CONFIGURE'. Under the 'NAME' tab, there is a 'Blueprint Summary' section. It contains two input fields: 'NAME' with the value 'My Terraform' and 'TYPE' with a dropdown menu set to 'Terraform'. At the bottom right, there are two buttons: 'PREVIOUS' (disabled) and 'NEXT' (active).

In the Blueprint Summary section, complete the following fields as needed:

- **NAME:** Enter a name for this Blueprint as it will appear in the Morpheus Blueprints list
- **DESCRIPTION:** An optional description field for your Blueprint
- **CATEGORY:** An optional category tag for your Blueprint, such as web, utility, or app
- **IMAGE:** An optional image icon to more easily identify your Blueprint from a list. If no image is uploaded, a default image will be used

The Terraform Configuration section is where the Terraform template file (.tf) is added or linked to the Blueprint. Using a Config Type of “Terraform (.tf)” or “Terraform JSON (.tf.json)”, you can write or paste your configuration directly into the new Blueprint. Alternatively, you can pull in config files from an integrated Git repository using the “Git Repository” Config type.

In the Terraform Configuration section, complete the following fields as needed when entering your configuration directly into the new Blueprint:

- **CONFIG TYPE:** “Terraform (.tf)” or “Terraform JSON (.tf.json)” to create or paste configuration directly into the new Blueprint
- **CONFIG:** Enter your configuration here
- **TFVAR SECRET:** Select an existing TFVar-formatted Cypher. See the Cyphers section or Morpheus docs for more information on Cyphers
- **OPTIONS:** Enter any additional options, such as a variable definition

In the Terraform Configuration section, complete the following fields as needed when syncing in configuration from a Git repository:

- **CONFIG TYPE:** “Git Repository”
- **SCM INTEGRATION:** If a pre-existing SCM integration is selected here, the available selections in the “Repository” dropdown menu will be filtered to show only those associated with the chosen SCM integration
- **REPOSITORY:** Select the repository in which your configuration resides
- **BRANCH OR TAG:** The branch in which your configuration resides
- **WORKING PATH:** The path to your configuration files
- **CONFIG:** Your selected config file

- **TFVAR SECRET:** Select an existing TFVar-formatted Cypher. See the Cyphers section of Morpheus docs for more information on Cyphers
- **OPTIONS:** Enter any additional options, such as a variable definition

Once finished, click *COMPLETE*.

Your new Terraform Blueprint is now saved and should be visible in the list of Blueprints. Blueprints are deployed in the Provisioning > Apps section of Morpheus. See the Apps section of Morpheus docs for more information on that process.

ARM Blueprints

ARM Blueprints provide a simple and repeatable way of deploying infrastructure-as-code to Azure Clouds. Objects and properties are defined in a JSON file and are provisionable on-demand in *Provisioning > Apps*

To create a new ARM Blueprint, navigate to Provisioning > Blueprints. Click + *ADD*.

On the Name tab of the New Blueprint modal, enter a name for your new Blueprint. In the Type dropdown menu, select ARM. *NEXT*

In the Blueprint Summary section, complete the following fields as needed:

- **NAME:** Enter a name for this Blueprint as it will appear in the Morpheus Blueprints list
- **DESCRIPTION:** An optional description field for your Blueprint
- **CATEGORY:** An optional category tag for your Blueprint, such as web, utility, or app
- **IMAGE:** An optional image icon to more easily identify your Blueprint from a list. If no image is uploaded, a default image will be used

The ARM template itself is defined in the ARM Configuration section. Using the Config Type dropdown menu, we can opt to write or paste JSON configuration directly into this modal, or we can choose to bring in a JSON which we're keeping under version control in a Git repository.

Depending on whether we need the Morpheus Agent installed and/or cloud-init enabled, mark the following boxes in the next section:

- **INSTALL AGENT**
- **CLOUD INIT ENABLED**

If writing or pasting your configuration JSON directly into the modal, fill out the following fields:

- **OS TYPE:** Identify the resources to be created as Linux or Windows
- **CONFIG TYPE:** ARM Template JSON (.json)
- **CONFIG:** Your JSON configuration template

If bringing in a template from a Git repository, fill out the following fields:

- **OS TYPE:** Identify the resources to be created as Linux or Windows
- **CONFIG TYPE:** "Git Repository"
- **SCM INTEGRATION:** If a pre-existing SCM integration is selected here, the available selections in the "Repository" dropdown menu will be filtered to show only those associated with the chosen SCM integration
- **REPOSITORY:** Select the repository in which your configuration resides
- **BRANCH OR TAG:** The branch in which your configuration resides
- **WORKING PATH:** The path to your configuration files

- **CONFIG:** Your selected config file

Once finished, click *COMPLETE*.

Your new ARM Blueprint is now saved and should be visible in the list of Blueprints. Blueprints are deployed in the Provisioning > Apps section of Morpheus. See the Apps section of Morpheus docs for more information on that process.

Cloud Formation Blueprints

CloudFormation Blueprints consume new or existing CloudFormation templates to create easily-deployable application stacks. CloudFormation templates in Morpheus are JSON-formatted text documents that declare all relevant AWS resources needed for the provisioned application. They can be created directly in the New Blueprint modal or pulled in from existing Git repositories.

If needed, Amazon has educational resources available for getting started with CloudFormation. They can be found in the [AWS CloudFormation documentation](#).

To create a new CloudFormation Blueprint, navigate to Provisioning > Blueprints. Click + *ADD*.

On the Name tab of the New Blueprint modal, enter a name for your new Blueprint. In the Type dropdown menu, select CloudFormation. Click *NEXT*

In the Blueprint Summary section, complete the following fields as needed:

- **NAME:** Enter a name for this Blueprint as it will appear in the Morpheus Blueprints list
- **DESCRIPTION:** An optional description field for your Blueprint
- **CATEGORY:** An optional category tag for your Blueprint, such as web, utility, or app
- **IMAGE:** An optional image icon to more easily identify your Blueprint from a list. If no image is uploaded, a default image will be used

Depending on whether we need the Morpheus Agent installed and/or cloud-init enabled, mark the following boxes in the next section:

- **INSTALL AGENT**
- **CLOUD INIT ENABLED**

In some cases, you must explicitly acknowledge that your template contains certain capabilities in order for the application to successfully be deployed. There is more information on this in Amazon's documentation [here](#). If any of the following capabilities are contained in your application, acknowledge them by marking any of the following boxes that apply:

- **CAPABILITY_IAM**
- **CAPABILITY_NAMED_IAM**
- **CAPABILITY_AUTO_EXPAND**

Continuing on with the CloudFormation Configuration section, complete the following fields as needed when entering your configuration directly into the new Blueprint:

- **CONFIG TYPE:** "CloudFormation Template JSON (.json)"
- **CONFIG:** Enter your configuration here

In the CloudFormation Configuration section, complete the following fields as needed when syncing in configuration from a Git repository:

- **CONFIG TYPE:** "Git Repository"

- **SCM INTEGRATION:** If a pre-existing SCM integration is selected here, the available selections in the “Repository” dropdown menu will be filtered to show only those associated with the chosen SCM integration
- **REPOSITORY:** Select the repository in which your configuration resides
- **BRANCH OR TAG:** The branch in which your configuration resides
- **WORKING PATH:** The path to your configuration files
- **CONFIG:** Your selected config file

Once finished, click *COMPLETE*.

Your new CloudFormation Blueprint is now saved and should be visible in the list of Blueprints. Blueprints are deployed in the Provisioning > Apps section of Morpheus. See the Apps section of Morpheus docs for more information on that process.

Kubernetes Blueprints

Morpheus allows you to store Kubernetes configuration YAML files for easy deployment on-demand. Kubernetes Blueprints can be built by pulling in Kubernetes spec stored as a Morpheus Spec Template object, those tracked under version control in a Git repository, or you can write them directly in the New Blueprint modal.

To create a new Kubernetes Blueprint, navigate to Provisioning > Blueprints. Click + *ADD*.

On the Name tab of the New Blueprint modal, enter a name for your new Blueprint. In the Type dropdown menu, select Kubernetes. *NEXT*

In the Cluster Summary section, complete the following fields as needed:

- **NAME:** Enter a name for this Blueprint as it will appear in the Morpheus Blueprints list
- **DESCRIPTION:** An optional description field for your Blueprint
- **CATEGORY:** An optional category tag for your Blueprint, such as web, utility, or app
- **IMAGE:** An optional image icon to more easily identify your Blueprint from a list. If no image is uploaded, a default image will be used

Complete the Kubernetes Configuration section as follows depending on your Config Type selection.

To consume a Morpheus Spec Template containing Kubernetes spec:

- **CONFIG TYPE:** “Kubernetes Spec”
- **SPEC TEMPLATE:** Use the typeahead field to locate the desired Spec Template

To draft or paste configuration directly in the New Blueprint modal:

- **CONFIG TYPE:** “Kubernetes Yaml Spec”
- **CONFIG:** Enter your YAML configuration template here

To consume YAML configuration files tracked in a Git repository:

- **CONFIG TYPE:** “Git Repository”
- **SCM INTEGRATION:** If a pre-existing SCM integration is selected here, the available selections in the “Repository” dropdown menu will be filtered to show only those associated with the chosen SCM integration
- **REPOSITORY:** Select the repository in which your configuration resides
- **BRANCH OR TAG:** The branch in which your configuration resides
- **WORKING PATH:** The path to your configuration files

- **CONFIG:** Your selected config file

Once finished, click *COMPLETE*.

Your new Kubernetes Blueprint is now saved and should be visible in the list of Blueprints. Blueprints are deployed in the Provisioning > Apps section of Morpheus. See the Apps section of Morpheus docs for more information on that process.

Helm Blueprints

If you're using Helm Charts to manage Kubernetes applications, Morpheus allows you to bring them in from a Git repository as a Blueprint. The selected repository must be integrated with Morpheus before creating the Blueprint.

To create a new Helm Blueprint, navigate to Provisioning > Blueprints. Click + *ADD*.

On the Name tab of the New Blueprint modal, enter a name for your new Blueprint. In the Type dropdown menu, select Helm. Click *guilabel:NEXT*.

In the Blueprint Summary section, complete the following fields as needed:

- **NAME:** Enter a name for this Blueprint as it will appear in the Morpheus Blueprints list
- **DESCRIPTION:** An optional description field for your Blueprint
- **CATEGORY:** An optional category tag for your Blueprint, such as web, utility, or app
- **IMAGE:** An optional image icon to more easily identify your Blueprint from a list. If no image is uploaded, a default image will be used

In the Helm Configuration section, complete the following fields as needed to sync in configuration from a Git repository:

- **CONFIG TYPE:** "Git Repository"
- **SCM INTEGRATION:** If a pre-existing SCM integration is selected here, the available selections in the "Repository" dropdown menu will be filtered to show only those associated with the chosen SCM integration
- **REPOSITORY:** Select the repository in which your configuration resides
- **BRANCH OR TAG:** The branch in which your configuration resides
- **CHART PATH:** The path to the folder within the repository containing your configuration files, enter *"/* if this is the top level folder within the repository
- **CONFIG:** Config files within your selected folder are displayed here for confirmation

Once finished, click *COMPLETE*.

Your new Helm Blueprint is now saved and should be visible in the list of Blueprints. Blueprints are deployed in the Provisioning > Apps section of Morpheus. See the Apps section of Morpheus docs for more information on that process.

Jobs

Jobs are for scheduled execution of Automation Tasks and Workflows. Jobs can be set to execute on a schedule, at one specific point in time, and/or execute manually (on-demand). Jobs are linked to existing Tasks or Workflows, and allow for custom configuration options. Jobs can be associated with Instances, Servers, or have no association, such as a job for an SSH task.

Jobs allow for scheduled execution of nearly anything as Task Types include Bash, Powershell, HTTP/API, Ansible, Chef, Puppet, Groovy, Python, jRuby, Javascript, and library scripts and templates, which can be configured for resource, remote, or local execution targets. If you need something to execute on a schedule, Morpheus Jobs can deliver.

Jobs are configured in the `JOBS` tab, and the `JOB EXECUTIONS` tab contains Job execution history with result output.

Jobs

Role Permissions

Provisioning: Jobs

- **None:** Cannot access `Provisioning > Jobs > Jobs` tab
- **Read:** Can access `Provisioning > Jobs > Jobs` tab but cannot create, edit, or delete Jobs
- **Full:** Full permissions to create, view, edit, and delete Jobs

Provisioning: Job Executions

- **None:** Cannot access `Provisioning > Jobs > Job Executions` tab
- **Read:** Can access and view `Provisioning > Jobs > Job Executions` tab including job execution history, status, and Job output

Creating Jobs

Note: Jobs require existing Tasks or Workflows. See the appropriate section of Morpheus docs for more on creating [Tasks](#) and [Workflows](#).

To create a new job:

1. Navigate to `Provisioning > Jobs`
2. Select `+ ADD`
3. Enter the following

NAME: Name of the Job in Morpheus **JOB TYPE:**

Task: Job will execute a selected Task Workflow: Job will execute a selected Workflow

ENABLED: When checked, the Job will run as scheduled

4. Select `NEXT`
5. Configure the Job

Task Jobs

TASK: Select target Task. If relevant to the Task, Option Type fields will be presented

SCHEDULE:

Manual: Job is not scheduled but can be executed from Provisioning > Jobs and selecting Actions > Execute

Date And Time: Job will be executed at one specific point in time and not again (unless rescheduled or executed manually)

Schedule: Select a configured Execution Schedule. Execution Schedules are created in Provisioning > Automation > EXECUTE SCHEDULING

Note: Morpheus provides two default execution schedules, Daily at Midnight and Weekly on Sunday at Midnight. Any additional schedules were created by a User. Additional schedules can be added in Provisioning > Automation > EXECUTE SCHEDULING

CONTEXT TYPE: Server or Instance

CONTEXT SERVER/INSTANCE: Select the Server or Instance you wish to target with the Job

RUN NOW: When checked, the Job will execute on save regardless of SCHEDULE setting.

Workflow Jobs WORKFLOW: Select target Workflow. If relevant to the Workflow, Option Type fields will be presented

SCHEDULE: Manual: Job is not scheduled but can be executed from Provisioning > Jobs and selecting Actions > Execute

Date And Time: Job will be executed at one specific point in time and not again (unless rescheduled or executed manually)

Schedule: Select a configured Execution Schedule. Execution Schedules are created in Provisioning > Automation > EXECUTE SCHEDULING

Note: Morpheus provides two default execution schedules, Daily at Midnight and Weekly on Sunday at Midnight. Any additional schedules were created by a User. Additional schedules can be added in Provisioning > Automation > EXECUTE SCHEDULING

CONTEXT TYPE: Server or Instance

CONTEXT SERVER/INSTANCE: Select the Server or Instance you wish to target with the Job

RUN NOW: When checked, the Job will execute on save regardless of SCHEDULE setting.

6. Select *NEXT*

7. Select *COMPLETE*

Creating and Running Security Scan Jobs

Security Scan Jobs allow users to create and schedule SCAP program (Security Content Automation Program) scans for groups of managed systems. These Jobs can call in existing SCAP packages and checklists, which are used to scan the targeted systems on-demand or on a scheduled basis. Historical data for these scans is saved in the Job Execution list and in the software section of server detail pages. Detailed scan reports can also be viewed for each system as needed once the scan is complete. See the [SCAP documentation](#) on the NIST website for information on developing your own scanning procedures.

Note: Creating and editing Security Scan Jobs requires the “Security: Scanning” Role permission set to Full. Viewing Security Scan Jobs and seeing the results for scanned servers requires at least a Read-level permission.

Add a new Security Package

1. Navigate to Provisioning > Jobs > Security Packages Tab
2. Click **+ADD > SCAP Package**
3. Provide a name in addition to a URL to source the package
4. Click **SAVE CHANGES**

Note: Currently URL is the only source option for security packages

ADD SECURITY PACKAGE

NAME

DESCRIPTION

☒ ENABLED

Security Package Files

SOURCE

URL

SAVE CHANGES

Add a new Security Scan Job

1. Navigate to Provisioning > Jobs > Jobs Tab
2. Click **+ADD**
3. Set the Job type to “Security Scan Job” and provide a friendly name for the Job
4. Click **NEXT**

NEW JOB

SETUP CONFIGURE REVIEW

Job Summary

NAME Scan Job

JOB TYPE Security Scan Job

☒ ENABLED

PREVIOUS NEXT

5. Select a security package, see the previous section to add a new one
6. Enter your Scan Checklist (XML document) and Security Profile (XCCDF document), more information on these can be found in the SCAP documentation linked above
7. Set a schedule or leave as Manual to only run this scan on-demand (new execution schedules can be created in Provisioning > Automation if needed)
8. Set the context, can be Instance or Server. Select as many Instances or Servers as needed for this scanning run
9. Click *NEXT*
10. After final review, click *COMPLETE*

EDIT JOB

SETUP CONFIGURE REVIEW

Job Configuration

SECURITY PACKAGE scap security bundle 0.1.51

SCAN CHECKLIST /scap-security-guide-0.1.51/ssg-ubuntu1804-ds.xml

SECURITY PROFILE xccdf_org.ssgproject.content_profile_cis

Execution Config

SCHEDULE Daily

☐ RUN NOW

CONTEXT TYPE Server

CONTEXT SERVER Search

de-ubuntu-1

PREVIOUS NEXT

Running Security Scan Jobs

Once created, Security Scan Jobs will run based on the configured schedule. They can also be run on-demand when needed:

1. Navigate to Provisioning > Jobs > Jobs Tab
2. Click *MORE*
3. Click “Execute”

Viewing Completed Security Scan Jobs

To view a list of completed Security Scan Jobs (and Jobs of other types):

1. Navigate to Provisioning > Jobs > Job Executions Tab
2. Additional details can be viewed by clicking (*i*)

To view scan results for specific servers:

1. Navigate to the server detail page (Infrastructure > Hosts > Virtual Machines tab > Selected server)
2. Click on the Software tab part way down the page, then click on the Security subtab
3. High level details on previous scans is viewable here

The screenshot shows the Morpheus Security interface. At the top, there is a navigation bar with tabs: Summary, Wiki, Storage, Network, Logs, **Software**, History, and Console. Below this, there are two sub-tabs: Software and Security. The Security tab is active, showing a search bar and a settings icon. Below the search bar is a table with the following columns: STATUS, NAME, TYPE, SCAN DATE, SCORE, and RESULTS. The table contains two rows of scan results for 'scap security bundle 0.1.51'.

STATUS	NAME	TYPE	SCAN DATE	SCORE	RESULTS
✓	scap security bundle 0.1.51	SCAP Package	10/15/2020 06:54 PM	28.12	29 passed, 39 failed, 3 other
✓	scap security bundle 0.1.51	SCAP Package	10/15/2020 06:48 PM	28.12	29 passed, 39 failed, 3 other

4. To view the full report, click (i)

Compliance and Scoring

The target system did not satisfy the conditions of 39 rules! Please review rule results and consider applying remediation.

Rule results



Severity of failed rules




Score

Scoring system	Score	Maximum	Percent
urn:xccdf:scoring:default	28.125000	100.000000	28.13%

Job Executions

The Job Executions tab contains execution history of completed Jobs, including any process outputs and error messages. Information included in the Job Executions list include:

- **JOB:** The name of the executed Job
- **DESCRIPTION:** When the Job Execution is expanded, the name of each executed task in the Job is listed in this column
- **TYPE:** The Job type, either Task or Workflow. When a Workflow Job is expanded, each individual Task making up the Workflow is identified as a Task in this column
- **START DATE:** The date and time the Job Execution kicked off. When expanded, the start date and time of each individual Task are also shown
- **ETA/TIME:** The time taken for the Job to complete. When expanded, the time to complete each individual Task is also shown
- **ERROR:** Any errors surfaced are shown here. When expanded, any surfaced errors for individual Tasks are also shown

Click the  icon at the end of the row for a Job Execution or individual Task (when a Job Execution is expanded) to view the Execution Detail modal which provides the following information:

- **Name of the Job or individual Task**
- **Description**
- **Start Date**

- **Created By**
- **Duration**
- **Status:** Completed, Running, or Failed
- **PROCESS OUTPUT:** Returned values and outputs from the completed Job
- **ERRORS:** Any errors surfaced from the completed Job

Automation

Provisioning -> Automation

The Automation section is composed of Tasks and Workflows. Tasks can be scripts added directly, scripts and blueprints from the Library section, recipes, playbooks, salt states, puppet agent installs, or http (api) calls. These Tasks are combined into workflows, which can be selected to run at provision time or executed on existing instances via `Actions -> Run Workflow`.

Tasks

Overview

There are many Task Types available, including scripts added directly, scripts and templates from the Library section, recipes, playbooks, salt states, puppet agent installs, and http (api) calls. Tasks are primarily created for use in Workflows, but a single Task can be executed on an existing instance via `Actions -> Run Task`.

Role Permissions

The User Role Permission 'Provisioning: Tasks FULL' is required to create, edit and delete tasks.

Tasks Types that can execute locally against the Morpheus Appliance have an additional Role Permission: `Tasks - Script Engines`. Script Engine Task Types will be hidden for users without `Tasks - Script Engines` role permissions.


Task Types

Table 7: Available Task Types

	Task Type	Task Description	Source Options	Execute Target Options	Configuration Requirements	Role Permissions Requirements
	Ansible	Runs an Ansible playbook. Ansible Integration required	Ansible Repo (Git)	Local, Resource	Existing Ansible Integration	Provisioning: Tasks
	Ansible Tower	Relays Ansible calls to Ansible Tower	Tower Integration	Local, Remote, Resource	Existing Ansible Tower Integration	Provisioning: Tasks
	Chef bootstrap	Executes Chef bootstrap and run list. Chef Integration required	Chef Server	Resource	Existing Chef Integration	Provisioning: Tasks
	Email	Send an email from a Workflow	Task Content	Local	SMTP Configured	Provisioning: Tasks
	Groovy script	Executes Groovy Script locally (on Morpheus app node)	Local, Repository, Url	Local	None	Provisioning: Tasks, Tasks - Script Engines
	HTTP	Executes REST call for targeting external API's.	Local	Local	None	Provisioning: Tasks
	Javascript	Executes Javascript locally (on Morpheus app node)	Local	Local	None	Provisioning: Tasks, Tasks - Script Engines
	jRuby Script	Executes Ruby script locally (on Morpheus app node)	Local, Repository, Url	Local	None	Provisioning: Tasks, Tasks - Script Engines
	Library Script	Creates a Task from an existing Library Script (Provisioning -> Library -> Scripts)	Library Script	Resource	Existing Library Script	Provisioning: Tasks
	Library Template	Creates a Task from an existing Library Template (Provisioning -> Library-> Templates)	Library Template	Resource	Existing Library Templates	Provisioning: Tasks
	PowerShell Script	Execute PowerShell Script on the Target Resource	Local, Repository, Url	Remote, Resource	None	Provisioning: Tasks
1.3. Hide Blueprint fields			Url			127
	Puppet Agent Install	Executes Puppet Agent bootstrap, writes puppet.conf and triggers agent checkin.	Puppet Mas-	Resource	Existing Puppet Integration	Provisioning: Tasks

Task Configuration

- **Ansible Playbook**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **ANSIBLE REPO:** Select existing Ansible Integration
- **GIT REF:** Specify tag or branch (Option, blank assumes default)
- **PLAYBOOK:** Name of playbook to execute, both `playbook` and `playbook.yml` format supported
- **TAGS:** Enter comma separated tags to filter executed tasks by (ie `--tags`)
- **SKIP TAGS:** Enter comma separated tags to run the playbook without matching tagged tasks (ie `--skip-tags`)

Important: Using different Git Refs for multiple Ansible Tasks in same Workflow is not supported. Git Refs can vary between Workflows, but Tasks in each Workflow must use the same Git Ref.

- **Chef bootstrap**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **CHEF SERVER:** Select existing Chef integration
- **ENVIRONMENT:** Populate Chef environment, or leave as `_default`
- **RUN LIST:** Enter Run List, eg `role[web]`
- **DATA BAG KEY:** Enter data bag key (will be masked upon save)
- **DATA BAG KEY PATH:** Enter data bag key path, eg `/etc/chef/databag_secret`
- **NODE NAME:** Defaults to Instance name, configurable
- **NODE ATTRIBUTES:** Specify attributes inside the `{ }`

- **Groovy script**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **CONTENT:** Contents of the Groovy script if not sourcing it from a repository

- Email



Email

- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **SOURCE:** Choose local to draft or paste the email directly into the Task. Choose Repository or URL to bring in a template from a Git repository or another outside source
- **EMAIL ADDRESS:** Email addresses can be entered literally or Morpheus automation variables can be injected, such as `<%=instance.createdByEmail%>`
- **SUBJECT:** The subject line of the email, Morpheus automation variables can be injected into the subject field
- **CONTENT:** The body of the email is HTML. Morpheus automation variables can be injected into the email body when needed
- **SKIP WRAPPED EMAIL TEMPLATE:** The Morpheus-styled email template is ignored and only HTML in the Content field is used

Tip: To whitelabel email sent from Tasks, select SKIP WRAPPED EMAIL TEMPLATE and use an HTML template with your own CSS styling

- HTTP (API)



HTTP

- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **URL:** An HTTP or HTTPS URL as the HTTP Task target
- **HTTP METHOD:** GET (default), POST, PUT, PATCH, HEAD, or DELETE
- **AUTH USER:** Username for username/password authentication
- **PASSWORD:** Password for username/password authentication
- **BODY:** Request Body
- **HTTP HEADERS:** Enter requests headers, examples below:

Authorization	Bearer <i>token</i>
Content-Type	application/json

- **IGNORE SSL ERRORS:** Mark when making REST calls to systems without a trusted SSL certificate

- Javascript



JavaScript

- **NAME:** Name of the Task

- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **SCRIPT:** Javascript contents to execute

- **jRuby Script**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **CONTENT:** Contents of the jRuby script is entered here if it's not being called in from an outside source

- **Library Script**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **SCRIPT:** Search for an existing script in the typeahead field

- **Library Template**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **TEMPLATE:** Search for an existing template in the typeahead field

- **Powershell Script**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **ELEVATED SHELL:** Run script with administrator privileges
- **IP ADDRESS:** IP address of the PowerShell Task target
- **PORT:** SSH port for PowerShell Task target (5985 default)
- **USERNAME:** Username for PowerShell Task target
- **PASSWORD:** Password for PowerShell Task target
- **Content:** Enter script to execute if not calling the script in from an outside source

- **Puppet Agent Install**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **PUPPET MASTER:** Select Puppet Master from an existing Puppet integration
- **PUPPET NODE NAME:** Enter Puppet node name. Variables supported eg. `<%= instance.name %>`
- **PUPPET ENVIRONMENT:** Enter Puppet environment, eg. `production`

- **Python Script**



Important: Beginning with Morpheus version 4.2.1, Python Tasks use virtual environments. For this reason, “virtualenv” must be installed on your appliances in order to work with Python tasks. Connect to the appliance node(s) and run “pip install virtualenv”.

- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **CONTENT:** Python script to execute is entered here if not pulled in from an outside repository
- **COMMAND ARGUMENTS:** Optional arguments passed into the Python script. Variables supported eg. `<%= instance.name %>`
- **ADDITIONAL PACKAGES:** Additional packages to be installed after `requirements.txt` (if detected). Expected format for additional packages: ‘packageName==x.x.x packageName2==x.x.x’, the version must be specified
- **PYTHON BINARY:** Optional binary to override the default Python binary

- **Restart**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references

- **Shell Script**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON

- **CONTENT:** Script to execute is entered here if not pulled in from an outside repository

- **vRealize Orchestrator Workflow**



- **NAME:** Name of the Task
- **CODE:** Unique code name for API, CLI, and variable references
- **RESULT TYPE:** Single Value, Key/Value Pairs, or JSON
- **vRO INTEGRATION:** Select an existing vRO integration
- **WORKFLOW:** Select a vRO workflow from the list synced from the selected integration
- **PARAMETER BODY (JSON):**

Task Management

Adding Tasks

1. Select the Provisioning link in the navigation bar.
2. Select Automation from the sub-navigation menu.
3. Click the *Add* button.
4. From the New Task Wizard input a name for the task.
5. Select the type of task from the type dropdown.
6. Input the appropriate details dependent on the task type you selected from the dropdown.
7. Save

Editing Tasks

1. Select the Provisioning link in the navigation bar.
2. Select Automation from the sub-navigation menu.
3. Click the Edit icon on the row of the task you wish to edit.
4. Modify information as needed.
5. Click the Save Changes button to save.

Deleting Tasks

1. Select the Provisioning link in the navigation bar.
2. Select Automation from the sub-navigation menu.
3. Click the Delete icon on the row of the task you wish to delete.

Task Results

Overview

Task Results allow Tasks to use the output from preceding Tasks in the same Workflow phase via results variables.

Results are available for all tasks executed in the same phase in a workflow. For example, instead of using just one Task's results in another Task, we can use all of the Task Results from the tasks in the same provision phase in a single task inside a workflow.

Configure Tasks

In script type tasks, if `RESULT TYPE` is set, Morpheus will store the Task's output as a variable.

Results Types

- **Single Value** Entire task output is stored in `<%=results.taskCode%>` or `<%=results["Task Name"]%>` variable.
- **Key/Value pairs** Expects `key=value,key=value` output. Entire task output is available with `<%=results.taskCode%>` or `<%=results["Task Name"]%>` variable (output inside []). Individual Values are available with `<%=results.taskCode.key%>` variables.
- **JSON** Expects `key:value,key:value` json formatted output. Entire task output is available with `<%=results.taskCode%>` or `<%=results["Task Name"]%>` variable (output inside []). Individual Values are available with `<%=results.taskCode.key%>` variables.

Important: The entire output of a script is treated as results, not just the last line. Ensure formatting is correct for the appropriate result type. For example, if Results Type is `json` and the output is not fully json compatible, the result would not return properly.

Important: Task results are not supported for Library Script task types

Script Config Examples

Single Value using Task Code

Source Task Config

```
NAME Var Code (single)
CODE single
RESULT TYPE Single Value
SCRIPT echo "string value"
```

Source Task Output `string value`

Results Task using task code in variable

```
Results Task Script echo "single: <%=results.single%>"
```

Results Task Output single: string value

Single Value using Task Name

Source Task Config

NAME Var Code

CODE none

RESULT TYPE Single Value

SCRIPT echo "string value"

Source Task Output string value

Results Task using task name in variable

Results Task Script echo "task name: <%=results["Var Code"]%>"

Results Task Output task name: test value

Key/Value Pairs

Source Task Config

NAME Var Code (keyval)

CODE keyval

RESULT TYPE Key/Value pairs

SCRIPT echo "flash=bang,ping=pong"

Source Task Output flash=bang,ping=pong

Results Task for all results

Results Task Script echo "keyval: <%=results.keyval%>"

Results Task Output keyval: [flash:bang, ping:pong]

Results Task for a single value)

Results Task Script echo "keyval value: <%=results.keyval.flash%>"

Results Task Output keyval value: bang

JSON

Source Task Config

NAME Var Code (json)

CODE json

RESULT TYPE JSON

SCRIPT echo "{\"ping\":\"pong\",\"flash\":\"bang\"}"

Source Task Output {"ping":"pong","flash":"bang"}

Results Task for all results

Results Task Script echo "json: <%=results.json%>"

Results Task Output json: [ping:pong, flash:bang]

Results Task for a single value

Results Task Script echo "json value: <%=results.json.ping%>"

Results Task Output json value: pong

Multiple Task Results

Results Task Script

```
echo "single: <%=results.single%>"
echo "task name: <%=results["Var Code"]%>"
echo "keyval: <%=results.keyval%>"
echo "keyval value: <%=results.keyval.flash%>"
echo "json: <%=results.json%>"
echo "json value: <%=results.json.ping%>"
```

Results Task Output

```
single: string value
task name: string value
keyval: [flash:bang, ping:pong]
keyval value: bang
json: [ping:pong, flash:bang]
json value: pong
```

Workflow Config

Add one or multiple tasks with Results Type configured to a workflow, and the results will be available to all tasks in the same phase of the workflow via the `<%=results.variables%>` during the workflow execution.

- Task Results are only available to tasks in the same workflow phase
- Task Results are only available during workflow execution

Workflows

Workflows are groups of Tasks, which are described in detail in the preceding section. Operational Workflows can be run on-demand against an existing Instance or server from the Actions menu on the Instance or server detail page. Additionally, they can be scheduled to run on a recurring basis through Morpheus Jobs (Provisioning > Jobs).

Provisioning Workflows are associated with Instances at provision time (in the Automation tab of the Add Instance wizard) or after provisioning through the Actions menu on the Instance detail page. Provisioning Workflows assign Tasks to various stages of the Instance lifecycle, such as Provision, Post Provision, and Teardown. When the Instance reaches a given stage, the appropriate Tasks are run. Task results and output can be viewed from the History tab of the Instance or server detail page.

Provisioning Workflow Execution Phases

- **Configuration:** Tasks are run prior to the initial provisioning
- **Pre Provision:** Tasks are run after the VM is running, for containers these Tasks are executed on the Docker host
- **Provision:** Tasks are run during provisioning, for many users this is the most commonly used phase
- **Post Provision:** Tasks are run after provisioning has completed
- **Start Service:** Tasks are run during start services
- **Stop Service:** Tasks are run during stop services

- **Pre Deploy:** Tasks are run prior to App deployment
- **Deploy:** Tasks are run during App deployment
- **Reconfigure:** Tasks are run during a reconfigure action on the Instance or host
- **Teardown:** Tasks are run during VM or container destroy

For VMs, Pre Provision and Provision phases execute after the VM is running. Pre Provision can be used for a Blueprint so it is added before a script which is set at the Provision phase executes. Pre Provision for scripts is mainly for Docker as you can execute on the host before the container is running. For Tasks that need to run prior to the start of provisioning, use the Configuration phase. Post Provision will execute after the entire provisioning process is complete.

Note: When adding a node to an Instance, Workflow Tasks in the Post Provision phase will be run on all nodes in the Instance after the new node is provisioned. This is because Post Provision operations may need to affect all nodes, such as when joining a new node to a cluster. Tasks in the Pre Provision and Provision phases would only be run on the new node.

Add Workflow

1. Select the Provisioning link in the navigation bar.
2. Select Automation from the sub-navigation menu.
3. Click the Workflows tab to show the Workflows tab panel.
4. Click the *Add* button.
5. From the New Workflow Wizard input a name for the workflow.
6. Optionally input a description.
7. Expand the execution phases to add tasks to, and type the name of a created task and click the task when it appears to add.
8. If multiple tasks are added to the same execution phase, their execution order can be changed by selecting the grip icon and dragging the task to the desired execution order.
9. For multi-tenant environments, select Public or Private visibility for the Workflow.
10. Click the *Save Changes* button to save.

Note: When setting Workflow visibility to Public in a multi-Tenant environment, Tenants will be able to see the Workflow and also execute it directly from the Workflows list (if it's an Operational Workflow). They will not be able to edit or delete the Workflow.

Edit Workflow

1. Select the Provisioning link in the navigation bar.
2. Select Automation from the sub-navigation menu.
3. Click the Workflows tab to show the workflows tab panel.
4. Click the Edit icon on the row of the workflow you wish to edit.
5. Modify information as needed.
6. Click the *Save Changes* button to save.

Delete Workflow

1. Select the Provisioning link in the navigation bar.
2. Select Automation from the sub-navigation menu.
3. Click the Workflows tab to show the workflows tab panel.
4. Click the Delete icon on the row of the workflow you wish to delete.

Executions

Automation - Executions contains execution status and history from Task and Operational Workflow Executions run from Automation - Tasks and Automation - Workflows.

Note: Tasks and Workflows executed from a Job or from Instance or Host Actions do not populate in Automation -> Executions, and can be referenced from the History tab on the target resource. All task and Workflow executions can be referenced in Operations -> Activity -> History

Execution results in the ui display:

NAME Name of the Task or Workflow Executed

TYPE Type of execution (Task or Workflow)

START DATE Date and time of execution

ETA/DURATION Estimate time of completion for executions in progress, or the total execution time for completed executions.

RESULTS Result status of execution (Succeeded, Failed, In-Progress or Pending)

Execution Detail (i) Click on the *i* to view process output results

Note: Job and automation executions can be expanded to show process details by clicking on the arrow icon immediately to the right of the *NAME* column.

Scale Thresholds

Scale Thresholds are pre-configured settings for auto-scaling Instances. When adding auto-scaling to an instance, existing Scale Thresholds can be selected to determine auto-scaling rules.

Creating Scale Thresholds

1. Navigate to Provisioning -> Automation -> Scale Thresholds
2. Select + *ADD*
3. Populate the following:

NAME Name of the Scale Threshold

AUTO UPSCALE Enable to automatically upscale per Scale Threshold specifications

AUTO DOWNSCALE Enable to automatically downscale per Scale Threshold specifications

MIN COUNT Minimum node count for Instance. Auto-scaling will not downscale below MIN COUNT, and will auto upscale if the MIN COUNT is not met)

MAX COUNT Maximum node count for Instance. Auto-scaling will not upscale past MAX COUNT, and will auto downscale if MAX COUNT is exceeded.

ENABLE MEMORY THRESHOLD Check to set auto-scaling by specified memory utilization threshold (%)

MIN MEMORY Enter MIN MEMORY % for triggering downscaling.

MAX MEMORY Enter MAX MEMORY % for triggering upscaling.

ENABLE DISK THRESHOLD Check to set auto-scaling by specified disk utilization threshold (%)

MIN DISK Enter MIN DISK % for triggering downscaling.

MAX DISK Enter MAX DISK % for triggering upscaling.

ENABLE CPU THRESHOLD Check to set auto-scaling by specified overall CPU utilization threshold (%)

MIN CPU Enter MIN CPU % for triggering downscaling.

MAX CPU Enter MAX CPU % for triggering upscaling.

Integrations

The Integrations section in Provisioning -> Automaton -> Integrations is for adding and managing Automation Integrations. Existing Automation Integrations from Administration -> Integrations are also populated and accessible from Provisioning -> Automaton -> Integrations and vice-versa.

Provisioning: Automation Integrations and **Admin:** Integrations are separate Role permissions, so Automations Integration access can be separated from the Administration Integrations section.

Automation Integrations

- Ansible
- Ansible Tower
- Chef
- Puppet
- Salt

Note: Automation integrations can be added and managed in Administration -> integrations as well. Adding and editing Integrations in Provisioning -> Automaton -> Integrations and Administration -> Integrations are the same dataset and additions and updates are reflected in both sections.

Note: Role access for Integrations: Ansible determines user access to Ansible Integration detail page, which contains Ansible command line and execution access.

Power Scheduling

Set weekly schedules for shutdown and startup times for Instances and VM's, apply Power Schedules to Instances pre or post-provisioning, apply Power Schedule policies on Group or Clouds, or use Guidance to automatically recommend and apply optimized Power Schedules.

Create Power schedules

1. Navigate to Provisioning -> Automation -> Power Scheduling
2. Select + *ADD*
3. Configure the following options:

NAME Name of the Power Schedule

DESCRIPTION Description for the Power Schedule

TIME ZONE Time Zone the Power Schedule times correlate to.

TYPE

Power On Power Up and then Down at scheduled times

Power off Power Down then Up at scheduled times

Enabled Check for Power Schedule to be Active. Uncheck to disable Power Schedule.

DAYS Slide the start and end time controls for each day to configure each days Schedule. Green sections indicate Power on, red sections indicate Power Off. Time indicated applies to selected Time Zone.

EDIT POWER SCHEDULE

X

NAME

Nighty Night

DESCRIPTION

Down at 6pm, up at at 6am weekdays. Off Weekends

TIME ZONE

America/Denver

▼

TYPE

Power On

▼

☒ Enabled

MONDAY

6:00

18:00

TUESDAY

6:00

18:00

WEDNESDAY

6:00

18:00

THURSDAY

6:00

18:00

FRIDAY

6:00

18:00

SATURDAY

0:00

SUNDAY

0:00

SAVE CHANGES

4. Select *SAVE CHANGES*

Tip: To view the Instances a power schedule is currently set on, select the name of a Power Schedule to go to the Power Schedule Detail Page.

Add Power Schedule to Instance

1. Navigate to Provisioning -> Instances
2. Select an Instance
3. Select *EDIT*
4. In the POWER SCHEDULE dropdown, select a Power Schedule.
5. Select *SAVE CHANGES*

Add Power Schedule to Virtual Machine

1. Navigate to Infrastructure -> Hosts -> Virtual Machines
2. Select a Virtual Machine
3. Select *EDIT*
4. Expand the Advanced Options section
5. In the POWER SCHEDULE dropdown, select a Power Schedule.
6. Select *SAVE CHANGES*

Add Power Schedule Policy

Note: Power Schedule Policies apply to Instances created after the Policy is enabled.

1. Navigate to Administration -> Policies
2. Select + *ADD*
3. Select TYPE *Power Schedule*
4. Configure the Power Schedule Policy:

NAME Name of the Policy

DESCRIPTION Add details about your Policy for reference in the Policies tab.

Enabled Policies can be edited and disabled or enabled at any time. Disabling a Power Schedule Policy will prevent the Power Schedule from running on the Clouds Instances until re-enabled.

ENFORCEMENT TYPE

- User Configurable: Power Schedule choice is editable by User during provisioning.
- Fixed Schedule: User cannot change Power Schedule setting during provisioning.

POWER SCHEDULE Select Power Schedule to use in the Policy. Power schedule can be added in Provisioning -> Automation -> Power Scheduling

SCOPE

Global Applies to all Instances created while the Policy is enabled

Group Applies to all Instances created in or moved into specified Group while the Policy is enabled

Cloud Applies to all Instances created in specified Cloud while the Policy is enabled

User Applies to all Instances created by specified User while the Policy is enabled

Role Applies to all Instances created by Users with specified Role while the Policy is enabled

Permissions- TENANTS Leave blank to apply to all Tenants, or search for and select Tenants to enforce the Policy on specific Tenants.

5. Select *SAVE CHANGES*

Execute Scheduling

Execute Scheduling creates time schedules for Jobs, including Task, Workflow and Backup Jobs. Jobs, which are discussed in greater detail in [another section](#) of Morpheus docs, combine either a Task or Workflow with an Execute Schedule to run the selected Task or Workflow at the needed time. Backup Jobs are a special type of Job configured in the Backups section which also use Execute Schedules to time backup runs as needed.

Schedules use CRON expressions, such as `0 23 * * 2` equalling *Executes every week on Tuesday at 23:00*. CRON expressions can easily be created by clicking the corresponding translation in the create or edit Execution Schedule modal below the Schedule field and selecting a new value.

Note: For more on writing CRON expressions, see our [KnowledgeBase article](#) on the topic.

Create Execution Schedules

NAME Name of the Execution Schedule

Note: When assigning Execution Schedules, the name value will appear in the selection drop-down. Using a name that references the time interval is often helpful

DESCRIPTION Description of the Execution Schedule for reference in the Execution Schedules list

TIME ZONE Time zone for execution

Enabled Check to enable the schedule. Uncheck to disable all associated executions and remove the schedule as an option for Jobs in the future

SCHEDULE Enter CRON expression for the Execution Schedule, for example `0 0 * * *` equals *Every day at 00:00*

SCHEDULE TRANSLATION The entered CRON schedule is translated below the SCHEDULE field. Highlighted values can be updated by selecting the value, and relevant options will be presented. The CRON expression will automatically be updated

Variables

A vast number of variables are available for use in Tasks, Scripts, Templates, Resource Names, Cloud-Init User Data and Option List configs.

Important: Variables are case sensitive

Pre-Provision Vars

A subset of variables are available for Instance, Host Name and Hostnames. These can be passed inside `${ }` blocks during provisioning or in relevant policy configs.

Instance Naming Policy example: `${userInitials}-${cloudCode}-${platform == 'windows' ? 'W' : 'L'}-${sequence}`

Commonly used variables for naming patterns include:

```

${groupName}
${groupCode}
${cloudName}
${cloudCode}
${type}
${accountId}
${account}
${accountType}
${platform}
${platform == 'windows' ? 'w':'l'} # results in `w` for Windows platforms and `l` for
↳Linux Platforms
${userId}
${username}
${userInitials}
${provisionType}
${instance.instanceContext} # Environment Code
${sequence} # results in 1
${sequence+100} # results in 101
${customOption.name}
${sequence.toString().padLeft(5,'0')} #results in 00001

```

An example Instance Name Policy using a naming pattern with User Initials, Cloud Code, Instance Type, and a sequential number starting at 3000 is `${userInitials}-${cloudCode}-${type}-${sequence+3000}`, resulting in an Instance Name of **md-vmwd3-centos-3001** for the first instance, followed by **md-vmwd3-centos-3002** and so on.

Syntax Examples

PowerShell Example: `$app_id = "<%= instance.metadata.app_id %>"`

Bash Example: `HOSTNAME="<%= container.server.hostname %>"`

Python Example: `hostname = morpheus['server']['hostname']`

HTTP Body Example: `{"name": "<%= instance.createdByUsername %>"}`

Note: customOptions values are defined from custom Option Types.

Common Examples

```
container.configGroup: <%=container.configGroup%>
container.configId: <%=container.configId%>
container.configPath: <%=container.configPath%>
container.configRole: <%=container.configRole%>
container.containerTypeCode: <%=container.containerTypeCode%>
container.containerTypeName: <%=container.containerTypeName%>
container.containerTypeShortName: <%=container.containerTypeShortName%>
container.cores: <%=container.cores%>
container.dataPath: <%=container.dataPath%>
container.dateCreated: <%=container.dateCreated%>
container.domainName: <%=container.domainName%>
container.environmentPrefix: <%=container.environmentPrefix%>
container.externalIp: <%=container.externalIp%>
container.hostMountPoint: <%=container.hostMountPoint%>
container.hostname: <%=container.hostname%>
container.image: <%=container.image%>
container.internalHostname: <%=container.internalHostname%>
container.internalIp: <%=container.internalIp%>
container.logsPath: <%=container.logsPath%>
container.memory: <%=container.memory%>
container.planCode: <%=container.planCode%>
container.provisionType: <%=container.provisionType%>
container.server: <%=container.server.serverTypeName%>
container.serverId: <%=container.serverId%>
container.sshHost: <%=container.sshHost%>
container.status: <%=container.status%>
container.storage: <%=container.storage%>
container.version: <%=container.version%>
customOptions: <%=customOptions.fieldName%>
evar: <%=evars.name%>
evars: <%=evars%>
group.code: <%=group.code%>
group.datacenterId: <%=group.datacenterId%>
group.location: <%=group.location%>
group.name: <%=group.name%>
instance.autoScale: <%=instance.autoScale%>
instance.configGroup: <%=instance.configGroup%>
instance.configId: <%=instance.configId%>
instance.configRole: <%=instance.configRole%>
instance.containers[0]: <%=instance.containers[0].containerTypeName%>
instance.cores: <%=instance.cores%>
instance.createdByEmail: <%=instance.createdByEmail%>
instance.createdByFirstName: <%=instance.createdByFirstName%>
instance.createdById: <%=instance.createdById%>
instance.createdByLastName: <%=instance.createdByLastName%>
instance.createdByUsername: <%=instance.createdByUsername%>
instance.deployGroup: <%=instance.deployGroup%>
instance.description: <%=instance.description%>
instance.displayName: <%=instance.displayName%>
instance.domainName: <%=instance.domainName%>
```

(continues on next page)

(continued from previous page)

```

instance.environmentPrefix: <%=instance.environmentPrefix%>
instance.expireDate: <%=instance.expireDate%>
instance.firewallEnabled: <%=instance.firewallEnabled%>
instance.hostname: <%=instance.hostname%>
instance.instanceContext: <%=instance.instanceContext%> (tip: instanceContext = Environment)
instance.instanceLevel: <%=instance.instanceLevel%>
instance.instanceTypeCode: <%=instance.instanceTypeCode%>
instance.instanceTypeName: <%=instance.instanceTypeName%>
instance.instanceVersion: <%=instance.instanceVersion%>
instance.memory: <%=instance.memory%>
instance.metadata: <%=instance.metadata%>
instance.name: <%=instance.name%>
instance.networkLevel: <%=instance.networkLevel%>
instance.plan: <%=instance.plan%>
instance.provisionType: <%=instance.provisionType%>
instance.status: <%=instance.status%>
instance.statusMessage: <%=instance.statusMessage%>
instance.storage: <%=instance.storage%>
instance.tags: <%=instance.tags%>
instance.userStatus: <%=instance.userStatus%>
server.agentInstalled: <%=server.agentInstalled%>
server.agentVersion: <%=server.agentVersion%>
server.apiKey: <%=server.apiKey%>
server.category: <%=server.category%>
server.commType: <%=server.commType%>
server.configGroup: <%=server.configGroup%>
server.configId: <%=server.configId%>
server.configRole: <%=server.configRole%>
server.consoleHost: <%=server.consoleHost%>
server.consolePort: <%=server.consolePort%>
server.consoleType: <%=server.consoleType%>
server.consoleUsername: <%=server.consoleUsername%>
server.dataDevice: <%=server.dataDevice%>
server.dateCreated: <%=server.dateCreated%>
server.description: <%=server.description%>
server.displayName: <%=server.displayName%>
server.domainName: <%=server.domainName%>
server.externalId: <%=server.externalId%>
server.externalIp: <%=server.externalIp%>
server.fqdn: <%=server.fqdn%>
server.hostname: <%=server.hostname%>
server.internalId: <%=server.internalId%>
server.internalIp: <%=server.internalIp%>
server.internalName: <%=server.internalName%>
server.internalSshUsername: <%=server.internalSshUsername%>
server.lastAgentUpdate: <%=server.lastAgentUpdate%>
server.lvmEnabled: <%=server.lvmEnabled%>
server.macAddress: <%=server.macAddress%>
server.managed: <%=server.managed%>
server.maxCores: <%=server.maxCores%>
server.maxMemory: <%=server.maxMemory%>
server.maxStorage: <%=server.maxStorage%>
server.name: <%=server.name%>
server.nodePackageVersion: <%=server.nodePackageVersion%>
server.osDevice: <%=server.osDevice%>
server.osType: <%=server.osType%>

```

(continues on next page)

(continued from previous page)

```

server.osTypeCode: <%=server.osTypeCode%>
server.parentServerId: <%=server.parentServerId%>
server.plan: <%=server.plan%>
server.platform: <%=server.platform%>
server.platformVersion: <%=server.platformVersion%>
server.powerState: <%=server.powerState%>
server.serialNumber: <%=server.serialNumber%>
server.serverModel: <%=server.serverModel%>
server.serverType: <%=server.serverType%>
server.serverTypeCode: <%=server.serverTypeCode%>
server.serverTypeName: <%=server.serverTypeName%>
server.serverVendor: <%=server.serverVendor%>
server.softwareRaid: <%=server.softwareRaid%>
server.sourceImageId: <%=server.sourceImageId%>
server.sshHost: <%=server.sshHost%>
server.sshPort: <%=server.sshPort%>
server.sshUsername: <%=server.sshUsername%>
server.status: <%=server.status%>
server.statusMessage: <%=server.statusMessage%>
server.tags: <%=server.tags%>
server.toolsInstalled: <%=server.toolsInstalled%>
server.visibility: <%=server.visibility%>
task.results (using task code): <%=results.taskCode%>
task.results (using task name): <%=results["Task Name"]%>
task.results.value: <%=results.taskCode.key%>
zone.agentMode: <%=zone.agentMode%>
zone.cloudTypeCode: <%=zone.cloudTypeCode%>
zone.cloudTypeName: <%=zone.cloudTypeName%>
zone.code: <%=zone.code%>
zone.domainName: <%=zone.domainName%>
zone.firewallEnabled: <%=zone.firewallEnabled%>
zone.location: <%=zone.location%>
zone.name: <%=zone.name%>
zone.regionCode: <%=zone.regionCode%>
zone.scalePriority: <%=zone.scalePriority%>
cypher: <%=cypher.read('secret/hello')%>

```

Instance

```

instance {
    autoScale,
    configGroup,
    configId,
    configRole
    containers:[],
    cores,
    deployGroup,
    description,
    displayName,
    domainName,
    environmentPrefix,
    evars:[],
    expireDate,
    firewallEnabled,

```

(continues on next page)

(continued from previous page)

```
hostname,  
instanceContext,  
instanceLevel,  
instanceTypeCode,  
instanceVersion,  
memory,  
metadata:[],  
name,  
networkLevel,  
plan,  
provisionType,  
status,  
statusMessage,  
storage,  
tags,  
tenantSubdomain,  
userStatus,  
instanceTypeName  
}
```

Container

```
container {  
    configGroup,  
    configId,  
    configPath,  
    configRole,  
    containerTypeCode,  
    containerTypeShortName,  
    cores,  
    dataPath,  
    dateCreated,  
    domainName,  
    environmentPrefix,  
    externalIp,  
    hostMountPoint,  
    hostname,  
    image,  
    internalHostname,  
    internalIp,  
    logsPath,  
    memory,  
    planCode,  
    provisionType,  
    server:{},  
    serverId,  
    sshHost,  
    status,  
    storage,  
    version,  
    containerTypeName  
}
```

Server

```
server {
    agentInstalled,
    agentVersion,
    apiKey,
    category,
    commType,
    configGroup,
    configId,
    configRole,
    consoleHost,
    consolePort,
    consoleType,
    consoleUsername,
    dataDevice,
    dateCreated,
    description,
    displayName,
    domainName,
    externalId,
    externalIp,
    fqdn,
    hostname,
    internalId,
    internalIp,
    internalName,
    internalSshUsername,
    lastAgentUpdate,
    lvmEnabled,
    macAddress,
    managed,
    maxCores,
    maxMemory,
    maxStorage,
    name,
    nodePackageVersion,
    osDevice,
    osType,
    osTypeCode,
    parentServerId,
    plan,
    platform,
    platformVersion,
    powerState,
    serialNumber,
    serverModel,
    serverType,
    serverTypeCode,
    serverTypeName,
    serverVendor,
    softwareRaid,
    sourceImageId,
    sshHost,
    sshPort,
    sshUsername,
    status,
```

(continues on next page)

(continued from previous page)

```
statusMessage,  
tags,  
toolsInstalled,  
visibility,  
volumes {  
    name  
    id  
    deviceName  
    maxStorage  
    unitNumber  
    displayOrder  
    rootVolume  
}  
}
```

Zone (Cloud)

```
zone {  
    agentMode,  
    cloudTypeCode,  
    cloudTypeName,  
    code,  
    datacenterId,  
    domainName,  
    firewallEnabled,  
    location,  
    name,  
    regionCode,  
    scalePriority  
}
```

Group (Site)

```
group {  
    code,  
    location,  
    datacenterId,  
    name  
}
```

Custom Options (Option Types)

```
customOptions {  
    customOptions.fieldName  
}
```

Global

ex: <%= morpheus.user.id %>

```
"morpheus":{
  "user":{
    "id":value,
    "account":{
      "id":value
    },
    "username":"value",
    "displayName":"value",
    "email":"value",
    "firstName":"value",
    "lastName":"value",
    "dateCreated":0000-00-00T00:00:00Z,
    "lastUpdated":0000-00-00T00:00:00Z,
    "enabled":true/false,
    "accountExpired":true/false,
    "accountLocked":false,
    "passwordExpired":false,
    "defaultGroupId":value,
    "defaultZoneId":value,
    "hasLinuxUser":true/false,
    "hasWindowsUser":true/false,
    "role":{
      "id":value
    },
    "instanceLimits":value
  },
}
```

Instance Map Example

```
"instance":{
  "poolProviderType":value,
  "isVpcSelectable":true/false,
  "smbiosAssetTag":value,
  "isEC2":true/false,
  "resourcePoolId":value,
  "hostId":value,
  "createUser":true/false,
  "nestedVirtualization":value,
  "vmwareFolderId":value,
  "expose":[

  ],
  "noAgent":value,
  "customOptions":value,
  "createBackup":true/false,
  "memoryDisplay":"MB/GB",
  "backup":{
    "veeamManagedServer":,
    "createBackup":true/false,
    "jobAction":"value",
  }
}
```

(continues on next page)

(continued from previous page)

```

    "jobRetentionCount":value
  },
  "expireDays":value,
  "layoutSize":value,
  "lbInstances":[

  ],
  "evars":{
    "evar1":{
      "value":value,
      "export":true/false,
      "masked":true/false,
      "name":"value"
    },
    "evar2":{
      "value":value,
      "export":true/false,
      "masked":true/false,
      "name":"value"
    }
  },
  "id":value,
  "instanceTypeName":"value",
  "instanceTypeCode":"value",
  "provisionType":"value",
  "layoutId":value,
  "layoutCode":value,
  "layoutName":"value",
  "instanceVersion":"value",
  "plan":value,
  "name":value,
  "displayName":value,
  "description":value,
  "environmentPrefix":value,
  "hostname":value,
  "domainName":"value",
  "assignedDomainName":,
  "firewallEnabled":true/false,
  "status":"value",
  "userStatus":"value",
  "scheduleStatus":"value",
  "networkLevel":"value",
  "instanceLevel":"value",
  "deployGroup":value,
  "instanceContext":value,
  "autoScale":true/false,
  "statusMessage":value,
  "expireDate":0000-00-00T00:00:00Z,
  "tags":"value",
  "storage":value (bytes),
  "memory":value (bytes),
  "cores":1,
  "configId":value,
  "configGroup":value,
  "configRole":value,
  "ports":value,
  "sslEnabled":true/false,

```

(continues on next page)

(continued from previous page)

```

    "sslCertId":value,
    "serviceUsername":value,
    "servicePassword":value,
    "adminUsername":value,
    "adminPassword":value,
    "createdByUsername":"value",
    "createdByEmail":"value",
    "createdByFirstName":"value",
    "createdByLastName":"value",
    "createdById":value,
    "metadata":{
    },
    "createdByUser":{
        "username":"value",
        "displayName":"value",
        "firstName":"value",
        "lastName":"value",
        "email":"value",
        "linuxUsername":"value",
        "windowsUsername":"value"
    },
    "containers":[
        {
            "maxMemory":value (bytes),
            "maxStorage":value (bytes),
            "maxCpu":value,
            "maxCores":value,
            "coresPerSocket":value,
            "poolProviderType":value,
            "isVpcSelectable":true/false,
            "smbiosAssetTag":value,
            "isEC2":true/false,
            "resourcePoolId":value,
            "hostId":value,
            "createUser":true/false,
            "nestedVirtualization":value,
            "vmwareFolderId":value,
            "expose":[
            ],
            "noAgent":true/false,
            "vm":true/false,
            "networkInterfaces":[
                {
                    "id":value,
                    "network":{
                        "id":value,
                        "group":value,
                        "subnet":value,
                        "dhcpServer":true/false,
                        "name":value,
                        "pool":{
                            "id":value,
                            "name":value
                        }
                    }
                }
            ],
        }
    ],

```

(continues on next page)

(continued from previous page)

```

        "ipAddress":value,
        "networkInterfaceTypeId":value,
        "ipMode":
    }
],
"volumes":[
    {
        "volumeCustomizable":true/false,
        "readonlyName":true/false,
        "controllerId":value,
        "maxIOPS":value,
        "displayOrder":value,
        "unitNumber":value,
        "minStorage":value(bytes),
        "configurableIOPS":true/false,
        "controllerMountPoint":0000:0:00:0,
        "vId":value,
        "size":value,
        "name":"root",
        "rootVolume":true/false,
        "storageType":value,
        "typeId":value,
        "id":value,
        "resizeable":true/false,
        "datastoreId":"value",
        "maxStorage":value(bytes)
    }
],
"storageController":value,
"datastoreId":value,
"networkId":value,
"cpuCount":value,
"memorySize":value,
"osDiskSize":value,
"publicKeyId":value,
"storagePodId":value,
"vmwareUsr":value,
"vmwarePwd":value,
"domainName":"value",
"hostname":value,
"networkType":value,
"ipAddress":value,
"netmask":value,
"gateway":value,
"dnsServers":value,
"resourcePool":value,
"folder":value,
"vmwareCustomSpec":value,
"hosts":{
    value
},
"evars":{

},
"id":value,
"name":value,
"containerTypeName":value,

```

(continues on next page)

(continued from previous page)

```

"containerTypeCode":value,
"containerTypeShortName":"value",
"containerTypeCategory":"value",
"provisionType":"value",
"dataPath":"value",
"logsPath":"value",
"configPath":"value",
"planCode":value,
"dateCreated":0000-00-00T00:00:00Z,
"status":"running",
"environmentPrefix":"value",
"version":"value",
"image":"value",
"internalHostname":value,
"storage":value (bytes),
"memory":value (bytes),
"cores":value,
"internalIp":value,
"externalIp":value,
"sshHost":value,
"hostMountPoint":value,
"configId":value,
"configGroup":value,
"configRole":value,
"certificatePath":value,
"certificateStyle":value,
"changeManagementExtId":value,
"changeManagementServiceId":value,
"serverId":value,
"server":{
  "poolProviderType":value,
  "isVpcSelectable":true/false,
  "smbiosAssetTag":value,
  isEC2:true/false,
  "resourcePoolId":value,
  "hostId":value,
  "createUser":true/false,
  "nestedVirtualization":value,
  "vmwareFolderId":value,
  "noAgent":value,
  "id":value,
  "uuid":value,
  "serverTypeName":"value",
  "serverTypeCode":"value",
  "computeTypeName":"value",
  "computeTypeCode":"value",
  "parentServerId":value,
  "plan":value,
  "visibility":"value",
  "osTypeCode":value,
  "sourceImageId":value,
  "name":value,
  "displayName":value,
  "internalName":value,
  "category":value,
  "description":value,
  "internalId":value,

```

(continues on next page)

(continued from previous page)

```

    "externalId":value,
    "platform":"value",
    "platformVersion":value,
    "agentVersion":value,
    "nodePackageVersion":value,
    "sshHost":value,
    "sshPort":value,
    "sshUsername":"value",
    "consoleType":value,
    "consoleHost":value,
    "consolePort":value,
    "consoleUsername":value,
    "internalSshUsername":"value",
    "internalIp":value,
    "externalIp":value,
    "osDevice":"value",
    "dataDevice":"value",
    "lvmEnabled":true/false,
    "apiKey":value,
    "softwareRaid":true/false,
    "status":"value",
    "powerState":"value",
    "dateCreated":0000-00-00T00:00:00Z,
    "lastAgentUpdate":0000-00-00T00:00:00Z,
    "serverType":"value",
    "osType":"value",
    "commType":"value",
    "managed":true/false,
    "agentInstalled":true/false,
    "toolsInstalled":true/false,
    "hostname":value,
    "domainName":value,
    "fqdn":value,
    "statusMessage":value,
    "maxStorage":value(bytes),
    "maxMemory":value(bytes),
    "maxCores":value,
    "macAddress":value,
    "serverVendor":value,
    "serverModel":value,
    "serialNumber":value,
    "tags":value,
    "configId":value,
    "configGroup":value,
    "configRole":value,
    "createdByUser":{
      "username":"value",
      "displayName":"value",
      "firstName":"value",
      "lastName":"value",
      "email":"value",
      "linuxUsername":"value",
      "windowsUsername":"value"
    },
    "volumes":[
      {
        "id":value,

```

(continues on next page)

(continued from previous page)

```

        "name": "value",
        "deviceName": "value",
        "maxStorage": value (bytes),
        "unitNumber": value,
        "displayOrder": value,
        "rootVolume": true/false
    }
]
},
"ports": [
    {
        "index": value,
        "external": value,
        "internal": value,
        "link": true/false,
        "loadBalance": true/false,
        "loadBalanceProtocol": value,
        "export": true/false,
        "exportName": value,
        "displayName": "value",
        "visible": true/false,
        "primaryPort": true/false,
        "protocol": value,
        "name": "value"
    }
],
"portMap": {
    "rpc": {
        "index": value,
        "external": value,
        "internal": value,
        "link": true/false,
        "loadBalance": true/false,
        "loadBalanceProtocol": value,
        "export": true/false,
        "exportName": value,
        "displayName": "value",
        "visible": true/false,
        "primaryPort": true/false,
        "protocol": value,
        "name": "value"
    }
},
    "internalPort": value,
    "externalPort": value
}
],
"container": {
    "maxMemory": value (bytes),
    "maxStorage": value,
    "maxCpu": value,
    "maxCores": value,
    "coresPerSocket": value,
    "poolProviderType": value,
    "isVpcSelectable": true/false,
    "smbiosAssetTag": value,
    isEC2: true/false,

```

(continues on next page)

(continued from previous page)

```

    "resourcePoolId":value,
    "hostId":value,
    "createUser":true/false,
    "nestedVirtualization":value,
    "vmwareFolderId":value,
    "expose":[

],
    "noAgent":true/false,
    "vm":true/false,
    "networkInterfaces":[
      {
        "id":value,
        "network":{
          "id":value,
          "group":value,
          "subnet":value,
          "dhcpServer":true/false,
          "name":value,
          "pool":{
            "id":value,
            "name":value
          }
        }
      },
      "ipAddress":value,
      "networkInterfaceTypeId":value,
      "ipMode":

    ]
  },
  "volumes":[
    {
      "volumeCustomizable":true/false,
      "readonlyName":true/false,
      "controllerId":value,
      "maxIOPS":value,
      "displayOrder":value,
      "unitNumber":value,
      "minStorage":value,
      "configurableIOPS":true/false,
      "controllerMountPoint":value,
      "vId":value,
      "size":value,
      "name":"root",
      "rootVolume":true/false,
      "storageType":value,
      "typeId":value,
      "id":value,
      "resizeable":true/false,
      "datastoreId":"autoCluster",
      "maxStorage":value (bytes)
    }
  ],
  "storageController":value,
  "datastoreId":value,
  "networkId":value,
  "cpuCount":value,
  "memorySize":value,

```

(continues on next page)

(continued from previous page)

```
"osDiskSize":value,
"publicKeyId":value,
"storagePodId":value,
"vmwareUsr":value,
"vmwarePwd":value,
"domainName":"value",
"hostname":value,
"networkType":value,
"ipAddress":value,
"netmask":value,
"gateway":value,
"dnsServers":value,
"resourcePool":value,
"folder":value,
"vmwareCustomSpec":value,
"hosts":{
  value
},
"evars":{

},
"id":value,
"name":value,
"containerTypeName":value,
"containerTypeCode":value,
"containerTypeShortName":"value",
"containerTypeCategory":"value",
"provisionType":"vmware",
"dataPath":"value",
"logsPath":"value",
"configPath":"value",
"planCode":value,
"dateCreated":"0000-00-00T00:00:00Z",
"status":"value",
"environmentPrefix":"value",
"version":"value",
"image":"value",
"internalHostname":value,
"storage":value(bytes),
"memory":value(bytes),
"cores":value,
"internalIp":value,
"externalIp":value,
"sshHost":value,
"hostMountPoint":value,
"configId":value,
"configGroup":value,
"configRole":value,
"certificatePath":value,
"certificateStyle":value,
"changeManagementExtId":value,
"changeManagementServiceId":value,
"serverId":value,
"server":{
  "poolProviderType":value,
  "isVpcSelectable":true/false,
  "smbiosAssetTag":value,
```

(continues on next page)

(continued from previous page)

```

isEC2:true/false,
"resourcePoolId":value,
"hostId":value,
"createUser":true/false,
"nestedVirtualization":value,
"vmwareFolderId":value,
"noAgent":value,
"id":value,
"uuid":value,
"serverTypeName":"value",
"serverTypeCode":"value",
"computeTypeName":"value",
"computeTypeCode":"value",
"parentServerId":value,
"plan":value,
"visibility":"value",
"osTypeCode":value,
"sourceImageId":value,
"name":value,
"displayName":value,
"internalName":value,
"category":value,
"description":value,
"internalId":value,
"externalId":value,
"platform":"value",
"platformVersion":value,
"agentVersion":value,
"nodePackageVersion":value,
"sshHost":value,
"sshPort":value,
"sshUsername":"value",
"consoleType":value,
"consoleHost":value,
"consolePort":value,
"consoleUsername":value,
"internalSshUsername":"value",
"internalIp":value,
"externalIp":value,
"osDevice":"value",
"dataDevice":"value",
"lvmEnabled":true/false,
"apiKey":value,
"softwareRaid":true/false,
"status":"provisioned",
"powerState":"on",
"dateCreated":0000-00-00T00:00:00Z,
"lastAgentUpdate":0000-00-00T00:00:00Z,
"serverType":"value",
"osType":"value",
"commType":"value",
"managed":true/false,
"agentInstalled":true/false,
"toolsInstalled":true/false,
"hostname":value,
"domainName":value,
"fqdn":value,

```

(continues on next page)

(continued from previous page)

```

    "statusMessage":value,
    "maxStorage":value,
    "maxMemory":value,
    "maxCores":value,
    "macAddress":value,
    "serverVendor":value,
    "serverModel":value,
    "serialNumber":value,
    "tags":value,
    "configId":value,
    "configGroup":value,
    "configRole":value,
    "createdByUser":{
        "username":"value",
        "displayName":"value",
        "firstName":"value",
        "lastName":"value",
        "email":"value",
        "linuxUsername":"value",
        "windowsUsername":"value"
    },
    "volumes":[
        {
            "id":value
            "name":"root",
            "deviceName":"value",
            "maxStorage":value(bytes),
            "unitNumber":value,
            "displayOrder":value,
            "rootVolume":true/false
        }
    ]
},
"ports":[
    {
        "index":0,
        "external":value,
        "internal":value,
        "link":true/false,
        "loadBalance":true/false,
        "loadBalanceProtocol":value,
        "export":true/false,
        "exportName":value,
        "displayName":"value",
        "visible":true/false,
        "primaryPort":true/false,
        "protocol":value,
        "name":"value"
    }
],
"portMap":{
    "rpc":{
        "index":0,
        "external":value,
        "internal":value,
        "link":true/false,
        "loadBalance":true/false,

```

(continues on next page)

(continued from previous page)

```

        "loadBalanceProtocol":value,
        "export":true/false,
        "exportName":value,
        "displayName":"value",
        "visible":true/false,
        "primaryPort":true/false,
        "protocol":value,
        "name":"value"
    },
    "internalPort":value,
    "externalPort":value
},
"apps":[
]
}

```

Virtual Images

Provisioning -> Virtual Images

Overview

The Virtual Image section displays a list of all images, local and synced, that are available to deploy. Morpheus includes a rich catalog of pre-configured System Images available for every cloud type. User Images are automatically synced from Cloud Integrations and added to the Virtual Images section. Images can also be uploaded directly into Morpheus via local file or url. Amazon and Azure Marketplace images can also be added to the Virtual Images Section.

Important: Invalid Image Settings cause provisioning failures. Morpheus syncs in as much meta-data as possible for synced images, but additional configuration may be needed to ensure successful provisioning.

Warning: Cloud-init is enabled by default for all Linux Images. If your Linux image does not have Cloud-init installed, *Cloud-init Enabled* must be unchecked before provisioning the image or it will fail immediately.

Image Types

Morpheus provides a vast *System Image* repo with pre-configured images for every Cloud. All other images are *User Images*. User images can be added directly to Morpheus, or automatically synced from integrated clouds. It is important to configure synced User Images for metadata, including specifying the Platform and User Credentials, prior to provisioning. Provisioning a User Image that has not been configured may result in failed provisioning.

Important: Synced User Images need to be configured prior to provisioning.

Configuring Virtual Images

System Images

System Virtual Images are pre-configured with metadata and have Cloud-Init or Cloudbase-Init installed. These images are ready to be provisioned with no configuration necessary, however it is required to populate *Administration -> Provisioning -> Cloud-Init* section with user data as well as User Profile(s) users data when creating additional users prior to provisioning, as the user data from these sections is required when provisioning System provided Virtual Images.

Note: System Images settings are not editable.

User Images

Typically Morpheus does not have sufficient metadata to successfully provision synced User Images. After integrating clouds and User Images have synced, it is highly recommended to configure the images prior to provisioning.

To edit and configure an existing Virtual Image:

1. Select *Actions - Edit* in the Virtual Images list, or *Edit* on a Virtual Image detail page.
2. Configure the following on the Image:

Name Name of the Virtual Image in Morpheus . This can be changed from the name of the Image, but editing will not change the name of the actual Image.

Operating System Specifies the Platform and OS of the image. All Windows images will need to have Operating System specified on the Virtual Image, as Morpheus will assign Linux as the Platform for all Images without Operating System specified.

Minimum Memory The Minimum Memory setting will filter available Service Plans options during provisioning. Service Plans that do not meet the Minimum Memory value set on the Virtual Image will not be provided as Service Plan choices.

Cloud Init Enabled? On by default, uncheck for any Image that does not have Cloud-Init or Cloudbase-Init installed.

Install Agent On by default, uncheck to skip Agent install. Note this will result in the loss of utilization statistics, logs, script execution, and monitoring. (Some utilization stats are collected for agent-less hosts and vm's from VMware and AWS clouds).

Username Existing Username on the Image. This is required for authentication, unless Morpheus is able to add user data, Cloud-Init, Cloudbase-Init or Guest Customizations. If Cloud-Init, Cloudbase-Init or Guest Customizations are used, credentials are defined in *Administration -> Provisioning* and *User Settings* . *If credentials are defined on the Image and Cloud-Init is enabled, \morpheus\ will add that user during provisioning, so ensure that user does not already exist n the image (aka ``root``)*. For Windows Guest Customizations, Morpheus will set the Administrator password to what is defined on the image if Administrator user is defined. Do not define any other user than Administrator for Windows Images unless using Cloudbase-init. Morpheus recommends running Guest Customizations for all Windows Images, which is required when joining Domains as the SID will change.

Password Password for the Existing User on the image if Username is populated.

Storage Provider Location where the Virtual Image will be stored. Default Virtual Image Storage location is */var/opt/morpheus/morpheus-ui/vms*. Additional Storage Providers can be configured in *Infrastructure -> Storage*.

Cloud-Init User Data Accepts what would go in `runcmd` and can assume bash syntax. Example use: Script to configure satellite registration at provision time.

Create Image Select FILE to select or drag and drop image file, or URL to download the image from an accessible URL. It is recommend to configure the rest of the settings below prior to uploading the source Image File(s).

Permissions Set Tenant permissions in a multi-tenant Morpheus environment. No impact on single-tenant environments.

Auto Join Domain? Enable to have instances provisioned with this image auto-join configured domains (Windows only, domain controller must be configure in *Infrastructure -> Network* and the configured domain set on the provisioned to Cloud or Network).

VirtIO Drivers Loaded? Enable if VirtIO Drivers are installed on the image for provisioning to KVM based Hypervisors.

VM Tools Installed? On by default, uncheck if VMware Tools (including OpenVMTools) are not installed on the Virtual Image. Morpheus will skip network wait during provisioning when deselected.

Force Guest Customization? VMware only, forces guest customizations to run during provisioning, typically when provisioning to a DHCP network where guest customizations would not run by default.

Trial Version Enable to automatically re-arm the expiration on Windows Trial Images during provisioning.

Enabled Sysprep? Applicable to Nutanix Only. Enable of the Windows Image has been sys-prepped. If enabled Morpheus will inject Unattend.xml through the Nutanix API (v3+ only)

3. Save Changes

Note: Cloud-Init is enabled by default on all Images. Images without Cloud-Init or Cloudbase-Init installed must have the *cloud-init* flag disabled on the Virtual Image setting or Provisioning may fail.

Provisioning Images

When provisioning a system image, Morpheus will stream the image from Amazon S3 to the target Cloud if the image is not local to the Cloud.

When using iamges that already exist in the destination Cloud, such as synced, marketplace, or previously copied images, no image stream from S3 through the Morpheus Appliance to the destination cloud will take place.

Note: The Morpheus Appliance must be able to download from Amazon S3 when provisioning system images.

Note: The Morpheus Appliance must be able reach and resolve the destination Host when provisioning System Images or uploaded Images for the first time. This included being able to resolve ESXi host names in VMware vCenter clouds, and reach the destination ESXi host over port 443.

Add Virtual Image

Virtual Images can be upload to Morpheus from local files or URL's. Amazon and Azure Marketplace metadata can also be added to the Virtual Images library, enabling the creation of custom catalog Instance Type from Marketplace images (no image is transferred to Morpheus when adding Marketplace images).

Warning: Be conscious of your Storage Provider selection. The default Storage Provider is the Morpheus Appliance at `/var/opt/morpheus/morpheus-ui/vms`. Uploading large images to the Morpheus Appliance when there is inadequate space will cause upload failures and impact Appliance functionality. Ensure there is adequate space on your selected Storage Provider. Additional Storage Provider can be added at *Infrastructure -> Storage*, which can be configured as the default Virtual Image Store or selected when uploading Images.

To Add Virtual Image:

1. Select + *Add* in the Virtual Images page.

2. Select Image format:

- Alibaba
- Amazon AMI
- Azure Marketplace
- Digital Ocean
- ISO
- PXE Boot
- QCOW2
- RAW
- VHD
- VMware (vmdk/ovf/ova)

3. Configure the following on the Virtual Image:

Name Name of the Virtual Image in Morpheus . This can be changed from the name of the Image, but editing will not change the name of the actual Image.

Operating System Specifies the Platform and OS of the image. All Windows images will need to have Operating System specified on the Virtual Image, as Morpheus will assign Linux as the Platform for all Images without Operating System specified.

Minimum Memory The Minimum Memory setting will filter available Service Plans options during provisioning. Service Plans that do not meet the Minimum Memory value set on the Virtual Image will not be provided as Service Plan choices.

Cloud Init Enabled? On by default, uncheck for any Image that does not have Cloud-Init or Cloudbase-Init installed.

Install Agent On by default, uncheck to skip Agent install. Note this will result in the loss of utilization statistics, logs, script execution, and monitoring. (Some utilization stats are collected for agent-less hosts and vm's from VMware and AWS clouds).

Username Existing Username on the Image. This is required for authentication, unless Morpheus is able to add user data, Cloud-Init, Cloudbase-Init or Guest Customizations. If Cloud-Init, Cloudbase-Init or Guest Customizations are used, credentials are defined in *Administration -> Provisioning* and *User Settings* .
If credentials are defined on the Image and Cloud-Init is enabled, \morpheus\ will add that user during

provisioning, so ensure that user does not already exist in the image (aka ``root``). For Windows Guest Customizations, Morpheus will set the Administrator password to what is defined on the image if Administrator user is defined. Do not define any other user than Administrator for Windows Images unless using Cloudbase-init. Morpheus recommends running Guest Customizations for all Windows Images, which is required when joining Domains as the SID will change.

Password Password for the Existing User on the image if Username is populated.

Storage Provider Location where the Virtual Image will be stored. Default Virtual Image Storage location is `/var/opt/morpheus/morpheus-ui/vms`. Additional Storage Providers can be configured in *Infrastructure -> Storage*.

Cloud-Init User Data Accepts what would go in `runcmd` and can assume bash syntax. Example use: Script to configure satellite registration at provision time.

Create Image Select FILE to select or drag and drop image file, or URL to download the image from an accessible URL. It is recommended to configure the rest of the settings below prior to uploading the source Image File(s).

Permissions Set Tenant permissions in a multi-tenant Morpheus environment. No impact on single-tenant environments.

Auto Join Domain? Enable to have instances provisioned with this image auto-join configured domains (Windows only, domain controller must be configured in *Infrastructure -> Network* and the configured domain set on the provisioned to Cloud or Network).

VirtIO Drivers Loaded? Enable if VirtIO Drivers are installed on the image for provisioning to KVM based Hypervisors.

VM Tools Installed? On by default, uncheck if VMware Tools (including OpenVMTools) are not installed on the Virtual Image. Morpheus will skip network wait during provisioning when deselected.

Force Guest Customization? VMware only, forces guest customizations to run during provisioning, typically when provisioning to a DHCP network where guest customizations would not run by default.

Trial Version Enable to automatically re-arm the expiration on Windows Trial Images during provisioning.

Enabled Sysprep? Applicable to Nutanix Only. Enable of the Windows Image has been sys-prepped. If enabled Morpheus will inject Unattend.xml through the Nutanix API (v3+ only)

Note: Default Storage location is `/var/opt/morpheus/morpheus-ui/vms`. Additional Storage Providers can be configured in *Infrastructure -> Storage*. Ensure local folders are owned by `morpheus-app.morpheus-app` if used.

Warning: Provisioning will fail if *Cloud init Enabled* is checked and Cloud-Init is not installed on the Image.

Note: Existing Image credentials are required for Linux Images that are not Cloud-Init enabled and for Windows Images when Guest Customizations are not used. Cloud-Init and Windows user settings need to be configured in *Administration -> Provisioning* when using Cloud-Init or Guest Customizations and new credentials are not set on the Virtual Image.

4. Upload Image

Images can be uploaded by File or URL:

File Drag and Drop the image file, or select *Add File* to select the image file.

Url Select the URL radio button, and enter URL of the Image.

Note: The Virtual Image configuration can be saved when using a URL and the upload will finish in the background. When selecting/drag and dropping a file, the image files must upload completely before saving the Virtual Image record or the Image will not be valid.

5. Save Changes.

Library

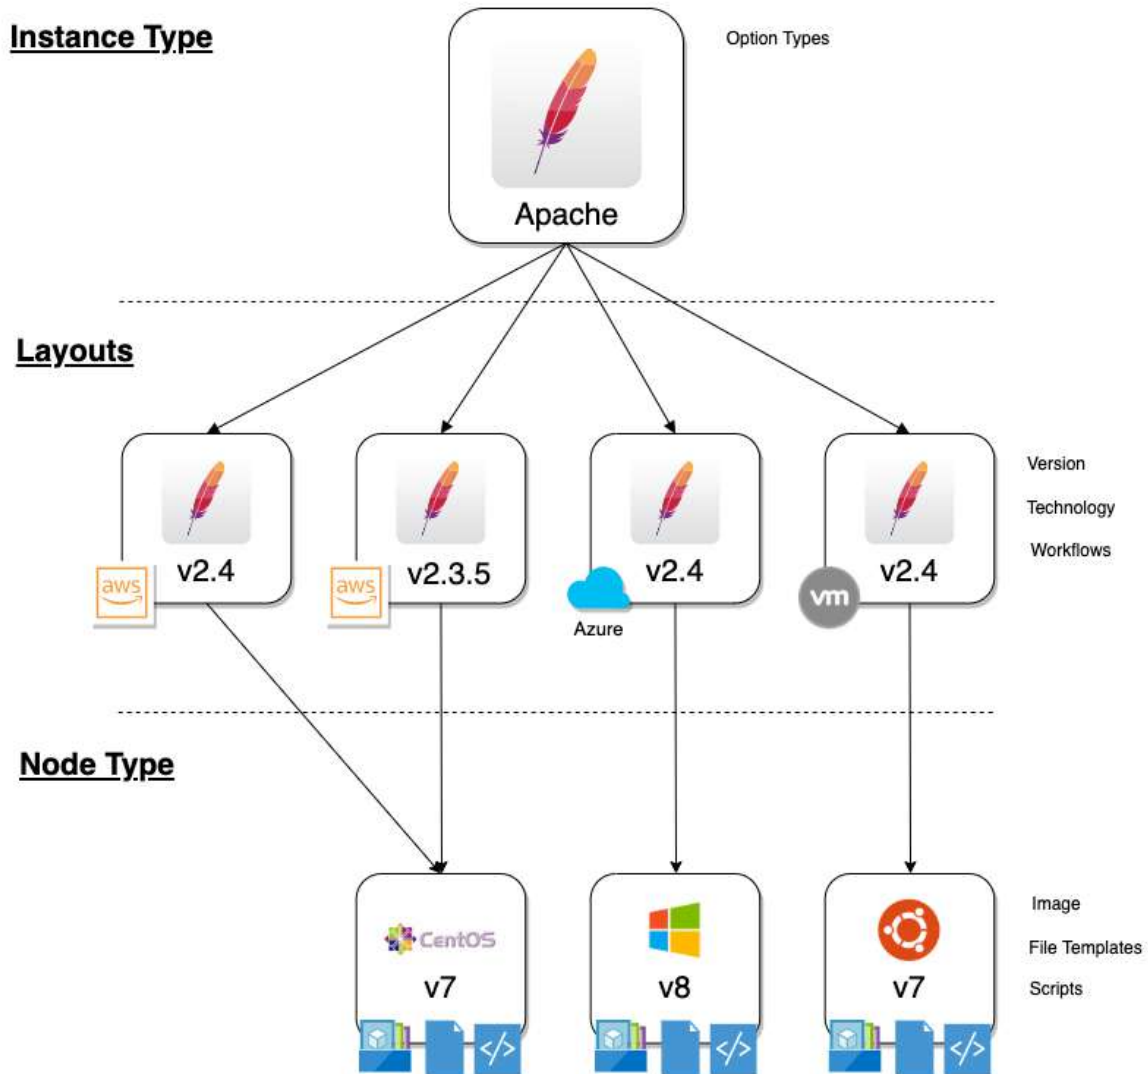
Overview

The Library section is used to add Virtual Images as custom Instance Types to the provisioning catalog. The Library section is composed of:

- Instance Types
- Layouts
- Node Types
- Option Types
- Option Lists
- File Templates
- Scripts
- Spec Templates
- Cluster Layouts

When provisioning, the User selects an INSTANCE TYPE from the provisioning wizard. At this stage, we can present custom OPTION TYPES to the User which alter deployment in ways the administrator predetermines. Based on the selected Cloud technology and Version, the User is presented with applicable LAYOUTS selections. LAYOUTS can take advantage of Workflows which automate Tasks and can utilize a wide range of DevOps automation technologies. One or more NODE TYPES is associated with the LAYOUT. NODE TYPES are the bridge between LAYOUTS and Images. NODE TYPES can also take advantage of File Templates for custom configuration and Scripts which can be queued to run at any stage of the Instance lifecycle.

Understanding Library Items



Instance Types

Adding an Instance Type creates a new Library item category. Multiple Layouts can be added to an Instance Type and these Layouts can have different Nodes attached. The Instance provisioning wizard will present the Layout options compatible with the selected Cloud. If Cloud selection is turned off, all Layouts will be presented for all Cloud types accessible by the User.

NEW LIBRARY ITEM

NAME: Salt Master

CODE: f5461b1a-f9f9-4198-80d5-4844597b
Useful shortcode for provisioning naming schemes and export reference.

DESCRIPTION: Salt Master Demo
239 Characters Remaining

CATEGORY: Utility

ICON: Choose File salt2_original.png
SALTSTACK
Suggested Dimensions: 150 x 51

Permissions

VISIBILITY: Private

Advanced

ENVIRONMENT PREFIX: SALT_MASTER
Used for exportable environment variables when tying instance types together in app contexts. If not specified a name will be generated.

ENVIRONMENT VARIABLES

Name	Value	OPTIONS
		+

☒ Enable Scaling (Horizontal)

☒ Supports Deployments
Requires a data volume be configured on each version. Files will be copied into this location.

SAVE CHANGES

Name Name of the Instance Type in the provisioning Library

Code A useful shortcode for provisioning naming schemes and export reference

Description The description of the Instance Type shown in the Provisioning Library. (255 characters max)

Category For filtering in Instance sections and provisioning wizard

- Web
- SQL
- NoSQL
- Apps
- Network
- Messaging
- Cache
- OS
- Cloud
- Utility

Icon An identifiable icon to display in-line with your Instance Type in the provisioning wizard (Suggested dimensions: 150 x 51)

Visibility

- **Private:** Only accessible by assigned Accounts/Tenants
- **Public:** Accessible by all Accounts/Tenants

Option Types Custom options presented to the user at provision time, Option Types are also created and stored in Morpheus Library

Environment Prefix Used for exportable environment variables when tying Instance Types together in App contexts. If not specified, a name will be generated

Environment Variables Name and value pairs for environment variables to be loaded on initialization

Enable Settings Allows for attachment of modifiable file templates to Node Types in a settings tab of the Instance after deployment

Enable Scaling (Horizontal) Enables load balancer assignment and auto-scaling features

Support Deployments Enables deployment features (Requires a data volume be configured on each version. Files will be copied into this location)

Upon saving, this Instance Type will be available in the provisioning catalog, per User Role access. However, we still need to add Layouts to the Instance Type, and prior to creating a Layout, we will add a Node Type.

Layouts

Layouts are attached to Instance Types. A Layout can only be attached to a single Instance Type and a single Technology. An Instance Type can have one or many Layouts attached to it, allowing for a single Instance Type to work with any Technology type. Node Types are added to Layouts. A Layout can have one or many node types attached to it. Node types can be shared across Layouts of matching Technology types.

Important: Once an Instance Type is defined on a Layout and saved, the Instance Type setting and Technology selections on the Layout cannot be changed.

Layout List View

The default page for Layouts is the Layout list view. Select + *ADD* to create a new Layout. Layouts can also be created from an Instance Type detail page.

The following fields are displayed for each Layout:

- **NAME:** Links to the Layout detail page
- **VERSION**
- **INSTANCE TYPE:** Links to the associated Instance Type
- **DESCRIPTION**

The Actions menu in each row reveals the following options:

- **Permissions:** Scope the Layout to Group(s) to narrow the list of available groups for a chosen Instance Type at provision time
- **Edit:** Edit the Layout

- **Delete:** Delete the Layout

Note: A Layout that is in use cannot be deleted.

Available Filters:

- **Technology:** Display Layouts by selected Cloud technology
- **Instance Type:** Display Layouts by the associated Instance Type

Layout Detail View

The Layout Detail view shows details on the Layout including Name, Short Name, Version, and Category. It also lists all associated Node Types.

- Select a Layout Name from the list page or Instance Type detail page to get to a Layout detail page.

Layout Configuration Options

Instance Type Select the Instance Type to add to the new Layout. Custom Instance Types must already be created and one Layout cannot be added to multiple instance types. The Instance Type also cannot be changed after creation.

Note: Layouts cannot be added to Morpheus pre-defined Instance Types

Name The name the Layout presents as in the Configuration Options dropdown of the provisioning wizard

Version The version number or name for the Layout. Layouts in an Instance Type with the same version will all show under the Configuration Options dropdown when that version is selected while provisioning

Description Description of the Layout, viewable on the Layout list view

Creatable When checked, this Layout will be selectable at provision time for the associated Instance Type (assuming the Layout is otherwise compatible with provisioning conditions). Instance Types with no Creatable Layouts will not be selectable from the provisioning wizard

Technology Technology determines which Cloud this layout will be available for and which Node Types can be added to it

Minimum Memory Defines the minimum amount of memory required for this Layout. Only Service Plans that meet the defined memory minimum will be available during provisioning when this Layout is selected. Custom memory values must also meet this minimum. Entering a minimum memory value of zero (the default value) indicates no minimum. This minimum memory value will override any Virtual Image minimum memory requirements

Workflow Select a Workflow to automatically run and be attached to associated Instances using this Layout. If a Workflow is defined, it is not presented in the provisioning wizard and is not user configurable

Supports Convert to Managed Enabled to allow users to select this layout when converting a Discovered workload to Managed

Enable Scaling (Horizontal) Enables Instances with this layout to use scaling features

Environment Variables Custom environment variables to be added to the Instance when provisioned

Option Types Search for and select one or multiple Option Types to add to the Layout. Option Type input fields (except for Hidden Option Types) will appear in Provisioning, App, Blueprint, and Cloning wizards when this layout is selected

Nodes Single or multiple nodes can be added to a Layout by searching for and selecting the Node(s)

Node Types

Node Types are the link between Images and Layouts.

Node Type Configuration Options

The following fields are for all Technology types:

Name Name of the Node Type in Morpheus

Short Name The short name is a lowercase name with no spaces used for display in your container list

Version Version for the Node Type. Examples: 7.5, 2012 R2, latest

Technology Select associated Technology. This will filter the available configuration options, images and which Layouts the Node Type can be added to

Environment Variables Add pre-set environment variables to the Node Type

Category Node Types of differing categories within the same Layout can have differing sizing

Technology-Specific Options

The Options fields will change depending on the Technology option selected. For VM provisioning technology options, select an image from the VM Image dropdown. This list is populated from the Morpheus Virtual Images section and will include images uploaded into Morpheus as well as synced images from added clouds.

Note: Amazon and Azure Marketplace Images can be added in the Virtual Images section for use as Node Types in custom library items.

Docker Options

For Docker, type in the name and version of the Docker Image, then select the integrated registry.

Service Ports To open ports on the node, enter the name and port to expose. The Load Balancer HTTP, HTTPS, or TCP setting is required when attaching to Load Balancers.

Defining an Exposed port will also create a hyperlink(s) on the container location (IP) in the VM or Container section of the associated Instance detail page.













Scripts Search for and select one or multiple scripts to be executed when the Node Type is provisioned

File Templates Search for and select one or multiple File Templates to be written when the Node Type is provisioned

Example port configuration:

Expose Ports

[Add Port](#)

NAME	PORT	LB	
HTTP	8000	HTTP 	
Collector	8088	None 	
Forwarder	9997	None 	
KVstore	8191	None 	
TCP	1514	None 	
Custom	1515	None 	

VMware Options

When VMware Technology Type is selected, EXTRA OPTIONS will be available in the VMware VM Options section. These allow defining Advance vmx-file parameters during provisioning.

Some Example include:

```
tools.setinfo.sizeLimit : 1048576
vmci0.unrestricted : FALSE
isolation.tools.diskWiper.disable : TRUE
```

Note: Not all parameters can be set using extra config parameters. A sample reference list can be found at <http://www.sanbarrow.com/vmx/vmx-advanced.html#vmx>

Important: Use caution when setting Extra Options. Malformed config files can break provisioning. Troubleshooting issues related to Extra Options defined are beyond the scope of Morpheus product support.

Option Types

Option Types are custom input fields that can be added to Instance Types and Layouts, then presented in Instance, App, and Cloning wizards. The resulting value is available in the Instance config map as `<%=customOptions.fieldName%>`. The `fieldName` and `value` can also be exported as Tags.

Create Option Type

Note: All possible fields listed. Displayed fields depend on `TYPE` selection

NAME Name of the Option Type

DESCRIPTION Description for reference in Option Type list view

FIELD NAME This is the input `fieldName` property that the value gets assigned to

EXPORT AS TAG Creates Tags for `fieldName/value` (key/value) on Instances

TYPE

- **Text:** Text Input Field
- **Select List:** Populated by Option Lists, presents a manual or REST-populated dropdown list
- **Checkbox:** Checkbox for `on` or `off` values
- **Number:** Input field allowing only numbers
- **Typeahead:** Populated by Option Lists: Rather than presenting a potentially-large dropdown menu, the user can begin typing a selection into a text field and choose the desired option. Multiple selections can be allowed with this type by marking the 'ALLOW MULTIPLE SELECTIONS' box

- **Hidden:** No field will be displayed, but the field name and default value will be added to the Instance config map for reference
- **Password:** An input field with suitable encryption for accepting passwords

LABEL This is the input label that typically shows to the left of a custom option

PLACEHOLDER Background text that populates inside a field for adding example values, does not set a value

DEFAULT VALUE Pre-populates field with a default value

REQUIRED Prevents User from proceeding without setting value

DEFAULT CHECKED For Checkbox types, when marked the Checkbox will be checked by default

OPTION LIST For Select List types, select a pre-existing Option List to set dropdown values

Note: Select List and Typeahead Option Types require creation and association of an Option List

Option Lists

Option Lists allow you to give the user more choices during provisioning to then be passed to scripts and/or automation. Option Lists, however, are pre-defined insofar as they are not free-form. They can be manually entered CSV or JSON, they can be dynamically compiled from REST calls via GET or POST requests, or populated by LDAP queries.

NEW OPTION LIST

NAME

DESCRIPTION

TYPE

REST

VISIBILITY

Private

If the account assigned is not the master tenant, visibility will automatically change to private.

Generic Option List Fields

The displayed fields in the create/edit Option List modal depend on the TYPE value selected.

NAME Name of the Option List

DESCRIPTION Description of the Option List for reference in Option List list view

TYPE

- **REST:** REST API call to populate Option List
- **Manual:** Manually entered dataset, CSV or JSON
- **Morpheus API:** Call to internal Morpheus API to populate the Option List
- **LDAP:** Searches and returns a list of Active Directory objects

VISIBILITY If the account currently signed in is not in the master tenant, visibility will automatically change to private

Manual Option List Fields

DATASET Appears only for manual Option Lists. Add your CSV or JSON list to this field

Note: JSON entries must be formatted like the following example: [{"name": "Test", "value": 1}, {"name": "Testing", "value": 2}]

REST Option List Fields

SOURCE URL A REST URL used to fetch list data which is cached in the appliance database

REAL TIME When checked, a REST call will be made to update the Option List at the time its presented to the User

IGNORE SSL ERRORS Do not fail API query for self-signed or invalid certs on REST call target

SOURCE METHOD GET or POST

SOURCE HEADERS Custom HTTP Headers to include in the source request

INITIAL DATASET Create an initial JSON or CSV dataset to be used as the collection for this option list. It should be a list containing objects with properties 'name' and 'value'

TRANSLATION SCRIPT Create a JS script to translate the result data object into an array containing objects with properties 'name' and 'value'. The input data is provided as 'data' and the result should be put on the global variable 'results'.

Example:

```
for(var x=0;x < data.length; x++) {
  results.push({name: data[x].title,value:data[x].id});
}
```

REQUEST SCRIPT Create a JS script to prepare the request. Return a data object as the body for a POST request, and return an array containing properties 'name' and 'value' for a GET request. The input data is provided as 'data' and the result should be put on the global variable 'results'

Example:

```
results.push({name: 'userId', value : data.users})
```

Morpheus API Option List Fields

OPTION LIST A list of available object types to return

TRANSLATION SCRIPT Create a JS script to translate the result data object into an array containing objects with properties 'name' and 'value'. The input data is provided as 'data' and the result should be put on the global variable 'results'.

Example:

```
var i=0;
results = [];
for(i; i<data.length; i++) {
  results.push({name: data[i].name, value: data[i].value});
}
```

Translation script inputs:

Clouds

- id: <Number>
- value: <Number> // id, convenience
- name: <String>
- displayName: <String>
- category: <String>
- description: <String>
- apiKey: <String>
- status: <String>
- hourlyPrice: <Number>
- hourlyCost: <Number>
- instanceType: <Object>
 - id: <Number>
 - name: <String>
- plan:<Object>
 - id: <Number>
 - name: <String>
- site:<Object>
 - id: <Number>
 - name: <String>

Environments

- id: <Number>
- value: <Number> // id, convenience attribute to avoid requiring translation

- code: <String>
- name: <String>

Groups

- id: <Number>
- value: <Number> // id, convenience attribute to avoid requiring translation
- name: <String>
- code: <String>
- uuid: <String>
- location: <String>
- datacenterId: <Number>

Instances

- id: <Number>
- value: <Number> // id, convenience
- name: <String>
- displayName: <String>
- category: <String>
- description: <String>
- apiKey: <String>
- status: <String>
- hourlyPrice: <Number>
- hourlyCost: <Number>
- instanceType: <Object>
 - id: <Number>
 - name: <String>
- plan: <Object>
 - id: <Number>
 - name: <String>
- site: <Object>
 - id: <Number>
 - name: <String>

Instances Wiki

- id: <Number>
- value: <Number> // id, convenience
- name: <String>
- urlName: <String>
- category: <String>

- instanceId: <String>
- content: <String>
- contentFormatted: <String>
- format: <String>
- createdByUsername: <String>
- updatedByUsername: <String>

Networks

- id: <Number>
- value: <Number> // id, convenience
- code: <String>
- category: <String>
- name: <String>
- status: <String>
- cloudId: <Number>
- groupId: <Number>
- networkType:<Object>
 - id: <Number>
 - code: <String>
 - name: <String>
- externalId: <String>
- externalNetworkType: <String>
- networkDomain: <Object>
 - id: <Number>
 - name: <String>
- networkPool: <Object>
 - id: <Number>
 - name: <String>
- createdBy: <String>

Plans

- id: <Number>
- value: <Number> // id, convenience
- code: <String>
- name: <String>
- storage: <Integer, bytes>
- memory: <Integer, bytes>
- cores: <Number>

Resource Pools

- id: <Number>
- value: <Number> // id, convenience
- code: <String>
- externalId: <String>
- name: <String>
- serverGroupId: <Number>
- status: <String>
- regionCode: <String>
- parentPoolId: <Number>
- type: <String>

Security Groups

- id: <Number>
- value: <Number> // id, convenience
- code: <String>
- name: <String>
- externalType: <String>
- externalId: <String>
- cloudId: <Number>
- scopeMode: <String>
- scopeId: <Number>

Servers

- id: <Number>
- value: <Number> // id, convenience
- name: <String>
- displayName: <String>
- description: <String>
- category: <String>
- osType: <String>
- powerState: <String>
- lastStats: <String>
- zone: <Object>
 - id: <Number>
 - name: <String>
- capacityInfo: <Object>
 - maxStorage: <Integer, bytes>

- maxMemory: <Integer, bytes>
- maxCores: <Number>
- usedMemory: <Integer, bytes>
- usedStorage: <Integer, bytes>
- computeServerType: <Object>
- id: <Number>
- name: <String>
- nodeType: <String>
- vmHypervisor: <String>
- containerHypervisor: <String>

Servers Wiki

- id: <Number>
- value: <Number> // id, convenience
- name: <String>
- urlName: <String>
- category: <String>
- serverId: <String>
- content: <String>
- contentFormatted: <String>
- format: <String>
- createdByUsername: <String>
- updatedByUsername: <String>

REQUEST SCRIPT The request script is used differently for Morpheus API Option List types. A Morpheus API option list type will use an internal API to return a list of objects instead of performing HTTP(S) requests to the Morpheus API. Due to this approach, the results object will not be used to generate query parameters or a JSON body. The results object will instead be used to contain a map of accepted key:value pairs that can be used to filter, sort and order the list of objects that get returned.

Below is a list of accepted key:value pairs for each object type: Generic options available for all object types

- max: <integer> // Maximum number of results to return. Default: 25
- offset: <integer> // Offset for returned results. Default: 0
- sort: <string> // Field to sort on. Default: 'name'
- order: <string> // Order of returned values. Accepted values: 'asc', 'desc'. Default: 'asc'

Example: results = {max: 5, order : 'desc'}

Networks

- zoneId

- `siteId`
- `planId`
- `provisionTypeId`: `<Number>` // Id of the provision type (technology), filters to only networks associated with this provision type
- `layoutId`: `<Number>` // Id of an Instance Layout, ignored if `provisionTypeId` is supplied, otherwise used to look up the provision type
- `poolId`: `<Number>` // Id of a network pool, filters to only networks within the specified network pool

Plans

- `zoneId`
- `siteId`
- `layoutId`
- `provisionTypeId`: `<Number>` // Id of the provision type (technology), filters to only plans associated with this provision type

Resource Pools

- `zoneId`
- `siteId`
- `planId`
- `layoutId`: `<Number>` // Id of an Instance Layout, used to get the associated provision type and filter to that provision type

Security Groups

- `zoneId` // required
- `poolId`

Clouds

- `zoneId` : `<integer>` // Database ID of cloud to return
- `tenantId` : `<integer>` // Database ID of tenant where clouds are added. Filters to only clouds added within the specified tenant. Only available in Master Tenant
- `zoneTypeId` : `<integer>` // Database ID of cloud type. Filters to only clouds with the specified cloud type
- `siteId` : `<integer>` // Database ID of group. Filters to only clouds within the specified group
- `tagName` : `<string>` // Filters to clouds with servers with tags containing the tagName
- `tagValue` : `<mixed>` // Requires tagName. Filters to clouds with servers that have tags containing the tagName and specified tagValue
- `phrase` : `<string>` // Fuzzy matches phrase on cloud name and description

Example: `results = {tenantId: 1, siteId: 1, tagName: "morpheus"}`

Instances

- `appsId` : `<integer>` // Database ID of app to filter by. Returns instances linked to the app
- `tenantId` : `<integer>` // Database ID of tenant where instances are located. Filters to only instances within the specified tenant. Only available in Master Tenant
- `serverId` : `<integer>` // Database ID of server. Filters to the instance that contains the specified server
- `tagName` : `<string>` // Filters to instances with tags containing the tagName
- `tagValue` : `<mixed>` // Requires tagName. Filters to instances with tags containing the tagName and specified tagValue
- `phrase` : `<string>` // Fuzzy matches phrase on instance name and description

Example: `results = {tenantId:1, phrase: "ha"}`

Groups

- `tenantId` : `<integer>` // Database ID of tenant where groups are located. Filters to only groups added within the specified tenant. Only available in Master Tenant
- `zoneTypeId` : `<integer>` Database ID of cloud type. Filters to only groups that contain clouds with the specified cloud type
- `zoneId` : `<integer>` // Database ID of cloud. Filters to only groups that contain the cloud with the specified ID
- `siteId` : `<integer>` // Database ID of group to return
- `phrase` : `<string>` // Fuzzy matches phrase on group name and location.

Servers

- `tenantId` : `<integer>` // Database ID of tenant where servers are located. Filters to only servers within the specified tenant. Only available in Master Tenant
- `serverId` : `<integer>` // Database ID of server. Filters to the server specified by the ID
- `siteZoneId` : `<integer>` // Database ID of cloud. Filters to servers contained within the specified cloud
- `serverType` : `<string>` // Type of server. Accepted values: 'host', 'baremetal', 'vm'
- `siteId` : `<integer>` // Database ID of group. Filters to only servers contained within clouds that are added in the specified group
- `tagName` : `<string>` // Filters to servers with tags containing the tagName
- `tagValue` : `<mixed>` // Requires tagName. Filters to servers with tags containing the tagName and specified tagValue
- `phrase` : `<string>` // Fuzzy matches phrase on server name and description.

Example: `results = {max: 50, siteZoneId : 3}`

instance-wiki: Contains same options for Instances Morpheus API type.

- `phrase` : `<string>` // Fuzzy matches phrase on wiki name, urlName and content

server-wiki: Contains same options for Servers Morpheus API type.

- `phrase` : `<string>` // Fuzzy matches phrase on wiki name, urlName and content

LDAP Option List Fields

LDAP URL The URL pointing to the LDAP server

USERNAME The fully qualified username (with @ suffix syntax) for the binding account

PASSWORD The password for the above account

LDAP Query The LDAP query to pull the appropriate objects. See the next section for an example use case

TRANSLATION SCRIPT Create a JS script to translate the result data object into an array containing objects with properties 'name' and 'value'. The input data is provided as 'data' and the result should be put on the global variable 'results'.

Note: Option Lists are set on one or multiple `Select List` or `Typeahead` Option Types. The Option Type is then set on an Instance Type, Layout, Cluster Layout, and/or Operational Workflow for input during provisioning or execution.

Creating an Option List Based on an LDAP Query

In Morpheus version 4.2.1 and higher, Option Lists can be populated from LDAP queries. This gives users the ability to search Active Directory, capture objects, and present them as custom options where needed.

It's recommended that you connect LDAP-type Option Lists to Typeahead-type Option Types as the list of returned selections can be very large. This also allows you to select multiple options from the list, presuming you've allowed for that when creating the Option Type.

Populating LDAP-type Option Lists requires knowledge of LDAP query syntax. This guide provides one example and there are many publicly-available resources for help writing additional queries.

1. Create a new Option List (Provisioning > Library > Option Lists > ADD)
2. Enter a name for the new LDAP Option List
3. Change the Type value to LDAP and the relevant fields will appear as shown in the screenshot:
4. Enter the LDAP URL in the following format (an example is also shown as a placeholder in the UI form field):

```
ldap[s]://<hostname>:<port>/<base_dn>
```

5. Enter the fully qualified username with @ suffix syntax, such as: *user@ad.mycompany.com*
6. Enter the account password
7. Enter your LDAP query. You can even inject variables into your query structure to query based on the value the user has entered into the typeahead field as shown in the example below:

```
(&(objectClass=user)(cn=<%=phrase%>*))
```

8. Finally, enter a translation script which will convert the returned LDAP object into a list of name:value pairs you can work with in Morpheus. The example script below shows the user DisplayName and sets the value to the SAMAccountName:

```
for(var x=0;x < data.length ; x++) {  
  
  var row = data[x];
```

(continues on next page)

(continued from previous page)

```
var a = {};  
  
if(row.displayName != null) {  
    a['name'] = row.displayName;  
  
} else {  
  
    a['name'] = row.sAMAccountName;  
  
}  
  
a['value'] = row.sAMAccountName;  
results.push;  
  
}
```

9. Click SAVE CHANGES

NEW OPTION LIST [X]

NAME

DESCRIPTION

TYPE

VISIBILITY

If the account assigned is not the master tenant, visibility will automatically change to private.

▼ Request Options

LDAP URL

USERNAME

PASSWORD

LDAP QUERY

LDAP Queries are standard LDAP formatted queries where different objects can be searched. Dependent parameters can be loaded into the query using the `<%=phrase%>` syntax.

TRANSLATION SCRIPT Create a js script to translate the result data object into an Array containing objects with properties `name`, and `value`. The input data is provided as `data` and the result should be put on the global variable `results`.

```
for(var x=0;x < data.length ; x++) {  
    var row = data[x];
```

SAVE CHANGES

File Templates

File Templates are for generating config files, such as my.cnf, elasticsearch.yml, morpheus.rb, or any text file. With full config map variable support, Template Files are dynamically generated during a Workflow phase or ad hoc via Instance actions.

File Templates can also be exposed on Instances in the Settings Tab. Ensure the Instance Type supports settings, and Category is defined in Advance Options on the Library Template config.

Note: Morpheus variables are supported in Library Templates using `<%= variable.var %>` format

Examples:

HA Proxy Config (haproxy.cfg)

- FILE NAME: haproxy.cfg
- FILE PATH: /config/haproxy.cfg
- PHASE: Pre Provision
- TEMPLATE:
- SETTING NAME: haproxyConfig
- SETTING CATEGORY: haproxy

```
#!/bin/bash

global
    maxconn 256
    log /dev/log local0 warning
    log-tag <%=logTag%>

defaults
    mode http
    timeout connect 5000ms
    timeout client 50000ms
    timeout server 50000ms
    log global

frontend http-in
    bind *:<%=container.externalPort%>
    default_backend servers

backend servers
    # server server1 127.0.0.1:80 maxconn 32
```

mysql config (mysqld.cnf)

- FILE NAME: mysqld.cnf
- FILE PATH: /config/mysqld.cnf
- PHASE: Pre Provision

```
#!/bin/bash

[mysqld]
pid-file= /var/run/mysqld/mysqld.pid
```

(continues on next page)

(continued from previous page)

```
socket= /var/run/mysqld/mysqld.sock
datadir= /var/lib/mysql
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0
explicit_defaults_for_timestamp = 1
```

Scripts

Scripts are bash and Powershell scripts that can be attached to Node Types to always execute at the selected phase when that Node Type is provisioned, added to Workflows as Library Script Tasks, and/or executed ad-hoc on Instances.

Creating Scripts

1. Navigate to Provisioning > Library > Scripts
2. Select + *ADD*
3. Enter the Following:

NAME Name of the Script in Morpheus

SCRIPT TYPE

- Bash
- Powershell

PHASE Select which phase the Script will execute when attached to a Node Type. When a script is attached to a Node Type, it will execute according to the selected phase:

Start Service Any time the Instance action *Start Service* is executed

Stop Service Any time the Instance action *Stop Service* is executed

Pre-Provision

- Containers: Script will execute against the container host before the container is provisioned
- Virtual Machines: Script will execute before any provision phase Scripts or Tasks

Provision

Script will execute once per new Instance node during the provision Phase. Provisioning will not be considered complete until all scripts and tasks in the provisioning phase are completed

Note: Any Script or Task set to the provision phase will be included in the total provision time and impact success/warn/failure provisioning status messages. As an example, your VM could be up and running but if your Script is in the provision phase and fails, provisioning will be marked as a failure.

Post-Provision Script will execute once per new Instance node after the provision phase is completed. Scripts and Tasks in the Post-Provision phase will show execution status and history, but are not considered part of the provision and do not impact provisioning status.

Pre-Deploy Script will execute on target Instance any time a deployment is run against the Instance. The Script will run prior to the deployment file(s) being written

Deploy Script will execute on target Instance any time a deployment is run against the Instance. The script will run after the deployment file(s) are written

Reconfigure Script will execute on target Instance any time a reconfigure is executed against the Instance.

Teardown Script will execute on target Instance upon Instance deletion. Script will execute against target Instance prior to the deletion/removal of resources.

SCRIPT Enter Bash or Powershell script.

Note: Morpheus variables are supported in Library Scripts using `<%= variable.var %>` format

RUN AS USER By default Scripts are execute as `morpheus-node`. To execute as another User, populate `RUN AS USER` and ensure proper user permissions & group access

SUDO Flag `SUDO` if `sudo` is required to execute the Script

To attach Scripts and templates that have been added to the Library to a Node Type, start typing the name and then select the script(s) and/or template(s).

- Multiple scripts and templates can be added to a Node Type
- Scripts and Templates can be added/shared among multiple Node Types
- The execution phase can be set for Scripts in the Scripts section
- Search will populate Scripts or Templates containing the characters entered anywhere in their name, not just the first letter(s) of the name

Spec Templates

Spec Templates allow Morpheus users to leverage several major Infrastructure-as-Code solutions. These are typically JSON or YAML-based configuration files which make creating and managing multiple resource types easier. Morpheus allows users to create and/or manage a collection of these templates for different solutions and from different sources.

Morpheus currently supports Spec Templates of the following types:

- Kubernetes Spec
- Helm Chart
- Terraform
- ARM Template
- CloudFormation Template
- OneView Server Profile Template
- UCS Service Profile Template

Morpheus also allows users to leverage templates pulled from URL sources, online repositories (such as GitHub), or you can write a template locally inside the “NEW SPEC TEMPLATE” modal.

Note: To see Morpheus Spec Templates in action, take a look at an example use case in our [KnowledgeBase](#) where a CloudFormation Spec Template is used to create a provisionable custom Instance Type in Morpheus.

Creating a Spec Template

1. Navigate to *Provisioning > Library > SPEC TEMPLATES*
2. Click + *ADD*
3. Complete the following fields, then click *SAVE CHANGES*:
 - **NAME**
 - **TYPE:** See the previous section for a complete list of Spec Template types
 - **SOURCE:** Local, Repository, or URL
 - **CONTENT:** If this is a local Spec Template, supply the template in this field. If the template is supplied through a URL or online repository, the CONTENT field will change to allow the user to point Morpheus to that resource.

Cluster Layouts

Morpheus provided Cluster Layouts:

NAME	VERSION	DESCRIPTION
Amazon Docker Host	16.04	This will provision a single docker host vm in amazon
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in amazon
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in amazon
Digital Ocean Docker Host	16.04	This will provision a single docker host vm in digitalOcean
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in digitalOcean
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in digitalOcean
Kubernetes 1.13 EKS	1.13	This will provision a single kubernetes master in amazon
Kubernetes 1.14 EKS	1.14	This will provision a single kubernetes master in amazon
ESXi Docker Host	16.04	This will provision a single docker host vm in esxi
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in esxi with
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in esxi with
Docker on Ubuntu 16.04	16.04	This will provision a single docker host vm in fusion
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in fusion
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in fusion
KVM on CentOS 7.5	7.5	This will provision a single kvm host vm in fusion
Google Docker Host	16.04	This will provision a single docker host vm in google
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in google
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in google
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in huawei
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in huawei
Openstack Docker Host	16.04	This will provision a single docker host vm in huawei
Hyper-V Docker Host	16.04	This will provision a single docker host vm in hyperv
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in hyperv
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in hyperv
IBM Docker Host	16.04	This will provision a single docker host vm in bluemix
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in bluemix
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in bluemix
Docker on Bare Metal Ubuntu 16.04	16.04	This will provision a single docker host
Docker on Ubuntu 16.04	16.04	This will provision a single docker host
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master with weave

Table 8 – continued from previous page

Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a kubernetes cluster with weave and
Manual Docker on Linux	1	This will add a single docker host
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in nutanix
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in nutanix
Nutanix Docker Host	16.04	This will provision a single docker host vm in nutanix
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in openstack
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in openstack
Openstack Docker Host	16.04	This will provision a single docker host vm in openstack
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in openstack
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in openstack
Openstack Docker Host	16.04	This will provision a single docker host vm in openstack
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in softlayer
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in softlayer
Softlayer Docker Host	16.04	This will provision a single docker host vm in softlayer
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in vcd with
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in vcd with
VCD Docker Host	16.04	This will provision a single docker host vm in vcd
Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in vmware
Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in vmware
KVM on CentOS 7.5	7.5	This will provision a single kvm host vm in vmware

Users can add new cluster layouts using the +ADD button. Morpheus-provided cluster layouts can be cloned for use in creating custom layouts. Custom cluster layouts can also be deleted or edited from the list view using the pencil or trash can icons.

LIBRARY

INSTANCES TYPES
LAYOUTS
NODE TYPES
OPTION TYPES
OPTION LISTS
FILE TEMPLATES
SCRIPTS
SPEC TEMPLATES
CLUSTER LAYOUTS

Search
SELECT: TECHNOLOGY
+ ADD

NAME	VERSION	DESCRIPTION	
Alibaba Custom Centos 7	Centos 7	do not modify	
Alibaba Docker Host	16.04	This will provision a single docker host vm in alibaba	
Kubernetes 1.14 Cluster on Ubuntu 16.04, Weave, OpenEBS	16.04	This will provision a single kubernetes master in alibaba with weave and openebs	
Kubernetes 1.14 HA Cluster on Ubuntu 16.04, Weave,	16.04	This will provision a single kubernetes master in alibaba with weave and openebs	

Deployments

The deployments section provides very useful PaaS like capabilities when it comes to deploying applications into the newly provisioned environment. These can be uploaded directly from the UI, pulled from a build server, pulled from a public or private Git repository or even via the API and the various plugins created, such as Jenkins, and Gradle to support continuous build / integration workflows.

A deployment can be considered a set of versions that relate to a particular project or application being deployed. This allows one to keep track of a history of versions and easily reuse these deployment versions across instances that may exist in different environments. An example might be to deploy a version from a deployment to a staging instance and (once approved) also deployed into production.

Getting Started

Getting started with deployments is easy. They can vary slightly for the application stack being deployed but the simplest phase of a deployment is adding a version and adding the appropriate files to the deployment archive that are needed for the application to run. This could be a single file like a *WAR* file for Tomcat, or it could be hundreds of files for stacks like *Ruby on Rails*.

There are a few ways to create a deployment. The first is to use the `Provisioning -> Deployments` section of the application to create them. Simply add a new deployment and give it a name representing the application that is being deployed. Once a deployment is created select the deployment to view its versions (which will be empty to start). Next, its time to add a version.

When adding a version there are several options. There are 3 types represented by the UI. These include File, Fetch, and Git respectively. A File deployment allows the user to simply drag their files into the file explorer presented by the dialog. This file explorer can take single files or entire file trees (If files exist in subfolders then only the Chrome browser is supported due to browser limitations at the time of this writing). This is also the common type that is represented when files are uploaded via the CLI, or available build tool integration plugins. Once the files have completed their upload simply save the version for use.

Git

For performing git based deploys Morpheus supports both public and private repositories. To utilize a private git repository the add version dialog will display a public keypair that can be added to the git service for authentication purposes. Currently this keypair is shared across the account and not specifically scoped to the user so it may be advisable to connect this integration to a deployment account in git. From here either a *ssh* or *https* git url can be entered along with a git branch or tag name. Once the version is saved, this repository will be copied down into the deployment archive for use.

Fetch

Fetch based deployments are pretty straightforward. Simply enter a url to a file representing the deployment. This can be a single file (in which case it will just be added to the deployment archive singularly) or it can be a zip file (which will automatically be expanded into the archive). HTTP Authentication options can also be entered if the url requires some form of basic authentication scheme for access by the appliance.

Deploying to an Instance

Now that a version has been added to a deployment it is easy to push that deploy out to any instance provisioned within Morpheus . Simply navigate to the specific Instance that needs deployed to. On the Instance detail page there is a tab called *Deploy*. From here simply add a deploy. The dialog will ask firstly from which deployment the deploy is from (or allow you to create a new one on the spot) , and secondly which version to deploy (also with the option to add one on the fly). The next step of the wizard will display any configuration options that might be specific to the instance type being deployed to (i.e. *CATALINA_OPTS* for Tomcat or *Java Command* for java) as well as the file explorer and deployment type selections for review (or use when creating a new version on the fly). Fill in the required items then simply hit complete. The deploy will now be asynchronously sent off to all of the virtual machines or containers within the instance in a rolling restart and the deployment status will be represented.

Tip: When deploying to an instance, the custom configuration options that were entered during the previous deployment are automatically carried forward allowing one to edit them or leave them as is.

Rolling Backwards and Forwards

Because of the tracked history of deployments kept within Morpheus , the deploy tab of instance detail makes it easy to choose a previously run deployment and jump back to it in the event of a failed deployment. The history will automatically be updated and the configuration, as well as data from the previous deployment state of the instance will be restored.

Offloading Storage

Since a full history of the backup builds are kept in Morpheus , as the appliance grows it becomes necessary to change where these are stored. On a fresh install these are stored on the local appliance in `/var/opt/morpheus` or wherever the master account may have changed the configuration to point to. It is also possible to adjust the deployment archive store by creating a *Storage Provider* tied to an S3 compatible object store, Openstack Swift object store, or any other type of mountpoint provided. This option can be adjusted in Admin -> Provisioning once a storage provider is created within the account.

Add Deployment

1. Select the `Provisioning` link in the navigation bar.
2. Select the `Deployments` link in the sub-navigation bar.
3. Click the *Add* button.
4. Enter a `Name` for the deployment and a description (optional)
5. Click the *Save Changes* button to save.

Add Version

1. Select the `Provisioning` link in the navigation bar.
2. Select the `Deployments` link in the sub-navigation bar.
3. Click the `Name` of the deployment you would like to add a version to.
4. Click the *Add Version* button.
5. From the Add Version Wizard select the deployment type.
6. Input the `Version` of the deployment.
7. Depending on the type of deployment selected perform one of the following:

Files Drag files into the file explorer presented by the dialog. This file explorer can take single files or entire file trees.

Fetch Enter a url to a file representing the deployment.

Git The add version dialog will display a public key pair that can be added to the git service for authentication purposes. Either a ssh or https git url can be entered along with a git branch or tag name.

8. Click the *Save Changes* button to save.

Edit Deployment

1. Select the `Provisioning` link in the navigation bar.
2. Select the `Deployments` link in the sub-navigation bar.
3. Click the Edit Deployment icon on the row of the deployment you wish to edit.
4. Modify information as needed
5. Click the *Save Changes* button to save.

Delete Deployment

1. Select the `Provisioning` link in the navigation bar.
2. Select the `Deployments` link in the sub-navigation bar.
3. Click the Delete Deployment icon on the row of the deployment you wish to delete.

Service Mesh

A service mesh is an infrastructure layer which handles network communication between application services. A service mesh ensures fast and reliable communication between containerized application services. Morpheus will consume and display information on service meshes.

To view this information, navigate to `Provisioning > Service Mesh`.

The SERVICES tab displays the following fields:

- **NAME**
- **CLOUD**
- **SERVICE ID**
- **LOCATION**

The DNS tab displays the following fields:

- **NAME**
- **DOMAIN**
- **TYPE**
- **TTL**

1.3.3 Infrastructure

The heart of Morpheus is the ability to manage provisioning across any infrastructure, from bare metal to virtualized clouds and all the way to public infrastructure.

Groups

Overview

Groups in Morpheus define what resources a user has access to. Group access is defined by User Roles. Clouds are added to groups, and a User can only access the Clouds that are in the Groups their Role(s) gives them access to. Resources such as Networks, Datastores, Resources Pools, and Folders have additional Group access settings.

Policies applied to a Group will be enforced on all Instances provisioned or moved into that Group.

Note: Groups are not multi-tenant. A group only exists in the tenant it is created in.

The Groups view displays all current groups, includes search feature, and also enables the addition of new groups.

To View Groups:

1. Select the Infrastructure link in the navigation bar
2. Click the Groups link

UI

1. Select the Infrastructure link in the navigation bar
2. Click the Groups link

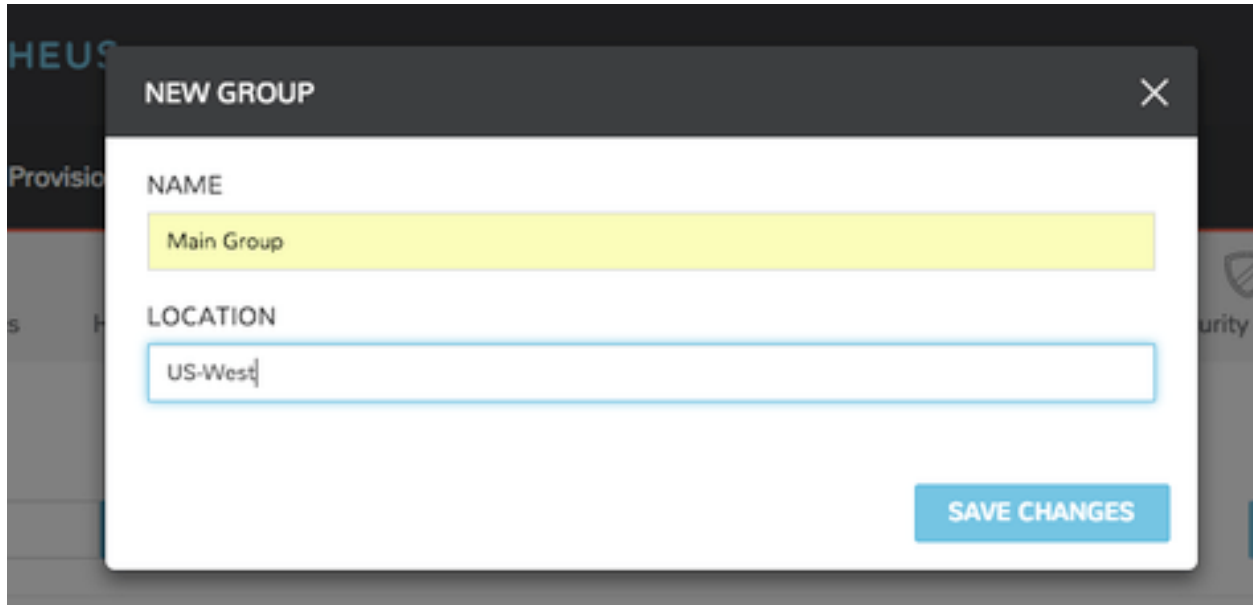
CLI

View all groups: `groups list` To use the group: `groups use <id>` or `groups use "group name"` Json output of a specific group: `groups get <id> -j` or `groups get "group name" -j`

API

View all groups: `curl https://api.gomorpheus.com/api/groups -H "Authorization: BEARER access_token"` View a specific group: `curl https://api.gomorpheus.com/api/groups/:id -H "Authorization: BEARER access_token"`

Adding Groups



The screenshot shows a 'NEW GROUP' dialog box. The title bar is dark gray with the text 'NEW GROUP' and a close button (X). The dialog body is white. It has two input fields: 'NAME' with the value 'Main Group' and 'LOCATION' with the value 'US-West'. A blue 'SAVE CHANGES' button is at the bottom right.

To add a group:

1. Select the Infrastructure link in the navigation bar
2. Click the Groups link
3. Click the Create Group button
4. Input out the Name and Location (optional) fields
5. Click the Save Changes button to save

Managing Groups

The screenshot displays the Morpheus web interface for managing groups. The top navigation bar includes links for Dashboard, Provisioning, Infrastructure (highlighted), Backups, Logs, Monitoring, Reports, and Admin. Below this, a secondary navigation bar shows links for Groups, Clouds, Hosts, Load Balancers, Storage, Key Pairs, Certificates, Boot, and Security. The main content area shows the 'Main Group' configuration page. It displays 'Clouds: 0', 'Servers: 0', and 'Location: US-West'. There are 'EDIT' and 'DELETE' buttons. Below this, there are tabs for Hosts, Virtual Machines, Bare Metal, and Clouds (selected). A message states 'You have no Clouds' and 'Click an Add button below to begin.' Below this, there are two columns: 'PUBLIC CLOUD' and 'PRIVATE CLOUD'. The 'PUBLIC CLOUD' column lists Azure, DigitalOcean, and Google Cloud Platform Live, each with a '+ ADD' button. The 'PRIVATE CLOUD' column lists Azure Stack, Hyper-V, and MORPHEUS, each with a '+ ADD' button.

To view a Group:

1. Select the Infrastructure link in the navigation bar
2. Click the Groups link
3. Click the Group name to view/modify

Available tabs in group view

Hosts Lists available hosts in the group and displays power, os, name, type, cloud, ip address, nodes, disc space, memory, and status. You can add a host from this tab panel by clicking Add Host.

Virtual Machines List all Virtual Machines in the Group.

Bare Metal List all Bare Metal Hosts added to the Group

Clouds Lists Clouds added to the Group. Existing Clouds or new Clouds can be added from the Group by clicking Add Cloud.

Policies Lists and allows creation or management of Policies applied to the Group.

Edit Group

To edit a group:

1. Select the Infrastructure link in the navigation bar.
2. Click the Groups link.
3. Click the name of the group you wish to edit.
4. Click the Edit button.
5. From the Edit Group Wizard modify information as needed.
6. Click the Save Changes button to save.

Delete Group

To delete a group:

1. Select the Infrastructure link in the navigation bar.
2. Click the Groups link.
3. Click the name of the group you wish to delete.
4. Click the Delete button.
5. Confirm

User Access

Important: User access to Groups is determined by their user Role(s). Group access for Roles can be configured in the Group Access section of a Roles Settings.

Clouds

Overview

Clouds are integrations or connections to public, private, hybrid clouds, or bare metal servers. Clouds can belong to many groups and contain many hosts. The clouds view includes clouds status, statistics, tenant assignment, and provides the option to add, edit, delete new clouds. Morpheus supports most Public Clouds and Private Clouds.

Supported Cloud Types

- Alibaba Cloud
- Amazon
- Azure (Public)
- Azure Stack (Private)
- Cloud Foundry
- Dell (Cloud type for PXE and manually added Dell EMC Hosts)

- DigitalOcean
- Google Cloud
- HPE (Cloud type for PXE and manually added HPE Hosts)
- HPE OneView
- Huawei
- Hyper-V
- IBM Cloud
- IBM Cloud Platform
- Kubernetes
- MacStadium
- Morpheus (Generic Cloud type for PXE/Bare Metal and manually added Hosts)
- Nutanix
- Open Telekom Cloud
- OpenStack
- Oracle Public Cloud
- Oracle VM
- Platform 9
- SCVMM
- SoftLayer
- Supermicro (Cloud type for PXE and manually added Supermicro Hosts)
- UCS
- UpCloud
- VMWare ESXi
- VMware Fusion
- VMWare on AWS
- VMware vCenter
- VMware vCloud Air
- VMware vCloud Director
- XenServer

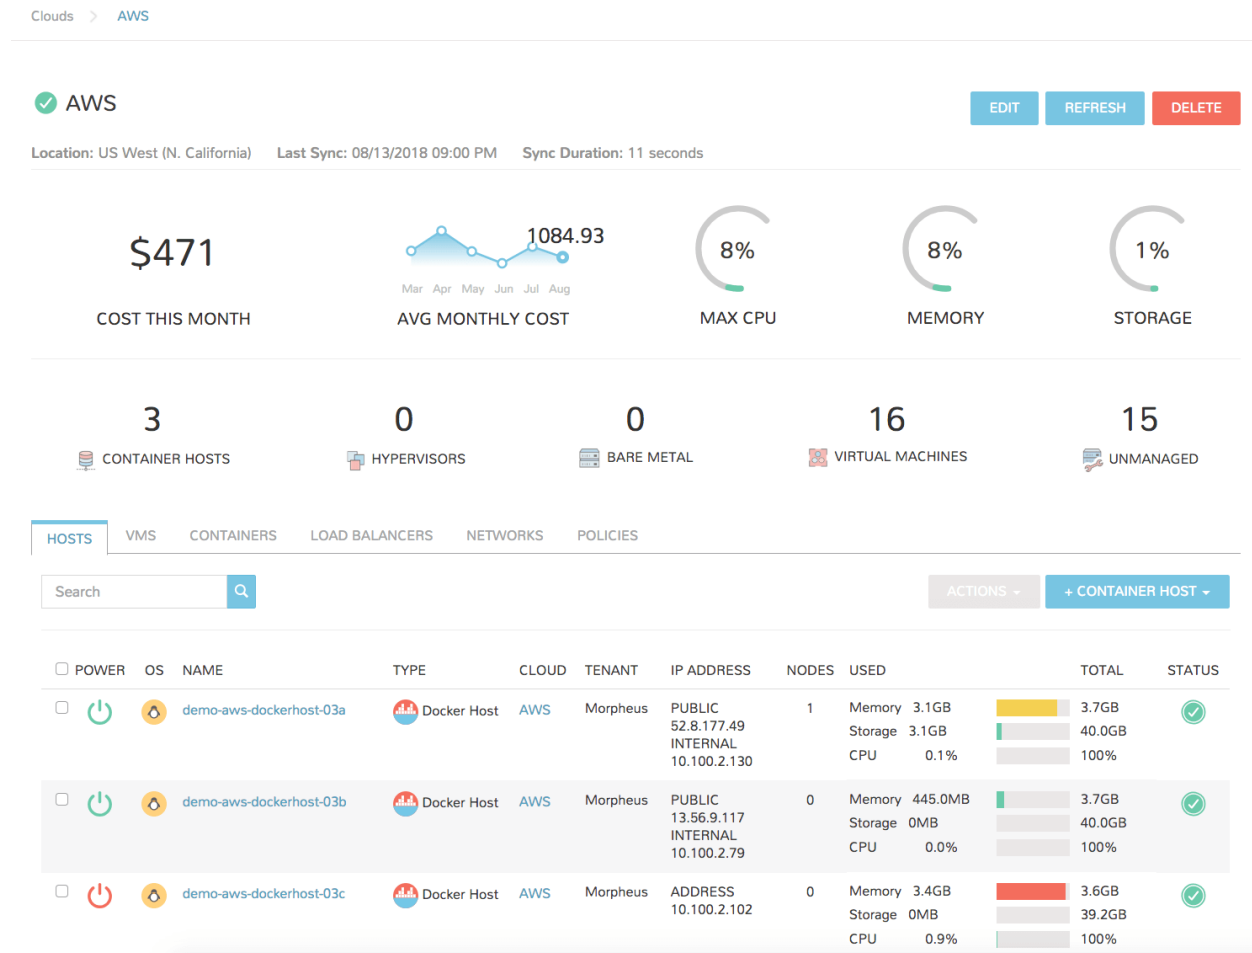
Information on each cloud type can be found in the [Guides](#) section.

Creating Clouds

Clouds can be added from *Infrastructure > Clouds* or in *Infrastructure -> Groups -> (select Group) -> Clouds*. Individual Guides for adding specific Cloud Types can be found in the [Guides](#) section.

Cloud Detail View

The Cloud Detail view shows metrics on health, sync status, current month costs, average monthly costs, resource utilization statistics, and resource counts for Container Hosts, Hypervisors, Bare Metal, Virtual Machines, and Un-managed resources.



To view the Cloud List View, select the name of a Cloud to display the clouds Detail View.

EDIT Edit the setup configuration of the Cloud.

REFRESH Force a sync with the Cloud. Last sync date, time and duration is shown under the Cloud name.

DELETE Delete the Cloud from Morpheus

Important: All Instances and managed Hosts and VM's associated with the Cloud must be removed prior to deleting a cloud.

Cloud Detail Tabs

Note: Not all tabs are available for all Cloud Types.

Clusters The Clusters tab displays clusters provisioned into the Cloud being viewed, including their status, type, name, layout, workers, and compute, memory, and storage stats. You can add a cluster by clicking *ADD CLUSTER*.

Hosts

The Hosts tab displays available hosts in the Cloud and displays power, OS, name, type, cloud, IP address, nodes, disk space, memory, and status. You can add a resource by clicking *ADD RESOURCE*, add a hypervisor host by clicking *ADD HYPERVISOR*, or perform action an action by selecting one or more Hosts and clicking *ACTIONS*.

VMs (Virtual Machines) Displays an inventory of existing Instances in your Cloud configuration and provides details such as power, OS, name, type, cloud, IP address, nodes, disk space, memory, and status.

Bare Metal Setup PXE Boot in the Boot section to add bare metal servers. Once set up you can view information such as power, OS, name, type, cloud, IP address, nodes, disk space, memory, and status.

Security Groups The Security Groups tab displays a list of existing security groups in the cloud. You can add a security group to this cloud by clicking *EDIT SECURITY GROUPS*.

Load Balancers The load balancers tab panel displays available load balancers in the cloud including the name, description, type, cloud and host. You can add a load balancer from this tab by clicking *ADD LOAD BALANCER*.

Networks Displays Networks synced or added to the Cloud, including their name, type, CIDR, pool, DHCP status, visibility and targeted Tenant.

Data Stores Displays Datastores synced or added to the Cloud, including their name, type, capacity, online status, visibility, and targeted Tenant.

Resources Displays Resource Pools synced from the Cloud, including their name, description, and targeted Tenant.

Policies Manages Policies enforced on the Cloud. Setting a policy on this tab is equal to creating a policy in Administration > Policies and scoping it to the selected Cloud.

Deleting Clouds

To delete a cloud:

1. Select the Infrastructure link in the navigation bar.
2. Select the Clouds link in the sub navigation bar.
3. Click the Delete icon of the cloud to delete.

Important: All Instances, managed Hosts and VMs must be removed prior to deleting a Cloud. To remove Instances, hosts and VMs from Morpheus without deleting the Cloud resources they represent, select Delete on the host or VM, unselect “Remove Infrastructure”, and select “Remove Associated Instances” if Instance are associated with the selected Hosts or VMs.

Clusters

Overview

Infrastructure -> Clusters is for creating and managing Kubernetes Clusters, Morpheus manager Docker Clusters, KVM Clusters, or Cloud specific Kubernetes services such as EKS. The Combo Cluster is a combination Kubernetes, KVM and Functions* Cluster, with all nodes supporting all three provision types.

Cluster Types

Name	Description	Supported Clouds	Provider Type
Kubernetes Cluster	Provisions by default a Kubernetes cluster consisting of 1 Kubernetes Master and 3 Kubernetes Worker nodes. Additional system layouts available including Master clusters. Custom layouts can be created.	All	Kubernetes
Docker Cluster	Provisions by default a Morpheus controlled Docker Cluster with 1 host. Additional hosts can be added. Custom layouts can be created. Existing Morpheus Docker Hosts are automatically converted to Clusters upon 4.0.0 upgrade.	All	Docker
EKS Cluster	Provisions a EKS master and 3 EC2 worker node.	AWS	Kubernetes
AKS Cluster		Azure	Kubernetes
KVM Cluster	Provisions by default a Morpheus controlled KVM Cluster with 1 host. Additional hosts can be added. Custom layouts can be created. Existing Morpheus KVM Hosts are automatically converted to Clusters upon 4.0.0 upgrade.	VMware, Bare Metal	KVM
Combo Cluster	Provisions by default a Morpheus controlled Docker, VM and Functions* Cluster with 1 host. Additional hosts can be added.	VMware, Bare Metal	Morpheus
Ext Kubernetes	Brings an existing (brownfield) Kubernetes cluster into Morpheus	All	Kubernetes

Requirements

- Morpheus Role permission Infrastructure: Clusters -> Full required for Viewing, Creating, Editing and Deleting Clusters.
- Morpheus Role permission Infrastructure: Clusters -> Read required for Viewing Cluster list and detail pages.

Cluster Permissions

- **Cluster Permissions** Each Cluster has Group, Tenant and Service Plan access permissions settings (“MORE” > Permissions on the Clusters list page).
- **Namespace Permissions** Individual Namespaces also have Group, Tenant and Service Plan access permissions settings

Kubernetes Clusters

Requirements

- Agent installation is required for Master and Worker Nodes. Refer to `**Morpheus Agent`_` section for additional information.**
- **Access to Cloud Front, Image copy access and permissions for System and Uploaded Images used in Cluster Layouts**
Image(s) used in Cluster Layouts must either exist in destination cloud/resource or be able to be copied to destination by Morpheus, typically applicable for non-public clouds. For the initial provision, Morpheus System Images are streamed from Cloud Front through Morpheus to target destination. Subsequent provisions clone the local Image.
- System Kubernetes Layouts require Master and Worker nodes to access to the following over 443 during K8s install and configuration:
 - Morpheus Appliance url (443)
 - <https://packages.cloud.google.com>
 - <https://storage.googleapis.com>
 - <https://docs.projectcalico.org>
 - <https://openebs.github.io>
 - <https://cloud.weave.works>
- Morpheus Role permission Infrastructure: Clusters -> Full required for Viewing, Creating, Editing and Deleting Clusters.
- Morpheus Role permission Infrastructure: Clusters -> Read required for Viewing Cluster list and detail pages.

Creating Kubernetes Clusters

Provisions a new Kubernetes Cluster in selected target Cloud using selected Layout.

System (Morpheus provided) Kubernetes Layouts:

Morpheus provides the following layouts for VMware vCenter, VMware Fusion, AWS, Openstack and Nutanix Clouds types.

Kubernetes Cluster 1.14 on Ubuntu 16.04, Weave, OpenEBS Kubernetes Master and 3 Worker Nodes

Kubernetes 1.14 on Ubuntu 16.04, Weave, OpenEBS Single Kubernetes Master

Kubernetes Cluster 1.17 on Ubuntu 18.04, Weave, OpenEBS Kubernetes Master and 3 Worker Nodes

Kubernetes 1.17 on Ubuntu 18.04, Weave, OpenEBS Single Kubernetes Master

Note: The minimum recommended memory size for a Kubernetes cluster is 8 GB.

To create a new Kubernetes Cluster:

1. Navigate to `Infrastructure > Clusters`
2. Select `+ ADD CLUSTER`
3. Select `Kubernetes Cluster`
4. Select a Group for the Cluster
5. Select `NEXT`
6. Populate the following:
 - CLOUD** Select target Cloud
 - CLUSTER NAME** Name for the Kubernetes Cluster
 - RESOURCE NAME** Name for Kubernetes Cluster resources
 - DESCRIPTION** Description of the Cluster
 - VISIBILITY**
 - Public** Available to all Tenants
 - Private** Available to Master Tenant
 - LABELS** Internal label(s)
7. Select `NEXT`
8. Populate the following:

Note: VMware sample fields provided. Actual options depend on Target Cloud

LAYOUT Select from available layouts. System provided layouts include Single Master and Cluster Layouts.

PLAN Select plan for Kubernetes Master

VOLUMES Configure volumes for Kubernetes Master

NETWORKS Select the network for Kubernetes Master & Worker VM's

CUSTOM CONFIG Add custom Kubernetes annotations and config hash

CLUSTER HOSTNAME Cluster address Hostname (cluster layouts only)

POD CIDR POD network range in CIDR format ie 192.168.0.0/24 (cluster layouts only)

WORKER PLAN Plan for Worker Nodes (cluster layouts only)

NUMBER OF WORKERS Specify the number of workers to provision

LOAD BALANCER Select an available Load Balancer (cluster layouts only) }

User Config

- CREATE YOUR USER** Select to create your user on provisioned hosts (requires Linux user config in Morpheus User Profile)
- USER GROUP** Select User group to create users for all User Group members on provisioned hosts (requires Linux user config in Morpheus User Profile for all members of User Group)

Advanced Options

DOMAIN Specify Domain override for DNS records

HOSTNAME Set hostname override (defaults to Instance name unless an Active Hostname Policy applies)

9. Select *NEXT*
10. Select optional Workflow to execute
11. Select *NEXT*
12. Review and select *COMPLETE*
 - The Master Node(s) will provision first.
 - **Upon successful completion of VM provision, Kubernetes scripts will be executed to install and configure Kubernetes**

Note: Access to the sites listed in the **Requirements** section is required from Master and Worker nodes over 443

- After Master or Masters are successfully provisioned and Kubernetes is successfully installed and configured, the Worker Nodes will provision in parallel.
 - **Provision status can be viewed:**
 - From the Status next to the Cluster in Infrastructure -> Clusters
 - Status bar with eta and current step available on Cluster detail page, accessible by selecting the Cluster name from Infrastructure -> Clusters
 - All process status and history is available - From the Cluster detail page History tab, accessible by selecting the Cluster name from Infrastructure -> Clusters and the History tab - From *Operations - Activity - History* - Individual process output available by clicking *i* on target process
13. Once all Master and Worker Nodes are successfully provisioned and Kubernetes is installed and configured, the Cluster status will turn green.

Important: Cluster provisioning requires successful creation of VMs, Agent Installation, and execution of Kubernetes workflows. Consult process output from ``Infrastructure -> Clusters - Details and morpheus-ui current logs at Operations - Health - Morpheus Logs for information on failed Clusters.

Adding Worker Nodes

1. Navigate to Infrastructure - Clusters
2. Select v MORE for the target cluster
3. Select ADD (type) Kubernetes Worker

NAME Name of the Worker Node. Auto=populated with `${cluster.resourceName}-worker-${seq}`

DESCRIPTION Description of the Worker Node, displayed in Worker tab on Cluster Detail pages, and on Worker Host Detail page

CLOUD Target Cloud for the Worker Node.

4. Select *NEXT*
5. Populate the following:

Note: VMware sample fields provided. Actual options depend on Target Cloud

SERVICE PLAN Service Plan for the new Worker Node

NETWORK Configure network options for the Worker node.

HOST If Host selection is enabled, optionally specify target host for new Worker node

FOLDER

Optionally specify target folder for new Worker node

Advanced Options

DOMAIN Specify Domain override for DNS records

HOSTNAME Set hostname override (defaults to Instance name unless an Active Hostname Policy applies)

6. Select *NEXT*
7. Select optional Workflow to execute
8. Select *NEXT*
9. Review and select *COMPLETE*


Note: Ensure there is a default StorageClass available when using a Morpheus Kubernetes cluster with OpenEBS so that Kubernetes specs or HELM templates that use a default StorageClass for Persistent Volume Claims can be utilised.

Kubernetes Cluster Detail Pages

- Cluster status check results icon
- Name of the Cluster
- Last sync date, time and duration
- **Edit, Delete and Actions buttons**
 - **Actions**
 - * **Refresh**
 - Sync the Cluster Status
 - * **Permissions** View and edit Cluster Group, Tenant and Service Plan Access
 - * **View API Token** Displays API Token for Cluster
 - * **View Kube Config** Displays Cluster Configuration
- Costs this month (to date, when Show Costing is enabled)
- Cluster resource utilization stats

- Counts for current Masters, Workers, Containers, Services, Jobs and Discovered Containers in the Cluster


SUMMARY


















Kubernetes Cluster summary tab contains:

- More Cluster metadata including Name, Type, Created By, Worker CPU, Worker Memory (used/max), Worker Storage (used/max), Enabled: Yes/No, and Description.
- Memory chart with total Cluster Free and Used Memory over last 24 hours
- Storage chart with total Cluster Reserved and Used Storage over last 24 hours
- CPU chart with total Cluster CPU Utilization over last 24 hours
- IOPS Chart with total Cluster IOPS over last 24 hours
- IOPS Chart with total Cluster Network utilization over last 24 hours

NAMESPACES

Search 

ACTIONS  CREATE NAMESPACE

<input type="checkbox"/>	NAME	DESCRIPTION	REGION	
<input type="checkbox"/>	default			 
<input type="checkbox"/>	jwKub40VmC23			
<input type="checkbox"/>	kube-node-lease			 
<input type="checkbox"/>	kube-public			 
<input type="checkbox"/>	kube-system			 
<input type="checkbox"/>	logging			 
<input type="checkbox"/>	openebs			 
<input type="checkbox"/>	wheeler			 

WIKI

vmwKubernetesCluster21

EDIT

Last Updated By: Jeff Wheeler on 07/06/2019 10:48 PM

ADD IMPORTANT INFORMATION, NOTES, HOW-TO'S AND ANY OTHER INFO ABOUT YOUR CLUSTER TO THE WIKI PAGE.

All Wiki pages are also accessible under **Operations -> Wiki** with proper Role Permissions.

- API Token can be accessed via **Actions -> View API Token**
- Cluster Config can be accessed via **Actions -> View Kube Config**

MASTERS

POWER	OS	NAME	TYPE	CLOUD	IP ADDRESS	COMPUTE	MEMORY	STORAGE
		jwKub40VmC23-master	Kube Master	vCenter 180	Address 10.30.20.141	<div>31</div>	<div>0</div>	<div>0</div>

WORKERS

POWER	OS	NAME	TYPE	CLOUD	IP ADDRESS	COMPUTE	MEMORY	STORAGE
		jwKub40VmC23-worker-1	Kube Worker	vCenter 180	Address 10.30.20.85	<div>17</div>	<div>0</div>	<div>0</div>
		jwKub40VmC23-worker-2	Kube Worker	vCenter 180	Address 10.30.20.159	<div>14</div>	<div>0</div>	<div>0</div>
		jwKub40VmC23-worker-3	Kube Worker	vCenter 180	Address 10.30.20.70	<div>18</div>	<div>0</div>	<div>0</div>

CONTAINERS

Search	Pods	Storage	All Workers	ACTIONS
<input type="checkbox"/> STATUS	NAME	COMPUTE	MEMORY	STORAGE
<input type="checkbox"/>	cstor-sparse-pool-4yfw-75499b5646-5cg87	0	0	0
<input type="checkbox"/>	cstor-sparse-pool-qx1q-9b8f89b8f-z7qmm	0	0	0
<input type="checkbox"/>	cstor-sparse-pool-w7k6-75547d7f7d-x58dn	0	0	0
<input type="checkbox"/>	maya-apiserver-d9589fbc6-zkz2b	0	0	0
<input type="checkbox"/>	openebs-ndm-7mfqp	0	0	0
<input type="checkbox"/>	openebs-ndm-pgnzh	0	0	0
<input type="checkbox"/>	openebs-ndm-x228k	0	0	0
<input type="checkbox"/>	openebs-provisioner-65fd45cf47-shnxj	0	0	0
<input type="checkbox"/>	openebs-snapshot-operator-86996d865f-qv6jd	0	0	0

HISTORY

NAME	DESCRIPTION	CREATED BY	START DATE	ETA/DURATION	STATUS	ERROR
jwKub40VmC23-worker-3	Server provision	Jeff Wheeler	06/28/2019 01:30 AM	00:28:36	Complete	
jwKub40VmC23-worker-1	Server provision	Jeff Wheeler	06/28/2019 01:30 AM	00:24:31	Complete	
jwKub40VmC23-worker-2	Server provision	Jeff Wheeler	06/28/2019 01:30 AM	00:26:12	Complete	
jwKub40VmC23-master	Server provision	Jeff Wheeler	06/28/2019 12:30 AM	00:43:26	Complete	

Docker Clusters

Provisions a new Docker Cluster managed by Morpheus.

To create a new Docker Cluster:

1. Navigate to Infrastructure > Clusters
2. Select + *ADD CLUSTER*
3. Select Docker Cluster
4. Populate the following:

CLOUD Select target Cloud

CLUSTER NAME Name for the Docker Cluster

RESOURCE NAME Name for Docker Cluster resources

DESCRIPTION Description of the Cluster

VISIBILITY

Public Available to all Tenants

Private Available to Master Tenant

LABELS Internal label(s)

5. Select *NEXT*

6. Populate the following (options depend on Cloud Selection and will vary):

LAYOUT Select from available layouts.

PLAN Select plan for Docker Host

VOLUMES Configure volumes for Docker Host

NETWORKS Select the network for Docker Master & Worker VM's

NUMBER OF HOSTS Specify the number of hosts to be created

User Config

CREATE YOUR USER Select to create your user on provisioned hosts (requires Linux user config in Morpheus User Profile)

USER GROUP Select User group to create users for all User Group members on provisioned hosts (requires Linux user config in Morpheus User Profile for all members of User Group)

Advanced Options

DOMAIN Specify Domain for DNS records

HOSTNAME Set hostname (defaults to Instance name)

7. Select *NEXT*

8. Select optional Workflow to execute

9. Select *NEXT*

10. Review and select *COMPLETE*

EKS Clusters

Provisions a new Elastic Kubernetes Service (EKS) Cluster in target AWS Cloud.

Note: EKS Cluster provisioning is different than creating a Kubernetes Cluster type in AWS EC2, which creates EC2 instances and configures Kubernetes, outside of EKS.

Create an EKS Cluster

1. Navigate to Infrastructure - Clusters

2. Select + *ADD CLUSTER*

3. Select *EKS Cluster*

4. Populate the following:

LAYOUT Select server layout for EKS Cluster (Kubernetes 1.13 EKS provided)

PUBLIC IP

Subnet Default Use AWS configured Subnet setting for Public IP assignment

Assigned EIP Assigned Elastic IP to Controller and Worker Nodes. Requires available EIP's

CONTROLLER ROLE Select Role for EKS Controller from synced role list

CONTROLLER SUBNET Select subnet placement for EKS Controller

CONTROLLER SECURITY GROUP Select Security Group assignment for EKS Controller

WORKER SUBNET Select Subnet placement for Worker Nodes

WORKER SECURITY GROUP Select Security Group assignment for Worker Nodes

WORKER PLAN Select Service Plan (EC2 Instance Type) for Worker Nodes

User Config

CREATE YOUR USER Select to create your user on provisioned hosts (requires Linux user config in Morpheus User Profile)

USER GROUP Select User group to create users for all User Group members on provisioned hosts (requires Linux user config in Morpheus User Profile for all members of User Group)

Advanced Options

DOMAIN Specify Domain for DNS records

HOSTNAME Set hostname (defaults to Instance name)

1. Select *NEXT*
2. Select optional Workflow to execute
3. Select *NEXT*
4. Review and select *COMPLETE*

Hosts

Overview

The *Infrastructure -> Hosts* section provides a universal stage for viewing and managing Hosts and Virtual Machines from all of your Clouds.

In this section you can:

- View & Manage all Hosts, Virtual Machines & Bare Metal
- Add Hypervisors
- Convert Hosts, Virtual Machines and Bare Metal to Managed

Important: When local firewall management is enabled, Morpheus will automatically set an IP table rule to allow incoming connections on tcp port 22 from the Morpheus Appliance.

Hosts

Hosts in Morpheus are hypervisors and Docker hosts that your VMs and Container are hosted on, such as ESXi, Hyper-V and Docker hosts. These hosts are populated from integrated clouds, hosts provisioned from Morpheus, or manually added hosts.

Provisioning new hosts takes place in the Infrastructure > Clusters section of Morpheus. For example, provisioning a new Docker cluster in that section will begin the process of creating a Morpheus-managed Docker cluster with one host (by default). Additional hosts and custom layouts can also be created. See the [Clusters section](#) of Morpheus docs for more information.

Virtual Machines

The Virtual Machines tab lists all managed and unmanaged VMs across Morpheus. Managed VMs are either provisioned by Morpheus, or are inventoried VMs that were converted to managed. Unmanaged VMs are from Cloud integrations with “Inventory Existing Instances” enabled in the Cloud settings.

Bare Metal

Bare Metal hosts are from PXE Boot or manually added in this section. Bare Metal hosts that are also Hypervisors will be listed in both the Bare Metal and Hypervisor sections.

Network

Networks

Infrastructure -> Network -> Networks

Overview

The Networks section is for configuring networks across all clouds in Morpheus. Existing networks from Clouds added in Morpheus will auto-populate in the Networks section.

Networks can be configured for DHCP or Static IP assignment, assigned IP pools, and configured for visibility and account assignment for multi-tenancy usage. Inactive Networks are unavailable for provisioning use. In addition, Morpheus allows administrators to restrict management of Morpheus-created Networks through Role permissions.

Configuring Networks

DHCP

To configure a network for DHCP:

1. Navigate to *Infrastructure -> Network -> Networks*
2. Search for the target network
3. Edit the Network by either:
 - Select *Actions -> Edit*
 - Select the Network, then select *Edit*

4. In the Network Config modal, set the DHCP flag as Active (default)
5. Save Changes

Important: The DHCP flag tells Morpheus this network has a DHCP server assigning IP Addresses to hosts. Morpheus does not act as the DHCP server, and provisioning to a network that has the DHCP server flag active in Morpheus, but no DHCP server actually on the network will in most cases cause the instance to not receive an IP address.

Note: When selecting a network with DHCP enabled during provisioning, “DHCP” will populate to the right of the selected network:

Static and IP Pools

To configure a network for Static IP Assignment:

1. Navigate to *Infrastructure -> Network -> Networks*
2. Search for the target network
3. Edit the Network by either:
 - Select *Actions -> Edit*
 - Select the Network, then select *Edit*
4. In the Network Config modal, add the following:
 - Gateway
 - DNS Primary
 - DNS Secondary
 - CIDR ex 10.10.10.0/22
 - VLAN ID (if necessary)
 - Network Pool * Leave as “choose a pool” for entering a static IP while provisioning * Select a Pool to use a pre-configured Morpheus or IPAM Integration IP Pool
 - The Permissions settings are used for Multi-Tenant resource configuration
 - Leave settings as default if used in a single-tenant environment (only one Tenant in your Morpheus appliance)
 - To share this network across all accounts in a multi-tenant environment, select the Master Tenant and set the Visibility to Public
 - To assign this network to be used by only one account in a multi-tenant environment, select the account and set visibility to Private
 - Active
 - Leave as enabled to use this network
 - Disable the active flag to remove this network from available network options
5. Save Changes

Note: When selecting a network with DHCP disabled and no IP Pool assigned during provisioning, an IP entry field will populate to the right of the selected network(s):

Note: When selecting a network with an IP Pool assigned during provisioning, the name of the IP pool will populate to the right of the selected network(s). IP Pools override DHCP.

Advanced Options (Scan Network)

When adding or editing a network there is an option to scan network. If checked scan network will ping the IP's in the network range, and if ping is successful Morpheus will quickly check for listening ports on the IP.

Important: Network scanning may cause network monitoring or other alerts

Subnets

Subnet details can be viewed from the *SUBNETS* tab on the detail page of a specific network. From the *SUBNETS* tab, Morpheus allows the user to search and edit existing subnets.

In an Azure VNet, you can also create new subnets with the *+ADD* button.

CREATE SUBNET

TYPE

Azure Subnet

SUBNET NAME

SUBNET CIDR

The subnet's address range in CIDR notation (e.g. 192.168.1.0/24). It must be contained by the address space of the virtual network.

☒ ACTIVE

CIDR

0.0.0.0/1

DHCP SERVER

ALLOW IP OVERRIDE

NETWORK POOL

Choose a pool

▶ Group Access

▶ Tenant Permissions

SAVE CHANGES

Network Groups

Overview

Network Groups are useful for grouping networks during provisioning and scaling or grouping availability subnets together such that during provisioning vm's within an instance can be round robin provisioned across availability zones.

Adding Network Groups

1. Navigate to *Infrastructure -> Network -> Networks Groups*
2. Select *ADD*
3. Enter the following:

Group info:

- *Name*: Name of the Network Group in Morpheus
- *Description*: Details of the Network Group in Morpheus

Networks

- Search for and select target Networks for the Network Group
- Search for and select target Subnets for the Network Group

Group Access

- Set Group Access and Defaults for the Network Group

Tenant Permissions

- Set Tenant Visibility for Network Group

2. Select *SAVE CHANGES*

Routers

Overview

Routers can be viewed, created, and managed from the Routers tab of the Infrastructure > Networks page. Morpheus supports the creation of the following router types depending on networks that are currently configured:

- Amazon Internet Gateway
- Huawei Router
- Neutron Router
- NSX Edge Gateway
- NSX Edge Logical Router
- NSX-T Cloud Tier0 Gateway
- NSX-T Cloud Tier1 Gateway
- NSX-T Tier0 Gateway
- NSX-T Tier1 Gateway

- Open Telekom Router

Create New Router

1. Navigate to Infrastructure > Networks > Routers tab
2. Click + *ADD*
3. Select the router type and complete the fields on the resulting modal
4. Once complete, click *ADD NETWORK ROUTER*

Editing Existing Routers

1. Navigate to Infrastructure > Networks > Routers tab
2. Click on the pencil icon for the appropriate router
3. After editing router fields, click *SAVE*

Deleting Existing Routers

1. Navigate to Infrastructure > Networks > Routers tab
2. Click on the trash can icon for the appropriate router
3. Acknowledge the pop-up banner ensuring you wish to delete the router

IP Pools

Infrastructure > Network > IP Pools

Overview

The IP Pools tab in the Networks section allows you to create Morpheus-type IP Pools (which is an IP address range Morpheus can use to assign available static IP addresses to Instances) and NSX-T IP Pools. The IP Pool section also displays pools synced from IPAM integrations like Infoblox, Bluecat and others.

To add a Morpheus Network Pool

1. Click + *ADD* in Infrastructure > Network > IP Pools
2. **Enter the following:**
 - Name** A friendly name for the IP Pool in Morpheus.
 - Pool Type** Currently Morpheus-type IP Pools and NSX-T IP Pools (with a configured integration) can be created directly from Morpheus
 - Starting Address** The starting IP address of the IP Pool address range. ex: 192.168.0.2
 - Ending Address:** The ending IP address of the IP Pool address range. ex: 192.168.0.255
3. Save Changes

Note: Multiple Address Ranges can be added to a pool by selecting the + icon next to the address range.

After saving the IP pool will be available for assignment to networks in the NETWORK POOL dropdown when adding or editing a network.

Domains

Infrastructure -> Network -> Domains

Overview

The Domains section is for creating and managing domains for use in Morpheus . Domains are used for setting FQDNs, joining Windows Instances to Domains, and creating A Records with DNS Integrations. The Domains section is also a multi-tenant endpoint for managing domain settings across multiple accounts

- Added and synced Domains are available for selection in the Domain dropdown when provisioning an Instance.
- Default domains can be set for Clouds and Networks in their Advanced Options sections.
- Images can be flagged to Auto-Join Domains in the *Provisioning -> Virtual Images* section.

Important: For an Instance to auto-join a Domain, a Domain must set in the Advanced Options section of the Cloud or Network used when provisioning.

Adding Domains

1. Navigate to *Infrastructure -> Network -> Domains*
2. Select + *Add*
3. Enter the following:

Domain Name Example demo.example.com

Description Descriptive meta-data for use in Morpheus

Public Zone Check for Public Zones, leave uncheck for Private Zones.

Join Domain Controller Enable to have Windows instances join a Domain Controller

Username Admin user for Domain Controller (in domain/username format)

Password Password for DC Username

DC Server (optional) Specify the URL or Path of the DC Server

OU Path (optional) Enter the OU Path for the connection string.

Permissions Configure Tenant permissions in Morpheus for the Domain (only applicable in Multi-tenant Morpheus setups)

Tenant Select the Tenant to set permissions to for the Domain.

Visibility

- Private: Only Accessible by the select Tenant

- Public: Available for use by all Tenants.

4. Save Changes

The Domain has been added and will be selectable in Domain dropdown during provisioning, and in Cloud and Network settings.

Note: Only resources assigned to the Master Tenant can be set as Publicly visible. If the Tenant assigned is not the master tenant, visibility will automatically change to private.

Editing and Removing Domains

- Domains can be edited by selecting the *Actions* dropdown for the Domain and selecting *Edit*.
- Added Domains can be removed from Morpheus by selecting the *Actions* dropdown for the Domain and selecting *Remove*.

Setting the default domain on a Cloud

1. Navigate to *Infrastructure -> Clouds*.
2. Edit the target Cloud.
3. Expand *Advanced Options* section.
4. In the *Domain* dropdown, select the Domain.
5. Save Changes

Setting the default domain on a Network

1. Navigate to *Infrastructure -> Network*.
2. Edit the target Network.
3. Expand *Advanced Options* section.
4. In the *Domain* dropdown, select the Domain.
5. Save Changes

Selecting a Domain while provisioning an instance

1. While creating an instance, in the *Configure* section, expand the *DNS Options*.
2. Select Domain from the *Domain* dropdown.

Proxies

Overview

In many situations, companies deploy virtual machines in proxy restricted environments for things such as PCI Compliance, or just general security. As a result of this Morpheus provides out of the box support for proxy connectivity. Proxy authentication support is also provided with both Basic Authentication capabilities as well as NTLM for Windows Proxy environments. Morpheus is even able to configure virtual machines it provisions to utilize these proxies by setting up the operating systems proxy settings directly (restricted to cloud-init based Linux platforms for now, but can also be done on windows based platforms in a different manner).

To get started with Proxies, it may first be important to configure the Morpheus appliance itself to have access to proxy communication for downloading service catalog images. To configure this, visit the Admin -> Settings page where a section labeled “Proxy Settings” is located. Fill in the relevant connection info needed to utilize the proxy. It may also be advised to ensure that the Linux environment’s `http_proxy`, `https_proxy`, and `no_proxy` are set appropriately.

Defining Proxies

Proxies can be used in a few different contexts and optionally scoped to specific networks with which one may be provisioning into or on a cloud integration as a whole. To configure a Proxy for use by the provisioning engines within Morpheus we must go to *Infrastructure -> Networks -> Proxies*. Here we can create records representing connection information for various proxies. This includes the host ip address, proxy port, and any credentials (if necessary) needed to utilize the proxy. Now that these proxies are defined we can use them in various contexts.

Cloud Communication

When morpheus needs to connect to various cloud APIs to issue provisioning commands or to sync in existing environments, we need to ensure that those api endpoints are accessible by the appliance. In some cases the appliance may be behind a proxy when it comes to public cloud access like Azure and AWS. To configure the cloud integration to utilize a proxy, when adding or editing a cloud there is a setting called “API Proxy” under “Advanced Options”. This is where the proxy of choice can be selected to instruct the Provisioning engine how to communicate with the public cloud. Simply adjust this setting and the cloud should start being able to receive/issue instructions.

Provisioning with Proxies

Proxy configurations can vary from operating system to operating system and in some cases it is necessary for these to be configured in the blueprint as a prerequisite. In other cases it can also be configured automatically. Mostly with the use of cloud-init (which all of our out of the box service catalog utilizes on all clouds). When editing/creating a cloud there is a setting for “Provisioning Proxy” in “Provisioning Options”. If this proxy is set, Morpheus will automatically apply these proxy settings to the guest operating system.

Overriding proxy settings can also be done on the Network record. Networks (or subnets) can be configured in *Infrastructure -> Networks* or on the Networks tab of the relevant Cloud detail page. Here, a proxy can also be assigned as well as additional options like the *No Proxy* rules for proxy exceptions.

Docker

When provisioning Docker based hosts within a Proxy environment it is up to the user to configure the docker hosts proxy configuration manually. There are workflows that can be configured via the Automation engine to make this automatic when creating docker based hosts. Please see documentation on Docker and proxies for specific information.

Proxy setups can vary widely from company to company, and it may be advised to contact support for help configuring morpheus to work in the proxy environment.

Security Groups

Infrastructure -> Network - Security Groups

Overview

A security group acts as a virtual firewall that controls the traffic for one or more Instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated Instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all Instances that are associated with the security group.

Important: The Host Level Firewall must be enabled for Security Groups to be applied. The Host Level Firewall can be enabled in *Administration -> Settings -> Host Level Firewall Enable/Disable*

Important: When local firewall management is enabled, Morpheus will automatically set an IP table rule to allow incoming connections on tcp port 22 from the Morpheus Appliance.

Add Security Group

1. Navigate to *Infrastructure -> Network - Security Groups*
2. Click the + *Add Security Group* button.
3. From the Security Group Wizard input a name, and description.
4. Save Changes

Add Security Group Rule

1. Navigate to *Infrastructure -> Network - Security Groups*
2. Click the name of the security group you wish to add a rule to.
3. From the security group page click the + *Add Rule* button.
4. From the Rule Wizard select the rule type and input source and depending on the type selected protocol and input a port range.
5. Save Changes

Edit security group rule

1. Navigate to *Infrastructure -> Network - Security Groups*
2. Click the name of the security group you wish to edit a rule in.
3. Click the edit icon on the row of the security group rule you wish to edit.
4. Modify information as needed.
5. Save Changes

Delete security group rule

1. Navigate to *Infrastructure -> Network - Security Groups*
2. Click the name of the security group you wish to delete a rule from.
3. Click the delete icon on the row of the security group rule you wish to delete.

Add Cloud Security Group

To add Cloud security group

1. Navigate to *Infrastructure -> Clouds*
2. Click the name of the desired cloud to add a security group
3. Click the Networks tab
4. Within the “Security Groups” section, click on a security group to edit its rules
5. Alternatively, click + *ADD SECURITY GROUP* to create a new one

Remove Cloud Security Group

1. Navigate to *Infrastructure -> Clouds*
2. Click the name of the cloud to remove the Security Group from.
3. Click the Security Groups tab.
4. Click the *Edit Security Groups* button.
5. Click the - (Minus) button of the Security Group from the Added Security groups list to remove.
6. Save Changes

Integrations

Overview

The Network Integrations section allows you to add and manage IPAM, DNS, and Service Registry integrations. These services can also be added in the *Administration -> Integrations* section.

The following integrations are currently supported:

Networking

- Cisco ACI
- VMWare NSX

IPAM

- Infoblox
- Bluecat
- phpIPAM

Security

- Cisco ACI

DNS

- Microsoft DNS
- Power DNS
- Route 53

Service Registry

- Consul

Scoping Services

NETWORKING Networking integrations are available in the *NETWORK MODE* dropdown located under the *Advanced Options* section in Cloud configurations.

IPAM IPAM integrations will populate pools in the IP Pool section, which are available for assignment to networks in the *NETWORK POOL* dropdown when configuring a network.

SECURITY Security integrations are available in the *SECURITY SERVER* dropdown located under the *Advanced Options* section in Cloud configurations.

DNS DNS integrations will populate domains in the *Infrastructure -> Network -> Domains* section, and are available in the *DOMAIN* dropdown located under the *Advanced Options* section in Cloud, Group, and Network configurations, as well as in the *Configure* section of the Create Instance wizard. DNS integrations are also available in the *DNS SERVICE* dropdown located under the *Advanced Options* section in Cloud and Group configurations.

Service Registry Service Registry integrations are available in the *SERVICE REGISTRY* dropdown located under the *Advanced Options* section in Cloud and Group configurations.

Note: Infoblox will also appear as a DNS INTEGRATION option in Clouds and Groups after adding Infoblox IPAM Integration.

Load Balancers

Infrastructure -> Load Balancers

Overview

Morpheus can provision VM or Container HaProxy Load Balancers, Amazon Elastic and Application Load Balancers, Azure Load Balancers, and integrates with several external Load Balancers, including F5, A10, Citrix, and AVI.

Once created or integrated, Load Balancers are available as an option to be added during provision time or post-provisioning.

Once a Load Balancer is added to an instance, you can manually scale or configure auto-scaling based on thresholds or schedules, and burst across clouds with cloud priority.

In the Load Balancers page there are two sections:

Load Balancers View or edit existing Load Balancers, add new Load Balancers.

Virtual Servers View and link to Instances that are attached to load balancers.

Load Balancers

The Load Balancers tab list currently available Load Balancers, which you can select, edit or delete, and is where you can create new or integrate with external Load Balancers.

Add a new Load Balancer

Select + LOAD BALANCER, chose an option, and fill in the required information:

A10 (aXAPI v3)

- API Host
- API Port
- Username
- Password
- Internal IP
- Public IP
- VIP Address
- VIP Port

Amazon ALB

- Scheme
- Internal
- Internet-Facing
- Amazon Subnets (Select + to add additional) * Specify the subnets to enable for your load balancer. You can specify only one subnet per Availability Zone. You must specify subnets from at least two Availability Zones to increase the availability of your load balancer.
- Amazon Security Groups (Select + to add additional)

AVI

- API Host
- API Port
- Username
- Password
- Internal IP
- Public IP
- VIP Address
- VIP Port

Azure Load Balancer

- Cloud
- Resource Group * Populated from cloud selection

Citrix NetScaler

- API Host
- API Port
- Username
- Password

F5 BigIP (v11.4+)

- API Host
- API Port
- Username
- Password
- Management URL

FortiADC

- API HOST
- API PORT
- USERNAME
- PASSWORD
- INTERFACE (syncd on auth)

HaProxy Container (Internal, will create a HaProxy container, must have available docker host to provision to)

- Group
- Cloud
- Name
- Description
- Plan * Select the size of HaProxy container to be provisioned

NSX-T Load Balancer

- NSX-T
- Name
- Description
- Enabled
- Admin State
- Size
- Tier-1 Gateways
- Log Level

Upon saving your new Load Balancer will be added to the Load Balancers list and available in the Load Balancer dropdown in the Provisioning Wizard Automation Section for Instance Types that have scaling enabled.

Load Balancer Detail Pages

In the main Load Balancer page, select an existing Load Balancer to go to that Load Balancers Details Page, which lists Stats, Settings, Actions and Virtual Servers for that load balancer.

Orchestrating Load Balancers

A large part of application orchestration and automation involves tying various web services and backend services into different load balancer configurations. If the automation tool is unable to communicate or integrate with this aspect of your infrastructure, a lot of gaps will be created in the full orchestrated flow of application deployment. This is why Morpheus provides deep integration with load balancers and explicit definitions with catalog items as to how they are connected to provisioned instances. Some of the functionality includes:

- Public Cloud Load Balancer Support
- Private Cloud Load Balancer Support
- Port Type definitions (Profiles like HTTP/HTTPS or UDP)
- SSL Certificate Management and SSL Certificate Upload
- SSL Passthrough or Forced Redirect

Not only does Morpheus have an ability to provision HAProxy based load balancer containers for easy consumption in development environments, but also has direct tie ins with several Load Balancer Types:

- F5 BigIP
- A10
- Netscaler
- AVI
- Amazon ELB
- Amazon ALB
- Azure Load Balancer
- Fortinet
- Openstack Octavia

- HA Proxy
- NSX-T

Morpheus exposes configuration options during provisioning of an Instance relevant and common to each supported LB Integration. In some cases, Morpheus also provides direct management and sync support for VIP configurations on the various Load Balancers (such as F5, and AVI), However in a day to day orchestrated workflow this would not be the ideal means by which a user should consume load balancer services.

By tying the Load Balancer associations into the provisioning of instances and the definition of the instance catalog item, the lifecycle of the VIP can more easily be maintained throughout the lifecycle of whatever application may be deployed.

Setting up an Instance for Load Balancer Consumption

Several of the provided Morpheus instance types are ready to go with load balancer orchestration out of the box (Apache, Nginx, Tomcat, Node.js, etc). It is also fairly easy to extend existing generic instance types during provisioning to be tied to load balancers or to set up said catalog items in advanced for such functionality.

When creating a custom Instance Type (in Provisioning -> Library), one can define a list of exposed ports that the node type within the instance exposes. When defining these exposed ports it prompts for a Name, Port Number, and LB Type. The LB Type is what enables load-balancer functionality. This can either be HTTP,HTTPS, or TCP. This specification helps build the correct profile for the VIP as well as setup the appropriate types of Health Monitors within the target load balancer integration.

Now, when a user consumes this custom instance type (either through single instance provisioning or full application blueprint provisioning), a section appears in the Automation phase of provisioning. Each port that is defined that exposes a load-balancer gets a dropdown to choose which load balancer integration attach to the exposed port and various prompts become available.

These prompts control features ranging from target VIP Address to selecting an SSL Certificate to be applied to the VIP. These SSL Certificates will even go so far as to create SSL Profiles in integrations for things like an F5 automatically for the application. There are also external integrations for SSL Certificate management with Venafi which allows for the consumption of certificates managed by that external system.

Once the instance is provisioned, as part of the final phase, the load balancer configuration will be applied and maintained on the instance. This association can be manipulated after the fact via the “Scale” tab found on the Instance Detail page.

Another benefit to associating load-balancers this way is that the pool members are automatically maintained during scaling events, either via auto-scaling thresholds or manual node additions / removals.

F5 Load Balancers

Add F5 Load Balancer

To add a F5 Load Balancer Integration:

1. Navigate to *Infrastructure -> Load Balancers*
2. Select + *ADD*
3. Select *F5 BigIP*
4. Fill in the following:

GROUP Select the Group the Load Balancer will be available for

CLOUD Select the Cloud the Load Balancer will be available for

NAME Name of the Load Balancer in Morpheus

DESCRIPTION Identifying information displayed on the Load Balancer list page.

VISIBILITY Define Multi-Tenant permissions

API HOST IP or resolvable hostname url.

API PORT Typically 8443

USERNAME API user

PASSWORD API user password

MANAGEMENT URL Example: `https://10.30.20.31:8443/xui/`

Advanced Options (optional)

- **VIRTUAL NAME**
- **POOL NAME**
- **SERVER NAME**

5. Save Changes

Virtual Servers

Instances attached to an F5 will be listed in the Virtual servers tab. Virtual servers can also be manually added in this section.

Add Virtual Server

1. Navigate to *Infrastructure -> Load Balancers*
2. Select F5 Integration name to drill into the detail page
3. Select + *ADD* in the **VIRTUAL SERVERS** tab
4. Fill in the following:
 - **NAME** Name of the Virtual Server in Morpheus
 - **DESCRIPTION** Description of the Virtual Server in Morpheus
 - **Enabled** Uncheck to keep the configuration but disable F5 availability in Morpheus
 - **VIP TYPE**
 - Standard
 - Forwarding (Layer 2)
 - Forwarding (IP)
 - Performance (HTTP)
 - Performance (Layer 4)
 - Stateless
 - Reject
 - DHCP

- Internal
 - Message Routing
- **VIP HOSTNAME** Enter Hostname of the VIP (optional)
- **VIP ADDRESS** Enter IP address for the VIP
- **VIP PORT** Enter port used for the VIP
- **SOURCE ADDRESS** Enter Virtual Server source address
- **PROTOCOL** tcp, udp, or sctp
- **PROFILES** Search for and select from available PROFILES
- **POLICIES** Search for and select from available POLICIES
- **IRULES** Search for and select from available RULE SCRIPTS
- **PERSISTENCE**
 - cookie
 - dest-addr
 - global-settings
 - hash
 - msrdp
 - sip
 - source-addr
 - ssl
 - universal
- **DEFAULT POOL** Select from available POOLS

5. Select *SAVE CHANGES*

Policies

Policies will be synced and listed in the Policies tab. These policies will be available options when creating Virtual Servers.

Pools

Create Pool

NAME Name of the POOL in Morpheus

DESCRIPTION Description of the POOL in Morpheus

BALANCE MODE

- Round Robin
- Least Connections

SERVICE PORT Specify SERVICE PORT for the POOL

MEMBERS Search for and select from available NODES

MONITORS Search for and select from available Monitors

Profiles

SSL Profiles are synced and will be created when an SSL Certificate is assigned in the Load balancer section when provisioning or editing a Load balancer on an Instance.

Monitors

Create Monitor

NAME Name of the MONITOR in Morpheus

DESCRIPTION Description of the MONITOR in Morpheus

PARENT MONITOR Select from available MONITORS

DESTINATION Specify Destination, such a *:443. Default is *:*

INTERVAL Specify Monitor Interval. Default is 5

TIMEOUT Specify Monitor Timeout. Default is 15

MONITOR CONFIG Enter monitor config.

Nodes

Create Node

NAME Name of the NODE in Morpheus

DESCRIPTION Description of the NODE in Morpheus

ADDRESS Enter node address

MONITOR Select from available MONITORS

SERVICE PORT Specify SERVICE PORT for the NODE

Rule Scripts

Rule Scripts will be synced and listed in the RULE SCRIPTS tab. These rules will be available options when creating Virtual Servers.

Citrix NetScaler



Add NetScaler Integration

To add a NetScaler Load Balancer Integration:

1. Navigate to *Infrastructure -> Load Balancers*
2. Select + *ADD*
3. Select *Citrix NetScaler*
4. Fill in the following:

GROUP * Select the Group the Load Balancer will be available for.

CLOUD * Select the Cloud the Load Balancer will be available for.

NAME * Name of the Load Balancer in Morpheus.

DESCRIPTION Identifying information displayed on the Load Balancer list page.

VISIBILITY

Define Tenant Visibility

- Public: Available to all Tenants.
- Private: Only available to specified Tenant.

Tenant If Visibility is set to private, define the Tenant the Load Balancer will be available in.

API URL *

URL of the NetScaler API

- Example: <http://10.30.21.55>

API PORT *

NetScaler API port

- Example: 80

USERNAME * NetScaler service account username

PASSWORD * NetScaler service account password

VIRTUAL NAME

Naming Pattern for new NetScaler Virtual Servers

- If blank, defaults to `morph_lb_${loadBalancer.id}`

SERVICE NAME

Naming Pattern for new NetScaler Services

- If blank, defaults to `morph_service_${container.id}`

SERVER NAME

Naming Pattern for new NetScaler Servers

- If blank, defaults to `morph_server_${server.id}`

Add Load Balancer to Instance

Load Balancers can be added to Instances during Provisioning or to existing Instances. For Load Balancer settings to appear during provisioning, or for the scale tab to be available on an Instance, the instances Node Type must have a LB port defined.

Note: For Load Balancer settings to appear during provisioning, or for the scale tab to be available on an Instance, the instances Node Type must have a LB port defined.

Add Load Balancer during Provisioning

In the Instance Provisioning wizard, Load Balancers can be configured in the Automation -> Load Balancer section.

1. Navigate to *Provisioning* -> *Instances*.
2. Select + *ADD*.
3. Select an Instance Type that supports scaling. (ENABLE SCALING (HORIZONTAL) flagged on Instance Type configuration)
4. Proceed with Instance configuration to the Automation section.
5. Fill in the following:

VIP ADDRESS

Define IP Address for the Virtual Server

- Example: 10.30.23.191

VIP PORT

Define port for the Virtual Server

- Example: 80

VIP HOSTNAME

Define hostname that will resolve to the VIP IP.

- Example: jwDemoHaApp59.den.example.com

VIRTUAL SERVICE NAME Define name for the Virtual Service. Defaults to `${instance.name}`

BALANCE MODE

Specify balance mode for the VIP

- Least Connections
- Round Robin

STICKY MODE

Specify sticky session options for the VIP

- Source IP

- Cookie

SHARED VIP ADDRESS Select if VIP is shared, then enter DIRECT VIP ADDRESS

SSL CERT

SSL Certificate that will be applied to the VIP.

- No SSL
- Select existing Certificate from Infrastructure -> Keys & Certs or from a Trust Provider Integration.

USE EXTERNAL ADDRESS FOR BACKEND NODES

- Select if traffic from LB to Backend Nodes needs to be sent to the External Addresses, or leave deselected to use Internal Addresses for Backend Nodes.

6. Optionally configure auto-scaling configuration in the `Scale` section

7. Select *NEXT* and provision the Instance.

After all nodes in the Instance are provisioned, the LB configuration will be added to the Instance and Virtual Servers, Services and Servers will be created and configured on the NetScaler. The Load Balancer settings and status will be visible in the Instance details page LOAD BALANCER section, with additional details, links, and configurations options available in the SCALE tab.

Storage

Note: In v3.5.2 STORAGE PROVIDERS has been split out into BUCKETS and FILE SHARES sections.

Overview

Infrastructure -> Storage is for adding and managing Storage Buckets, File Shares, Volumes, Data Stores and Storage Servers for use with other Services in Morpheus.

Role Requirements

There are two Role permissions for the *Infrastructure -> Storage* section: *Infrastructure: Storage* and *Infrastructure: Storage Browser*. *Infrastructure: Storage* give Full, Read or No access to the *Infrastructure -> Storage* sections, while *Infrastructure: Storage Browser* is specific to *Buckets* and *Files Shares*. Full *Infrastructure: Storage Browser* permissions allows *Buckets* and *Files Shares* to be browsed and files and folders to be added, downloaded and deleted from the *Buckets* and *Files Shares*. Read *Infrastructure: Storage Browser* permissions allows *Buckets* and *Files Shares* to be browsed only.

Default Storage

The default Storage path for Virtual Images, Backups, Deployment Archives, Archive Service, and Archived Snapshots is *var/opt/morpheus/morpheus-ui/*. It is recommended to add Storage Buckets and File Shares for these targets in the *Infrastructure -> Storage* section to avoid running out of disk space on the Morpheus Appliance.

Storage Buckets

Storage Buckets are for Backup, Archives, Deployment and Virtual Images storage targets. Buckets can be browsed and files and folders can be uploaded, downloaded or deleted from the Bucket section. Retention Policies can be set on Storage Buckets for files to be deleted or backed up to another bucket after a set amount of time.

Supported Bucket Types

- Alibaba
- Amazon S3
- Azure
- Openstack Swift
- Rackspace CDN

Alibaba Buckets

To Add an Alibaba Storage Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Alibaba* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

ACCESS KEY Alibaba Access Key

SECRET KEY Alibaba Secret Key

REGION Enter Alibaba Region for the Bucket

BUCKET NAME Enter existing Alibaba Bucket name, or to add a new Bucket enter a new name and select *Create Bucket*.

Create Bucket Enable if the Bucket entered in BUCKET NAME does not exist and needs to be created.

Default Backup Target Sets this Bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and then select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

Amazon S3 Buckets

To Add an Amazon S3 Storage Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Amazon S3* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

ACCESS KEY AWS IAM Access Key

SECRET KEY AWS IAM Secret Key

BUCKET NAME Enter existing S3 Bucket name, or to add a new Bucket enter a new name and select *Create Bucket*.

CREATE BUCKET Enable if the Bucket entered in BUCKET NAME does not exist and needs to be created. If enabled, select an AWS Region to create the Bucket in.

ENDPOINT URL Optional endpoint URL if pointing to an object store other than amazon that mimics the Amazon S3 APIs.

Default Backup Target Sets this Bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and then select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

Azure Buckets

To Add an Azure Storage Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Azure* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

STORAGE ACCOUNT Name of the Storage Account in Azure for the Bucket

STORAGE KEY Storage Key provided from Azure

SHARE NAME Enter existing Azure Storage Share name, or to add a new Share enter a new name and select *Create Bucket* below.

CREATE BUCKET Enable if the Share entered in SHARE NAME does not exist and needs to be created.

Default Backup Target Sets this bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this Bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and then select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

Dell EMC ECS Buckets

Note: A Dell EMC ECS Storage Server must be configured in *Infrastructure - Storage - Servers* prior to adding a Dell EMC ECS Bucket.

To Add a Dell EMC ECS Storage Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Dell EMC ECS Bucket* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

STORAGE SERVICE Select existing Dell EMC ECS Storage Server (configured in *Infrastructure - Storage - Servers*)

BUCKET NAME Enter a name for the new Dell EMC ECS bucket.

USER Dell EMC ECS User

SECRET KEY Dell EMC ECS Secret key

NAMESPACE Select Dell EMC ECS Namespace for the Bucket

STORAGE GROUP Select a Dell EMC ECS Storage Group

Default Backup Target Sets this bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this Bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and then select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

Openstack Swift Buckets

To Add an Azure Storage Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Openstack Swift* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

USERNAME Openstack Swift Username

API KEY Openstack Swift API Key

BUCKET NAME Enter existing Openstack Swift Bucket name, or to add a new Bucket enter a new name and select *Create Bucket* below.

IDENTITY URL Openstack Swift Identity URL

Create Bucket Enable if the name entered in BUCKET NAME does not exist and needs to be created.

Default Backup Target Sets this bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this Bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and then select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

Rackspace CDN Buckets

To Add a Rackspace CDN Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Rackspace CDN* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

USERNAME Rackspace CDN Username

API KEY Rackspace CDN API Key

REGION Enter Rackspace CDN Region

BUCKET NAME Enter existing Rackspace CDN Bucket name, or to add a new Bucket enter a new name and select *Create Bucket* below.

Create Bucket Enable if the name entered in BUCKET NAME does not exist and needs to be created.

Default Backup Target Sets this bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this Bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and then select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

File Shares

File Shares are for Backup, Archives, Deployment and Virtual Images storage targets. File Shares can be browsed and files and folders can be uploaded, downloaded or deleted from the File Shares section. Retention Policies can be set on Storage File Shares for files to be deleted or backed up to another File Share after a set amount of time.

Supported File Share Types

- CIFS (Samba Windows File Sharing)
- Dell EMC ECS Share
- Dell EMC Isilon Share
- Local Storage
- NFSv3

CIFS File Shares

To Add a CIFS File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *CIFS (Samba Windows File Sharing)* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

HOST

Enter host IP or resolvable hostname Example: 192.168.200.210

USERNAME CIFS Share Username

PASSWORD CIFS Share User Password

SHARE PATH

Enter CIFS Share Path Example: *cifs*

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

Dell EMC ECS File Shares

To Add a Dell EMC ECS File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *Dell EMC ECS Share* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

STORAGE SERVICE Select existing Dell EMC ECS Storage Server (configured in *Infrastructure - Storage - Servers*)

SHARE PATH

Enter Dell EMC ECS Share Path Example: `ecs-file-share-1`

USER Dell EMC ECS User

SECRET KEY Dell EMC ECS Secret key

Volume Size Specify volume size for the File Share (in MB)

Allowed IP's

Specify IP Addresses to limit accessibility to the File Share

Leave blank for open access Click the + symbol to the right of the first ALLOWED IPS field to add multiple IP's

NAMESPACE Select Dell EMC ECS Namespace (syncd)

STORAGE GROUP Select Dell EMC ECS Storage Group (syncd)

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

Dell EMC Isilon File Shares

To Add a Dell EMC Isilon File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *Dell EMC Isilon Share* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

STORAGE SERVICE Select existing Dell EMC Isilon Storage Server (configured in *Infrastructure - Storage - Servers*)

SHARE PATH

Enter Dell EMC Isilon Share Path Example: `ecs-file-share-1`

Volume Size Specify volume size for the File Share (in MB)

Allowed IP's

Specify IP Addresses to limit accessibility to the File Share

Leave blank for open access Click the + symbol to the right of the first ALLOWED IPS field to add multiple IP's

NAMESPACE Select Dell EMC Isilon Namespace (syncd)

STORAGE GROUP Select Dell EMC Isilon Storage Group (syncd)

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

Local Storage File Shares

Important: Local Storage refers to local to the Morpheus Appliance and the path must be owned by *morpheus-app*. Please be conscious of storage space. High Availability configurations require Local Storage File Shares paths to be shared storage paths between the front end Morpheus Appliances.

Note: To change the owner of a file path to be used as a Local Storage File Share, run `chown morpheus-app . morpheus-app /path` on the Morpheus Appliance.

Note: Morpheus will validate path and ownership of the File Share Path.

To Add a Local Storage File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *Local Storage Share* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

STORAGE PATH

Enter the File Share path on the local Morpheus Appliance. Example: `/var/opt/morpheus/morpheus-ui/vms/virtual-images`

Important: High Availability configurations require Local Storage File Shares paths to be shared storage paths between the front end Morpheus Appliances.

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

NFSv3 File Shares

Note: Configure access to the NFS folder on the NFS Provider prior to adding the NFSv3 File Share.

Note: Upon save Morpheus will create a persistent mount owned by `morpheus-app.morpheus-app` on the Morpheus Appliance for the NFSv3 File Share.

To Add a NFSv3 File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *NFSv3* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

HOST Enter the File Share path on the local Morpheus Appliance.

EXPORT FOLDER Enter the NFSv3 Folder

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

Volumes

Volumes sync or created in Morpheus can be viewed in *Infrastructure- Storage - Volumes*. Volumes can be added for Storage Servers integrated with Morpheus in the *Infrastructure- Storage - Servers* section.

Volumes Types

The available Volume Types list and filterable by are:

- 3Par Volume
- Alibaba Cloud SSD
- Alibaba Efficiency Disk
- Alibaba Cloud Disk
- AWS gp2
- AWS io1
- AWS sc1
- AWS st1
- Azure Volume
- Azure Disk
- Bluemix Disk
- Bluemix SAN
- Bluemix SAN
- CD ROM
- DO Disk
- ECS Block Storage
- ECS Object Storage
- ECS Shared File System
- Floppy Disk
- Google Standard

- HP Enclosure Disk
- Oracle iSCSI
- Isilon NFS Volume
- Nutanix IDE
- Nutanix SATA
- Nutanix SCSI
- Open Telekom Volume
- Openstack Disk
- Openstack Volume
- Oracle Block Volume
- Oracle Disk
- Oracle Virtual Volume
- SCVMM Datastore
- Softlayer Disk
- Softlayer SAN
- Softlayer SAN
- Disk
- UpCloud Disk
- UpCloud MaxIOPS
- NULL
- NULL
- VMWare Datastore
- VMWare Disk

CREATE VOLUME

At least one Storage Server Integration from *Infrastructure- Storage - Servers* is required to create volumes from *Infrastructure- Storage - Volumes*.

3par

To Add a 3Par Volume:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the VolumeS tab, Click the + *ADD* button.
4. Select *3Par* from the dropdown list
5. From the CREATE VOLUME Wizard input the following:

SELECT TYPE

STORAGE SERVER Name of the 3par Storage Server added in *Infrastructure- Storage - Servers*

GROUP Select Storage Group

VOLUME TYPE 3Par Volume

Click NEXT Select *NEXT*

CONFIGURE

NAME Name of the Volume

VOLUME SIZE Specify size of the Volume (in MB)

PROVISION TYPE

- FULL
- TPVV
- SNP
- PEER
- UNKNOWN
- TDVV

Click COMPLETE Select *COMPLETE*

Dell EMC ECS

To Add a Dell EMC ECS Volume:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the VolumeS tab, Click the + *ADD* button.
4. Select *Dell EMC ECS* from the dropdown list
5. From the CREATE VOLUME Wizard input the following:

SELECT TYPE

STORAGE SERVER Name of the DELL EMC ECS Storage Server added in *Infrastructure- Storage - Servers*

GROUP Select Storage Group

VOLUME TYPE ECS Block Storage ECS Object Storage ECS Shared File System

Click NEXT Select *NEXT*

CONFIGURE

NAME Name of the Volume

Click COMPLETE Select *COMPLETE*

Dell EMC Isilon

To Add a Dell EMC ECS Volume:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the VolumeS tab, Click the + *ADD* button.
4. Select *Dell EMC Isilon* from the dropdown list
5. From the CREATE VOLUME Wizard input the following:

SELECT TYPE

STORAGE SERVER Name of the Dell EMC Isilon Storage Server added in *Infrastructure- Storage - Servers*

GROUP Select Storage Group

VOLUME TYPE Isilon NFS Volume

Click NEXT Select *NEXT*

CONFIGURE

NAME Name of the Volume

ALLOWED IP's

Specify IP Addresses to limit accessibility to the File Share

Leave blank for open access Click the + symbol to the right of the first ALLOWED IPS field to add multiple IP's

VOLUME SIZE Specify size of the Volume (in MB)

Click COMPLETE Select *COMPLETE*

Servers

Add Storage Server

Adding 3Par Storage Server

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the SERVERS tab, Click the + *ADD* button.
4. From the ADD STORAGE SERVER wizard input the following:

NAME Name of the Storage Server in Morpheus

TYPE Select *3Par*

URL URL Of 3Par Server Example : *https://192.168.190.201:8008*

USERNAME Add your administrative user account.

PASSWORD Add your administrative password.

5. Select *SAVE CHANGES*

The 3Par Storage Server will be added and displayed in the Buckets tab. Buckets, Files Shares and Storage Groups will be synced in.

Adding Dell EMC ECS Storage Server

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the SERVERS tab, Click the + *ADD* button.
4. From the ADD STORAGE SERVER wizard input the following:

NAME Name of the Storage Server in Morpheus

TYPE Select *Dell EMC ECS*

URL

URL Of DELL EMC ECS Server Example : *https://192.168.190.200:4443*

Tip: The port 4443 is the api port for ECS api. This may be different depending on your configuration

USERNAME Add your administrative user account.

PASSWORD Add your administrative password.

S3 SERVICE URL (Optional) Add your S3 service url Example: *http://192.168.190.220:9020*

Note: S3 SERVICE URL is not required if you are not planning on using ECS S3.

5. Select *SAVE CHANGES*

The Dell EMC ECS Storage Server will be added and displayed in the Buckets tab. Buckets, Files Shares and Storage Groups will be synced in.

Adding Dell EMC Isilon Storage Server

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the SERVERS tab, Click the + *ADD* button.
4. From the ADD STORAGE SERVER wizard input the following:

NAME Name of the Storage Server in Morpheus

TYPE Select *Dell EMC Isilon*

URL URL Of Dell EMC Isilon Server Example : *https://192.168.190.202:8080*

USERNAME Add your administrative user account.

PASSWORD Add your administrative password.

PROVISION USER Select Provision User

PROVISION GROUP Select Provision Group

ROOT PATH

Enter Root Path Example : \

5. Select *SAVE CHANGES*

The Dell EMC Isilon Storage Server will be added and displayed in the Buckets tab. Buckets, Files Shares and Storage Groups will be synced in.

Key Pairs & Certificates

Key Pairs

The Key Pairs section enables the following actions: Add and Delete key pairs. Key Pairs are commonly used by Morpheus for accessing instances via SSH. Morpheus stores key pairs to simplify administration and access across both private and public clouds.

Add Key Pair

To Add Key Pair:

1. Navigate to Infrastructure > Keys & Certs
2. On the Key Pairs tab, click + *ADD*
3. From the Add Key Pair wizard input the following as needed:
 - Name
 - Public Key
 - Private Key
 - Passphrase

Note: Certain features do not require storage of the private key.

Delete Key Pair

To Delete Key Pair:

1. Navigate to Infrastructure > Keys & Certs
2. On the Key Pairs tab, select the trash can icon at the end of any row
3. Acknowledge that you wish to delete the selected key pair

SSL Certificates

SSL certificates authenticate the identity of web servers and encrypt the data being transmitted. Morpheus stores SSL certificates to simplify administration and application of SSL certificates to Morpheus-managed resources.

Add SSL Certificate

1. Navigate to Infrastructure > Keys & Certs
2. On the SSL Certificates tab, click + *ADD*
3. From the Add SSL Certificate wizard input the following as needed:
 - Name
 - Domain Name
 - Key File
 - Cert File
 - Root Cert

Delete SSL Certificate

To Delete SSL Certificate:

1. Navigate to Infrastructure > Keys & Certs
2. On the SSL Certificates tab, select the trash can icon at the end of any row
3. Acknowledge that you wish to delete the selected SSL Certificate

Trust Integrations

Some organizations may use outside technologies to manage their key and certificates. Morpheus allows users to integrate with Venafi for trust management. Trust management integrations can be managed from the Integrations tab on the Infrastructure > Keys & Certs page. Additionally, they can be managed in Administration > Integrations.

Currently, Morpheus supports trust integration Venafi. For more detailed information on integrating Venafi with Morpheus, take a look at our [integration guide](#).

PXE Boot

Overview

Morpheus includes a built in PXE Server to enable easy and rapid bare metal provisioning.

Prerequisites

- DHCP server with following config added to dhcpd.conf

```
allow booting;
allow bootp;
option option-128 code 128 = string;
option option-129 code 129 = text;
next-server morpheus-appliance-ip;
filename "pxelinux.0";
```

Note: Replace `morpheus-appliance-ip` in the `dhcpd.conf` file with your Morpheus appliance IP address.

- Internal Appliance URL (PXE) set in *Administration - Settings*. For PXE-Boot your appliance needs to be routable directly with minimal NAT masquerading. This allows one to override the default appliance url endpoint for use by the PXE Server. If this is unset, the default appliance url will be used instead.
- Mac or IP addresses of PXE target mapped in {morpheus} *Infrastructure -> Boot - Mapping*
- Target host configured for Network boot in BIOS

Note: On the Morpheus Appliance, PXE is enabled by default and port 69 is forwarded to the Internal PXE port 6969. These settings are configurable in in the `pxe:` section of `/opt/morpheus/conf/application.yml`.

Mapping

Add Mapping

1. Select the Mapping tab then click the Add Mapping button.

2. From the New Mapping Wizard input the following information:

Match Pattern Mac address separated by ':' or an ip address filter

Description(optional) Description of the new mapping.

Active Flag to denote the mapping as active or disabled.

Operating System List of operating systems for the mapping.

Boot Image Lists available PXE boot images.

Answer File Lists available answer files.

Cloud Lists the available clouds.

Server Mode List of server modes:: unmanaged, Managed, Bare metal host, Container host, VM host, and Container & VM host.

3. Save

Once the mapping is added, and the target host is powered on, the {morpheus} PXE menu will load and PXE boot will start.

Edit Mapping

1. Click the edit icon on the row of the mapping you wish to edit.
2. Modify information as needed.
3. Click the Save Changes button to save.

Delete Mapping

1. Click the delete icon on the row of the mapping you wish to delete.

Boot Menus

System-seeded Boot Menus are displayed and user-created Boot Menus can be edited and deleted. User-created Boot Menus are edited or deleted by clicking on the pencil or trash can icon in the appropriate row.

Adding a Boot Menu

To begin, click + *ADD*. Available fields include:

- NAME
- DESCRIPTION
- ENABLED
- DEFAULT MENU
- ROOT MENU
- MENU NAME
- BOOT IMAGE
- ANSWER FILE
- MENU CONTENT
- SUB MENUS

Click *SAVE CHANGES*

Answer Files

Answer files are like lists of answers for questions that you know the setup program is going to ask but the user is not prepared to answer. They contain one or more sections, and each section contains one or more properties in the form name=value. Morpheus provides Answer Files for ESXi, CentOS, Ubuntu and XenServer, and user can add their own.

Add Answer Files

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar
3. Select the Answer Files tab then click the Add Answer File button.
4. From the New Answer File Wizard input the following information

Name Name of the answer file.

Description(optional) Description of the new answer file.

Active Flag to denote the mapping as active or disabled.

Script Name Name of the new answer file.

Script Version Version of the new answer file.

Script The script for the new answer file.

5. Save

Edit Answer File

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar
3. Select the Answer Files tab
4. Click the edit icon on the row of the answer file you wish to edit.
5. Modify information as needed.
6. Save Changes

Delete Answer File

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar
3. Select the Answer Files tab.
4. Click the delete icon on the row of the answer file you wish to delete.

Images

Morpheus provides Images for ESXi, CentOS, Ubuntu and XenServer, and user can add their own Images.

Add Images

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar
3. Select the Images tab then click the Add Image button.
4. From the Upload Virtual Image Wizard input the following information
 - Name** Name of the Image.
 - Operating System** List of available operating systems.
 - Storage Provider** List of available storage providers.
 - Image Path** Path of the image.
 - Visibility** Private or Public
 - Account** List of accounts to allow permission to this image.
5. Save Changes

Edit Image

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar
3. Select the Images tab
4. Click the actions drop down and select edit.
5. Modify information as needed.
6. Click the Save Changes button to save.

Convert Image

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar.
3. Select the Images tab
4. Click the *Actions* drop and select *Convert*.

Download Image

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar.
3. Select the Images tab
4. Click the *Actions* drop and select *Download*.

Remove Image

1. Click the Infrastructure link in the navigation bar.
2. Click the Boot link in the sub navigation bar.
3. Select the Image tab.
4. Click the *Actions* drop and select *Remove*.

1.3.4 Administration

There are several administrative integrations built into Morpheus that make it great to work with within any organization ranging from small to large. Especially, with its built in white label support and multitenancy capabilities, managed service providers have a wide range of capabilities when it comes to managing customer accounts and users.

Tenants

Overview

A Tenant in Morpheus is an isolated environment with unique users and workloads. The Master Tenant is the default Tenant in Morpheus, created upon installation. All other Tenants outside of the Master Tenants are Subtenants.

- The Master Tenant is the default Tenant created during the installation of Morpheus
- All Tenants created after installation are Subtenants. Only one Master Tenant can exist
- The Master Tenant creates and controls all Subtenants.
- Tenants are isolated environments with unique users, workloads, and Groups
- The Master Tenant can share or assign its resources to Subtenants
- Subtenants cannot share their resources with other Tenants
- Subtenants cannot see resources from other Subtenants
- Subtenants can only access Master Tenant resources that have been set to Public visibility or specifically assigned to the Subtenant

Roles

There are two Role types in Morpheus, Tenant Roles and User Roles. Understanding these Role types is key to effectively administering Role permissions in Morpheus. These two Role types are discussed in greater detail in this section.

Tenant Roles

Tenant Roles set the maximum permission levels for Users in the Tenant. User Role permissions will not exceed the permissions of the Tenant Role.

- Tenant Roles set the maximum permissions for a Tenant
- User Roles in a Tenant cannot exceed the permissions of the Tenant Role
- A Tenant Role can be assigned to one or multiple Tenants
- Tenant Roles determine Cloud access for the Subtenant such that all Clouds in the Master Tenant which have visibility set to Public will show as options in the Tenant Role Cloud Access tab
- Only Master Tenant Clouds given access in the Tenant Role will be accessible in the Subtenant

Important: Tenant Roles cap permissions on all Subtenant User Roles. User Roles can be created in the Subtenant with lesser permissions than the Tenant Role allows. Tenant Roles are designed for a Master Tenant Admin to set max permissions for the Subtenant, and a Subtenant Admin to configure User Roles inside the Subtenant.

User Roles

User Roles determine feature, Group, and Instance Type access for all Users. In a multi-Tenant environment, there are two types of User Roles: Single-Tenant User Roles and Multi-Tenant User Roles.

- **Single-Tenant User Roles:** These exist solely in the Tenant they are created in. All Roles created in a Subtenant will be Single-Tenant User Roles
- **Multi-Tenant User Roles:** The Master Tenant can create Multi-Tenant User Roles. These Roles are automatically seeded into Subtenants and can be assigned to Subtenant Users. Changes to Multi-Tenant User Roles made in the Master Tenant are propagated to all Subtenants. However, once a Multi-Tenant User Role is edited inside a Subtenant, it is no longer linked to the Multi-Tenant User Role and becomes its own unique Role. It will no longer receive propagated changes.

Note: Multi-Tenant User Roles are intended to make Subtenant User Role creation easier, so Master Tenant Users do not have to re-create the same base Subtenant Users Roles for every Subtenant. Multi-Tenant User Roles are not a single Role across Tenants, but more like a template that creates new Subtenant User Roles that can then be managed in the Sub Tenant.

Tenants

The Tenants page displays a list of all Tenants. This page enables users to Create, Edit, and Delete Tenants. The list of Tenants displays the Tenant Name, Role, Total Instances, Total Users, and the Created Date.

Click the Tenant Name to drill into the Tenant View where you can again Edit or Delete the Tenant, as well as Create Users, Edit Users, and Delete Users users belonging to the Tenant.

Create Tenants

To Create Tenants:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the Create Tenant button.
4. From the New Tenant wizard input:
 - Name
 - Description (optional)
 - Subdomain
 - Base Role
 - Currency
5. Within the Advanced Options section, track customer data related to the Tenant if needed:
 - Account Number
 - Account Name
 - Customer Number
6. Click the *Save Changes*

Edit Tenant

To edit a Tenant:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the Edit pencil icon on the row of the Tenant to edit.
4. Edit the Edit Tenant settings.

Disabling Tenant

When disabling a tenant, they are not able to login and cannot be impersonated by another tenant. However all of their information will still remain in Morpheus and they may still receive notifications and alerts.

To disable a Tenant:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the Edit pencil icon on the row of the Tenant to edit.
4. Uncheck the Enabled box.

Delete Tenant

To delete a Tenant:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the Delete trashcan icon on the row of the Tenant to delete.
4. Confirm

Tenant Users

The Tenant View displays a list of users belonging to the Tenant and their Name, Username, Email, and Role.

From this page: Create, Edit, and Delete users within the Tenant.

Important: In versions 3.1.1 and 2.12.5 and later, a Multi-Tenant User Role must be created prior to adding Subtenant Users or the User will not save. In previous versions a default Multi-Tenant Role was seeded. Due to customer requests, the seeded role was removed and a Multi-Tenant Role must be created by the Master Tenant for Subtenant Users.

Create Tenant User

To create a Tenant User:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the Tenant Name on the row of the Tenant where the user will be added.
4. Click + *ADD USER*
5. From the New User wizard, input the fields below:
 - First Name
 - Last Name
 - Username
 - Email address
 - Role (to be inherited by the user)
 - Password
 - Any default Windows or Linux credentials

Click *SAVE CHANGES*

Important: In versions 3.1.1 and 2.12.5 and later, a Multi-Tenant User Role must be created prior to adding Subtenant Users or the User will not save. In previous versions a default Multi-Tenant Role was seeded. Due to customer requests, the seeded role was removed and a Multi-Tenant Role must be created by the Master Tenant for Subtenant Users.

Edit a Tenant User

To edit a User:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the specific Tenant name from the row of available Tenants.
4. Click the Edit pencil icon for your selected Tenant.
5. Edit User information

Note: Name, Username, Passwords and e-mail addresses cannot be edited on Users created from Identity Source Integrations.

Click *SAVE CHANGES*

Delete Tenant User

To delete a Tenant User:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the Tenant Name from the row for the Tenant containing the user.
4. Click the Delete trashcan icon of the row of the user to delete.
5. Confirm

Subtenant User Login

Subtenant Users can have the same Username as the User on the Master Tenant or any other Tenant. Subtenant Users will now have to login using the subdomain prefix.

Important: Subtenant users will no longer be able to login from the main login page without specifying their subdomain.

Example: I have a username `subuser` that belongs to a tenant with the subdomain `subaccount`. When logging in from the main login url, I would now need to enter in: `subaccount\subuser`

Configuring Tenants and Resources for Multi-Tenancy

A very common scenario for Managed Service Providers is the need to provide access to resources on a customer by customer basis. Several administrative features are available in Morpheus to ensure customer resources are properly scoped and isolated. With its built multi-tenancy capabilities and white label support, managed service providers have a wide range of capabilities when it comes to managing customer Tenants and users.

Tenants

There are essentially two types of Tenants in Morpheus

- Master Tenant
- Sub Tenants

During the initial setup of a Morpheus Appliance, the Master Tenant is created. All Tenants created in addition to this Master Tenant are sub-Tenants. There can only be one Master Tenant, and sub-Tenants cannot become the Master Tenant. The delineation between the Master Tenant and sub-Tenants is important to understand for properly scoping resources across Tenants.

Creating Tenants

The Master Tenant is created during the initial appliance setup. Additional sub-Tenants can be created in the *Administration -> Tenants* section.

The Tenants page displays a list of all Tenants. This page enables users to: Create, Edit, and Delete Tenants. The list of Tenants displays the Tenant Name, Role, Total Instances, Total Users, Status (active or inactive) and the Created Date. Click the Tenant Name to drill into the Tenant View where you can edit or delete the Tenant, as well as create, edit and delete users belonging to the Tenant.

Note: At least one Tenant in addition to the Master Tenant is required to scope resources across Tenants.

To create a new sub-Tenant

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Click the *+Create Tenant* button.
4. From the New Tenant wizard input * Name (Required) * Description * Base Role * Currency (for pricing)

The Base Role defines a role set from which all roles created within the Tenant will inherit.

Note: In prior versions, we could set Limits when creating a Subtenant. These could restrict the amount of storage, memory, and CPUs that can be collectively provisioned by all users in the Tenant. In more recent versions, this functionality has been rolled into Policies (Administration > Policies). When creating a Policy, we are able to specify a Tenant to which the Policy should apply.

Click the *Save Changes* button.

NEW TENANT

NAME

DESCRIPTION

☒ ENABLED

SUBDOMAIN

The SubDomain is used for creating a direct login url in Identity Sources or as a login prefix to identify the tenant i.e. 'subdomain\username'

BASE ROLE

Select

This is the primary role of the tenant. All roles created within the tenant must inherit this role as the base role.

CURRENCY

USD

SAVE CHANGES

Viewing Tenants

To View an individual Tenant page, select the Tenant name from the main Tenants section.

OperationsProvisioningInfrastructureBackupsLogsMonitoringToolsAdministration

TenantsPlans & PricingRolesUsersIntegrationsPoliciesProvisioningMonitoringBackupsLogsSettings

Tenants > Example Tenant

Example Tenant

IDENTITY SOURCESACTIONS EDITDELETE

USERS

Search

+ ADD USER

NAME	USERNAME	EMAIL	ROLE	
Chris Jones	chrisjones	chrisjones@morpheusdata.com	Tenant Admin User	ACTIONS
John Smith	jsmith	johnsmith@morpheusdata.com	Tenant Admin User	ACTIONS

From inside the Tenant view, we can edit or delete the Tenant, as well as click into any of the Tenant’s users.

Tenant Users

To create a new user within the Tenant:

Click the *CREATE USER* button, then from the New User wizard input the fields below:

- First Name
- Last Name
- Username
- Email
- Role
- Password
- Confirm Password

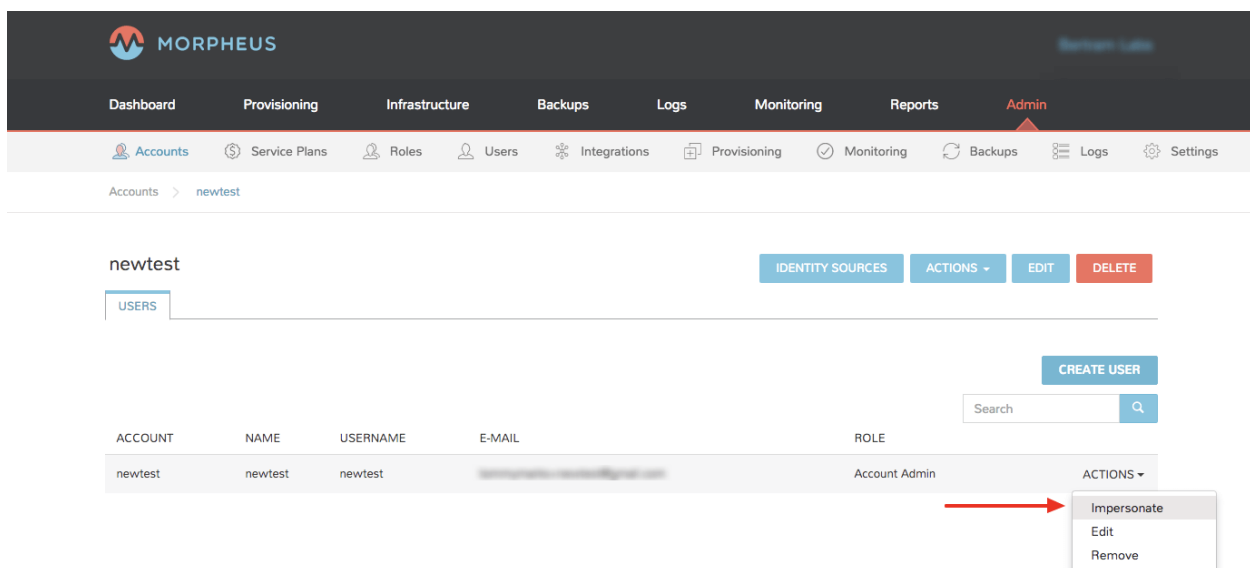
Click *Save Changes*.

Note: Users are specific to each Tenant. Users created in the Master Tenant or other sub-Tenants will only have access to the Tenant they are created in.


Impersonate Tenant User

Morpheus allows admin users in the Master Tenant to impersonate any user in the Subtenants to see the application as if they are that user. To impersonate a user, you must be logged in as a user with the “Impersonate User” feature enabled in the assigned role.

From inside a Tenant detail page (containing the list of that Tenant’s users), and in the specific user’s ACTIONS drop down, select “Impersonate”.



This will log you in as that user in their respective Tenant. To log out of the impersonate users Tenant, select the username in the header, and then select “Quit Impersonating”



Resources

In the Master Tenant, resources can be configured with private or public visibility:

- Private Visibility: Only available to the assigned Tenant.
- Public Visibility (option available in Master Tenant only): Available across all Tenants

Resources in the Master Tenant can also be assigned directly to Subtenants. When a resource is assigned to a Subtenant, it is only available for that Subtenant, and its visibility is automatically set to private. Public visibility is not an option for any resource assigned to or created in a Subtenant.

From the Master Tenant, the following resources can be configured for public visibility across all Tenants, or assigned to individual sub-Tenants

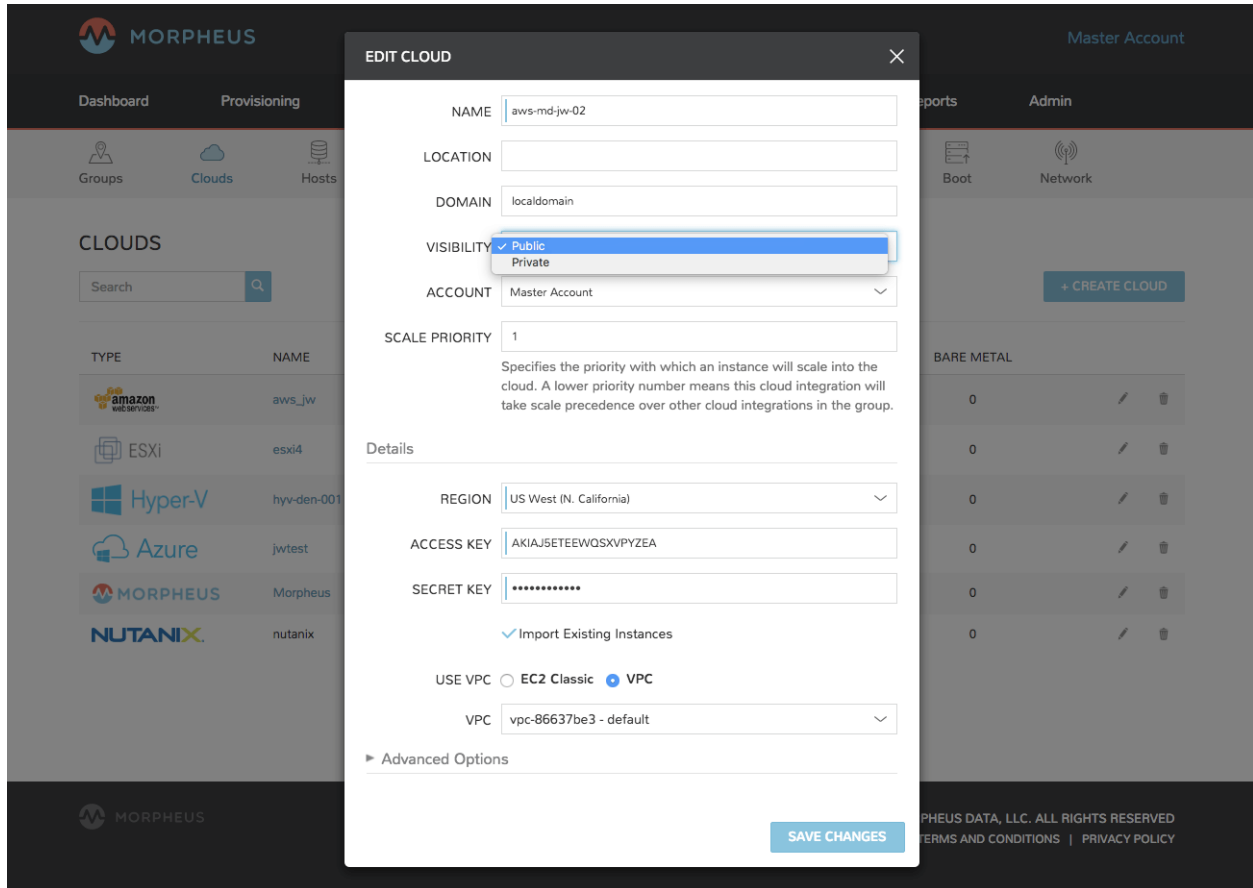
- Clouds
- Hosts
- Virtual Machines
- Networks
- Datastores
- Resource Pools
- Folders
- Virtual Images
- Library Instance Types
- Pricing
- Policies
- Workflows
- Roles

Note: Virtual Image Blueprints can be made available to multiple select Tenants when set to private.

Cloud Visibility & Assignment

To set the visibility of a Cloud to Public (shared across all Tenants) or Private (only available to the assigned Tenant):

1. Navigate to Infrastructure > Clouds
2. Select either the pencil/edit icon on the end of the cloud row, or click the name of the cloud and select “Edit” in the cloud page.
3. From the “Visibility” drop down, select either “Public” or “Private”
4. Select *Save Changes* in the footer of the Edit Cloud modal.



When a cloud is set to Public visibility, it is available to be added to Subtenants. All Subtenants created after a Master Tenant cloud is set to public will automatically have clouds with public visibility added, and a group will be created for each available cloud matching the cloud name in the new Subtenant(s).

For Tenants created prior to a Master Tenant cloud being set to public visibility, the Subtenant will have the option to add that cloud but it will not automatically be added.

While the cloud will be available for Subtenants, the resources available in that cloud to the Subtenant(s) depends on the visibility or assignment of the individual resources.

Note: A Subtenant user must have sufficient role permissions and cloud access to add publicly available clouds. Master Tenant clouds settings cannot be edited from Subtenants.

Assign a Cloud to an Tenant

Important: When assigning a Cloud to a Tenant, all resources for that Cloud will only be available to the assigned Tenant. If a cloud is created in the Master Tenant and assigned to a sub-Tenant, it will no longer be available for use by the Master Tenant or any other sub-Tenants, although it can be assigned back to the Master Tenant, or to another sub-Tenant.

It may be preferable for service providers to share or assign their cloud resources, such as specific hosts, networks, resources pools and datastores, across sub-Tenants, rather than an entire cloud.

To assign a cloud from the Master Tenant to a Sub-Tenant

1. Navigate to Infrastructure, Clouds
2. Select either the pencil/edit icon on the end of the cloud row, or click the name of the cloud and select “Edit” in the cloud page.
3. From the “Tenant” drop down, select the Tenant to assign the cloud to. The visibility will automatically be set to “Private” when a cloud is assigned to a sub-Tenant.
4. Select *Save Changes* in the footer of the Edit Cloud modal.

The screenshot shows the Morpheus web interface with the 'EDIT CLOUD' modal open. The modal contains the following fields and options:

- NAME:** aws-md-jw-02
- LOCATION:** (empty)
- DOMAIN:** localdomain
- VISIBILITY:** Private (dropdown menu)
- ACCOUNT:** A dropdown menu is open showing options: Select, Master Account, Sub-Account 1 (highlighted), Sub-Account 2, and Sub-Account 3.
- SCALE PRIORITY:** (empty)
- Details:**
 - REGION:** US West (N. California)
 - ACCESS KEY:** AKIAJ5ETEEWQSKVPYZEA
 - SECRET KEY:** (masked with dots)
 - ☒ Import Existing Instances
 - USE VPC:** ☐ EC2 Classic ☒ VPC
 - VPC:** vpc-86637be3 - default
 - Advanced Options:** (collapsed)
- SAVE CHANGES:** (button)

The background interface shows the 'CLOUDS' section with a table listing various cloud integrations (aws_jw, esxi4, hyv-den-001, jwtest, Morpheus, nutanix) and a 'BARE METAL' section with a table of instances.

When a cloud is assigned to a sub-Tenant, or assigned to the Master Tenant with private visibility, that cloud and all of its resources are only available to the assigned Tenant. The Master Tenant still maintains control and visibility, and can edit the cloud settings or re-assign the cloud.

Individual Resource Visibility & Assignment

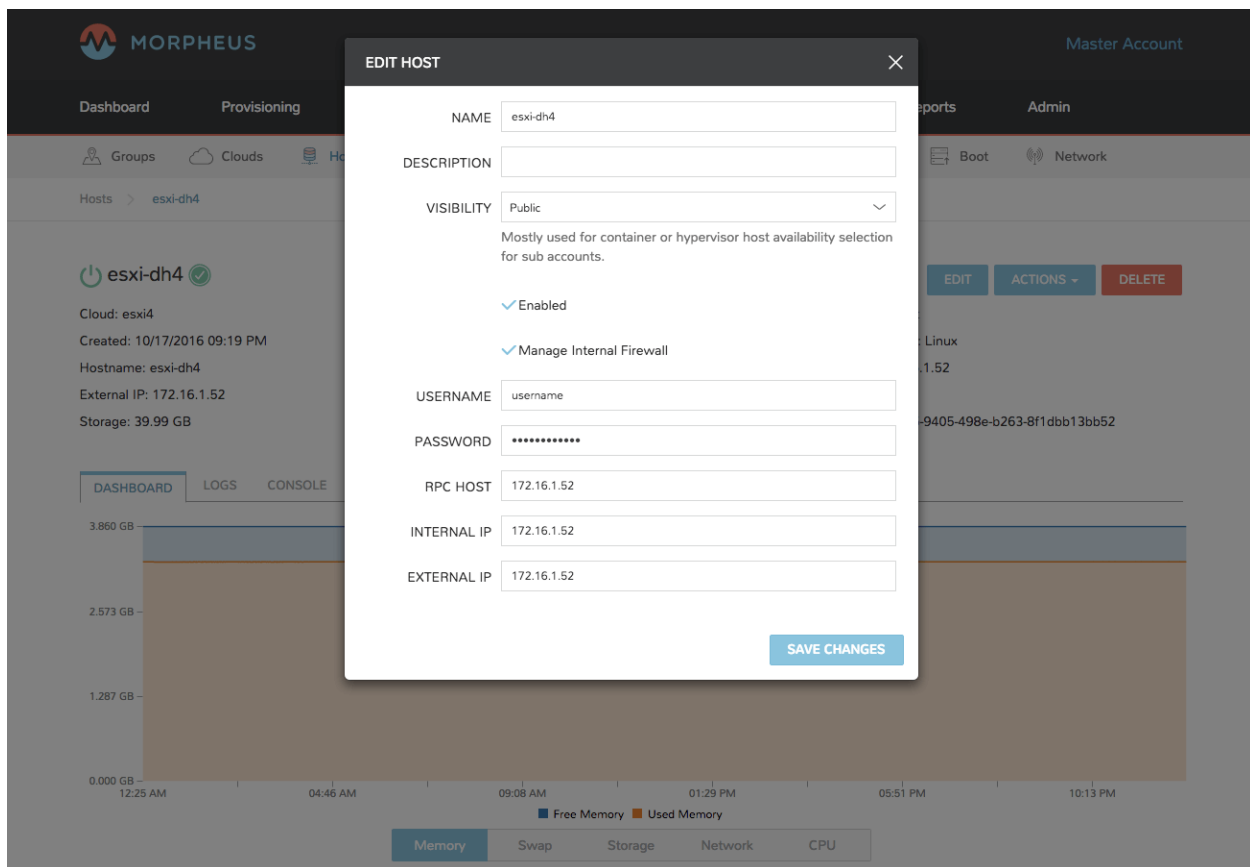
Similar to clouds, individual resources from the Master Tenant can be set to public and available to sub-Tenants, or assigned to sub-Tenants.

By default, any host, virtual machine, bare metal server, network, resource pool, datastore or blueprint added, created or inventoried by an Tenant is assigned to that Tenant. If these resources are in the Master Tenant, they can be assigned to sub Tenants. Assigning one of these resources will make it unavailable to the Master Tenant, but it will still be visible and editable by the Master Tenant. This allows Master Tenant resources to be isolated for use by sub-Tenants while still under the control of the Master Tenant.

Resources assigned to sub-Tenants from the Master Tenant will be visible and available for use by that sub-Tenant, however they cannot be edited or re-assigned by the sub-tenant.

Set the Visibility of a Host, Virtual Machine or Bare metal Server to Public or Private


1. From the Master Tenant, navigate to Infrastructure, Hosts
2. Select either the Hosts, Virtual Machines or Bare Metal tab
3. Click the name of the resource
4. Select *Edit* in the resource page to bring up the config modal
5. From the “Visibility” drop down, select either “Public” or “Private”
6. Select *Save Changes*



Assigning a Host, Virtual Machine, or Bare Metal server to an Tenant

1. From the Master Tenant, navigate to Infrastructure, Hosts


2. Select either the Hosts, Virtual Machines or Bare Metal tab
3. Click the name of the resource
4. From the “Actions” dropdown in the the resource page, select Assign Tenant
5. In the Assign Tenant modal, select the Tenant to assign the resource to.
6. Select *Execute* in the modal



The resource will now be assigned and available for use by the assigned Tenant. If assigned to a sub-Tenant, the Master Tenant will maintain visibility and control.


Set the Visibility of a Network to Public or Private

1. From the Master Tenant, navigate to Infrastructure, Network
2. Select either the pencil/edit icon in the network row, or click the name of the network and select “Edit” in the network page.
3. From the “Visibility” drop down, select either “Public” or “Private”
4. Select *Save Changes* in the modal



Assign a Network to a Tenant

1. From the Master Tenant, navigate to Infrastructure, Network
2. Select either the pencil/edit icon in the network row, or click the name of the network and select “Edit” in the network page.
3. From the “Tenant” drop down, select an Tenant to assign the network to.
4. Select *Save Changes* in the lower the modal



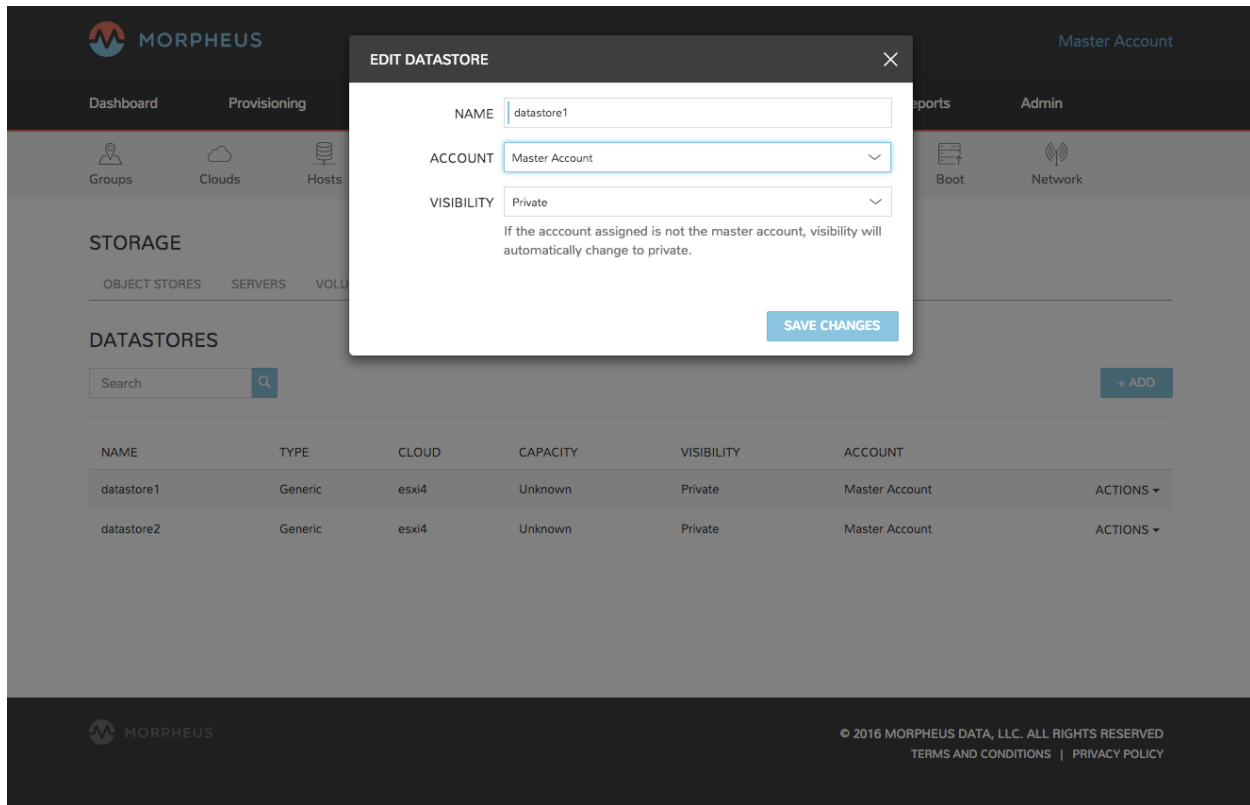
The Network will now be assigned and available for use by the assigned Tenant. If assigned to a sub-Tenant, the Master Tenant will maintain visibility and control.

Set the Visibility or assign a datastore to an Tenant

1. From the Master Tenant, navigate to Infrastructure, Storage
2. Select the “Data Stores” tab
3. Select Edit from the “Actions” dropdown in the datastores row
4. From the “Visibility” drop down, select either “Public” or “Private”
5. From the “Tenant” drop down, select the Tenant to assign the datastore to.

Note: If assigned to a sub-tenant, the visibility will be automatically set to private.

6. Select *Save Changes* in the modal

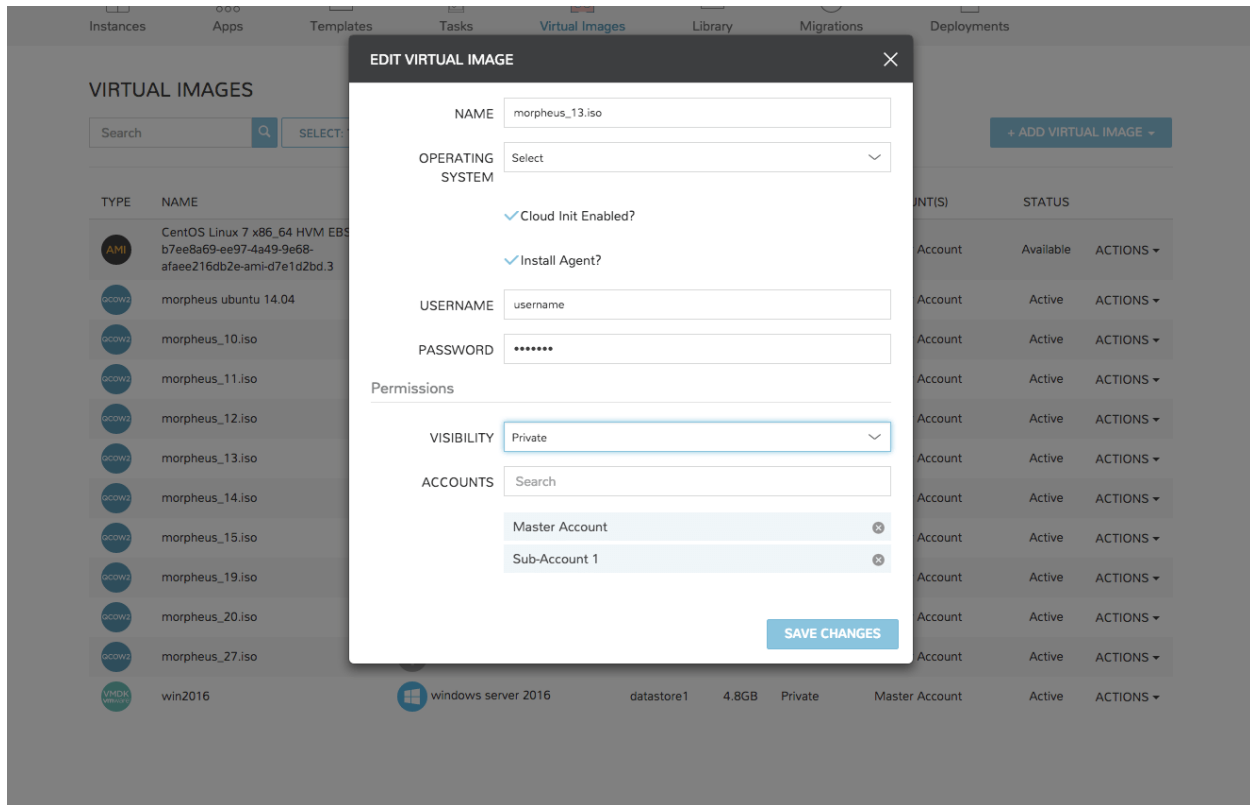


Set the Visibility or assign a Virtual Image to an Tenant

1. From the Master Tenant, navigate to Provisioning, Virtual Images
2. Select Edit from the “Actions” dropdown in the Virtual Images row
3. From the “Visibility” drop down, select either “Public” or “Private”. Public will share the
4. From the “Tenant” field, start typing the name of the Tenant to assign the Virtual Image to. Matching Tenants will populate, then select the Tenant to add.

Note: Virtual Images can be set to Private, but accessible to more than one Tenant

#. Repeat step 4 for all Tenants requiring access to the virtual image. .. To remove access for an Tenant, click the “x” next to the Tenant name #. Select *Save Changes* in the modal



The Virtual Image will now be available for use by the assigned Tenants.

Identity Sources

Administration > Tenants > (Selected Tenant) > Identity Sources Administration
> Users > Identity Sources

Overview

Morpheus can integrate with many of the most common identity source technologies, such as Active Directory, Okta, and many others. These can be configured via the *Identity Sources* button on any Tenant detail page (Administration > Tenants > Selected Tenant) or on the Users list page (Administration > Users). These integrations map roles within these sign-on tools to equivalent roles in Morpheus so at first log in users are assigned the appropriate role.

Active Directory

Overview

Active Directory is Microsoft's primary authentication service widely used in Enterprise organizations and even via Microsoft's cloud services. While Active Directory also supports LDAP protocol support (which Morpheus can integrate with as well), the main Active Directory integration can also be utilized. It is even possible to map Active Directory groups to equivalent Roles within Morpheus. Morpheus will connect over port 389 for non-secure LDAP and port 636 for secure LDAP.

Note: To use Active Directory, a valid / trusted SSL certificate must be in place on the Active Directory services (self signed will not work).

Adding an Active Directory Integration

1. Navigate to Administration -> Tenants
2. Select a Tenant
3. Select *IDENTITY SOURCES*
4. Select + *IDENTITY SOURCE*
5. Choose “Active Directory”
6. Populate the following:

Name Unique name for authentication type.

AD Server Hostname or IP address of AD Server.

Domain Domain name of AD Domain.

Binding Username Service account username for bind user.

Binding Password Password for bind service account.

Required Group The AD group users must be in to have access (optional)

Default Role The default role a user is assigned if no group is listed under AD user that maps under Role Mappings section.

Service Account Holder This is the admin account type in Morpheus and an AD group can be created and populated to a user that this role should be assigned. Roles are assigned dynamically based on group membership.

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

1. Select *SAVE CHANGES*.

Now allowed AD users can login to Morpheus via their Active Directory credentials and a User will be automatically generated to Morpheus with matching metadata and mapped Role permissions.

Note: Only the username is required with password, not the `username@domain`.

Note: Sub-tenant Morpheus API authentication for Active Directory generated users is not currently supported.

Azure Active Directory SSO (SAML)

Azure Active Directory Single Sign-on can be added as a Identity Source in Morpheus using the SAML Identity Source Type. The Azure AD SSO configuration is slightly different than other SAML providers, and this guide will assist in adding a Azure AD SSO Identity Source.

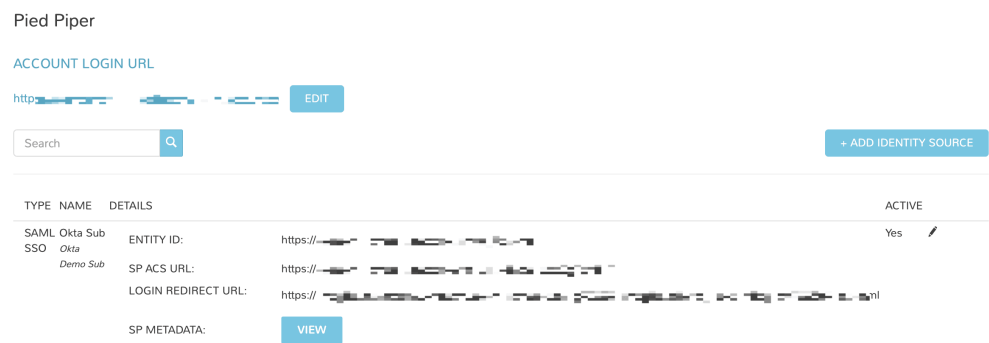
Create a Azure AD SAML Integration

Azure requires inputting the *Identifier (Entity ID)* and *Reply URL (Assertion Consumer Service URL)* in the Azure SSO configuration before it provides the Endpoints and Certificate necessary to add the Integration into Morpheus. In order to get the *Identifier (Entity ID)* and *Reply URL (Assertion Consumer Service URL)* to input into Azure SSO config, we need to create a base SAML Integration in Morpheus first.

To add a base SAML integration:

- 1. Navigate to Administration -> Tenants
- 2. Select a tenant.
- 3. Select IDENTITY SOURCES in the Tenant detail page
- 4. Select + ADD IDENTITY SOURCE.
- 5. Select SAML SSO from the TYPE field
- 6. Add a Name, optional Description and any value in the LOGIN REDIRECT URL field. Since we do not have the LOGIN REDIRECT URL from Azure yet, type any text such as test into the LOGIN REDIRECT URL field so the Identity Source Integration can be saved and the *Identifier (Entity ID)* and *Reply URL (Assertion Consumer Service URL)* generated. We will edit the Integration with the proper LOGIN REDIRECT URL after configuring SSO in Azure.
- 7. Select SAVE CHANGES.

Upon save, the *Entity ID (Identifier (Entity ID))* and *SP ACS URL (Reply URL (Assertion Consumer Service URL))* will be provide in the Identity Source list view. Copy these for use in Azure SSO config.



Configure Azure SSO

This guide assumes an Azure AD Application has already been created in Azure with a subscription level high enough to configure SSO in the application. Please refer to Azure documentation if this has not already been configured.

- Next, in the Azure Active Directory Application details page, select `Single sign-on`, then enter the following:
 - Single Sign-on Mode dropdown** Select `SAML-based Sign-on`
 - Identifier (Entity ID)** Enter the `Entity ID` URL from the Morpheus Identity Source Integration above.
 - Reply URL (Assertion Consumer Service URL)** Enter the `SP ACS URL` from the Morpheus Identity Source Integration above.
- Save and click the `Test SAML Settings` button. Azure will confirm connection with Morpheus
- In Azure's `User Attributes & Claims` settings (step 2), select `Add a group claim` with value `user.groups [SecurityGroup]`

User Attributes & Claims config

Table 9: Required Claim

Claim name	Value
Unique User Identifier (Name ID)	<code>user.userprincipalname [nameid-format:emailAddress]</code>

Table 10: Additional Claims

Claim name	Value
<code>http://schemas.microsoft.com/ws/2008/06/identity/claims/groups</code>	<code>user.groups [SecurityGroup]</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress</code>	<code>user.mail</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname</code>	<code>user.givenname</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name</code>	<code>user.userprincipalname</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname</code>	<code>user.surname</code>

- Copy or keep available for reference the the Claim Names/Namespace URLs for entering Role Attribute Values in the Morpheus Identity Source Integration.
- In Azure SSO config, if one has not been generated, select `Create new certificate` to generate a new SAML Signing Certificate.
- Enter a valid email address to receive certificate expiration notifications (these are not Morpheus-generated email).
- In Azure SSO config, select `Configure {AD App Name}`
- In the `Configure sign-on` pane, copy the following:
 - SAML Single Sign-On Service URL** This will be used for the LOGIN REDIRECT URL in the Morpheus Identity Source Integration settings
 - Sign-Out URL** This will be used for the LOGOUT REDIRECT URL in the Morpheus Identity Source Integration settings

- Click on the **SAML XML Metadata** link, open the xml file, and copy the key between the **<X509Certificate>** and **</X509Certificate>**. This will be used for the *Public Key* value in the SAML RESPONSE section of the Morpheus Identity Source Integration settings

Example Key (this key is an example and is not valid):

```
MIIC8ECCAdigAwIBAgIQEOZX1Nx5wY9Dc6Ow1sKEMzANBgkqhkiG9w0BAQsFADA0MTIwMAYDVQQDEylNaWNyb3NvZnQ
→V6GcBpRkoxJd0DLbhubwd0kp65LD9IIh5PUY2ohBHvrFAy3SZ04mXoH7LWvY3oNrxaNAksbYF6phOkONf/
→XeTdZor14xdGnTuD9zRqPsJHHisyfFBUG/CxYxzO6w9fAPzJGLzc0Y7o5lMW2OjINaQI4R/
→pqp3qw+nYf7DXSzY6tf1Sspk64jffZDt1jSVjD7upMITKPeOCRmeBUCnebJzwXqFBO7l4Vf5gloEJyftT7Wpr4VVmo
→pH6xzQVRz0GZQpol9ViQJJbJJqhLm4LjWT9VU2lYqdi0NdgtK7QthZo4J0ZFdUG6qfFTfPKqVn0AEHxiM4JWxfigz
→y56+ksYSRP87XdOcVvTftHYmQnDOfoqKrpqMK7LtmsEwqc7rKX7nTCenZnBEOCFDBVH4QEzMrAznpEPdJnQs9nJZB
→sec
```

9. Save the SSO config in Azure AD app and return to Morpheus

Edit the existing Azure AD SAML Integration

Now that we have the required information, we can finalize the Azure AD SAML Integration in Morpheus

1. Edit the existing Azure AD SAML Integration created in the first step and populate the following:

LOGIN REDIRECT URL Add the SAML Single Sign-On Service URL copied from Azure SSO config.

LOGOUT REDIRECT URL Add the Sign-Out URL copied from Azure SSO config.

SAML RESPONSE Set to “Validate Assertion Signature”, then in the “Public Key” field enter the Public Key value we discussed in the last section

GIVEN NAME ATTRIBUTE NAME (May have to click “show” to see hidden SAML Assertion Attribute Names fields)

Enter the givenname Namespace url from Azure SSO config: <http://schemas.xmlsoap.org/ws/2005/05/identity/claims>

SURNAME ATTRIBUTE NAME Enter the emailAddress Namespace url from Azure SSO config: <http://schemas.xmlsoap.org/ws/2005/05/identity/claims>

EMAIL ATTRIBUTE NAME (May need to scroll down within the SAML Assertion Attribute Names section see this field)

Enter the surname Namespace url from Azure SSO config: <http://schemas.xmlsoap.org/ws/2005/05/identity/claims>

Configure Role Mappings

Role mappings will map Azure AD Groups to Morpheus Roles. Azure AD users will be assigned Roles in Morpheus upon signing in based on their Group Membership in Azure AD.

Important: Use an Azure Groups Object ID, not Group name, when entering Role Mappings. Example: 7626a4a2-b388-4d9b-a228-72ce9a33bd4b

DEFAULT ROLE Role a Azure AD user will be assigned by default upon signing in to Morpheus using this Identity Source.

REQUIRED AZURE AD GROUP OBJECT ID Object ID of Azure AD Group a user must be a member of to be authorized to sign in to Morpheus. Users not belonging to this Group will not be authorized to login to Morpheus. This field is optional, and if left blank, any user from the Azure AD App will be able to sign in to Morpheus and will be assigned the Default Role if no Role Mappings match AD Group membership.

GROUP ASSERTION ATTRIBUTE NAME Enter `http://schemas.microsoft.com/ws/2008/06/identity/claims/groups` for Azure AD SSO

Additional Role Mappings The existing Roles in Morpheus will be listed. To map a Morpheus Role to an Azure AD Group, enter the Object ID of the desired Azure AD Group in the *Role Attribute Value* field for the corresponding Morpheus Role.

Important: Use an Azure Groups Object ID, not Group name, when entering Role Mappings. Example: 7626a4a2-b388-4d9b-a228-72ce9a33bd4b

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).


Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

Once populated, select *SAVE CHANGES* and the SAML identity source integration will be added. The Identity Source can be edited anytime to deactivate or change Role Mappings or other values.

Note: If Role mappings are edited after Azure AD SSO users have signed into Morpheus, currently logged in users will need to log out of Morpheus for the new Role mappings to take effect, when applicable.

Signing In to Morpheus

When there is an active SAML/Azure AD SSO Identity Source Integration, a new button will appear on the Morpheus login page with the name of the Identity Source Integration as the button title. Example: *ADFS*. Another button titled “USERNAME AND PASSWORD” is also added for Morpheus account authentication outside of an Identity Source.



- **SAML/Azure AD SSO users can log into Morpheus by clicking the SAML button** This will redirect the User to Azure AD app sign in url. If they are currently signed into Azure and authorized, the user will be instantly signed into Morpheus.
- Local Morpheus users can select “USERNAME AND PASSWORD” to sign in with their local credentials as before.

Note: If no local users other than the System Admin have been created, “USERNAME AND PASSWORD” option will not be displayed, only the SAML option.

Okta

Overview

Morpheus allows users to integrate an Okta deployment for user management and authentication. In Morpheus, identity sources are added on a per-Tenant basis and Morpheus allows you to map Okta user groups to Morpheus user groups. User accounts are automatically created with matching metadata and role permissions when users are authenticated.

Adding an Okta Integration

1. Navigate to Administration -> Tenants
2. Select a Tenant
3. Select *IDENTITY SOURCES*
4. Select + *IDENTITY SOURCE*
5. Choose TYPE: “Okta”
6. Populate the following, then select *SAVE CHANGES*:

Name Unique name for authentication type

Description A description for your new Okta Identity Source

Okta URL Your Okta URL

Administrator API Token Your Okta Administrator API Token

Required Group The Okta group that users must be in to have access (optional)

Default Role The default role a user is assigned if no group is listed under an Okta user that maps within the Morpheus Role Mappings section

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

Now, allowed Okta users can log into Morpheus via their Okta credentials and a user will be automatically generated within Morpheus with matching metadata and mapped Role permissions.

Note: If you’ve created multi-tenant roles, these will also appear here and can be mapped to Okta user groups allowing you to map users to equivalent user groups in Morpheus.

OneLogin

Adding OneLogin Identity Source Integration

1. Navigate to Administration -> Tenants
2. Select the Tenant to add the Identity Source Integration
3. Select *IDENTITY SOURCES*
4. Select + *IDENTITY SOURCE*
5. Enter the following:

TYPE OneLogin

NAME

Name of the Identity Source Integration in Morpheus

DESCRIPTION Optional Description of the Identity Source

ONELOGIN SUBDOMAIN

example: morpheus-dev

Warning: Please verify the subdomain carefully. An invalid subdomain will cause authentication attempts by OneLogin users to fail.

ONELOGIN REGION Specify US or EU region

API CLIENT SECRET OneLogin API Client Secret from the Settings - API section in OneLogin portal

API CLIENT ID OneLogin API Client ID from the Settings - API section in OneLogin portal

REQUIRED ROLE Enter a role if OneLogin users logging into morpheus must have at least this OneLogin role to gain access to Morpheus.

DEFAULT ROLE The default Morpheus Role applied to users created from OneLogin Integration if no other role mapping is specified below

ROLE MAPPINGS Existing Morpheus Roles will be listed with fields to enter OneLogin Roles to map to. Users with OneLogin roles matching the role mappings will be assigned the appropriate Role(s) in Morpheus when signing in.

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

6. Select *SAVE CHANGES* and the OneLogin Integration will be added.

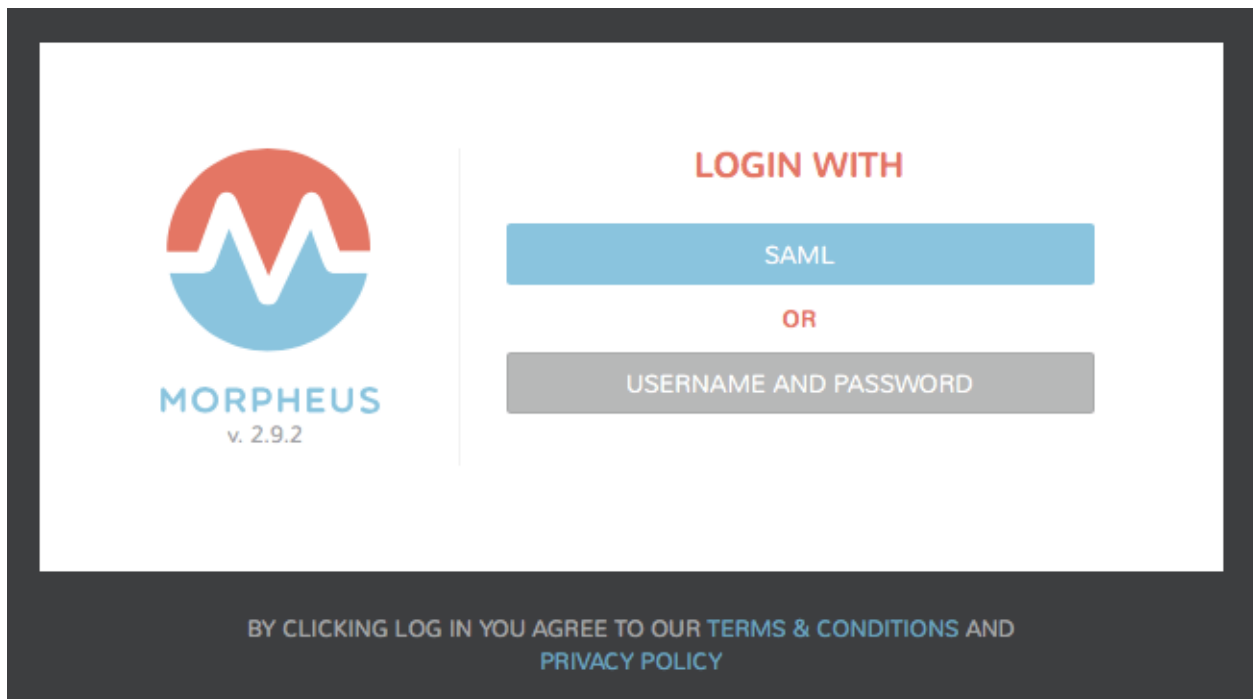
Users can now login to Morpheus with OneLogin credentials. The first Login will create a user in Morpheus matching the Username, email and Password from OneLogin. If a REQUIRED ROLE is specified in the Identity Source settings, only users with that Role in OneLogin will be able to login to Morpheus.

Important: OneLogin users will not authenticate in Morpheus if there is an existing Morpheus User with matching username or email address.

SAML Integration

Overview

The Morpheus SAML identity source integration allows customers to add user SSO to Morpheus , authenticated by external login SAML providers.



Adding a SAML Integration

To add a SAML integration:

1. Navigate to Administration -> Tenants
2. Select a tenant.
3. Select IDENTITY SOURCES in the Tenant detail page
4. Select + *ADD IDENTITY SOURCE*.
5. Select SAML (external login) from the TYPE field
6. Add a Name and optional Description for the SAML integration

NEW IDENTITY SOURCE

Identity Source

TYPE: SAML

NAME:

DESCRIPTION:

SAML Configuration

LOGIN REDIRECT URL:

☐ Do not include SAMLRequest parameter

LOGOUT POST URL:

SIGNING PUBLIC KEY:

☒ Do not validate SAMLResponse signatures

ADVANCED VALIDATION OPTIONS Show/Hide

Role Attribute Value Show/Hide

Role Mappings

DEFAULT ROLE: Developer

ROLE ATTRIBUTE NAME: memberOf

REQUIRED ROLE ATTRIBUTE VALUE:

DEVELOPER: Role Attribute Value

DIPESH-008USER: Role Attribute Value

IMAGE DEVELOPER: Role Attribute Value

LIPSCOMBE USER: Role Attribute Value

MORPHEUS ADMIN: Role Attribute Value

MORPHEUS SALES: Role Attribute Value

SAVE CHANGES

There are 3 sections with fields that need to be populated depending on the desired configuration:

- SAML Configuration
- Role Mappings
- User Attribute Names

SAML Configuration

LOGIN REDIRECT URL This is the SAML endpoint Morpheus will redirect to when a user signs into Morpheus via SAML.

LOGOUT POST URL The url morpheus will post to when a SAML user log out of Morpheus to log out of the SAML provider as well.

SIGNING PUBLIC KEY Add the X.509 Certificate public key from the SAML provider.

Role Mappings

DEFAULT ROLE Role a saml user will be assigned by default when no role is mapped

ROLE ATTRIBUTE NAME The name of the attribute field that will map to morpheus roles, such as MemberOf

REQUIRED ROLE ATTRIBUTE VALUE Role attribute value that a user must be assigned/a member of to be authorized, such as group or role in the SAML SP.

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

The rest of the Role Mapping Fields will be the existing Roles in morpheus with a Role Attribute Value field.

User Attribute Names

GIVEN NAME ATTRIBUTE NAME SAML SP field value to map to Morpheus user First Name

SURNAME ATTRIBUTE NAME SAML SP field value to map to Morpheus user Last Name

EMAIL ATTRIBUTE NAME SAML SP field value to map to Morpheus user email address

EDIT IDENTITY SOURCE



Identity Source

TYPE	SAML [Beta]
ACTIVE	Yes
NAME	SAML
DESCRIPTION	onelogin SAML

SAML Configuration

LOGIN REDIRECT URL	https://morpheusdata-dev.onelogin.com/trust/saml2/http- <input type="checkbox"/> Do not include SAMLRequest parameter
LOGOUT POST URL	https://morpheusdata-dev.onelogin.com/trust/saml2/http-
SIGNING PUBLIC KEY	MIIEFzCCAv+gAwIBAgIUayYdMuoXBTGcalAARanxhRJwwtQwDQYJKoZIhvcNAQEF

Role Mappings

DEFAULT ROLE	System Admin
ROLE ATTRIBUTE NAME	MemberOf
REQUIRED ROLE ATTRIBUTE VALUE	dev
LEGACY ACCOUNT ADMIN	Role Attribute Value

User Attribute Names

Show/Hide

GIVEN NAME ATTRIBUTE NAME	firstName
SURNAME ATTRIBUTE NAME	lastName
EMAIL ATTRIBUTE NAME	Email

DEACTIVATE

DELETE

SAVE CHANGES

Once populated, select SAVE CHANGES and the SAML identity source integration will be added.

In the Identity Sources section, important information for configuration of the SAML integration is provided. Use the SP ENTITY ID and SP ACS URL for configuration on the external login SAML provider side.

- SP ENTITY ID
- SP ACS URL*
- IDP LOGIN REDIRECT URL
- IDP LOGOUT POST URL
- SP METADATA

Master Account

Search		+ ADD IDENTITY SOURCE	
TYPE	NAME	DETAILS	ACTIVE
SAML [Beta]	SAML <i>onelogin SAML</i>	<div>SP ENTITY ID: https://someip.com/saml/CDWPjmZt</div> <div>SP ACS URL: https://someip.com/saml/CDWPjmZt</div> <div>IDP LOGIN REDIRECT URL: https://someip.com/saml/CDWPjmZt</div> <div>IDP LOGOUT POST URL: https://someip.com/saml/CDWPjmZt</div> <div>SP METADATA: VIEW</div>	Yes

Sample Metadata code output:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?><EntityDescriptor entityID=
  ↪ "https://someip.com/saml/CDWPjmZt" xmlns="urn:oasis:names:tc:SAML:2.0:metadata">
  ↪ <SPSSODescriptor AuthnRequestsSigned="false" WantAssertionsSigned="true"
  ↪ protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol"><NameIDFormat>
  ↪ urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified</NameIDFormat>
  ↪ <AssertionConsumerService index="0" isDefault="true" Binding=
  ↪ "urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Location="https://someip.com/
  ↪ externalLogin/callback/CDWPjmZt" /></SPSSODescriptor></EntityDescriptor>
```

Note: Different SAML providers will have different field names and requirements. A onelogin SAML Test Connector (IdP w/attr) was used for the example integration this article.

Onelogin SAML SSO

For Onelogin SAML integration, the following fields are mapped:

- LOGIN REDIRECT URL : SAML 2.0 Endpoint (HTTP)
- LOGOUT POST URL : SLO Endpoint (HTTP)
- SIGNING PUBLIC KEY : X.509 Certificate
- SP ENTITY ID: ACS (Consumer) URL Validator
- SP ACS URL: ACS (Consumer) URL

Plans & Pricing

Overview

The Plans & Pricing page displays a list of all of your available service plans. From the service plans page you will be able to Create, Edit, and Delete service plans, as well as review basic plan details. The list of plans displayed on this page displays the plan name, description, Instances layout, memory, storage, and cost, as well as an action column to edit and delete. A default set of Service Plans are created in Morpheus. They provide a means to set predefined tiers on memory, storage, cores, and cpu. Price tables can also be applied to these so estimated cost per virtual machine can be tracked as well as pricing for customers.

Service Plans

Create Service Plan

To create service plan

1. Select the Administration link in the navigation bar.
2. Select the Plans & Pricing link in the sub navigation bar.
3. Click the Create Service Plan button.
4. From the New Service Plan wizard, input:
 - Name
 - Code used as a unique identifier in the API and CLI.
 - Storage size in megabytes.
 - Memory size in megabytes.
 - Cost is internal cost of plan.
 - Price is what the service offering will be priced at.
 - Instance Types that will be associated with this plan.
 - Click the Save Changes button to save.

Edit Service Plan

By default, these options are fixed sizes but can be configured for dynamic sizing. A service plan can be configured to allow a custom user entry for memory, storage, or cpu. To configure this, simply edit an existing Service Plan. These all can be easily managed from the Admin -> Service Plans section.

To edit service plan:

1. Select the Administration link in the navigation bar.
2. Select the Plans & Pricing link in the sub navigation bar.
3. Click the Edit pencil icon on the row of the plan to edit.
4. Edit the following Edit Service Plan.
5. Click the Save Changes button to save.

Delete Service Plan

To delete a service plan

1. Select the Administration link in the navigation bar.
2. Select the Plans & Pricing link in the sub navigation bar.
3. Click the Delete trashcan icon on the row of the plan to delete.
4. Confirm

Pricing

Price Sets

Price Sets combine Prices and then attach to Plans. Prices must be created prior to creating Price Sets, but it is recommended to review the Price Set Type options prior to creating Prices.

Price Unit Select the Price Unit to use for the Price Set.

- Minute
- Hour
- Day
- Month
- Year
- Two Year
- Three Year
- Four Year
- Five Year

Note: Only Prices configured with matching Price Units can be used in a Price Set.

Note: Month is equivalent to 30 days by default. For AWS, month is 30.5 days. For Azure, month is 30.4 days.

Type Price Set types determine which Prices are available to make up the set. This selection will filter the values returned in the Prices field at the bottom of the modal.

Note: It's helpful to make note of the Prices options below before creating Price Sets.

- **Everything:** 'Everything' price sets require 1 or more 'Everything' price types and may include 'Platform' or 'Software' price types.
- **Compute + Storage:** 'Compute + Storage' price sets require at least one of each 'Memory', 'CPU', and 'Disk Only' price types and may include 'Platform' or 'Software' price types.
- **Component:** 'Component' price sets require at least one of each 'Memory', 'Cores', 'CPU', and 'Storage' price types and may include 'Platform' or 'Software' price types.

Apply Price Changes to Usage If marked, when saving a Price Set (new Price Set or saving changes to an existing one), usage records will be restarted for servers affected by the pricing change.

Prices Search for and select Prices to be added to the Price Set. One of each Price Type required for the Price Set Type selected must be added for the Price Set to save.

Prices

Price Types

- Everything: One price for all resources Storage, CPU, Memory, and Disks
- Memory + CPU
- Memory Only (per MB)
- Cores Only (per core)
- Disk Only (per GB)
- Platform
- Software
- Datastore (per GB)

Price Units

- Minute
- Hour
- Day
- Month
- Year
- Two Year
- Three Year
- Four Year
- Five Year

Currency

- AUD
- CHF
- DKK
- EUR
- GBP
- IDR
- ILS
- MAD
- NOK
- NZD

- ROL
- SEK
- TRL
- USD
- XAF
- XCD
- XOF
- XPF
- ZAR (South African Rand)

Cost The base cost of the resource(s). The Price will match the Cost unless a Price Adjustment is added.

Price Adjustment

- **None:** Default, no markup added and Price will match Cost
- **Fixed Markup:** A fixed amount added to the Cost. Price will equal $\text{Cost} + \text{Markup}$.
- **Percentage Markup:** Adds a percentage markup to Cost. Price equals $\text{Cost} + (\text{Cost} \times \text{Markup} \%)$
- **Custom Price:** Sets a Price independent from the Cost. If the Cost changes, a Custom Price will not.

Price A computed value of the final price including the cost plus any applicable markup.

Apply Price Changes to Usage If marked, when saving a Price Set (new Price Set or saving changes to an existing one), usage records will be restarted for servers affected by the pricing change.

Roles

Overview

Within Morpheus is a wide array of role based access control capabilities. These roles can be managed within the Admin -> Roles section of the Morpheus UI as well as through the API or CLI. They are designed to be robust enough to fit within a wide array of enterprise and managed service provider scenarios so they can be a bit hard to grasp at first, but should make sense once a few simple concepts are explained. There are two types of roles within Morpheus called Tenant and User based roles. Both sets of roles allow restrictions to be imposed on a user at the feature access level. Entire sections within the appliance UI can be hidden based on the specified access levels for features within Morpheus. Features have different access scopes that can be selected from and can range depending on the specific feature. The most common scope set involves none, read, and full. Instance Type access is also common among both role types which allow the administrator to restrict which service catalog items they are allowed to provision within Morpheus .

There are several handy tricks for creating new roles within Morpheus and users can be assigned more than one role. When a user is assigned more than one role, permissions are granted by the role with the highest level of scope access. This allows roles to be built with small subsets of features and combined to grant different individuals relevant permission control.

Note: Feature access control not only applies to the Morpheus UI but also applies to the public developer API. It is sometimes necessary to logout and back in for changes to a users feature access level to be respected.

Role Types

Tenant Roles

A Tenant based role (formerly called an Account based role) is used to ensure access control enforcement across an entire tenant with many sub-users. This allows the subtenant to manage their own set of internal user based roles without worrying master tenant involvement in setting them up. The master tenant is the only tenant able to create and manage these types of roles. When editing a Tenant, a singular tenant role can be assigned to the account. Users within the tenant can be assigned roles but those user based roles will never be able to supersede the level of access granted by the tenant role. This allows a super administrator the ability to restrict access at the department or organization level without having to worry about per user access control within said tenant.

Tenant roles also have an additional section not in User based roles related to Cloud Access. Cloud Access allows the master tenant the ability to assign cloud integration resources to specific subtenants or groups of subtenants. An example would be granting access to a specific VMware cluster only to a subset of tenants using the tenant based role control.

User Roles

User roles can be created by any tenant given permission at the tenant role level. These allow tenants to manage their own sets of users and their levels of access. They also allow tenants to control which users have access to specific “Groups” for provisioning into within Morpheus. Groups are not cross tenant and therefore need to be controlled within the individual tenant in Morpheus.

Master tenants are able to create a special type of user role called a multi-tenant user role. A multi-tenant user role is copied / duplicated down to all subtenants within Morpheus. These can be viewed as pre-canned role templates available to new tenants when their account is first created. Any changes made to the main role are propagated down to the subtenants version of the shared role so long as the subtenant has not previously adjusted/changed that role. The moment a subtenant makes adjustments to the shared role within their account, it is unlinked from the parent role and treated entirely independently.

Another note about user roles is that when a user role is copied down to a subtenant, the permission scopes cannot supersede the tenants assigned tenant role. If they do they are automatically downgraded when propagated to the specific tenant. Any changes made to the tenant role will automatically ensure roles within the tenant are downgraded appropriately.

Multi-Tenant User Role Lock

As discussed above, multi-tenanted user roles are made available within all subtenants as ‘canned’ user role sets. Master tenant administrators can prevent changes to these ‘canned’ user roles by marking the box labeled ‘MULTITENANT LOCKED’ when creating or editing the role. In addition to preventing subtenant administrators from modifying permissions of these roles within their subtenancy, this option also ensures master tenant administrators can propagate new changes to that role. Modification of the role by the subtenant (if allowed) breaks the link back to the master tenant and the copy of the role within the subtenant will become its own unlinked role.

Note: Multi-tenant role lock applies only to permission sets on the ‘FEATURE ACCESS’ tab. Permissions in the ‘GROUP ACCESS’, ‘INSTANCE TYPE ACCESS’, and ‘BLUEPRINT ACCESS’ tabs are not locked. Similarly, changes made to the role on these tabs in the master tenant are not synced down.

Roles and Identity Sources

It is very common for large Enterprises to have an existing identity source that they would like to plug in to Morpheus for authentication. This includes services like LDAP, Active Directory, OKTA, Jump Cloud, One Login, and SAML. When using these services it becomes important to configure a role mapping between the Morpheus role assignments to the equivalent identity source groups/roles the user belongs to. This is configurable within the identity source management UI. Sections are provided allowing things like LDAP groups to be directly mapped to specific roles within Morpheus. If a user matches more than one LDAP/role group then both sets of roles are applied to the user automatically. Configuring Identity Sources is done in Tenant management or user management in Administration > Tenants or Administration > Users, and has to be configured on a per-tenant basis. Additionally, administrators may opt to lock users to their mapped role in Morpheus or keep the roles unlocked to manually administer roles in one-off scenarios.

Role Permissions

Note: Permission options for sub-tenant user roles will only list options permitted by the Tenant role applied to the sub-tenant. Sub-Tenant user roles permissions cannot exceed permissions set by the overriding Tenant Role.

User Role Permission Sections

FEATURE ACCESS Controls Tenant and User access level for sections and features in Morpheus.

GROUP ACCESS Controls User access level for Groups. (Groups are not Multi-Tenant.)

CLOUD ACCESS Controls Sub-Tenant access level for Master Tenant publicly visible Clouds.

INSTANCE TYPE **User only has access to Objects they have created/own.** Controls Tenant and User access level for Instance Types.

BLUEPRINT ACCESS Controls Tenant and User access level for Blueprints during App provisioning.

Feature Access Permissions

Feature Access settings control permissions for sections and objects in Morpheus. Permission options include:

None Hidden or inaccessible for user

Read User can view but cannot edit or create

Full User has full access

User User can access Objects they have created or own.

Group User can access Objects assigned to or shared with Groups the User has access to.

Remote Console: Provisioned Remote Console tab will only appear after instance is successfully provisioned.

Remote Console: Auto Login RDP and SSH only, controls if user is auto-logged in to Remote Console or presented with login prompt.

Role Mappings Gives User Access to Role Mappings config in `/admin/roles` for configuring Identity Source Role Mappings without providing Access to other Identity Source configuration settings.

Permission Name	Permission Options	Feature Access
Admin: Appliance Settings	None, Full	Allows or disallows access to the Appliance
Admin: Backup Settings	None, Full	Allows or disallows access to Administration
Admin: Environment Settings	None, Full	Allows or disallows access to the Environment
Admin: Guidance Settings	None, Full	Allows or disallows access to the Guidance t
Admin: Health	None, Read	Determines access to the Operations > Health
Admin: Identity Source	None, Role Mappings, Full	Allows or disallows access to create, edit, or
Admin: Integrations	None, Read, Full	This allows or disallows full or read access to
Admin: License Settings	None, Full	Allows or disallows access to the Licenses ta
Admin: Log Settings	None, Full	Allows or disallows access to the Administra
Admin: Message of the day	None, Full	Allows or disallows access to create and edit
Admin: Monitoring Settings	None, Full	Allows or disallows access to Administration
Admin: Plugins	None, Full	Allows or disallows access to the Plugins tab
Admin: Policies	None, Read, Full	This setting determines the level of access to
Admin: Provisioning Settings	None, Full	Allows or disallows access to the Settings ta
Admin: Roles	None, Read, Full	This setting determines access to the Admini
Admin: Service Plans	None, Read, Full	This setting determines access to Administra
Admin: Tenant	None, Read, Full	This setting determines access to the Admini
Admin: Tenant - Impersonate Users	None, Full	This setting allows or disallows access to imp
Admin: Users	None, Read, Full	This setting determines access to the Admini
Admin: Whitelabel Settings	None, Full	Allows or disallows access to the Whitelabel
API: Billing	None, Read, Full	Allows or disallows access to invoices and pr
API: Execution Request	None, Full	Allows or disallows access to an API endpoi
Backups	None, View, Read, User, Full	Determines access to the Backups section of I
Backups: Integrations	None, Read, Full	Determines access to the Backups > Integrati
Infrastructure: Boot	None, Read, Full	Determines access to the Integrations > Boot
Infrastructure: Certificates	None, Read, Full	Determines access to the SSL Certificates tab
Infrastructure: Clouds	None, Read, Full	Determines access to the Infrastructure > Clo
Infrastructure: Clusters	None, Read, Group, Full	Determines access to the Infrastructure > Clu
Infrastructure: Groups	None, Read, Full	Determines access to the Infrastructure > Gro
Infrastructure: Hosts	None, Read, Full	Determines access to the Infrastructure > Ho
Infrastructure: Keypairs	None, Read, Full	Determines access to the Key Pairs tab on the
Infrastructure: Load Balancers	None, Read, Full	Determines access to the Infrastructure > Lo
Infrastructure: Network Domains	None, Read, Full	Determines access to the Domains tab on the
Infrastructure: Network Firewalls	None, Read, Full	Determines access to the Firewall tab on app
Infrastructure: Network Integration	None, Read, Full	Determines access to the Integrations tab on
Infrastructure: Network IP Pools	None, Read, Full	Determines access to the IP Pools tab on the
Infrastructure: Network Proxies	None, Read, Full	Determines access to the Proxies tab on the I
Infrastructure: Network Router DHCP Pool	None, Read, Full	Determines access to the DHCP tab on the d
Infrastructure: Network Router Firewalls	None, Read, Full	Determines access to Firewall tabs on Router
Infrastructure: Network Router Interfaces	None, Read, Full	Determines access to Interfaces tabs on Rout
Infrastructure: Network Router NAT	None, Read, Full	Determines access to the NAT tab on Router
Infrastructure: Network Router Redistribution	None, Read, Full	Determines access to Route Redistribution ta
Infrastructure: Network Router Routes	None, Read, Full	Determines access to Routing tabs on Router
Infrastructure: Network Routers	None, Read, Group, Full	Determines access to the Routers tab on the I
Infrastructure: Networks	None, Read, Group, Full	Determines access to the Infrastructure > Ne
Infrastructure: Policies	None, Read, Full	Determines access to the Policies tabs on the
Infrastructure: Security Groups	None, Read, Full	Determines access to the Security Groups tab
Infrastructure: State	None, Read, Full	Determines access to the power state toggle o
Infrastructure: Storage	None, Read, Full	Determines access to the Infrastructure > Sto

Permission Name	Permission Options	Feature Access
Infrastructure: Storage Browser	None, Read, Full	Determines file browsing access to buckets and objects
Infrastructure: Trust Integrations	None, Read, Full	Determines access to the Integrations tab of the Infrastructure section
Integrations: Ansible	None, Full	Determines access to Ansible integrations on the Integrations tab
Logs	None, Read, User, Full	Determines level of access to the Logs section
Monitoring	None, Read, User, Full	Determines level of access to the Monitoring section
Operations: Activity	None, Read	Determines access to the Activity and History tabs
Operations: Alarms	None, Read, Full	Determines access to the Alarms tab in the Operations section
Operations: Analytics	None, Read, Full	Determines access to the Operations > Analytics tab
Operations: Approvals	None, Read, Full	Determines access to the Operations > Approvals tab
Operations: Budgets	None, Read, Full	Determines access to the Operations > Budgets tab
Operations: Dashboards	None, Read	Determines access to the Operations > Dashboards tab
Operations: Guidance	None, Read, Full	Determines access to the Operations > Guidance tab
Operations: Invoices	None, Read, Full	Determines access to the Invoices tab in the Operations section
Operations: Reports	None, Read, Full	Determines access to the Operations > Reports tab
Operations: Usage	None, Read, Full	Determines access to the Usage tab on the Operations section
Operations: Wiki	None, Read, Full	Determines access to the Operations > Wiki tab
Projects	None, Read, Full	Determines access to Projects through Morpheus
Provisioning: Administrator	None, Full	When editing an Instance (Provisioning > Instance)
Provisioning: Advanced Node Type Options	None, Full	This allows or disallows access to the “Extra Options” section
Provisioning: Allow Force Delete	None, Full	This allows or disallows access to the “Force Delete” checkbox
Provisioning: Apps	None, Read, User, Full	Determines access to the Provisioning > Apps tab
Provisioning: Automation Integrations	None, Read, Full	Determines access to the Integrations tab on the Provisioning section
Provisioning: Blueprints	None, Read, Full	Determines access to the Provisioning > Blueprints tab
Provisioning: Blueprints - ARM	None, Provision, Full	Determines access to ARM-type Blueprints on the Provisioning > Blueprints tab
Provisioning: Blueprints - CloudFormation	None, Provision, Full	Determines access to CloudFormation-type Blueprints on the Provisioning > Blueprints tab
Provisioning: Blueprints - Helm	None, Provision, Full	Determines access to Helm-type Blueprints on the Provisioning > Blueprints tab
Provisioning: Blueprints - Kubernetes	None, Provision, Full	Determines access to Kubernetes-type Blueprints on the Provisioning > Blueprints tab
Provisioning: Blueprint - Terraform	None, Provision, Full	Determines access to Terraform-type Blueprints on the Provisioning > Blueprints tab
Provisioning: Clone Instance	None, Full	Determines access to the Clone Instance selection
Provisioning: Deployment Integrations	None, Read, Full	Determines access to the Integrations tab on the Provisioning section
Provisioning: Deployments	None, Read, Full	Determines access to the Deployments tab on the Provisioning section
Provisioning: Execute Script	None, Full	Determines access to the Run Script and App button
Provisioning: Execute Task	None, Full	Determines access to the Run Task selection
Provisioning: Execute Workflow	None, Full	Determines access to the Run Workflow selection
Provisioning: Import Image	None, Full	Determines access to the Import as Image and Import as Template buttons
Provisioning: Instances	None, Read, User, Full	Determines access to the Provisioning > Instances tab
Provisioning: Job Executions	None, Read	Determines access to the Job Executions tab
Provisioning: Jobs	None, Read, Full	Determines access to the Jobs tab on the Provisioning section
Provisioning: Library	None, Read, Full	Determines access to the Provisioning > Library tab
Provisioning: Scheduling - Execute	None, Read, Full	Determines access to the Execute Scheduling tab
Provisioning: Scheduling - Power	None, Read, Full	Determines access to the Power Scheduling tab
Provisioning: Service Mesh	None, Read, User, Full	Determines access to the Provisioning > Service Mesh tab
Provisioning: Tasks	None, Read, Full	Determines access to the Tasks, Workflows, and Pipelines tabs
Provisioning: Tasks - Script Engines	None, Full	Determines access to the Tasks tab of the Provisioning section
Provisioning: Thresholds	None, Read, Full	Determines access to the Scale Thresholds tab
Provisioning: Virtual Images	None, Read, Full	Determines access to the Provisioning > Virtual Images tab
Remote Console	None, Provisioned, Full	Determines access to the console on a Host or Container
Remote Console: Auto Login	No, Yes	This allows or disallows the ability to automatically login to the console
Security: Scanning	None, Read, Full	Determines access to the Security Packages section

Permission Name	Permission Options	Feature Access
Service Catalog: Catalog	None, Full	Determines access to the Catalog page of the
Service Catalog: Dashboard	None, Read	Determines access to the Dashboard page of the
Service Catalog: Inventory	None, Full	Determines access to the Inventory page of the
Snapshots	None, Read, Full	Determines access to the “Create Snapshot”
Tools: Archives	None, Read, Full	Determines access to the Tools > Archives page
Tools: Cypher	None, Read, User, Full, Full Decrypt	Determines access to the Tools > Cypher page
Tools: Image Builder	None, Read, Full	Determines access to the Tools > Image Builder
Tools: Kubernetes	None, Read, User, Full	Allows for the management of Kubernetes cl
Tools: Migrations	None, Read, Full	Determines access to the Tools > Migrations
Tools: Self Service	None, Read, Full	Determines access to the Tools > Self Service
None - No Permissions	None	When all permissions are set to None, the fol

Creating Roles

User Roles

User Roles can be single or multitenant. A Multitenant User Role is automatically copied into all existing subtenants as well as placed into a subtenant when created. Useful for providing a set of predefined roles a Customer can use. The Multitenant Locked option prevent subtenant from modifying FEATURE ACCESS settings in the Role. Note Group, Instance Type and Blueprint Access settings will still be editable as Groups are unique per Tenant, and Instance and Blueprints can be a mix of unique and shared items.

Important: Multitenant Roles still need to be configured/managed by each subtenant, as Groups are unique per Tenant, and Instance and Blueprints can be a mix of unique and shared items.

Note: User Roles cannot exceed Tenant Role permissions. If a Multitenant User Role has higher permissions than the Tenant Role assigned to a subtenant, the Multitenant User Role permissions in that Tenant will automatically be reduced to match the Tenant Role permissions.

Create a Single Tenant User Role

1. In the Master Account, navigate to Administration -> Roles
2. Select + *CREATE ROLE*
3. Enter a name for the Role and optional Description
4. For TYPE, select “User Role”
5. Leave the “Multi-tenant Role” checkbox blank.
6. Optionally select an existing Role to copy in the COPY FROM ROLE dropdown. * This will configure the new Role with the same configuration as the selected role to copy. A new role that is not copied from another role will be generated with all permissions set to NONE.
7. Select *SAVE CHANGES*

After saving the Role will be created, and you will be redirected to the Roles Permissions settings.

Create a MultiTenant User Role

A Multitenant User Role is automatically copied into all existing subtenants as well as placed into a subtenant when created. Useful for providing a set of predefined roles a Customer can use. The Multitenant Locked option prevent subtenant from modifying FEATURE ACCESS settings in the Role. Note Group, Instance Type and Blueprint Access settings will still be editable as Groups are unique per Tenant, and Instance and Blueprints can be a mix of unique and shared items.

1. In the Master Account, navigate to Administration -> Roles
2. Select + *CREATE ROLE*
3. Enter a name for the Role and optional Description
4. For TYPE, select "User Role"
5. Optionally select an existing Role to copy in the COPY FROM ROLE dropdown. * This will configure the new Role with the same configuration as the selected role to copy. A new role that is not copied from another role will be generated with all permissions set to NONE.
6. Select the MULTITENANT ROLE checkbox
7. Optionally select the MULTITENANT LOCKED checkbox * When enabled, the FEATURE ACCESS settings in the Role will not be editable by subtenants. Group, Instance Type and Blueprint Access settings will still be editable as Groups are unique per Tenant, and Instance and Blueprints can be a mix of unique and shared items.
8. Select *SAVE CHANGES*

After saving the Role will be created, and you will be redirected to the Roles Permissions settings.

Important: Multitenant Roles still need to be configured/managed by each subtenant, as Groups are unique per Tenant, and Instance and Blueprints can be a mix of unique and shared items.

Note: User Roles cannot exceed Tenant Role permissions. If a Multitenant User Role has higher permissions than the Tenant Role assigned to a subtenant, the Multitenant User Role permissions in that Tenant will automatically be reduced to match the Tenant Role permissions.

Tenant Roles

A Tenant Role sets the highest possible permissions for a Tenant. User Roles within that Tenant cannot exceed those of the Tenants assigned Tenant Role. Tenant Roles can be assigned to single or multiple Tenants, and do not apply to the Master Account.

To create a Tenant Role:

1. In the Master Account, navigate to Administration -> Roles
2. Select + *CREATE ROLE*
3. Enter a name for the Role and optional Description
4. For TYPE, select "Tenant Role"

5. Optionally select an existing Role to copy in the COPY FROM ROLE dropdown. * This will configure the new Role with the same configuration as the selected role to copy. A new role that is not copied from another role will be generated with all permissions set to NONE.
6. Select *SAVE CHANGES*

After saving, the Role will be created and you will be redirected to the Roles Permissions settings.

Users & User Groups

Users

Overview

The Users page displays a list of all Users. The following fields are surfaced for each User:

- Tenant
- Display Name
- Username
- Email
- Role

Users which are grayed out in the list are currently inactive and cannot log in. From the Actions menu in each User row, the option is given to Impersonate the User, Edit, or Remove the User.

In Morpheus 4.2.1 and higher, click on the hyperlinked Display Name of the User to see a page detailing their effective Role permissions. This is especially useful for Users in multiple Roles where it might otherwise be difficult to determine their exact rights. This page looks identical to a User Role create/edit page except none of the fields are editable. Edit the User Role permissions for the User if changes need to be made.

Note: Some User data created through an Identity Source integration (such as Active Directory) is not editable in Morpheus, as it is synced from the Identity Source.

Create User

Users can be created from *Administration -> Users* or *Administration -> Tenants -> (selected Tenant) -> Users tab*.

Note: Authorized Identity Source Users will be automatically created upon first sign in.

To create a User:

1. Navigate to either *Administration -> Users* or *Administration -> Tenants -> Select a Tenant*.
2. Select + *CREATE USER*.
3. From the New User Wizard input:

Username & Email

- First Name
- Last Name

- Username
- Email address

Receive Notifications Enable to receive Provisioning and Policy email notifications.

Roles Role(s) to be inherited by the user. If multiple roles are selected, the higher permission levels of one role will override the other role(s).

Password Password must contain at least one uppercase letter, one lowercase letter, a number, and a symbol.

Enabled If unchecked, the user will no longer be able to sign into Morpheus, but their user data will remain.

Password Expired If enabled, the User will be forced to create a new password upon next login. The expired password cannot be used again.

Linux Settings Creates a User with the supplied Username, Password and/or Key-pair on Linux Instances when “Create my User” is selected during provisioning, or a User Group is added to an Instance of which this Morpheus user is a member of.

Windows Settings Creates a User with the supplied Username, Password and/or Key-pair on Windows Instances when “Create my User” is selected during provisioning, or a User Group is added to an Instance of which this Morpheus user is a member of.

Important: Please ensure password entered is allowable by Windows.

Note: Instance Resource Limits for a user are now configured through [Policies](#)

1. Select *SAVE CHANGES*.

Edit User

User settings can be edited from *Administration -> Users*, *Administration -> Tenants -> Select a Tenant -> Users tab*, or from *User Settings*.

Note: Some User data from Users created via an Identity Source Integration such as Active Directory is not editable in Morpheus, as it is synced with the Identity Source.

To edit a User from the *Administration -> Users* Section:

1. Select the Administration link in the navigation bar.
2. Select the Users link in the sub navigation bar.
3. Click **ACTIONS** on the row of the user to edit.
4. Select **EDIT** in the ACTIONS dropdown.
5. Make changes.
6. Select *SAVE CHANGES*.

To edit a User from the *Administration -> Tenants -> Select a Tenant -> Users tab*:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.

3. Select a Tenant
4. Click **ACTIONS** on the row of the user to edit.
5. Select **EDIT** in the ACTIONS dropdown.
6. Make changes.
7. Select *SAVE CHANGES*.

User Settings

Additional settings for a User can be found in the User Settings section, including:

- User Photo
- Default Group
- Default Cloud
- API Access

To access User Settings:

1. Select your name in the header
2. Select *User Settings*

To edit the User you are currently logged in as from *User Settings*:

1. Select your name in the header
2. Select *User Settings*
3. Make changes.
4. Select *SAVE*.

API Access

API and CLI Access Tokens can be regenerated from the *User Settings* section.

To regenerate a CLI or API Access Token:

1. Select your name in the header
2. Select *User Settings*.
3. Select *API ACCESS* under the *Windows Settings* section.
4. Select *ACTIONS* for the Client ID the token will be generated for.
5. Select *Regenerate*.
6. Copy the Generated Access Token.

Important: The Access Token will be masked after User Setting are saved.

7. Select *SAVE*.

Delete User

To delete a User from the *Administration -> Users* Section:

```
#. Select the Administration link in the navigation bar.
#. Select the Users link in the sub navigation bar.
#. Select **ACTIONS** on the row of the user to delete.
#. Select **REMOVE** in the ACTIONS dropdown.
#. Confirm
```

To delete a User from the *Administration -> Tenants -> Select a Tenant -> Users* tab:

1. Select the Administration link in the navigation bar.
2. Select the Tenants link in the sub navigation bar.
3. Select a Tenant
4. Click **ACTIONS** on the row of the user to delete.
5. Select **REMOVE** in the ACTIONS dropdown.
6. Confirm

User Groups

Overview

User Groups can be selected during provisioning to add each group members credentials to the Instance. User Groups can be configured for sudo access and in Linux will assign Group members to a groupId in linux.

Creating User Groups

1. Navigate to *Administration -> Users*
2. Select the USER GROUPS tab.
3. Select + *CREATE USER GROUP*.
4. Enter the following:
 - NAME** Name of the User Group
 - DESCRIPTION** Optional User Group Description
 - SERVER GROUP** Name of the groupId to assign Group members to in linux.
 - SUDO ACCESS** Enable to give Group members sudo access
 - USERS** Search for and select existing Users to add to the User Group.
5. Select *SAVE CHANGES*.

Editing User Groups

1. Navigate to *Administration -> Users*
2. Select the **USER GROUPS** tab.
3. Select the **ACTIONS** dropdown next to the target User Group.
4. Select **EDIT**
5. Make changes, add or remove users from the group.
6. Select *SAVE CHANGES*.

Adding a User Group when Provisioning

1. When provisioning, in the **CONFIG** section expand the **USER** section.
2. Select an existing Group from the **USER GROUP** dropdown.
3. Users will be created for members in the selected User Group on the provisioned Instance(s).

Integrations

Integrations

Administration > Integrations

To add an integration select + *ADD* and choose your integration. Many Morpheus-supported integrations can be configured in this section, though not all. Some integrations, such as networking integrations, must be configured within their own areas of the application. The following integrations can be configured in this section:

- Chef
- Puppet
- Ansible
- Salt Masters
- Ansible Tower
- vRealize Orchestrator
- Microsoft DNS
- PowerDNS
- Route 53
- Git
- Github
- Docker Repositories
- Consul
- Jenkins
- ServiceNow
- Cherwell

- Remedy
- ACI
- Venafi

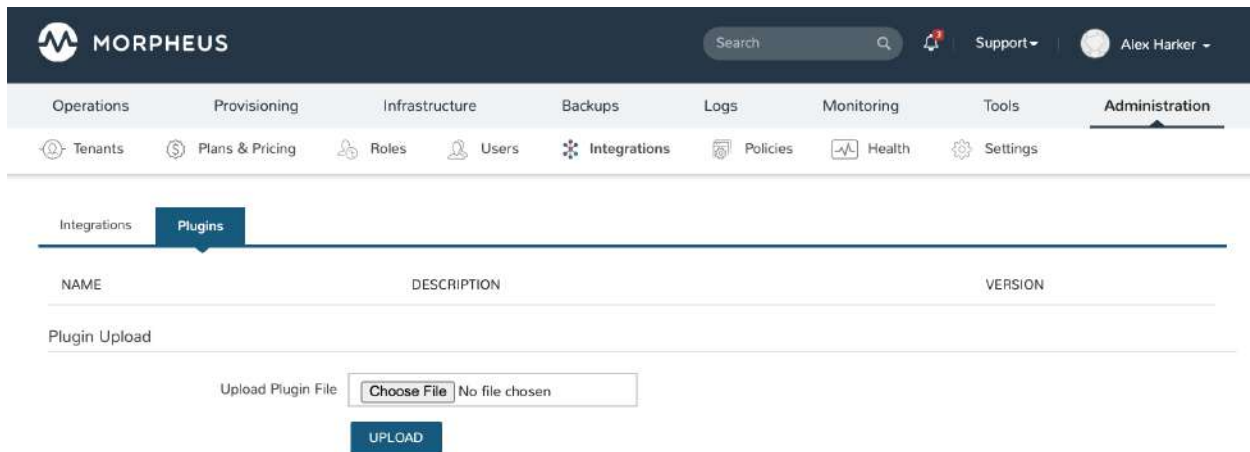
Please see the [Guides](#) for each specific integration type for more detailed information on setup steps and features supported by the integration.

Plugins

Overview

Morpheus is extendable with custom plugins for Task types, UI tabs, approvals, cypher, and more. Plugins are added from the Plugins tab of the Integration page under Administration (Administration > Integrations > Plugins). Simply browse for a local plugin file (.jar) to add it to the UI. Custom plugins can also be deleted by clicking on the trash can icon in the corresponding row.

Please visit the [Morpheus Developer Portal](#) for Plugin Architecture SDK documentation and help getting started with custom Plugin development.



Policies

Overview

Policies add governance, ease of use, cost-savings, and auditing features to Morpheus. Morpheus enables end users to create Policies scoped to Users, Roles, Groups, Clouds, Tenants and Global scoping to give Admins full control and governance over their environments! Policy generation is a role permission.

Policy Types

Backup Creation Disable or enable the ability to create a backup when provisioning an instance.

Backup Target A master account can determine storage provider options for backups with Backup Targets policies.

Budget Sets a maximum total combined price for all instances in the Group, Cloud, Tenant or owned by the User this policy is applied to.

Delayed Removal Delayed Removals allow for soft deletion of Instances and Apps. Instead of deleting immediately, Instances and Apps with a Delayed Removal policy applied will be shutdown upon deletion request and hidden by default from the ui. The Instance/App will then be in Pending Removal status.

Expiration Sets an expiration timeframe in days after which the Instance will be deleted. Extensions can be auto-approved or require approval immediately or after x amount of auto-extensions using Morpheus Approvals or an Approval Integration.

File Share Storage Quota Sets a Storage Quota for File Share usage (in GB) to scoped User, Role, Tenant or Global.

Host Name The name of the virtual machine. Pre-populates a fixed or editable name for Hosts and Virtual Machines using \${variable} naming patterns and/or text.

Hostname The `hostname` or `computer` name which is set in the OS and DNS. On some platforms, hostnames are restricted by length, spaces, and/or special characters. Pre-populates a fixed or editable name for hostnames/machine names using \${variable} naming patterns and/or text.

Instance Name Pre-populates a fixed or editable name for Instance Names using \${variable} naming patterns and/or text.

Max Containers Sets the max number of Containers for the Group or Cloud the Policy is added to.

Max Cores Sets the max number of total of Cores combined for Instances in the Group or Cloud the Policy is added to.

Max Hosts Sets the max number of total Hosts in the Group or Cloud the Policy is added to.

Max Memory Sets the max number of total of RAM combined for Instances in the Group or Cloud the Policy is added to.

Max Storage Sets the max number of total of Storage combined for Instances in the Group or Cloud the Policy is added to.

Max VMs Sets the max number of Virtual Machines for the Group or Cloud the Policy is added to.

Message of the Day (MOTD) Message of the Day Policy for displaying Alerts in Morpheus. Configurable as a pop-up or full-page notification with Info, Warning and Critical message types.

Note: Requires role permission: Admin: Message Of the Day set to "Full" to create and manage MOTD Policies.

Network Quota Limits the number of networks that can be created within the policy's scope

Object Storage Quota Sets a Storage Quota for Object Storage usage (in GB) to scoped User, Role, Tenant or Global.

Power Scheduling Adds a Power Schedule for the Instances in a Group or Cloud. Power Schedules can be created in Provisioning -> Automation -> Power Scheduling

Provision Approval Sets an Approval requirement for Provisioning into a Group or Cloud using Morpheus Approvals or an Approval Integration such a Service Now.

Router Quota Limits the number of routers that can be created within the policy's scope

Shutdown Sets a shutdown timeframe in days upon provision after which the Instance will be stopped. Extensions can be auto-approved or require approval immediately or after x amount of auto-extensions using Morpheus Approvals or an Approval Integration.

Storage Server Storage Quota Sets a Storage Quota for selected Storage Server (in GB), applied Globally or per specified Tenants.

Tags Requires the user to add compliant Tags at provision time, this can be enforced on a strict or passive basis

Note: Tag scanning and enforcement is currently only available for Azure, Amazon, Google, and VMware clouds. For a more comprehensive guide on implementing Tag Policies, see the associated article in our [KnowledgeBase](#).

User Creation Controls the “CREATE YOUR USER” flag in the User Config options during provisioning do be always disabled, always enabled, enabled by default, or disabled by default.

User Group Creation Forces User Creation of members in the selected User Group during Provisioning.

Workflow Forces execution of selected Workflow for Instance Provisioning.

Creating Policies

Policies can be created in three different locations.

- Administration -> Policies
- Infrastructure -> Groups -> Group -> Policies
- Infrastructure -> Clouds -> Cloud -> Policies

Policies can be disabled and re-enabled at anytime.

Important: Precedence is applied to matching or conflicting Policies in the following order: Cloud > Group > Role > User > Global.

To create a Global Policy:

1. Navigate to Administration -> Policies
2. Select + *ADD Policy* and choose from the available policy types.
3. Refer to Policy Type sections below for Configuration options.
4. Under Filter next to scope select *Global*
5. Select *SAVE CHANGES*

To create a Policy for a User:

1. Navigate to Administration -> Policies
2. Select + *ADD Policy* and choose from the available policy types.
3. Refer to Policy Type sections below for Configuration options.
4. Under filter next to scope select *User* a drop down menu will appear below allowing you to select a user
5. Select *SAVE CHANGES*

To create a Policy for a Role:

1. Navigate to Administration -> Policies
2. Select + *ADD Policy* and choose from the available policy types.
3. Refer to Policy Type sections below for Configuration options.
4. Under filter next to scope select *Role* a drop down menu will appear below allowing you to select a Role
5. **For APPLY INDIVIDUALLY TO EACH USER IN ROLE**
 - Select for Max Resource/Quota Policies to be calculated per user
 - Leave unselected to calculate by total usage of all users within that Role.
6. Select *SAVE CHANGES*

To create a Policy for a Cloud:

Note: Resource Limitation Policies apply to all Instances in the Cloud the Policy is added to. Approval, Naming, Power, Shutdown and Expiration Policies apply to Instances created or moved into the Group after the Policy is enabled.

1. Navigate to Infrastructure -> Clouds
2. Select a Cloud by clicking on the name of the Cloud to go to the Cloud Detail page.
3. Select the **POLICIES** tab in the Cloud Detail page.
4. Select + *ADD* and choose from the available policy types.
5. Refer to Policy Type sections below for Configuration options.
6. Select *SAVE CHANGES*

To create a Policy for a Group:

Note: Resource Limitation Policies apply to all Instances in the Group the Policy is added to. Approval, Naming, Power, Shutdown and Expiration Policies apply to Instances created after the Policy is enabled.

1. Navigate to Infrastructure -> Groups
2. Select a Group by clicking on the name of the Group to go to the Group Detail page.
3. Select the POLICIES tab in the Group Detail page.
4. Select + ADD and choose from the available policy types.
5. Refer to Policy Types sections below for Configuration options.
6. Select *SAVE CHANGES*

Policy Types

Expiration Policies

Expiration policies set an expiration timeframe for any instance provisioned into the cloud, role, group or by the user the policy is added to. When an instance expires, it is terminated and deleted.

Configuration options for expiration policies:

Expiration Type

- User Configurable- expiration timeframe is editable during provisioning
- Fixed Expiration- user cannot change expiration timeframe

Expiration Days Configures the number of days the instance is allowed to exist before being removed.

Renewal Days If the instance is renewed, this is the number of days by which the expiration date is increased.

Notification Days This allows an email notice to be sent out X days before the instance is set to expire.

Notification Message Customizable message for notification emails. The default message is Instance `${instance?.name}` is set to expire on `${instance?.expireDate}`

Auto Approve Extensions Enable this to auto-approve extension requests, bypassing approval workflows.

Instances with expirations show the time until expiration in the instance detail pane. Instances with active expiration policies can be extended by selecting the EXTEND NOW button in the instance detail pane. The extension length is set in the policy by the RENEWAL DAYS field.

Expirations can also be added to any instance during provisioning by entering the number of days in the EXPIRATION DAYS field in the Lifecycle section of the automation section of the provisioning wizard. Expiration can be added to any instance even if no policies have been created.

Note: Expiration and Shutdown Policies will be enforced on Instances moved into a Group with an Active Policy or Instances created when converting an unmanaged host to managed.

Instance and Host Names

Naming Policies will populate a fixed or editable name for instances, hosts and hostnames. The Name Pattern field uses `${variable}` string interpolation.

NAMING TYPE

User Configurable Naming pattern will pre-populate during provisioning but can be edited by the user.

Fixed Name Naming pattern will pre-populate during provisioning and cannot be changed.

NAME PATTERN The Name Pattern field uses Static text and/or `${variable}` string interpolation, such as `morpheus${cloudCode}${type}${sequence+3000}`

An example Instance Name Policy using a naming pattern with User Initials, Cloud Code, Instance Type, and a sequential number starting at 3000 is `${userInitials}-${cloudCode}-${type}-${sequence+3000}`, resulting in an Instance Name of **md-vmwd3-centos-3001** for the first instance, followed by **md-vmwd3-centos-3002** and so on.

Commonly used variables for naming patterns include:

```

${groupName}
${groupCode}
${cloudName}
${cloudCode}
${type}
${accountId}
${account}
${accountType}
${platform}
${platform == 'windows' ? 'w':'l'} # results in `w` for Windows platforms and `l`
↳for Linux Platforms
${userId}
${username}
${userInitials}
${provisionType}
${instance.instanceContext} # Environment Code
${sequence} # results in 1
${sequence+100} # results in 101
${sequence.toString().padLeft(5,'0')} #results in 00001

```

Cloud codes and Group codes are fields found in their respective configuration panes.

AUTO RESOLVE CONFLICTS Morpheus will automatically resolve naming conflicts by appending a sequential -number to the name when enabled.

Shutdown Policies

Shutdown policies dictate the number of days an instance is allowed to run before it is shut down. Shutdown is consistent across cloud types i.e.: in VMware, a VM is powered off. In AWS, an instance is stopped. Etc.

Configuration options for shutdown policies:

Shutdown Type

User Configurable Shutdown timeframe is editable during provisioning.

Fixed Expiration User cannot change shutdown timeframe during provisioning.

Expiration Days Configures the number of days the instance is allowed to exist before being shut down.

Renewal Days If the instance is renewed, this is the number of days by which the shutdown date is increased.

Notification Days This allows an email notice to be sent out X days before the instance is set to shut down.

Notification Message Customizable message for notification email.

Auto Approve Extensions Enable this to auto-approve extension requests, bypassing approval workflows.

Note: Expiration and Shutdown Policies will be enforced on Instances moved into a Group with an Active Policy or Instances created when converting an unmanaged host to managed.

Provision Approval

Morpheus Provision Approvals enable an approval workflow via internal Morpheus approval or via ServiceNow workflow. If a ServiceNow integration is present, the ServiceNow option is enabled. The Approval workflow to be selected is dynamically created by querying the ServiceNow Workflow table in the integrated ServiceNow instance.

This ServiceNow approval integration enables users to use the Morpheus Self-Service provisioning portal to provision new instances and still respect the required ServiceNow business approval workflow.

Power Schedules

Power Schedules set daily times to shutdown and startup instances. Power schedule can be created and managed in Provisioning -> Automation -> Power Scheduling

Note: Power Schedule Policies will apply to Instances created in a Group or Cloud after the Policy is enabled, and will not apply to pre-existing Instances.

Configuration options for Power Schedule Policies:

DESCRIPTION Add details about your Policy for reference in the Policies tab.

Enabled Policies can be edited and disabled or enabled at any time. Disabling a Power Schedule Policy will prevent the Power Schedule from running on the Groups Instances until re-enabled.

ENFORCEMENT TYPE

- User Configurable: Power Schedule choice is editable by User during provisioning.
- Fixed Schedule: User cannot change Power Schedule setting during provisioning.

POWER SCHEDULE Select Power Schedule to use in the Policy. Power schedule can be added in Provisioning -> Automation -> Power Scheduling

TENANTS Leave blank for the Policy to apply to all Tenants, or search for and select Tenants to enforce the Policy on specific Tenants.

Max Resources

Max Resource policies allow setting quotas for Clouds, Groups, Roles or Users for maximum amount of Memory, Storage, Cores, Hosts, VM's, or Containers that can be created in the Cloud, Group, Role or by the User the Policy is assigned to.

Configuration options for Max Resources Policies:

Max Containers Sets the maximum combined total of Containers in Instances per Policy Scope.

Max Cores Sets the maximum combined total of Cores in Instances per Policy Scope.

Max Hosts Sets the maximum total of Hosts per Policy Scope.

Max Memory Sets the maximum combined total of RAM (capacity) for Instances per Policy Scope.

Max Storage Sets the maximum combined total of Storage (capacity) for Instances per Policy Scope.

Max VMs Sets the maximum total of managed Virtual Machines per Policy Scope.

TENANTS Leave blank for the Policy to apply to all Tenants, or search for and select Tenants to enforce the Policy on specific Tenants.

User Creation

The User Creation policy controls the “CREATE YOUR USER” flag in the User Config options during provisioning do be always disabled, always enabled, enabled by default, or disabled by default.

Configuration options for User Creation Policies:

TYPE User Creation

DESCRIPTION Description to identify the policy config

Enabled Policies enforcement can be disabled or enabled at any time.

ENFORCEMENT TYPE

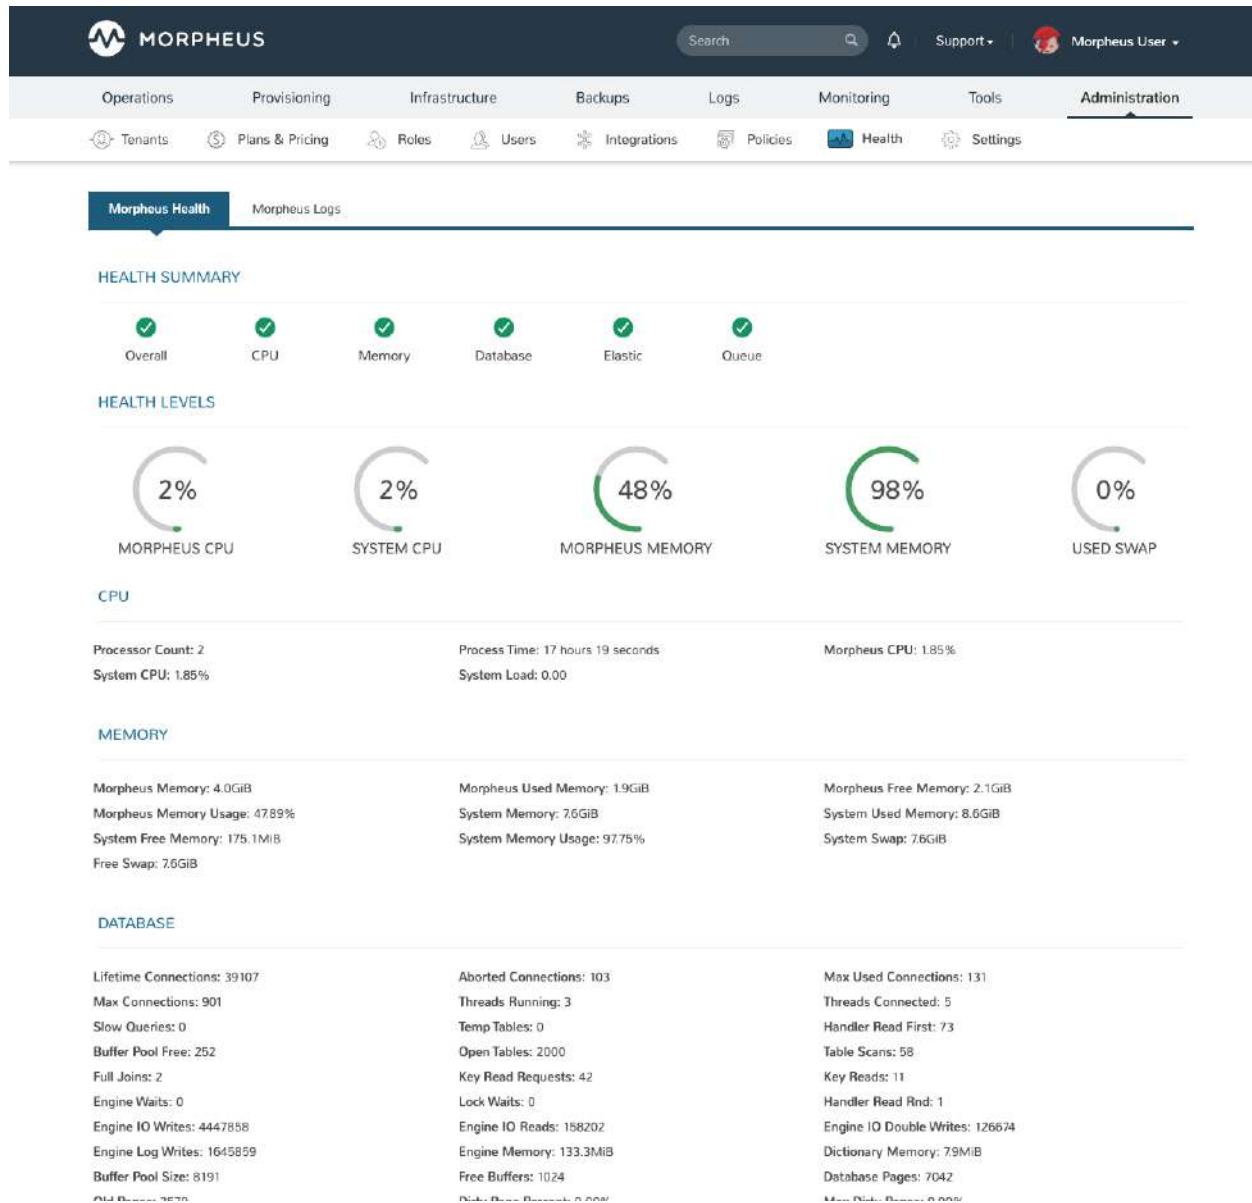
- User Configurable: User Creation choice is editable by User during provisioning.
- Fixed: User cannot change User Creation setting during provisioning.

CREATE USER Check to allow or force user creation. Uncheck to disable by default or force no user creation.

TENANTS Leave blank for the Policy to apply to all Tenants, or search for and select Tenants to enforce the Policy on specific Tenants.

Health

Morpheus Health



The Morpheus Health section provides an overview of the health of your Morpheus appliance. It includes data on the following:

- Health Levels
- CPU
- Memory
- Database
- Elastic
- Queues

Note: An Elasticsearch warning status is typical for single node Appliances due to a single elasticsearch node and

default replica count exceeding available nodes.

HEALTH LEVELS include

- Morpheus CPU
- System CPU
- Morpheus Memory
- System Memory
- Used Swap

CPU include

- Processor Count
- Process Time
- Morpheus CPU
- System CPU
- System Load

MEMORY includes

- Morpheus Memory
- Morpheus Used Memory
- Morpheus Free Memory
- Morpheus Memory Usage
- System Memory
- System Used Memory
- System Free Memory
- System Memory Usage
- System Swap
- Free Swap

DATABASE includes

- Lifetime Connections
- Aborted Connections
- Max Used Connections
- Max Connections
- Threads Running
- Threads Connected
- Slow Queries
- Temp Tables
- Key Reads
- Handler Reads

- Buffer Pool Free
- Open Tables
- Table Scans
- Full Joins
- Key Read Requests
- Key Reads
- Engine Waits
- Lock Waits
- Handler Reads
- Engine IO Writes
- Engine IO Reads
- Engine IO Double Writes
- Engine Log Writes
- Engine Memory
- Dictionary Memory
- Buffer Pool Size
- Free Buffers
- Database Pages
- Old Pages
- Dirty Page Percent
- Max Dirty Pages
- Pending Reads
- Insert Rate
- Update Rate
- Delete Rate
- Read Rate
- Buffer Hit Rate
- Read Write Ratio
- Uptime

ELASTIC includes

- Status
- Cluster
- Node Count
- Data Nodes
- Shards
- Primary Shards

- Relocating Shards
- Initializing
- Unassigned
- Pending Tasks
- Active Shards

Note: Warning status is typical for Elasticsearch

Elastic Nodes include

- Node
- Master
- Location
- Heap Usage
- Memory Usage
- CPU Usage
- 1M Load
- 5M Load
- 15M Load

Elastic Indices include

- Health
- Index
- Status
- Primary
- Replicas
- Doc
- Count
- Primary
- Size
- Total Size

QUEUES INCLUDE

- Queue Count
- Busy Queues
- Error Queues

Settings

The *Administration > Settings* section sets global configuration parameters for the Morpheus appliance, whitelabeling, provisioning, monitoring, backups, logs, software licenses, and the license for Morpheus itself.

Appliance

Appliance Settings

Host Level Firewall Enabled Enables or Disables the host level firewall. This must be Enabled to use Morpheus Security Groups.

Appliance URL The default URL used for Agent install and Agent functionality. All Instances and Hosts must be able to resolve and reach this URL over 443 for successful agent install and communication.

Note: Alternate Appliance URLs can be configured per Cloud in the *Edit Cloud -> Advanced Options* section.

Internal Appliance URL (PXE) For PXE-Boot your appliance needs to be routable directly with minimal NAT masquerading. This allows one to override the default appliance url endpoint for use by the PXE Server. If this is unset, the default appliance url will be used instead.

API Allowed Origins Specifies which origins are allowed to access the Morpheus API.

Cloud Sync Interval Data is refreshed through cloud integrations at the interval specified here in seconds, the default value is 300 seconds (five minutes). Appliances managing a very large number of clouds may be adversely affected by setting this value too low.

Blacklisted Hosts Provide a comma-separated list of IP addresses or hostnames which should be blocked when using HTTP Task types and/or REST-sourced Option Lists

Enable SSL Verification of Agent (Communications) Enabling SSL Verification of Agent Communications requires a valid Certificate be installed on the Appliance.

Disable SSH Password Authentication Only allow ssh login using SSH keys. When true, SSH Password Authentication will not be enabled for VM's and Hosts provisioned after the setting is enabled.

Tenant Management Settings

Registration Enabled If enabled, the appliance login screen will have a “NEED AN ACCOUNT? SIGN UP HERE” link added, enabling new Tenant registration.

Default Tenant Role Sets the default Tenant Role applied to Tenants created from Tenant Registration.

Default User Role Sets the default User Role applied to the User created from a Tenant Registration.

User Management Settings

Min Password Length User passwords must at least be as many characters in length as the entered value

Min Password Uppercase User passwords must include at least as many uppercase characters as the entered value

Min Password Numbers User passwords must include at least as many numerals as the entered value

Min Password Symbols User passwords must include at least as many special characters as the entered value

Session Expires (Minutes) A user session is forcibly logged out after the entered number of minutes of inactivity

Session Warning (Minutes) A pop-up warning is shown to the user when they have been inactive for the number of minutes entered. Example: If sessions are set to expire after 90 minutes, warn the user after 60 minutes if you intend to provide 30 minutes advance warning

Expire Password After (Days) User account passwords will expire after the entered number of days. Enter 0 or leave the field empty to opt out of this feature.

Disable User After Attempts (Number of Attempts) Disable a User account after a specified number of failed login attempts. Enter 0 or leave the field empty to opt out of this feature.

Disable User If Inactive For (Days) Disable a User account if inactive for the entered number of days. The User will not be able to log into the appliance again until another User with sufficient rights enables the account. Enter 0 or leave the field empty to opt out of this feature.

Send warning email before deactivating (Days) Enter the number of days prior to account deactivation that a warning email should be sent. For example, enter “5” to warn the User when they are five days short of the deactivation time entered in the prior field. Enter 0 or leave the field empty to opt out of this feature.

Email Settings

A default installation of Morpheus uses an online service called postmarkapp. Morpheus api requests to the postmarkapp service to send notification e-mails.

To add your own SMTP server you will need to go to the Administration and Settings of your Morpheus appliance. You will then need to provide Morpheus the following information, your mail server systems administrator should provide you with the below information and the preferred encryption method.

- From Address
- SMTP Server
- SMTP Port
- SSL Enabled
- TLS Encryption
- SMTP User
- SMTP Password

We recommend that you add your Morpheus server to your SMTP white list as well as using user authentication as an additional security measure.

Once you have added your SMTP server information into Morpheus scroll down the Administration and Settings page and press the blue save button which can be found under enabled clouds.

When you have saved your SMTP server settings in the Morpheus appliance you will then need to restart the Morpheus-ui. To restart the Morpheus-ui connection to your Morpheus server via ssh and run the below command.

```
sudo morpheus-ctl restart morpheus-ui
```

Important: If you do not restart the Morpheus-ui the notifications will be sent by the original notification service postmarkapp. Please note it can take up to 3 minutes for the ui to become reachable again. has a built in SMTP server for email notifications and alerts. An alternate SMTP server can be specified below:

Add an alternate SMTP Server:

- From Address
- SMTP Server
- SMTP Port
- SSL Enabled
- TLS Encryption
- SMTP User
- SMTP Password

Proxy Settings

The Morpheus Appliance can be configured to communicate through a Proxy server for Cloud API's and Agent communication back to the Appliance.

Note: Additional Proxy configuration is available in the *Infrastructure > Network > Proxies* section. Added Proxies can be scoped to Clouds in the *Edit Cloud > Advanced Options* section of the Cloud.

Add a Global Proxy server by entering the following:

- Proxy Host
- Proxy Port
- Proxy User
- Proxy Password
- Proxy Domain
- Proxy Workstation

Currency Settings

In Morpheus, Tenants are separate environments which can be defined as using currencies that are unique from one Tenant to the next. In addition, these currencies may be different from the currency in which Price Sets have been defined. In order to present pricing to Subtenant users in their designated currency, Morpheus allows for integration with currency conversion services “open exchange rates” and “fixer.io”. This article goes through the process of setting up the integration and how it works to determine pricing conversions.

Integrating With a Currency Exchange Provider

1. Navigate to Administration > Settings > APPLIANCE
2. Under the Currency Settings heading, make a “Currency Provider” selection
3. Enter your “Provider API Key”

The service is now integrated and can be used as described in the next section.

Consuming Currency Exchange in Morpheus

Currency exchange data is synced from the integrated provider once every 12 hours. When needed, Morpheus will use this cached data to present currency conversions rather than hitting the API directly each time. This limits the total number of API hits and reduces costs.

Exchanged currency values will be shown under conditions similar to the following scenario:

A user is working in a Subtenant configured for Currency B. The user is attempting to provision an instance with pricing sets that have only been defined in Currency A. Morpheus will convert the pricing data from currency A to Currency B for this user (and all users in this Subtenant) since price conversion has been enabled.

Enabled Clouds (Types)

Controls which types of Cloud can be created.

- When a Cloud type is disabled, it will be removed from the available options when adding new Clouds in Infrastructure > Clouds. Existing Clouds are not affected by changes to this setting.

Whitelabel

Whitelabel Settings

Overview

Morpheus Tenants can be WhiteLabeled with custom Logos, Colors, Copy, and custom CSS. Sub-Tenants can be individually white-labeled, or the Master Tenant Whitelabel can apply to all Sub-Tenants.

Enable Whitelabel Turns on the configured Whitelabel settings. Disabling will return the Appliance to the default colors and logos, but the configured options will remain saved and will apply if Whitelabel is re-enabled.

Appliance Name Replaces Morpheus in page titles.

Header Logo Top left header logo. Uploaded image is resized to 38 pixels high with a proportional width at that height.

Disable Support Menu Enable this flag to hide the support dropdown menu in the header.

Support Menu Links Customize support links. Label Code can be used for translations and is optional. Be sure to specify fully qualified url if linking to external sites.

Security Banner

The Security Banner section in `/admin/settings#!whitelabel` displays content on the login screen for Security and Consent

- Applicable at Global and Tenant levels

- Security Banner input field accepts plain text and markdown
- Content is displayed below login section in scoped /login/auth pages.

Footer Logo Footer Logo in bottom left. Uploaded image is resized to 27 pixels high with a proportional width at that height.

Login Logo Logo shown on Login screen. Uploaded image is resized to 192 pixels wide with an unbound height proportional to that locked width.

Favicon Must be a .ico file type.

Reset When selected and Whitelabel settings are saved, associated logo is returned to blank default value.

Colors

Update Colors by entering HEX value or selecting the Color Selector pop-up next to each field and selecting a color.

- Header Background
- Header Foreground
- Nav Background
- Nav Foreground
- Nav Hover
- Primary Button Bg
- Primary Button Fg
- Primary Button Hover Bg
- Primary Button Hover Fg
- Footer Background
- Footer Foreground
- Login Background

Override CSS

Override CSS settings by entering CSS in *Override CSS* field.

Example: (this will add one continues background image to the Header)

```
header #topHeader {
    background-image: url(http://image_url.png);
}
header {
    background-image: url(http://image_url.png);
}
```


Reuse Sequence Numbers Enable for sequence numbers to always increment and never be reused. When disabled, sequence numbers will be reused.

Deployment Archive Store Default Storage Provider for storing Deployment Archives.

Note: Storage Providers can be configured and managed in the *Infrastructure -> Storage* section.

Cloud-Init Settings

Morpheus can add global users for Linux and Windows at provision time. Cloud-init/Cloudbase-Init or VMware Tools installed on the provisioned virtual images is required.

Linux

- **Username:** Enter User to be added to Linux Instances during provisioning.
- **Password:** Enter password to be set for the above Linux user.
- **KeyPair:** Select KeyPair to be added for the above Linux user.

Note: Either a password, keypair, or both can be populated for the Linux user. Keypairs can be added in the *Infrastructure > Keys & Certs* section.

Windows Settings

- **Administrator Password:** Enter password to be set for the Windows Administrator User during provisioning.

PXE Boot Settings

Default Root Password Enter the default password to be set for Root during PXE Boots.

Overview

The Environments section is where you create and manage your environment labels, which are available in the *Environment* dropdown during Instance or App provisioning. An Instance's environment label can be changed by editing the Instance.

Creating Environments

1. Select + *Create Environment*
2. Populate the following for the New Environment:

Name The friendly name for the environment in Morpheus

Code Shortcode used for API and CLI

Description Environment description displayed on the Environments list page

Display Order The order in which environments are presented when provisioning, a value of "0" will position the environment at the top of the list

Visibility

- *Private*: Available only in the Tenant the environment is created in
- *Public*: Available for all Tenants. Public is only applicable for environments created in the the Master Tenant.

Note: User-created environments can be edited, hidden, or removed from the Actions menu on the environments list page. Morpheus-default environments can only be hidden from users during provisioning.

Overview

The License section is for automating the application of Licenses to Instances while provisioning. Licenses can be added to Morpheus and then attached to images. Morpheus will then apply the license to Instances provisioned using the images with license attached. Licenses can be configured for single or multiple Tenants.

Creating Licenses

1. Select + *Create License*
2. In the New License modal, enter the following:
 - **License Type** Windows
 - **Name** Name of the License in Morpheus
 - **License Key** Enter the License Key
 - **Org Name** The Organization Name (if applicable) related to the license key
 - **Full Name** The Full Name (if applicable) related to the license key
 - **Version** The License Version
 - **Copies** The Number of copies available on the License
 - **Description** License description displayed in the Licenses list in Morpheus, helpful for identifying the License after creation
 - **Virtual Images**

Search for existing Virtual Images by name and select to attach the image to the license.

Note: Virtual Images are synced from Clouds or added in the *Provisioning -> Virtual Images* section.

- **Tenant Permissions** Search for and select the Tenant(s) the License will be available for. Multiple Tenants can be added.
3. Save Changes

Provisioning with Licenses

When a Virtual Image is added to a license, Morpheus will automatically apply the License to Instances configured with the Virtual Image during provisioning, including Instance Types with a Node Type that is configured with the Virtual Image, or if the image is selected when using generic Cloud Instances types (VMware, AWS, Nutanix, Openstack etc). Virtual Images can be removed from a License by editing the License.

Managing Licenses

Created Licenses details are displayed in the License page, including the number of copies applied per License, the Tenants added to the License, and the Virtual Images attached to the License.

The Name, Version, Copies, Description, Virtual Images and Tenant Permissions are editable but selecting the *Actions* dropdown on a License.

Note: License Types, Keys, Org Names and Full Names are not editable after a license has been created.

License can also be removed using the *Actions* dropdown on a License.

App Blueprint Settings

Determines the Default Blueprint Type selected in new App Wizard

- Morpheus
- ARM Template
- Cloud Formation
- Terraform
- Kubernetes Spec
- Helm Chart

Monitoring

Morpheus Monitoring Settings

Auto Create Checks When enabled a Monitoring Check will automatically be create for Instances and Apps.

Availability Time Frame The number of days availability should be calculated for. Changes will not take effect until your checks have passed their check interval.

Availability Precision The number of decimal places availability should be displayed in. Can be anywhere between 0 and 5.

Default Check Interval The default interval to use when creating new checks.

Note: Monitoring Checks can be manually configured if *Auto Create Checks* is disabled.

AppDynamics

AppDynamics Monitoring Integration Settings

Enabled Enables the AppDynamics Integration

Controller Host This is the host name or the IP address of the AppDynamics Controller. This is the same host that you use to access the AppDynamics browser-based user interface.

Controller Port This is the HTTP(S) port of the AppDynamics Controller. This is the same port that you use to access the AppDynamics browser-based user interface. If the Controller SSL Enabled property is set to true, specify the HTTPS port of the Controller; otherwise specify the HTTP port.

Controller SSL Enabled This property specifies whether the agent should use SSL (HTTPS) to connect to the Controller. If SSL Enabled is true, set the Controller Port property to the HTTPS port of the Controller.

Tenant Name This is the account name used to authenticate with the Controller.

Access Key This is the account access key used to authenticate with the Controller.

Controller Version This is the controller version and can be obtained at the bottom of the controller login page.

Application Name This is the name of the logical business application. Note that this is not the deployment name(ear/war/jar) on the application server. (Maximum of 30 numbers or letters)

Tier Name This is the name of the logical tier. (Maximum of 30 numbers or letters)

Controller User A user that can login to the Controller ui and upload a dashboard.

Controller Password Password for the Controller User.

Service Now

ServiceNow Monitoring Integration Settings

Note: A ServiceNow Integration must be already configured in *Administration -> Integrations* to enable the ServiceNow Monitoring Integration.

Enabled Enables the ServiceNow Monitoring Integration

Integration Select from a ServiceNow Integration added in *Administration -> Integrations*

New Incident Action The Service Now action to take when a Morpheus incident is created.

Close Incident Action The Service Now action to take when a Morpheus incident is closed.

Incident Severity Mapping

Morpheus Severity	ServiceNow Impact
Info	Low/Medium/High
Warning	Low/Medium/High
Critical	Low/Medium/High

New Relic

New Relic Integration Settings

Enabled Enables the New Relic Monitoring Integration

License Key License Key to be used when installing the New Relic agent in order for the agent to report data to your New Relic account

Note: The License Key is the 40-character hexadecimal string that New Relic provides when you sign up for your account.

Backups

Backup Settings

The Backup settings page allows you enable or disable scheduled backups, select a default backup bucket, and administer global settings related to backups. Changes to global settings only affect new backups going forward and do not affect existing backups.

Morpheus Backup Settings

Scheduled Backups Enable automatic scheduled backups for provisioned instances

Create Backups When enabled, Morpheus will automatically configure instances for manual or scheduled backups

Backup Appliance When enabled, a backup will be created for the Morpheus appliance database. Select the [Backup](#) text link to view or edit settings related to the appliance backup

Default Backup Bucket Select an existing bucket as the default for future backup runs. Click the [Infrastructure Storage](#) text link to add a new storage bucket to Morpheus if needed

Default Backup Schedule Choose a default schedule interval for automated backups. The available selections in this dropdown menu are Execution Schedules defined in [Provisioning > Automation > Execute Scheduling](#)

Default Backup Retention Choose the default number of backups to be retained for automated Instance and appliance backup jobs

Logs

Logging Settings

Overview

Morpheus contains a built-in logging solution that aggregates logs from hosts and services. Logs are displayed, searchable, and filterable in the Instance, App, Host and overall Logs sections. Logs can also be forwarded using Syslog Forward rules to any external solution that supports syslogs.

The logs displayed in the Instance, App, Host and overall Logs sections are only from Managed VM's and Hosts that have the Morpheus agent installed. Instances can be configured to show additional logs by configuring the LOG

FOLDER in the Library NODE TYPE. Logs from any .log file in the specified folder will be forwarded by the Morpheus agent to the Morpheus appliance or forwarded with Syslog Forward rules.

Note: The *Logs* section does not contain Morpheus appliance logs, which can be found in */var/log/morpheus/* and in 3.5.2+ in *Operations - Health*.

Logs are stored in ElasticSearch and retention can be set by adjusting the Availability Time Frame in the *Administration* -> *Logs* section.

Morpheus also has built in Integrations with 3rd Party solutions. When configured, the Morpheus agent will forward logs to the integrated platforms automatically.

Logging Settings for the build-in Logging, Syslog forwards, and 3rd Party Integrations are configurable in the *Administration* -> *Logs* section.

Morpheus Logging

Morpheus contains a built-in logging solution that aggregates logs from hosts and services. Logs are displayed, searchable, and filterable in the Instance, App, Host and overall Logs sections. Logs can also be forwarded using Syslog Forward rules to any external solution that supports syslogs.

Splunk

To configure Splunk, create a syslog listener configuration in Splunk. Then it is simply a matter of expanding the section in Logging settings pertaining to Splunk and filling out the host and port of the appender. Once saved, all hosts managed by Morpheus will be configured to forward logs to the target Splunk listener.

LogRhythm

Configuring LogRhythm is much like configuring Splunk. Simply toggle the enabled flag in the LogRhythm section to enabled and fill in the Host, and Port information for the LogRhythm listener.

Guidance

Overview

Morpheus guidance is an important tool that makes recommendations for resource and cost optimization. It analyzes CPU, memory, and storage activities over time to make intelligent recommendations on sizing and power state. These recommendations can free up resources and save organizations significant amounts of money over time. Out of the box, Morpheus is configured for sensible thresholds used in making these recommendations but they can be edited here if needed.

Power Settings

Morpheus will recommend shutting down a resource if *all* three of the baselines in this section are exceeded:

- **Average CPU (%):** Shutdown will be recommended if the average CPU usage is below this value. Values over 100% are possible as this value factors the number of CPU cores. Default value: 75
- **Maximum CPU (%):** Shutdown will be recommended if the CPU usage never exceeds this value. Values over 100% are possible as this value factors the number of CPU cores. Default value: 500
- **Network Threshold (bytes):** Shutdown will be recommended if the average network bandwidth is below this value. Default value: 2000 bytes/second

CPU Up-size Settings

CPU up-size will be recommended when *both* of the following baselines are exceeded for a resource:

- **Average CPU (%):** CPU up-size is recommended if the average CPU percentage exceeds this value (and other criteria are also met). Default value: 50
- **Maximum CPU (%):** CPU up-size is recommended if the maximum CPU percentage exceeds this value. Default value: 99

Memory Up-size Settings

Memory up-size will be recommended when both of the following thresholds are met for a resource:

- **Minimum Free Memory (%):** Memory up-size will be recommended if free memory dips below this value. Default value: 10

Memory Down-size Settings

Memory down-size will be recommended when both of the following thresholds are met for a resource:

- **Average Free Memory (%):** Memory down-size is recommended if the average free memory is above this value. Default value: 60
- **Maximum Free Memory (%):** Memory down-size is recommended if free memory has never dipped below this value. Default value: 30

Environments

Overview

The Environments section is where you create and manage your environment labels, which are available in the *Environment* dropdown during Instance or App provisioning. An Instance's environment label can be changed by editing the Instance.

Creating Environments

1. Select + *Create Environment*

2. Populate the following for the New Environment:

Name The friendly name for the environment in Morpheus

Code Shortcode used for API and CLI

Description Environment description displayed on the Environments list page

Display Order The order in which environments are presented when provisioning, a value of “0” will position the environment at the top of the list

Visibility

- *Private*: Available only in the Tenant the environment is created in
- *Public*: Available for all Tenants. Public is only applicable for environments created in the the Master Tenant.

Note: User-created environments can be edited, hidden, or removed from the Actions menu on the environments list page. Morpheus-default environments can only be hidden from users during provisioning.

Software Licenses

Overview

The License section is for automating the application of Licenses to Instances while provisioning. Licenses can be added to Morpheus and then attached to images. Morpheus will then apply the license to Instances provisioned using the images with license attached. Licenses can be configured for single or multiple Tenants.

Creating Licenses

1. Select + *Create License*

2. In the New License modal, enter the following:

- **License Type** Windows
- **Name** Name of the License in Morpheus
- **License Key** Enter the License Key
- **Org Name** The Organization Name (if applicable) related to the license key
- **Full Name** The Full Name (if applicable) related to the license key
- **Version** The License Version
- **Copies** The Number of copies available on the License
- **Description** License description displayed in the Licenses list in Morpheus, helpful for identifying the License after creation
- **Virtual Images**

Search for existing Virtual Images by name and select to attach the image to the license.

Note: Virtual Images are synced from Clouds or added in the *Provisioning -> Virtual Images* section.

- **Tenant Permissions** Search for and select the Tenant(s) the License will be available for. Multiple Tenants can be added.

3. Save Changes

Provisioning with Licenses

When a Virtual Image is added to a license, Morpheus will automatically apply the License to Instances configured with the Virtual Image during provisioning, including Instance Types with a Node Type that is configured with the Virtual Image, or if the image is selected when using generic Cloud Instances types (VMware, AWS, Nutanix, Openstack etc). Virtual Images can be removed from a License by editing the License.

Managing Licenses

Created Licenses details are displayed in the License page, including the number of copies applied per License, the Tenants added to the License, and the Virtual Images attached to the License.

The Name, Version, Copies, Description, Virtual Images and Tenant Permissions are editable but selecting the *Actions* dropdown on a License.

Note: License Types, Keys, Org Names and Full Names are not editable after a license has been created.

License can also be removed using the *Actions* dropdown on a License.

License

Overview

Morpheus requires a valid license for provisioning new Instances, Apps and Hosts, and converting existing Instances and Hosts to managed. Licenses can be applied and updated in this section, and the current license status can be checked.

Note: Morpheus is licensed for a certain number of concurrent workload elements (WLEs) that may be managed or inventoried at any one time. See our [Knowledge Base](#) for specific information on the types of WLEs that count against Morpheus licensing.

Current License

If a License Key has already been applied, the License status is shown in the *Current License* section:

Tenant Name Company name the License was generated for.

Start Date Date and time the current License started.

End Date Date and time the current License expires.

Space Amount of used and unused Managed RAM under the current License.

EXAMPLE: On a 1 TB License with 182 GB of RAM under management, the Space section will show *Used Space 182.9GB Unused Space 841.0GB*

Note: Once a current License expires or has reached its Space limit, users will no longer be able to provision new Instances, Apps, Hosts, or Bare Metal, or convert existing Hosts, Virtual Machines, or Bare Metal to managed. Morpheus will otherwise continue to function.

Upgrade License Key

To add a new or update an existing License:

1. Copy the License Key into the License Key field
2. Click *UPDATE*

If valid, the new License will be applied.

Request new License

Licenses can be requested at <https://morpheushub.com>, or by contacting support@ or sales@ morpheusdata.com.

Utilities

System administrators have access to a utilities panel with the following options:

- **Reindex all searchable data:** Execute
- **Toggle Maintenance Mode:** Enable

Note: Maintenance mode cleanly places Morpheus into a state where maintenance can be performed on the appliance. This drains any active sessions and queues so an auto-scaling group can scale down. It also drains active sessions across services.

User Settings

User settings are accessed by clicking on your display name in the far upper-right corner of the application window. In this dropdown menu, click on the “USER SETTINGS” link.

User Photo

Upload a custom image for your user avatar that is displayed in the top header and user administration sections.
Suggested Photo Dimensions: 128 x 128

User Settings

The fields included in this section are described below. By entering any new values in these fields and clicking *SAVE*, the existing value will be overwritten.

- **Username:** Your Morpheus username
- **First Name:** Your first name (together with Last Name makes up your display name)
- **Last Name:** Your last name (together with First Name makes up your display name)
- **Email:** Your email address
- **Password:** Enter a value and save changes to update your password. The value in the Confirm field below must match
- **Confirm:** Confirm the new password you’ve entered
- **RECEIVE NOTIFICATIONS** Determines if Provisioning notifications are emailed to this Users.

User Photo



Suggested Photo Dimensions:
128 x 128

UPLOAD PHOTO

User Settings

Username morphUser

First Name morpheus

Last Name user

Email morpheus@morpheus.cloud

Password

Confirm

☒ RECEIVE NOTIFICATIONS

API ACCESS

Preferences

- **Default Group:** Sets the default Group selection when provisioning.
- **Default Cloud:** Sets the default Cloud selection when provisioning.
- **Default Persona:** Sets the default Persona used when logging in.

Preferences

Default Group	<div>Labs</div>
Default Cloud	<div>Azure Labs</div>
Default Persona	<div>Standard</div>


Linux Settings

When provisioning a Linux-based resource and opting to have your user created during the provisioning process, the credentials entered in this section will be used to seed that user into the provisioned resource.

- **Username:** The username that will be used with your Linux user
- **Password:** The password that will be used with your Linux user (optional if specifying key)
- **Confirm:** Confirm your entered password. These must match in order for the new password value to be saved
- **SSH Key:** Select a pre-existing SSH key pair object in Morpheus. Required if not specifying password and creating your user during provisioning, or required if ssh password authentication has been disabled.

Warning: If your users Linux Settings password and/or key are not defined, and 'Create User' is enabled during provisioning (default), a random password will be generated but not exposed and you will not be able to login with your user.

Linux Settings

Username	<input type="text" value="morphUser"/>
Password	<input type="password" value="....."/>
Confirm	<input type="password" value="....."/>
SSH Key	<input type="text" value="morpheus-3-D4XB8L"/> 

Windows Settings

When provisioning a Windows-based resource and opting to have your user created during the provisioning process, the credentials entered in this section will be used to seed that user into the provisioned resource.

- **Username:** The username that will be used with your Windows accounts
- **Password:** The password that will be used with your Windows accounts
- **Confirm:** Confirm your entered password. These must match in order for the new password value to be saved

Warning: If your users Windows Settings password is not defined, and ‘Create User’ is enabled during provisioning (default), a random password will be generated but not exposed and you will not be able to login with your user.

Windows Settings

Username

Password

Confirm

API Access

Click the *API Access* button to expand the “API ACCESS” modal. In this modal you can generate or refresh access tokens that can be used with Morpheus API and Morpheus CLI.

If no token yet exists for a particular “CLIENT ID”, click *ACTIONS* and then Generate. If a token has expired, we can also regenerate that token by clicking *ACTIONS* and then Regenerate. After regenerating a particular token, you would need to ensure any scripts using those tokens are updated.

After navigating away from the User Settings page, the complete access and refresh tokens will be masked from view. If these are lost or compromised, you can eliminate a token completely by clicking *ACTIONS* and then Clear. If you need to generate a new token for the same Client ID, click *ACTIONS* and then Regenerate.

API ACCESS ×				
CLIENT ID	ACCESS TOKEN	REFRESH TOKEN	EXPIRES	
morph-api			09/25/2021 10:42 PM	ACTIONS ▾
morph-automation			09/25/2020 11:06 PM	ACTIONS ▾
morph-cli			06/12/2020 09:16 PM	ACTIONS ▾
morph-customer			09/25/2021 10:33 PM	ACTIONS ▾

DONE

Note: Access Tokens are only displayed/available after generation. Copy new Tokens and store appropriately before navigating from /user-settings, they will not be displayed again.

1.3.5 Monitoring

Overview

Morpheus provides great monitoring features out of the box. Anything provisioned within Morpheus automatically gets a check created in the monitoring service. These checks are organized hierarchically in “Groups” and “Apps”. This makes it easy to gain a perspective as to what a customer or full stack facing impact is in the event of a particularly instance failure. This also takes into account redundancy layers when it comes to calculating the applications overall uptime percentage.

There are also several integrations built into the monitoring subsystem of Morpheus including App Dynamics , New Relic, and even Service Now integration.

Apps

App monitors are very useful for seeing an aggregation of failures or impact based on a set of checks and groups. App monitors typically correlate to Apps provisioned from Morpheus Blueprints but can also be manually created and organized. They can be great for visualizing the customer impact a failure might have or even keeping up on a screen in a NOC. To create an App monitor:

Name A friendly name for the new app monitor in Morpheus

Description An optional description value to identify the app monitor

Max Severity The maximum severity incident a failed app may create. This setting overrides check and group max severity settings

Affects Availability When checked, this failed app impacts system-wide availability calculations

App Checks Use the typeahead field to select as many checks as needed to complete the app monitor. Checks are created in Monitoring > Checks and must exist prior to creating the app monitor

Checks

The Monitoring system is composed of individual checks. A check is created for every container or vm that is provisioned through Morpheus . One interesting thing about these checks is they are type aware. There are several different built in check types that are selected based on the service or instance type that is being provisioned. These range from database type checks to web checks and message checks. They are highly configurable and also feature fallback check types for those more generic use cases.

Checks can be customized to run custom queries, check queue sizes, or even adjust severity levels and check intervals. All of these things can be controlled from the Checks sub tab within Monitoring.

Health

A check can have 3 health states. They are Failed, Warning (Recovering), and Healthy. When a check test fails the system automatically reattempts the check after 30 seconds to eliminate false positives. This will convert the check into a *Failed* state and raise the appropriate severity incident depending on the grouping of the check. When a check recovers it automatically goes into a Warning state. This will remain in the warning state until 10 successful check runs have completed.

Options

All check types have several core options and some of these default options can be configured in *Admin -> Monitoring*. This includes the default check interval time. By default a check is run every 5 minutes. This can however be changed to run as frequently as once every minute.

- *Max Severity*: The maximum severity level impact for a created incident that can occur if the check fails (defaults to Critical).
- *Check Interval*: The frequency with which a check is run (default 5 minutes).
- *Affects Availability*: Whether or not this check impacts overall system availability calculations.

SSH Tunneling

In many cases when it comes to monitoring databases, and services they may not be fronted on the public ip's for external monitoring. To reach these safely, and securely Morpheus provides an SSH Tunneling mechanism for its check servers. This allows the check to be confirmed via an ssh port tunnel securely using a keypair.

Check Servers

On a base installation of Morpheus a single *check server* is installed on the appliance. This is used for running any custom user checks. This service connects to the provided rabbitmq services and can be moved off or even scaled horizontally onto sets of check servers. All other checks that are related to provisioned containers or VMs are executed by the installed agent on the guest OS or Docker host.

Check types

Web Check

A web check is useful to identify if a url is reachable and the text to match check criteria confirms if the website is loading with the expected values. The text to match character should be within the first few lines of the page source.

Use case:

Adding a check to make sure morpheus demo environment is functioning. The below check will login to the morpheus demo environment.

Values to be added in Check:

- Name: "<enter name>"
- Type: Web Check
- Interval: 5 mins (Select an interval)
- Max severity: Critical
- Check the box for affects availability
- Web Url: <https://demo.morpheusdata.com/operations/dashboard> (Note: this page will load only if my login is successful. Enter the login details in Username and password fields)
- Request Method: GET
- Basic Authentication: * User: <username> * Password: <password>

- Text to Match: “Morpheus” (Login to the url and on the page of dashboard, right click and select view page source. In the first few lines, look for a text that you want this check to verify)
- Save Changes

Push API Check

This check can be used to send an API call to morpheus from a platform to check if the push api is working. A push Check is not polled regularly by the standard monitoring system. Instead it is expected that an external API push updates as to the status of the check timed closely with the configured check interval setting. This is used to throttle the push from performing too many status updates.

Note: If a check is not heard from within the check intervals, It’s status will be updated to error and an incident will be raised as if it failed.

Use Case:

Send an API call from an app to make sure the API is not cluttered and can send checks in a 2 mins interval.

Values to be added to the check:

- Name: “<enter name>”
- Type: “Push API Check”
- Interval: 5 mins (Select an interval)
- Max severity: Critical
- Check the box for affects availability
- Copy the curl command and schedule to send this via your API. For testing we used postman to send the api call at an interval of 4 mins.
- Save Changes

MySQL Check

This check is used to run a query on a host running mysql.

Use Case:

Query localhost running mysql to query a table to check if there is any status as requested. If the status has a count

Values to be added to the check:

- Name: “<enter name>”
- Type: “MySQL Check”
- Interval: 5 mins (Select an interval)
- Check the box for affects availability
- Host: 127.0.0.1
- Port: 3306

- DB Name: morpheus
- User: <db user name>
- Password: <password>
- Query: “select count(*) as count from request_reference where status = ‘requested’;”
- Operator: Equal
- Check results: 1
- Save Changes

Groups

Group monitors can only contain checks and can be edited or created in *Monitoring > Groups*. Besides simply adding and removing checks to a group there are a few other useful options that can be customized in a group:

Name A friendly name for the group monitor in Morpheus

Min Checks This specifies the minimum number of checks within the group that must be happy to keep the group from becoming unhealthy

Max Severity The maximum severity incident a failed check may create. This setting overrides a check’s max severity setting

Affects Availability If checked, a failing group monitor impacts system-wide availability calculations

Checks Use the typeahead field to select pre-existing checks for the group monitor. If check(s) need to be created, this can be done in *Monitoring > Checks*

Note: Some useful information can also be seen on the detail page of a check. For example, the average response time of all checks within the group, or an aggregated check history can be viewed.

Incidents

Incident management is very important in any IT Operations environment. The ability to notify the appropriate people of an outage that requires immediate attention is critical to reducing recovery time and even preventing potential customer facing impacts. Because of this, Morpheus provides incident management features as well as external integrations out of the box.

Incidents can be found in the *Monitoring->Incidents* section. When a check fails, an incident is automatically raised. These can vary in severity based on the user configured check severities as well as the group hierarchy (representative of redundancy).

Incidents are also grouped. If an application is impacted and multiple checks fail for that application they automatically get grouped together in one Incident that can fluctuate or escalate in severity as time progresses. These incidents can be muted so as not to affect availability and they can also be resolved manually with an option to detail resolution information.

There are also integrations and API’s for integrating with existing corporate workflows when it comes to incident management.

Contacts

To configure user notifications, a contact must first be created in `Monitoring > Contacts`. These contacts can be one of a few types:

- **Contact:** Used for either Email or SMS
- **Web Hook:** Used for posting a notification to a web endpoint or Alert Ops
- **Slack Hook:** Used for posting notifications to a Slack channel (ex.
- **VictorOps:** Provides a web post format consistent with the required notification format for Victor Ops.

Most of these options provide convenient examples and information when configuring the contact. Once they are configured, contacts can freely be used to build *Alert Rules*.

Alert Rules

Alert Rules provide a powerful means to configure who gets notified in various scenarios. These scenarios include targeting specific checks, groups, or apps, and adding the appropriate recipients to be notified during a situation in which those filters are impacted.

- **Min Duration:** This setting delays notification to the recipients by the entered number of minutes required for the incident to be opened.
- **Min Severity:** Some executives might want to be notified of an outage but only if the severity impact goes above a certain level. This is very useful for scoping escalations.

To add recipients to a rule just start typing their name in the Recipients section towards the bottom of the edit form. An auto-complete list will start populating with contact names. Once one is selected a delivery method can be selected as well as whether or not they should be notified of any escalation changes and/or closed incidents. The delivery methods available depend on the type of contact information configured for your contact. If needed, contacts can be created or edited in `Monitoring > Contacts`.

Tip: A recipient can be in multiple alert rules and can even be configured to be notified via different methods depending on the rule. A useful example might be to alert someone via email for lower severity incidents but SMS for critical severity levels.

1.3.6 Logs

Logs

Overview

The logging architecture backing Morpheus uses the latest and greatest technologies and standards to be able to service large amounts of log traffic as well as facilitate easy viewing. Utilizing elasticsearch behind the scenes and buffered log transmission protocols Morpheus provides a highly efficient and highly scalable solution for capturing log data from anything provisioned via the system. By utilizing common formats (syslog) it is also very easy to forward logs to external third party log services.

Configuration

Logging configuration can be setup in the `Admin > Logs` section. There are useful settings here, including customizing the retainment policy (7 days by default). This could be expanded to years for PCI compliance purposes or other requirements an organization might have.

Note: When increasing the retainment policy of the logging system, it may be necessary to scale out the elasticsearch cluster. Please refer to the relevant information with regards to scaling elasticsearch and advanced installation options for externalizing the elasticsearch cluster.

The Log administration section also provides options for setting custom syslog forward rules. These rules are applied on each individual host therefore keeping the Morpheus appliance itself out of the data plane. For information on different syslog formatting rules please refer to the <http://www.rsyslog.com/sending-messages-to-a-remote-syslog-server/> [{}rsyslog] documentation.

Usage

Morpheus automatically sets up and configures logging for all of the standard catalog items provisioned through morpheus. This includes both Docker containers as well as virtual machines. Simple view instance specific logs in instance detail via the “Logs” tab.

There are several filtering capabilities built into the logging ui with more being added continually. Easily toggle log level filters from the dropdown or change the date range filter using the handy date filter component. A chart is also displayed above logs representing the log counts by level over the selected time range (default last 24 hours). A handy pattern search is also available with some rather capable features based on Lucene search syntax.

Tip: It may be useful to review the Lucene search query syntax for powerful use cases: https://lucene.apache.org/core/2_9_4/queryparsersyntax.html [{}Syntax Guide]

There are several other places logs can be viewed. Not only can they be viewed across an application in app detail but also across all instances in the account. The main level `Logs` section provides an ability to query all logs produced by the system. It is also possible to view host specific logs on a docker host by viewing the host detail page via `Infrastructure`.

Note: New features are on the roadmap for the main logs section including saved searches, and handy charting dashboards for garnering insights out of log data.

Integrations

While the built in logging solution provided by Morpheus is sufficient for most, there are some scenarios in which a more advanced logging system may be desired or already in place. To facilitate this Morpheus makes it easy to add custom syslog rules as well as built in direct integrations with Splunk and LogRhythm. All integrations pertaining to logging can be configured in the `Administration -> Logging` section.

Splunk

To configure Splunk simply create a syslog listener configuration in Splunk. Then it is simply a matter of expanding the section in Logging settings pertaining to Splunk and filling out the host and port of the appender. Once saved, all hosts managed by Morpheus will be configured to forward logs to the target Splunk listener.

LogRhythm

Configuring LogRhythm is much like configuring Splunk. Simply toggle the enabled flag in the LogRhythm section to enabled and fill in the Host, and Port information for the LogRhythm listener.

Exporting Logs

Log Settings

There are three main log areas in Morpheus

- Agent Logs
- Morpheus Server Logs
- Activity / Audit Logs

Agent Logs

When Instances are deployed through Morpheus, the installed Agent captures application logs and sends them back to the Morpheus server.

In most cases, the built-in Morpheus logging features are sufficient for tracking and reviewing Agent logs. However, if needed, Morpheus supports integration with advanced logging systems. See the [log integration section](#) above for more information.

Morpheus Server Logs

The main Morpheus server log is in `/var/log/morpheus/morpheus-ui` and the latest log file is named `current`. This log is archived every 24hrs. There are a number of other log files for the individual infrastructure components as well.

If you wish to export these to an external syslog platform, do the following:

1. Once you have configured your syslog destination (edit `rsyslog.conf`), create a `morpheus-syslog.conf` file in the `/etc/rsyslog.d` directory and add the following entries

```
module(load="imfile" PollingInterval="50")
input(type="imfile" File="/var/log/morpheus/morpheus-ui/current" Tag="morpheus-ui"
↪ " ReadMode="2" Severity="info" StateFile="morpheus-ui")
input(type="imfile" File="/var/log/morpheus/check-server/current" Tag="check-
↪ server" ReadMode="2" Severity="info")
input(type="imfile" File="/var/log/morpheus/guacd/current" Tag="guacd" ReadMode="2"
↪ " Severity="info")
input(type="imfile" File="/var/log/morpheus/elasticsearch/current" Tag=
↪ "elasticsearch" ReadMode="2")
input(type="imfile" File="/var/log/morpheus/mysql/current" Tag="mysql" ReadMode="2"
↪ " Severity="info")
```

(continues on next page)

(continued from previous page)

```
input(type="imfile" File="/var/log/morpheus/nginx/current" Tag="nginx" ReadMode="2
↪" Severity="info")
input(type="imfile" File="/var/log/morpheus/rabbitmq/current" Tag="rabbitmq"
↪ReadMode="2" Severity="info")
```

2. Restart rsyslog

The logfiles will now be to the destination you have defined.

This configuration is valid for an ‘all-in-one’ Morpheus server. If the infrastructure components are running on separate servers /clusters, you will need to create the relevant redirects for the logs on those boxes.

Activity Log

The final log type that may require export is the Morpheus Activity log. This tracks system changes made by users, for example create and delete instances etc.

1. To set up CEF/SIEM auditing export, you should edit the following file: `logback.groovy` located at `/opt/morpheus/conf/logback.groovy`.
2. Copy the below configuration to the bottom of the `logback.groovy` configuration file, save and then exit.

```
appender("AUDIT", RollingFileAppender) {
  file = "/var/log/morpheus/morpheus-ui/audit.log"
  rollingPolicy(TimeBasedRollingPolicy) {
    fileNamePattern = "/var/log/morpheus/morpheus-ui/audit_%d{yyyy-MM-dd}.%i.log"
    timeBasedFileNamingAndTriggeringPolicy(SizeAndTimeBasedFNATP) {
      maxFileSize = "50MB"
    }
    maxHistory = 30
  }
  encoder(PatternLayoutEncoder) {
    pattern = "[%d] [%thread] %-5level %logger{15} - %maskedMsg %n"
  }
}

logger("com.morpheus.AuditLogService", INFO, ['AUDIT'], false)
```

3. Once you have done this, you need to restart the Morpheus Application server. To do this, do the following:

```
morpheus-ctl stop morpheus-ui
```

Note: Please be aware this will restart the web interface for Morpheus.

4. Once the service has stopped enter the following at the shell prompt to restart (if the service does not stop, replace stop with graceful-kill and retry)

```
morpheus-ctl start morpheus-ui
```

5. To know when the UI is up and running you can run the following command

```
morpheus-ctl tail morpheus-ui
```

Once you see the ASCII art show up you will be able to log back into the User Interface. A new audit file will have been created called audit.log and will be found in the default Morpheus log path which is `/var/log/morpheus/morpheus-ui/`

Instead of writing the output to a logfile, you could create an Appender definition for your SIEM audit database product

morpheus-ssl nginx logs

Note: Morpheus does not put a logrotate in for Morpheus-ssl access logs

svlogd will only rotate the current file, nginx is setup to write the access logs to separate files and not stdout.

Implementation of a log rotate is left up to end users for files outside of the services. This is done in case end users have a log management solution.

Below is what a suggested configuration looks like for the file `/etc/logrotate.d/morpheus-nginx`:

```
/var/log/morpheus/nginx/morpheus*access.log {
    daily
    rotate 14
    compress
    delaycompress
    missingok
    notifempty
    create 644 morpheus-app morpheus-app
    postrotate
        [ ! -f /var/run/morpheus/nginx/nginx.pid ] || kill -USR1
        ↪`cat /var/run/morpheus/nginx/nginx.pid`
    endscript
}
```

1.3.7 Backups

Morpheus built-in Backup solution provides VM, Container, Host, Database, File, Directory, Volume and Storage Provider Backup, Snapshot and Replication capabilities. Backups can be automatically configured during provisioning or manually created at any time. Backup Jobs with custom Execution Schedules and retention counts can be created and used across all environments in conjunction with configured Storage Providers. Backups can be restored over current Instances or as new Instances, and downloaded or deleted from Morpheus.

Morpheus also integrates with external services to automate availability with other providers.

Initial Backups Setup

Global Backup settings, Storage Providers and Execution Schedules should be configured prior to creating backups.

Global Backups Settings

Morpheus Backups can be enabled under *Administration -> Backups*.

Scheduled Backups When enabled, configured Backups will automatically run on the set Schedule. If disabled, backups need to be manually ran.

Create Backups When enabled, Morpheus will automatically configure backup jobs for Instances.

Backup Appliance When enabled, a Backup will be created to backup the Morpheus appliance database. Select the [Backup](#) text link to edit Appliance Backup Settings and view existing Appliance Backups.

Default Backup Storage Provider Storage Providers can be configured and managed in the *Infrastructure -> Storage* section.

Default Backup Schedule Schedules can be configured and managed in the *Provisioning -> Automation -> Execution Schedules*

Backup Retention Count Default maximum number of successful backups to retain.

Backup Schedules

Backup Execution Schedules can be configured and managed in *Provisioning > Automation > EXECUTE SCHEDULING*. An execution schedule stores only the interval at which some execution should be run. To create a new backup job with this schedule, navigate to *Backups > Backups* and click “+ADD”. In the final step of creating the backup job we are able to select any of our created execution schedules. The Default Backup Schedule set in *Administration -> Backups* will be selected when creating a backup job and not specifying an execution schedule.

Configuring Backups during Provisioning

When Backups are enabled, Backup options are presenting in the Automation tab of the Provisioning wizard.

Note: The Backup options presented in the Automation tab can be disabled using a “Create Backup” policy. See [Policies](#)

BACKUP TYPE Select the type for the Backup. Backup Types displayed will be filtered by available options per selected Instance Layout.

BACKUP NAME Defaults to Instance name

BACKUP TARGET Select Storage Provider target for the Backup (when applicable).

BACKUP JOB TYPE Create New, Clone, or Add to existing Job

JOB Name Defaults to Instance name

RETENTION COUNT Maximum number of successful backups to retain.

BACKUP SCHEDULE Select the schedule the Backup Job will be executed.

Backup Types displayed will be filtered by available options per selected Instance Layout. Backup Job Types include:

- File Backup
- Directory Backup
- MySQL
- MongoDB

- LVM Snapshot
- LVM Image
- LVM Migration
- Windows Migration
- Postgres
- Tar Directory Backup
- Amazon VM Snapshot
- VMWare VM Snapshot
- Fusion VM Snapshot
- Xen VM Snapshot
- SqlServer
- Veeam VMWare VM Backup
- Veeam Hyper-V VM Backup
- Google VM Snapshot
- Commvault File/Directory Backup
- Azure VM Snapshot
- Morpheus Appliance
- Openstack VM Snapshot
- DigitalOcean VM Snapshot
- Nutanix VM Snapshot
- Softlayer VM Snapshot
- Hyper-V VM Snapshot
- VMWare VM Snapshot
- SCVMM VM Snapshot
- UpCloud VM Snapshot
- Bluemix VM Snapshot
- Alibaba VM Snapshot
- Oracle Cloud VM Snapshot
- KVM VM Snapshot
- Container Backup
- VM Backup
- Object Storage Backup

Summary

The Backups Summary section shows the following metrics

- Number of Configured Backups trend
- Backup Success Rate
- Number of Completed Backups
- Number of Failed Backups
- Total Size of Backups (MB) trend
- Upcoming and In Progress Backups

If a User's Role permission for Backups is set to *User*, the user will only see metrics for backups they own.

Backups

In the *Backups* -> *Backups* section, currently configured Backups can be viewed and managed, and new Instance, Host and Provider backups be configured.

Note: Role permissions for Backups determine which backups will be accessible per user.

Manage an existing Backup

1. Select the Backups link in the navigation bar.
2. Select the Backups link in the sub navigation bar.
3. Select the name of the Backup to view the Backups detail page.

Create Instance Backup

To create instance backup

1. Select the Backups link in the navigation bar.
2. Select the Backups link in the sub navigation bar.
3. Click the Add Backup button.
4. From the Create Backup Wizard select the radio button Instance, then click Next.
5. Input the following:
 - Name** Name of the backup job being created.
 - Instance** Select an instance to backup from the dropdown.
6. Click Next.
7. Depending on the instance type selected in the previous step, enter additional details such as:
 - Database Name
 - Username
 - Password

- Container
 - etc..
8. Click the Next button.
 9. Schedule the backup Days, Time, Storage Provider & Retention Count.
 10. Click Complete to save.

Managing Backups

Overview

Backups are automatically configured and performed on each new Morpheus -provisioned Instance. Users can edit the frequency of backups. Administrators can define destination targets where backups are stored and perform all user-based tasks.

To View Backups:

Select the Backups link in the navigation bar.

Note: If backups are disabled, they are still created upon instance provisioning and can be executed manually. However, backups will not be executed on a schedule automatically. Scheduled backups must be enabled by an administrator to run automatically. To review how to enable/disable backups see [here](#).

Backup View

Review information about configuration such as: schedule, target details, total amount and successfully run backups, total and average size of backups from the Backup Page.

To Display Backup

1. Select the Backups link in the navigation bar.
2. Select the Backups link in the sub navigation bar.
3. Clicking the backup name to review its details.

Create Instance Backup

To create instance backup

1. Select the Backups link in the navigation bar.
2. Select the Backups link in the sub navigation bar.
3. Click the Add Backup button.
4. From the Create Backup Wizard select the radio button Instance, then click Next.
5. Input the following:

Name Name of the backup job being created.

Instance Select an instance to backup from the dropdown.

6. Click Next.
7. Depending on the instance type selected in the previous step, enter additional details such as:
 - Database Name
 - Username
 - Password
 - Container
 - etc..
8. Click the Next button.
9. Schedule the backup Days, Time, Storage Provider & Retention Count.
10. Click Complete to save.

1.3.8 Operations

Dashboard

The Dashboard is a single pane of glass showing quick, easy to read performance and configuration information about the Morpheus Environment.

Status There are four gauges across the top of the dashboard page showing quick system stats for Instances, Monitoring Status, Log Errors, and Backups. Each gauge also serves as a quick link for each section.

My Instances The My Instances section shows quick information about 5 favorite instances like Type, IP and Port. Click *View All* to be taken directly to the instances page.

Monitoring The Monitoring section displays an overall health, availability statistics, as well as response time and any open incidents requiring action.

Recent Activity Recent Activity is displayed on the right side of the dashboard page. Items like instance provisioning and deletion, backups, and alerts are displayed here.

Logs All Morpheus logs are application aware. Log information from hypervisors, servers, and applications are pushed up into the Morpheus controller node and made searchable and actionable. Choose a timeframe from the Logs pane to view statistics or click List to view all log information.

Backups The backup pane at the bottom of the page shows statistics about Morpheus backups. Information about success and failure rates and the number of backups run versus scheduled is available here. Click on the List button to be taken directly to the backups page where you can view and configure backups.

Reports

Overview

Morpheus offers 28 different report types which are designed to slice up costing and usage across Clouds, Tenants, and more. Reports can be run on-demand as needed or can be scheduled to run on certain intervals to be viewed at a later time. The list of available report types can be viewed at Operations > Reports.

The screenshot shows the Morpheus Reports interface. At the top, there's a navigation bar with the Morpheus logo, a search bar, a bell icon, a 'Support' dropdown, and a user profile 'Alex Harker'. Below this is a secondary navigation bar with tabs: Operations (selected), Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration. Under 'Operations', there are sub-tabs: Dashboard, Reports (selected), Analytics, Guidance, Wiki, Costing, Approvals, and Activity.

The main section is titled 'REPORTS' and displays three statistics: 28 REPORT TYPES, 634 REPORTS RUN, and 1 REPORTS SCHEDULED. Below these statistics is a filter bar with tabs: Report Types (selected), Results, and Scheduled. There is a search bar and a 'Select' dropdown.

The list of report types includes:

- TENANT INVENTORY SUMMARY**
Category: Account Inventory
View an inventory of VMs and containers by tenant.
Buttons: RUN NOW, SCHEDULE
- CLOUD USAGE**
Category: Cloud Usage
View a cloud's usage of storage, memory, and compute resources.
Buttons: RUN NOW, SCHEDULE
- CLOUD USAGE APP SUMMARY**
Category: Cloud Usage
View a cloud's usage of storage, memory, and compute resources by app.
Buttons: RUN NOW, SCHEDULE
- CLOUD USAGE INSTANCE TYPE SUMMARY**
Category: Cloud Usage
View a cloud's usage of storage, memory, and compute resources by instance type.
Buttons: RUN NOW, SCHEDULE
- TENANT USAGE**
Category: Cloud Usage
View a cloud's usage of storage, memory, and compute resources by tenant.
Buttons: RUN NOW, SCHEDULE
- AMAZON RESERVATION COVERAGE**
Category: Cloud Cost
View how Amazon EC2 instance types are covered by Reserved Instances.
Buttons: RUN NOW, SCHEDULE
- AMAZON RESERVATION UTILIZATION**
Buttons: RUN NOW, SCHEDULE

Report Types

ACCOUNT INVENTORY

- Tenant Inventory Summary

CLOUD USAGE

- Cloud Usage
- Cloud Usage App Summary

- Cloud Usage Instance Type Summary
- Tenant Usage

CLOUD COST

- Amazon Reservation Coverage
- Amazon Reservation Utilization
- Amazon Savings Inventory Summary
- Amazon Savings Plan Coverage
- Amazon Savings Plan Utilization
- Application Cost
- Cloud Cost
- Group Cost
- Instance Cost
- Invoice Details
- Tenant Cost
- Time Series Cost

INFRASTRUCTURE INVENTORY

- Cloud Inventory Summary
- Container Host Inventory Summary
- Group Inventory Summary
- Guidance
- Hypervisor Inventory Summary
- Migration Planning
- Tenant Resource Allocation

PROVISIONING INVENTORY

- Instance Inventory Summary
- Software Inventory
- Software Inventory By Server
- Virtual Machine Inventory Summary
- Workload Summary

By clicking into a report type, users can see any previous runs and active report schedules. New on-demand runs and new schedules of the selected report type can be created, edited, or deleted from here. The next few sections go into creating, editing, scheduling, viewing, and deleting reports in greater detail.

Create Reports

To create a new report, navigate to the report type list page (Operations > Reports). Click *RUN NOW* to the right of the specific report type to bring up the wizard to run that particular report. The required and optional fields to run the selected report type will appear, for example, the configuration panel for the Instance Cost report is shown below:

RUN INSTANCE COST REPORT [X]

START: 08/21/2020 END: 08/24/2020

GROUP: Select

CLOUD: Select

TAGS: [] +

EXCLUDE TAGS: [] +

TENANT: All

RUN

In this case, we can choose to scope the report by start and end dates, Groups, Clouds, Tenants, and can specific include or omit Instances based on tags. Once the report is run, it will be visible in the list of Instance Cost reports and all reports until deleted.

Schedule Reports

In addition to running on-demand reports, Morpheus also allows reports to be scheduled. This allows you to save report configuration and have access to refreshed information on the schedule you need.

The process of scheduling a report is nearly identical to running on on-demand. From the report type list page (Operations > Reports) click *SCHEDULE* to the right of the report type you wish to schedule. The required and optional fields to schedule the selected report type will appear, for example, the configuration panel for the Instance Cost report is shown below:

SCHEDULE INSTANCE COST REPORT

×

NAME

START

08/21/2020

END

08/24/2020

GROUP

Select

CLOUD

Select

TAGS

+

EXCLUDE TAGS

+

TENANT

All

SCHEDULE

☒ Select

☐ Date And Time

☐ Daily

☐ Daily At Midnight

☐ Hourly

☐ Weekly On Sunday At Midnight

In this case, we can choose to scope the report by start and end dates, Groups, Clouds, Tenants, and can specific include or omit Instances based on tags. Additionally, we select the time schedule on which this report should automatically run.

Note: Morpheus includes three schedules by default: Date and Time (run once at the specified time), Daily at Midnight, and Weekly on Sunday at Midnight. Any other listed scheduling periods are user-configured execution schedules (Provisioning > Automation > Execute Scheduling). Create a new execution schedule if none of the existing schedules work for your reporting needs.

Viewing Results

A list of all report runs is viewable on the Results tab of the report types list page (Operations > Reports). To view the report itself, click on the hyperlinked report filters. Only reports that are ready for viewing will have an active hyperlink on their filters. In addition to report filters, the run date, report type, creating user, and run status are shown. Click on any of these headers to filter the report list by that column in either ascending or descending order. Any report can be deleted by clicking on the trash can icon at the end of its row.

The screenshot shows the Morpheus Reports page. At the top, there's a navigation bar with the Morpheus logo, a search bar, and user information (Alex Harker). Below this is a secondary navigation bar with tabs: Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration. Under 'Operations', there are sub-tabs: Dashboard, Reports, Analytics, Guidance, Wiki, Costing, Approvals, and Activity. The 'Reports' tab is active.

The main content area is titled 'REPORTS' and displays three summary cards:

- 28** REPORT TYPES
- 636** REPORTS RUN
- 1** REPORTS SCHEDULED

Below the summary cards, there's a section for 'Report Types' with three tabs: 'Report Types', 'Results', and 'Scheduled'. The 'Results' tab is selected. This section includes a search bar, a dropdown for 'All Types', and a settings icon.

The 'Results' tab displays a table of report runs with the following columns: FILTERS, DATE RUN, REPORT TYPE, CREATED BY, and STATUS. The table contains six rows of data:

FILTERS	DATE RUN	REPORT TYPE	CREATED BY	STATUS
Aug 24, 2020	08/24/2020 12:56 PM	Tenant Inventory Summary	Alex Harker	Failed
Jul 2020 - Aug 2020 AWS Prod Service Aug 24, 2020	08/24/2020 12:55 PM	Time Series Cost	Alex Harker	Failed
Aug 24, 2020	08/23/2020 08:00 PM	Tenant Inventory Summary	Alex Harker	Ready
Aug 23, 2020	08/22/2020 08:00 PM	Tenant Inventory Summary	Alex Harker	Ready
Aug 22, 2020 All Clouds	08/22/2020 12:33 PM	Software Inventory	Alex Harker	Generating
Aug 22, 2020 All Clouds	08/22/2020 01:49 AM	Software Inventory	Alex Harker	Generating

Viewing Schedules

A list of all scheduled report runs can be viewed in the Scheduled tab of the report types list page (Operations > Reports). The friendly name of the report schedule is displayed along with the report type, last run time, next run time, and success status of the previous run. Schedules can be edited or deleted by clicking on the pencil or trash can icon, respectively. We can also view the most recent run of a given schedule (if it was successful) by clicking on the hyperlinked “last run” value.

Analytics

Overview

The Morpheus Analytics engine gives administrators the tools to break down costs and usage, then filter the results by relevant delineations including Groups, Clouds, Tenants or even tag values. Analytics dashboards can be organized into three primary categories based on their measurement intentions: costing, utilization, and workloads. Each dashboard type is discussed in further detail below.

Costing: Cloud Costing

CLOUD COST ANALYSIS

31

CLOUDS

\$4,808

MONTH TO DATE

\$7,822

PROJECTED


16

CLOUD TYPES

8


REGIONS

TOTAL COST BY CLOUD TYPE




Amazon	\$2,732.64
Azure (Public)	\$1,038.98
private	\$1,036.69

TOTAL COST BY CLOUD



AWS MorpheusDemo	\$2,732.60
Morpheus Azure	\$1,030.92
Nutanix Labs 2	\$511.96
VMware Demo Cluster	\$395.17
	\$79.82
	\$35.51

TOTAL COST BY REGION



ec2.us-west-1.amazonaws.com	\$2,732.60
westus	\$1,038.41
unassigned	\$1,036.69

Clouds

TYPE	NAME	LOCATION	REGION	MONTH TO DATE	PROJECTED TOTAL
	AWS MorpheusDemo	US West (N. Cali)	ec2.us-west-1.amazonaws.com	\$2,732.60	\$4,345.97
	Morpheus Azure	West US	westus	\$1,030.92	\$1,742.40
	Nutanix Labs 2			\$511.96	\$859.41

The Cloud Costing dashboard breaks down total costs by cloud type, cloud region, and individual cloud. This includes reporting the total number of clouds that meet your filters, the month-to-date running total, the projected monthly spend, among other useful metrics.

Filters

Filter the Clouds pulled into the dashboard by one or more of the following fields:

- Cloud (all matched by search)

- Group
- Cloud (selected from dropdown)
- Tenant
- Tag (Key)
- Value (tag value)

Data Displayed

The following aggregate totals are compiled for all Clouds that meet set filters:

- **CLOUDS:** The total number of Clouds that meet set filters
- **MONTH TO DATE:** The total spend in the current month for all Clouds meeting dashboard filters
- **PROJECTED:** The projected total spend for the current month for all Clouds meeting dashboard filters
- **CLOUD TYPES:** The number of distinct cloud types that meet the dashboard filters, such as Amazon AWS, Microsoft Azure, VMware, or any other Morpheus-supported Cloud types
- **REGIONS:** The total number of regions represented by Clouds meeting the dashboard filters

In addition to the totals described above, graphs visualize the percentage of these totals accounted for by specific Clouds, Cloud types, and Cloud regions.

Cloud List

Each Cloud that meets set filters is listed at the bottom of the dashboard, the following data points are revealed for each individual Cloud:

- **TYPE:** The Cloud type, such as Amazon AWS, Microsoft Azure, VMware, or any other Morpheus-supported Cloud type
- **NAME:** The name given to the Cloud in Morpheus at the time of integration
- **LOCATION:** The Cloud location (if available)
- **REGION:** The Cloud region (if available)
- **MONTH TO DATE:** The current month-to-date spend for the individual Cloud listed
- **PROJECTED TOTAL:** The projected total spend for the current month for the individual Cloud listed

Costing: Group Costing

GROUP COST ANALYSIS

Search

Q

Select Group

▼

Select Cloud

▼

▼ MORE

Select Tenant

▼

42
GROUPS

\$25,024
MONTH TO DATE

\$25,124
PROJECTED

TOTAL COST BY GROUP

TOP GROUPS



Groups

NAME	LOCATION	DATACENTER	MONTH TO DATE	PROJECTED TOTAL
ServiceNow			\$24,865.05	\$24,865.05
IBM			\$61.86	\$104.31
IBM			\$30.49	\$30.49
IBM			\$26.04	\$43.21
IBM	Denver		\$11.72	\$11.72
IBM	Denver		\$0.16	\$0.16

The Group Costing dashboard aggregates cost totals for all Groups that meet filters set on the dashboard. This allows administrators to sort Groups by their total costs and anticipate monthly total costs by Group.

Filters

Filter the Groups pulled into the dashboard by one or more of the following fields:

- Group (all matched by search)
- Group (selected from dropdown)
- Cloud
- Tenant

Data Displayed

The following aggregate totals are compiled for all Groups that meet set filters:

- **GROUPS:** The total number of Groups that meet set filters
- **MONTH TO DATE:** The total spend in the current month for all Groups meeting dashboard filters
- **PROJECTED:** The projected total spend for the current month for all Groups meeting dashboard filters

Group List

Each Group that meets set filters is listed at the bottom of the dashboard, the following data points are revealed for each individual Group:

- **NAME:** The name given to the Group in Morpheus at the time of creation
- **LOCATION:** The Group location (if available)
- **DATACENTER:** The Group datacenter (if available)
- **MONTH TO DATE:** The current month-to-date spend for the individual Group listed
- **PROJECTED TOTAL:** The projected total spend for the current month for the individual Group listed

Costing: Tenant Costing

TENANT COST ANALYSIS

58

TENANTS

\$4,476

MONTH TO DATE


\$7,776

PROJECTED

\$7,848

LAST MONTH

TOTAL COST BY TENANT



Tenants

NAME	DESCRIPTION	MONTH TO DATE	PROJECTED TOTAL
morpheus		\$2,388.69	\$4,227.03
		\$749.24	\$1,254.15
		\$620.15	\$1,082.77
		\$576.38	\$962.72

The Tenant Costing dashboard aggregates costing totals across all Tenants that meet the filters set on the dashboard. This information helps administrators track the current spend of each Tenant for the current monthly period. It also helps identify the costliest Tenants and to anticipate month-end costs for each individual Tenant.

Filters

Filter the Tenants pulled into the dashboard by one or more of the following fields:

- Tenant (all matched by search)

- Cloud
- Tenant (selected from dropdown)

Data Displayed

The following aggregate totals are compiled for all Tenants that meet set filters:

- **TENANTS:** The total number of Tenants that meet set filters
- **MONTH TO DATE:** The total spend in the current month for all Tenants meeting dashboard filters
- **PROJECTED:** The projected total spend for the current month for all Tenants meeting dashboard filters
- **LAST MONTH:** The total spend in the prior month for all Tenants meeting dashboard filters

Tenant List


Each Tenant that meets set filters is listed at the bottom of the dashboard, the following data points are revealed for each individual Tenant:

- **NAME:** The name given to the Tenant in Morpheus at the time of creation
- **DESCRIPTION:** The Tenant description (if available)
- **MONTH TO DATE:** The current month-to-date spend for the individual Tenant listed
- **PROJECTED TOTAL:** The projected total spend for the current month for the individual Tenant listed

Costing: User Costing

USER COST ANALYSIS



Select Group 

Select Cloud 

Current 

 MORE

Select tenant 

66

USERS


\$893

MONTH TO DATE

\$1,446

PROJECTED

TOTAL COST BY USER



	discovered	\$673.01
	...	\$61.86
	...	\$49.52
	...	\$39.72
	...	\$25.00
	...	\$14.22
	...	\$8.16
	...	\$7.47

Users

USERNAME	MONTH TO DATE	PROJECTED TOTAL
discovered	\$673	\$1,092
...	\$62	\$104

The User Costing dashboard allows administrators to analyze costs for a group of Users that meet specific filters. Once the group is selected, total costs by User for the current month and projected totals are displayed. Administrators can identify their costliest Users and anticipate the total cost by User for budgeting purposes.

Filters

Filter the Groups pulled into the dashboard by one or more of the following fields:

- User (all matched by search)
- Group
- Cloud
- Period (Current month, last three months, last six months, or last 12 months)
- Tenant

Data Displayed

The following aggregate totals are compiled for all Users that meet set filters:

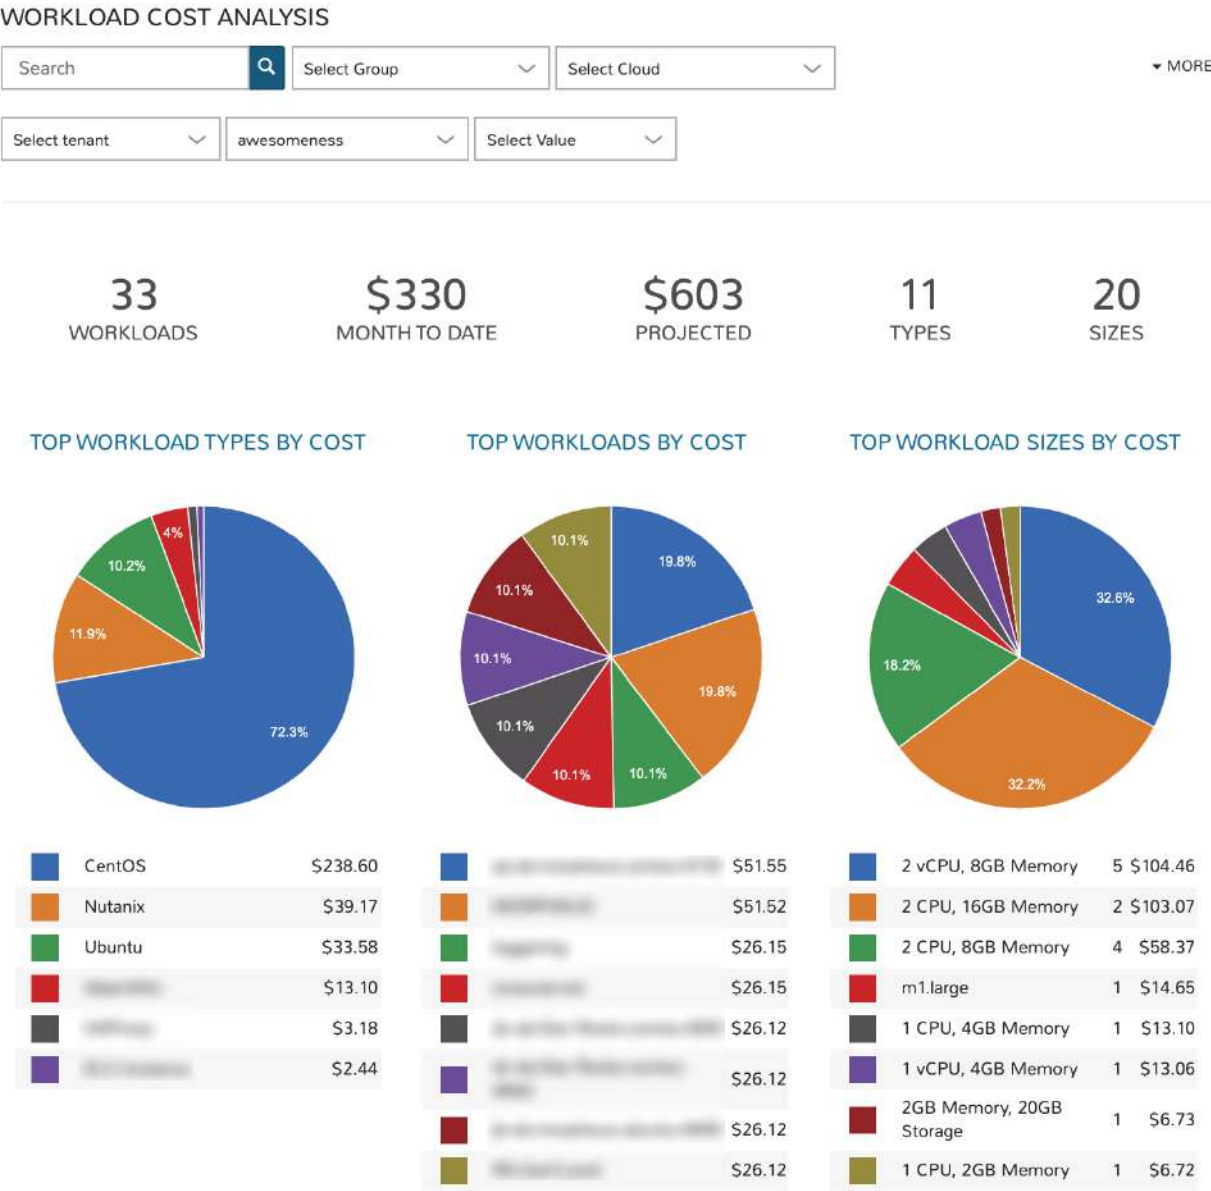
- **USERS:** The total number of Users that meet set filters
- **MONTH TO DATE:** The total spend in the current month for all Users meeting dashboard filters
- **PROJECTED:** The projected total spend for the current month for all Users meeting dashboard filters

User List

Each User that meets set filters is listed at the bottom of the dashboard, the following data points are revealed for each individual User:

- **USERNAME:** The username given to the User in Morpheus at the time of creation
- **MONTH TO DATE:** The current month-to-date spend for the individual User listed
- **PROJECTED TOTAL:** The projected total spend for the current month for the individual User listed

Costing: Workload Costing



The Workload Costing dashboard allows administrators to look at all or a subset of Morpheus-managed workloads to analyze their cost impact. Filters can be set to isolate a specific group of workloads and cost breakdowns are shown. Graphs are generated to reveal cost breakdowns of individual workloads or certain groups of workloads.

Filters

Filter the workloads pulled into the dashboard by one or more of the following fields:

- Workload name (all matched by search)
- Group
- Cloud
- Tenant

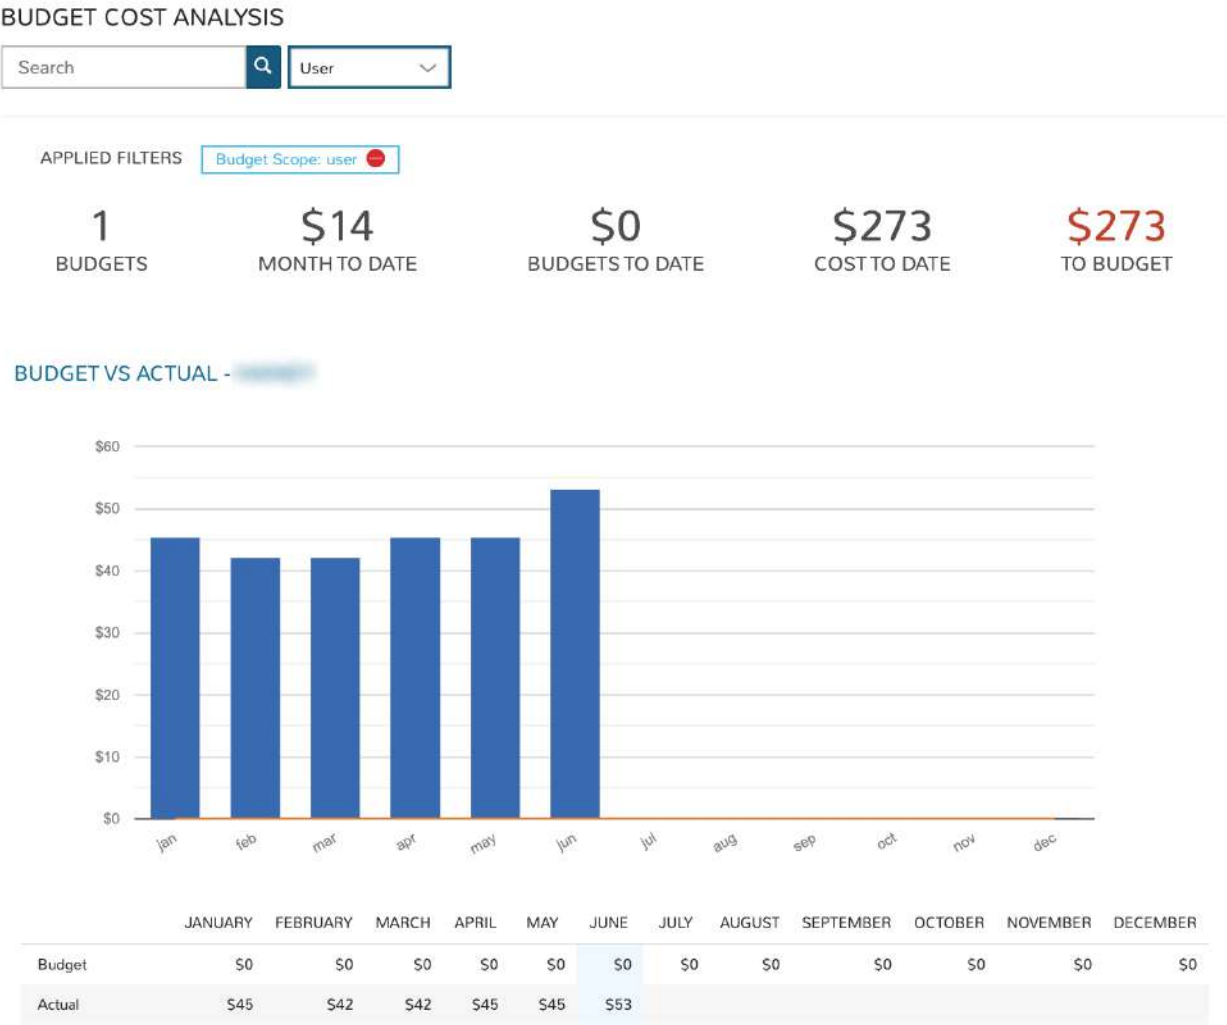
- Tag (Key). This is a required field and the top key in the list will be pre-selected by default
- Value (tag value)

Data Displayed

The following aggregate totals are compiled for all workloads that meet set filters:

- **WORKLOADS:** The total number of workloads that meet set filters
- **MONTH TO DATE:** The total spend in the current month for all workloads meeting dashboard filters
- **PROJECTED:** The projected total spend for the current month for all workloads meeting dashboard filters
- **TYPES:** The total number of workload types represented among workloads meeting set filters
- **SIZES:** The total number of unique workload sizes represented among workloads meeting set filters

Costing: Budget Analysis



The Budget Analysis dashboard allows administrators to filter and view budgets in one place in order to keep track of progress against budget over time. Budgets in Morpheus (Operations > Budgets) are tied to a specific scope (Account,

Tenant, Cloud, Group, or User) and budgets of the same scope are viewed together in this dashboard. A scope filter must be set in order for data to be populated into the dashboard. Once a scope is selected, the search bar can be utilized to return only budgets within the selected scope whose “Name” meets the search terms.

Filters

Filter the budgets pulled into the dashboard by one or more of the following fields:

- Budget name (all matched by search)
- Scope (This is a required field, data is not populated into the dashboard until a scope is specified)

Data Displayed

The following aggregate totals are compiled for all budgets that meet set filters:

- **BUDGETS:** The total number of budgets that meet set filters
- **MONTH TO DATE:** The total spend in the current month against the selected budgets
- **BUDGETS TO DATE:** The total amount budgeted to date among budgets selected by the dashboard filters (from the start of the year to the end of the current interval)
- **TO BUDGET:** The difference between the COST TO DATE and BUDGETS TO DATE value, should be close to \$0 if the costs are appropriately tracking against the budgeted amounts

Budget List

Each budget with its own graph and breakdown is displayed going down the page. The format of the information presented depends on the interval that the specific budget is configured for.

Costing: Tag Costing

TAG COST ANALYSIS

Count

Select Group

Select Cloud


▼ MORE

20
TAGS

\$1,384
MONTH TO DATE


\$2,024
PROJECTED

TOP USED TAGS



	Morpheus Server Id	37
	Morpheus Id	21
	costCenter	21
	Morpheus Instance Id	17
	Morpheus Environment	11
	...	3
	...	3
	...	3

TOP TAG VALUES



	Morpheus Server Id	21
	Morpheus Id	6
	costCenter	3
	Morpheus Instance Id	3
	Morpheus Environment	3
	...	3
	...	3
	...	3

Tags

TAG	WORKLOADS	MEMORY	STORAGE	CPU	PRICE	VALUES
Morpheus Server Id	37	4.6TiB	73.8TiB	1	\$299.32	▼ MORE
Morpheus Id	21	2.4TiB	25.4TiB	1	\$172.56	▼ MORE

The Tag Analysis dashboard creates groups of workloads based on the presence of specific tags and meeting other filters. This workload group can be analyzed for total cost and projected costs.

Filters

Filter the workloads pulled into the dashboard by one or more of the following fields:

- Tag key (all matched by search)
- Metric (apply to see the top tag values by workload count, price, memory, storage, or CPU cores)

- Group
- Cloud
- Tenant
- Tag (Key)

Data Displayed

The following aggregate totals (by tag) are compiled for workloads that meet set filters:

- **TAGS:** The total number of unique tag keys for workloads meeting dashboard filters
- **MONTH TO DATE:** The total spend in the current month for selected workloads
- **PROJECTED:** The total projected current-month spend for selected workloads

Tags List

A list of each tag (key) represented on selected workloads is displayed in a list below the dashboard graphs. We also see the total number of workloads associated with the tag, the total memory, total storage, total CPU cores, and total price. If we click the “MORE” link at the end of each row, we can see a list of all tag values associated with the key.

Utilization: Utilization vs Cost

UTILIZATION VS COST

Select Group

Select Cloud

Current

MORE

Summary

733

COUNT

34

CLOUD COUNT

\$3,823


MONTH TO DATE

\$6,452

PROJECTED

27%

AVERAGE UTILIZATION



Lowest Utilization per Cost

POWER	OS	NAME	TYPE	CLOUD	PRICE	CPU UTILIZATION
		aa-sb-morpheus-centos-4110	Managed VM	VMWare vCenter	\$51.66	<div><div>3</div>CPU</div> <div><div>6</div>MEMORY</div> <div><div>25</div>SCORE</div>
		labs_west-chef-web-01a	Discovered	AWS MorpheusDemo	\$51.80	<div><div></div></div> <div><div></div></div> <div><div></div></div>

The Utilization vs Cost dashboard is designed to reveal workloads which are underutilized (expensive and seldom-used) and which are very cost-efficient (inexpensive and frequently-used). Administrators can filter the workloads considered by the dashboard through the use of filters and potentially identify areas of cost savings by decommissioning seldom-used machines.

Filters

Filter the workloads pulled into the dashboard by one or more of the following fields:

- Workload name (all matching search terms)
- Time period (Current, one-day average, one-week average, one-month average, three-month average, six-month average, or one-year average)
- Type (virtual machines, hosts, or bare metal)
- Tenant
- Tag (Key)
- Value (Tag value)

Data Displayed

The following aggregate totals are compiled for workloads that meet set filters:

- **COUNT:** The total number of workloads that meet dashboard filters
- **CLOUD COUNT:** The total number of Clouds represented by the selected workloads
- **MONTH TO DATE:** The total spend in the current month for selected workloads
- **PROJECTED:** The total projected current-month spend for selected workloads
- **AVERAGE UTILIZATION:** The computed average utilization figure for all workloads selected by dashboard filters

Utilization List

In addition to the totals and graph displayed, two workload lists are given showing the least utilized workloads by cost (lowest utilization per cost dollar) and the least utilized workloads overall (lowest utilization overall). These workloads are listed with links to the Instance or server detail pages, along with other details related to price and resource utilization.

Workloads: Instance Type Usage

INSTANCE TYPE USAGE

Count ▾

Select Group ▾

Select Cloud ▾

▼ MORE

13

TYPES

35

INSTANCES

\$346

MONTH TO DATE

\$556

PROJECTED


165.3GB

MEMORY

1.9TB

STORAGE

TOP INSTANCE TYPES BY COUNT



Instance Type Usage

TYPE	NAME	GROUPS	CLOUDS	COUNT
CentOS	CentOS	11	5	16
HAProxy	HAProxy	3	1	3
CentOS		1	1	2
amazon web services™	EC2 Instance	1	1	2
NUTANIX	Nutanix	2	2	2

The Instance Type Usage dashboard organizes workloads meeting dashboard filters by their Instance type. In addition to counts, administrators can view resource consumption and cost figures by Instance type groupings as well.

Filters

Filter the workloads pulled into the dashboard by one or more of the following fields:

- Instance type name (all matching search terms, Morpheus-default Instance types are not included when using the search filter)
- Metric (apply to see the top Instance types by workload count, price, memory, storage, or CPU cores)
- Group

- Cloud
- Tenant
- Tag (Key)
- Value (Tag value)

Data Displayed

The following aggregate totals are compiled for workloads that meet set filters:

- **TYPES:** The total number of Instance types represented among workloads meeting the dashboard filters
- **INSTANCES:** The total number of Instances represented in the data
- **MONTH TO DATE:** The total spend in the current month for selected workloads
- **PROJECTED:** The total projected current-month spend for selected workloads
- **MEMORY:** The total memory allotted to selected workloads
- **STORAGE:** The total storage allotted to selected workloads

Instance Type Usage List

Each Instance type represented in the dashboard is listed below the graph. For each Instance type shown, we see the number of Groups the Instance type is represented in, the number of Clouds the Instance type has been provisioned into, and the total amount of memory allotted to workloads of each Instance type.

Workloads: Instance Usage

INSTANCE USAGE

Count

Select Group

Select Cloud

8

CLOUDS

36

INSTANCES

\$346

MONTH TO DATE

\$556

PROJECTED


167.3GB

MEMORY

1.9TB

STORAGE

INSTANCE USAGE BY CLOUD



Instance Usage

CLOUD	COUNT	MEMORY	STORAGE	CPU	PRICE
VMware Demo Cluster	19	72.4GiB	842.0GiB	12	\$143.94
Nutanix Labs 2	7	52.5GiB	720.0GiB	8	\$91.82
AWS MorpheusDemo	5	7.0GiB	60.0GiB	1	\$17.41
VMware Demo Cluster	1	7.5GiB	90.0GiB	4	\$15.08
Nutanix Labs 2	1	2.0GiB	30.0GiB	1	\$0.00
AWS MorpheusDemo	1	2.0GiB	10.0GiB	1	\$0.00
VMware Demo Cluster	1	8.0GiB	120.0GiB	1	\$26.23
Nutanix Labs 2	1	16.0GiB	160.0GiB	1	\$51.78

The Instance Usage dashboard shows Instance counts, resource utilization, and cost breakdowns by Cloud. Administrators can set filters to limit the workloads that are considered for dashboard analysis and then see the results given by Cloud groupings.

Filters

Filter the workloads pulled into the dashboard by one or more of the following fields:

- Instance name (all matching search terms)
- Metric (apply to see the top Clouds by workload count, price, memory, storage, or CPU cores)

- Group
- Cloud
- Tenant
- Tag (Key)
- Value (Tag value)

Data Displayed

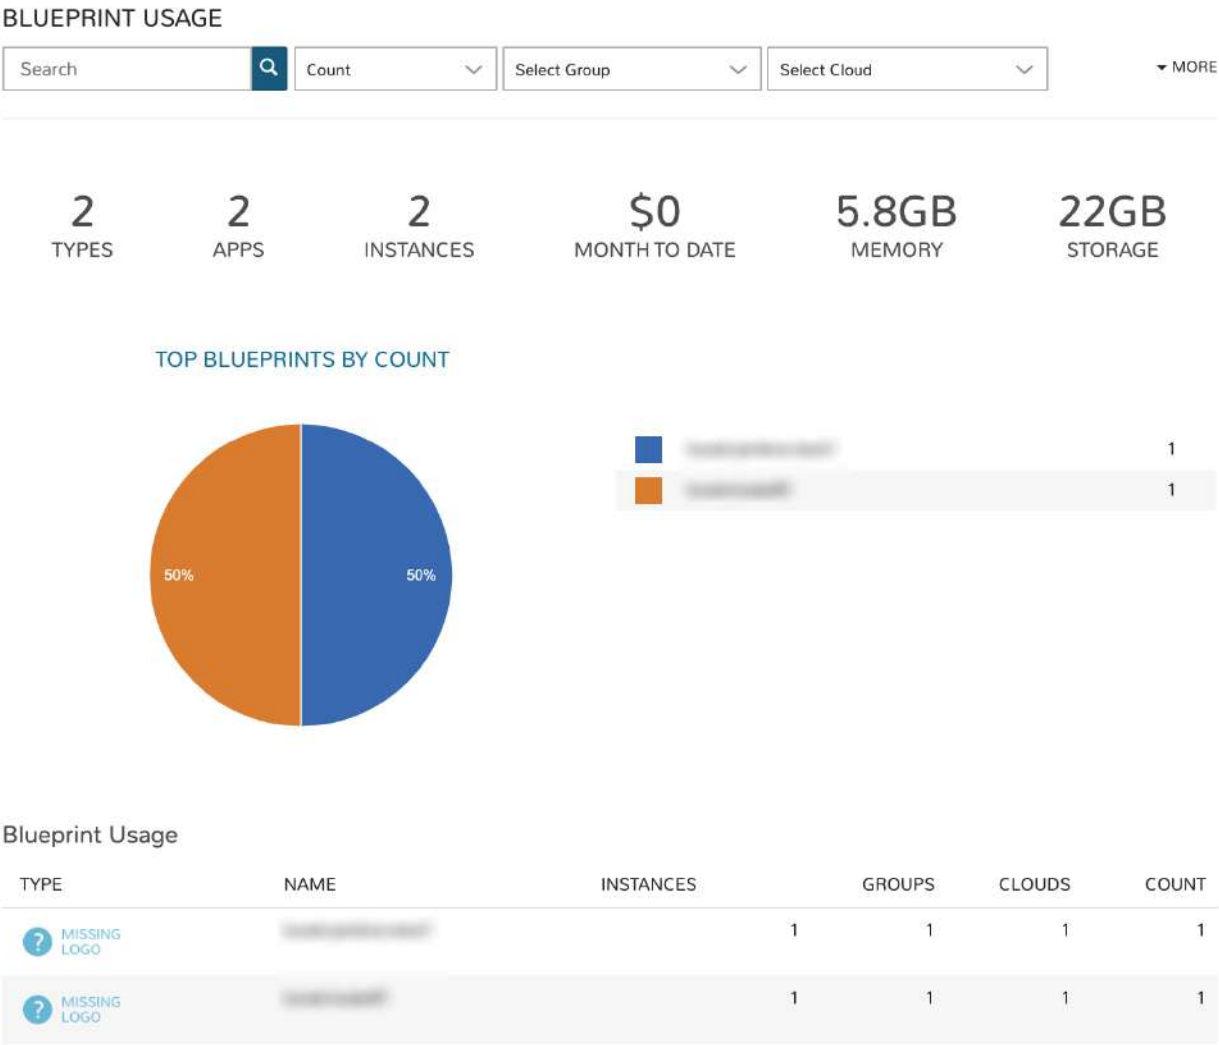
The following aggregate totals are compiled for workloads that meet set filters:

- **CLOUDS:** The total number of Clouds represented among workloads meeting the dashboard filters
- **INSTANCES:** The total number of Instances represented in the data
- **MONTH TO DATE:** The total spend in the current month for selected workloads
- **PROJECTED:** The total projected current-month spend for selected workloads
- **MEMORY:** The total memory allotted to selected workloads
- **STORAGE:** The total storage allotted to selected workloads

Instance Usage List

All Clouds represented in the dashboard are listed here. For each Cloud, we see the total Instance count, total memory allotted, total storage allotted, total CPU cores, and the total price.

Workloads: Blueprint Usage



The Blueprint Usage dashboard lists all provisioned Apps that meet filters set on the dashboard. Once the desired group of Apps is filtered into the dashboard, administrators will see the total provisioned from each Blueprint, total number of Instances created from the Apps, and costing details.

Filters

Filter the Apps pulled into the dashboard by one or more of the following fields:

- App name (all matching search terms)
- Metric (apply to see the top Clouds by workload count, price, memory, storage, or CPU cores)
- Group
- Cloud
- Tenant
- Tag (Key)
- Value (Tag value)

Data Displayed

The following aggregate totals are compiled for Apps that meet set filters:

- **TYPES:** The total number of App types represented among Apps meeting the dashboard filters
- **APPS:** The total number of Apps represented in the dashboard
- **INSTANCES:** The total number of Instances contained in all Apps meeting dashboard filters
- **MONTH TO DATE:** The total month-to-date spend for all Apps shown in the dashboard
- **MEMORY:** The total memory allotted to selected Apps
- **STORAGE:** The total storage allotted to selected Apps

Blueprint Usage List

All Blueprints which have a currently-existing App provisioned from them and selected in the dashboard filters are listed here. The name and type of the Blueprint is listed along with the total number of Instances across all provisionings, total Groups, total Clouds, and the total count of all Apps from that Blueprint.

Workloads: Apps Usage

APPS USAGE

Count ▾

Select Group ▾

Select Cloud ▾

➤ MORE

1

CLOUDS

2

APPS

3

INSTANCES

\$0

TOTAL COST


5.8GB

MEMORY

22GB

STORAGE

APP USAGE BY CLOUD



Instance Usage

CLOUD	COUNT	MEMORY	STORAGE	CPU	PRICE
VMware	2	5.9GiB	22.0GiB	1	\$0.07

The Apps Usage dashboard lists all provisioned Apps that meet a set of filters and organizes them by Cloud. Totals for cost and resource usage of all relevant Apps can be viewed with a per-Cloud breakdown.

Filters

Filter the Apps pulled into the dashboard by one or more of the following fields:

- App name (all matching search terms)
- Metric (apply to see the top Clouds by workload count, price, memory, storage, or CPU cores)
- Group
- Cloud
- Tenant
- Tag (Key)
- Value (Tag value)

Data Displayed

The following aggregate totals are compiled for Apps that meet set filters:

- **CLOUDS:** The total number of Clouds represented among Apps meeting the dashboard filters
- **APPS:** The total number of Apps represented in the dashboard
- **INSTANCES:** The total number of Instances contained in all Apps meeting dashboard filters
- **TOTAL COST:** The total cost of all selected Apps
- **MEMORY:** The total memory allotted to selected Apps
- **STORAGE:** The total storage allotted to selected Apps

Instance Usage List

All Clouds with a currently-provisioned App which is selected in the dashboard filters are listed here. The name of the Cloud is listed along with its App count, the total memory, total storage, total CPU cores and price of the Apps provisioned in that Cloud are also listed.

Guidance

Overview

The `Operations > Guidance` section shows recommendations for resource and cost optimization. By analyzing the CPU, RAM, and Storage activity of Instances and Hosts, Morpheus can recommend actions for sizing and power state. Guidance is customizable to show recommendations based on 30, 60, or 90 day periods, this dropdown toggle is visible on the Guidance list page (`Operations > Guidance`). Guidance thresholds are also customizable, they can be edited in `Administration > Settings > Guidance`.

Configuration

Guidance is configured per Cloud and is turned off by default.

To turn on Morpheus Guidance for a Cloud:

1. Navigate to *Infrastructure > Clouds*
2. Click *EDIT*
3. Expand the *Advanced Options* section in the Edit Cloud modal
4. In the *Guidance* dropdown, select *Manual*

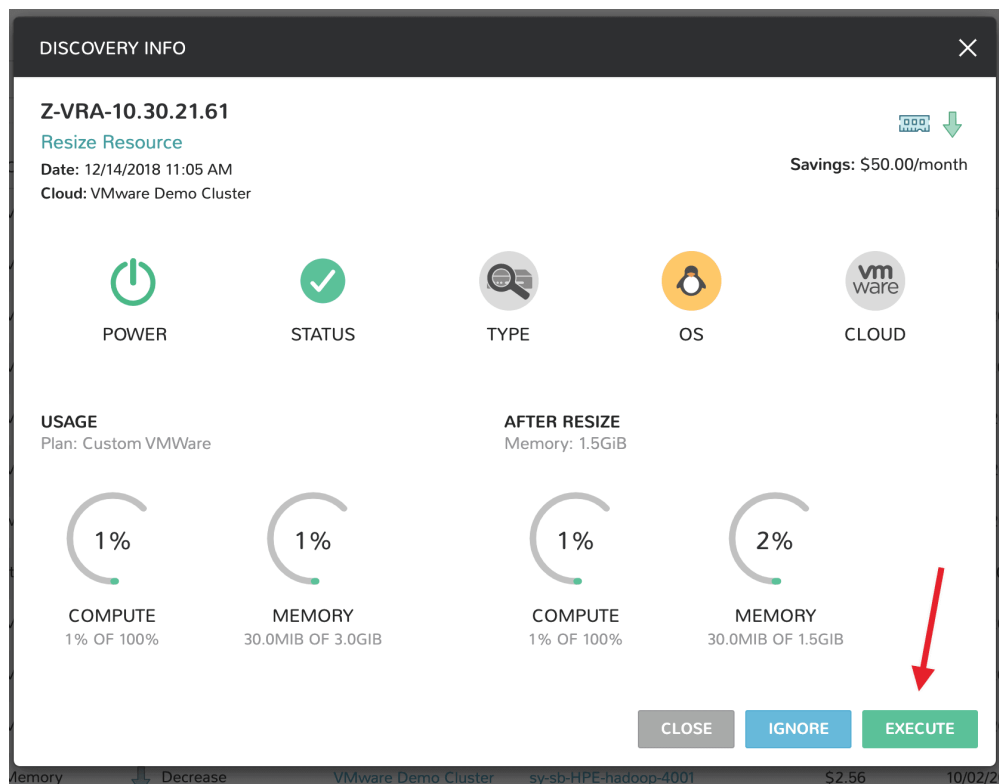
5. Click *SAVE CHANGES*

Guidance recommendations will begin to appear in the guidance section when generated.

Once Guidance has been turned on for a Cloud, Morpheus will determine if a guidance recommendation should be made once every 30 minutes. In the event that no recommendations can be made, no entry will be added to the list of suggested guidances. As the guidance list continues to grow, sorting and filtering may become necessary to focus on the recommendations that are relevant to the user at the time.

It's important to note that acting on recommendations is entirely manual at this time. In many cases, Morpheus provides one-click action to take the recommended steps but Guidance recommendations cannot be taken automatically. This is a feature that is being explored for a future update but has not been tagged for any specific release version at this time. In addition, it's recommended that Morpheus Agent be installed to maximize the benefits of Guidance. While Guidance will still work without installing the Agent, the greatly-enhanced statistics provided by the Agent will significantly improve Guidance recommendations.

To see more detail on a Guidance recommendation in your list, click on the (i) button at the far right side of each list row. This will open a detail modal that gives additional information on the Guidance entry. In some cases, such as with sizing and shutdown recommendations, the user can simply click the "EXECUTE" button to take the recommended action as shown below.



Other types of Guidance recommendations, such as reserve compute recommendations, must be taken in the cloud and Morpheus does not offer the execute button.

DISCOVERY INFO

Morpheus Azure
Reserve Compute
Date: 10/11/2019 06:05 PM

Savings: \$34.40/month

STATUS **CLOUD**

CURRENT
Current Cost: \$109.63

AFTER RESERVATIONS
Proposed Cost: \$75.23

\$109.63 **\$75.23** **\$34.40** **31.38%**

ON-DEMAND COST PROPOSED COST MONTHLY SAVINGS SAVINGS PERCENT

NAME	REGION	TERM	CURRENT COST	QUANTITY	PROPOSED COST	SAVINGS
Standard_D2	westus	P1Y	\$109.63	1	\$75.23	\$34.40

CLOSE **IGNORE**

Note: The IGNORE button will remove the recommendation from the UI. Subsequent recommendations of the same type will NOT display for the same object (VM, Cloud etc) again unless the original recommendation is resolved.

Recommendations

To view and act on Guidance recommendations, navigate to *Operations -> Guidance*.

The Guidance list contains the following details:

Severity Icon Indicates the severity of the recommended action.

Type Recommended action Type

Metric Guidance Metric used for recommended action.

Action Recommended Action for the Instance or Host, such as “Reduce Host memory” or “Shutdown Instance”

RESOURCE The Instance or Host targeted

SAVINGS Shows projected Monthly Costs savings if recommended action is taken.

DATE Date and Time stamp the recommended action was generated.

Information Link Click to view details on the recommendation.

Note: Guidance Actions are not automatically triggered at this time.

Filters

Search Search for Guidance recommendations

Type Filter by Sizing or Shutdown Guidance Types.

Severity Filter by Guidance Severity of All, Info, Warning, or Critical.

Metric Filter by All, Memory, CPU, or Power Guidance Metrics.

Wiki

The Morpheus Wiki is a tenant wide, RBAC controlled, audit-able Wiki that allows easy UI, API and CLI access to information, notes, config or any other data needed to be referenced or shared with others. Wiki pages encompass individual Clouds, Groups, Servers, Instances, Clusters, and other pages can be manually created. Wiki pages from resources are accessible from `Operations - Wiki` or in within individual resource detail pages in the Wiki tab.

- Main Wiki section is at `Operations - Wiki`
- Wiki tabs are on Clouds, Groups, Instances, Hosts, VM's, Bare Metal, and Clusters.
- Additional Wiki Pages and Categories can be created from `Operations - Wiki`.
- When a Wiki tab is populated, a Page is automatically added and accessible to `Operations - Wiki`.
- Wiki's are per Tenant. There is no multi-tenant access to Wikis.
- The Wiki is accessible from the UI, CLI and API.
- RBAC controlled via the Operations: Wiki User and Tenant Role permission (None, Read and Full).
- Page updates contain Updated by User and Date stamps.
- Wiki pages can be searched from `/operations/wiki` or navigated from `/operations/wiki-page/page-index`.
- All wiki pages are encrypted using AES-256 bit encryption.

Note: The Wiki replaces Notes. Notes are automatically migrated to corresponding Wiki pages when upgrading Morpheus to 4.0.

The Wiki service ties into assets throughout the environment. Create pages for Instances, hosts, groups, clouds, and even clusters directly on their detail pages. Or, just create general notes pages in the centralized Wiki section in Markdown format.

Creating your first page is as simple as clicking the Create Page button. Write down some content, give the page a title, and click the save button. The wiki will also keep track of who last edited a page and when. The beauty of this wiki is its clean and easy to write down notes related to various parts of your application deployment or infrastructure without going to an external tool.

All wiki pages are encrypted using AES-256 bit encryption. Though we don't advise storing passwords in a wiki document (theres services like Cypher for that), role based access control also can properly restrict access to content related to instances or hosts the user may not have access to.

To get started, simply create a page with the title Home. When that page is created, it will become the new default page for the Wiki and visible by all the users that have access to this section.

Costing

Budgets

Budgets provide insight into spending across entire accounts, allowing users to create and plan a budget scoped to their account, clouds, tenants, users, or groups.

Creating A Budget

1. Navigate to `Operations > Budget`
2. Create a new budget and enter in the following:
 1. **Name**
 2. **Description**
 3. **Scope:** Here you can choose what this budget is tied to
 4. **Period**
 5. **Year:** Set future budgets
 6. **Interval:** Choose Month, Quarter, Year then fill in the budget for that interval
3. *SAVE CHANGES*

CREATE BUDGET

×

NAME

Demo Budget

DESCRIPTION

☒ Enabled

SCOPE

Account

▼

PERIOD

Year

▼

YEAR

2019

▼

INTERVAL

Month

▼

JANUARY

2500

FEBRUARY

2500

MARCH

2500

APRIL

2500

MAY

2500

JUNE

2500

JULY

2500

AUGUST

2500

SEPTEMBER

2500

OCTOBER

2500

NOVEMBER

2500

DECEMBER

2500

SAVE CHANGES

Cloud Budgets

If you scope a budget to a cloud visit the cloud summary page in Infrastructure > Clouds > Select Cloud > Summary for a detailed breakdown of the costing

SUMMARY

HOSTS

VMS


CONTAINERS

LOAD BALANCERS

NETWORKS

RESOURCES

POLICIES



Appliance URL: default
Domain: bw.bertramlabs.com
Container Mode: docker
Network Mode: AWS

Time Zone: default
Datacenter ID:
Storage Mode: lvm
Security Mode: off

Scale Priority: 1
Agent Install Mode: cloudinit
Guidance: manual

AMAZON RESOURCES

Instances: 24

Running Instances: 19


Volumes: 31


Load Balancers: 2

Subnets: 7

Security Groups: 61


GUIDANCE

 Sizing: Memory

 Decrease

Savings: \$20.20

06/03/2018 11:05 AM



COSTING

\$2,500
BUDGET

\$7,277
MONTH TO DATE

\$11,702
ESTIMATED SPEND

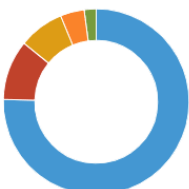
\$12,952
LAST MONTH

(\$1,250)
CHANGE FROM LAST MONTH

AWS SPEND	JANUARY	FEBRUARY	MARCH	APRIL	MAY	JUNE	JULY	AUGUST	SEPTEMBER	OCTOBER	NOVEMBER	DECEMBER
Budget	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00	\$2,500.00
Actual	\$616	\$14,629	\$12,952	\$11,702								

SERVICES BREAKDOWN

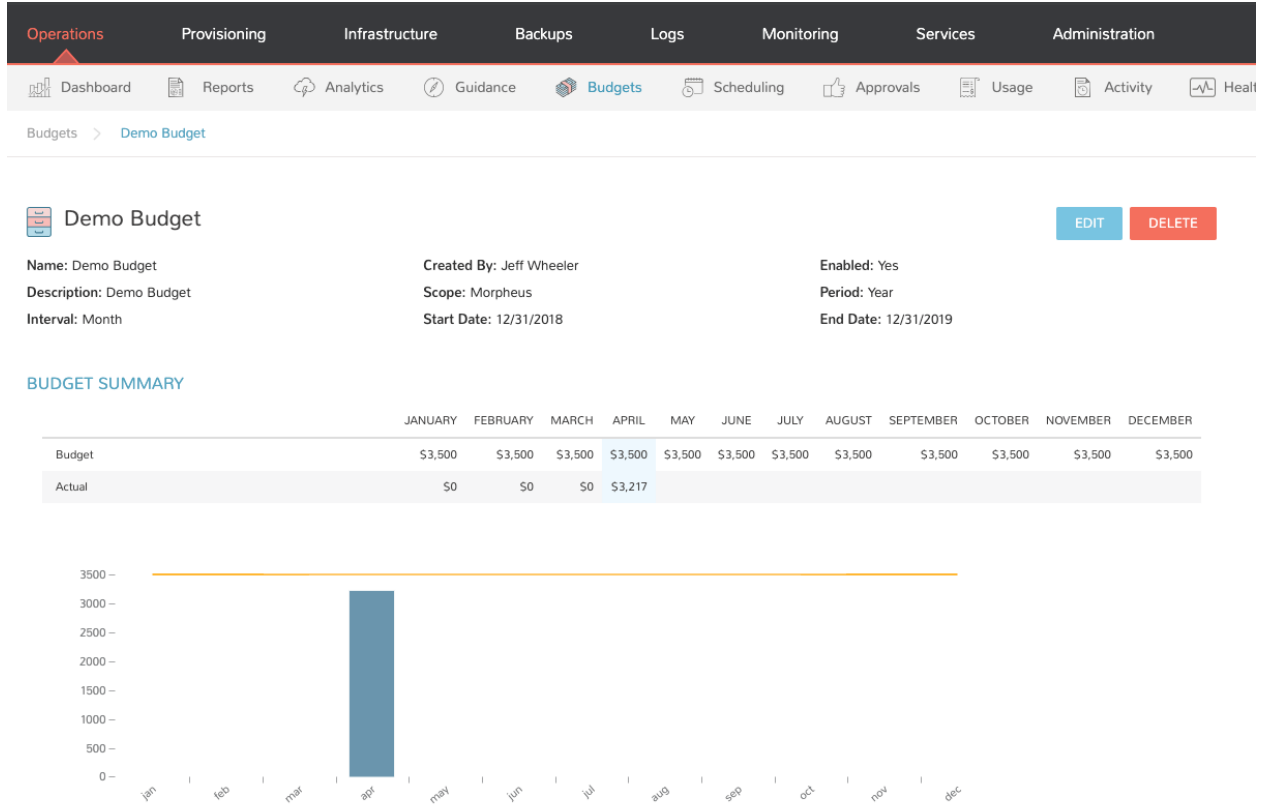
EC2	\$5,486.34
EFS	\$152.01
OTHER	\$759.81
S3	\$303.90
SUPPORT	\$574.45
TOTAL	\$7,276.51



View Budget Summary

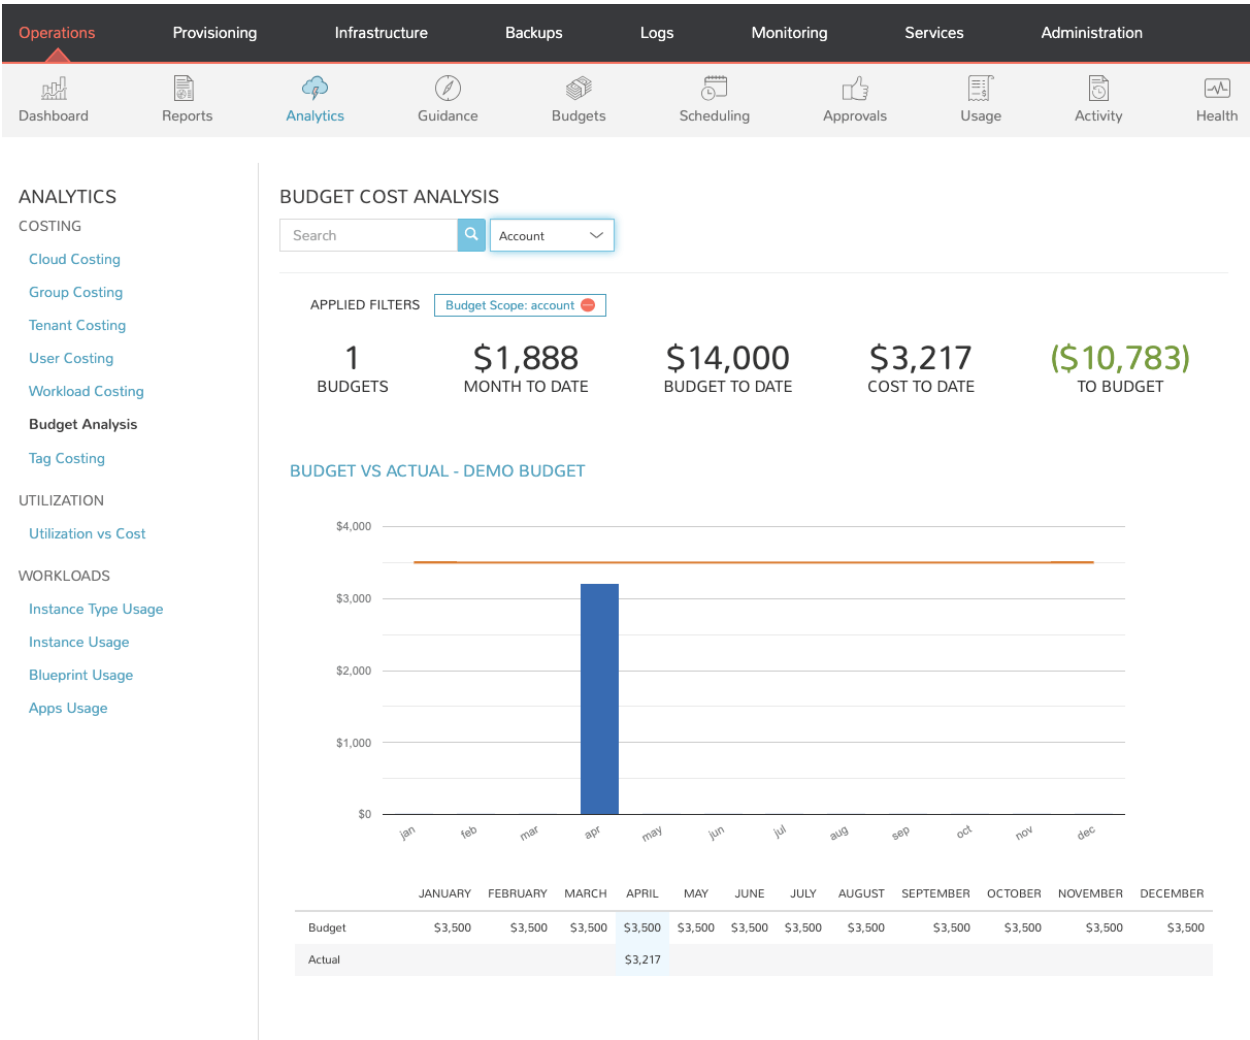
To view the budget summary, click into the budget to see the actual vs budgeted spend for the interval selected.

To edit the budget just select *EDIT*



Budget Analytics

In Operations > Analytics > Budget Analysis select scope (Account, Tenant, Cloud, Group, User) to view the budget analysis.



Invoice List Page

The invoices list page shows high-level aggregate data on a group of invoices and allows you to locate a specific invoice for a deeper look. By default, invoices from the last three months across all clouds and types are shown here. Depending on settings, this page can automatically refresh itself so the list of displayed invoices and the aggregate information on them is up to date.

The screenshot shows the Morpheus Invoices page. At the top is a navigation bar with the Morpheus logo, a search bar, and user information (Alex Harker). Below this is a secondary navigation bar with tabs for Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration. Under 'Operations', there are sub-tabs for Dashboard, Reports, Analytics, Guidance, Wiki, Costing, Approvals, and Activity. The 'Invoices' sub-tab is selected. Below the sub-tabs, there are filters for Search, All Types, All Clouds, Start Date, and End Date. A settings gear icon is also present. The summary section displays four metrics: 3 Months (PERIODS), 24,261 (TOTAL INVOICES), \$39,251 (MTD COST), and \$64,290 (TOTAL COST). Below this is a table of invoices with columns: INVOICE ID, TYPE, REFERENCE, CLOUD, PERIOD, START DATE, END DATE, MTD COST, and TOTAL COST. The table lists 10 invoices, mostly for September 2020, with various cloud providers like Nutanix, Amazon, and AWS.

INVOICE ID	TYPE	REFERENCE	CLOUD	PERIOD	START DATE	END DATE	MTD COST	TOTAL COST
1327105	Server	[REDACTED]	QA Nutanix	September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0.97	\$50.00
1327172	User	[REDACTED]		September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0	\$0
1327326	Server	[REDACTED]	QA Amazon	September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0.79	\$40.41
1327327	Server	[REDACTED]	QA Amazon	September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0.79	\$40.41
1327324	Server	[REDACTED]	QA Amazon	September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0.79	\$40.41
1327323	Server	[REDACTED]	QA Amazon	September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0.79	\$40.41
1327206	User	[REDACTED]		September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0	\$0
1327204	User	[REDACTED]		September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$0	\$0
1327303	Server	[REDACTED]	AWS Prod	September 2020	9/1/20 12:00 AM	9/30/20 11:59 PM	\$1.57	\$80.81

Aggregated Invoice Data

The following aggregate totals are displayed for all invoices picked up by set filters:

- **Periods:** The total number of months in the period determined by your start and end date filters. If no start and end dates are set, a three month period will be shown. If a one-month period is selected, the name of that month (ex. Aug 2020) is shown
- **Total Invoices:** The total number of invoices captured by current filters
- **MTD Cost:** The combined month-to-date cost for all invoices captured by current filters
- **Total Cost:** The expected total month-end cost for all invoices captured by current filters

Creating Views

Invoice list views are highly-customizable allowing for virtually limitless combinations of output columns and filtering. A common set of output columns is provided in a default view but users can add and remove columns from a large list of data points and even save multiple views for easy access to different data sets. Any of your stored views can be set as the default to be displayed each time the invoices list page is active.

To create a new invoices view:

1. Click the gear icon ()
2. Click on one or more columns from the “Columns” list to add or remove an output field
3. Click “+ add view”
4. Provide a name and a description value for your new view
5. If desired, mark the box to set the new view as your default view so it appears automatically each time the invoices list page is accessed
6. Click **SAVE**

The screenshot displays the Morpheus web interface. The top navigation bar contains the Morpheus logo and a search bar. Below it, a secondary navigation bar lists various modules: Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration. Under 'Operations', there are sub-links for Dashboard, Reports, Analytics, Guidance, Wiki, Costing, Approvals, and Activity. The main content area is titled 'Invoices' and shows a summary for '3 Months' with '273' total invoices, '\$668' MTD cost, and '\$699' total cost. A table below this summary lists individual invoices with columns: INVOICE ID, CLOUD, TYPE, REF ID, PERIOD, SERVER ID, INSTANCE ID, SERVER, and INSTANCE. A gear icon in the top right corner of the table area opens a configuration menu. This menu has three sections: 'Options' with a checked 'auto refresh' checkbox; 'Views' with 'default' selected, '+ add view', and 'manage views'; and 'Columns' with a list of available columns: 'invoice id', 'type', 'ref id', 'reference', 'cloud id', and 'cloud'. The 'Columns' section is highlighted with a red rectangular box.

Available Output Columns

When creating an invoices view, there are many output columns available to select. Consult the list below for more details on each of the available columns:

Available Output Columns: **Expand for Complete List**

- **Invoice ID:** The unique ID in Morpheus for the invoice
- **Type:** The invoice type; Cloud, Container, Group, Server, Instance, Resource, User, or Volume
- **Ref ID:** An ID for the reference object tied to the invoice (server, instance, cloud, etc.). Reference IDs are reused across invoice types so invoices referring to identical Ref IDs may not necessarily refer to the same reference object
- **Reference:** The name of the reference object (server, cloud, user, group, etc.) tied to the invoice

- **Cloud ID:** The internal ID for a Cloud integration in Morpheus. This field will be blank unless the invoice references a Cloud
- **Cloud:** The name for a Cloud integration in Morpheus. This field will be blank unless the invoice references a Cloud
- **Instance ID:** The internal ID for an Instance in Morpheus. This field will be blank unless the invoice references an Instance
- **Instance:** The name for an Instance in Morpheus. This field will be blank unless the invoice references an Instance
- **Server ID:** The internal ID for a server in Morpheus. This field will be blank unless the invoice references a server
- **Server:** The name for a server in Morpheus. This field will be blank unless the invoice references a server
- **Cluster ID:** The internal ID for a cluster in Morpheus. This field will be blank unless the invoice references a cluster
- **Cluster:** The name for a cluster in Morpheus. This field will be blank unless the invoice references a cluster
- **Plan ID:** The internal ID for a service plan in Morpheus. This field will be populated only for invoices that reference an object which would be associated with a service plan (server, Instance, container, etc.).
- **Plan:** The name for a service plan in Morpheus. This field will be populated only for invoices that reference an object which would be associated with a service plan (server, Instance, container, etc.).
- **Group ID:** The internal ID for a Group in Morpheus. This field will be blank unless the invoice references a Group
- **Group:** The name for a Group in Morpheus. This field will be blank unless the invoice references a Group
- **User ID:** The internal ID for a User in Morpheus. This field will be blank unless the invoice references a User.
- **User:** The name for a User in Morpheus. This field will be blank unless the invoice references a User.
- **Tenant ID:** The internal ID for the Morpheus Tenant which owns the reference object
- **Tenant:** The name of the Morpheus Tenant which owns the reference object
- **Period:** The monthly period during which the invoice was generated
- **Interval:** The length of the invoice billing period, currently all invoices are generated at a one-month interval
- **Start Date:** The start date and time for the invoice period, typically the first of the month at midnight
- **End Date:** The end date and time for the invoice period, typically the last day of the month at midnight
- **Ref Start:** The date and time the reference object is created or the start of the invoicing period if the reference object existed prior to the start of the invoicing period
- **Ref End:** The date and time the reference object is decommissioned or the end of the invoicing period if the reference object still existed at the end of the period
- **Compute Cost:** The actual compute costs for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Storage Cost:** The actual storage costs for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Network Cost:** The actual network costs for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Extra Cost:** The actual additional costs for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)

- **MTD Cost:** The actual month-to-date costs for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Total Cost:** The actual total costs for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Metered Compute Cost:** Compute costs determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Storage Cost:** Storage costs determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Network Cost:** Network costs determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Extra Cost:** Additional costs determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered MTD Cost:** Month-to-date costs determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Total Cost:** Total costs determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Compute Price:** The actual compute price (cost plus markup) for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Storage Price:** The actual storage price (cost plus markup) for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Network Price::** The actual network price (cost plus markup) for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Extra Price:** The actual additional price (cost plus markup) for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **MTD Price:** The actual month-to-date price (cost plus markup) for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Total Price:** The actual total price (cost plus markup) for the invoice (from public cloud costing API when available, otherwise mirrored metered cost)
- **Metered Compute Price:** Compute price (cost plus markup) determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Storage Price:** Storage price (cost plus markup) determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Network Price:** Network price (cost plus markup) determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Extra Price:** Additional price (cost plus markup) determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered MTD Price:** Month-to-date price (cost plus markup) determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Metered Total Price:** Total price (cost plus markup) determined by Morpheus usage and pricing data (when live pricing data from a public cloud is not available, such as with an on-prem cloud)
- **Active:** Indicates whether or not the reference object is currently existing and active
- **Date Created:** The date and time the invoice is created
- **Last Updated:** The date and time the invoice was last updated

Invoice Types

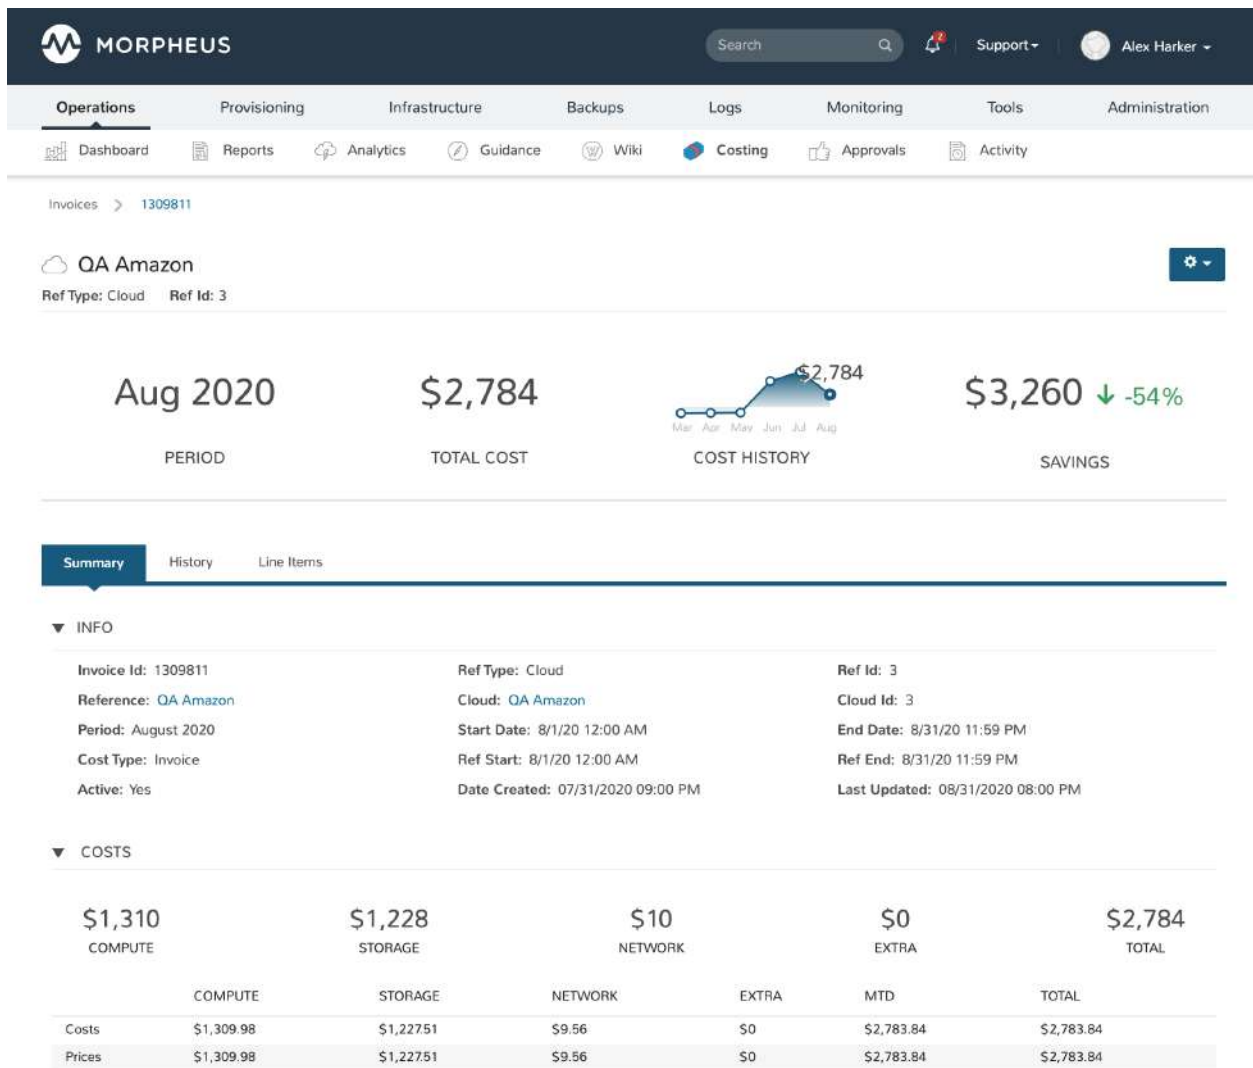
Invoices can reference any of the Morpheus workload element types or resource reference types in the list below. Some invoice types are broader and may account for resource costs calculated in other narrower invoice types. For instance, a container-type invoice returns costs for a single node of an Instance while an Instance-type invoice for the same period may be including costs for that same single node. The invoices list view can be filtered to show just one type or all types. Complete descriptions of each invoice type are included below:

- **Cloud:** In Morpheus, a Cloud is any connection into a public cloud, private cloud, hybrid cloud, or bare metal server
- **Container:** A single node of a service, in other words, a single node of a Morpheus Instance. This could be a virtual machine or Docker container which is part of a Morpheus-managed Instance
- **Group:** In Morpheus, Groups define which resources a user has access to through their role. Clouds are added to Groups and users access Clouds to which their roles give access
- **Server:** A server refers to any individual host, virtual machine, or bare metal server that is inventoried or managed by Morpheus. This can include servers which are parts of Morpheus-managed Instances or inventoried servers from integrated Clouds
- **Instance:** A set of containers or virtual machines which correlate to a single horizontally-scalable entity. This could be a single VM or it could be many VMs operating as a service
- **Resource:** Resource-type invoices are generated when Morpheus cannot determine that the referenced costs belong to any of the other resource reference types in this list
- **User:** User-type invoices aggregate the costs of resources owned by a specific Morpheus user during the invoicing period
- **Volume:** When possible, costs will be tied to known volumes and a volume-type invoice is generated as a result

Invoice Detail Page

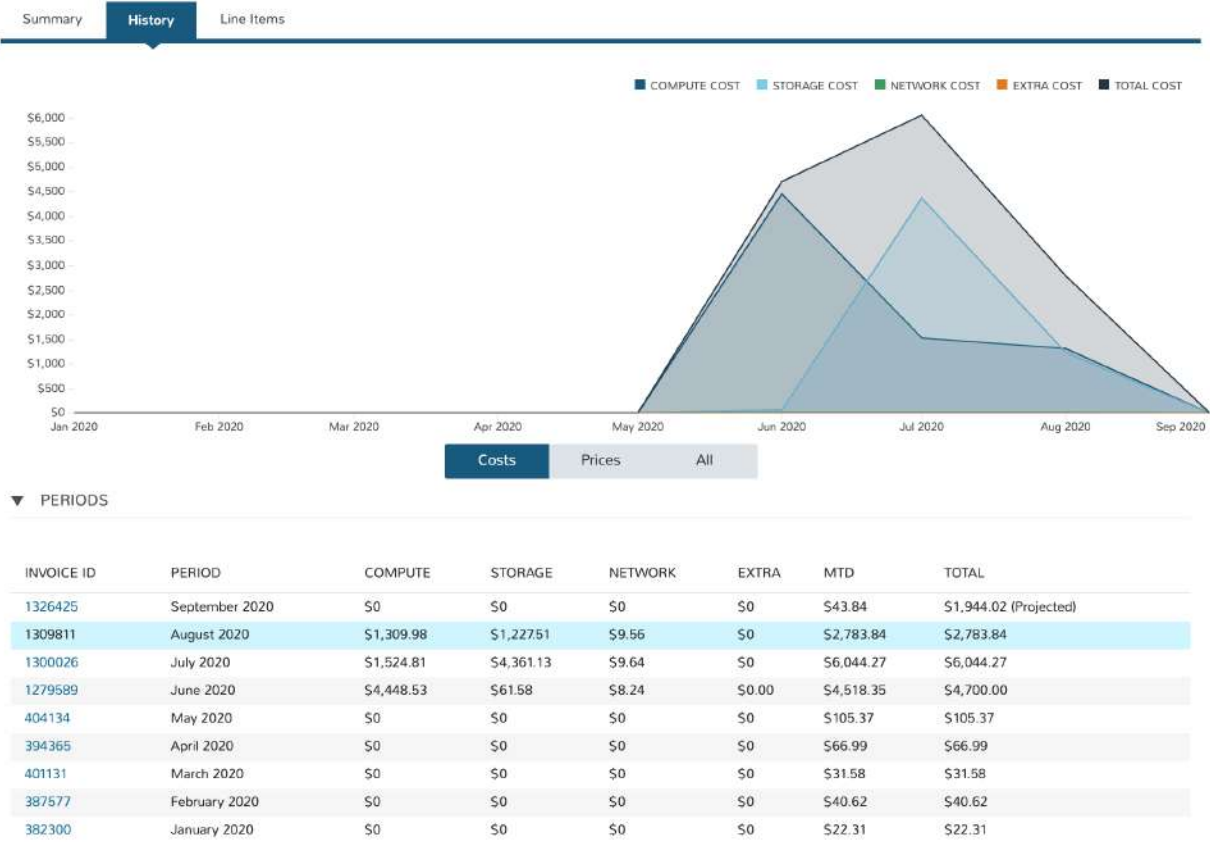
Summary

The summary tab of the invoice detail page displays a great deal of reference information about the resource identified by the invoice. This will vary depending on the type of resource. In addition, total and projected costs are displayed along with cost breakdowns for compute, storage, network, and other categories. Month-to-date totals and final month projections are given.



History

The history tab displays the costs and prices for the associated resource over time. This tab is especially valuable for resources that have existed through at least a few invoicing periods to show changes over time. In addition, cost breakdowns for compute, storage, network, and other categories are shown for each invoicing period. These costs can be displayed visually through graphs.



Line Items

For supported resource types, Morpheus includes a tab to display all costing line items. This provides the user with a complete list of line items that make up the costing totals on the invoice.

Summary History Line Items								
ID	START DATE	END DATE	USAGE TYPE	USAGE	RATE	UNIT	COST	ITEM NAME
1612407	8/1/20 12:00 AM	8/30/20 1:00 PM	USW1-CW-Requests	1.0078647E7	1.0E-5	Requests	\$100.79	QA Amazon
1614136	8/1/20 6:00 PM	8/1/20 7:00 PM	USW1-DataTransfer-Out-Bytes	9.4892E-6	0.0897631338	GB	\$0.00	QA Amazon
1614137	8/1/20 6:00 PM	8/1/20 7:00 PM	USW1-DataTransfer-Out-Bytes	7.7188E-6	0.0897631338	GB	\$0.00	QA Amazon
1614138	8/1/20 6:00 PM	8/1/20 7:00 PM	USW1-LCUUsage	1.833333E-4	0.007999995	LCU-Hrs	\$0.00	QA Amazon
1614139	8/1/20 6:00 PM	8/1/20 7:00 PM	USW1-LoadBalancerUsage	1.0	0.028	Hrs	\$0.03	QA Amazon
1614140	8/1/20 6:00 PM	8/1/20 7:00 PM	USW1-LoadBalancerUsage	1.0	0.0252	Hrs	\$0.03	QA Amazon
1619670	8/1/20 7:00 PM	8/24/20 10:00 PM	us-west-1-KMS-Requests	39.0	3.0E-6	Requests	\$0.00	QA Amazon
1617343	8/3/20 3:00 PM	8/30/20 1:00 PM	USW1-LoadBalancerUsage	646.0	0.0252	Hrs	\$16.28	QA Amazon
1620003	8/5/20 4:00 PM	8/5/20 5:00 PM	USW1-DataProcessing-Bytes	6.1E-7	0.008	GB	\$0.00	QA Amazon
1620006	8/5/20 4:00 PM	8/5/20 5:00 PM	USW1-LCUUsage	2.055556E-4	0.007999995	LCU-Hrs	\$0.00	QA Amazon
1620008	8/5/20 4:00 PM	8/5/20 5:00 PM	USW1-LoadBalancerUsage	1.0	0.0252	Hrs	\$0.03	QA Amazon
1620005	8/5/20 4:00 PM	8/5/20 6:00 PM	USW1-DataTransfer-Out-Bytes	6.9551000000000005E-6	0.0897631338	GB	\$0.00	QA Amazon
1620004	8/5/20 4:00 PM	8/6/20 7:00 AM	USW1-DataTransfer-Out-Bytes	3.4650800000000004E-5	0.0897631338	GB	\$0.00	QA Amazon
1620007	8/5/20 4:00 PM	8/6/20 7:00 AM	USW1-LoadBalancerUsage	2.0	0.028	Hrs	\$0.06	QA Amazon
1620726	8/6/20 6:00 AM	8/6/20 7:00 AM	USW1-DataTransfer-Out-Bytes	7.5875E-6	0.0897631338	GB	\$0.00	QA Amazon
1620727	8/6/20 6:00 AM	8/6/20 7:00 AM	USW1-LCUUsage	2.055556E-4	0.007999995	LCU-Hrs	\$0.00	QA Amazon
1620728	8/6/20 6:00 AM	8/6/20 7:00 AM	USW1-LoadBalancerUsage	1.0	0.0252	Hrs	\$0.03	QA Amazon
1621067	8/6/20 3:00 PM	8/6/20 4:00 PM	USW1-DataTransfer-Out-Bytes	2.2820000000000002E-7	0.0897631338	GB	\$0.00	QA Amazon
1623759	8/11/20 8:00 AM	8/11/20 9:00 AM	USW1-DataTransfer-Out-Bytes	1.08387E-5	0.0897631338	GB	\$0.00	QA Amazon
1623762	8/11/20 9:00 AM	8/11/20 9:00 AM	USW1-LoadBalancerUsage	1.0	0.028	Hrs	\$0.03	QA Amazon

Usage

The *Operations > Costing > Usage* section shows billing information for Instances and hosts that have pricing configured on their Service Plan.

Important: Pricing must be enabled in *Administration > Settings > Provisioning* and Service Plans configured with price sets in *Administration > Plans & Pricing* for pricing to show in the Usage section.

View Usage

All Instances are listed by default, with the most recent usage information showing first.

Usage details can be filtered by Cloud and date:

Cloud Default view is for all Clouds. Select a Cloud to show Instance and host usage for only one Cloud.

Date Default view shows most current Usage. Select the Date filter to scope to a different date range.

API & CLI

Usage information can also be extracted via the Morpheus API and CLI, including the ability to extract usage per Tenant.

Note: Appropriate Role permissions for *Operations: Usage* are required to view the Usage section.

Approvals

Morpheus and Service Now Approvals

Overview

Policies can be created for Groups and Clouds to require approvals for actions with the built-in Morpheus approvals engine, or via a ServiceNow integration. Approvals can be configured for Provisioning and Lifecycle extensions.

Configuring Approvals

Configuring Morpheus for Approvals

To configure Morpheus for approvals:

1. Configure Roles for Approval access
2. Optionally configure a ServiceNow Integration for ServiceNow approvals.
 - Please note ServiceNow integration is not required for Internal Approvals.
3. Create approvals policies for:
 - Internal Approvals
 - SNOW Approvals

Configure Roles

Configure User Role access settings in Administration -> Roles -> (Role) -> Operations: Approvals.

- All Users with a Role applied containing Operations: Approvals set to Full will have approval authority, and be able to Approve, Deny or Cancel approval requests.
- All Users with a Role applied that has Operations: Approvals set to Read will be able to view Approval requests and history, but will not be able to Approve, Deny or Cancel approval requests.
- All Users with a Role applied that has Operations: Approvals set to None will not have access to the Operations: Approvals section, and such will not be able to see or act on approval requests.
- Regardless of Role settings, any instance or app provisioned by any user to a group or cloud with an active Approval policy applied will require approval before the instance or app will provision.

ServiceNow Approvals

Configure ServiceNow integration for SNOW Approvals

1. Navigate to Admin -> Integrations
2. Select + **NEW INTEGRATION**
3. Select **ServiceNow** from the Type dropdown in the Integration modal and enter:
 - **Name** Name of the integration in Morpheus
 - **Enabled** Leave checked to enable the integration.
 - **Host** URL of the ServiceNow host (ex: <https://ven0000.service-now.com>)

- **User** A User in ServiceNow that is able to access the REST interface and create/update/delete incidents, requests, requested items, item options, catalog items, workflows, etc.
- **Password** Password for User above

4. Save Changes

Morpheus then configures the integration with ServiceNow, syncs ServiceNow workflows which are available when creating approvals policies. (This process can take up to 5 minutes depending on the size of the workflow table in ServiceNow.)

Create Approval Policies

- Policies applied to a Group are created in Infrastructure -> Groups -> (group) -> Policies tab.
- Policies applied to a Cloud are created in Infrastructure -> Clouds -> (cloud) -> Policies tab.

To create an Approval policy:

1. Navigate to the Policies tab in the Group or Cloud to which the policy will apply.
2. Select + ADD POLICY to open the New Policy wizard
3. Select Provision Approval from the Type dropdown
4. Add an optional description
5. **Leave Enabled selected for this Policy to be active once saved.**
 - Enabled can be deselected to disable to policy.
6. In the config section, select either Internal Approvals or ServiceNow Approvals:
 - **Internal Approvals** Approval requests will be managed within Morpheus via the Operations: Approvals section.
 - **ServiceNow Approvals** Approval requests will be managed with ServiceNow (SNOW). Please note a ServiceNow integration (Admin: Integrations) must be configured prior to SNOW Approval policy generation.
 - For ServiceNow Approvals, select the appropriate ServiceNow workflow for this policy. Please note the workflows presented are created in ServiceNow and synced with Morpheus .
7. Add the Morpheus Accounts to which this policy will apply, or leave the Accounts field blank to apply to all accounts.
8. Save

Upon saving, a new policy is created in the Group or Cloud Policies tab.

Note: SNOW Approvals will take a few moments to save as the policy is generated.

Managing Approval Requests

Once Instance approval policies are added to a Group or Cloud, any Instance or App provisioned into that Group or Cloud will create an approval request entry in the *Operations -> Approvals* section.

Note: User Role permission *Operations: Approvals -> FULL* is required to manage Approvals.

- To Approve, Deny, or Cancel an internal approval request, select the request and use the Actions dropdown.
- To Cancel a ServiceNow Approval request, select the request and use the Actions dropdown. ServiceNow approvals are managed in ServiceNow.

Note: Instances requiring provisioning approval will have a PENDING status until approved.

Each Approval Request will have:

- Name: A name for the approval in Morpheus
- Monthly Price (Est.): The estimated monthly price of the Instance to be provisioned
- Request Type: What is being requested, such as Instance approval
- External Name: For ServiceNow integrations, the name of the approval in the ServiceNow console
- Type: Internal or ServiceNow
- Status
- Date Created
- Requested By: The Morpheus user making the request
- **Actions dropdown (for internal approval requests)**
 - Approve
 - Deny
 - Cancel
- **Actions dropdown (for ServiceNow requests)**
 - Cancel

Note: The Approvals list view can be sorted by NAME, REQUEST TYPE, EXTERNAL NAME, DATE CREATED, and REQUESTED BY

Internal approval requests

To Approve, Deny or Cancel an Internal approval request:

1. Navigate to *Operations -> Approvals*
2. Select the Name of the Approval request
3. Select Actions on the far right of the request
4. Select Approve, Deny, or Cancel from the Actions dropdown

5. Select OK on the confirmation modal

- When an Internal request is approved, the related instance will begin to provision immediately and the request will show approved.
- When an Internal request is denied, the related instances status will change to Denied and the request will show Rejected in the Approvals section.
- When an Internal request is canceled, the related related instances status will change to Cancelled and the request will be canceled.

ServiceNow Approval requests

ServiceNow approval request are managed in ServiceNow. The process of approving or rejecting requests is determined by the ServiceNow Workflow selected when configuring the SNOW Approval policy. These Workflows are configured in ServiceNow.

Important: Morpheus syncs with ServiceNow every 5 minutes. Once an Approval Request is Approved or Rejected in Service Now, it will take up to 5 minutes for the instance to respond accordingly, and the status for the approval request in the Approvals section in Morpheus to update.

Activity

The Activity section displays a recent activity report for Auditing. Morpheus defines an activity as any major action performed on an instance or server, such as, but not limited to adding a server, deleting a server, provisioning an instance, deleting an instance, creating a backup, etc. . . This view can be searched and filtered by type, user, and date range.

Activity

There are 5 types of activities that are displayed in the Activity Reports:

- Provisioning
- Monitoring
- Alert
- Backups
- Logs

To View a Recent Activity report:

1. Select the Reports link in the navigation bar.
2. Click the tab Recent Activity.

Recent activity is displayed in order from recent to oldest. This view can be searched and filtered by type, user, and date range.

Review

To review the item the activity occurred on, click the name of the activity and it will go to a new page and display that item.

Note: Deleted activities are displayed as an alert and do not contain a link to the event item. If the activity is not a deletion event we provide a link on the activity name to go to the item the activity occurred on.

To Filter:

1. Click the filter drop down of type of filter you want to apply.
2. Select the appropriate filter.

Alarms

The *ALARMS* section shows Operation notifications from Cloud and other Service Integrations. Cloud and other Service Integration Alarms are not generated by Morpheus but synced and displayed for visibility in Morpheus.

Morpheus Logs

The Logs displayed in `Operations - Health - Morpheus Logs` are from `/var/log/morpheus/morpheus-ui/current`. These logs show all ui activity and are useful for troubleshooting and auditing.

Note: Stack traces in `Operations - Health - Morpheus Logs` are filtered for Morpheus services. Complete stack traces can be found in `/var/log/morpheus/morpheus-ui/current`.

History

1.3.9 Tools

Cypher

Overview

Cypher at its core is a secure Key/Value store. But what makes cypher useful is the ability to securely store or generate credentials to connect to your instances. Not only are these credentials encrypted but by using a cypher you don't have to burn in connection credentials between instances into your apps.

Cypher keys can be revoked, either through lease timeouts or manually. So even if somebody were to gain access to your keys you could revoke access to the keys and generate new ones for your applications.

Keys can have different behaviors depending on the specified mountpoint.

Mountpoints

password Generates a secure password of specified character length in the key pattern (or 15) with symbols, numbers, upper case, and lower case letters (i.e. password/15/mypass generates a 15 character password).

tfvars This is a module to store a tfvars file for terraform app blueprints.

secret This is the standard secret module that stores a key/value in encrypted form.

uuid Returns a new UUID by key name when requested and stores the generated UUID by key name for a given lease timeout period.

key Generates a Base 64 encoded AES Key of specified bit length in the key pattern (i.e. key/128/mykey generates a 128-bit key)

- Key lease times are entered in milliseconds and default to 32 days (2764800000 ms).
 - Quick MS Time Reference:
 - Day: 86400000
 - Week: 604800000
 - Month (30 days): 2592000000
 - Year: 31536000000

Creating Cypher Keys

1. Navigate to Services - Cypher and select “+ ADD KEY”
2. Configure one of the following types of Keys:

Password

A Cypher password generates a secure password of specified character length in the key pattern (or 15) with symbols, numbers, upper case, and lower case letters (i.e. password/15/mypass generates a 15 character password).

Key Pattern “password/character_length/key”

Example: password/10/mypassword

Value Leave the Value field blank for a password, as it will be generated.

Lease Enter lease time in milliseconds (ex. 604800000 for one week)

Save changes and the password will be generated and available for use.

If your user role has Cypher: Decrypt permissions, a “DECRYPT” button will be available in the Cypher section to view the generated password.

To delete the password key, select *Actions* -> *Remove* and confirm.

Secret

A Cypher secret is the standard secret module that stores a key/value in encrypted form.

Key Pattern “secret/key”

- EXAMPLE: secret/mysecret

Value Add the secret value to be encrypted

Lease Enter lease time in milliseconds (ex. 604800000 for one week)

Save changes and the secret will be encrypted and available for use.

If your Morpheus user role has Cypher: Decrypt permissions, a “DECRYPT” button will be available in the Cypher section to view the secret.

To delete the secret, select *Actions* -> *Remove* and confirm.

UUID

A Cypher UUID Returns a new UUID by key name when requested and stores the generated UUID by key name for a given lease timeout period.

Key Pattern “uuid/key”

- Example: uuid/myuuid

Value Leave the Value field blank for UUID, as it will be generated.

Lease Enter lease time in milliseconds (ex. 604800000 for one week)

Save changes and the UUID will be generate and available for use.

If your user role has Cypher: Decrypt permissions, a “DECRYPT” button will be available in the Cypher section to view the generate UUID.

To delete the UUID, select *Actions* -> *Remove* and confirm.

Key

A Cypher Key generates a Base 64 encoded AES Key of specified bit length in the key pattern (i.e. key/128/mykey generates a 128-bit key).

Key Pattern “key/bit_length/key”

- Example: key/256/mykey

Value Leave the Value field blank for key, as it will be generated.

Lease Enter lease time in milliseconds (ex. 604800000 for one week)

Save changes and the AES Key will be generate and available for use.

If your user role has Cypher: Decrypt permissions, a “DECRYPT” button will be available in the Cypher section to view the generate AES Key.

To delete the UUID, select *Actions* -> *Remove* and confirm.

Using Cypher Keys in Scripts

To use a cypher Key in a script, use the following syntax:

```
<%=cypher.read('var_name') %>
```

Example: `PASSWORD=<%=cypher.read('secret/myuserpassword') %>`

Note: You can reference the original owner of a workflow so that keys can be used in a sub-tenant. Example `PASSWORD=<%=cypher.read('secret/myuserpassword') %>` could be changed to `PASSWORD=<%=cypher.read('secret/myuserpassword',true) %>` within a library or a workflow and the true means OWNER true. This will keep that key in the master tenants cypher store.

Archives

Overview

Archives provides a way to store your files and make them available for download by your Scripts and Users. Archives are organized by buckets. Each bucket has a unique name that is used to identify it in URLs and Scripts.

Storage Provider

Archive buckets are assigned a Storage Provider (Object Store). This is where the bucket will write its files. A Storage Provider can be configured to use the local appliance file system (Local), an Amazon S3 bucket, etc.

Every archive bucket generates and uses a random File Path to store its files under. This ensures two different archive buckets will not contend for the same backend storage location.

Permissions

Visibility

Visibility determines whether your files are secure or not.

Private This secures your files. Only authorized users of the Owner and Tenants account may view the bucket and download its files. This is the default.

Public This makes your files available to the public. Anyone, including anonymous users/scripts can download these files without any authentication.

Warning: Be careful not to store sensitive files in a Public archive.

Users of the Owner account may fully manage the files in a bucket.

Tenants

Users of the Owner account may fully manage the files in a bucket. Users of the Tenant account(s) will have read-only access. They may browse and download files in the bucket.

Both Owner and Tenants must have the Services: Archives permission to access a Private bucket. READ level access allows browsing and downloading files in the bucket.

FULL access allows full management of the bucket and its files. This includes modifying files and links, bucket settings and deleting it.

Files

To add a file to a bucket, click on the bucket name, and then click the + ADD FILE button. Once added, click on the file name to access the links, history and script section for the file.

Links

You can create a Link to download a Private file without any authentication. Links may be configured to expire after a period of time.

Scripts

Morpheus automatically generates syntax for creating a link to a file in your Scripts. When the Script is generated, it will create a temporary link to download the file and return the URL of that link. This link is made available to the public. It is accessible to any user or script that can reach the appliance. Downloading the file only requires knowing the URL, which includes a secret token parameter. You can specify the number of seconds before the link expires. The default value is 1200 (20 minutes).

Image Builder

The Morpheus Image Builder tool creates vmdk, qcow2, vhd and raw Images from scratch. The Image Builder creates a blank VM in VMware, attaches an OS ISO, executes a boot script on the VM at startup via VNC which calls a preseed script which runs the unattended OS installation and configuration. Morpheus then executes an OVA export of the completed vmdk to target Storage provider, and converts the image to all other specified formats. The new Virtual Image records are automatically added to Morpheus and the Images are then available for use.

Requirements

- **DHCP must be enabled on the network specified for the VM in Morpheus, and network settings configured for DHCP in M**
The blank VM must get network configuration via DHCP upon boot. Static IP assignment is not possible.
- **Hypervisor Console must be enabled on the Target Cloud** Morpheus utilizes VNC to pass the boot script to the VM.
- **VM must be able to reach and resolve the Morpheus appliance URL over 443** The boot script calls to the Morpheus appliance to get the preseed script.
- **Valid Linux ISO set for the Virtual Image. Windows is not supported.** The ISO can exist in the target cloud or be uploaded to Morpheus

Note: `cloud-init` enabled must be disabled on the iso Virtual Image settings.

- **Access to target ESXi host(s) over 443 and ESXi hostname dns resolution**
 - Same requirement as Hypervisor Console and Image upload/download to/from vCenter.
- Valid Boot Script
- Valid Pre-seed script
- Valid Storage Provider configure for ova export of generated image.

Sample Scripts

Sample Boot Script

```
<wait5><tab> text ks=<%=preseedUrl%><enter>
```

Note: `<%=preseedUrl%>` is a Morpheus variable that will populate with the Morpheus appliance url.

Sample Preseed Script

```
# CentOS 7.x kickstart file - ks.cfg
#
# For more information on kickstart syntax and commands, refer to the
# CentOS Installation Guide:
# https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/
# ↪Installation_Guide/sect-kickstart-syntax.html
#
# For testing, you can fire up a local http server temporarily.
# cd to the directory where this ks.cfg file resides and run the following:
# $ python -m SimpleHTTPServer
# You don't have to restart the server every time you make changes. Python
# will reload the file from disk every time. As long as you save your changes
# they will be reflected in the next HTTP download. Then to test with
# a PXE boot server, enter the following on the PXE boot prompt:
# > linux text ks=http://<your_ip>:8000/ks.cfg
#
# Required settings
lang en_US.UTF-8
keyboard us
rootpw password
authconfig --enablesshadow --enablemd5
timezone UTC
#
# Optional settings
install
cdrom
user --name=cloud-user --plaintext --password password
unsupported_hardware
network --bootproto=dhcp
```

(continues on next page)

(continued from previous page)

```
firewall --disabled
selinux --permissive
bootloader --location=mbr --append="biosdevname=0 net.ifnames=0"
text
skipx
zerombr
clearpart --all --initlabel
autopart --type=plain
firstboot --disabled
reboot

%packages --nobase --ignoremissing --excludedocs
openssh-clients
# Prerequisites for installing VMware Tools guest additions.
# Put in kickstart to ensure first version installed is from install disk,
# not latest from a mirror.
kernel-headers
kernel-devel
gcc
make
perl
curl
wget
bzip2
dkms
patch
net-tools
git
# Core selinux dependencies installed on 7.x, no need to specify
# Other stuff
sudo
nfs-utils
open-vm-tools
-fprintd-pam
-intltool
-biosdevname

# unnecessary firmware
-aic94xx-firmware
-atmel-firmware
-b43-openfwfw
-bfa-firmware
-ipw*-firmware
-irqbalance
-ivtv-firmware
-iwl*-firmware
-libertas-usb8388-firmware
-ql*-firmware
-rt61pci-firmware
-rt73usb-firmware
-xorg-x11-drv-ati-firmware
-zd1211-firmware
%end

%post
# configure vagrant user in sudoers
echo "%cloud-user ALL=(ALL) NOPASSWD: ALL" >> /etc/sudoers.d/cloud-user
```

(continues on next page)

(continued from previous page)

```

chmod 0440 /etc/sudoers.d/cloud-user
cp /etc/sudoers /etc/sudoers.orig
sed -i "s/^\(.*requiretty\)$/#\1/" /etc/sudoers
# keep proxy settings through sudo
echo 'Defaults env_keep += "HTTP_PROXY HTTPS_PROXY FTP_PROXY RSYNC_PROXY NO_PROXY"' >>
↪ /etc/sudoers
%end

```

Migrations

Migration Types

Hypervisor to Hypervisor

Store Morpheus will create a snapshot of existing VM and upload the snapshot to virtual image directory. Images that have been uploaded to the Virtual Images library can be converted to VHD, QCOW2, RAW and VMDK formats and then re-provisioned.

New Morpheus will create a snapshot of an existing VM, convert from source format to required destination format, and then provision the VM into the target environment.

Source VMWare, Openstack, Xen, Nutanix* Azure* Hyper-V* (*in-development)

Destination Softlayer, Openstack, Metapod, Xen, Amazon, VMWare, ESXi, Nutanix, Hyper-V Supported OS Type: Windows or Linux

Service Impact Disruptive Migration

Virtual Image Extract

The Virtual Image extract capabilities allow for a virtual image to be extracted and stored in the virtual image repository or the image can be migrated into a cloud.

Source Any Cloud

Destination SoftLayer (Only)

Supported OS Type Windows

Service Impact Non Disruptive

Requirements Requires a separate disk or network share to store the image during conversion process. Capacity of the disk or network share should be sized appropriately to support the data that will be exported.

Live Stream

Note: Live Stream is deprecated

Live Stream is a linux only streaming process that will take a snapshot of a volume and allow it to be streamed to a destination linux system that is either existing or new. The destination linux must already exist and it can either be a managed or unmanaged VM in Morpheus . The destination will be overwritten from a root level perspective.

Source Any Cloud

Destination Morpheus

Supported OS Type Linux (Only)

Service Impact Non Disruptive

Requirements Requires the Linux host/guest to be configured for LVM and that free space of the capacity to be streamed is available. A destination linux host/guest must be available to receive the stream.

Add Migration

1. Select the Tools link in the navigation bar.
2. Select the Migrations link in the sub-navigations bar.
3. Click the Add Migration button.
4. From the Create Migration Wizard select the type of migration, then click the *Next* button.

Depending on the Migration Type selected input the following, then click the *Next* button.

- Hypervisor to Hypervisor * Select Cloud, and Server * Input Host, Remote Port, Username, and Password
 - Virtual Image Extract * Select Platform, Existing or New, Cloud, and Server. * Input Host, WinRM Port, WinRM User, WinRM Password, and Snapshot path.
 - Live Stream * Select Platform, Existing or New, Cloud, and Server * Input Host, SSH Port, SSH User, SSH Password, Public Key, and Logical Volume Device. * Enter Destination details, then click the *Next* button.
5. Finalize your configuration if needed, then click the complete button.

Manually Start Migration

If you chose to not run your migration in the Create Migration Wizard then you will be able to manually start the migration.

1. Select the Tools link in the navigation bar.
2. Select the Migrations link in the sub-navigations bar.
3. Click the actions dropdown of the row of the migration you wish start, and select Run.

Remove Migration

1. Select the Tools link in the navigation bar.
2. Select the Migrations link in the sub-navigations bar.
3. Click the actions dropdown of the row of the migration you wish remove, and select Remove.

VMware to AWS Migration

Requirements

When performing a Hypervisor to Hypervisor migration from VMware to AWS, there are some requirements that must be met:

1. Add S3 Storage Provider to Morpheus
2. Set Image Transfer Store in you AWS cloud(s) settings in Morpheus
3. Create VM Import Service roles in your AWS account (not in Morpheus)
4. Storage Provider selected for migration destination must be set as a Local Storage Provider (not AWS)

Add S3 Storage Provider

An AWS S3 bucket is required for VMware - AWS migrations. S3 buckets created in AWS are automatically synced into Morpheus. S3 buckets can also be created from Morpheus from Infrastructure -> Storage -> Buckets

Set Image Transfer Store

Under Infrastructure -> Clouds, select your AWS cloud and click *EDIT*. Expand the Advanced Options section and for *IMAGE TRANSFER STORE* select the target AWS S3 Bucket and then Save.

Add VM Import Service

Tip: Refer to the AWS document below to add the required VM Import Service role in AWS: <http://docs.aws.amazon.com/vm-import/latest/userguide/import-vm-image.html>

VM Import requires a role to perform certain operations in your account, such as downloading disk images from an Amazon S3 bucket. You must create a role named `vmimport` with a trust relationship policy document that allows VM Import to assume the role, and you must attach an IAM policy to the role.

To create the service role

Create a file named `trust-policy.json` with the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": { "Service": "vmie.amazonaws.com" },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:Externalid": "vmimport"
        }
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
}
}
]
}
```

You can save the file anywhere on your computer. Take note of the location of the file, because you'll specify the file in the next step.

Use the `create-role` command to create a role named `vmimport` and give VM Import/Export access to it. Ensure that you specify the full path to the location of the `trust-policy.json` file.

```
aws iam create-role --role-name vmimport --assume-role-policy-document file://trust-
↪policy.json
```

Create a file named `role-policy.json` with the following policy, where `disk-image-file-bucket` is the bucket where the disk images are stored:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::disk-image-file-bucket"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::disk-image-file-bucket/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:ModifySnapshotAttribute",
        "ec2:CopySnapshot",
        "ec2:RegisterImage",
        "ec2:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

Use the following `put-role-policy` command to attach the policy to the role created above. Ensure that you specify the full path to the location of the `role-policy.json` file.

```
aws iam put-role-policy --role-name vmimport --policy-name vmimport --policy-document ↪
↪file://role-policy.json
```

For more information about IAM roles, see IAM Roles in the IAM User Guide.

Storage Providers

Set the “Storage Provider” in the migration wizard destination as a Local Storage type, or leave as Select to use the Morpheus Appliance.

A local image must be created by Morpheus prior to S3 upload. A Local Storage provider can be used if one had been added in the Infrastructure -> Storage -> File Shares section. Simply leaving the Storage Provider setting as “select” will create an image on the Morpheus appliance, provided sufficient storage existing on the Morpheus appliance in /tmp.

Important: Setting AWS as the Destination Storage Provider will result in a migration failure.

These settings will allow a successful migration from VMware to AWS using the Morpheus migration wizard.

Self Service

The Self Service catalog (Tools > Self Service) is where administrators can create easily-deployable items for consumption by users operating under the “Service Catalog” Persona in Morpheus. Catalog items can be fully-configured Morpheus Instances or Blueprints, complete with user input through Morpheus Option Types, automation Workflows, and more. The catalog items are presented in a simplified interface for ease of deployment without sacrificing configurability for administrators. All available catalog items are built in the Self Service area and users will see relevant items in their catalogs based on Role permissions.


Note: For more on Personas and using the Service Catalog persona, see the Personas sections of our documentation.

Access is granted to the Self Service section through the Tools: Self Service Role permission. Users with Read rights can view the catalog while users with full rights can create and edit catalog items. Users without any rights to Self Service will not be able to access the page at all.

Additionally, a user’s Role determines access to the standard and/or service catalog persona and which service catalog items are available for a user under the service catalog persona. See the Roles section of Morpheus documentation for more information on administering Roles.

Viewing the Self Service Catalog

The complete Self Service catalog can be viewed by clicking on Self Service from the Tools menu. The complete list of items available for the Self Service catalog are shown here but users working in the Service Catalog Persona will see only those allowed based on their user role. In addition to the name and type of each catalog item, we can also see a description and whether not the catalog item is featured or active. Featured items are given special visibility in the Service Catalog Persona and inactive items will not appear as provisioning options.



Catalog Items

Search		Q	+ ADD	⚙
TYPE	NAME	FEATURED		
Resource	A BP 1		/	🗑
Resource	A BP 2		/	🗑
Resource	A BP 3		/	🗑
APACHE	Apache Docker		/	🗑
NGINX	Custom	✓	/	🗑
MISSING LOGO	Custom CentOS		/	🗑
NGINX	Fully Loaded		/	🗑
APACHE	lll-lanarha-armv		/	🗑

Building Catalog Instances

An Instance in Morpheus is a set of one or more containers or virtual machines that correlate to a single, horizontally-scalable entity or service suite. From the Self Service section, we can pre-configure Morpheus Instances and present them to users viewing the Service Catalog Persona for one-click deployment.

From the Catalog Items List Page (Tools > Self Service), click **ADD**. From the dropdown menu, select **Instance**. The modal window will appear to configure and add a new catalog Instance.

ADD CATALOG ITEM

×

NAME

Apache on Ubuntu

DESCRIPTION

Description

☒

ENABLED

☒


FEATURED

VISIBILITY

Private

⌵

LOGO



⌵

CONFIG

CONFIGURATION WIZARD

Option Types

Search option types

SAVE CHANGES

Configure the following:

- **NAME:** A friendly name for the catalog item in Morpheus
- **DESCRIPTION:** An optional description identifying the catalog item
- **ENABLED:** When checked, this catalog item will be available for provisioning
- **FEATURED:** When checked, this catalog item will be given special visibility in the Service Catalog Persona view
- **VISIBILITY:** Set to private to keep the catalog item available only to users in the current Tenant. Master Tenant

administrators may set catalog items to public to make them viewable and usable by Subtenant users

- **LOGO:** Select or upload a logo to be associated with this catalog item
- **CONFIG:** Enter, view, or edit Instance config here. Click *CONFIGURATION WIZARD* to build this catalog item through the Morpheus Add Instance wizard
- **CONTENT:** Optionally include documentation content for this Catalog Item. Markdown-formatted text is accepted and displayed appropriately when the item is ordered from the Service Catalog. A new Catalog Item-type Wiki entry will also be added containing this information.
- **OPTION TYPES:** If desired, select Option Types to present users with mandatory or optional selections prior to provisioning

Once done, click *SAVE CHANGES*

Tip: Building catalog items through the configuration wizard is similar to the typical provisioning process for Instances in Morpheus. For more details on selections available in the configuration wizard, take a look at other sections of Morpheus docs on provisioning Instances.

Building Catalog Blueprints

Morpheus Blueprints allow for full multi-tier application deployment. In the Self Service catalog, user can create catalog items based on pre-existing App Blueprints. If new Blueprints need to be created for the Service Catalog, see other sections of Morpheus docs on building App Blueprints of various supported types. Just like with catalog Instances, we can pre-configure Blueprints and present them to users viewing the Service Catalog Persona view for easy, one-click deployment.

From the Catalog Items List Page (Tools > Self Service), click *ADD*. From the dropdown menu, select Blueprint. The modal window will appear to configure and add a new catalog Blueprint.

Configure the following:

- **NAME:** A friendly name for the catalog item in Morpheus
- **DESCRIPTION:** An optional description identifying the catalog item
- **ENABLED:** When checked, this catalog item will be available for provisioning
- **FEATURED:** When checked, this catalog item will be given special visibility in the Service Catalog Persona view
- **VISIBILITY:** Set to private to keep the catalog item available only to users in the current Tenant. Master Tenant administrators may set catalog items to public to make them viewable and usable by Subtenant users
- **LOGO:** Select or upload a logo to be associated with this catalog item
- **BLUEPRINT:** Select a pre-configured Blueprint (Provisioning > Blueprints) to associate with this catalog item
- **APP SPEC:** Inject App spec here for any fields required to provision an App from your Blueprint. You may also inject any overrides to the existing Blueprint spec that are desired. App Spec configuration must be YAML, a simple example that names the App and sets the Group and Cloud is included below:

```
#Example App Spec

name: '<%= customOption.appName %>'
group:
  name: Dev Group
environment: Dev
```

(continues on next page)

(continued from previous page)

```
tiers:
  Web:
    instances:
      - instance:
          type: nginx
          cloud: Dev AWS
  App:
    instances:
      - instance:
          type: apache
          cloud: Dev AWS
```

- **CONTENT:** Optionally include documentation content for this Catalog Item. Markdown-formatted text is accepted and displayed appropriately when the item is ordered from the Service Catalog. A new Catalog Item-type Wiki entry will also be added containing this information.
- **OPTION TYPES:** If desired, select Option Types to present users with mandatory or optional selections prior to provisioning

Note: App spec custom option variables should be single quoted in YAML: `cloud: '<%= customOption.cloud %>'`

Once done, click *SAVE CHANGES*

Building Catalog Workflows

From the Catalog Items List Page (Tools > Self Service), click *ADD*. From the dropdown menu, select Workflow. The modal window will appear to configure and add a new catalog Workflow.

Configure the following:

- **NAME:** A friendly name for the catalog item in Morpheus
- **DESCRIPTION:** An optional description identifying the catalog item
- **ENABLED:** When checked, this Workflow item will be available for selection in the Service Catalog
- **FEATURED:** When checked, this catalog item will be given special visibility in the Service Catalog Persona view
- **VISIBILITY:** Set to private to keep the catalog item available only to users in the current Tenant. Master Tenant administrators may set catalog items to public to make them viewable and usable by Subtenant users
- **LOGO:** Select or upload a logo to be associated with this catalog item
- **WORKFLOW:** Select an existing Workflow to be associated with this Catalog Item, new Workflows are created in Provisioning > Automation
- **CONTEXT TYPE:** Optionally restrict users to a specific target context, Instance, Server, or None
- **CONTENT:** Optionally include documentation content for this Catalog Item. Markdown-formatted text is accepted and displayed appropriately when the item is ordered from the Service Catalog. A new Catalog Item-type Wiki entry will also be added containing this information.

Once done, click *SAVE CHANGES*

Editing and Deleting from the Self Service Catalog

Once created, Service Catalog items can be edited or deleted from the Catalog Items list view (Tools > Self Service). Click the pencil icon in the relevant row to edit the Service Catalog item or the trash can icon to delete it. Alternatively, Service Catalog items can be made inactive to remove them as provisioning options rather than deleting them.

1.3.10 Personas

Personas are alternate views in Morpheus UI. A user's access to certain Personas is controlled by Role permissions. At present, there are two Persona types: Standard and Service Catalog. The Standard Persona is the typical default view. The Service Catalog Persona is a simplified view where users are presented with pre-configured Instance types and Blueprints to choose from based on their Role. The rest of this section goes through the use of the Service Catalog Persona and how administrators can curate the selection their users see in this area.

Configuring Persona Access

Access to Personas is controlled by a user's Role. Additionally, Persona access can be configured on the Tenant Role to set maximum Persona access for any user in the Tenant. By default, new Roles and Roles that existed prior to the creation of Personas will only have access to the Standard Persona. If desired, new Roles can be configured to have access to one or both Personas and existing Roles can be edited in the same way.

Tip: It's recommended to set access to both Personas to "None" if you intend not to use Personas at all. Under this configuration, Morpheus gives access only to the Standard Persona and hides the Persona selection menu from the user. New Roles and Roles that existed prior to creation of the Personas feature are pre-configured in this way.

Edit Persona access on a Role with the following steps:

1. Navigate to Administration > Roles
2. Select the desired Role to edit
3. Go to the Personas tab
4. Allow access to one or both Personas as needed, changes are saved automatically

The screenshot displays the Morpheus UI interface for configuring role permissions. The top navigation bar includes the Morpheus logo, a search bar, a notification bell, a support link, and the user profile 'Morpheus User'. The main navigation menu has tabs for Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration. The 'Administration' tab is active, showing a sub-menu with 'Roles' selected. The 'Roles' page shows the 'Engineer' role selected. The 'PERMISSIONS' section for the 'Engineer' role has the 'Personas' tab selected. Below the tabs, there is a 'Default Persona' dropdown set to 'Select'. A table lists the personas and their access levels:

NAME	ACCESS
Standard	Full
Service Catalog	Full

Configuring Catalog Item Access

Within the Service Catalog Persona, users are presented with Catalog Items based on their User Role. Additionally, Catalog Item access can be set on the Tenant Role to restrict certain items from all users in the Tenant. By default, User Roles have no access to any catalog items (and no access to the Service Catalog Persona). When enabling Service Catalog Persona access for User Roles, you will also need to give access to some or all Catalog Items.

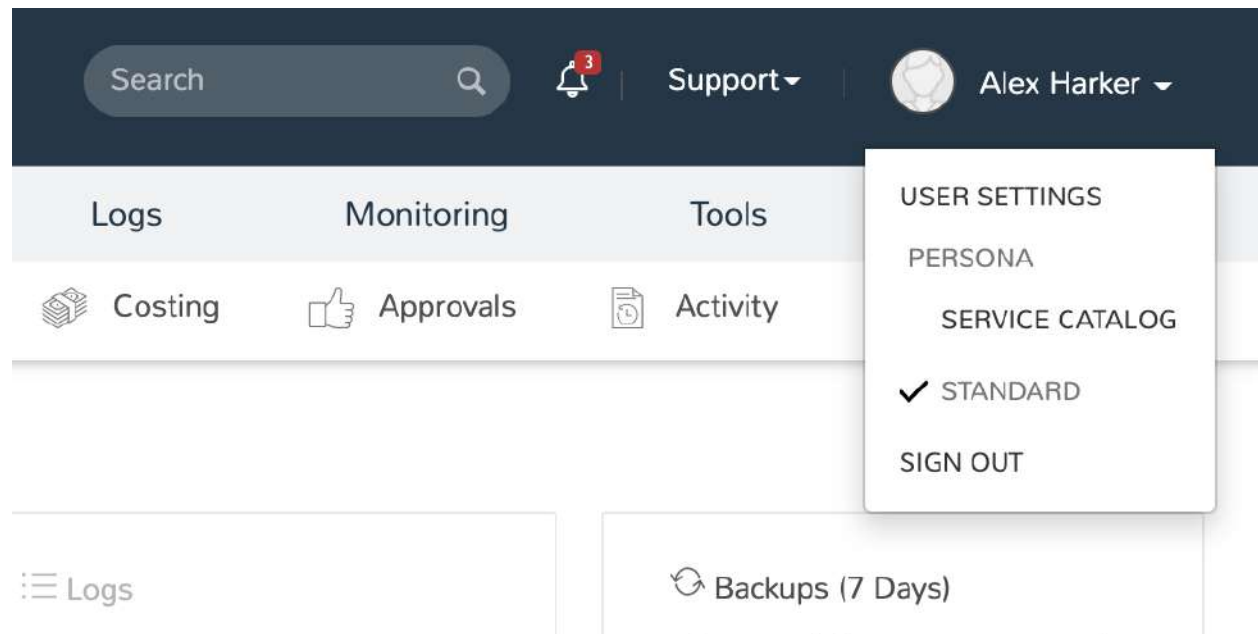
Configuring Global Access:

- **Full:** Gives access to all Catalog Items
- **Custom:** Gives access to individually-selected items from the list below
- **None:** No access is given to any Catalog Items

Tip: When giving Custom access, be sure to set access on some of the individual catalog items to Full in order to reveal those items to the Role group.

Accessing Alternate Personas

Switch Personas by clicking on your name in the upper-right corner of the application window. If your Role gives you access to any additional Personas, they will be listed here.

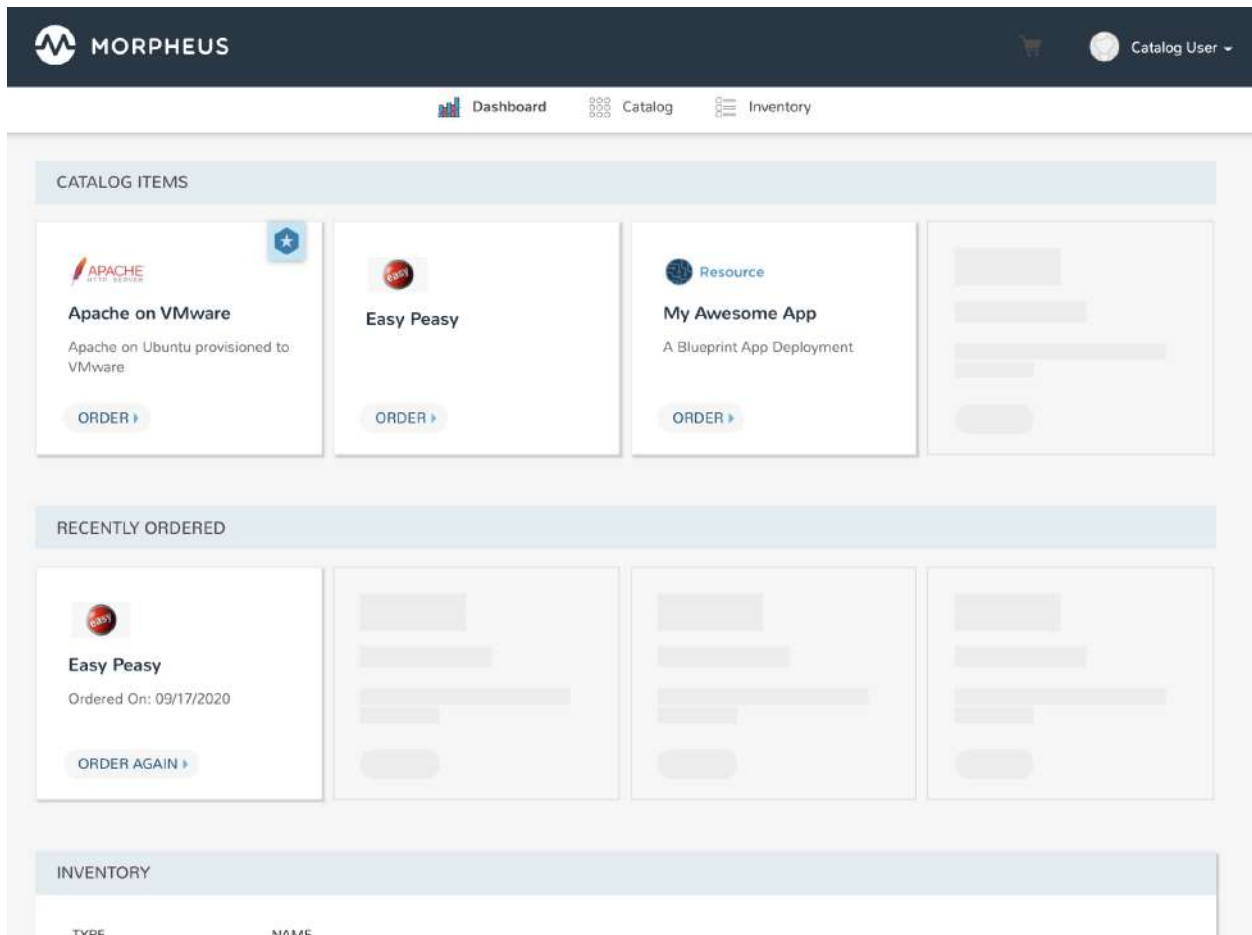


1.3.11 Service Catalog Persona

The Service Catalog Persona presents a simplified catalog where users can select and deploy Instances or Blueprints with pre-defined configuration with just a few clicks and without presenting an overwhelming list of options.

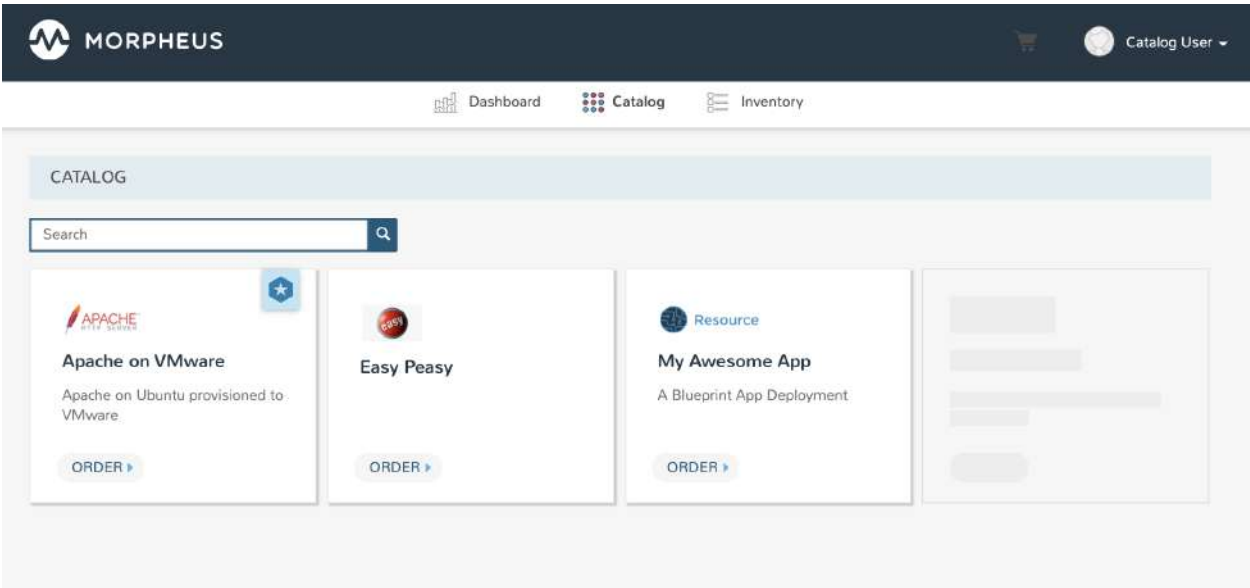
Dashboard

The default page for the Service Catalog Persona is the Dashboard. The Dashboard shows a selection of featured catalog items, a listing of the last few items the user has ordered, and a selection from the user's inventory.



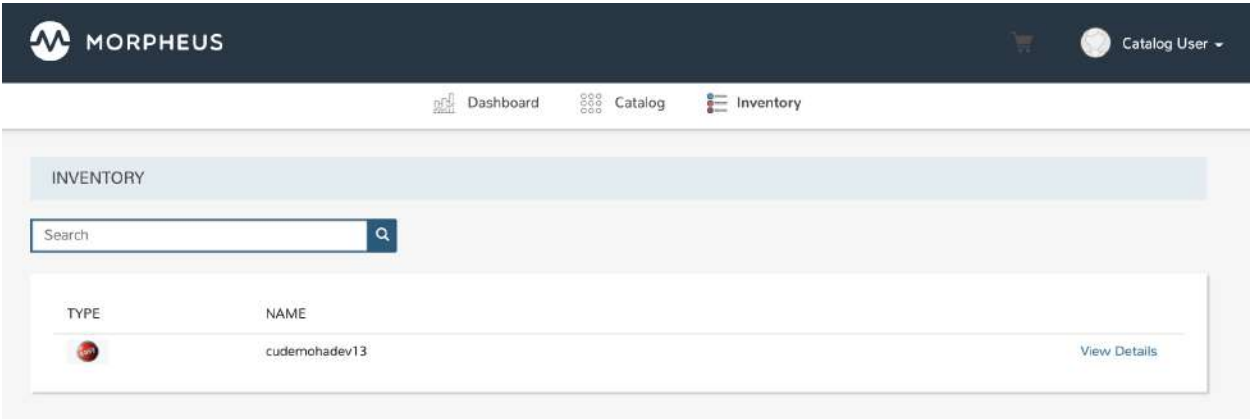
Catalog

The catalog shows the complete list of pre-defined catalog items available to the user for provisioning. Catalog items are not created in the Service Catalog Persona. For more on creating catalog items, see the [Self Service section](#) (Tools > Self Service) of Morpheus docs.



Inventory

The Inventory Page reveals the complete list of items which have been ordered by the user and provisioned. Users will only see their own items in this section.



Inventory Detail

Access the detail page for an item in your inventory by clicking the View Details link. The page displayed will look very similar to an Instance or App detail page in the Standard Persona. Highly detailed information on the health of the Instance or App are shown here. We can also take actions against Instances such as running Tasks or Workflows, reconfiguring the Instance, controlling the power state, and more.

The screenshot shows the Morpheus web interface. At the top is a dark navigation bar with the Morpheus logo and a user profile 'Catalog User'. Below this is a breadcrumb trail: 'Inventory > cudemohadev13'. The main content area has a header for 'cudemohadev13' with a star icon, environment 'Env: Dev', and plan 'Plan: 1 CPU, 512MB Memory'. To the right are 'ACTIONS' and 'DELETE' buttons. Below the header is a row of metrics: STATUS (warning icon), HEALTH (info icon), LAST BACKUP (minus icon), AVAILABILITY (100.00%), RESPONSE TIME (N/A), MAX CPU (0%), MEMORY (3%), and STORAGE (0%).

The 'INFO' section shows details for a CentOS instance:

- Group: Labs
- Owner: Catalog User
- Cores: 1
- Price: \$33.50 / Month
- Cloud: vCenter 180
- Layout: VMware VM
- Memory: 512.0MiB
- Source Image: Morpheus CentOS 7.5 v4
- Date Created: 09/17/2020 02:27 PM
- Version: 7.5
- Total Storage: 10.0GiB
- Provision Time: 4 minutes 28 seconds

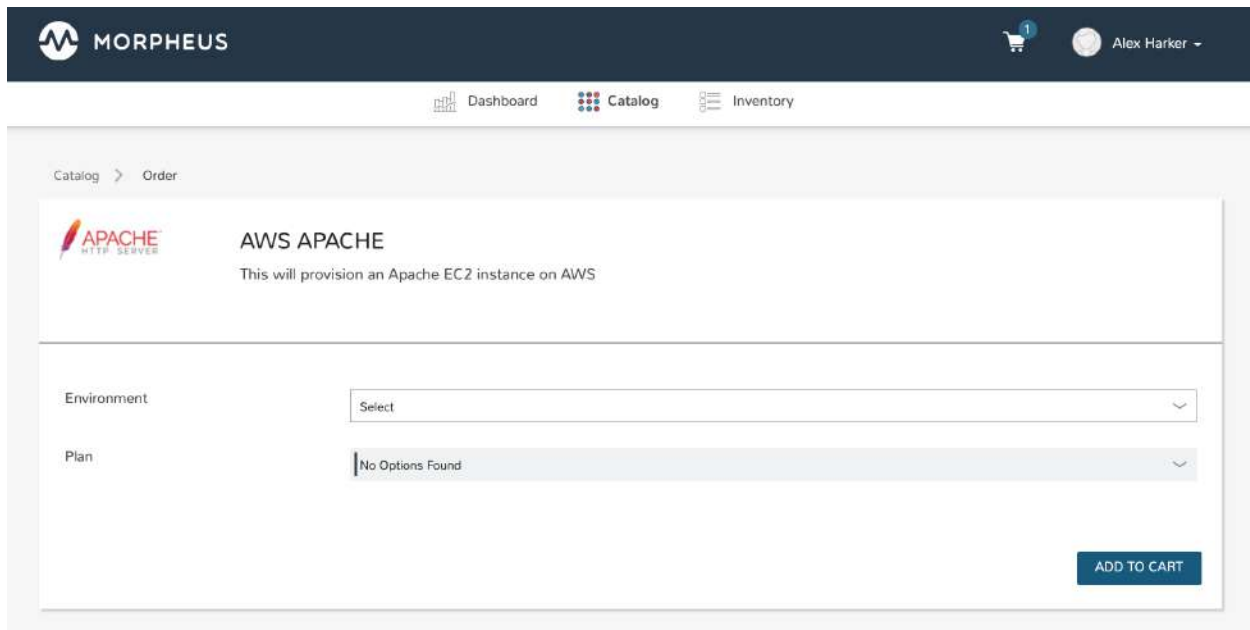
The 'VMS' section contains a table with one entry:

<input type="checkbox"/>	STATUS	NAME	TYPE	CLOUD	LOCATION	COMPUTE	MEMORY	STORAGE	ACTIONS
<input type="checkbox"/>		cudemohadev13	CentOS 7.5	vCenter 180	10.30.20.163:22	0	3	0	

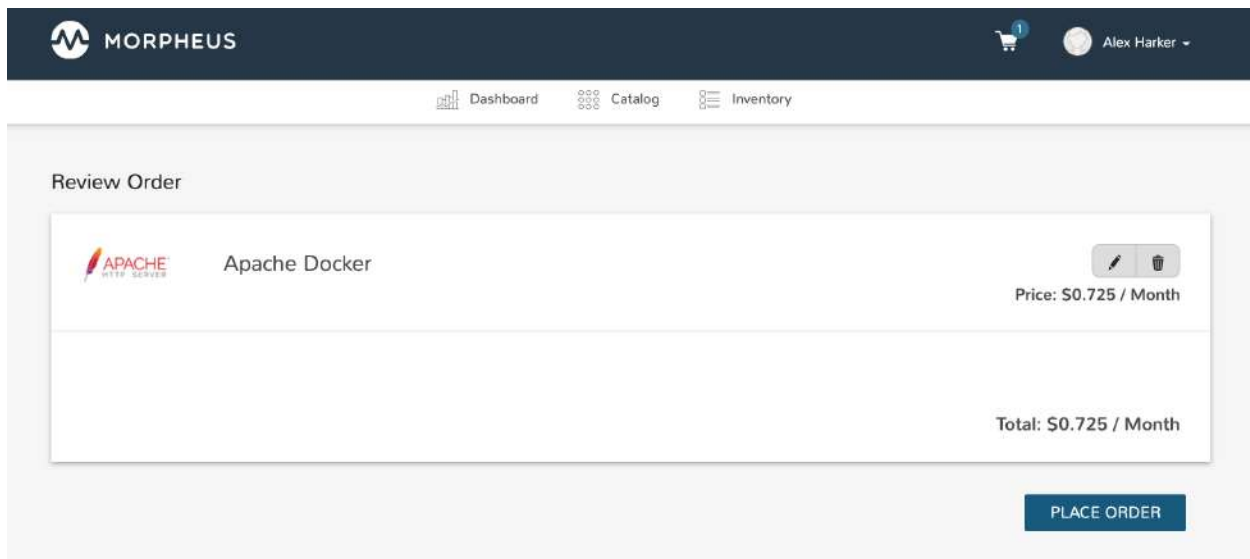
Below the table is a tabbed interface with 'Summary' selected. Other tabs include Wiki, Storage, Network, Logs, Backups, Environment, History, and Console. At the bottom, a memory usage bar shows '0.50 GiB' with a legend for USED MEMORY (black), CACHE MEMORY (red), and MAX MEMORY (grey).

Placing Orders

From the Catalog page, select the tile for your chosen item to see any custom options that may need to be set prior to provisioning.



Once the item is in the cart, make any additional selections to complete the order. Once finished, proceed to the cart by clicking on the cart icon at the top of the application window. Click “Review Order”. When reviewing your order, each selected item is listed along with its estimated cost. The total estimated cost for the entire order is also computed.



Once *PLACE ORDER* is clicked, the provisioning process will begin and the user is redirected back to the catalog page. Any new orders can be viewed in Inventory and additional details can be accessed through the Inventory item detail page.

1.3.12 Troubleshooting

Ansible Troubleshooting

When a workflow is executed manually, the Ansible run output is available in the Instance History tab. Select the **i** bubble next to the Ansible task to see the output. You can also see the run output in UI logs at `/var/log/morpheus/morpheus-ui/current`. These can be tailed by running `morpheus-ctl tail morpheus-ui`.

Verify Ansible is installed on the Morpheus Appliance

Ansible should be automatically installed but certain OS or network conditions can prevent the automated install. You can confirm installation by running `ansible --version` in the Morpheus appliance, or by viewing the Ansible integration details page (Administration > Integrations > Select Ansible Integration). We also see it in the Ansible tab of a Group or Cloud scoped to Ansible, just run `--version` as `ansible` is already included in the command.

If Ansible is not installed

Follow these instructions to install, or use your preferred installation method

Ubuntu:

```
sudo apt-get install software-properties-common
sudo apt-add-repository ppa:ansible/ansible
sudo apt-get update
sudo apt-get install ansible python-requests
```

CentOS:

```
sudo yum install epel-release
sudo yum install ansible python-requests
```

Then create the working Ansible directory for Morpheus:

```
sudo mkdir /opt/morpheus/.ansible
sudo chown morpheus-app.morpheus-app /opt/morpheus/.ansible
```

Validate Git repo authorization and the configured paths

The public and private SSH keys need to be added to the Morpheus appliance via Infrastructure > Keys & Certs and the public key needs to be added to the Git repo via user settings. If both are set up correctly, you will see the playbooks and roles populate in the Ansible Integration details page.

The Git Ref field on playbook tasks is to specify a different git branch than default. It can be left to use the default branch. If your playbooks are in a different branch you can add the branch name in the Git Ref field.

When running a playbook that is in a workflow, the additional playbooks fields do not need to be populated, they are for running a different playbook than the one set in the Ansible task in the Workflow, or using a different Git Ref.

Note: If you are manually running Workflows with Ansible tasks on existing Instances through Actions > Run Workflow and not seeing results, set the Provision Phase on the Ansible task to Provision as there may be issues with executing tasks on other phases when executing manually.

Attaching Logs to Case

When submitting a case it is critical to attach the relevant logs. The logs can be found at `/var/log/morpheus/morpheus-ui/current`. Logs can be attached to the case at anytime.

When submitting logs please reproduce the error right before capturing and sending the log file. This will ensure the activity that took place and resulted in an error is contained in the logs.

Log rotation takes the current file each night or after it's a certain size and compresses them. The `*.s` files in the current directory are rotated and zipped logs that can be sent as is.

The logs can also be captured from the Morpheus UI. Under Operations -> Health -> Morpheus Logs. Please copy relevant logs and add to case as an attachment.

Level	Timestamp	Message
ERROR	05/28/2019 01:53 PM	(1) [pool-2-thread-7] listScaleGroups error: com.amazonaws.services.autoscaling.model.AmazonAutoScalingClientException: User: arn:aws:iam::370459551696:user/morpheus is not authorized to perform the action: autoscaling:DescribeAutoScalingGroups
ERROR	05/28/2019 01:53 PM	(1) [pool-2-thread-7] listInstanceProfiles error: java.net.MalformedURLException: java.net.MalformedURLException: no protocol: https://api.github.com/user/repos
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-17] refresh zone error: java.net.MalformedURLException: java.net.MalformedURLException: no protocol: https://api.github.com/user/orgs
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-20] Error in adding image [properties:[storageProfile:[osDisk:[osType:Windows,
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-20] Error in adding image [properties:[storageProfile:[osDisk:[osType:Linux, c
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-20] Error in adding image [properties:[storageProfile:[osDisk:[osType:Windows,
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-20] Error in adding image [properties:[storageProfile:[osDisk:[osType:Windows,
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-20] Error in adding image [properties:[storageProfile:[osDisk:[osType:Windows,
ERROR	05/28/2019 01:51 PM	(1) [pool-2-thread-20] Error in adding image [properties:[storageProfile:[osDisk:[osType:Windows,
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-5] listSubscriptions error: java.lang.RuntimeException: Not authorized java.la
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-5] path: /3c29d137-45a9-43e1-9ba2-ee3bd81c3221/oauth2/token error: 401 - {"err
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-18] remoteGitFetch error: org.eclipse.jgit.api.errors.TransportException: gitf
INFO	05/28/2019 01:50 PM	(1) [pool-2-thread-11] online: bitbucket.org bitbucket.org true
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-18] remoteGitFetch error: org.eclipse.jgit.api.errors.TransportException: gitf
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-10] error: 401 - https://api.github.com/user/repos [message:Requires authenti
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-10] error: 401 - https://api.github.com/user/orgs [message:Requires authentic
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-4] listDnsHostedZones error: com.amazonaws.services.route53.model.AmazonRoute53ClientException: User: arn:aws:iam::370459551696:user/morpheus is not authorized to perform the action: route53:ListHostedZones
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-6] Error Occurred calling web API: http://10.30.20.175:8080 overallLoad/api/js
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-6] Error Occurred calling web API: jenkins.bertramlabs.com overallLoad/api/jsc
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-6] path: /v2.0/lbaas/healthmonitors error: 404 - 404 Not Found The resource cc
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-18] remoteGitFetch error: org.eclipse.jgit.api.errors.TransportException: gitf
ERROR	05/28/2019 01:50 PM	(1) [pool-2-thread-5] path: /v2.0/lbaas/healthmonitors error: 404 - 404 Not Found The resource cc

Blank Dashboard

Problem A blank dashboard or 500 error after installing morpheus

Note: A blank or 500 error on just the dashboard is different than the entire morpheus-ui not loading. Please see UI note loading article for troubleshooting the ui not loading after an upgrade.

Cause Elasticsearch restarting prior to being fully bootstrapped during the initial install.

Solution To fix, purge elasticsearch by running the following on the Morpheus Appliance:

```
curl -XDELETE http://localhost:9200/*
morpheus-ctl restart elasticsearch
morpheus-ctl restart morpheus-ui
```

Another option is:

```
sudo rm -rf /var/opt/|morpheus| /elasticsearch/data/morpheus
morpheus-ctl restart elasticsearch
morpheus-ctl restart morpheus-ui
```

If you get a term/timeout on ui restart, run

```
morpheus-ctl kill morpheus-ui
morpheus-ctl start morpheus-ui
```

Note: The morpheus-ui may take a few minutes to load and be available after being restarted

CLI Troubleshooting

If you have installed the Morpheus CLI successfully and get a successful login but see this error Error Communicating with the Appliance. SSL_connect returned=1 errno=0 state=error: certificate verify failed

run the command

```
morpheus remote update {appliancename} --insecure
```

Cannot Login

Forgot password

If a user forgets their password, they can use the *FORGOT PASSWORD?* link on the login page. They can then enter their username or email address to send a reset password email to the email address defined on the user.

If the default or user added SMTP server is not functioning or blocked, a System Admin user can impersonate that user and update their password.

If the System Admin user password needs to be reset and the default or user added SMTP server is not functioning or blocked, please contact Morpheus support for assistance.

Sub-Tenant user cannot login after 3.4.0 upgrade

Morpheus v3.4.0 added support for all subtenant users to login via the main tenant url using subtenant id or subdomain prefix, ie `tenantId\username` or `subdomain\username`.

Note: Tenant subdomains can be defined by editing Tenant settings and updating the *SUBDOMAIN* field.

Important: Subtenant local users will no longer be able to login from main login url without using their subtenant id or subdomain prefix.

The login requirements were added in v3.4.0 to allow subtenant users with identity source integration generated user accounts to be able to login to the master tenant, gain API and CLI access, and remove the requirement for usernames to be unique across all tenants.

Previously subtenant users that had local/morpheus generated user accounts could login to their tenant via the master tenant url, while subtenant users that had identity source integration generated user accounts had to use the subtenant specific login url.

In v3.4.0+ all subtenant users can login via the master tenant url by specifying their tenant id or subdomain prefix, \, then username. Subtenants can still use the tenant specific login url as well.

Example: I have a username `subuser` that belongs to a tenant with the subdomain `acme` and tenant id `58`. When logging in from the main login url, I now need to enter in: `acme\subuser` and the password. Alternatively the tenant ID can be used, ie `58\subuser`

Active Directory user suddenly cannot Login

In Morpheus v3.4.0 and prior, OU changes in Active Directory can disable logins for AD users who had previously authenticated/have existing user accounts in Morpheus. If an Active Directory user cannot login to Morpheus after their OU was changed in AD, please contact Morpheus support for a resolution. The OU association for the user(s) can also be manually updated in the database. This issue is resolved in Morpheus versions 3.4.1 and higher.

Common Ports & Requirements

The following chart is useful for troubleshooting Agent install, Static IP assignment, Remote Console connectivity, and Image transfers.

Table 12: Common Ports & Requirements

Feature	Method	OS	Source	Destination	Port	Requirement
Agent Communication	All	All	Node	Appliance	443	DNS Resolution from node to appliance url
Agent Install	All	Linux	Node	Appliance	80	Used for appliance yum and apt repos
	SSH	Linux	Appliance	Node	22	DNS Resolution from node to appliance url Virtual Images configured SSH Enabled on Virtual Image
	WinRM	Windows	Appliance	Node	5985	DNS Resolution from node to appliance url Virtual Images configured WinRM Enabled on Virtual Image(<i>winrm quickconfig</i>)
	Cloud-init	Linux				Cloud-init installed on template/image Cloud-init settings populated in User Settings or in <i>Admin -> Provisioning</i> Agent install mode set to Cloud-Init in Cloud Settings
	Cloudbase-init	Windows				Cloudbase-init installed on template/image Cloud-init settings populated in User Settings or in <i>Admin -> Provisioning</i> Agent install mode set to Cloud-Init in Cloud Settings
	VMtools	All				VMtools installed on template Cloud-init settings populated in Morpheus user settings or in <i>Administration -> Provisioning</i> when using Static IP's Existing User credentials entered on Virtual Image when using DHCP RPC mode set to VMtools in VMware cloud settings.
Static IP Assignment & IP Pools	Cloud-Init	All				Network configured in Morpheus (Gateway, Primary and Secondary DNS, CIDR populated, DHCP disabled) Cloud-init/Cloudbase-init installed on template/image Cloud-init settings populated in Morpheus user settings or in <i>Administration -> Provisioning</i>
422	VMware Tools	All				Chapter 1. v5.2.0 Highlights Network configured in Morpheus (Gateway, Primary and Secondary DNS, CIDR populated, DHCP disabled)

Deleting Instances

It is important to know the difference between deleting an Instance from the Provisioning section, and deleting a VM from the Infrastructure section.

Important: Deleting an Instance a with Virtual Machines in it will always try to delete the actual Virtual Machines.

Instances are managed resources that may have one or multiple Virtual Machines associated. Since the vm's in the Instance are managed by Morpheus, deleting an Instance a with Virtual Machines in it will always try to delete the actual Virtual Machines.

There are scenarios where deleting, or attempting to delete the associated Virtual Machines is not desired:

- The Instance needs to be deleted, but the actual Virtual Machines need to remain.
- The actual Virtual Machines have already been deleted outside of Morpheus, so only the records in Morpheus need to be removed.

Deleting an Instance without deleting Infrastructure

It is not possible to delete an Instance from the Provisioning section without removing the associated Infrastructure/VM's. However this can be accomplished from the Infrastructure section by deselecting "Remove Infrastructure" when deleting the VM:

1. Navigate to the Virtual Machine record by clicking on the VM's name in the Virtual Machines section in the Instances details section, or by navigating to *Infrastructure - Hosts - Virtual Machines* and selecting the VM.
2. Click "DELETE"
3. In the delete confirmation modal:
 - Uncheck "Remove Infrastructure"
 - Check "Remove Associated Instances"

DELETE HOST?

×

Warning! Deleting this Host with "Remove Infrastructure" enabled will permanently delete it from the Cloud. To only delete the record but leave in the Cloud, uncheck "Remove from Infrastructure". If "Inventory Existing Instances" is enabled on the Cloud, the host will be re-synced as discovered.

☐ Remove Infrastructure

DO NOT SELECT

☐ Preserve Volumes

☒ Remove Associated Instances

☐ Force Delete

CANCEL

DELETE

Important: Ensure "Remove Infrastructure" is NOT checked if you do not want to delete the actual Virtual Machine.

4. Select DELETE

This will delete the Virtual Machine record as well as the Instance record, but leave the Infrastructure/VM in place. If the VM is in a Cloud that is being inventoried, it will s

Deleting an Instance/VM that does not exist anymore

Deleting a managed resource outside of Morpheus is not recommended as it will leave stranded record in Morpheus and cause deleting the records in Morpheus to get stuck on delete when Morpheus tries to remove infrastructure that is no longer there.

To select an Instance and/or VM record in Morpheus for a Virtual Machine that no longer exists:

1. Navigate to the Virtual Machine record by clicking on the VM's name in the Virtual Machines section in the Instances details section, or by navigating to *Infrastructure - Hosts - Virtual Machines* and selecting the VM.
2. Click "DELETE"
3. In the delete confirmation modal:
 - Uncheck "Remove Infrastructure"
 - Check "Remove Associated Instances"

DELETE HOST?
✕

Warning! Deleting this Host with "Remove Infrastructure" enabled will permanently delete it from the Cloud. To only delete the record but leave in the Cloud, uncheck "Remove from Infrastructure". If "Inventory Existing Instances" is enabled on the Cloud, the host will be re-synced as discovered.

☐ Remove Infrastructure DO NOT SELECT

☐ Preserve Volumes

☒ Remove Associated Instances

☐ Force Delete

CANCEL
DELETE

Important: Ensure “Remove Infrastructure” is NOT checked. If it is checked, Morpheus will try to delete the actual VM, and since it is not there anymore, the delete will not complete successfully since Morpheus will not be able to verify successful deletion of the Infrastructure.

4. Select DELETE

The key point is when deleting an Instance, or when selecting “Remove Infrastructure” when deleting a VM record, Morpheus will always try to remove the Infrastructure. If the Infrastructure/VM no longer exists, or you do not want to remove it, simply delete from the Infrastructure section and uncheck “Remove Infrastructure”.

Note: When deleting a managed VM, if that VM is the only VM inside the associated Instance, the Associated Instance must also be removed.

How to un-manage an Instance/VM/Host

Description

A managed VM (and associated Instance) needs to be unmanaged and returned to Discovered type.

Solution

Delete the record from the `Infrastructure - Hosts` (! not from `Provisioning - Instances`) selection with the following configuration in the Delete modal:

- Remove `Infrastructure` **UNCHECKED**
- Remove `Associated Instances` **Must** be checked if the server has an associated Instance, as deleting the VM but not the Instance would result in an abandoned Instance thus not allowed.
- Force Delete **UNCHECKED**

The most important items to be aware of when “un-managing” an Instance/VM/Host are:

1. The “Remove from Infrastructure” flag when deleting a VM or Host in Morpheus determines if the actual VM is deleted from the target Infrastructure.
 - Checking “Remove Infrastructure” means you **WANT TO DELETE THE ACTUAL VM**. Typing “DELETE” in the confirmation field is required when “Remove From Infrastructure” is enabled.
 - Unchecking “Remove Infrastructure” means you only want to delete the record in Morpheus but leave the actual VM untouched.
2. Deleting an Instance will always remove Infrastructure.

Important: REPEAT: Deleting an Instance from the `Provisioning` section will always remove the VM aka Infrastructure.

3. After removing the record from Morpheus, the VM must be in a Cloud with Inventory enabled to automatically be re-discovered.

Process

Steps to delete a managed VM from Morpheus and, when necessary, remove the associated Instance:

1. Navigate to the VM (not Instance) detail page at `Infrastructure - Hosts - VMs`

Note: VM's inside an Instance can be navigated to inside the Instance Details page by selecting the VM in the VM's section on the Instance Details page.

2. Select *DELETE*
3. Configure the DELETE HOST modal with the following settings:

DELETE HOST?

×

Warning! Deleting this Host with "Remove Infrastructure" enabled will permanently delete it from the Cloud. To only delete the record but leave in the Cloud, uncheck "Remove from Infrastructure". If "Inventory Existing Instances" is enabled on the Cloud, the host will be re-synced as discovered.

☐ Remove Infrastructure
 ☒ Remove Associated Instances
 ☐ Force Delete

CANCEL

DELETE

- Remove Infrastructure **UNCHECKED**
- Remove Associated Instances Must be checked if the server has an associated Instance, as deleting the VM but not the Instance would result in an abandoned Instance thus not allowed.
- Force Delete **UNCHECKED**

Important: If you have to type DELETE that means the Remove Infrastructure flag is selected and you are confirming deletion of the actual VM. Ensure Remove Infrastructure is **UNCHECKED** when you want to leave the VM intact!

4. Select *DELETE*
5. The VM and associated Instance will be removed from Morpheus but the actual VM will remain.
6. Wait up to 5 min or click *REFRESH* on the associated Clouds details page to force a cloud sync.

Note: Inventory must be enabled on the associated cloud for the VM to automatically be re-discovered by Morpheus.

7. The VM is now back in Morpheus as discovered/unmanaged. To manage and create a new Instance from the VM, select *ACTIONS* : Convert To Managed.

MySQL Too many connections error

If you see the following error in the Morpheus UI logs:

```
SQLExceptionHelper - Data source rejected establishment of connection, message from ↵  
↵server: "Too many connections"
```

it means the number connections between Morpheus application and mysql have reached the *max_connections* limit set in mysql (default is 151), or the *max_active* setting, which limits the number of connections on the Morpheus end (default is 100), and the limit needs to be raised, either in Morpheus or mysql, or both depending on the number of connections and configuration.

Note: The *max_connections* setting in mysql and the maximum used connections between an app node and mysql can be viewed in the Morpheus ui in the *Operations - Health* section under Database.

Important: In Single Morpheus app node configurations, the *max_active* setting on the app node must be less than the *max_connections* setting in mysql.

Important: In HA configurations, the *max_active* setting is per app node, and the *max_connections* setting in mysql must be greater than all app nodes *max_active* values combined, ie $(\text{max_active} * \text{number_of_app_nodes}) \leq \text{max_connections}$.

Morpheus max_active setting

Edit `/etc/morpheus/morpheus.rb` and add `mysql['max_active'] = $value` replacing *\$value* with desired number of maximum connections allowed by Morpheus to mysql. For example, to set *max_active* at 150:

```
mysql['max_active'] = 150
```

Replacing 100 with the desired number of maximum connections allowed by Morpheus to mysql.

Run `morpheus-ctl reconfigure` for the setting to be applied. Reconfigure will not restart the ui unless additional ram has been added to the appliance host since the previous reconfigure. To edit the *max_active* without a reconfigure, update the *max_active* setting in `/opt/morpheus.conf.application.yml`. Please note the default setting of 100 will be applied upon the next reconfigure unless *max_active* is defined as instructed above in the `morpheus.rb` file.

mysql max_connections setting

Important: Customers are responsible for configuring and maintaining external databases used by Morpheus. This explains how to set the *max_connections* setting, but the value for the setting needs to be established by a customers qualified db admin.

In mysql prompt,

```
run mysql> SET GLOBAL max_connections = $value;
```

This will immediately write the variable, however it is only a temporary setting that will be overwritten upon restart of the mysql service.

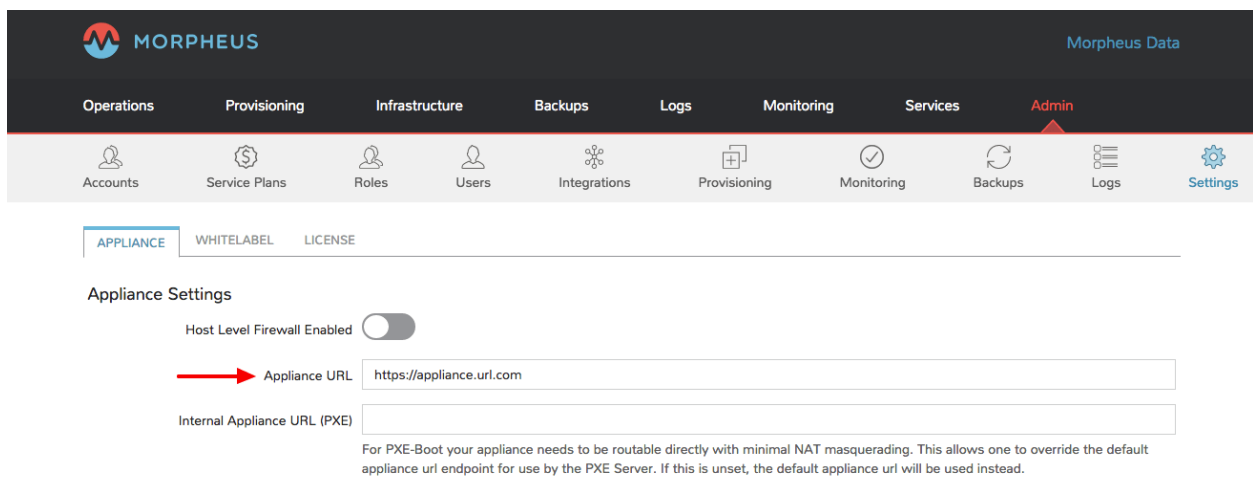
To persist the `max_connections` setting, edit `my.cnf`, and add `max_connections = $value` replacing `$value` with desired value, ie to set `max_connections` at 300 in `my.cnf`, add:

```
max_connections = 300
```

Morpheus Agent Install Troubleshooting

When provisioning an Instance, there are network and configuration requirements to consider in order to successfully install the Morpheus Agent. Typically, when a VM Instance is still in the provisioning phase long after the VM is up, the Instance is unable to reach Morpheus. Depending on the Agent install mode, it could also mean Morpheus is unable to reach the Instance.

The most common reason an Agent install fails is the provisioned Instance cannot reach the Morpheus Appliance via the Appliance URL set in Administration > Settings over port 443. When an Instance is provisioned from Morpheus, it must be able to reach the Morpheus appliance via the Appliance URL or the Agent will not be installed.



In addition to the main Appliance URL in Administration > Settings, additional Appliance URLs can be set per Cloud in the Advanced Options section of the Cloud configuration modal when creating or editing a Cloud. When this field is populated, it will override the main Appliance URL for anything provisioned into that Cloud.

Tip: The Morpheus UI current log, located at `/var/log/morpheus/morpheus-ui/current`, is very helpful when troubleshooting Agent installations.

Agent Install Methods

Morpheus Agent installation supports multiple install methods.

- SSH/WinRM
- VM Tools
- Cloud-Init & Cloudbase-Init
- Windows Unattended
- Manual

For All Agent Install Methods

When an Instance is provisioned and the Agent does not install, verify the following for any Agent install mode:

- The Morpheus Appliance URL (Administration > Settings) is both reachable and resolvable from the provisioned node
- The Appliance URL begins with <https://>, not <http://>

Note: Be sure to use <https://> even when using an IP address for the appliance.

- Inbound connectivity access to the Morpheus appliance from provisioned VMs and container hosts on port 443 (needed for Agent communication)
- Private (non-Morpheus provided) VM images and templates must have their credentials stored. These can be entered or edited in the Provisioning > Virtual Images section by clicking the Actions dropdown on an imaged detail page and selecting Edit.

Note: Administrator user is required for Windows Agent install.

- The Instance does not have an IP address assigned. For scenarios without a DHCP server, static IP information must be entered by selecting the Network Type: Static in the Advanced Options section during provisioning. IP Pools can also be created in the Infrastructure > Networks > IP Pools section and added to Cloud network sections for IPAM
- DNS is not configured and the node cannot resolve the appliance. If DNS cannot be configured, the IP address of the Morpheus appliance can be used as the main or Cloud appliance

SSH

- Port 22 is open for Linux images, and SSH is enabled
- Credentials set on the image if using a custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section

WinRM

- Port 5985 must be open and WinRM enabled for Windows images
- Credentials have been entered on the image if using a custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section

Note: Administrator user is required for Windows Agent install.

VMware Tools (vmtools)

- VMware Tools is installed on the template(s)
- Credentials have been entered on the image if using custom or synced image. Credentials can be entered on images in the Provisioning > Virtual Images section
- Sudo privileges required for Linux
- Administrator User required for Windows (SID 500)

Cloud-Init

- Cloud-Init settings configured in Administration > Provisioning section
- Cloud-Init installed on Virtual Image
- Cloud-Init enabled on Virtual Image config

Cloudbase-Init

- Windows Administrator Password defined in Administration > Provisioning section
- Cloudbase-Init installed on Virtual Image
- Cloud-Init enabled on Virtual Image config

Note: Unattend Agent Installation and customizations are recommended over Cloudbase-Init

Windows Unattended

- Windows Administrator Password defined in Administration > Provisioning section
- VMware: Force Guest Customizations set to forced on Virtual Image config when using DHCP (Static Assignment will already force Guest Customizations)
- Nutanix & SCVMM: Virtual Image is sysprepped and shutdown, Sysprep Enabled flagged on Virtual Image config

Manual

- Agent Install scripts can be downloaded from Morpheus and ran manually on the target host when required via Actions -> Download Agent Script on the managed Resource. Please note the script will be unique per managed Resource.

Restarting the Morpheus Agent

In some situations, it may necessary to restart the Morpheus Agent on the host to re-sync communication from the Agent to the Morpheus appliance.

Linux

On the target host, run `sudo morpheus-node-ctl restart morphd` and the Morpheus agent will restart. `morpheus-node-ctl status` will also show the agent status.

Windows

The Morpheus Windows Agent service can be restarted in Administrative Tools -> Services.

Tip: The Morpheus Remote Console is not dependent on Agent communication and can be used to install or restart the Morpheus agent on an Instance.

Uninstall Morpheus Agent

Linux Agents

You can use the following to uninstall the linux agent (contains commands for both rpm and deb agents)

```
sudo rm /etc/apt/sources.list.d/morpheus.list \  
sudo morpheus-node-ctl kill \  
sudo apt-get -y purge morpheus-node \  
sudo apt-get -y purge morpheus-vm-node \  
sudo yum -y remove morpheus-node \  
sudo yum -y remove morpheus-vm-node \  
sudo yum clean all \  
sudo systemctl stop morpheus-node-runsvdir \  
sudo rm -f /etc/systemd/system/morpheus-node-runsvdir.service \  
sudo systemctl daemon-reload \  
sudo rm -rf /var/run/morpheus-node \  
sudo rm -rf /opt/morpheus-node \  
sudo rm -rf /etc/morpheus \  
sudo rm -rf /var/log/morpheus-node \  
sudo pkill runsv \  
sudo pkill runsvdir \  
sudo pkill morphd \  
sudo usermod -l morpheus-old morpheus-node \  

```


Windows Agents

```
$app = Get-WmiObject -Class Win32_Product
                -Filter "Name = 'Morpheus Windows Agent'"
$app.Uninstall()
```

CentOS/RHEL 7 Images

For custom CentOS 7 images we highly recommend setting up Cloud-Init and fixing the network device names. More information for custom CentOS images can be found in the CentOS 7 image guide.

Morpheus UI not loading after upgrade or reconfigure

Problem: The Morpheus ui does not load after performing an upgrade.

Common Causes:

1. The morpheus-ui has not finished loading
2. The morpheus-ui was not fully stopped before reconfigure, or not started after reconfigure
3. Morpheus was forced to restart or shut down while the database schema was being migrated during an upgrade

Solutions:

1. The morpheus-ui has not finished loading.

An easy way to see when the ui is finished loading and running is to tail the ui current file and look for the morpheus logo with version and start time

```
morpheus-ctl tail morpheus-ui
```

Note: After running *morpheus-ctl start morpheus-ui*, the Morpheus ui takes around 3 minutes to run depending on hardware.

1. The morpheus-ui was not fully stopped before reconfigure, or not started after reconfigure

The morpheus ui must be stopped prior to running morpheus-ctl reconfigure when upgrading. Sometimes running morpheus-ctl stop morpheus-ui will timeout and the ui is not actually stopped. If stopping the ui does timeout, run morpheus-ctl kill morpheus-ui prior to reconfigure, and be sure to run morpheus-ctl start morpheus-ui after reconfigure is completed.

If you ran a reconfigure before stopping the ui, run:

```
sudo morpheus-ctl kill morpheus-ui
sudo morpheus-ctl reconfigure
sudo morpheus-ctl start morpheus-ui
```

Wait for the ui to come up.

2. Morpheus was forced to restart or shut down while the database schema was being migrated during an upgrade

If the ui fails to start and you see the error `Invocation of init method failed; nested exception is liquibase.exception.LockException: Could not acquire change log lock. Currently locked by morpheus` it likely means morpheus was forced to restart or shut down while the database schema was being migrated during an upgrade, and the lock was not released.

To release the lock, you will need to run a mysql query. You will need to install mysql-client on the morpheus appliance, and grab the password for morpheus mysql. The username and db name are both morpheus. The password to login to mysql can be found in the `application.yml` file located at `/opt/morpheus/conf/application.yml`

Then run the following:

```
mysql -u morpheus -p -h 127.0.0.1 morpheus
```

At the prompt, enter the mysql password from the `application.yml`

Then run:

```
DELETE FROM DATABASECHANGELOGLOCK;
```

Then restart morpheus-ui:

```
sudo morpheus-ctl restart morpheus-ui
```

If the restart timesout, run:

```
sudo morpheus-ctl kill morpheus-ui
sudo morpheus-ctl start morpheus-ui
```

Remote Console

Morpheus has a built in Remote Console for Instances, Hosts, Virtual Machines and Bare Metal. The following information reviews the Roles Settings, Protocols, and Requirements necessary to configure and troubleshoot Remote Console access.

Role Settings

User Role settings determine if the Console tab or `Open Console Action` appear for a user, and if a login prompt is presented or the user is automatically logged in when using the Console.

- **Remote Console (None, Provisioned, Full)**

None The user will not have access to remote console.

Provisioned The user will only have remote console access for Instances they provisioned.

Full The user will have remote console access for all instances they have access to.

- **Remote Console: Auto Login (No, Yes)**

No A login prompt will be present in the console for Linux platforms, and the main login screen will present for Windows platforms.

Yes Morpheus will automatically login to the remote console using the credentials defined on the VM or Host. For provisioned Instances, the credentials are defined either from the credentials defined on the Virtual Image used, added via cloud-init or VMware Tools using the global cloud-init settings (Administration - Provisioning) or the Linux or Windows settings defined in User Settings. For

Instances created when converting a VM or Host to managed, the credentials are entered when converting to managed. These credentials can be changed by editing the underlying VM or Host of the Instance.

Note: If the credentials defined on the VM or Host are not valid, and the `Remote Console: Auto Login Role` setting is set to `Yes`, the console will not be able to connect and no console window or login prompt will be presented. The credentials on the underlying VM or Host must be edited or `Remote Console: Auto Login Role` setting can be set to `No` for a login prompt to present in the console. Credentials cannot be changed from an Instance view, only in the Infrastructure VM or Host view.

Protocols

Platform Type and Cloud Settings determines the protocol and port used for Remote Console connections.

- **SSH** The SSH protocol will be used for Linux and OSX platform types, and 22 is the default port used.
- **RDP** The RDP (Remote Desktop) protocol will be used for Windows platform types over port 3389 by default.
- **VNC** The VNC protocol will be used for all platform types in Clouds with the `Hypervisor Console` option enabled in cloud settings. VNC connection are made directly to the Hypervisor Host over port 443.

Note: Alternative ports can be configured per VM or Host by editing the VM or Host and editing the Port field in the RPC host section.

SSH

For all Linux and OSX platform types, Morpheus will use the SSH protocol via port 22 by default for Remote Console connections, unless the `Hypervisor Console`` option is enabled for VMware type clouds.

Morpheus will SSH using the username, password, RPC Host IP address and Port defined in the VM or Host record.

Default Requirements for SSH Connectivity

- SSH Enabled on the target VM or Host
- Port 22 incoming open on the target VM or Host firewalls and security groups from the Morpheus Appliance (not from the users IP address)
- An IP address defined on the VM or Host record that is routable from the Morpheus Appliance.
- Valid credentials defined on the VM or Host record in the RPC host field.
- *Remote Console* Role Permissions set to *Provisioned* or *Full* if the User provisioned the instance, or *Full* if the user did not provision the instance.

RDP

For all Windows platform types, Morpheus will use the RDP protocol via port 3389 by default for Remote Console connections, unless the *Hypervisor Console`* option is enabled for VMware type clouds.

Morpheus will RDP using the username, password, RPC Host IP address and Port defined in the VM or Host record.

Default Requirements for RDP Connectivity

- Remote Access enabled on the target VM or Host and Remote Desktop enabled in the Windows Firewall settings. If the VM or Host is on a different network than the Morpheus appliance, public access for Remote Desktop must be enabled in the Firewall settings.
- Port 3389 incoming open on the target VM or Host firewalls and security groups from the Morpheus Appliance (not from the users IP address)
- An IP address defined on the VM or Host record that is routable from the Morpheus Appliance.
- Valid credentials defined on the VM or Host record in the RPC host field.
- *Remote Console* Role Permissions set to *Provisioned* or *Full* if the User provisioned the instance, or *Full* if the user did not provision the instance.

Note: If *Remote Console: Auto Login* is set to *No* in a users Role permissions, *Allow connections only from computers running Remote Desktop with Network Level Authentication* in the *Windows System Properties -> Remote* settings must be **DISABLED** for Remote Console to connect.

VNC (VMware Hypervisor Console)

When the `Hypervisor Console` option is enabled in cloud settings, the VNC protocol will be used for all platform types that Cloud.

When using VNC Hypervisor Console, the Morpheus Appliance connects directly to the host the VM is on, not directly to the VM.

Morpheus features Remote Console support directly to hypervisors. To enable this feature a few prerequisites must be met:

- The Morpheus Appliance must have network access to the host the VM is on over 443.
- The Morpheus Appliance must be able to resolve the hypervisor hostnames.

Note: VNC connections for VMs and Hosts in VMware type clouds are made directly to the ESXi hosts, not vCenter.

Unlike SSH and RDP, valid credentials do not need to be set on the VM or Host records in Morpheus for VNC hypervisor console connections. An IP address is also not required on the VM or Host for VNC hypervisor console connections. Morpheus will be able to connect to the VM or Host as soon as the `Host (Hypervisor)` record is set, which can be viewed in the Info section on the VM or Host detail page.

Note:

- Auto-login is not supported for Hypervisor Console. Auto-login role settings do not apply to console connecting when using Hypervisor Console. Please note Hypervisor Console sessions persist on the ESXi host and once a user manually logs in to the VM they will continue to be logged in, even if the console tab/window in Morpheus is closed, until they manually log out.

- Copy and Paste and Text selection in Linux terminals is not supported when using VNC (VMware Hypervisor Console).
 - In Morpheus versions 3.2.0 and higher, a newer Guacamole version is installed that is not compatible with MacOS Platform Types over VNC.
-

Copy and Paste

Note: Copy and Paste for Text is supported for SSH and RDP protocols only.

To Copy text from the console:

1. Select text in the Console window.
2. Click the COPY button at the top of the Console window.
3. The selected text is copied to the users clipboard.

To Paste text into console:

1. Copy text on the local computer to you clipboard
2. Right click into the “Paste Text Here” field at the top of the Console window. The field will the display “Text Copied, Use Console to Paste.”
3. Right click into the console window.
4. The text is pasted into the VM.

Guacamole

Overview

Morpheus uses Apache Guacamole, a clientless remote console. Guacamole is installed on the Morpheus Appliance during the initial reconfigure. In Morpheus versions 3.2.0 and higher, Guacamole 0.9.14 is automatically installed. On Morpheus versions older than 3.2.0, 0.9.9 is installed. The 0.9.14 version is required for VNC Hypervisor Console functionality on ESXi v6.5 and later.

The Guacamole proxy daemon, `guacd`, is used for all Remote Console connections and must be running for Remote Console functionality.

Troubleshooting guacd

If all console connections are not functioning, the Guacamole proxy daemon (`guacd`) process may not be running or have a stuck process preventing console connections. This is evident when only the header appears in the console tab/window, and no console window appears below the header and no connection status is show in the console header. The following commands can be used on the Morpheus Appliance to restore console functionality.

morpheus-ctl status Lists all local Morpheus services including `guacd` and their states. If `guacd` is stopped, it will need to be started again for Remote Console to function.

morpheus-ctl start guacd Starts the `guacd` process

morpheus-ctl stop guacd Stops the `guacd` process

morpheus-ctl kill guacd Forcefully kills the guacd process

morpheus-ctl restarts guacd Restarts the guacd process

morpheus-ctl tail guacd Tails the guacd current and state logs, located by default at `/var/log/morpheus/guacd/`. This log is useful when troubleshooting console connections, guacamole service status, and to determine the protocol being used for the Remote Console connection.

If guacd continues to stop even after being started, or if guacd is running and no properly configured console connections are functioning, there may be a stuck guacd or multiple guacd processes running, which will need to be killed and guacd started again.

To kill all guacd processes on the Morpheus Appliance and start guacd again:

1. Kill the morpheus guacd process: `morpheus-ctl kill guacd`
2. Grep for all running guacd processes: `sudo ps -aux | grep guacd` and note the guacd pid(s) (minus the process from the grep)
3. Kill all running guacd processes: `kill -9 pid` replacing *pid* with the pid(s) of the target processes
4. Start guacd again: `morpheus-ctl start guacd`
5. Tail the guacd logs to verify guacd is started and listening: `morpheus-ctl tail guacd` The log output will resemble below when guacd is properly running:

```
guacd[16899]: INFO:      Guacamole proxy daemon (guacd) version 0.9.14 started
guacd[16899]: INFO:      Listening on host 127.0.0.1, port 4822
```

6. Additional information in the guacd logs appears when Morpheus is making a console connection. A successful connection will resemble:

```
guacd[24725]: INFO: Creating new client for protocol "ssh"
guacd[24725]: INFO: Connection ID is "$24f67856-f050-4a17-83eb-9101g0cd8869"
guacd[24743]: INFO: Current locale does not use UTF-8. Some characters may not
↳render correctly.
guacd[24743]: INFO: User "@63102f19-eff4-412e-b1f9-718405f55782" joined
↳connection "$24f67856-f050-4a17-83eb-9101g0cd8869" (1 users now present)
guacd[24743]: INFO: Auth key successfully imported.
guacd[24743]: INFO: SSH connection successful.
```

Guacamole Version

In Morpheus versions 3.2.0 and higher, Guacamole version 0.9.14 is automatically installed. On Morpheus versions older than 3.2.0, 0.9.9 is installed. The 0.9.14 version is required for VNC Hypervisor Console functionality on ESXi v6.5 and later.

Note Guacamole version 0.9.14 is not compatible with MacOS Platform Types over VNC on ESXi v6.0 or prior (6.5 is supported). If necessary, the guacamole version can be reverted to 0.9.9.

To revert the guacamole version from 0.9.14 to 0.9.9.

1. Kill guacd - `morpheus-ctl kill guacd`
2. Check if any guacd processes are still running `ps -aux | grep guac`
3. If so, kill the processes `kill -9 pid` with *id* being the actual process id, like 16101.
4. Go to the guac 0.9.9 directory: `cd /var/opt/morpheus/guacamole-server-0.9.9`
5. Run: `make install`

6. Start guacd: `morpheus-ctl start guacd`

SSL Self-signed Certificate Regeneration

When Morpheus is deployed it generates a 10 year self-signed non-trusted SSL certificate. Below details the process to regenerate this certificate and key.

Replacing both the certificate and private key

1. Delete the certificate and key files in `/etc/morpheus/ssl/` that end in `.crt` and `.key`
2. Run Reconfigure `morpheus-ctl reconfigure`
3. Restart NGINX `morpheus-ctl restart nginx`

Replacing only the certificate

1. Delete the certificate file in `/etc/morpheus/ssl/` it ends in `.crt`
2. Run Reconfigure `morpheus-ctl reconfigure`
3. Restart NGINX `morpheus-ctl restart nginx`

Unable to Delete Tenant

Problem When trying to delete a tenant, a message stating managed resources must be removed or other error occurs and the tenant is not deleted. The tenant may be stuck in a deleting status or return to OK status after delete attempt.

Cause All managed resources must be removed from a tenant in order for that tenant to be deleted. This includes instances and their underlying managed vm's

Solution

1. Login or impersonate that an Admin user inside the tenant
2. Navigate to Infrastructure > Hosts
3. Under Hosts and VM's, delete any managed resources
 - Uncheck `remove infrastructure` when deleting a VM to only remove it from Morpheus but not from the underlying hypervisor/cloud
 - You must check `remove associated instances` if the VM has an associated instance
 - If the VM no longer exists but there is still a record in Morpheus, uncheck `remove infrastructure` and check `force delete`
4. Once all managed resources are removed from the tenant, the tenant can then be deleted
5. In certain situations other components may prevent a tenant from being deleted. If you have removed all managed resources from a tenant and the tenant still cannot be deleted, please contact Morpheus support

Warning: Managed resources can also be removed by deleting instances, but be aware this will delete VM's associated with the instance from the underlying hypervisor/cloud

Unable to Provision a Custom Image

Prior to provisioning an custom image, the image must be configured in the Provisioning -> Virtual Images section by selecting Edit on the Actions dropdown of the Virtual Image.

In the Edit Virtual Image pane:

1. Select “Cloud Init Enabled?” only if the Virtual Image is a linux image with cloud init installed.
2. Enter the username and password that are set on the Virtual Image.

Note: When using Static IP’s or IP Pools in VMware, VMware tools must also be installed on the template in order for Morpheus to set the static IP address when provisioning.

Note: Morpheus agents only support 64-bit vm’s prior to versions 2.12.3 and 3.0.2

Variables

A vast number of variables are available for use in Tasks, Scripts, Templates, Resource Names, Cloud-Init User Data and Option List configs.

Important: Variables are case sensitive

Pre-Provision Vars

A subset of variables are available for Instance, Host Name and Hostnames. These can be passed inside `${ }` blocks during provisioning or in relevant policy configs.

Instance Naming Policy example: `${userInitials}-${cloudCode}-${platform == 'windows' ? 'W' : 'L'}-${sequence}`

Commonly used variables for naming patterns include:

```
${groupName}
${groupCode}
${cloudName}
${cloudCode}
${type}
${accountId}
${account}
${accountType}
${platform}
${platform == 'windows' ? 'w':'l'} # results in `w` for Windows platforms and `l` for
↳Linux Platforms
${userId}
${username}
${userInitials}
${provisionType}
${instance.instanceContext} # Environment Code
${sequence} # results in 1
${sequence+100} # results in 101
```

(continues on next page)

(continued from previous page)

```

${customOption.name}
${sequence.toString().padLeft(5,'0')} #results in 00001

```

An example Instance Name Policy using a naming pattern with User Initials, Cloud Code, Instance Type, and a sequential number starting at 3000 is `${userInitials}-${cloudCode}-${type}-${sequence+3000}`, resulting in an Instance Name of **md-vmwd3-centos-3001** for the first instance, followed by **md-vmwd3-centos-3002** and so on.

Syntax Examples

PowerShell Example: `$app_id = "<%= instance.metadata.app_id %>"`

Bash Example: `HOSTNAME="<%= container.server.hostname %>"`

Python Example: `hostname = morpheus['server']['hostname']`

HTTP Body Example: `{"name": "<%= instance.createdByUsername %>"}`

Note: customOptions values are defined from custom Option Types.

Common Examples

```

container.configGroup: <%=container.configGroup%>
container.configId: <%=container.configId%>
container.configPath: <%=container.configPath%>
container.configRole: <%=container.configRole%>
container.containerTypeCode: <%=container.containerTypeCode%>
container.containerTypeName: <%=container.containerTypeName%>
container.containerTypeShortName: <%=container.containerTypeShortName%>
container.cores: <%=container.cores%>
container.dataPath: <%=container.dataPath%>
container.dateCreated: <%=container.dateCreated%>
container.domainName: <%=container.domainName%>
container.environmentPrefix: <%=container.environmentPrefix%>
container.externalIp: <%=container.externalIp%>
container.hostMountPoint: <%=container.hostMountPoint%>
container.hostname: <%=container.hostname%>
container.image: <%=container.image%>
container.internalHostname: <%=container.internalHostname%>
container.internalIp: <%=container.internalIp%>
container.logsPath: <%=container.logsPath%>
container.memory: <%=container.memory%>
container.planCode: <%=container.planCode%>
container.provisionType: <%=container.provisionType%>
container.server: <%=container.server.serverTypeName%>
container.serverId: <%=container.serverId%>
container.sshHost: <%=container.sshHost%>
container.status: <%=container.status%>
container.storage: <%=container.storage%>

```

(continues on next page)

(continued from previous page)

```

container.version: <%=container.version%>
customOptions: <%=customOptions.fieldName%>
evar: <%=evars.name%>
evars: <%=evars%>
group.code: <%=group.code%>
group.datacenterId: <%=group.datacenterId%>
group.location: <%=group.location%>
group.name: <%=group.name%>
instance.autoScale: <%=instance.autoScale%>
instance.configGroup: <%=instance.configGroup%>
instance.configId: <%=instance.configId%>
instance.configRole: <%=instance.configRole%>
instance.containers[0]: <%=instance.containers[0].containerTypeName%>
instance.cores: <%=instance.cores%>
instance.createdByEmail: <%=instance.createdByEmail%>
instance.createdByFirstName: <%=instance.createdByFirstName%>
instance.createdById: <%=instance.createdById%>
instance.createdByLastName: <%=instance.createdByLastName%>
instance.createdByUsername: <%=instance.createdByUsername%>
instance.deployGroup: <%=instance.deployGroup%>
instance.description: <%=instance.description%>
instance.displayName: <%=instance.displayName%>
instance.domainName: <%=instance.domainName%>
instance.environmentPrefix: <%=instance.environmentPrefix%>
instance.expireDate: <%=instance.expireDate%>
instance.firewallEnabled: <%=instance.firewallEnabled%>
instance.hostname: <%=instance.hostname%>
instance.instanceContext: <%=instance.instanceContext%> (tip: instanceContext = Environment)
instance.instanceLevel: <%=instance.instanceLevel%>
instance.instanceTypeCode: <%=instance.instanceTypeCode%>
instance.instanceTypeName: <%=instance.instanceTypeName%>
instance.instanceVersion: <%=instance.instanceVersion%>
instance.memory: <%=instance.memory%>
instance.metadata: <%=instance.metadata%>
instance.name: <%=instance.name%>
instance.networkLevel: <%=instance.networkLevel%>
instance.plan: <%=instance.plan%>
instance.provisionType: <%=instance.provisionType%>
instance.status: <%=instance.status%>
instance.statusMessage: <%=instance.statusMessage%>
instance.storage: <%=instance.storage%>
instance.tags: <%=instance.tags%>
instance.userStatus: <%=instance.userStatus%>
server.agentInstalled: <%=server.agentInstalled%>
server.agentVersion: <%=server.agentVersion%>
server.apiKey: <%=server.apiKey%>
server.category: <%=server.category%>
server.commType: <%=server.commType%>
server.configGroup: <%=server.configGroup%>
server.configId: <%=server.configId%>
server.configRole: <%=server.configRole%>
server.consoleHost: <%=server.consoleHost%>
server.consolePort: <%=server.consolePort%>
server.consoleType: <%=server.consoleType%>
server.consoleUsername: <%=server.consoleUsername%>
server.dataDevice: <%=server.dataDevice%>

```

(continues on next page)

(continued from previous page)

```

server.dateCreated: <%=server.dateCreated%>
server.description: <%=server.description%>
server.displayName: <%=server.displayName%>
server.domainName: <%=server.domainName%>
server.externalId: <%=server.externalId%>
server.externalIp: <%=server.externalIp%>
server.fqdn: <%=server.fqdn%>
server.hostname: <%=server.hostname%>
server.internalId: <%=server.internalId%>
server.internalIp: <%=server.internalIp%>
server.internalName: <%=server.internalName%>
server.internalSshUsername: <%=server.internalSshUsername%>
server.lastAgentUpdate: <%=server.lastAgentUpdate%>
server.lvmEnabled: <%=server.lvmEnabled%>
server.macAddress: <%=server.macAddress%>
server.managed: <%=server.managed%>
server.maxCores: <%=server.maxCores%>
server.maxMemory: <%=server.maxMemory%>
server.maxStorage: <%=server.maxStorage%>
server.name: <%=server.name%>
server.nodePackageVersion: <%=server.nodePackageVersion%>
server.osDevice: <%=server.osDevice%>
server.osType: <%=server.osType%>
server.osTypeCode: <%=server.osTypeCode%>
server.parentServerId: <%=server.parentServerId%>
server.plan: <%=server.plan%>
server.platform: <%=server.platform%>
server.platformVersion: <%=server.platformVersion%>
server.powerState: <%=server.powerState%>
server.serialNumber: <%=server.serialNumber%>
server.serverModel: <%=server.serverModel%>
server.serverType: <%=server.serverType%>
server.serverTypeCode: <%=server.serverTypeCode%>
server.serverTypeName: <%=server.serverTypeName%>
server.serverVendor: <%=server.serverVendor%>
server.softwareRaid: <%=server.softwareRaid%>
server.sourceImageId: <%=server.sourceImageId%>
server.sshHost: <%=server.sshHost%>
server.sshPort: <%=server.sshPort%>
server.sshUsername: <%=server.sshUsername%>
server.status: <%=server.status%>
server.statusMessage: <%=server.statusMessage%>
server.tags: <%=server.tags%>
server.toolsInstalled: <%=server.toolsInstalled%>
server.visibility: <%=server.visibility%>
task.results (using task code): <%=results.taskCode%>
task.results (using task name): <%=results["Task Name"]%>
task.results.value: <%=results.taskCode.key%>
zone.agentMode: <%=zone.agentMode%>
zone.cloudTypeCode: <%=zone.cloudTypeCode%>
zone.cloudTypeName: <%=zone.cloudTypeName%>
zone.code: <%=zone.code%>
zone.domainName: <%=zone.domainName%>
zone.firewallEnabled: <%=zone.firewallEnabled%>
zone.location: <%=zone.location%>
zone.name: <%=zone.name%>
zone.regionCode: <%=zone.regionCode%>

```

(continues on next page)

(continued from previous page)

```
zone.scalePriority: <%=zone.scalePriority%>
cypher: <%=cypher.read('secret/hello')%>
```

Instance

```
instance {
    autoScale,
    configGroup,
    configId,
    configRole
    containers:[],
    cores,
    deployGroup,
    description,
    displayName,
    domainName,
    environmentPrefix,
    evars:[],
    expireDate,
    firewallEnabled,
    hostname,
    instanceContext,
    instanceLevel,
    instanceTypeCode,
    instanceVersion,
    memory,
    metadata:[],
    name,
    networkLevel,
    plan,
    provisionType,
    status,
    statusMessage,
    storage,
    tags,
    tenantSubdomain,
    userStatus,
    instanceTypeName
}
```

Container

```
container {
    configGroup,
    configId,
    configPath,
    configRole,
    containerTypeCode,
    containerTypeShortName,
    cores,
    dataPath,
    dateCreated,
```

(continues on next page)

(continued from previous page)

```
    domainName,  
    environmentPrefix,  
    externalIp,  
    hostMountPoint,  
    hostname,  
    image,  
    internalHostname,  
    internalIp,  
    logsPath,  
    memory,  
    planCode,  
    provisionType,  
    server:{},  
    serverId,  
    sshHost,  
    status,  
    storage,  
    version,  
    containerTypeName  
}
```

Server

```
server {  
    agentInstalled,  
    agentVersion,  
    apiKey,  
    category,  
    commType,  
    configGroup,  
    configId,  
    configRole,  
    consoleHost,  
    consolePort,  
    consoleType,  
    consoleUsername,  
    dataDevice,  
    dateCreated,  
    description,  
    displayName,  
    domainName,  
    externalId,  
    externalIp,  
    fqdn,  
    hostname,  
    internalId,  
    internalIp,  
    internalName,  
    internalSshUsername,  
    lastAgentUpdate,  
    lvmEnabled,  
    macAddress,  
    managed,  
    maxCores,
```

(continues on next page)

(continued from previous page)

```
maxMemory,
maxStorage,
name,
nodePackageVersion,
osDevice,
osType,
osTypeCode,
parentServerId,
plan,
platform,
platformVersion,
powerState,
serialNumber,
serverModel,
serverType,
serverTypeCode,
serverTypeName,
serverVendor,
softwareRaid,
sourceImageId,
sshHost,
sshPort,
sshUsername,
status,
statusMessage,
tags,
toolsInstalled,
visibility,
volumes {
    name
    id
    deviceName
    maxStorage
    unitNumber
    displayOrder
    rootVolume
}
}
```

Zone (Cloud)

```
zone {
    agentMode,
    cloudTypeCode,
    cloudTypeName,
    code,
    datacenterId,
    domainName,
    firewallEnabled,
    location,
    name,
    regionCode,
    scalePriority
}
```

Group (Site)

```
group {
  code,
  location,
  datacenterId,
  name
}
```

Custom Options (Option Types)

```
customOptions {
  customOptions.fieldName
}
```

Global

ex: <%= morpheus.user.id %>

```
"morpheus":{
  "user":{
    "id":value,
    "account":{
      "id":value
    },
    "username":"value",
    "displayName":"value",
    "email":"value",
    "firstName":"value",
    "lastName":"value",
    "dateCreated":0000-00-00T00:00:00Z,
    "lastUpdated":0000-00-00T00:00:00Z,
    "enabled":true/false,
    "accountExpired":true/false,
    "accountLocked":false,
    "passwordExpired":false,
    "defaultGroupId":value,
    "defaultZoneId":value,
    "hasLinuxUser":true/false,
    "hasWindowsUser":true/false,
    "role":{
      "id":value
    },
    "instanceLimits":value
  },
}
```

Instance Map Example

```
"instance":{
  "poolProviderType":value,
  "isVpcSelectable":true/false,
  "smbiosAssetTag":value,
  "isEC2":true/false,
  "resourcePoolId":value,
  "hostId":value,
  "createUser":true/false,
  "nestedVirtualization":value,
  "vmwareFolderId":value,
  "expose":[

  ],
  "noAgent":value,
  "customOptions":value,
  "createBackup":true/false,
  "memoryDisplay":"MB/GB",
  "backup":{
    "veeamManagedServer":,
    "createBackup":true/false,
    "jobAction":"value",
    "jobRetentionCount":value
  },
  "expireDays":value,
  "layoutSize":value,
  "lbInstances":[

  ],
  "evars":{
    "evarl":{
      "value":value,
      "export":true/false,
      "masked":true/false,
      "name":"value"
    },
    "evar2":{
      "value":value,
      "export":true/false,
      "masked":true/false,
      "name":"value"
    }
  },
  "id":value,
  "instanceTypeName":"value",
  "instanceTypeCode":"value",
  "provisionType":"value",
  "layoutId":value,
  "layoutCode":value,
  "layoutName":"value",
  "instanceVersion":"value",
  "plan":value,
  "name":value,
  "displayName":value,
  "description":value,
  "environmentPrefix":value,
```

(continues on next page)

(continued from previous page)

```

    "hostname":value,
    "domainName":"value",
    "assignedDomainName":,
    "firewallEnabled":true/false,
    "status":"value",
    "userStatus":"value",
    "scheduleStatus":"value",
    "networkLevel":"value",
    "instanceLevel":"value",
    "deployGroup":value,
    "instanceContext":value,
    "autoScale":true/false,
    "statusMessage":value,
    "expireDate":0000-00-00T00:00:00Z,
    "tags":"value",
    "storage":value (bytes),
    "memory":value (bytes),
    "cores":1,
    "configId":value,
    "configGroup":value,
    "configRole":value,
    "ports":value,
    "sslEnabled":true/false,
    "sslCertId":value,
    "serviceUsername":value,
    "servicePassword":value,
    "adminUsername":value,
    "adminPassword":value,
    "createdByUsername":"value",
    "createdByEmail":"value",
    "createdByFirstName":"value",
    "createdByLastName":"value",
    "createdById":value,
    "metadata":{

    },
    "createdByUser":{
        "username":"value",
        "displayName":"value",
        "firstName":"value",
        "lastName":"value",
        "email":"value",
        "linuxUsername":"value",
        "windowsUsername":"value"
    },
    "containers":[
        {
            "maxMemory":value (bytes),
            "maxStorage":value (bytes),
            "maxCpu":value,
            "maxCores":value,
            "coresPerSocket":value,
            "poolProviderType":value,
            "isVpcSelectable":true/false,
            "smbiosAssetTag":value,
            "isEC2":true/false,
            "resourcePoolId":value,

```

(continues on next page)

(continued from previous page)

```

    "hostId":value,
    "createUser":true/false,
    "nestedVirtualization":value,
    "vmwareFolderId":value,
    "expose":[
    ],
    "noAgent":true/false,
    "vm":true/false,
    "networkInterfaces":[
      {
        "id":value,
        "network":{
          "id":value,
          "group":value,
          "subnet":value,
          "dhcpServer":true/false,
          "name":value,
          "pool":{
            "id":value,
            "name":value
          }
        },
        "ipAddress":value,
        "networkInterfaceTypeId":value,
        "ipMode":
      }
    ],
    "volumes":[
      {
        "volumeCustomizable":true/false,
        "readonlyName":true/false,
        "controllerId":value,
        "maxIOPS":value,
        "displayOrder":value,
        "unitNumber":value,
        "minStorage":value(bytes),
        "configurableIOPS":true/false,
        "controllerMountPoint":0000:0:00:0,
        "vId":value,
        "size":value,
        "name":"root",
        "rootVolume":true/false,
        "storageType":value,
        "typeId":value,
        "id":value,
        "resizeable":true/false,
        "datastoreId":"value",
        "maxStorage":value(bytes)
      }
    ],
    "storageController":value,
    "datastoreId":value,
    "networkId":value,
    "cpuCount":value,
    "memorySize":value,
    "osDiskSize":value,

```

(continues on next page)

(continued from previous page)

```

    "publicKeyId":value,
    "storagePodId":value,
    "vmwareUsr":value,
    "vmwarePwd":value,
    "domainName":"value",
    "hostname":value,
    "networkType":value,
    "ipAddress":value,
    "netmask":value,
    "gateway":value,
    "dnsServers":value,
    "resourcePool":value,
    "folder":value,
    "vmwareCustomSpec":value,
    "hosts":{
      value
    },
    "evars":{

    },
    "id":value,
    "name":value,
    "containerTypeName":value,
    "containerTypeCode":value,
    "containerTypeShortName":"value",
    "containerTypeCategory":"value",
    "provisionType":"value",
    "dataPath":"value",
    "logsPath":"value",
    "configPath":"value",
    "planCode":value,
    "dateCreated":"0000-00-00T00:00:00Z",
    "status":"running",
    "environmentPrefix":"value",
    "version":"value",
    "image":"value",
    "internalHostname":value,
    "storage":value (bytes),
    "memory":value (bytes),
    "cores":value,
    "internalIp":value,
    "externalIp":value,
    "sshHost":value,
    "hostMountPoint":value,
    "configId":value,
    "configGroup":value,
    "configRole":value,
    "certificatePath":value,
    "certificateStyle":value,
    "changeManagementExtId":value,
    "changeManagementServiceId":value,
    "serverId":value,
    "server":{
      "poolProviderType":value,
      "isVpcSelectable":true/false,
      "smbiosAssetTag":value,
      isEC2:true/false,

```

(continues on next page)

(continued from previous page)

```
"resourcePoolId":value,
"hostId":value,
"createUser":true/false,
"nestedVirtualization":value,
"vmwareFolderId":value,
"noAgent":value,
"id":value,
"uuid":value,
"serverTypeName":"value",
"serverTypeCode":"value",
"computeTypeName":"value",
"computeTypeCode":"value",
"parentServerId":value,
"plan":value,
"visibility":"value",
"osTypeCode":value,
"sourceImageId":value,
"name":value,
"displayName":value,
"internalName":value,
"category":value,
"description":value,
"internalId":value,
"externalId":value,
"platform":"value",
"platformVersion":value,
"agentVersion":value,
"nodePackageVersion":value,
"sshHost":value,
"sshPort":value,
"sshUsername":"value",
"consoleType":value,
"consoleHost":value,
"consolePort":value,
"consoleUsername":value,
"internalSshUsername":"value",
"internalIp":value,
"externalIp":value,
"osDevice":"value",
"dataDevice":"value",
"lvmEnabled":true/false,
"apiKey":value,
"softwareRaid":true/false,
"status":"value",
"powerState":"value",
"dateCreated":0000-00-00T00:00:00Z,
"lastAgentUpdate":0000-00-00T00:00:00Z,
"serverType":"value",
"osType":"value",
"commType":"value",
"managed":true/false,
"agentInstalled":true/false,
"toolsInstalled":true/false,
"hostname":value,
"domainName":value,
"fqdn":value,
"statusMessage":value,
```

(continues on next page)

(continued from previous page)

```

    "maxStorage":value(bytes),
    "maxMemory":value(bytes),
    "maxCores":value,
    "macAddress":value,
    "serverVendor":value,
    "serverModel":value,
    "serialNumber":value,
    "tags":value,
    "configId":value,
    "configGroup":value,
    "configRole":value,
    "createdByUser":{
      "username":"value",
      "displayName":"value",
      "firstName":"value",
      "lastName":"value",
      "email":"value",
      "linuxUsername":"value",
      "windowsUsername":"value"
    },
    "volumes":[
      {
        "id":value,
        "name":"value",
        "deviceName":"value",
        "maxStorage":value(bytes),
        "unitNumber":value,
        "displayOrder":value,
        "rootVolume":true/false
      }
    ]
  },
  "ports":[
    {
      "index":value,
      "external":value,
      "internal":value,
      "link":true/false,
      "loadBalance":true/false,
      "loadBalanceProtocol":value,
      "export":true/false,
      "exportName":value,
      "displayName":"value",
      "visible":true/false,
      "primaryPort":true/false,
      "protocol":value,
      "name":"value"
    }
  ],
  "portMap":{
    "rpc":{
      "index":value,
      "external":value,
      "internal":value,
      "link":true/false,
      "loadBalance":true/false,
      "loadBalanceProtocol":value,

```

(continues on next page)

(continued from previous page)

```

        "export":true/false,
        "exportName":value,
        "displayName":"value",
        "visible":true/false,
        "primaryPort":true/false,
        "protocol":value,
        "name":"value"
    }
},
"internalPort":value,
"externalPort":value
}
],
"container":{
    "maxMemory":value(bytes),
    "maxStorage":value,
    "maxCpu":value,
    "maxCores":value,
    "coresPerSocket":value,
    "poolProviderType":value,
    "isVpcSelectable":true/false,
    "smbiosAssetTag":value,
    "isEC2":true/false,
    "resourcePoolId":value,
    "hostId":value,
    "createUser":true/false,
    "nestedVirtualization":value,
    "vmwareFolderId":value,
    "expose":[

    ],
    "noAgent":true/false,
    "vm":true/false,
    "networkInterfaces":[
        {
            "id":value,
            "network":{
                "id":value,
                "group":value,
                "subnet":value,
                "dhcpServer":true/false,
                "name":value,
                "pool":{
                    "id":value,
                    "name":value
                }
            }
        },
        "ipAddress":value,
        "networkInterfaceTypeId":value,
        "ipMode":

    ]
},
"volumes":[
    {
        "volumeCustomizable":true/false,
        "readonlyName":true/false,
        "controllerId":value,

```

(continues on next page)

(continued from previous page)

```

        "maxIOPS":value,
        "displayOrder":value,
        "unitNumber":value,
        "minStorage":value,
        "configurableIOPS":true/false,
        "controllerMountPoint":value,
        "vId":value,
        "size":value,
        "name":"root",
        "rootVolume":true/false,
        "storageType":value,
        "typeId":value,
        "id":value,
        "resizeable":true/false,
        "datastoreId":"autoCluster",
        "maxStorage":value (bytes)
    }
},
"storageController":value,
"datastoreId":value,
"networkId":value,
"cpuCount":value,
"memorySize":value,
"osDiskSize":value,
"publicKeyId":value,
"storagePodId":value,
"vmwareUsr":value,
"vmwarePwd":value,
"domainName":"value",
"hostname":value,
"networkType":value,
"ipAddress":value,
"netmask":value,
"gateway":value,
"dnsServers":value,
"resourcePool":value,
"folder":value,
"vmwareCustomSpec":value,
"hosts":{
    value
},
"evars":{

},
"id":value,
"name":value,
"containerTypeName":value,
"containerTypeCode":value,
"containerTypeShortName":"value",
"containerTypeCategory":"value",
"provisionType":"vmware",
"dataPath":"value",
"logsPath":"value",
"configPath":"value",
"planCode":value,
"dateCreated":"0000-00-00T00:00:00Z",
"status":"value",

```

(continues on next page)

(continued from previous page)

```
"environmentPrefix":"value",
"version":"value",
"image":"value",
"internalHostname":value,
"storage":value(bytes),
"memory":value(bytes),
"cores":value,
"internalIp":value,
"externalIp":value,
"sshHost":value,
"hostMountPoint":value,
"configId":value,
"configGroup":value,
"configRole":value,
"certificatePath":value,
"certificateStyle":value,
"changeManagementExtId":value,
"changeManagementServiceId":value,
"serverId":value,
"server":{
  "poolProviderType":value,
  "isVpcSelectable":true/false,
  "smbiosAssetTag":value,
  isEC2:true/false,
  "resourcePoolId":value,
  "hostId":value,
  "createUser":true/false,
  "nestedVirtualization":value,
  "vmwareFolderId":value,
  "noAgent":value,
  "id":value,
  "uuid":value,
  "serverTypeName":"value",
  "serverTypeCode":"value",
  "computeTypeName":"value",
  "computeTypeCode":"value",
  "parentServerId":value,
  "plan":value,
  "visibility":"value",
  "osTypeCode":value,
  "sourceImageId":value,
  "name":value,
  "displayName":value,
  "internalName":value,
  "category":value,
  "description":value,
  "internalId":value,
  "externalId":value,
  "platform":"value",
  "platformVersion":value,
  "agentVersion":value,
  "nodePackageVersion":value,
  "sshHost":value,
  "sshPort":value,
  "sshUsername":"value",
  "consoleType":value,
  "consoleHost":value,
```

(continues on next page)

(continued from previous page)

```

    "consolePort":value,
    "consoleUsername":value,
    "internalSshUsername":"value",
    "internalIp":value,
    "externalIp":value,
    "osDevice":"value",
    "dataDevice":"value",
    "lvmEnabled":true/false,
    "apiKey":value,
    "softwareRaid":true/false,
    "status":"provisioned",
    "powerState":"on",
    "dateCreated":0000-00-00T00:00:00Z,
    "lastAgentUpdate":0000-00-00T00:00:00Z,
    "serverType":"value",
    "osType":"value",
    "commType":"value",
    "managed":true/false,
    "agentInstalled":true/false,
    "toolsInstalled":true/false,
    "hostname":value,
    "domainName":value,
    "fqdn":value,
    "statusMessage":value,
    "maxStorage":value,
    "maxMemory":value,
    "maxCores":value,
    "macAddress":value,
    "serverVendor":value,
    "serverModel":value,
    "serialNumber":value,
    "tags":value,
    "configId":value,
    "configGroup":value,
    "configRole":value,
    "createdByUser":{
      "username":"value",
      "displayName":"value",
      "firstName":"value",
      "lastName":"value",
      "email":"value",
      "linuxUsername":"value",
      "windowsUsername":"value"
    },
    "volumes":[
      {
        "id":value
        "name":"root",
        "deviceName":"value",
        "maxStorage":value(bytes),
        "unitNumber":value,
        "displayOrder":value,
        "rootVolume":true/false
      }
    ]
  },
  "ports":[

```

(continues on next page)

(continued from previous page)

```

    {
      "index":0,
      "external":value,
      "internal":value,
      "link":true/false,
      "loadBalance":true/false,
      "loadBalanceProtocol":value,
      "export":true/false,
      "exportName":value,
      "displayName":"value",
      "visible":true/false,
      "primaryPort":true/false,
      "protocol":value,
      "name":"value"
    }
  ],
  "portMap":{
    "rpc":{
      "index":0,
      "external":value,
      "internal":value,
      "link":true/false,
      "loadBalance":true/false,
      "loadBalanceProtocol":value,
      "export":true/false,
      "exportName":value,
      "displayName":"value",
      "visible":true/false,
      "primaryPort":true/false,
      "protocol":value,
      "name":"value"
    }
  },
  "internalPort":value,
  "externalPort":value
},
"apps":[
]
}

```

1.3.13 Guides

Getting started with Morpheus and AWS

Introduction

This guide is designed to help you get started and quickly get the most out of Morpheus with AWS. By the end, you will integrate your first cloud, configure networking, prepare and consume images, provision instances, and get started with automation. We will briefly discuss installation and account setup but will provide links to additional resources for those very first steps. For the most part, this guide assumes you are able to get Morpheus installed and are ready to move forward from that point. There is a lot more to see and do in Morpheus that is beyond the scope of this guide. For more, consult the complete Morpheus documentation or take part in our user community forum.

Installation & Setup

In the simplest configuration, Morpheus needs one appliance server which will contain all the components necessary to orchestrate virtual machines and containers. Full requirements, including storage and networking considerations, can be found in Morpheus documentation [here](#). In order to provision any new instances, hosts, or applications, (or convert any discovered resources to managed resources) you will need a valid license. If you don't have one, Morpheus will automatically set up a lab license on installation. A lab license is a time-unlimited license for Morpheus that limits you to 25 managed and discovered workloads. If you have a timed trial or a paid license, the license can be applied in Administration > Settings > LICENSE.

Groups

Groups in Morpheus define which resources a user has access to. Clouds are added to groups and a user can only access clouds that are in the groups to which their roles give them access. More information on Morpheus groups is [here](#). A deep dive into groups goes beyond the scope of this guide but it's often useful to create a group that contains all clouds for testing purposes. We will create that group now so that we can add our first cloud into this group in the next section.

Navigate to *Infrastructure > Groups*. Here we will see a list of all configured groups but, of course, this will be empty immediately after installation. Click "+CREATE". Give your group a name, such as "All Clouds". The "CODE" field is used when calling Morpheus through Morpheus API or Morpheus CLI. It's useful in most cases to have an "All Clouds" group for testing purposes so this will likely help you down the road.

NEW GROUP

Configuration

NAME

All Clouds

CODE

LOCATION

► Advanced Options

SAVE CHANGES

Click "SAVE CHANGES". Your group is now ready to accept clouds.

Integrating Your First Cloud

Clouds in Morpheus consist of any consumable endpoint whether that be On-Prem, Public clouds, or even bare metal. In this guide, we will focus on integrating and working with AWS.

To get started, we will navigate to *Infrastructure > Clouds*. This is the cloud detail page which lists all configured clouds. It will be empty if you've just completed installation and setup of Morpheus but soon we will see our integrated AWS cloud here.

Click the “+ADD” button to pop the “CREATE CLOUD” wizard. Select “AMAZON” and click the “NEXT” button.

On the “CONFIGURE” tab, give your cloud a “NAME”. Select your Amazon region and enter your credentials in the “ACCESS KEY” and “SECRET KEY” fields. We can also set a number of advanced options here that may be relevant to your AWS use case.

Once you're satisfied with your selections, click “NEXT”

We have now arrived at the “GROUP” tab. In this case, we will mark the radio button to “USE EXISTING” groups if you wish to use the group we configured earlier.

Once you've selected the group, click “NEXT”

On the final tab of the “CREATE CLOUD” wizard, you'll confirm your selections and click “COMPLETE”. The new cloud is now listed on the cloud detail page. After a short time, Morpheus will provide summary information and statistics on existing virtual machines, networks, and other resources available in the cloud.

Viewing Cloud Inventory

Now that we've integrated our first AWS cloud, we can stop for a moment to review what Morpheus gives us from the cloud detail page. We can see that Morpheus gives us estimated costs and cost histories, metrics on used resources, and also lists out resource counts in various categories including container hosts, hypervisors, and virtual machines. We can drill into these categories to see lists of resources in the various categories individual resources within them by clicking on the category tabs. We can link to the detail page for any specific resource by clicking on it from its resource category list.

Configuring Resource Pools

With our AWS cloud configured, Morpheus will automatically sync in available resource pools and data stores.

Resource pools, once Morpheus has had time to ingest them, will be visible from the cloud detail page. Navigate to *Infrastructure > Clouds > (your AWS cloud) > RESOURCES tab*. In here, we are able to see and control access to the various resource pools that have been configured in Amazon. For example, we can restrict access to a specific resource pool within Morpheus completely by clicking on the “ACTIONS” button, then clicking “Edit”. If we unmark the “ACTIVE” button and then click “SAVE CHANGES” we will see that the resource pool is now grayed out in the list. The resources contained in that pool will not be accessible for provisioning within Morpheus.

SUMMARYCLUSTERSHOSTSVMSCONTAINERSLOAD BALANCERSNETWORKSRESOURCESPOLICIESWIKI

POOLS

Search

+ ADD RESOURCE POOL

NAME	DESCRIPTION	VISIBILITY	DEFAULT	TENANT	
My Resource Pool		Private		morpheus	ACTIONS ▾

Often our clients will want to make specific blocks of resources available to certain groups. This can be easily and conveniently controlled through the same “EDIT RESOURCE POOL” dialog box we were just working in. If we expand the “Group Access” drawer, we are able to give or remove access to each pool to any group we’d like. We can also choose to make some or all of our resource pools available to every group. Specific resource pools can also be defined as the default for each group if needed.

EDIT RESOURCE POOL

NAME: labs

☐ MOVE SERVERS

☒ ACTIVE

☐ DEFAULT

▼ Group Access

GROUP	ACCESS	DEFAULT
all	<input checked="" type="checkbox"/>	
AA	<input type="checkbox"/>	<input type="checkbox"/>
AA-DB	<input type="checkbox"/>	<input type="checkbox"/>

Additionally, we may choose to allow only certain service plans to be provisioned into a specific pool of resources. For example, perhaps a specific cluster is my SQL cluster and only specific services plans should be consumable within it. We can control that through this same dialog box.

Configuring Network for Provisioning

When configuring networking, we can set global defaults by going to *Infrastructure > Network > NETWORKS* tab. Here we can add or configure networks from all clouds integrated into Morpheus. Depending on the number of clouds Morpheus has ingested, this list may be quite large and may also be paginated across a large number of pages. In such a case, it may be easier to view or configure networks from the specific cloud detail page so that networks from other clouds are not shown.

NETWORKS

NETWORKS						
NAME	TYPE	CLOUD	CIDR	POOL	DHCP	VISIBILITY TENANTS
10.30.20.0	OracleVM Network	Morpheus Oracle VM	10.30.20.0/22	✓	Public	morpheus
172.31.0.0/20 (subnet-63dffb13b)	Amazon Subnet	ah-only	172.31.0.0/20	✓	Private	morpheus
172.31.16.0/20 (subnet-27110ed6)	Amazon Subnet	ah-only	172.31.16.0/20	✓	Private	morpheus

Still in *Infrastructure > Network*, make note of the “INTEGRATIONS” tab. It’s here that we can set up any integrations that may be relevant, such as IPAM integrations. Generally speaking, when adding IPAM integrations, we simply need to name our new integration, give the API URL, and provide credentials. There’s more information in the [IPAM integration](#) section of Morpheus Docs.

ADD IPAM INTEGRATION

×

NAME

☒ ENABLED

URL

USERNAME

PASSWORD

THROTTLE RATE
ms

In *Infrastructure > Networking* we can also set up IP address pools from the IP Pools tab. These pools can be manually defined, known as a Morpheus-type IP pool, or they can come from any IPAM integrations you’ve configured. As instances are provisioned, Morpheus will assign IP addresses from the pool chosen during provisioning. When the instance is later dissolved, Morpheus will automatically release the IP address to be used by another instance when needed. When adding or editing a network, we can opt to scope the network to one of these configured IP address pools.

CREATE NETWORK POOL

×

NAME

POOL TYPE

Morpheus

▼

IP Ranges

STARTING ADDRESS

ENDING ADDRESS

192.168.0.2

–

192.168.0.255

+

SAVE CHANGES

Since this guide is focused on working within the AWS cloud that we integrated at the start, we will take a look at our network configurations on the cloud detail page as well. Navigate to *Infrastructure > Clouds > (your AWS cloud) > NETWORKS tab*. Just as with resource pools, we have the ability to make certain networks inactive in Morpheus, or scope them to be usable only for certain groups or tenants.

SUMMARY	CLUSTERS	HOSTS	VMS	CONTAINERS	LOAD BALANCERS	NETWORKS	RESOURCES	POLICIES	WIKI
---------	----------	-------	-----	------------	----------------	----------	-----------	----------	------

NETWORKS

Search

Q

ACTIONS +

<input type="checkbox"/> NAME	TYPE	CIDR	POOL	DHCP VISIBILITY	TENANT	
<input type="checkbox"/> labs_lb_1a (subnet-21dd1a44)	Amazon Subnet	10.100.5.0/24		✓ Private	morpheus	ACTIONS ▼
<input type="checkbox"/> labs_lb_1c (subnet-dcb0b49a)	Amazon Subnet	10.100.4.0/24		✓ Private	morpheus	ACTIONS ▼

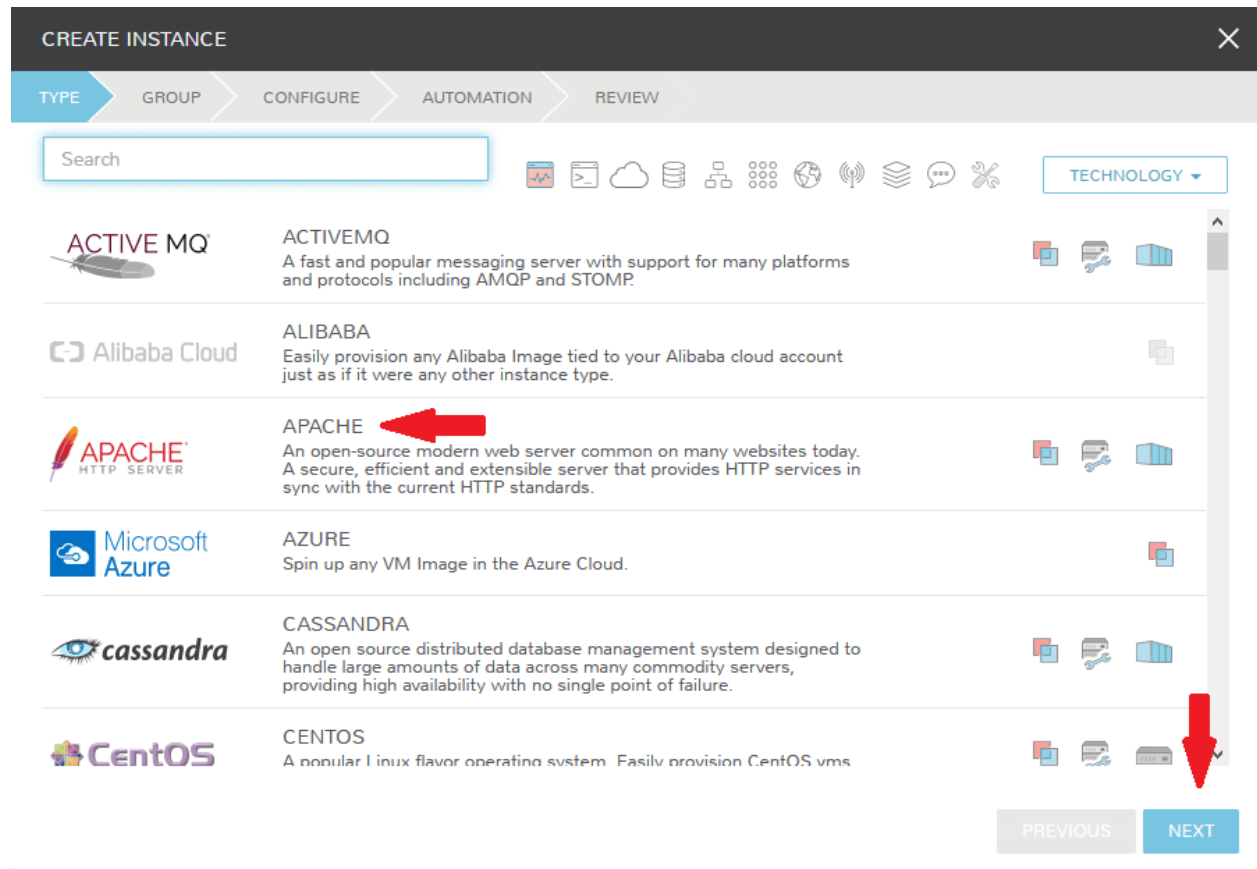
Prepping an Image

As we'll discuss and try out in the next section, Morpheus comes out of the box with a default set of blueprints that are relevant to many modern deployment scenarios. For the most part, these are base operating system images with a few additional adjustments. However, in many on-premise deployments, there are often custom image and networking requirements. We will work with the images included in Morpheus by default but have guides in Morpheus Docs for [creating Windows and Linux images](#) which are consumable in Morpheus.

Provisioning Your First Instance

At this point, we are ready to provision our first image. As a first instance, we'll provision an Apache web server to our AWS cloud.

Navigate to *Provisioning > Instances*. If any instances are currently provisioned, we will see them listed here. To start a new instance we click the “+ADD” button to pop the “CREATE INSTANCE” wizard. We'll scroll down to and select the Apache instance type and click “NEXT”.



First, we'll specify the group to provision into which determines the clouds available. If you've followed this guide to this point, you should at least have a group that houses all of your clouds which you can select here. This will allow us to select the AWS cloud from the “CLOUD” dropdown menu. Provide a unique name to this instance and then click “NEXT”

From the “CONFIGURE” tab, we're presented with a number of options. The options are cloud and layout-specific, more generalized information on creating instances and available options is [here](#). For our purposes, we'll select the following options:

- **LAYOUT:** Includes options such as the base OS, custom layouts will also be here when available
- **PLAN:** Select the resource plan for your instance. Some plans have minimum resource limits, Morpheus will only show plans at or above these limits. User-defined plans can also be created in *Administration > Plans & Pricing*.
- **VOLUMES:** The minimum disk space is set by the plan, this value may be locked if you've selected a custom plan that defines the volume size
- **NETWORKS:** Select a network, note that IP pools must be linked with the networks defined in VMware in order to assign static IP addresses

• SECURITY GROUPS

Under the “User Config” drawer, mark the box to “CREATE YOUR USER”. Click “NEXT”.

Note: “CREATE YOUR USER” will seed a user account into the VM with credentials set in your Morpheus user account settings. If you’ve not yet defined these credentials, you can do so by clicking on your username in the upper-right corner of the application window and selecting “USER SETTINGS”.

For now, we’ll simply click “NEXT” to move through the “AUTOMATION” tab but feel free to stop and take a look at the available selections here. There is more information later in this guide on automation and even more beyond that in the rest of Morpheus docs.

Review the settings for your first instance and click “COMPLETE”.

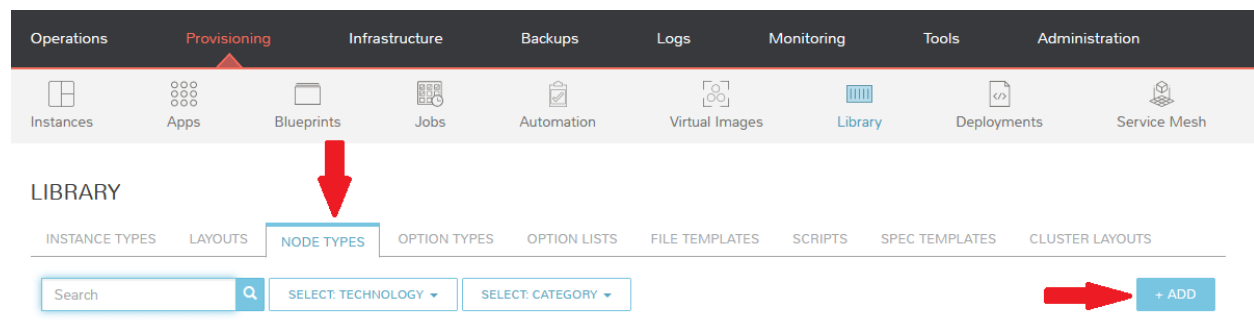
We are now dropped back onto the instances list page. We can see a new entry in the list at this point with a status indicator that the new machine is being launched (rocket icon in the status field). We can double click on the instance in the list to move to the instance detail page. For now we will see a progress bar indicating that the instance is being created and is starting up. The exact amount of time this process will take depends selections made when provisioning the instance. For more detailed information on the status of various provisioning processes, we can scroll down and select the “HISTORY” tab. The “STATUS” icon will change from the blue rocket to a green play button when the instance is fully ready. Furthermore, we can click on the hyperlinked IP address in the “VMS” section of this page to view a default page in a web browser to confirm success.

Creating Your First Library Item

In the prior section, we manually provisioned our first instance. However, Morpheus allows you to build a catalog of custom provisionable items to simplify and speed provisioning in the future. In this section, we’ll build a catalog item and show how that can translate into quick instance provisioning after configuration.

Note: Before starting this process, it’s important to decide which virtual image you plan to use. If you’re not using a Morpheus-provided image, you’ll want to ensure it’s uploaded. You will not be able to complete this section without selecting an available image. In this example we will use Morpheus Redis 3.0 on Ubuntu 14.04.3 v2.

Navigate to *Provisioning > Library > NODE TYPES* and click “+ADD”.



In this example, I am going to set the following options in the “NEW NODE TYPE” wizard:

- **NAME**
- **SHORT NAME**
- **VERSION:** 1 (In this particular case, the version is not important)

- **TECHNOLOGY:** Amazon
- **AMI IMAGE:** Morpheus Redis 3.0 on Ubuntu 14.04.3 v2

NEW NODE TYPE

NAME

SHORT NAME
The short name is a name with no spaces used for display in your container list.

VERSION

TECHNOLOGY

ENVIRONMENT VARIABLES

Name	Value

Amazon Options

AMI IMAGE

LOG FOLDER

DEPLOY FOLDER

With the new node created, we'll now add a new instance type which will be accessible from the provisioning wizard once created. Move from the "NODE TYPES" tab to the "INSTANCE TYPES" tab and click "+ADD".

LIBRARY

INSTANCE TYPES | LAYOUTS | NODE TYPES | OPTION TYPES | OPTION LISTS | FILE TEMPLATES | SCRIPTS | SPEC TEMPLATES | CLUSTER LAYOUTS

Search

In the "NEW INSTANCE TYPE" wizard, I'll simply enter a **NAME** and **CODE** value. Click "SAVE CHANGES".

NEW INSTANCE TYPE

NAME

NewInstanceType

CODE

newtype

Useful shortcode for provisioning naming schemes and export reference.

DESCRIPTION

255 Characters Remaining

CATEGORY

Web



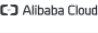


ICON

Browse

Suggested Dimensions: 150 x 51

Now that we've created a new instance type, access it by clicking on the name in the list of custom instances you've created. In my case, I've given the name "NewInstanceType".

LIBRARY

LIBRARY				
<div> <div>INSTANCE TYPES</div> <div>LAYOUTS</div> <div>NODE TYPES</div> <div>OPTION TYPES</div> <div>OPTION LISTS</div> <div>FILE TEMPLATES</div> <div>SCRIPTS</div> <div>SPEC TEMPLATES</div> <div>CLUSTER LAYOUTS</div> </div>				
<div> <div>Search</div> <div>SELECT TECHNOLOGY</div> <div>SELECT CATEGORY</div> <div>+ ADD</div> </div>				
	NAME	TECHNOLOGY	CATEGORY	FEATURED
	NewInstanceType		Web	ACTIONS
	ActiveMQ	Mixed	Messaging	ACTIONS
	Alibaba	Alibaba	Cloud	ACTIONS
	Amazon Api		Apps	ACTIONS
	AmazonMQ		Messaging	ACTIONS

Once we've opened the new instance type, by default, we should be on the "LAYOUTS" tab. Click "+ADD LAYOUT". I've set the following fields on my example layout:

- **NAME**
- **VERSION:** This is the version number of the layout itself, which is labeled 1.0 in the example
- **TECHNOLOGY:** Amazon

- **Nodes:** Select the node we created earlier, if desired you can specify multiple nodes

Click “SAVE CHANGES”.

At this point we’ve completed the setup work and can now provision the instance we’ve created to our specifications. Navigate to *Provisioning > Instances* and click “+ADD”. From the search bar we can search for the new instance type we’ve created. In the example case, we called it “newinstancetype”. Click “NEXT”.

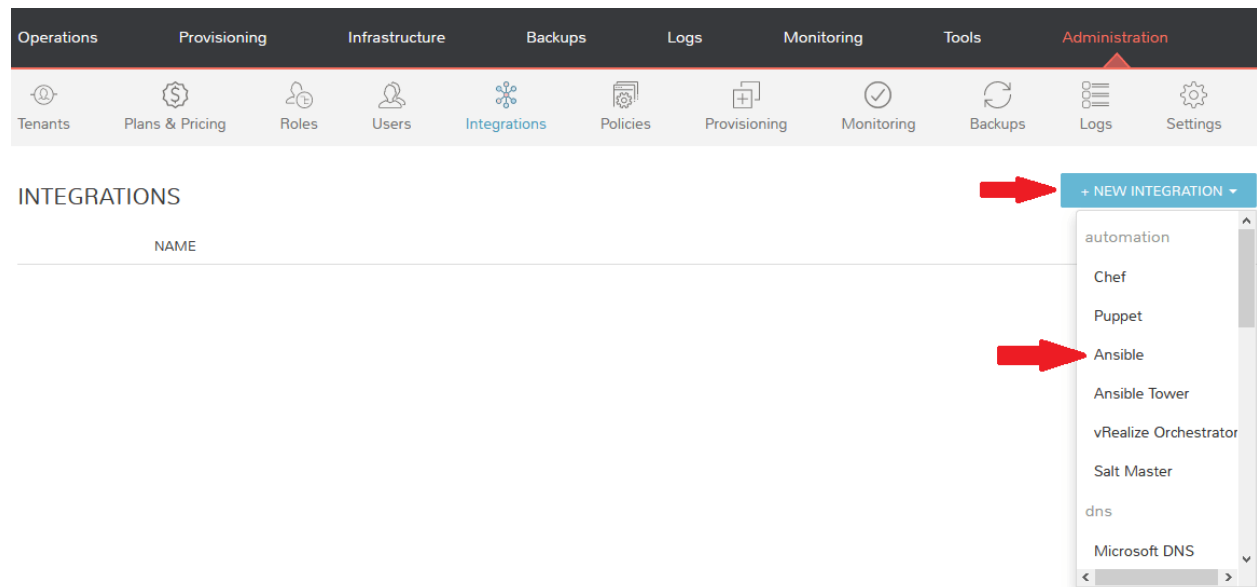
As before, we can select a group and cloud to provision this new instance. Click “NEXT”. On the “CONFIGURE” tab, make note that the layout and plan are already selected because they were configured as part of creating the new instance type. Select a network and click “NEXT”. Once again we will also click “NEXT” through the “AUTOMATION” tab. Finally, click “COMPLETE”.

As before when we manually provisioned an instance, Morpheus will now begin to spin up the new VM. Once the provisioning process has completed, open the instance detail page in Morpheus and click on the “CONSOLE” tab. You’ll be logged in with your user account and are then able to confirm the machine is ready and available.

Automation and Configuration Management

Morpheus automation is composed of Tasks and Workflows. A task could be a script added directly, scripts or blueprints pulled from the Morpheus Library, playbooks, recipes, or a number of other things. The complete list of task types can be found in the [Automation section](#) of Morpheus docs. Tasks can be executed individually but they are often combined into workflows. We can opt to run a workflow at provision time or they can be executed on existing instances through the Actions menu.

In this guide we will set up an Ansible integration, create a task, add the task to a workflow, and run the workflow against a new and existing instance. If you’ve worked through this guide to this point, you should already have an Apache instance running. If you don’t yet have that, provision one before continuing with this guide and ensure it’s reachable on port 80.



We’ll first set up the Ansible integration, you can integrate with the sample repository referenced here or integrate with your own. Go to ‘Administration > Integrations’. Click “+NEW INTEGRATION” and select Ansible from the dropdown menu. Fill in the following details:

- **NAME**
- **ANSIBLE GIT URL:** <https://github.com/ncelebic/morpheus-ansible-example>, or enter the URL for your own Ansible git repository

- **PLAYBOOKS PATH**
- **ROLES PATH**
- Mark the box to “USE MORPHEUS AGENT COMMAND BUS”

Note: If your git repository requires authentication, you should create a keypair and use the following URL format: `git@github.com:ncelebic/morpheus-ansible-example.git`.

NEW ANSIBLE INTEGRATION

NAME

NewAnsibleIntegration

☒ ENABLED

ANSIBLE GIT URL

https://github.com/ncelebic/morpheus-ansible-example

KEY PAIR

PLAYBOOKS PATH

/

ROLES PATH

/roles

GROUP VARIABLES PATH

HOST VARIABLES PATH

☐ USE ANSIBLE GALAXY

☐ ENABLE VERBOSE LOGGING

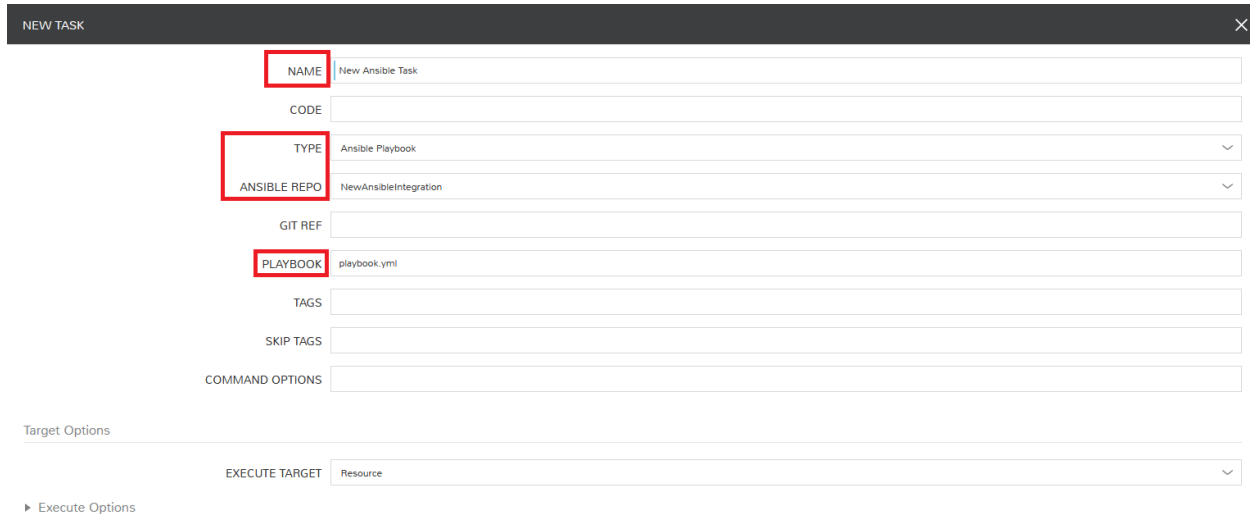
☒ USE MORPHEUS AGENT COMMAND BUS

SAVE CHANGES

Click “SAVE CHANGES”. You’ll now see our new Ansible integration listed among any other configured integrations. If we click on this new integration to view detail, a green checkmark icon indicates the git repository has been fully synced.

With the Ansible integration set up, we can now create a task that includes our playbook. Go to *Provisioning > Automation*, click “+ADD”. We’ll first set our “TYPE” value to Ansible Playbook so that the correct set of fields appear in the “NEW TASK” wizard. Set the following options:

- **NAME**
- **ANSIBLE REPO:** Here we will choose the Ansible integration that we just created
- **PLAYBOOK:** In our example case, enter ‘playbook.yml’



NEW TASK

NAME New Ansible Task

CODE

TYPE Ansible Playbook

ANSIBLE REPO NewAnsibleIntegration

GIT REF

PLAYBOOK playbook.yml

TAGS

SKIP TAGS

COMMAND OPTIONS

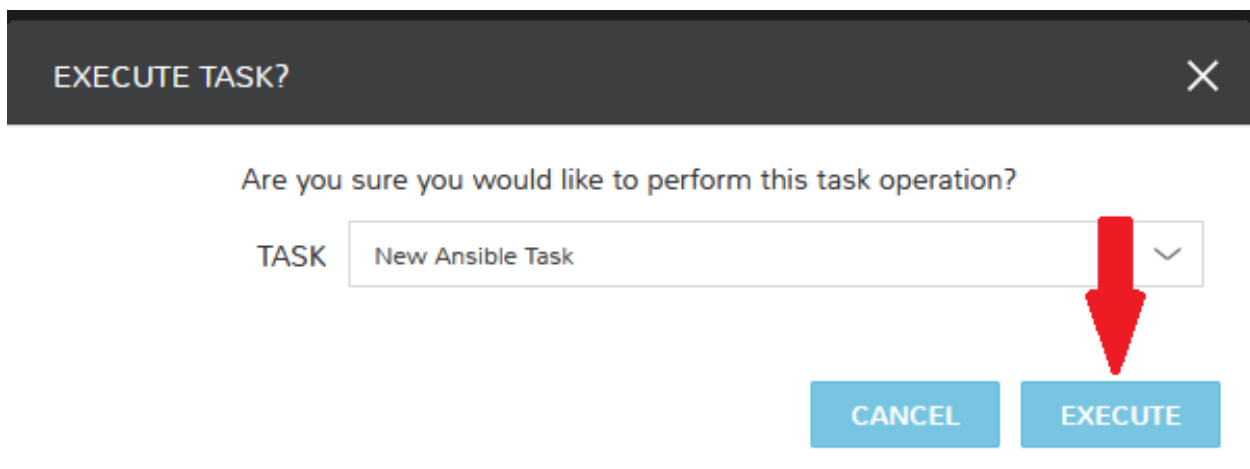
Target Options

EXECUTE TARGET Resource

► Execute Options

SAVE CHANGES

Click “SAVE CHANGES” to save our new task. We can test the new task on our Apache VM now by going to *Provisioning > Instances* and clicking into our VM. From the “ACTIONS” menu select “Run Task”. From the “TASK” dropdown menu, select the task we just added and click “EXECUTE”.



EXECUTE TASK?

Are you sure you would like to perform this task operation?

TASK New Ansible Task

CANCEL EXECUTE

To see the progress of the task, click on the “HISTORY” tab and click on the (i) button to the right of each entry in the list. In this case, we can also see the results of the task by clicking on the link in the “LOCATION” column of the “VMS” section.

Now that our task is created, we can put it into a workflow. Back in *Provisioning > Automation* we will click on the

“WORKFLOWS” tab. Click “+ADD” and select Provisioning Workflow. We’ll give the new workflow a name and expand the Post Provision section. As we begin to type in the name of the task we’ve created, it should appear as a selection. Click “SAVE CHANGES”.

NEW WORKFLOW

NAME

New Workflow

DESCRIPTION

PLATFORM

All

Tasks

► Pre Provision

▼ Provision

Search Tasks

▼ Post Provision

new Ansible Task

► New Ansible Task

► Stop Service

► Pre Deploy

► Deploy

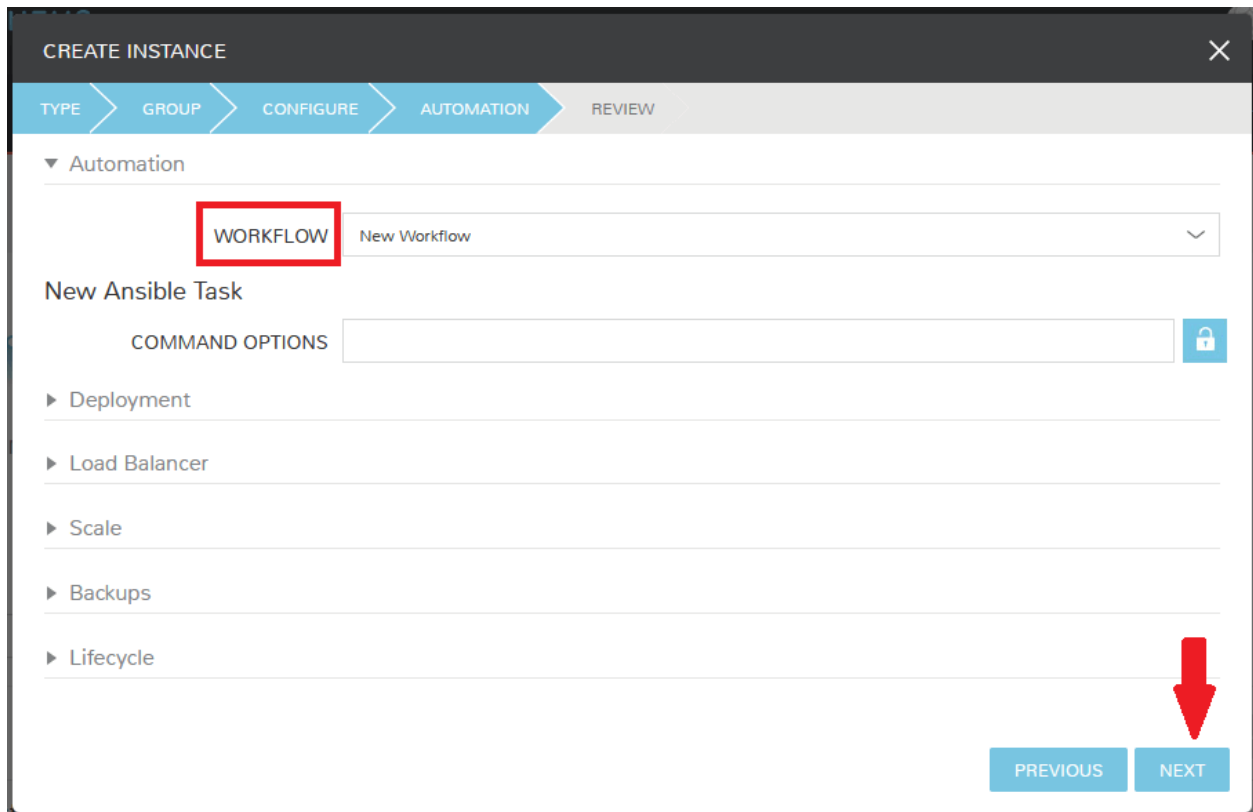
► Reconfigure

► Teardown

SAVE CHANGES

Now that we have a workflow, return to *Provisioning > Instances* and begin to provision another Apache instance.

More detailed instructions on provisioning a new Apache instance are included earlier in this guide if needed. Now, when you reach the “AUTOMATION” section of the “CREATE INSTANCE” wizard, we have a workflow to select. From the “WORKFLOW” dropdown menu, select the workflow we just created and complete provisioning of the new instance.



The screenshot shows the 'CREATE INSTANCE' wizard in the 'AUTOMATION' step. The 'WORKFLOW' dropdown menu is highlighted with a red box. Below it, there is a 'New Ansible Task' section with a 'COMMAND OPTIONS' input field. A list of automation types (Deployment, Load Balancer, Scale, Backups, Lifecycle) is shown with expandable arrows. At the bottom right, a red arrow points to the 'NEXT' button.

As the instance is provisioning, we can go to the “HISTORY” tab and see Morpheus executing the tasks that were contained in our workflow.

This is just one example of using Morpheus to automate the process of configuring and instance to your needs. There are a number of other automation types that can be built into your workflows as well. For further information, take a look at the [automation integrations](#) guide in Morpheus docs.

Conclusion

At this point you should be up and running in Morpheus, ready to consume AWS. This guide only scratches the surface, there is a lot more to see and do in Morpheus. Take a look at the rest of [Morpheus Docs](#) for more information on supported integrations and other things possible.

Getting started with Morpheus and VMware

Introduction

This guide is designed to help you get started and quickly get the most out of Morpheus with VMWare. By the end, you will integrate your first cloud, configure networking, prepare and consume images, provision instances, and get started with automation. We will briefly discuss installation and account setup but will provide links to additional resources for those very first steps. For the most part, this guide assumes you are able to get Morpheus installed and are ready to move forward from that point. There is a lot more to see and do in Morpheus that is beyond the scope of this guide. For more, consult the complete Morpheus documentation or take part in our user community forum.

Installation & Setup

In the simplest configuration, Morpheus needs one appliance server which will contain all the components necessary to orchestrate virtual machines and containers. Full requirements, including storage and networking considerations, can be found in Morpheus documentation [here](#). In order to provision any new instances, hosts, or applications, (or convert any discovered resources to managed resources) you will need a valid license. If you don't have one, you can request a lab license for free at [Morpheus Hub](#). Once obtained, the license can be applied in Administration > Settings > LICENSE.

Groups

Groups in Morpheus define which resources a user has access to. Clouds are added to groups and a user can only access clouds that are in the groups to which their roles give them access. More information on Morpheus groups is [here](#). A deep dive into groups goes beyond the scope of this guide but it's often useful to create a group that contains all clouds for testing purposes. We will create that group now so that we can add our first cloud into this group in the next section.

Navigate to *Infrastructure > Groups*. Here we will see a list of all configured groups but, of course, this will be empty immediately after installation. Click "+CREATE". Give your group a name, such as "All Clouds". The "CODE" field is used when calling Morpheus through Morpheus API or Morpheus CLI. It's useful in most cases to have an "All Clouds" group for testing purposes so this will likely help you down the road.

NEW GROUP

Configuration

NAME

All Clouds

CODE

LOCATION

► Advanced Options

SAVE CHANGES

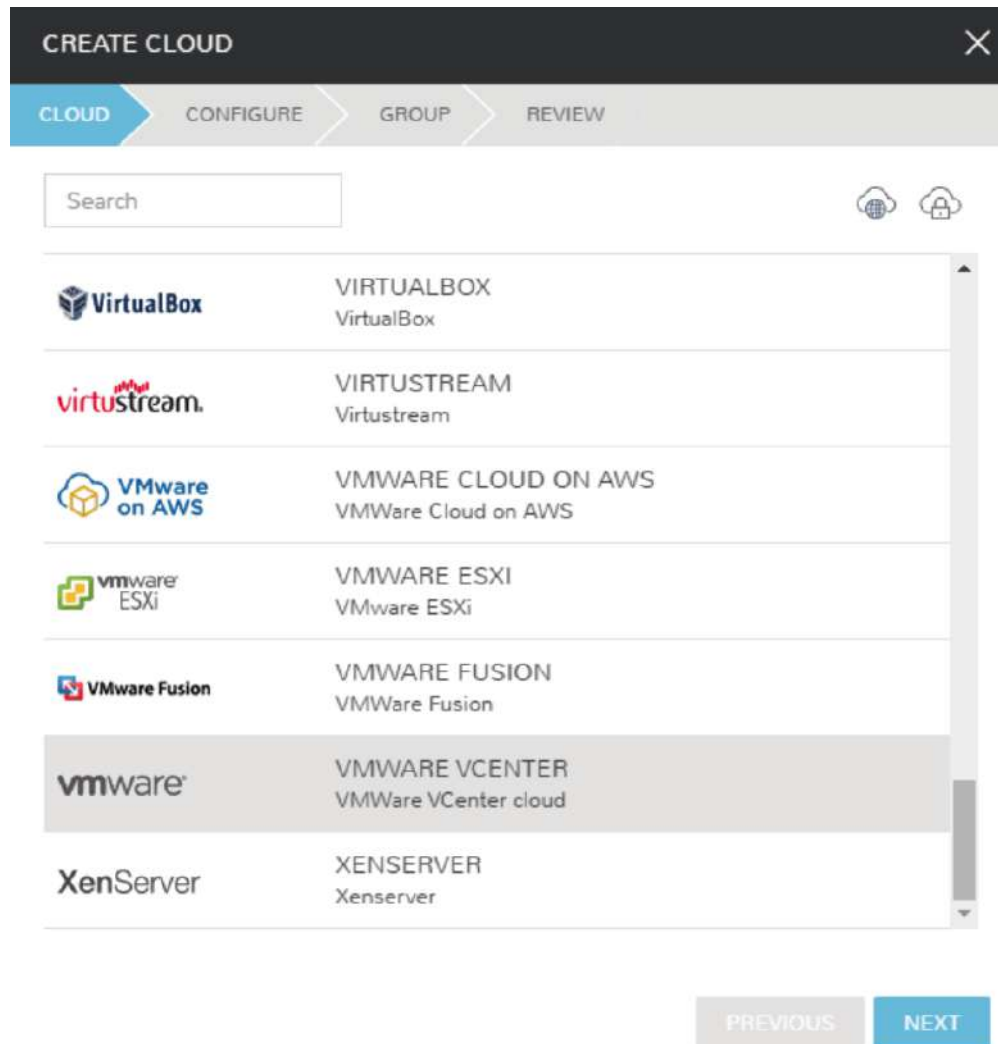
Click “SAVE CHANGES”. Your group is now ready to accept clouds.

Integrating Your First Cloud

Clouds in Morpheus consist of any consumable endpoint whether that be On-Prem, Public clouds, or even bare metal. In this guide, we will focus on integrating and working with VMWare vCenter.

To get started, we will navigate to *Infrastructure > Clouds*. This is the cloud detail page which lists all configured clouds. It will be empty if you’ve just completed installation and setup of Morpheus but soon we will see our integrated vCenter cloud here.

Click the “+ADD” button to pop the “CREATE CLOUD” wizard. Select “VMWARE VCENTER” and click the “NEXT” button.



On the “CONFIGURE” tab, we’re asked to set the initial connection strings into vSphere. The **API URL** should be in the following format: <https://<URL>/sdk>. The **USERNAME** should be in user@domain format.

CREATE CLOUD [X]

CLOUD > **CONFIGURE** > GROUP > REVIEW

NAME: Lab vCenter

CODE: labvc

LOCATION: Denver, CO

Details

API URL: https://192.168.0.110/sdk

USERNAME: administrator@vsphere.local

PASSWORD:

VERSION: 6.7+ ▼

VDC: Home ▼

CLUSTER: All ▼

Morpheus allows vCenter clouds to be scoped to the **VDC** and **CLUSTER** or even the specific **RESOURCE POOL** if you choose. Once you’ve entered your URL and credentials, these dropdown menus will become populated.

The **RPC MODE** setting determines how Morpheus will connect to VMs and make configuration and scripting calls if **Morpheus Agent** is not installed. In a VMware environment we have the additional option to select VMware Tools if WinRM/SSH are not available.

Additionally, we can opt to **INVENTORY EXISTING INSTANCES** to begin polling VMs for statistics and right-sizing recommendations as well as **ENABLE HYPERVISOR CONSOLE** to use native vSphere console with port 443 connectivity between Morpheus and ESXi hosts.

To move on, expand the “Advanced Options” section.

Within the “Advanced Options” drawer are additional configurations to consider for your first cloud. Some of these won’t be usable until they reference additional configured integrations. Common settings to consider are **DOMAIN**, **STORAGE TYPE**, **APPLIANCE URL** (overrides the Morpheus URL for external systems), **GUIDANCE** (setting “Manual” will make recommendations for rightsizing), and **AGENT INSTALL MODE**.

▼ Advanced Options

☐ ENABLE DISK TYPE SELECTION☐ ENABLE NETWORK INTERFACE TYPE SELECTIONSTORAGE TYPE DOMAIN SCALE PRIORITY

Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL TIME ZONE

Once you're satisfied with your selections, click "NEXT"

We have now arrived at the "GROUP" tab. In this case, we will mark the radio button to "USE EXISTING" groups if you wish to use the group we configured earlier.

CREATE CLOUD [X]

CLOUD > CONFIGURE > **GROUP** > REVIEW

☒ USE EXISTING ☐ CREATE NEW

GROUP

PREVIOUS NEXT

Once you've selected the group, click "NEXT"

On the final tab of the "CREATE CLOUD" wizard, you'll confirm your selections and click "COMPLETE". The new cloud is now listed on the cloud detail page. After a short time, Morpheus will provide summary information and statistics on existing virtual machines, networks, and other resources available in the cloud.

Viewing Cloud Inventory

Now that we've integrated our first VMware cloud, we can stop for a moment to review what Morpheus gives us from the cloud detail page. We can see that Morpheus gives us estimated costs and cost histories, metrics on used resources, and also lists out resource counts in various categories including container hosts, hypervisors, and virtual machines. We can drill into these categories to see lists of resources in the various categories individual resources within them by clicking on the category tabs. We can link to the detail page for any specific resource by clicking on it from its resource category list.

Configuring Resource Pools

With our VMware cloud configured, Morpheus will automatically sync in available resource pools and data stores.

For resource pools, once Morpheus has had time to ingest them, then will be visible from the cloud detail page. Navigate to *Infrastructure > Clouds > (your VMware cloud) > RESOURCES tab*. In here, we are able to see and control access to the various resource pools that have been configured in vCenter. For example, we can restrict access to a specific resource pool within Morpheus completely by clicking on the "ACTIONS" button, then clicking "Edit". If we unmark the "ACTIVE" button and then click "SAVE CHANGES" we will see that the resource pool is now grayed out in the list. The resources contained in that pool will not be accessible for provisioning within Morpheus.

POOLS

<input type="text" value="Search"/>				
NAME	DESCRIPTION	VISIBILITY	DEFAULT	TENANT
Demo-vSAN		Private		morpheus
Demo		Private		morpheus
Pied Piper		Private		morpheus
Prod		Private		morpheus
Sandbox		Private		morpheus

Often our clients will want to make specific blocks of resources available to their own customers. This can be easily and conveniently controlled through the same "EDIT RESOURCE POOL" dialog box we were just working in. If we expand the "Group Access" drawer, we are able to give or remove access to each pool to any group we'd like. We can also choose to make some or all of our resource pools available to every group. Specific resource pools can also be defined as the default for each group if needed.

EDIT RESOURCE POOL

×

NAME
Demo-vSAN

☐ MOVE SERVERS

☐ ACTIVE

☐ DEFAULT

▼ Group Access

GROUP	ACCESS	DEFAULT
all	<input checked="" type="checkbox"/>	
AA	<input type="checkbox"/>	<input type="checkbox"/>
AA-DB	<input type="checkbox"/>	<input type="checkbox"/>

Additionally, we may choose to allow only certain service plans to be provisioned into a specific pool of resources. For example, perhaps a specific cluster is my SQL cluster and only specific services plans should be consumable within it. We can control that through this same dialog box.

Configuring Data Stores

To take a look at data stores, we'll move from the "RESOURCES" tab to the "DATA STORES" tab on our cloud detail page.

Morpheus gives the user similar control with data stores to what we saw with our resources pools earlier. Just like with resource pools, we can disable access within Morpheus completely by clicking on "ACTIONS" and then "Edit". If we unmark the "ACTIVE" checkbox and click "SAVE CHANGES", you will see that specific data store has been grayed out.

SUMMARY	CLUSTERS	HOSTS	VMS	CONTAINERS	LOAD BALANCERS	NETWORKS	DATA STORES	RESOURCES	POLICIES	WIKI
<div> <div>Search</div> <div>Q</div> <div>ACTIONS +</div> </div>										
<input type="checkbox"/>	NAME	TYPE CAPACITY ONLINE VISIBILITY TENANT								
<input type="checkbox"/>	ds-65-root	Vmfs	461.1GiB	Yes	Private	morpheus	ACTIONS ▼			
<div> <div>morpheuscamusara</div> <div>morpheus Morpheus Shalimar EMC IT SANI AM troublemaker Camusel Anish rtalder KDE Lineromha</div> </div>										

Just like with resource pools, we are also able to scope data stores to specific groups. This ensures that the members of each group are only able to consume the data stores they should have access to.

EDIT DATA STORE

×

NAME

ds-65-root

ACTIVE
☐

▼ Group Access

GROUP	ACCESS
all	<input checked="" type="checkbox"/>
AA	<input type="checkbox"/>
AA-DB	<input type="checkbox"/>

Configuring Network for Provisioning

When configuring networking, we can set global defaults by going to *Infrastructure > Network > NETWORKS* tab. Here we can add or configure networks from all clouds integrated into Morpheus. Depending on the number of clouds Morpheus has ingested, this list may be quite large and may also be paginated across a large number of pages. In such a case, it may be easier to view or configure networks from the specific cloud detail page so that networks from other clouds are not shown.

NETWORKS

NETWORKS

NETWORK GROUPS

ROUTERS

IP POOLS

DOMAINS

PROXIES

SECURITY GROUPS

INTEGRATIONS

Search

+

ADD

NAME	TYPE	CLOUD	CIDR	POOL	DHCP	VISIBILITY	TENANTS	
10.30.20.0	OracleVM Network	Morpheus Oracle VM	10.30.20.0/22	✓	Public	morpheus		ACTIONS ▼
172.31.0.0/20 (subnet-63dff13b)	Amazon Subnet	ah-only	172.31.0.0/20	✓	Private	morpheus		ACTIONS ▼
172.31.16.0/20 (subnet-22110ed6)	Amazon Subnet	ah-only	172.31.16.0/20	✓	Private	morpheus		ACTIONS ▼

Still in *Infrastructure > Network*, make note of the “INTEGRATIONS” tab. It’s here that we can set up any integrations that may be relevant, such as IPAM integrations. Generally speaking, when adding IPAM integrations, we simply need to name our new integration, give the API URL, and provide credentials. There’s more information in the [IPAM integration](#) section of Morpheus Docs.

ADD IPAM INTEGRATION

×

NAME

☒ ENABLED

URL

USERNAME

PASSWORD

THROTTLE RATE

ms

In *Infrastructure > Networking* we can also set up IP address pools from the IP Pools tab. These pools can be manually defined, known as a Morpheus-type IP pool, or they can come from any IPAM integrations you’ve configured. As instances are provisioned, Morpheus will assign IP addresses from the pool chosen during provisioning. When the instance is later dissolved, Morpheus will automatically release the IP address to be used by another instance when needed. When adding or editing a network, we can opt to scope the network to one of these configured IP address pools.

CREATE NETWORK POOL

×

NAME

POOL TYPE

Morpheus

▼

IP Ranges

STARTING ADDRESS	ENDING ADDRESS
<input type="text" value="192.168.0.2"/>	<input type="text" value="192.168.0.255"/>








-

+

SAVE CHANGES

Since this guide is focused on working within a VMware cloud that we integrated at the start, we will take a look at our network configurations on the cloud detail page as well. Navigate to *Infrastructure > Clouds > (your VMware cloud) > NETWORKS tab*. Just as with resource pools and data stores, we have the ability to make certain networks inactive in Morpheus, or scope them to be usable only for certain groups or tenants.

SUMMARY	CLUSTERS	HOSTS	VMS	CONTAINERS	LOAD BALANCERS	NETWORKS	DATA STORES	RESOURCES	POLICIES	WIKI
---------	----------	-------	-----	------------	----------------	----------	-------------	-----------	----------	------

NETWORKS						
Search 			ACTIONS 			
<input type="checkbox"/> NAME	TYPE	CIDR	POOL	DHCP VISIBILITY	TENANT	
<input type="checkbox"/> default	KVM Host Bridge	192.168.122.1/24		 Private	morpheus	ACTIONS 
<input type="checkbox"/> docker bridge	Docker Bridge			 Private	morpheus	ACTIONS 
<input type="checkbox"/> Garf-Network	Overlay	0.0.0.0/10	Vanishing-IP-Test	Private	morpheus	ACTIONS 

Prepping an Image

As we'll discuss and try out in the next section, Morpheus comes out of the box with a default set of blueprints that are relevant to many modern deployment scenarios. For the most part, these are base operating system images with a few additional adjustments. However, in many on-premise deployments, there are often custom image and networking requirements. We will work with images included in Morpheus by default in this guide but it's important to discuss how to prep custom images as well.

Creating a Windows Image

The following versions of Windows Server are supported:

- 2008 R2
- 2012
- 2012 R2
- 2016
- 2019

To start, create a new Windows machine in vCenter using a base version of your selected Windows build.

Note: It's recommended to make the VMDK drive as small as possible for your purposes as this generally speeds cloning and deploy times. Morpheus provisioning and post-deploy scripts allow to to expand the drive to any size that you need.

Once the machine is created, ensure VMtools is installed on the operating system. Then, apply all updates and service packs. Next, configure WinRM and open the firewall:

```
winrm quickconfig
```

Note: WinRM configuration is optional if using VMtools RPC mode for agent install and Morpheus Agent for guest exec.

Next, we'll install .NET 4.5 or higher. Ensure Windows Firewall will allow WinRM connections and shut down the virtual instance. Finally, convert it to a template.

Note: Morpheus will Sysprep images based on the "Force Guest Customizations" flag under VM settings when using DHCP. If this flag is enabled or if using static IP addresses or IP pools when provisioning, ensure a Sysprep has not been performed. In such cases, guest customization will always be performed and a Sysprep will be triggered.

Creating a CentOS/RHEL Image

Create a new machine in vCenter and install a base version of your preferred Linux distro.

Note: If you are using cloud-init as part of your image, you will need to ensure your virtual machine has a cdrom.

Before installing the operating system, set up a single ext or xfs partition without a swap disk. Next, install the distro applying any updates to the operating system or security updates. Once the operating system is running and updated, install the following:

```
yum install cloud-init
yum install cloud-utils-growpart
yum install open-vm-tools
yum install git
yum install epel-release
```

Set selinux to permissive as the enforced setting can cause problems with cloud-init:

```
sudo vi /etc/selinux/config
```

Cloud-Init

We'll get started by installing cloud-init using the following command:

```
yum -y install epel-release
yum -y install git wget ntp curl cloud-init dracut-modules-growroot
rpm -qa kernel | sed 's/^kernel-//' | xargs -I {} dracut -f /boot/initramfs-{}.img {}
```

Note: The above command will install some core dependencies for cloud-init and automation later as you work with your provisioned instances. For example, we install Git here as it is used for Ansible automation. If you had no plans to use Ansible, this installation could be skipped. The dracut-modules-growroot is responsible for resizing the root partition upon initial boot which was potentially adjusted during provisioning.

One key benefit of using cloud-init is that we don't have to lock credentials into the blueprint. We recommend configuring a default cloud-init user that will get created automatically when the VM is booted by cloud-init. We can define that default user in *Administration > Provisioning > Cloud-Init*.

Network Interfaces

As of CentOS 7, network interface naming conventions have changed. You can check this by running *ifconfig* and noting that the primary network interface has some value similar to "ens2344". The naming is dynamic and typically set based on hardware ID. We don't want this to fluctuate when provisioning this blueprint in our VMware environments. To accomplish this end, we will rename the interface back to "eth0" using the steps below.

First, adjust the bootloader to disable interface naming:

```
sed -i -e 's/quiet/quiet net.ifnames=0 biosdevname=0/' /etc/default/grub
grub2-mkconfig -o /boot/grub2/grub.cfg
```

The next step is to adjust network scripts in CentOS. Start by confirming the presence of a file called */etc/sysconfig/network-scripts/ifcfg-eth0*. Once confirmed, run the following script:

```
export iface_file=$(basename "$(find /etc/sysconfig/network-scripts/ -name 'ifcfg*' -
↪not -name 'ifcfg-lo' | head -n 1)")
export iface_name=${iface_file:6}
```

(continues on next page)

(continued from previous page)

```
echo $iface_file
echo $iface_name
sudo mv /etc/sysconfig/network-scripts/$iface_file /etc/sysconfig/network-scripts/
↪ifcfg-eth0
sudo sed -i -e "s/$iface_name/eth0/" /etc/sysconfig/network-scripts/ifcfg-eth0
sudo bash -c 'echo NM_CONTROLLED="no" >> /etc/sysconfig/network-scripts/ifcfg-eth0'
```

This script tries to confirm there is a new *ifcfg-eth0* config created to replace the old config file. Confirm this config exists after running and if not you will have to build your own:

```
TYPE=Ethernet
DEVICE=eth0
NAME=eth0
ONBOOT=yes
NM_CONTROLLED="no"
BOOTPROTO="dhcp"
DEFROUTE=yes
```

For more on CentOS/RHEL image prep, including additional configurations for specific scenarios, take a look at the [VMware image prep](#) page in Morpheus Docs.

Creating an Ubuntu Image

Create a new machine in vCenter and install a base version of your preferred Linux distro.

Note: If you are using cloud-init as part of your image, you will need to ensure your virtual machine has a cdrom.

Before installing the operating system, set up a single ext partition without a swap disk. Install the distro and apply any operating system and security updates. Ensure you’ve set a root password.

Install cloud-init and cloud-utils-growpart:

```
sudo apt install cloud-init
sudo apt install cloud-utils
```

Install desired hypervisor drivers, such as Virto or Open-VM Tools

Install Git:

```
sudo apt install git
```

Since Debian 9 includes network manager, ensure this is disabled. You can do this by editing the configuration file at */etc/NetworkManager/NetworkManager.conf*. Within that file, update the “managed” flag to false:


```
managed=false
```

We also recommend setting the network adapter to “eth0”. This process is described above in the “Network Interfaces” section of the CentOS image prep guide above.

Provisioning Your First Instance

At this point, we are ready to provision our first image. As a first instance, we'll provision an Apache web server to our vCenter cloud.

Navigate to *Provisioning > Instances*. If any instances are currently provisioned, we will see them listed here. To start a new instance we click the “+ADD” button to pop the “CREATE INSTANCE” wizard. We'll scroll down to and select the Apache instance type and click “NEXT”.



First, we'll specify the group to provision into which determines the clouds available. If you've followed this guide to this point, you should at least have a group that houses all of your clouds which you can select here. This will allow us to select the vCenter cloud from the “CLOUD” dropdown menu. Provide a unique name to this instance and then click “NEXT”

From the “CONFIGURE” tab, we're presented with a number of options. The options are cloud and layout-specific, more generalized information on creating instances and available options is [here](#). For our purposes, we'll select the following options:

- **LAYOUT:** Includes options such as the base OS, custom layouts will also be here when available
- **PLAN:** Select the resource plan for your instance. Some plans have minimum resource limits, Morpheus will only show plans at or above these limits. User-defined plans can also be created in *Administration > Plans & Pricing*.
- **VOLUMES and DATASTORES:** The minimum disk space is set by the plan, this value may be locked if you've selected a custom plan that defines the volume size
- **NETWORKS:** Select a network, note that IP pools must be linked with the networks defined in VMware in order to assign static IP addresses

Under the “User Config” drawer, mark the box to “CREATE YOUR USER”. Click “NEXT”.

CREATE INSTANCE [X]

TYPE > GROUP > **CONFIGURE** > AUTOMATION > REVIEW

Configuration Options

LAYOUT ESXi Apache on Ubuntu 14.04

PLAN 1 Core, 512MB Memory
Cores: 1 Memory: 512 MB

VOLUMES root 10 GB Auto - Datastore +

NETWORKS VM Network DHCP +

▼ User Config

☒ **CREATE YOUR USER**

USER GROUP Select

► Network Options

► Advanced Options

► Metadata

► Environment

PREVIOUS NEXT

Note: “CREATE YOUR USER” will seed a user account into the VM with credentials set in your Morpheus user account settings. If you’ve not yet defined these credentials, you can do so by clicking on your username in the upper-right corner of the application window and selecting “USER SETTINGS”.

For now, we’ll simply click “NEXT” to move through the “AUTOMATION” tab but feel free to stop and take a look at the available selections here. There is more information later in this guide on automation and even more beyond that in the rest of Morpheus docs.

Review the settings for your first instance and click “COMPLETE”.

CREATE INSTANCE

TYPE

GROUP

CONFIGURE

AUTOMATION

REVIEW

apachetest

admin (esxtank)

Summary

INSTANCE OPTIONS

NAME: apachetest

GROUP: admin

CLOUD: esxtank

TYPE: Apache

PLAN: 1 Core, 512MB Memory

Cores: 1 Memory: 512MB

VERSION: 2.4

LAYOUT: ESXi Apache on Ubuntu 14.04

VOLUMES

ROOT: 10 GB Auto - Datastore SCSI 0

NETWORKS

VM NETWORK: DHCP

Options

PREVIOUS

COMPLETE

We are now dropped back onto the instances list page. We can see a new entry in the list at this point with a status indicator that the new machine is being launched (rocket icon in the status field). We can double click on the instance in the list to move to the instance detail page. For now we will see a progress bar indicating that the instance is being created and is starting up. The exact amount of time this process will take depends on your environment and selections made when provisioning the instance. Initially, Morpheus will guess as to how long this will take and the progress bar may not be accurate. Over time, Morpheus will learn how long these processes take and progress bar accuracy will improve. For more detailed information on the status of various provisioning processes, we can scroll down and select the “HISTORY” tab. The “STATUS” icon will change from the blue rocket to a green play button when the instance is fully ready. Furthermore, we can click on the hyperlinked IP address in the “VMS” section of this page to view a default page in a web browser to confirm success.

INFO

Group: admin

Created By: Nick Celebic

Cores: 1

Source Image: Morpheus Apache 2.4 on Ubuntu 14.04.3

Cloud: esxtank

Layout: ESXi Apache on Ubuntu 14.04

Memory: 512.0MiB

Date Created: 11/13/2019 01:44 PM

Version: 2.4

Total Storage: 10.0GiB

VMS

	STATUS	NAME	TYPE	CLOUD	LOCATION	COMPUTE	MEMORY	STORAGE	ACTIONS
<input type="checkbox"/>		apachetest	Apache 2.4	esxtank		<div>0</div>	<div>0</div>	<div>0</div>	

SUMMARY

WIKI

DEPLOY

STORAGE

NETWORK

LOGS

BACKUPS

ENVIRONMENT

SCALE

HISTORY


	NAME	DESCRIPTION	CREATED BY	START DATE	ETA/DURATION	STATUS	ERROR
<input type="checkbox"/>	apachetest	Provision	Nick Celebic	11/13/2019 01:44 PM	00:18:44	<div></div>	
		Prepare Resources		11/13/2019 01:44 PM	00:18:44	Complete	
		Prepare Image		11/13/2019 01:44 PM	00:00:01	Complete	
		Configure Instance		11/13/2019 01:44 PM	260ms	Complete	
		Deploy Instance		11/13/2019 01:44 PM	00:18:44	<div></div>	

Creating Your First Library Item

In the prior section, we manually provisioned our first instance. However, Morpheus allows you to build a catalog of custom provisionable items to simplify and speed provisioning in the future. In this section, we'll build a catalog item and show how that can translate into quick instance provisioning after configuration.

Note: Before starting this process, it's important to decide which virtual image you plan to use. If you're not using a Morpheus-provided image, you'll want to ensure it's uploaded. You will not be able to complete this section without selecting an available image. In this example we will use Morpheus Redis 3.0 on Ubuntu 14.04.3 v2.

Navigate to *Provisioning > Library > NODE TYPES* and click “+ADD”.



In this example, I am going to set the following options in the “NEW NODE TYPE” wizard:

- **NAME**
- **SHORT NAME**
- **VERSION:** 1 (In this particular case, the version is not important)
- **TECHNOLOGY:** VMware
- **VM IMAGE:** Morpheus Redis 3.0 on Ubuntu 14.04.3 v2

Note: Within the “VMware VM Options” section you should add anything that will always be used for this node, regardless of the specific deployment details. This can include LDAP Authentication, bash scripts that should run on installation, among other things.

NAME

customnodetype

SHORT NAME

customnodetype

The short name is a name with no spaces used for display in your container list.

VERSION

1

TECHNOLOGY

VMware

ENVIRONMENT VARIABLES

Name

Value

⚙️

+

VMware VM Options

VM IMAGE

Morpheus Redis 3.0 on Ubuntu 14.04.3 v2

LOG FOLDER

CONFIG FOLDER

DEPLOY FOLDER

(Optional) If using deployment services, this mount point will be replaced with the contents of said deployments.

EXTRA OPTIONS

Name

Value

🗑️

+

SERVICE PORTS

port

⬆️⬆️

name

No LB

⌵

+

SCRIPTS

Search


FILE TEMPLATES

Search

▶ Advanced Options

SAVE CHANGES

With the new node created, we'll now add a new instance type which will be accessible from the provisioning wizard once created. Move from the "NODE TYPES" tab to the "INSTANCE TYPES" tab and click "+ADD".



In the “NEW INSTANCE TYPE” wizard, I’ll simply enter a **NAME** and **CODE** value. Click “SAVE CHANGES”.

NEW INSTANCE TYPE [X]

NAME

CODE

Useful shortcode for provisioning naming schemes and export reference.

DESCRIPTION

255 Characters Remaining

CATEGORY

ICON

Suggested Dimensions: 150 x 51

Now that we’ve created a new instance type, access it by clicking on the name in the list of custom instances you’ve created. In my case, I’ve given the name “NewInstanceType”.

LIBRARY				
INSTANCE TYPES				
LAYOUTS				
NODE TYPES				
OPTION TYPES				
OPTION LISTS				
FILE TEMPLATES				
SCRIPTS				
SPEC TEMPLATES				
CLUSTER LAYOUTS				
Search				
SELECT TECHNOLOGY				
SELECT CATEGORY				
+ ADD				
	NAME	TECHNOLOGY	CATEGORY	FEATURED
	NewInstanceType		Web	ACTIONS
	ActiveMQ	Mixed	Messaging	ACTIONS
	Alibaba	Alibaba	Cloud	ACTIONS
	Amazon Api		Apps	ACTIONS
	AmazonMQ		Messaging	ACTIONS

Once we’ve opened the new instance type, by default, we should be on the “LAYOUTS” tab. Click “+ADD LAYOUT”.

I’ve set the following fields on my example layout:

- **NAME**
- **VERSION:** This is the version number of the layout itself, which is labeled 1.0 in the example
- **TECHNOLOGY:** VMware

- **Nodes:** Select the node we created earlier, if desired you can specify multiple nodes

Click “SAVE CHANGES”.

NEW LAYOUT

NAME

customlayout

VERSION

1.0

DESCRIPTION

☒ CREATABLE

TECHNOLOGY

VMware

MINIMUM MEMORY

0

MB

This will override any memory requirement set on the virtual image

WORKFLOW

☐ SUPPORTS CONVERT TO MANAGED

ENVIRONMENT VARIABLES

Name

Value

+

Option Types

Search option types

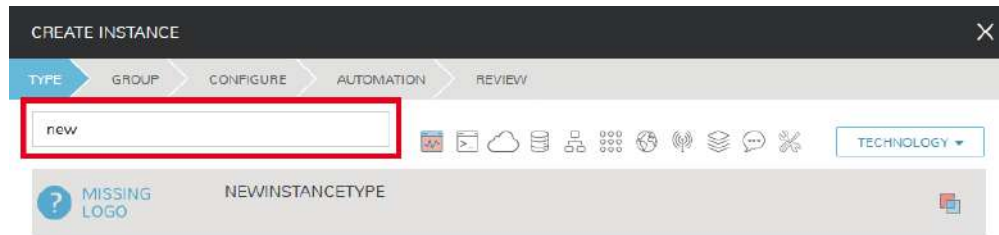
Nodes

customnodetypevm

customnodetypevm (1)

SAVE CHANGES

At this point we’ve completed the setup work and can now provision the instance we’ve created to our specifications. Navigate to *Provisioning > Instances* and click “+ADD”. From the search bar we can search for the new instance type we’ve created. In the example case, we called it “newinstancetype”. Click “NEXT”.



CREATE INSTANCE

TYPE GROUP CONFIGURE AUTOMATION REVIEW

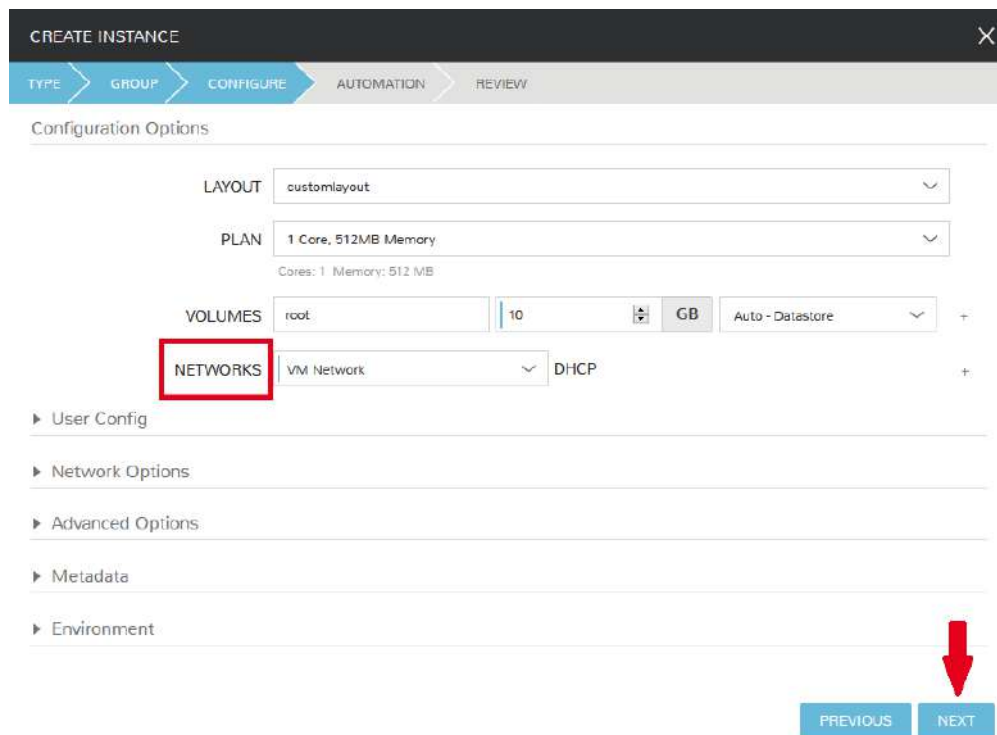
new

MISSING LOGO NEWINSTANCETYPE

TECHNOLOGY



As before, we can select a group and cloud to provision this new instance. Click “NEXT”. On the “CONFIGURE” tab, make note that the layout and plan are already selected because they were configured as part of creating the new instance type. Select a network and click “NEXT”. Once again we will also click “NEXT” through the “AUTOMATION” tab. Finally, click “COMPLETE”.



CREATE INSTANCE

TYPE GROUP CONFIGURE AUTOMATION REVIEW

Configuration Options

LAYOUT customlayout

PLAN 1 Core, 512MB Memory

Cores: 1 Memory: 512 MB

VOLUMES root 10 GB Auto - Datastore

NETWORKS VM Network DHCP

User Config

Network Options

Advanced Options

Metadata

Environment

PREVIOUS NEXT

As before when we manually provisioned an instance, Morpheus will now begin to spin up the new VM. How long this will take depends on your environment but Morpheus will predict how long this process will take and represent that on a progress bar. Over time, Morpheus begins to learn how long these processes take and becomes more accurate

in predicting spin-up time.

Once the provisioning process has completed, open the instance detail page in Morpheus and click on the “CONSOLE” tab. You’ll be logged in with your user account and are then able to confirm the machine is ready and available.

The screenshot shows the Morpheus instance detail page for an instance named 'newinstance'. At the top, there are buttons for 'EDIT', 'ACTIONS', and 'DELETE'. Below this, a row of metrics shows: STATUS (green play button), HEALTH (blue question mark), LAST BACKUP (grey minus), AVAILABILITY (100.00%), RESPONSE TIME (N/A), MAX CPU (1%), MEMORY (39%), and STORAGE (27%).

The 'INFO' section displays:

- Group:** admin
- Created By:** Nick Celebic
- Cloud:** esxtank
- Date Created:** 11/14/2019 12:21 PM
- Cores:** 1
- Layout:** customlayout
- Version:** 1.0
- Source Image:** Morpheus Redis 3.0 on Ubuntu 14.04.3 v2
- Memory:** 512.0MB
- Provision Time:** 14 minutes 29 seconds
- Total Storage:** 10.0GiB

The 'VMS' section contains a table with columns: STATUS, NAME, TYPE, CLOUD, LOCATION, COMPUTE, MEMORY, STORAGE, and ACTIONS. The table has one row for 'newinstance' with a green play button icon in the STATUS column. A red arrow points to the 'CONSOLE' tab below the table.

The 'CONSOLE' tab shows a terminal window with the following text:


```
(Connected)
Welcome to Ubuntu 14.04.3 LTS (GNU/Linux 3.19.0-80-generic x86_64)
* Documentation:  https://help.ubuntu.com/
Last login: Thu Nov 14 17:55:47 2019 from 192.168.88.61
ncelebic@newinstance:~$
```

Automation and Configuration Management

Morpheus automation is composed of Tasks and Workflows. A task could be a script added directly, scripts or blueprints pulled from the Morpheus Library, playbooks, recipes, or a number of other things. The complete list of task types can be found in the [Automation section](#) of Morpheus docs. Tasks can be executed individually but they are often combined into workflows. We can opt to run a workflow at provision time or they can be executed on existing instances through the Actions menu.

In this guide we will set up an Ansible integration, create a task, add the task to a workflow, and run the workflow against a new and existing instance. If you’ve worked through this guide to this point, you should already have an Apache instance running. If you don’t yet have that, provision one before continuing with this guide and ensure it’s reachable on port 80.

The screenshot shows the Morpheus Administration page. The top navigation bar includes: Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration (highlighted). Below this, a row of icons represents various management areas: Tenants, Plans & Pricing, Roles, Users, Integrations (highlighted), Policies, Provisioning, Monitoring, Backups, Logs, and Settings.

The 'INTEGRATIONS' section is displayed, showing a table with a 'NAME' column. A red arrow points to the '+ NEW INTEGRATION' button. A dropdown menu is open, showing a list of integration options: automation, Chef, Puppet, Ansible (highlighted with a red arrow), Ansible Tower, vRealize Orchestrator, Salt Master, dns, and Microsoft DNS.

We'll first set up the Ansible integration, you can integrate with the sample repository referenced here or integrate with your own. Go to 'Administration > Integrations'. Click "+NEW INTEGRATION" and select Ansible from the dropdown menu. Fill in the following details:

- **NAME**
- **ANSIBLE GIT URL:** <https://github.com/ncelebic/morpheus-ansible-example>, or enter the URL for your own Ansible git repository
- **PLAYBOOKS PATH**
- **ROLES PATH**
- Mark the box to "USE MORPHEUS AGENT COMMAND BUS"

Note: If your git repository requires authentication, you should create a keypair and use the following URL format: [git@github.com:ncelebic/morpheus-ansible-example.git](https://github.com:ncelebic/morpheus-ansible-example.git).

The screenshot shows a 'NEW ANSIBLE INTEGRATION' form. The following fields and options are highlighted with red boxes:

- NAME:** NewAnsibleIntegration
- ENABLED:** ☒ ENABLED
- ANSIBLE GIT URL:** <https://github.com/ncelebic/morpheus-ansible-example>
- KEY PAIR:** (Dropdown menu)
- PLAYBOOKS PATH:** /
- ROLES PATH:** /roles
- GROUP VARIABLES PATH:** (Empty text field)
- HOST VARIABLES PATH:** (Empty text field)
- USE ANSIBLE GALAXY:** ☐
- ENABLE VERBOSE LOGGING:** ☐
- USE MORPHEUS AGENT COMMAND BUS:** ☒

A large red arrow points down to the **SAVE CHANGES** button at the bottom right of the form.

Click "SAVE CHANGES". You'll now see our new Ansible integration listed among any other configured inetegra-

tions. If we click on this new integration to view detail, a green checkmark icon indicates the git repository has been fully synced.

With the Ansible integration set up, we can now create a task that includes our playbook. Go to *Provisioning > Automation*, click “+ADD”. We’ll first set our “TYPE” value to Ansible Playbook so that the correct set of fields appear in the “NEW TASK” wizard. Set the following options:

- **NAME**
- **ANSIBLE REPO:** Here we will choose the Ansible integration that we just created
- **PLAYBOOK:** In our example case, enter ‘playbook.yml’

The screenshot shows the 'NEW TASK' form with the following values:

- NAME: New Ansible Task
- TYPE: Ansible Playbook
- ANSIBLE REPO: NewAnsibleIntegration
- PLAYBOOK: playbook.yml
- EXECUTE TARGET: Resource

A red arrow points to the 'SAVE CHANGES' button.

Click “SAVE CHANGES” to save our new task. We can test the new task on our Apache VM now by going to *Provisioning > Instances* and clicking into our VM. From the “ACTIONS” menu select “Run Task”. From the “TASK” dropdown menu, select the task we just added and click “EXECUTE”.

The screenshot shows the 'EXECUTE TASK?' dialog box with the following details:

- Question: Are you sure you would like to perform this task operation?
- TASK dropdown: New Ansible Task
- Buttons: CANCEL, EXECUTE

A red arrow points to the 'EXECUTE' button.

To see the progress of the task, click on the “HISTORY” tab and click on the (i) button to the right of each entry in the list. In this case, we can also see the results of the task by clicking on the link in the “LOCATION” column of the “VMS” section.

Now that our task is created, we can put it into a workflow. Back in *Provisioning > Automation* we will click on the “WORKFLOWS” tab. Click “+ADD” and select Provisioning Workflow. We’ll give the new workflow a name and expand the Post Provision section. As we begin to type in the name of the task we’ve created, it should appear as a selection. Click “SAVE CHANGES”.

NEW WORKFLOW ✕

NAME

DESCRIPTION

PLATFORM ▼

Tasks

- ▶ Pre Provision
- ▼ Provision
-
- ▼ Post Provision
-
- ▶ **New Ansible Task**
- ▶ Stop Service
- ▶ Pre Deploy
- ▶ Deploy
- ▶ Reconfigure
- ▶ Teardown

SAVE CHANGES

Now that we have a workflow, return to *Provisioning > Instances* and begin to provision another Apache instance. More detailed instructions on provisioning a new Apache instance are included earlier in this guide if needed. Now, when you reach the “AUTOMATION” section of the “CREATE INSTANCE” wizard, we have a workflow to select. From the “WORKFLOW” dropdown menu, select the workflow we just created and complete provisioning of the new instance.

The screenshot shows the 'CREATE INSTANCE' dialog with the 'AUTOMATION' tab selected. The 'WORKFLOW' dropdown is highlighted with a red box. Below it, there is a 'New Ansible Task' section with a 'COMMAND OPTIONS' input field and a lock icon. A list of automation types is shown: Deployment, Load Balancer, Scale, Backups, and Lifecycle. At the bottom right, there are 'PREVIOUS' and 'NEXT' buttons, with a red arrow pointing to the 'NEXT' button.

As the instance is provisioning, we can go to the “HISTORY” tab and see Morpheus executing the tasks that were contained in our workflow.

This is just one example of using Morpheus to automate the process of configuring and instance to your needs. There are a number of other automation types that can be built into your workflows as well. For further information, take a look at the [automation integrations](#) guide in Morpheus docs.

Conclusion

At this point you should be up and running in Morpheus, ready to consume VMware. This guide only scratches the surface, there is a lot more to see and do in Morpheus. Take a look at the rest of [Morpheus Docs](#) for more information on supported integrations and other things possible.

Getting started with Morpheus and Azure

Introduction

This guide is designed to help you get started and quickly get the most out of Morpheus with Microsoft Azure public cloud. By the end, you will integrate your first cloud with Morpheus, configure networking, prepare and consume images, provision instances, and get started with automation. We will briefly discuss installation and account setup but will provide links to additional resources for those very first steps. For the most part, this guide assumes you are able to get Morpheus installed and are ready to move forward from that point. There is a lot more to see and do in Morpheus that is beyond the scope of this guide. For more, consult the complete Morpheus documentation or take part in our [Reddit user community forum](#).

Installation & Setup

In the simplest configuration, Morpheus needs one appliance server which will contain all the components necessary to orchestrate virtual machines and containers. Full requirements, including storage and networking considerations, can be found in Morpheus documentation [here](#). In order to provision any new Instances, hosts, or applications (or convert any discovered resources to managed resources) you will need a valid license. If you don't have one, you can request a community edition license for free at [Morpheus Hub](#). Once obtained, the license can be applied in Administration > Settings > LICENSE. For more, take a look at our community edition [welcome package](#).

Groups

Groups in Morpheus define which resources a user has access to. Clouds are added to Groups and a user can only access Clouds that are in the Groups to which their roles give them access. More information on Morpheus Groups is [here](#). A deep dive into Groups goes beyond the scope of this guide but it's often useful to create a Group that contains all Clouds for testing purposes. We will create that group now so that we can add our first Cloud into this Group in the next section.

Navigate to *Infrastructure > Groups*. Here we will see a list of all configured groups but, of course, this will be empty immediately after installation. Click "+CREATE". Give your group a name, such as "All Clouds". The "CODE" field is used when calling Morpheus through Morpheus API or Morpheus CLI. It's useful in most cases to have an "All Clouds" group for testing purposes so this will likely help you down the road.

Click *SAVE CHANGES*. Your Group is now ready to accept Clouds.

NEW GROUP

×

Configuration

NAME

All Clouds

CODE

LOCATION

► Advanced Options

SAVE CHANGES

Integrating Your First Cloud

Clouds in Morpheus consist of any consumable endpoint whether that be on-prem, public clouds, or even bare metal. In this guide, we will focus on integrating and working with Microsoft Azure public cloud.

To get started, we will navigate to *Infrastructure > Clouds*. This is the Cloud list page which lists all configured Clouds. It will be empty if you've just completed installation and setup of Morpheus but soon we will see our integrated Azure cloud here.

Click the “+ADD” button to pop the “CREATE CLOUD” wizard. Select “AZURE (PUBLIC)” and click the “NEXT” button.

On the “CONFIGURE” tab, we're asked to provide Azure-specific details to connect to the cloud. Morpheus Azure integration requires Owner or Contributor access to subscription via App Registration. Adding an Azure Cloud or Clouds to Morpheus will require the following:

- Azure Subscription ID
- Directory (tenant) ID
- Application (client) ID
- Application (client) Secret
- Application (client) must be Owner or Contributor of Subscription

CSP Accounts require the additional following input:

- CSP Directory (tenant) ID
- CSP Application (client) ID
- CSP Application (client) SECRET

CREATE CLOUD

×

CLOUD

CONFIGURE

GROUP

REVIEW

NAME

CODE

LOCATION

VISIBILITY

Private

▼

TENANT

morpheus

▼

☒ ENABLED

☒ AUTOMATICALLY POWER ON VMS

Details

CLOUD TYPE

Global

▼

SUBSCRIPTION ID

TENANT ID

CLIENT ID

CLIENT SECRET

LOCATION

▼

RESOURCE GROUP

No Locations found: verify credentials above.

▼

☐ INVENTORY EXISTING INSTANCES

INVENTORY LEVEL

Basic

▼

ACCOUNT TYPE

Standard

▼

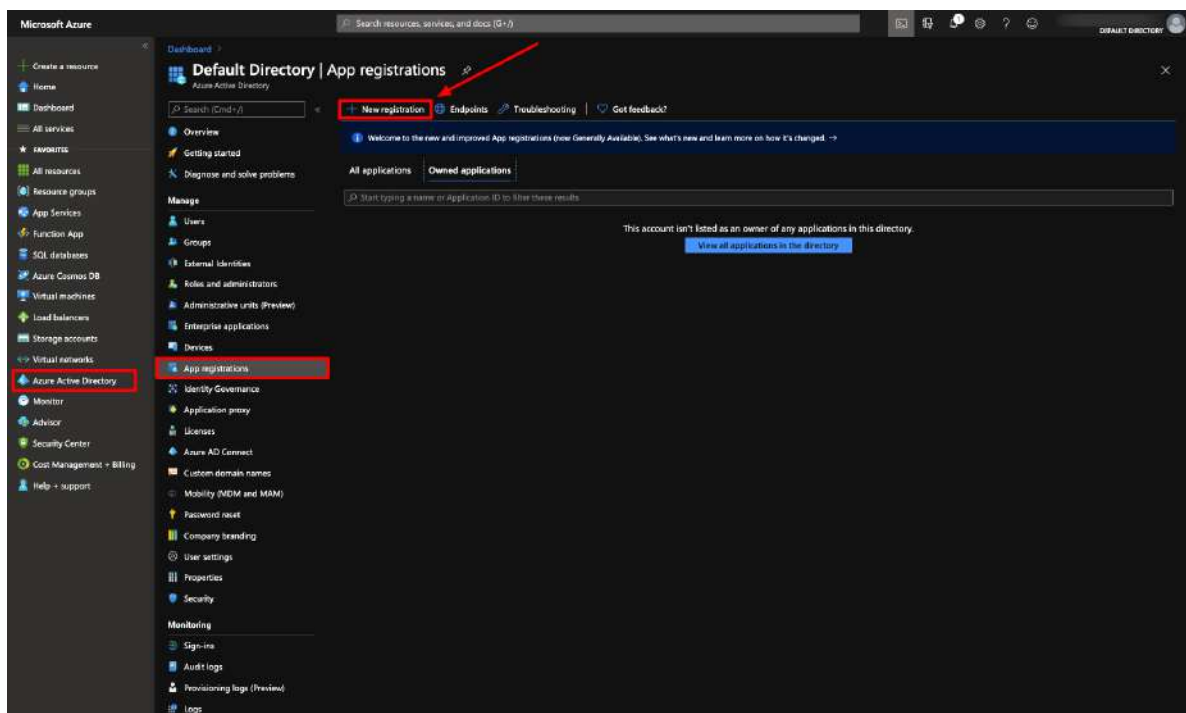
Create App Registration

Morpheus authenticates with Azure via an App Registration with an Owner or Contributor Role on a Subscription. Use the steps below to create and collect the required credentials and assign the required permissions to integrate Azure with Morpheus.

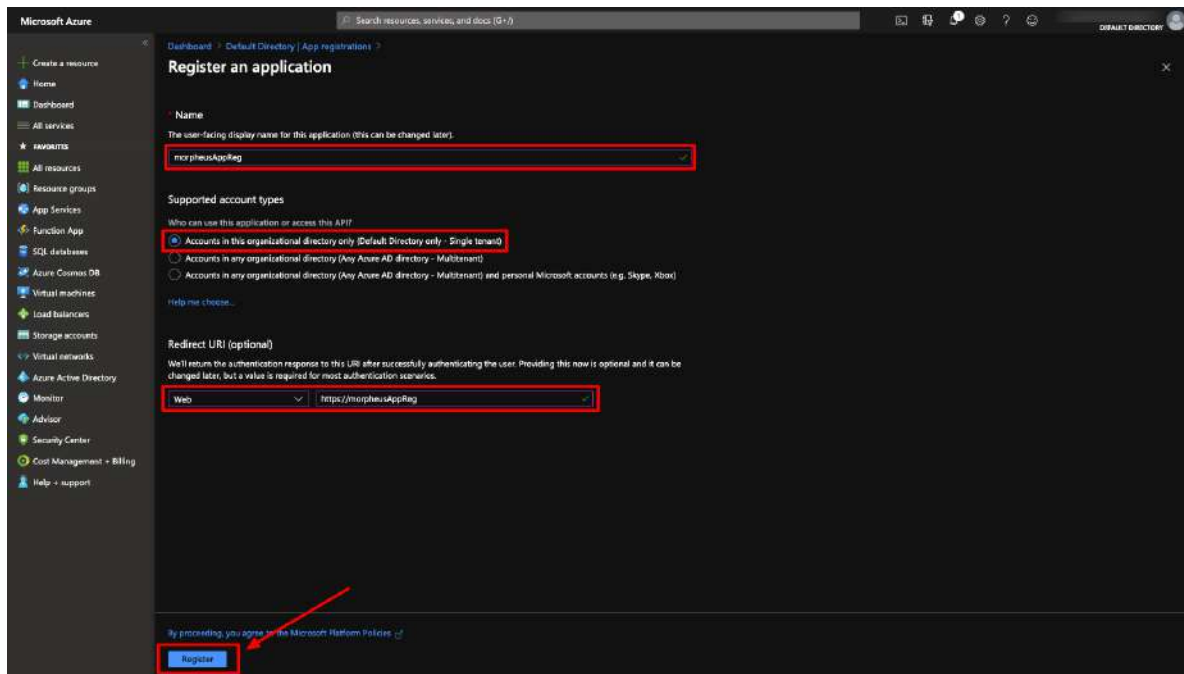
Warning: Using an App Registration (service principal) that has selective resource permissions and is not an Owner or Contributor of the Subscription is not supported and will cause failures/issues. Please confirm the App Registration you use to integrate Azure with Morpheus has Owner or Contributor permissions on the specified Subscription.

If you do not have an existing Azure Active Directory App Registration, or you wish to use a new one for Morpheus, you will need to create one using the steps below. If you already have one you wish to use, continue to the next section.

1. Log into the Azure portal
2. Select “Azure Active Directory”
3. Select “App Registrations”
4. Select “New Registration”



5. Next, give the app a name, specify Web app / API for the type (default) and enter any URL for the Sign-on URL:
6. Click Create and your new App Registration will be created.

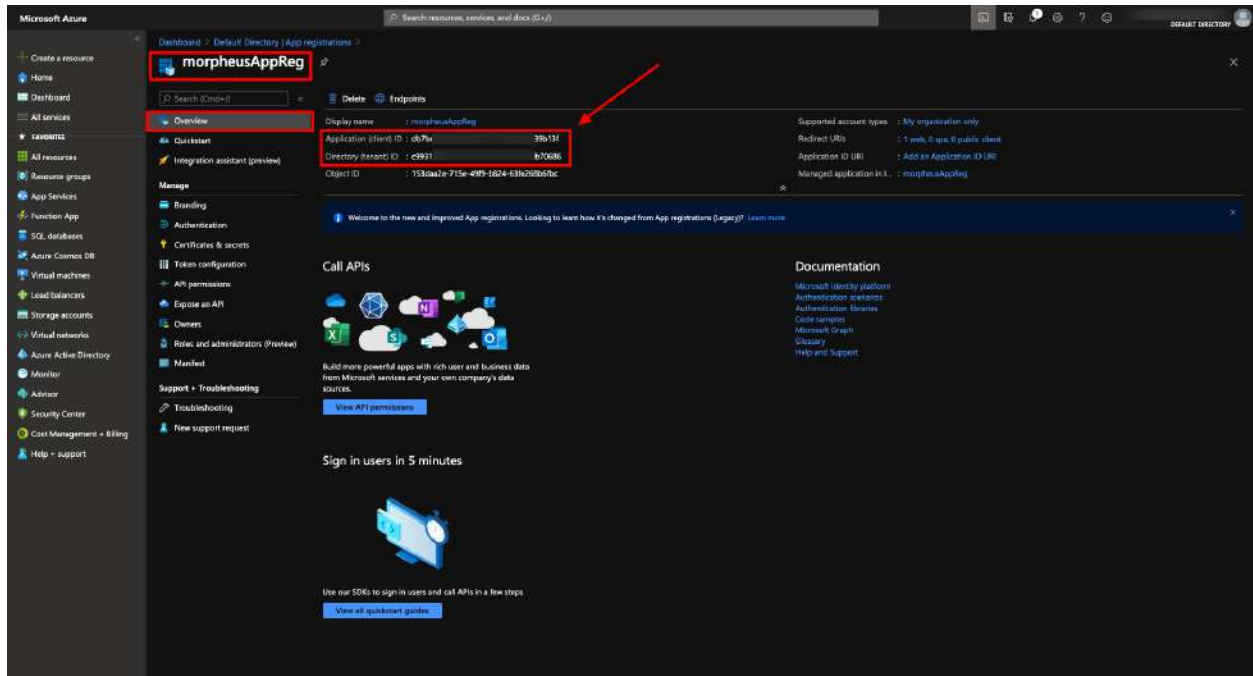


Now that we have our App Registration, we will gather the credentials required for the Morpheus Azure integration in the next section.

Copy Directory (tenant) and Application (client) IDs

The App Registration Directory (tenant) and Application (client) ID are required for the Morpheus Azure integration. Both can be found in the overview section of the App Registration.

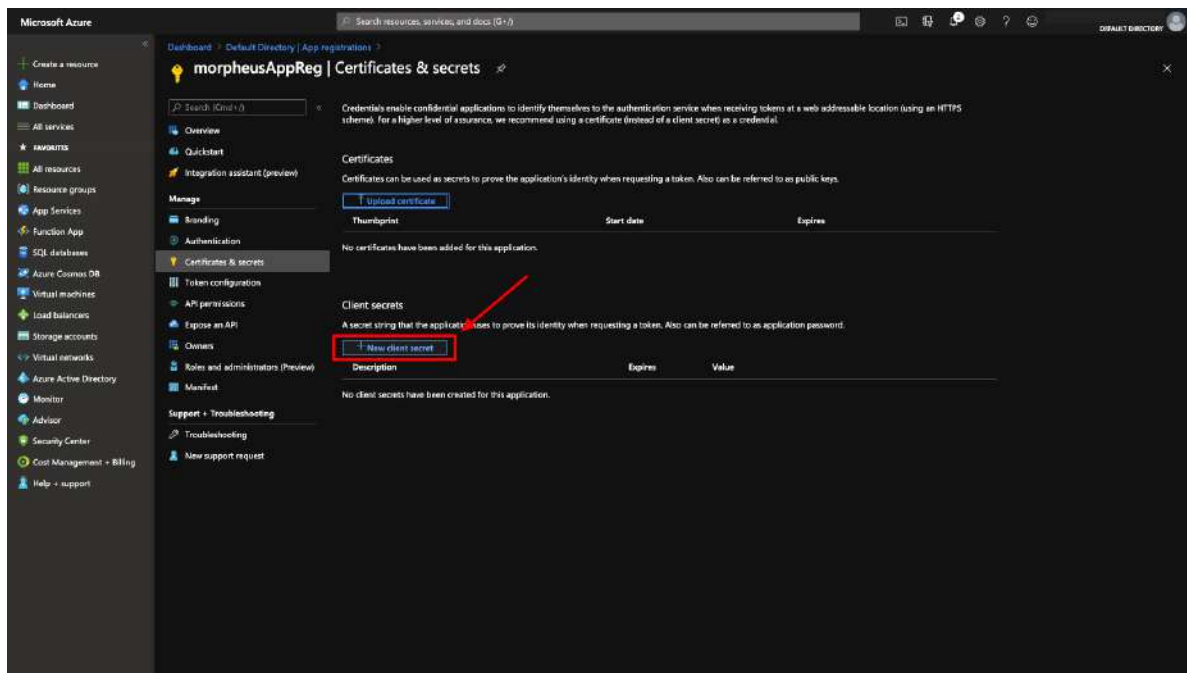
1. Go to the Overview section of your App Registration
2. Copy the Directory (tenant) ID
3. Store/Paste for use as the Tenant ID when adding your Azure cloud in Morpheus
4. Copy the Application (client) ID
5. Store/Paste for use as the Client ID when adding your Azure cloud in Morpheus



Generate a Client Secret

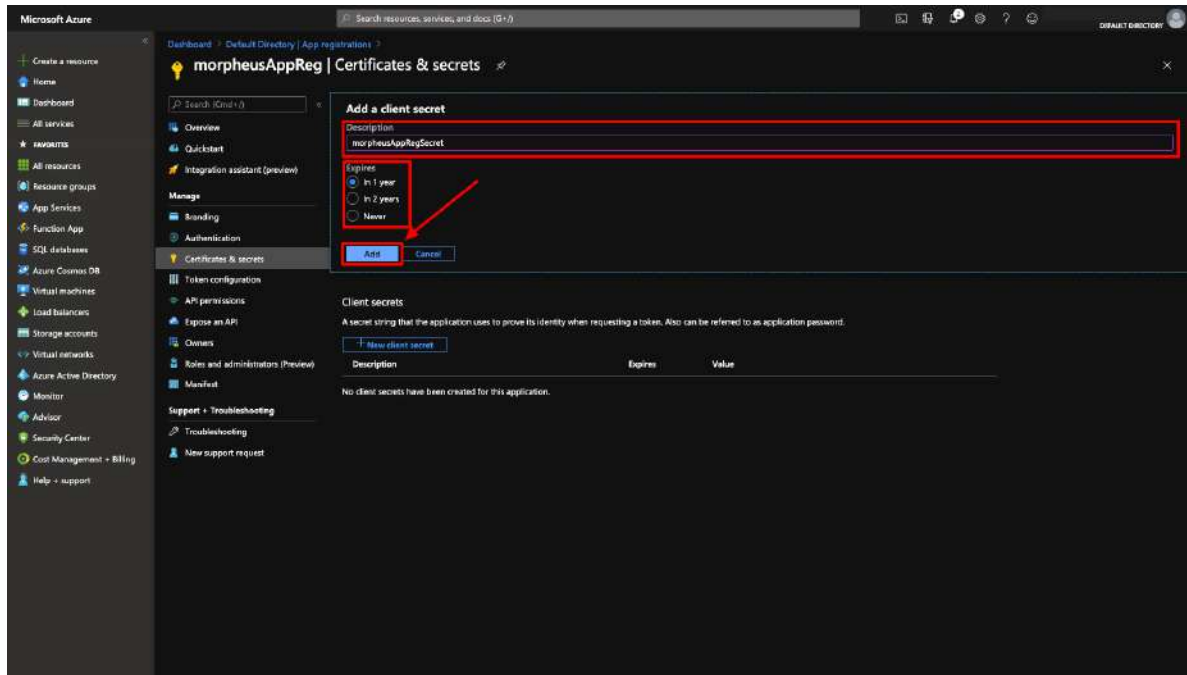
While still in your App Registration:

1. Select “Certificates & secrets” in the Manage section
2. Select + New client secret



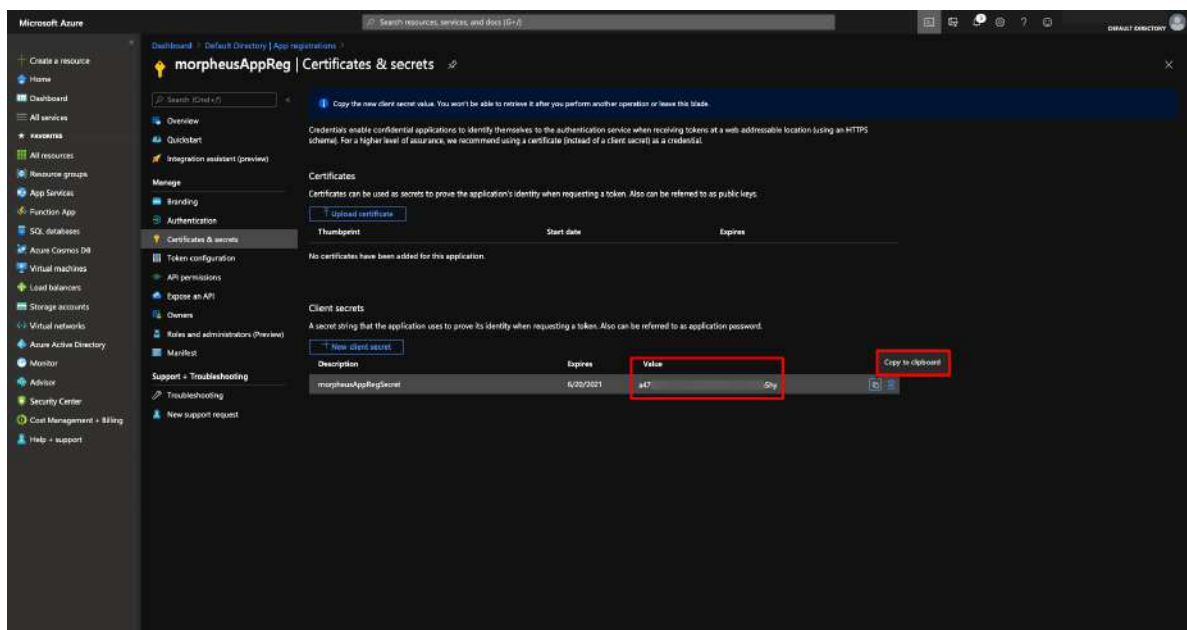
3. The “Add a client secret” modal will come up
4. Add a description to help identify the secret in the future

5. Select an expiration duration
6. Click *Add*



7. Copy the newly-generated client secret value.

Important: Copy the client secret value before continuing as it will not be viewable again later.



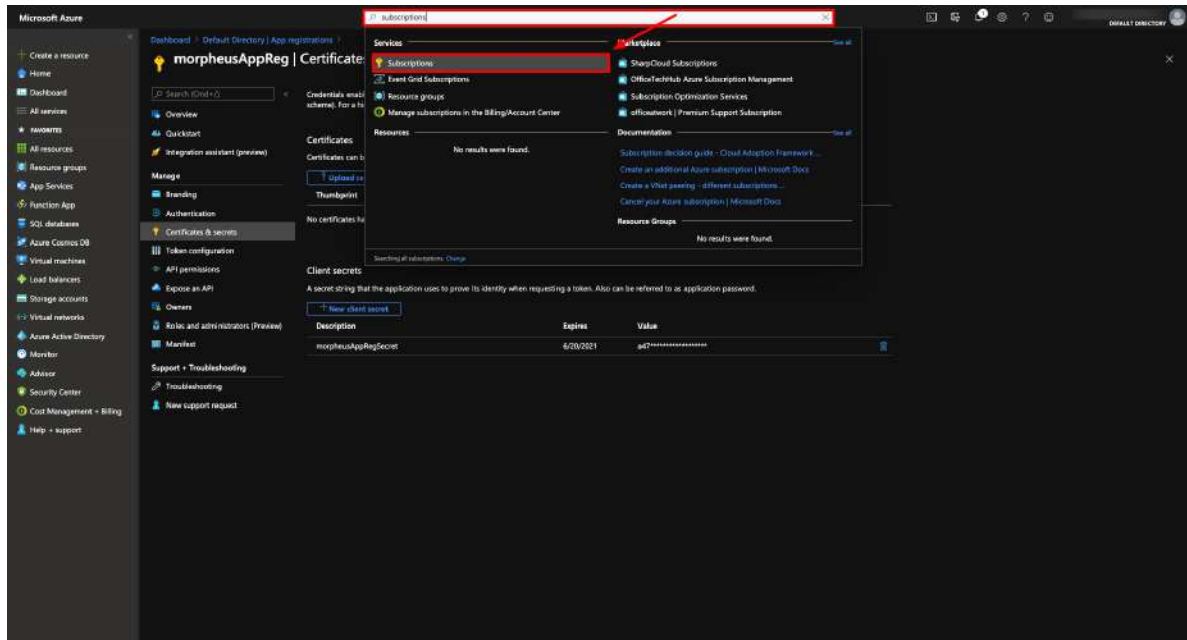
8. Store/Paste client secret for use later when adding your Azure cloud in Morpheus

You now have three of the four credentials required for Morpheus Azure cloud integration. The last credential required is the Azure Subscription ID which we will gather in the next section.

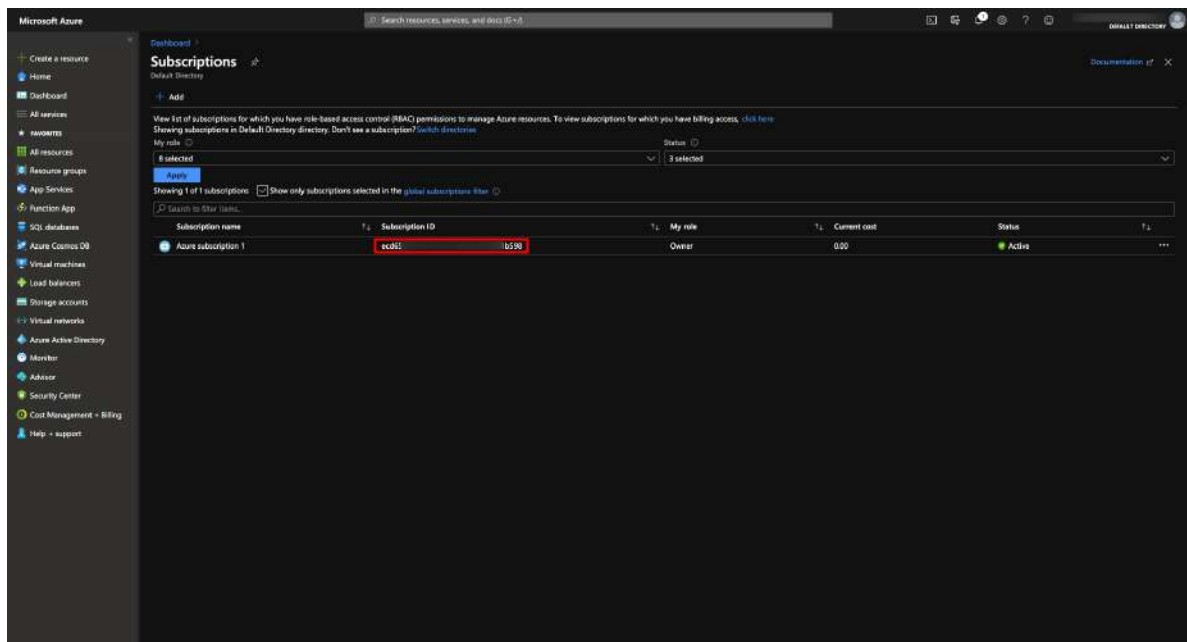
Subscription ID

To get the Azure Subscription ID:

1. Navigate to the Subscriptions section. The search function can help to locate these sections if they aren't immediately apparent in the UI menu



2. In the Subscriptions section, copy the Subscription ID

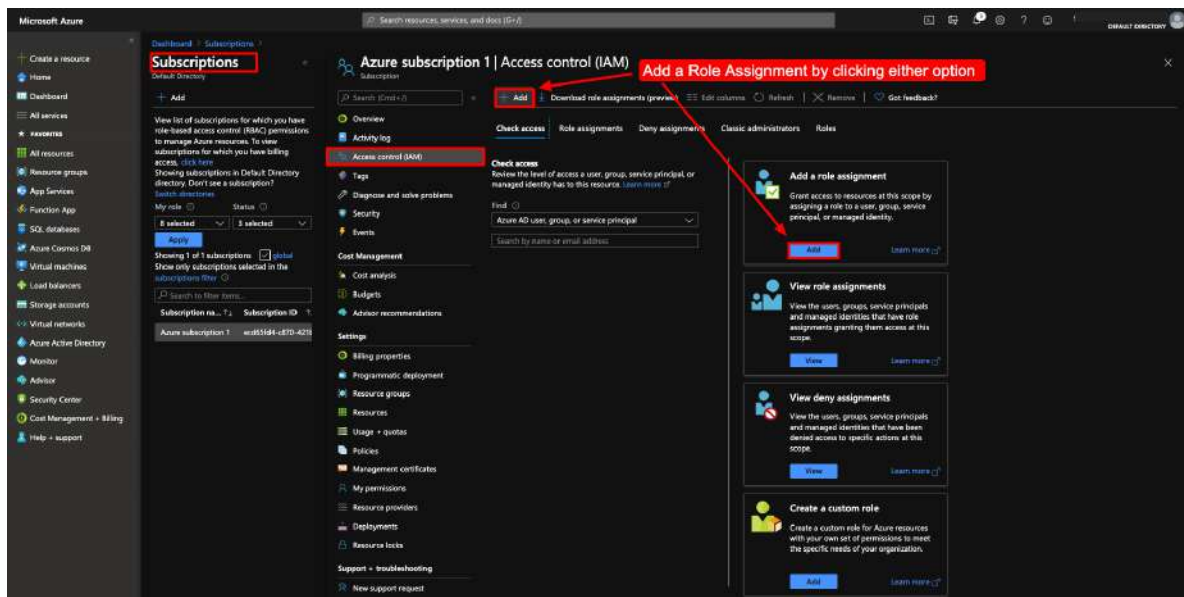


3. Store/Paste for use as the Subscription ID when adding your Azure cloud in Morpheus

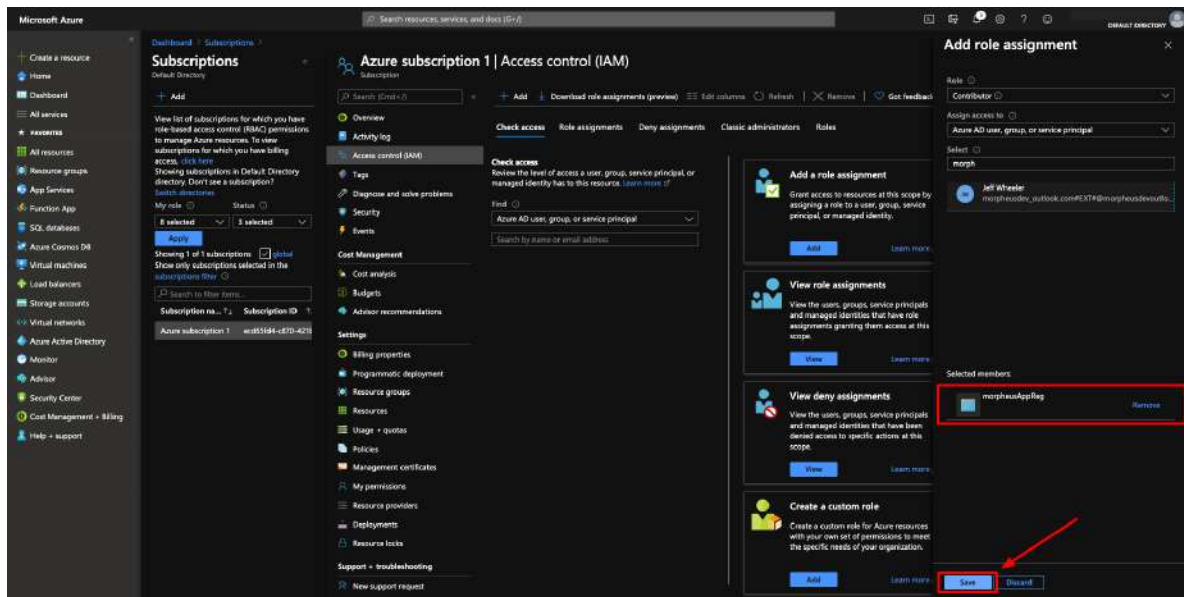
Make App Registration owner or contributor of Subscription

The App Registration used needs to be an owner of the Azure Subscription used for the Morpheus cloud integration. If lesser permissions are given or permissions are assigned at individual resource levels, Morpheus will not be able to properly inventory existing cloud resources, create resources or remove them.

1. In the Subscriptions section in Azure, select the Subscription
2. In the Subscription pane, select “Access Control (IAM)”
3. Either Click :guilabel` + Add`, and then “Add Role Assignment” OR simply select “Add a role assignment”



4. In the right pane, select “Owner” or “Contributor” Role type
5. Search for the name of the App Registration used for the Morpheus integration
6. Select the App Registration in the search results
7. Select “Save”



You now have the required credentials and permissions to add an Azure Cloud integration into Morpheus. Continue on with the next sections of this guide to complete the integration from the Morpheus side.

Complete the Add Cloud Process in Morpheus

If you've followed this guide from the start, you will already have a Cloud integration modal open in Morpheus UI. If you still need to open that wizard, navigate to Infrastructure > Clouds > + ADD > Azure (Public) and click **NEXT**. Fill in the following fields with the information gathered in the steps above:

- Subscription ID
- Tenant ID
- Client ID
- Client Secret
- Location
- Resource Group
- Inventory Existing Instances
- Inventory Level
- Account Type

Once valid credentials are populated in the appropriate fields, the **LOCATION** dropdown menu will be populated. Select an available region, this is also a helpful check to ensure you've correctly provided working credentials. In addition, we can scope the cloud integration to all resource groups in the region (All) or can select a specific resource group to limit Morpheus resource inventorying and creation to just that resource group.

By checking **INVENTORY EXISTING INSTANCES**, Morpheus will automatically onboard existing cloud resources which are scoped to the region and resource group indicated. If this box is checked, we will also need to select either basic inventorying, which syncs name, IP addresses, platform types, power status, and sizing data (storage, CPU, and RAM) OR full (API heavy) inventorying which syncs resource utilization metrics (storage, CPU, and RAM) when available in addition to what we get with basic inventorying.

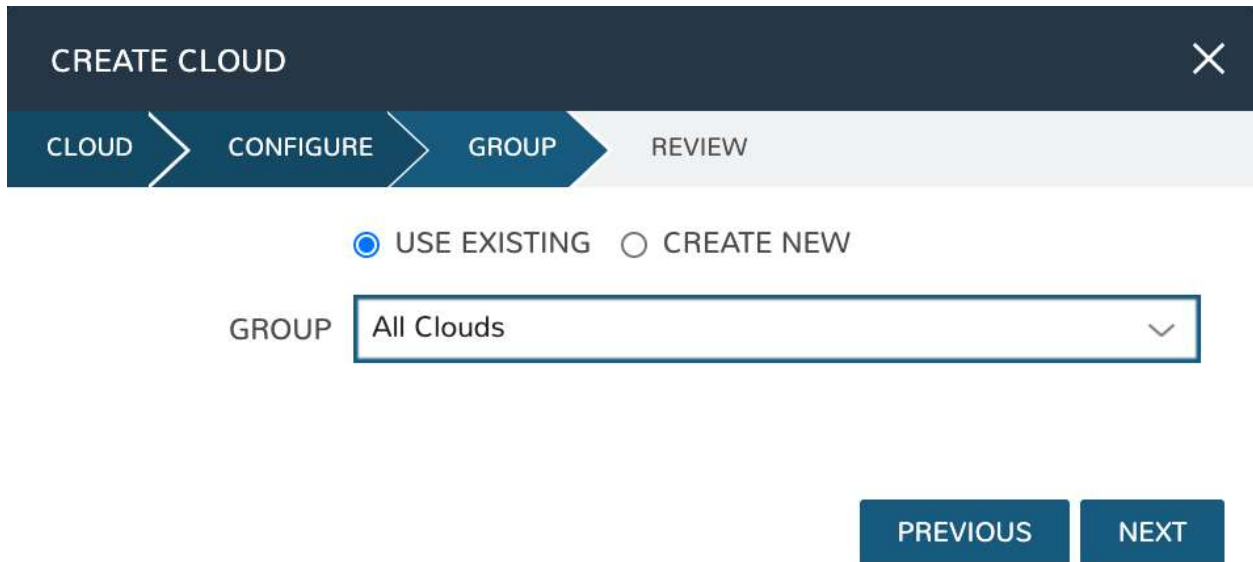
To move on, expand the "Advanced Options" section.

Note: CSP accounts will also need to enter CSP TENANT ID, CSP CLIENT ID, and CSP CLIENT SECRET in the Advanced Options section.

Within the “Advanced Options” drawer are additional configurations to consider for your first Cloud. Some of these won’t usable until they reference additional configured integrations. Common settings to consider are **DOMAIN**, **STORAGE TYPE**, **APPLIANCE URL** (overrides the Morpheus URL for external systems), **GUIDANCE** (setting “Manual” will make recommendations for rightsizing), **COSTING**, **DNS INTEGRATION**, **CMDB**, and **AGENT INSTALL MODE**.

Once you’re satisfied with your selections, click “NEXT”

We have now arrived at the “GROUP” tab. In this case, we will mark the radio button to “USE EXISTING” Groups if you wish to use the Group we configured earlier. Alternatively, you can create a new one here.



The screenshot shows the 'CREATE CLOUD' wizard interface. At the top is a dark blue header with the title 'CREATE CLOUD' and a close button (X). Below the header is a progress bar with four steps: 'CLOUD', 'CONFIGURE', 'GROUP' (which is the active step and highlighted with a dark blue arrow), and 'REVIEW'. Below the progress bar, there are two radio buttons: 'USE EXISTING' (which is selected) and 'CREATE NEW'. Below the radio buttons is a label 'GROUP' followed by a dropdown menu that currently displays 'All Clouds'. At the bottom right of the form are two buttons: 'PREVIOUS' and 'NEXT'.

Once you’ve selected or created the Group, click “NEXT”

On the final tab of the “CREATE CLOUD” wizard, you’ll confirm your selections and click “COMPLETE”. The new Cloud is now listed on the Cloud list page. After a short time, Morpheus will provide summary information and statistics on existing virtual machines, networks, and other resources available in the Cloud.

Viewing Cloud Inventory

Now that we've integrated our first Azure cloud, we can stop for a moment to review what Morpheus gives us from the Cloud detail page. We can see that Morpheus gives us estimated costs and cost histories, metrics on used resources, and also lists out resource counts in various categories including container hosts, hypervisors, and virtual machines. We can drill into these categories to see lists of resources in the various categories by clicking on the category tabs. We can link to the detail page for any specific resource by clicking on it from its resource category list.

Configuring Resource Pools

With our Azure Cloud configured, Morpheus will automatically sync in available resource pools and data stores.

For resource pools, once Morpheus has had time to ingest them, then will be visible from the cloud detail page. Navigate to *Infrastructure > Clouds > (your Azure cloud) > Resources tab*. In here, we are able to see and control access to the various resource pools that have been configured in Azure. For example, we can restrict access to a specific resource pool within Morpheus completely by clicking on the "ACTIONS" button, then clicking "Edit". If we unmark the "ACTIVE" button and then click "SAVE CHANGES" we will see that the resource pool is now grayed out in the list. The resources contained in that pool will not be accessible for provisioning within Morpheus if it is not configured as active.

The screenshot shows the Morpheus UI interface. At the top, there's a navigation bar with tabs: Operations, Provisioning, Infrastructure (selected), Backups, Logs, Monitoring, Tools, and Administration. Below this is a sub-navigation bar with icons for Groups, Clouds, Clusters, Hosts, Network, Load Balancers, Storage, Keys & Certs, and Boot. The main content area shows the 'Clouds > AdamAzure' detail page. It includes a status bar with a green checkmark, 'Last Sync: 07/28/2020 02:39 PM', and 'Sync Duration: 6 seconds'. Below this are five circular gauges: 'COST THIS MONTH' (\$0), 'AVG MONTHLY COST' (0.0), 'MAX CPU' (0%), 'MEMORY' (0%), and 'STORAGE' (0%). Underneath these are five resource count cards: 'CONTAINER HOSTS' (0), 'HYPERVISORS' (0), 'BARE METAL' (0), 'VIRTUAL MACHINES' (1), and 'DISCOVERED' (1). At the bottom, there's a 'POOLS' section with a search bar and a '+ ADD RESOURCE POOL' button. Below the search bar is a table with columns: NAME, DESCRIPTION, VISIBILITY, DEFAULT, TENANT, and ACTIONS. The table shows one resource pool with the name 'morpheus', visibility 'Private', and tenant 'morpheus'.

NAME	DESCRIPTION	VISIBILITY	DEFAULT	TENANT	ACTIONS
morpheus		Private		morpheus	ACTIONS

Often our clients will want to make specific blocks of resources available to their own customers. This can be easily

and conveniently controlled through the same “EDIT RESOURCE POOL” dialog box we were just working in. If we expand the “Group Access” drawer, we are able to give or remove access to each pool to any Group we’d like. We can also choose to make some or all of our resource pools available to every Group. Specific resource pools can also be defined as the default for each Group when needed.

EDIT RESOURCE POOL

NAME

☐ MOVE SERVERS

☒ ACTIVE

☐ DEFAULT

▼ Group Access

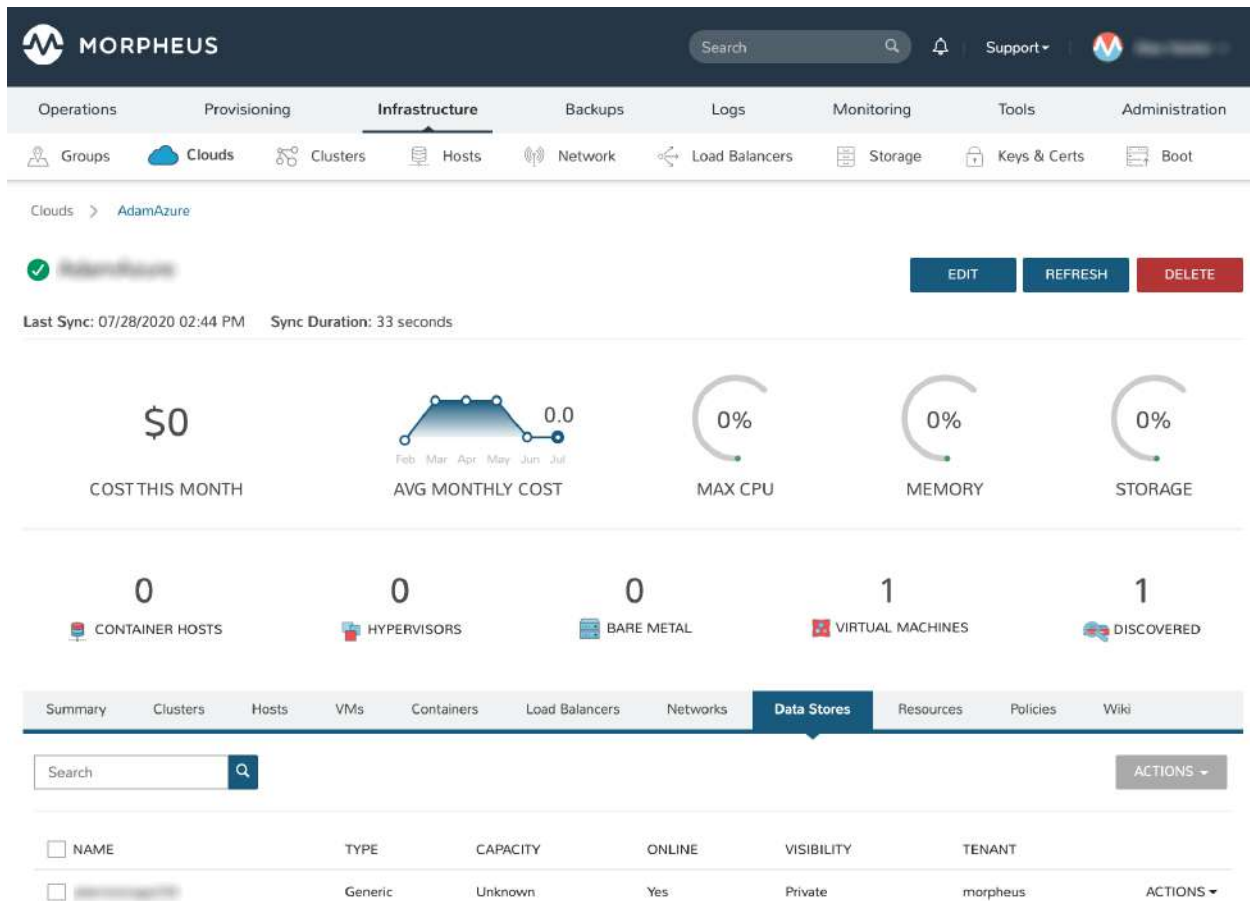
GROUP	ACCESS	DEFAULT
all	<input checked="" type="checkbox"/>	
	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="checkbox"/>	<input type="checkbox"/>
	<input type="checkbox"/>	<input type="checkbox"/>

Additionally, we may choose to allow only certain service plans to be provisioned into a specific pool of resources. For example, perhaps a specific cluster is my SQL cluster and only specific services plans should be consumable within it. We can control that through this same dialog box.

Configuring Data Stores

To take a look at data stores, we’ll move from the “Resources” tab to the “Data Stores” tab on our Cloud detail page.

Morpheus gives the user similar control with data stores to what we saw with our resources pools earlier. Just like with resource pools, we can disable access within Morpheus completely by clicking on “ACTIONS” and then “Edit”. If we unmark the “ACTIVE” checkbox and click “SAVE CHANGES”, you will see that specific data store has been grayed out.



Just like with resource pools, we are also able to scope data stores to specific Groups. This ensures that the members of each Group are only able to consume the data stores they should have access to.

Configuring Network for Provisioning

When configuring networking, we can set global defaults by going to *Infrastructure > Network > NETWORKS* tab. Here we can add or configure networks from all Clouds integrated into Morpheus. Depending on the number of clouds Morpheus has ingested, this list may be quite large and may also be paginated across a large number of pages. In such a case, it may be easier to view or configure networks from the specific Cloud detail page so that networks from other Clouds are not shown.

Still in *Infrastructure > Network*, make note of the “INTEGRATIONS” tab. It’s here that we can set up any integrations that may be relevant, such as IPAM integrations. Generally speaking, when adding IPAM integrations, we simply need to name our new integration, give the API URL, and provide credentials. There’s more information in the [IPAM integration](#) section of Morpheus Docs.

In *Infrastructure > Networking* we can also set up IP address pools from the IP Pools tab. These pools can be manually defined, known as a Morpheus-type IP pool, or they can come from any IPAM integrations you’ve configured. As Instances are provisioned, Morpheus will assign IP addresses from the pool chosen during provisioning. When the Instance is later dissolved, Morpheus will automatically release the IP address to be used by another Instance when needed. When adding or editing a network, we can opt to scope the network to one of these configured IP address pools.

Since this guide is focused on working within an Azure cloud that we integrated at the start, we will take a look at our network configurations on the cloud detail page as well. Navigate to *Infrastructure > Clouds > (your Azure cloud) >*

NETWORKS tab. Just as with resource pools and data stores, we have the ability to make certain networks inactive in Morpheus, or scope them to be usable only for certain Groups or Tenants.

The screenshot displays the Morpheus web interface. At the top, the 'Infrastructure' tab is selected, showing a navigation bar with options like Groups, Clouds, Clusters, Hosts, Network, Load Balancers, Storage, Keys & Certs, and Boot. Below this, the 'Networks' section is active, showing a summary of network resources. The summary includes a table of networks and a section for host security groups.

Networks Summary:

NAME	TYPE	CIDR	POOL	DHCP VISIBILITY	TENANT	ACTIONS
default	Subnet			Private	morpheus	ACTIONS

Host Security Groups:

The Host Level Firewall is not currently enabled. These Security Groups will not be applied unless this setting is turned on in the cloud settings host firewall.

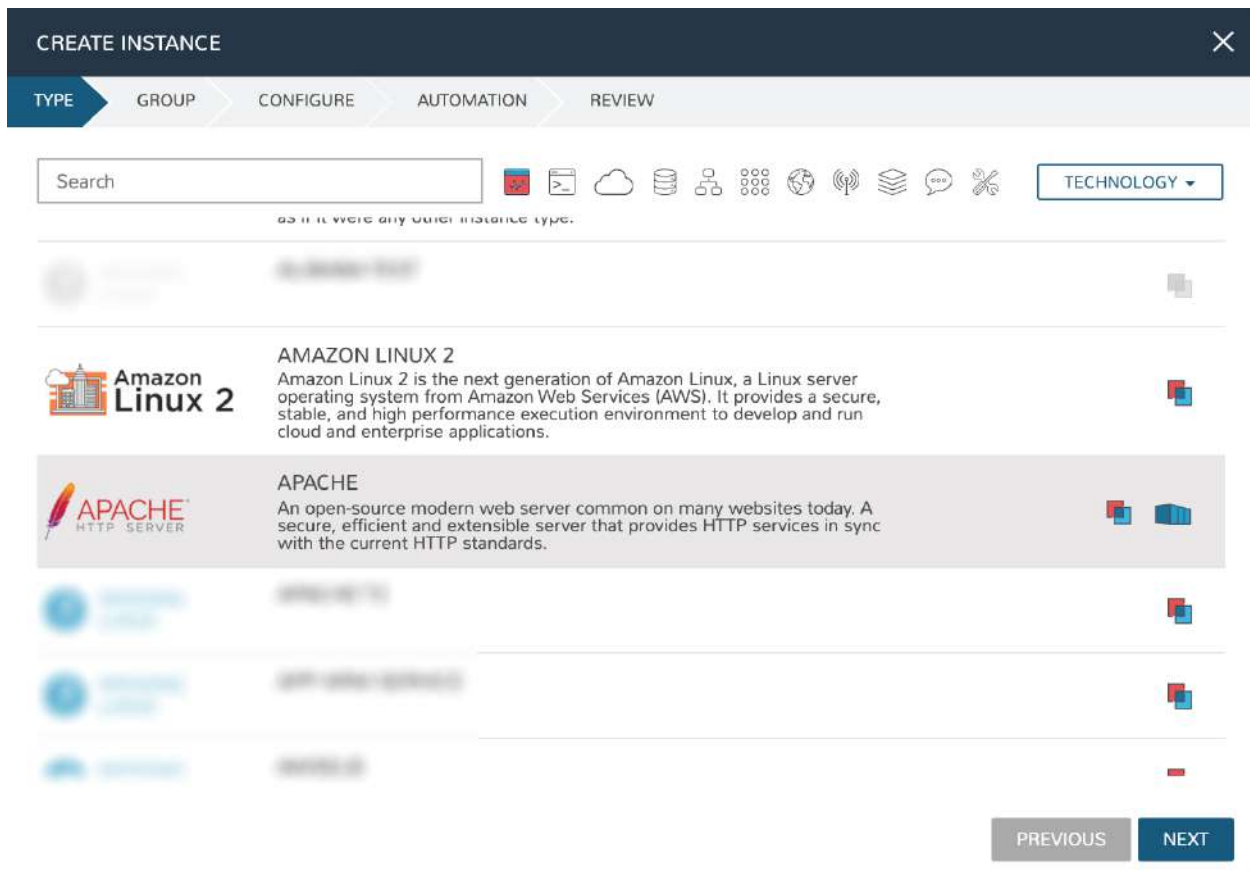
[EDIT SECURITY GROUPS](#)

Provisioning Your First Instance

At this point, the groundwork is laid and we are ready to attempt our first new provisioning. As a first Instance, we'll provision an Apache web server to our Azure cloud. Morpheus includes a very robust catalog of pre-configured Instance types. We'll use one of these included catalog items for this guide but you'll likely also need to prep your own custom images and Instance types to make available to your users. Much more on this can be found elsewhere in Morpheus documentation.

Navigate to *Provisioning > Instances*. If any Instances are currently provisioned, we will see them listed here. To start a new Instance we click + **ADD** to open the "CREATE INSTANCE" wizard. We'll scroll down to and select the

Apache instance type and click “NEXT”.



First, we'll specify the Group to provision into which determines the Clouds available. If you've followed this guide to this point, you should at least have a Group that houses all of your Clouds which you can select here. This will allow us to select the Azure cloud from the “CLOUD” dropdown menu. Provide a unique name to this instance and then click “NEXT”

From the “CONFIGURE” tab, we're presented with a number of options. The options are cloud and layout-specific, more generalized information on creating Instances and available options is [here](#). For our purposes, we'll select the following options:

- **LAYOUT:** Includes options such as the base OS, custom layouts will also be here when available
- **PLAN:** Select the resource plan for your instance. Some plans have minimum resource limits, Morpheus will only show plans at or above these limits. User-defined plans can also be created in *Administration > Plans & Pricing*.
- **VOLUMES:** The minimum disk space is set by the plan, this value may be locked if you've selected a custom plan that defines the volume size
- **NETWORKS:** Select a network

Under the “User Config” drawer, mark the box to “CREATE YOUR USER”. Click *NEXT*.

CREATE INSTANCE

×

TYPE

GROUP

CONFIGURE

AUTOMATION

REVIEW

Configuration Options

LAYOUT

Azure Apache

▼

PLAN

Standard_B2ms (2 Core, 8GB Memory) (westus)

▼

Cores: 2 Memory: 8 GB Price: \$73.0816 / Month ⓘ

RESOURCE POOL

▼

VOLUMES

root

Page Blob

▼

morpheusfuorhvcynyobjtra

▼

+

NETWORKS

▼

DHCP

+

SECURITY GROUP

Select

▼

AVAILABILITY SET

No Availability Set

▼

☒ ASSIGN PUBLIC IP

► User Config

► Advanced Options

PREVIOUS

NEXT

Note: “CREATE YOUR USER” will seed a user account into the VM with credentials set in your Morpheus user account settings. If you’ve not yet defined these credentials, you can do so by clicking on your username in the upper-right corner of the application window and selecting “USER SETTINGS”.

For now, we’ll simply click *NEXT* to move through the “AUTOMATION” tab but feel free to stop and take a look at the available selections here. There is more information later in this guide on automation and even more beyond that in the rest of Morpheus docs.

Review the settings for your first instance and click *COMPLETE*.

We are now dropped back onto the Instances list page. We can see a new entry in the list at this point with a status indicator that the new machine is being launched (rocket icon in the status field). We can double click on the Instance in the list to move to the Instance detail page. For now we will see a progress bar indicating that the Instance is being created and is starting up. The exact amount of time this process will take depends on selections made when provisioning the Instance. Initially, Morpheus will guess as to how long this will take and the progress bar may not be accurate. Over time, Morpheus will learn how long these processes take and progress bar accuracy will improve. For more detailed information on the status of various provisioning processes, we can scroll down and select the “HISTORY” tab. The “STATUS” icon will change from the blue rocket to a green play button when the Instance is fully ready. Furthermore, we can click on the hyperlinked IP address in the “VMS” section of this page to view a

default page in a web browser to confirm success.

Creating Your First Library Item

In the prior section, we manually provisioned our first Instance. However, Morpheus allows you to build a catalog of custom provisionable items to simplify and speed provisioning in the future. In this section, we'll build a catalog item and show how that can translate into quick Instance provisioning after configuration.

Note: Before starting this process, it's important to decide which virtual image you plan to use. If you're not using a Morpheus-provided image, you'll want to ensure it's configured. You will not be able to complete this section without selecting an available image. In this example we will use a CentOS image that was previously configured in the Morpheus library. If you need to configure your own images prior to starting this section, navigate to Provisioning > Virtual Images and click + *ADD*. A deeper dive into image prep and virtual image configuration goes beyond the scope of this guide.

Provisionable elements in Morpheus combine a Node Type(s), Layout(s), and an Instance Type. The [Overview section](#) of Morpheus docs discusses these objects and how they work together in greater detail. Our first step here will be to create a Node Type which wrap the image itself with additional configuration, templates, and scripts. While not strictly required, creating the Node Type, Instance Type, and then the Layout is often a good workflow for creating Library items. That is the order we will follow in this guide.

Navigate to *Provisioning > Library > NODE TYPES* and click + *ADD*

In this example, I am going to set the following options in the "NEW NODE TYPE" wizard:

- **NAME:** *Example Azure CentOS 7*
- **SHORT NAME:** *eac7* (Identifies the Node Type in Morpheus API/CLI)
- **VERSION:** *7* (Ensures the correct Node Types are used when tying Layouts with multiple images to the same Instance Type)
- **TECHNOLOGY:** *Azure*
- **VM IMAGE:** *Azure-Centos-7*

Click *SAVE CHANGES*

EDIT NODE TYPE

NAME SHORT NAME

The short name is a name with no spaces used for display in your container list.

VERSION ENVIRONMENT
VARIABLES

Azure Options

VM IMAGE LOG FOLDER DEPLOY FOLDER

(Optional) If using deployment services, this mount point will be replaced with the contents of said deployments.

SERVICE PORTS



SCRIPTS

FILE TEMPLATES

► Advanced Options

SAVE CHANGES

With the new Node Type created, we'll now add a new Instance type which will be accessible from the provisioning wizard once created. Move from the "NODE TYPES" tab to the "INSTANCE TYPES" tab and click + *ADD*.

In the "NEW INSTANCE TYPE" wizard, I'll simply enter a **NAME** and **CODE** value. Click *SAVE CHANGES*. You could also provide a description, icon, and category for easier identification from the provisioning wizard later.

EDIT INSTANCE TYPE

×

NAME

*Example Azure CentOS 7

CODE

hc7

Useful shortcode for provisioning naming schemes and export reference.

DESCRIPTION

CATEGORY

OS

▼

ICON

Browse

Suggested Dimensions: 150 x 51

Now that we've created a new Instance type, access it by clicking on the name in the list of custom Instances you've created. In my case, I've given the name "*Example Azure CentOS 7*".

Once we've opened the new Instance type, by default, we should be on the "LAYOUTS" tab. Click + *ADD LAYOUT*. I've set the following fields on my example layout:

- **NAME:** *Example Azure CentOS 7*
- **VERSION:** 7 (This is the version number of the layout itself, which is labeled 7 in the example)
- **TECHNOLOGY:** Azure
- **Nodes:** Select the Node Type we created earlier, if desired you can specify multiple nodes

Click *SAVE CHANGES*.

EDIT LAYOUT

NAME

Harker Azure CentOS 7

VERSION

7

DESCRIPTION

☒ CREATABLE

TECHNOLOGY

Azure

MINIMUM MEMORY

1024

MB

This will override any memory requirement set on the virtual image

WORKFLOW

Select Workflow

☐ SUPPORTS CONVERT TO MANAGED

ENVIRONMENT VARIABLES

Name

Value

Option Types

Search option types

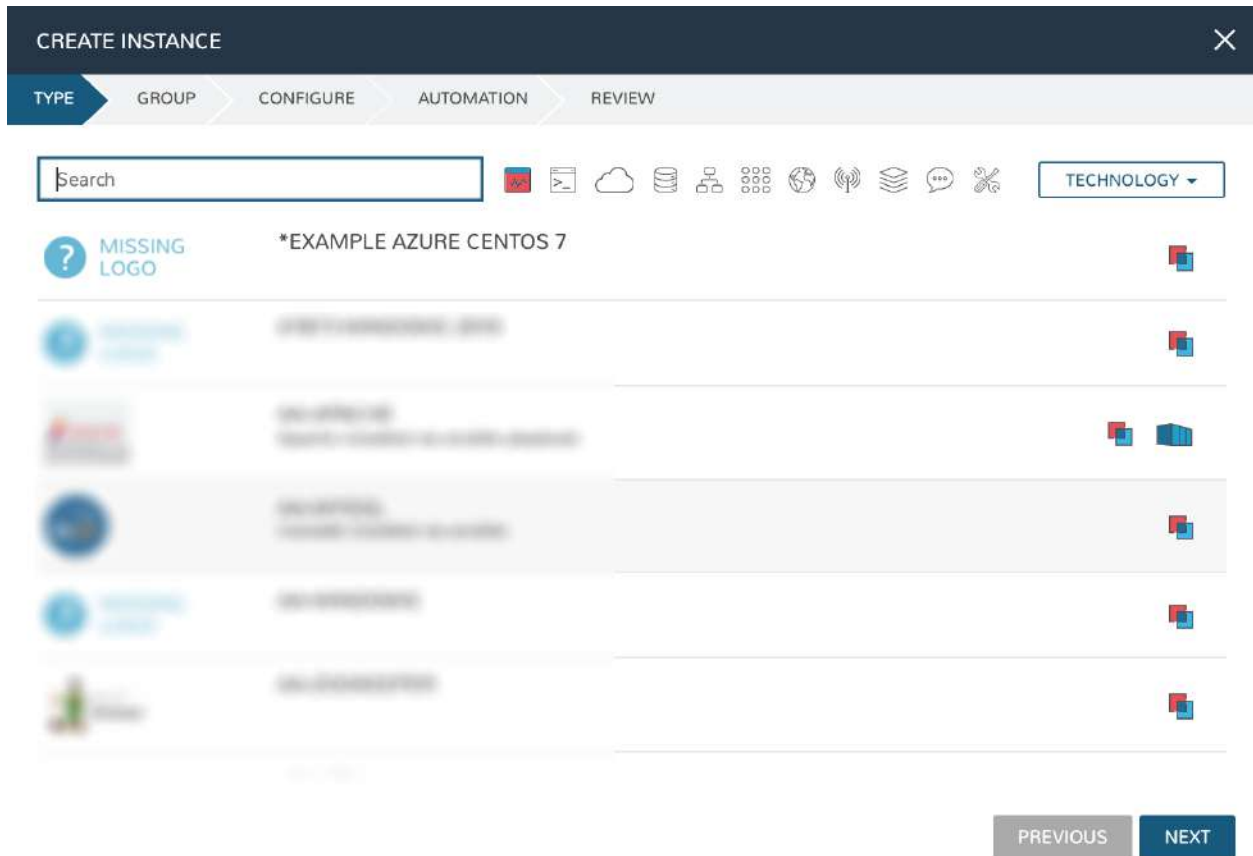
Nodes

Search nodes

*Example Azure CentOS 7 (7)

SAVE CHANGES

At this point we've completed the setup work and can now provision the Instance we've created to our specifications. Navigate to *Provisioning > Instances* and click + *ADD*. From the search bar we can search for the new Instance type we've created.



As before, we can select a Group and Cloud to provision this new Instance. Click *NEXT*. On the “CONFIGURE” tab, make note that the layout and plan are already selected because they were configured as part of creating the new Instance type. Select a network and click *NEXT*. Once again we will also click *NEXT* through the “AUTOMATION” tab. Finally, click *COMPLETE*.

As before when we provisioned a pre-existing Instance from the default catalog, Morpheus will now begin to spin up the new VM. How long this will take depends on the configuration and environmental factors but Morpheus will predict how long this process will take and represent that on a progress bar. Over time, Morpheus begins to learn how long these processes take and becomes more accurate in predicting spin-up time.

Once the provisioning process has completed, open the Instance detail page in Morpheus and click on the “CONSOLE” tab. You’ll be logged in with your user account and are then able to confirm the machine is ready and available, assuming the image and your custom catalog item were configured to seed user accounts and connect back to the Morpheus appliance.

Automation and Configuration Management

Morpheus automation is composed of Tasks and Workflows. A Task could be a script added directly, scripts or Blueprints pulled from the Morpheus Library, playbooks, recipes, or a number of other things. The complete list of Task types can be found in the [Automation section](#) of Morpheus docs. Tasks can be executed individually but they are often combined into workflows. We can opt to run a workflow at provision time or they can be executed on existing instances through the Actions menu.

In this guide we will set up an Ansible integration, create a Task, add the Task to a Workflow, and run the Workflow against a new and existing Instance. If you've worked through this guide to this point, you should already have an Apache instance running. If you don't yet have that, provision one before continuing with this guide and ensure it's reachable on port 80.

We'll first set up the Ansible integration, you can integrate with the sample repository referenced here or integrate with your own. Go to 'Administration > Integrations'. Click **+NEW INTEGRATION** and select Ansible from the dropdown menu. Fill in the following details:

- **NAME**
- **ANSIBLE GIT URL:** <https://github.com/ncelebic/morpheus/-ansible-example>, or enter the URL for your own Ansible git repository
- **PLAYBOOKS PATH**
- **ROLES PATH**
- Mark the box to "USE Morpheus AGENT COMMAND BUS"

Note: If your git repository requires authentication, you should create a keypair and use the following URL format: <git@github.com:ncelebic/morpheus/-ansible-example.git>.

Click **SAVE CHANGES**. You'll now see our new Ansible integration listed among any other configured integrations. If we click on this new integration to view detail, a green checkmark icon indicates the git repository has been fully synced.

With the Ansible integration set up, we can now create a task that includes our playbook. Go to *Provisioning > Automation*, click **+ ADD**. We'll first set our "TYPE" value to Ansible Playbook so that the correct set of fields appear in the "NEW TASK" wizard. Set the following options:

- **NAME**
- **ANSIBLE REPO:** Here we will choose the Ansible integration that we just created
- **PLAYBOOK:** In our example case, enter 'playbook.yml'

Click "SAVE CHANGES" to save our new task. We can test the new task on our Apache VM now by going to *Provisioning > Instances* and clicking into our VM. From the "ACTIONS" menu select "Run Task". From the "TASK" dropdown menu, select the task we just added and click "EXECUTE".

To see the progress of the task, click on the "HISTORY" tab and click on the (i) button to the right of each entry in the list. In this case, we can also see the results of the task by clicking on the link in the "LOCATION" column of the "VMS" section.

Now that our task is created, we can put it into a workflow. Back in *Provisioning > Automation* we will click on the "WORKFLOWS" tab. Click **+ADD** and select Provisioning Workflow. We'll give the new workflow a name and expand the Post Provision section. As we begin to type in the name of the task we've created, it should appear as a selection. Click **SAVE CHANGES**.

Now that we have a Workflow, return to *Provisioning > Instances* and begin to provision another Apache instance. More detailed instructions on provisioning a new Apache instance are included earlier in this guide if needed. Now,

when you reach the “AUTOMATION” section of the “CREATE INSTANCE” wizard, we have a workflow to select. From the “WORKFLOW” dropdown menu, select the workflow we just created and complete provisioning of the new instance.

As the instance is provisioning, we can go to the “HISTORY” tab and see Morpheus executing the tasks that were contained in our workflow.

This is just one example of using Morpheus to automate the process of configuring an instance to your needs. There are a number of other automation types that can be built into your Workflows as well. For further information, take a look at the [automation integrations](#) guide in Morpheus docs.

Conclusion

At this point you should be up and running in Morpheus, ready to consume Azure public cloud. This guide only scratches the surface, there is a lot more to see and do in Morpheus. Take a look at the rest of [Morpheus Docs](#) for more information on supported integrations and other things possible.

Backing Up and Restoring Morpheus Appliance

Morpheus includes built-in tools for backing up managed Instances as well as the appliance itself. Use this guide to configure a location and schedule for backing up your Morpheus appliance. This guide also includes steps for restoring or migrating your appliance from the created backup. The steps are the same whether your appliance is deployed in a single node or distributed architecture.

The built-in Morpheus appliance backup functionality backs up the MySQL data. In addition to the database, it’s advisable to back up your shared storage (at `/var/opt/morpheus/morpheus-ui`) and the `morpheus.rb` configuration file.

Note: The destination Morpheus appliance must be running the same version as that which the backup was taken from.

Create A Backup Job

A Backup Job in Morpheus holds the schedule timing and retention count for automated backups. If you already have a Job configured, you can move on to the next section. By default, Morpheus includes two execution schedules: Daily at Midnight and Weekly on Sunday at Midnight. If currently-existing options do not make sense for your backup needs, create a new execution schedule:

1. Navigate to Provisioning > Automation
2. Click on the “Execute Scheduling” tab
3. Click + *ADD*
4. Enter schedule timing using `cron` notation
5. Click *SAVE*

With the execution schedule created, we can move on to creating the Backup Job itself. A Backup Job includes both the backup retention count and an execution schedule (which we just created).

1. Navigate to Backups > Jobs
2. Click + *ADD*
3. Name the Job, then configure the retention count and the schedule

4. Click *SAVE*

Integrate a Bucket or File Share

When configuring a Morpheus appliance backup, a storage location is selected. If you already have the destination bucket or file share integrated with Morpheus, skip to the next section.

1. Navigate to Infrastructure > Storage
2. Click on the Buckets or File Shares tab depending on your chosen storage type
3. Click + *ADD*
4. Select the appropriate bucket or file share type
5. Complete the required fields and click *SAVE CHANGES*

Note: Additional guidance on integrating each of the supported bucket and file share types can be found elsewhere in Morpheus documentation.

Configuring Morpheus Appliance Backup

With the groundwork laid in the previous sections, we're ready to enable and configure Morpheus appliance backup.

1. Navigate to Administration > Backups
2. Slide the switch labeled "Backup Appliance"
3. Click *SAVE*

On saving this change, a text link labeled "Backup" will be activated which will take you directly to the automatically-generated appliance backup job. Click this link to continue.

1. Click *EDIT*
2. Enter a name for the appliance backup job
3. Select an integrated storage bucket or file share
4. Choose a pre-created backup job. If you do not have an existing backup job that fits, a retention count and schedule can be manually created in this modal. If you manually configure retention counts and schedules in addition to associating a Job, the Job values will override any manual settings.
5. Click *SAVE CHANGES*

At this point, your appliance will be automatically backed up on the schedule you chose and stored in the selected location. An appliance backup will store backup copies of the appliance MySQL database. Should you need to restore or migrate your database from backup, follow the steps in the next section of this guide.

Restoring an Appliance from Backup

Begin by ensuring the Morpheus UI service is stopped on all of the application servers:

```
[root@app-server-new ~] morpheus-ctl stop morpheus-ui
```

To access the MySQL shell we will need the password for the Morpheus DB user. We can find this in the morpheus-secrets file:

```
[root@app-server-old ~] cat /etc/morpheus/morpheus-secrets.json | grep morpheus_
↪password
"morpheus_password": "451e122cr5d122asw3de5e1b", <---- this one
"morpheus_password": "9b5vdj4de5awf87d",
```

Make note of the first morpheus_password value as indicated above.

Copy the SQL database backup from the backup bucket or file share to an appliance node at /tmp/morpheus_backup.sql. Then, you can import the MySQL dump into the target database using the embedded MySQL binaries, specifying the database host, and entering the password for the morpheus user when prompted:

```
[root@app-server-new ~] /opt/morpheus/embedded/mysql/bin/mysql -u morpheus -h 10.1.2.
↪2 morpheus -p < /tmp/morpheus_backup.sql
Enter password:
```

The data from the old appliance is now replicated on the new appliance. Simply start the UI to complete the process:

```
[root@app-server-new ~] morpheus-ctl start morpheus-ui
```

Cloud Resource Tagging with Morpheus

Introduction

As organizations scale their cloud environments, they often need to devise methodologies for organizing those resources. Tags consist of key and optional value pairs which make it easier to search for or filter your cloud resources based on categories relevant to the organization. While each public or private cloud handles tagging slightly differently, Morpheus removes that complexity and differentiation by handling tags in a consistent way across clouds. In addition, the Morpheus policy engine gives administrators tools to set up guard rails and ensure cloud resource tagging is handled consistently with each provisioning.

This guide will go through Morpheus tagging features and best practices, as well as include provisioning examples from some of the most commonly integrated clouds.

Note: Tag syncing is bi-directional in Morpheus for supported clouds. Tag syncing is currently supported in Amazon, Azure, Google, VMware, and Alibaba Clouds.

Tagging on Provisioning

In the simplest use case, tags can be entered manually in most provisioning scenarios. For example, on the Configure tab of the Create Instance wizard, the user can expand the Tags section and enter as many Tags as he or she needs to comply with organization practices.

CREATE INSTANCE [X]

TYPE > GROUP > **CONFIGURE** > AUTOMATION > REVIEW

Configuration Options

LAYOUT: Amazon VM

PLAN: Amazon T2 Micro - 1 Core, 1GB Memory
Cores: 1 Memory: 1 GB Price: \$8.4912 / Month

RESOURCE POOL: [Select]

VOLUMES: root | 20 | GB | gp2 | +

NETWORKS: [Select] | DHCP | +

SECURITY GROUPS: default | +

PUBLIC IP: Subnet Default

▸ User Config

▸ Network Options

▸ Advanced Options

▼ Tags

TAGS	Application	App1	
			🗑️ +

▸ Environment

PREVIOUS NEXT

Once the resource is deployed, Tags are synced and applied to the provisioned machine in the relevant cloud. Tags are shown on the Instance detail page in Morpheus and can also be confirmed in the cloud console if desired. Tag syncing is also a two-way street in that any tag updates applied within Morpheus or within the cloud console will be reflected everywhere.

Note: Tag updates made in the cloud console may take up to five minutes to be reflected in Morpheus UI following the next sync of cloud data.

Custom Instance Types and Tagging

Manually tagging resources as described in the previous section will work in some cases but many administrators will likely need to pre-seed the provisioning wizards with tagging prompts or build dropdown lists for tag values. This is accomplished in Morpheus through the use of custom Instance Types, Option Types, and Option Lists.

To get started, first create an Option List to hold dropdown or typeahead sets of available key values if needed based on organizational tagging policy. If the tag value field should be manually entered at provision time or if the value field is to be left blank, this step can be skipped. Option Lists are created and stored in the Morpheus Library (Provisioning > Library). They can be populated manually by entering CSV or JSON datasets as shown in the example below. They can also be dynamically populated through Morpheus API or REST calls.

EDIT OPTION LIST

×

NAME

Tag Application

DESCRIPTION

TYPE

Manual

▼

DATASET

Create an initial JSON or CSV dataset to be used as the collection for this option list. It should be a list containing objects with properties 'name', and 'value'.

'App1' , 'App1 '

'App2' , 'App2 '

'App3' , 'App3 '

'App4' , 'App4 '

'App5' , 'App5 '

SAVE CHANGES

With the set of possible values defined (if needed), we next create an Option Type to prompt the user for tag input when provisioning relevant Instance types. Option Types are also housed in the Morpheus Library (Provisioning > Library).

It's important to note the value entered for the FIELD NAME on the new option type will be set as the tag key. The EXPORT AS TAG box should also be marked. By default, the TYPE value is Text. This is appropriate when the user should be prompted with a free text field at provision time to enter a tag value. To tie this Option Type to the Option List that was just created (if needed), change the TYPE value to Select List or Typeahead. Typeahead works best for very long lists while Select List is often a better user experience for lists of a more manageable size. Set the Option List we created in the previous step in the OPTION LIST value (if needed).

EDIT OPTION TYPE

×

NAME

*Tag Application

DESCRIPTION

FIELD NAME

application

This is the input fieldName property that the value gets assigned to.

☒ EXPORT AS TAG

TYPE

Select List

LABEL

Tag Application

This is the input label that shows typically to the left of a custom option.

DEFAULT VALUE

☒ REQUIRED

OPTION LIST

Tag Application

DEPENDENT FIELD

A fieldName that will trigger reloading this option list

SAVE CHANGES

At this point, we are ready to add this Option Type to any custom Instance Types or Layouts. When those Instance Types or Layouts are provisioned, the values input by the user become tags associated with the created cloud resources. By setting the Option Type on an Instance Type, the tag selection appears when provisioning all associated Layouts. Alternatively, if the Option Type is set on individual Layouts, it will only appear when those Layouts are provisioned.

Instance Types and Layouts are also stored in the Morpheus Library (Provisioning > Library). By opening up any custom Instance Type, we can add the Option Type we just created when editing the Instance Type. Additionally, we can drill into associated Layouts and apply the Option Type to selected Layouts if that's more appropriate.

EDIT INSTANCE TYPE

×

NAME

*Harker CentOS

CODE

harkerCent

Useful shortcode for provisioning naming schemes and export reference.

DESCRIPTION

CATEGORY

OS

▼

ICON

Browse

Suggested Dimensions: 150 x 51

Option Types

Search option types

*Tag Application (select)	≡	×
*Tag Compliance (text)	≡	×
*Tag Cost Center (text)	≡	×
*Tag Department (text)	≡	×
*Tag Environment (text)	≡	×
*Tag Owner (text)	≡	×

Going forward, each time the chosen Instance Type or Layout is provisioned, the user will be prompted to enter

relevant tag selections. We can even require the user input these values or govern their inputs through tagging policies which will be discussed in the next section.

Instituting Tagging Policies

If needed, Morpheus allows cloud resource tagging to be governed through its native policy engine. Like other policies, tag policies are added from Administration > Policies. By creating a new policy and setting the TYPE to Tags, the relevant fields are revealed.

Note: At the time of this writing (Morpheus 4.2.1), tag policy scanning and enforcement is only functional in Azure, Amazon AWS, VMware, and Google Cloud Platform clouds.

With a tag policy, we can choose to enforce the policy on a strict or passive basis by marking or unmarking the STRICT ENFORCEMENT box. Strictly enforced tagging policies will not allow provisioning to proceed in supported clouds if the policy requirements are not met. If we opt to enforce the policy passively, a warning banner will appear on the detail page of any server that does not meet policy requirements. Additionally, existing servers in supported clouds will be scanned and those which do not meet policy requirements will also receive the warning banner.

A tag policy must be given a KEY value. If we define only a KEY value, the policy will look for a tag with that key and any (or no) value. Alternatively, we can select any pre-existing Option List as the VALUE LIST to require the tag contain a value that exists in that list.

Finally, like other Morpheus policies, we can choose to scope it globally, by group, by cloud, or by user. Master Tenant administrators can also choose to scope the policy to one or more Subtenants.

NEW POLICY ✕

TYPE

Tags

NAME

DESCRIPTION

☒ **ENABLED**

Tag Policy scanning and enforcement is only currently functional for three cloud types. Azure, Amazon, and VMware.

Config

STRICT ENFORCEMENT

KEY

Tag Key

VALUE LIST

Select

Choose from any configured Option List

Filter

SCOPE

Global

SAVE CHANGES

Tagging in Action

With the prep work complete, we can take a look at our Option Types in action at provision time. In this example case, several Option Types have been created and applied to one custom Instance Type. The example Instance Type has three associated CentOS Layouts, one for AWS, one for VMware, and one for Azure. Regardless of the selected Layout, users are prompted to fill the same tag fields and our tagging remains consistent regardless of the user who is provisioning a new resource at the time.

Tagging and AWS

When provisioning my CentOS Instance Type with an Amazon Layout, the tag prompts are shown in the provisioning wizard.

CREATE INSTANCE [X]

TYPE > GROUP > **CONFIGURE** > AUTOMATION > REVIEW

Configuration Options

LAYOUT:

PLAN:
Cores: 1 Memory: 1 GB Price: \$8.4912 / Month

RESOURCE POOL:

VOLUMES: +

NETWORKS: DHCP +

SECURITY GROUPS: +

TAG APPLICATION:

COMPLIANCE:

COST CENTER:

DEPARTMENT:

ENVIRONMENT:

OWNER:

PUBLIC IP:

► User Config

In the AWS web console, we can see the same tags are applied. We also have two-way tag sync going forward. When tags are updated in Morpheus, the changed is synced to the AWS web console. The opposite is also true.

Description

Status Checks

Monitoring

Tags

Add/Edit Tags

Key	Value	
Morpheus ID	64748	Show Column
Morpheus Instance ID	6488	Show Column
Morpheus Server ID	15749	Show Column
Tags	an instance named Cent 1001	Show Column
application	App1	Show Column
compliance	PCI	Show Column
costCenter	10001	Show Column
department	Security	Show Column
environment	Dev	Show Column
owner	John User	Show Column

Tagging and VMware

When provisioning my CentOS Instance Type with a VMware Layout, the tag prompts are shown in the provisioning wizard.

CREATE INSTANCE

TYPE

GROUP

CONFIGURE

AUTOMATION

REVIEW

Configuration Options

LAYOUT

vmware_cent

PLAN

1 CPU, 1GB Memory

Cores: 1 Memory: 1 GB Price: \$5.726 / Month

RESOURCE POOL

VOLUMES

root

10

GB

Auto - Cluster

+

NETWORKS

Select Network

+

TAG APPLICATION

App1

COMPLIANCE

PCI

COST CENTER

10001

DEPARTMENT

Security

ENVIRONMENT

Dev

OWNER

John User

HOST

Select

FOLDER

User Config

In the VMware console, we can see the same tags are applied. We also have two-way tag sync going forward. When tags are updated in Morpheus, the changed is synced to VMware. The opposite is also true.

Tagging and Azure

When provisioning my CentOS Instance Type with an Azure Layout, the tag prompts are shown in the provisioning wizard.

CREATE INSTANCE

TYPE > GROUP > **CONFIGURE** > AUTOMATION > REVIEW

Configuration Options

LAYOUT: azure_cent

PLAN: Basic_A1 (1 Core, 1.75GB Memory) (eastus)
Cores: 1 Memory: 1.8 GB Price: \$17.4556 / Month

RESOURCE POOL: Select

VOLUMES: root Page Blob blabsuseastsa +

NETWORKS: Select Network DHCP +

TAG APPLICATION: Select

COMPLIANCE: PCI

COST CENTER: 10001

DEPARTMENT: Security

ENVIRONMENT: Dev

OWNER: John User

SECURITY GROUP: Select

AVAILABILITY SET: No Availability Set

In the Azure console, we can see the same tags are applied. We also have two-way tag sync going forward. When tags are updated in Morpheus, the changed is synced to Azure. The opposite is also true.

Automation

Ansible

Overview

Ansible is a configuration management engine that is rapidly growing in popularity in the IT and DevOPS community. While it lacks some of the benefits at scale that solutions such as Salt, Chef, or Puppet offer. It is very easy to get started and allows engineers to develop tasks in a simplistic markup language known as YAML. Morpheus integrates with an existing repository of playbooks as the master in a master-slave Ansible architecture.

Morpheus not only supports Ansible but greatly enhances Ansible to do things that it could not do in its native form. For example, Ansible can now be configured to run over the Morpheus agent communication bus. This allows playbooks to be run against instances where SSH/WinRM access may not be feasible due to networking restrictions or other firewall constraints. Instead it can run over the Morpheus Agent which only requires port 443 access back to the Morpheus appliance URL.

This integration supports both Linux-based and Windows platforms for playbook execution and can also be configured to query secrets from Morpheus Cypher services (similar to Vault).

Requirements

- Minimum Ansible Version Requirement is 2.7.x
- For agentless non-commandbus, sshpass is required
- For Windows non-agent command bus, pywinrm is required
- **Integrations:** Ansible User Role Permission required for access to Ansible Details Pages and Ansible tabs in Groups and Clouds

Note: Installing Ansible on the Morpheus appliance is a requirement. In most cases, this is handled automatically but in certain situations you may have to install manually. See the section below on [troubleshooting Ansible](#) for installation steps.

Add Ansible Integration

1. Navigate to *Provisioning > Automation > Integrations* and select + *New Integration*
2. Select Integration Type “Ansible”
3. Populate the following fields:

Name Name of the Ansible Integration in Morpheus

Enabled Enabled by default

Ansible Git URL https or git url format of the Ansible Git repo to use

Keypair For private Git repos, a keypair must be added to Morpheus and the public key added to the git account.

Playbooks Path Path of the Playbooks relative to the Git url.

Roles Path Path of the Roles relative to the Git url.

Group Variable Path Path of the Group Variables relative to the Git url.

Host Variables Path Path of the Host Variables relative to the Git url.

Use Ansible Galaxy Install roles defined in `requirements.yml`

Enable Verbose Logging Enable to output verbose logging for Ansible task history

Use Morpheus Agent Command Bus Enable for Ansible Playbooks to be executed via Morpheus Agent Command Bus instead of SSH

4. Save Changes

Once you have completed this section and saved your changes you can set up a Cloud or Group to utilize this integration.

Ansible on Windows

When executing Ansible playbooks on Windows platforms, a few requirements must be met:

- `pywinrm` may need to be installed on the Morpheus Appliance via `pip install pywinrm`
- An Ansible Integration must be scoped to a Group or Cloud for Ansible to execute on Windows, as Morpheus assumes Ansible local when no group or cloud is scoped to Ansible. The playbooks do not need to be executed solely in the Group or Cloud, one just needs to be scoped to an Ansible Integration for Ansible Windows to run properly.

Scope Ansible Integration to a Cloud

1. Navigate to *Infrastructure -> Clouds*
2. Edit the target Cloud
3. Expand the *Advanced Options* section
4. In the *Config Management* dropdown, select the Ansible Integration.
5. Save Changes

Once an Ansible integration is added to a Cloud, a new “ANSIBLE” tab will appear on the Cloud details page, populated with the Ansible integrations Playbook and Roles, as well as an editable Inventory list.

Scope Ansible Integration to a Group

1. Navigate to *Infrastructure -> Groups*
2. Edit the target Group
3. Expand the *Advanced Options* section
4. In the *Config Management* dropdown, select the Ansible Integration.
5. Save Changes

Once an Ansible integration is added to a Group, a new “ANSIBLE” tab will appear on the Group details page, populated with the Ansible integrations Playbook and Roles, as well as an editable Inventory list.

Provisioning Options

When provisioning Instances into a Cloud or Group with a Ansible Integration added, an *Ansible* section will appear in the Config section of the provisioning wizard. By default, Ansible is enabled, but can be disabled by expanding the *Ansible* section and unchecking *Enable Ansible*.

Ansible Integration Provisioning options:

Enable Ansible Select to bootstrap

Ansible Group Ansible Inventory Group. Use existing group or enter a new group name to create a new group. Leaving this field blank will place instance in the “unassigned” inventory group.

Note: An instance can belong to multiple groups by separating group names with a comma

Playbook Playbook(s) to run. The `.yaml` extension is optional.

Running Playbooks

Playbooks can also be run on all inventory groups, individual groups, or added as a task and ran with workflows.

To run Ansible on all or a single inventory group, in the Ansible tab of the Morpheus Group page, select the *Actions* dropdown and click *Run*.

In the *Run Ansible* modal, you can then select all or an individual group, and then all or a single Playbook, as well as add custom tags.

Playbook's can also be added as tasks to workflows in the *Provisioning -> Automation* section, and then selected in the Automation pane during provisioning of new instances, when creating app blueprints, or ran on existing instances using the *Actions -> Run Workflow* on the Instance or Host pages.

Using variables

Morpheus variables can be used in playbooks.

Use Case:

Create a user as instance hostname during provisioning.

Below is the playbook. Add this playbook to a task and run it as a workflow on the instance.

```
---
- name: Add a user
  hosts: all
  gather_facts: false
  tasks:
    - name: Add User
      win_user:
        name: "{{ morpheus['instance']['hostname'] }}"
        password: "xxxxxxx"
        state: present
```

Note: `{{ morpheus['instance']['hostname'] }}` is the format of using Morpheus Variables

Create a user with a name which you enter during provisioning using a custom Instance type.

This instance type has a *Text* Option type that provides a text box to enter a username. The fieldName of the option

```
---
- name: Add a user
  hosts: all
  gather_facts: false
  tasks:
    - name: Add User
      win_user:
        name: "{{ morpheus['customOptions']['username'] }}"
        password: "xxxxxxx"
        state: present
```

Note: `{{ morpheus['customOptions']['username'] }}` will be the format.

Using Secrets

Another great feature with using Ansible and Morpheus together is the built in support for utilizing some of the services that Morpheus exposes for automation. One of these great services is known as Cypher (please see documentation on [Cypher](#) for more details). Cypher allows one to store secret data in a highly encrypted way for future retrieval. Referencing keys stored in cypher in your playbooks is a matter of using a built-in lookup plugin for ansible.

```
- name: Add a user
  win_user:
    name: "myusername"
    password: "{{ lookup('cypher','secret=password/myusername') }}"
    state: present
```

By using the `{{ lookup('cypher','secret=password/myusername') }}` syntax. One can grab the value directly out of the key for use. This lookup plugin also supports a few other fancy shortcuts. In this above example the `password/` mountpoint is capable of autogenerating passwords if they have not previously been defined and storing them within cypher for reference later.

Another capability is accessing properties from within a key in cypher. The value of a key can also be a JSON object which can be referenced for properties within. For example:

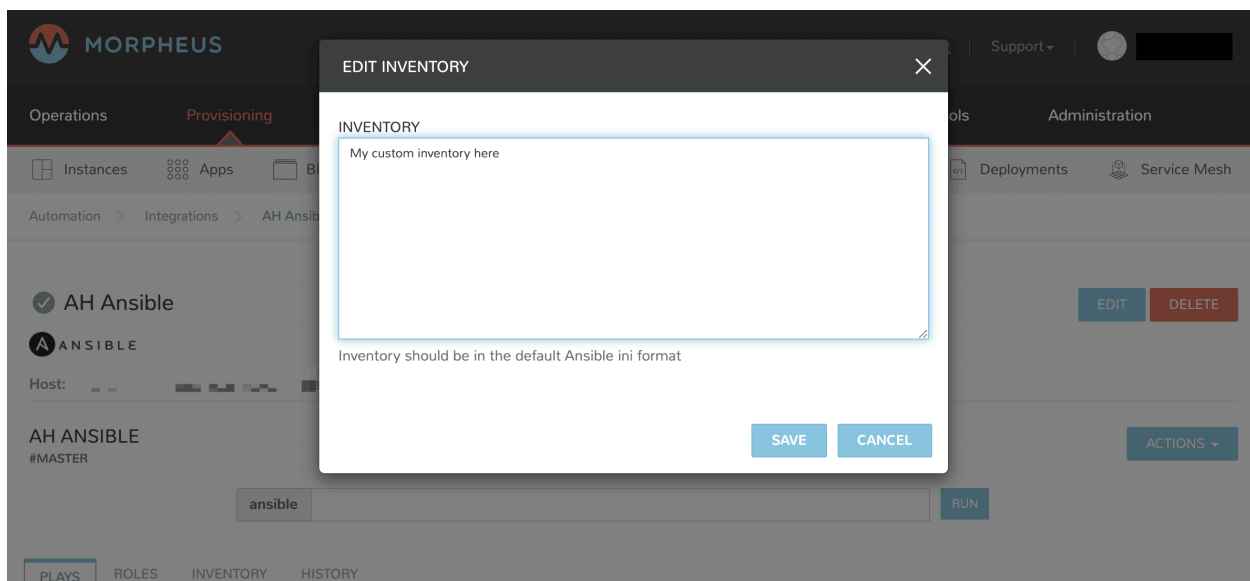
```
{{ lookup('cypher','secret=secret/myjsonobject:value') }}
```

This would grab the `value` property off the nested json data stored within the key.

Cypher is very powerful for storing these temporary or permanent secrets that one may need to orchestrate various tasks and workflows within Ansible.

Custom Inventory Entries

With Morpheus it is possible to add custom inventory entries that exist outside of morpheus host/server entry. This is global across cloud or group and is done on the integration details page of the Ansible integration. To add a custom inventory entry navigate to Provisioning > Automation > Integrations > (Your specific Ansible integration). Click on the ACTIONS button, then click EDIT INVENTORY. Inventory should be in the default Ansible ini format.



Using Ansible over the Morpheus Agent Command Bus

In many environments, there may be security restrictions on utilizing SSH or WinRM to run playbooks from an Ansible server on the appliance to a target machine. This could be due to being a customer network (in the environment of an MSP), or various security restrictions put in place by tighter industries (i.e. Government, Medical, Finance).

Ansible can get one in trouble in a hurry. It is limited in scalability due to its fundamental design decisions that seem to bypass concepts core to all other configuration management frameworks (i.e. SaltStack, Chef, and Puppet). Because of its lack of an agent, the Ansible execution binary itself has to handle all the load and logic of executing playbooks on all the machines in the inventory of an Ansible project. This differs from other tools where the workload is distributed across the agents of each vm. Because of this (reaching out) approach, Ansible is very easy to get started with, but can be quite a bit slower as well as harder to scale up. However, Morpheus offers some solutions to help mitigate these issues and increase scalability while, at the same time improving security.

How does the Morpheus Agent Command Bus Work?

One of the great things about Morpheus is its Agent Optional approach. This means that this functionality can work without the Agent, however the agent is what adds the security benefits being represented here. When an instance is provisioned (or converted to managed) within Morpheus, an agent can be installed. This agent opens a secure websocket back to the Morpheus appliance (over port 443). This agent is responsible for sending back logs, guest statistics, and a command bus for automation. Since it is a WebSocket, bidirectional communication is possible over a STOMP communication bus.

When this functionality is enabled on an Ansible integration, a *connection_plugin* is registered with Ansible of type *morpheus* and *morpheus_win*. These direct bash or powershell commands, in their raw form, from Ansible to run over a Morpheus api. The Ansible binary sends commands to be executed as an https request over the API utilizing a one time execution lease token that is sent to the Ansible binary. File transfers can also be enacted by this API interface. When Morpheus receives these commands, they are sent to the target instances agent to be executed. Once they have completed a response is sent back and updated on the *ExecutionRequest* within Morpheus. Ansible polls for the state and output on these requests and uses those as the response of the execution. This means Ansible needs zero knowledge of a machines target ip address, nor its credentials. These are all stored and safely encrypted within Morpheus.

It has also been pointed out that this execution bus is dramatically simpler than utilizing *pywinrm* when it comes to orchestrating Windows as the winrm configurations can be cumbersome to properly setup, especially in tightly secured Enterprise environments.

Using Ansible Galaxy

Morpheus can use a `requirements.yml` file to define Ansible roles to download prior to running your playbook. Place `requirements.yml` into the root of your Git repository and make sure *Use Ansible Galaxy* is checked in the integration. Roles will be installed in the root of the repository if a directory is not defined in *Roles Path*.

- Example requirements.yml:

```
- src: https://github.com/geerlingguy/ansible-role-java
  name: java
```

- Example playbook.yml:

```
- hosts: all
  gather_facts: true
  roles:
    - java
```

Troubleshooting Ansible

- When a workflow is executed manually, the Ansible run output is available in the Instance History tab. Select the **i** bubble next to the Ansible task to see the output. You can also see the run output in the ui logs in `/var/log/morpheus/morpheus-ui/current` which can be tailed by running `morpheus-ctl tail morpheus-ui`.
- Verify Ansible is installed on the Morpheus Appliance.

Ansible should be automatically installed but certain OS or network conditions can prevent the automated install. You can confirm installation by running `ansible --version` in the Morpheus appliance, or by viewing the Ansible integration details page (Administration > Integrations > Select Ansible Integration). We also see it in the Ansible tab of a Group or Cloud scoped to Ansible, just run `--version` as ansible is already included in the command.

If Ansible is not installed, follow these instructions to install, or use your preferred installation method:

Ubuntu:

```
sudo apt-get install software-properties-common
sudo apt-add-repository ppa:ansible/ansible
sudo apt-get update
sudo apt-get install ansible
```

CentOS:

```
sudo yum install epel-release
sudo yum install ansible
```

Then create the working Ansible directory for Morpheus:

```
sudo mkdir /opt/morpheus/.local/.ansible
sudo chown morpheus-local.morpheus-local /opt/morpheus/.local/.ansible
```

- Validate the git repo is authorizing and the paths are configured correctly.
The public and private ssh keys need to be added to the Morpheus appliance via “Infrastructure -> Keys & Certs” and the public key needs to be added to the git repo via user settings. If both are set up right, you will see the playbooks and roles populate in the Ansible Integration details page.
- The Git Ref field on playbook tasks is to specify a different git branch than default. It can be left to use the default branch. If your playbooks are in a different branch you can add the branch name in the Git Ref field.
- When running a playbook that is in a workflow, the additional playbooks fields do not need to be populated, they are for running a different playbook than the one set in the Ansible task in the Workflow, or using a different Git Ref.
- If you are manually running Workflows with Ansible tasks on existing Instances through *Actions -> Run Workflow* and not seeing results, set the Provision Phase on the Ansible task to Provision as there may be issues with executing tasks on other phases when executing manually.

Ansible Tower

Overview

Morpheus supports Ansible Tower for configuration management. Morpheus accomplishes this by integrating with an existing instance running Ansible Tower (AT) 3.3.0-1 and earlier. The username and password required for integration can be a user with admin access or a user with project admin access. Morpheus will import the current Inventory, Templates, Hosts, Groups and Projects. In the integration view it will add a Job tab which will have information of all the jobs executed from Morpheus. Note: It will not import data of the jobs which are not executed from Morpheus.

Add Ansible Tower Integration

1. Navigate to *Administration* -> *Integrations* and select + *New Integration*
2. Select Integration Type “Ansible Tower”
3. Populate the following fields:
 - Name: Name of the Ansible Tower Integration in Morpheus
 - Enabled: Enabled by default it is enabled. To disable the integration, uncheck this option and save.
 - Ansible Tower URL: This would be an https or http Ansible tower url.
 - Username: The user morpheus would use to communicate with Ansible Tower.
 - Password: Enter the password. Password is encrypted and saved in DB.
 - API Version: This drop down has one option v2 for now but may have others in future.
4. Save Changes

Once you have completed this section and saved your changes you can set up a Cloud or Group to utilize this integration.

Scope Ansible Tower Integration to a Cloud

All instances provisioned in this cloud will have the Ansible Tower config option during provisioning. See below the Provisioning Options for more details about the options.

1. Navigate to *Infrastructure* -> *Clouds*
2. Edit the target Cloud
3. Expand the *Advanced Options* section
4. In the *Config Management* dropdown, select the Ansible Tower Integration.
5. Save Changes

Scope Ansible Tower Integration to a Group

All instances provisioned in this Group will have the Ansible Tower config option during provisioning in any cloud part of the Group. See below the Provisioning Options for more details about the options.

1. Navigate to *Infrastructure -> Groups*
2. Edit the target Group
3. Expand the *Advanced Options* section
4. In the *Config Management* dropdown, select the Ansible Tower Integration.
5. Save Changes

Provisioning Options

When provisioning Instances into a Cloud or Group with a Ansible Tower Integration added, an *Ansible Tower* section will appear in the Config section of the provisioning wizard. By default, Ansible Tower is enabled, but can be disabled by expanding the *Ansible Tower* section and unchecking *Enable Ansible Tower*.

Ansible Integration Provisioning options:

Enable Ansible Tower Select to bootstrap

Inventory A list of Inventory available in Ansible Tower will appear in the drop down. Select an existing inventory. The instance will be added to the inventory selected.

Ansible Group Enter the name of an existing Group in the inventory selected above.

Template

Select an existing template or select the option 'Create New Template'. If 'Create New Template' is selected below fields w

Template Name Enter the template name

Project Select an existing project from the drop down options

Playbook Select a playbook from the dropdown to be associated with the template. Note: Morpheus doesn't store a local copy of the playbooks visible in Ansible Tower. SCM or local path for playbooks should be maintained in Ansible Tower.

Execute Mode

Select one of the options from the dropdown

Limit to instance This will execute the template on the instance provisioned.

Limit to Group This will execute the template on all hosts attached to the group entered in the 'Ansible Group' field.

Run for all This will execute the template on all hosts in the inventory

Skip execution This will skip the execution of the template on the instance provisioned.

Use Case

You have Job template(s) in Ansible Tower to do post build config after the OS is deployed. The playbook with roles and tasks to do post build will add specific users and groups, install required packages, remove packages, disable services, change config for ntp, resolv, hosts etc. You want to add the instance to an existing Group/Inventory in Tower.

You can achieve this by adding the Ansible Tower Integration and then scope it to a Cloud or Group. While provisioning an instance, in the config stage you have the Ansible Tower section with option to select the post build job template, select the Inventory and provide an existing Group Name or if the Group doesn't exist Morpheus will create it and submit for provisioning.

Morpheus will provision the instance, once it is in the finalize state where the instance has an ip and has completed domain join if required, added user(s) or User Groups if specified then Morpheus will add the instance to the inventory and Group and run the Template which will do all the post build of the server.

The output of the post build template execution can be seen under Instance history.

Chef

Overview

Morpheus integrates with one or multiple Chef servers to be used for bootstrapping while provisioning or as tasks in workflows in the Automation section. These workflows can then be run during provisioning in the provisioning wizard Automation pane, or on an existing instance by selecting Actions->Run Workflow. Workflows can also be added to instances in the blueprint and app sections.

Add Chef Integration

1. Navigate to *Administration -> Integrations* and select + *New Integration*
2. Select Integration Type "Chef"
3. Populate the following fields:
 - Name: Name of the Chef Integration in Morpheus
 - Chef Endpoint: url of chef server api endpoint in <https://api.example.com> format. Do not add /organization/xxxx here, which is populated in the Chef Organization field
 - Chef Version: 12.3.0 by default, can be changed to use a different/more recent version of chef
 - Chef Organization: Chef Server Organization
 - Chef User: Chef Server User
 - User Private Key: The private key of the user with access to this chef server
 - Organization Validator: Validator key for the organization
4. Save Changes

The added Chef Integration is now available for use in Morpheus. The Chef Integration can be added to Clouds or Groups to auto-bootstrap nodes and specify Environment, Node ID, Runlist, Attributes and Tags when creating instances. The Chef integration can also be selected in the Chef Server dropdown when creating a Chef Bootstrap type task.

Scope Chef Integration to a Cloud

1. Navigate to *Infrastructure -> Clouds*
2. Edit the target Cloud
3. Expand the *Advanced Options* section
4. In the *Config Management* dropdown, select the Chef Integration.
5. Save Changes

Scope Chef Integration to a Group

1. Navigate to *Infrastructure -> Groups*
2. Edit the target Group
3. Expand the *Advanced Options* section
4. In the *Config Management* dropdown, select the Chef Integration.
5. Save Changes

Provisioning Options

When provisioning Instances into a Cloud or Group with a Chef Integration added, a *Chef* section will appear in the Config section of the provisioning wizard. By default, Chef is enabled, but can be disabled by expanding the *Chef* section and unchecking *Enable Chef*.

Chef Integration Provisioning options:

Enable Chef Select to bootstrap

Chef Environment Populate Chef environment, or leave as *_default*

Chef Node ID Defaults to instance name, configurable.

Chef Runlist Add Runlist

CHEF ATTRIBUTES Add Chef Attributes

CHEF TAGS Add Chef tags

Puppet

Add Puppet Integration

1. Navigate to Administration > Integrations
2. Click + *NEW INTEGRATION*
3. Select Integration type “Puppet”
4. Populate the following fields
 - Name: A friendly name for this Puppet integration in Morpheus
 - Puppet Master (Hostname): The resolvable DNS name to the Puppet Master, communicating on port 443 by default

- Allow Immediate Execution: Yes or No

5. Click *SAVE CHANGES*

Salt

Overview

Morpheus integrates with an existing Salt Master for seamless deployment of Salt States to Minions provisioned from Morpheus .

Add Salt Integration

To get started browse to Admin -> Integrations from within Morpheus .

Once there simply add a New Integration

	NAME	TYPE		
	Labs Salt Master	saltMaster		
	morpheus-ansible	ansible		
	Labs Puppet Master	puppet		
	Labs Private	docker.registry		
	SNOW Approvals	serviceNow		
	Labs Chef	chef		
	Consul	consulRegistry		
	Labs Git A	git		
	My Private Docker Repo	docker.registry		
	Labs Salt Master 2	saltMaster		

And then scope the integration to your existing Salt Master by ip address. Make sure that the username entered is one with proper escalation privileges for running Salt, and point the Working Directory at the directory on your Master where your States live.

Note: Morpheus will allow you to run States from a git backend, but in v2.10 you will not see states from a git backend within Morpheus

INTEGRATION

TYPE

Salt Master

NAME

Labs Salt Master

SALT MASTER

192.168.162.51

SSH PORT

22

USERNAME

morpheus

PASSWORD

.....

KEY PAIR

Select

WORKING DIRECTORY

/srv/salt

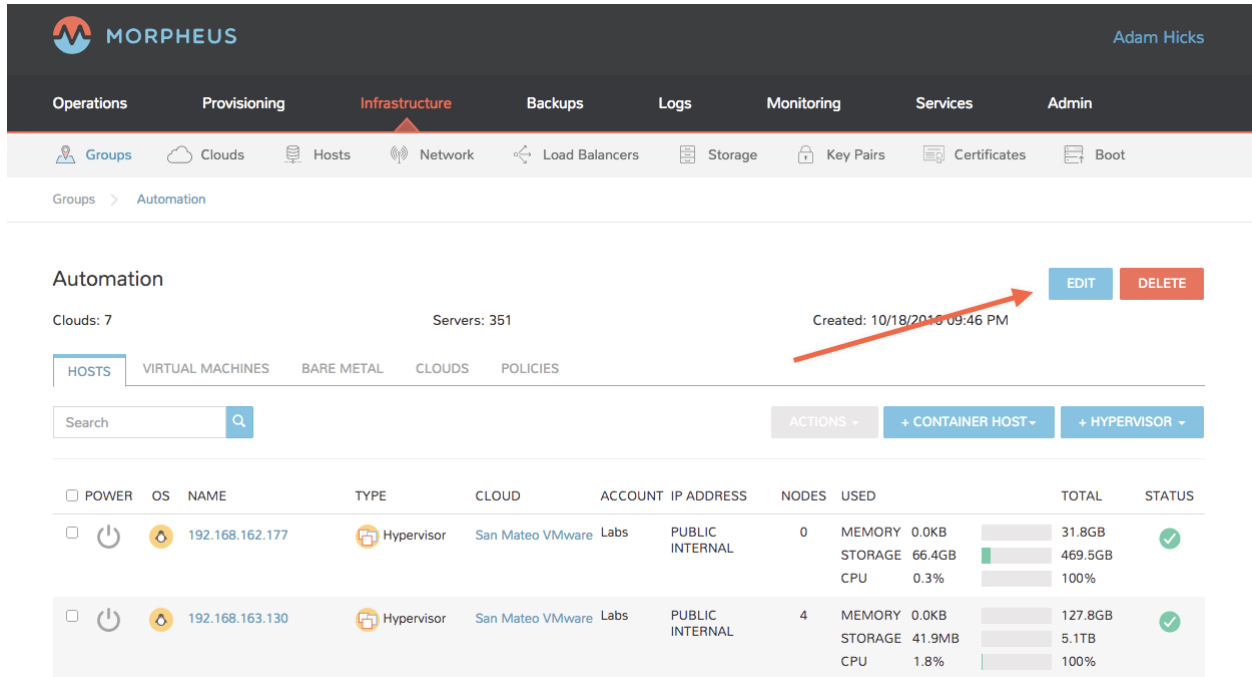
SALT VERSION

SAVE CHANGES

Scope Salt Integration to Group Or Cloud

Configuration Management integrations like Saltstack apply to the Infrastructure Group abstraction in Morpheus . To tie yours in, browse to Infrastructure -> Groups in Morpheus and select the group that you would like to tie to your Salt Master.

From here select *Edit*



The screenshot shows the Morpheus web interface. The top navigation bar includes 'Operations', 'Provisioning', 'Infrastructure' (selected), 'Backups', 'Logs', 'Monitoring', 'Services', and 'Admin'. Below this is a sub-navigation bar with 'Groups', 'Clouds', 'Hosts', 'Network', 'Load Balancers', 'Storage', 'Key Pairs', 'Certificates', and 'Boot'. The main content area is titled 'Automation' and shows 'Clouds: 7' and 'Servers: 351'. A red arrow points to the 'EDIT' button for the first host in the list.

POWER	OS	NAME	TYPE	CLOUD	ACCOUNT	IP ADDRESS	NODES	USED	TOTAL	STATUS
<input type="checkbox"/>		192.168.162.177	Hypervisor	San Mateo VMware Labs	PUBLIC INTERNAL	0	MEMORY 0.0KB STORAGE 66.4GB CPU 0.3%	31.8GB 469.5GB 100%		
<input type="checkbox"/>		192.168.163.130	Hypervisor	San Mateo VMware Labs	PUBLIC INTERNAL	4	MEMORY 0.0KB STORAGE 41.9MB CPU 1.8%	127.8GB 5.1TB 100%		

And from the options toggle Advanced Options and select your Saltstack integration in the Config Management drop-down.

EDIT GROUP

×

Configuration

NAME

Automation

CODE

LOCATION

▼ Advanced Options

SERVICE REGISTRY

Please Select

CONFIG

MANAGEMENT

Please Select

✓ Labs Salt Master

morpheus-ansible

Labs Puppet Master

Labs Chef

Labs Salt Master 2

SAVE CHANGES

After a page refresh you should see your Saltstack tab in your group page

MORPHEUS

Adam Hicks

Operations

Provisioning

Infrastructure

Backups

Logs

Monitoring

Services

Admin

Groups

Clouds

Hosts

Network

Load Balancers

Storage

Key Pairs

Certificates

Boot

Groups > Automation

Automation

Clouds: 7

Servers: 351

Created: 10/18/2016 09:46 PM

EDIT

DELETE

HOSTS

VIRTUAL MACHINES

BARE METAL

CLOUDS

SALTSTACK

POLICIES

Search

ACTIONS +

+ CONTAINER HOST +

+ HYPERVISOR +

POWER	OS	NAME	TYPE	CLOUD	ACCOUNT	IP ADDRESS	NODES	USED	TOTAL	STATUS
<input type="checkbox"/>		192.168.162.177	Hypervisor	San Mateo VMware Labs		PUBLIC INTERNAL	0	MEMORY 0.0KB STORAGE 66.4GB CPU 0 3%	31.8GB 469.5GB 100%	

Clicking on it will reveal a page that includes:

1. An interface to run Salt Master commands
2. Parsed Top File
3. Available States

Automation

Clouds: 7 Servers: 351 Created: 10/18/2016 09:46 PM

HOSTS VIRTUAL MACHINES BARE METAL CLOUDS **SALTSTACK** POLICIES

LABS SALT MASTER
@ 192.168.162.51

ACTIONS ▾

salt

RUN

▼ Top File

```
base:
  '*':
    - git
    - cmatrix

'web*':
  - apache

'web*2':
  - webserver
```

▼ States

top.sls

webserver.sls

The classic example of running

```
salt '*' test.ping
```

will return empty unless there are existing Minions with accepted keys on the Master. However, provisioning Minions via Morpheus is extremely easy.

Provisioning with Saltstack

To do so, provision as usual and Instances within the Group tied to the Saltstack Integration will now show additional options on the Configure pane

CREATE INSTANCE

TYPE

GROUP

CONFIGURE

AUTOMATION

REVIEW

Configuration Options

VERSION

14.04

INSTANCE CONFIGURATION

VMware VM

PLAN

Memory: 512MB Storage: 10GB

VOLUMES

root

10

GB

SCSI 0:0

Auto - Cluster

+

NETWORKS

Select Network

E1000

+

PLAN PRICE

\$0.00 / Month

RESOURCE POOL

Resources

PUBLIC KEY

Select

► Create User (optional)

▼ Advanced Options

STORAGE CONTROLLERS

Type: IDE

Bus: 0

+

Type: IDE

Bus: 1

Type: SCSI LSI Logic Parallel

Bus: 0

DOMAIN NAME

CUSTOMIZATION SPEC

Select

HOSTNAME

webserver

MINION ID

STATE

PREV

NEXT

Minion ID defaults to the hostname, and a State can be applied directly at provision time.

Note: Only States served from the Master's Working Directory can be applied at provision, not States from a git backend

Once your instance is provisioned and key negotiation has completed you will be able to access it and run commands via the integrated Salt command center in your Group.

Automation

Clouds: 7 Servers: 359 Created: 10/18/2016 09:46 PM

EDITDELETE

HOSTS VIRTUAL MACHINES BARE METAL CLOUDS SALTSTACK POLICIES

LABS SALT MASTER

@ 192.168.162.51

ACTIONS ▾

salt

test.ping

RUN

RUN RESULTS

```
[WARNING ] The file_roots parameter is not properly formatted, using defaults
webserver:
  True
```

▼ Top File

```
base:
  test:
    - test
```

If you did not apply a state at provision time now you will be able to run State commands through Morpheus .

LABS SALT MASTER

@ 192.168.162.51

ACTIONS ▾

salt '*' state.highstate

RUN

RUN RESULTS

[WARNING] The file_roots parameter is not properly formatted, using defaults
webserver:

```

ID: GIT software
Function: pkg.installed
Result: True
Comment: The following packages were installed/updated: python-git
         The following packages were already installed: git
Started: 18:52:07.812386
Duration: 23912.153 ms
Changes:

```

```

python-async:

```

```

  new:
    0.6.1-1
  old:

```

```

python-git:

```

```

  new:
    0.3.2aRC1-1

```

▼ Top File

In our example the Apache State from a git backend was applied successfully to our newly created vm.

192.168.163.86

Apache2 Ubuntu Default Page

ubuntu

It works!

This is the default welcome page used to test the correct operation of the Apache2 server after installation on Ubuntu systems. It is based on the equivalent page on Debian, from which the Ubuntu Apache packaging is derived. If you can read this page, it means that the Apache HTTP server installed at this site is working properly. You should **replace this file** (located at `/var/www/html/index.html`) before continuing to operate your HTTP server.

If you are a normal user of this web site and don't know what this page is about, this probably means that the site is currently unavailable due to maintenance. If the problem persists, please contact the site's administrator.

Configuration Overview

Ubuntu's Apache2 default configuration is different from the upstream default configuration, and split into several files optimized for interaction with Ubuntu tools. The configuration system is **fully documented in `/usr/share/doc/apache2/README.Debian.gz`**. Refer to this for the full documentation. Documentation for the web server itself can be found by accessing the **manual** if the `apache2-doc` package was installed on this server.

The configuration layout for an Apache2 web server installation on Ubuntu systems is as follows:

```

/etc/apache2/
|-- apache2.conf
|   |-- ports.conf
|-- mods-enabled
|   |-- *.load
|   |-- *.conf
|-- conf-enabled
|   |-- *.conf
|-- sites-enabled
|   |-- *.conf

```

- `apache2.conf` is the main configuration file. It puts the pieces together by including all remaining configuration files when starting up the web server.
- `ports.conf` is always included from the main configuration file. It is used to determine the listening ports for incoming connections, and this file can be customized anytime.
- Configuration files in the `mods-enabled/`, `conf-enabled/` and `sites-enabled/` directories contain particular configuration snippets which manage modules, global configuration fragments, or

Terraform

Requirements

Role Access

- In order to see the Terraform Blueprint type option and create Terraform App Blueprints in *Provisioning > Blueprints*, the Morpheus user must have Role permissions for *Provisioning: Blueprints - Terraform* set to *Full*.
- In order to provision Terraform Apps in *Provisioning > Apps*, the Morpheus user must have Role permissions for *Provisioning: Blueprints > Terraform* set to *Provision* or *Full*.
- Existing Terraform Blueprints must be added before they can be provisioned from *Provisioning > Apps*.

Github/Git Repo

- To use .tf files from a Git repo a Git or Github integration needs to be configured in *Administration > Integrations*. If one is not configured .tf or .tf.json files can be manually added to Terraform App Blueprints.

Supported App Provisioning Targets

- VMware
- Amazon AWS
- Microsoft Azure
- Oracle Cloud

Note: Additional clouds will be available in later releases.

Terraform Installation

Morpheus will automatically install Terraform locally upon the first Terraform App provision. It is possible on some operating system configurations for the automated terraform installation to fail, in which case it can be manually installed (run `terraform --version` to verify).

To manually install and configure terraform on the Morpheus Appliance:

1. Run the following curl on the Morpheus Appliance to install Terraform:

```
curl -k -s "https://applianceServerUrl/api/server-script/terraform-install?  
↪local=true" | bash
```

Note: Replace `applianceServerUrl` with your Morpheus appliance url or ip.

2. Create a working directory for Terraform, and change owner to `morpheus-app`.


```
sudo mkdir /var/opt/morpheus/morpheus-ui/terraform
sudo chown morpheus-app.morpheus-app /var/opt/morpheus/morpheus-ui/terraform
```

The default location is `/var/opt/morpheus/morpheus-ui/terraform` but can be changed.

3. Add the Terraform working path to `/opt/morpheus/conf/application.yml`

```
sudo vi /opt/morpheus/conf/application.yml
```

Add the following to the `application.yml` config below and in-line with the `repo` section:

```
terraform:
  location: '/var/opt/morpheus/morpheus-ui/terraform'
```

Example `application.yml` config with Terraform location added:

```
repo:
  git:
    location: '/var/opt/morpheus/morpheus-ui/repo/git'
  local:
    location: '/var/opt/morpheus/morpheus-ui/repo/local'
terraform:
  location: '/var/opt/morpheus/morpheus-ui/terraform'
bitcan:
  backup:
    destination:
      root: '/var/opt/morpheus/bitcan/backup'
      working: '/var/opt/morpheus/bitcan/working'
```

Important: Uses spaces not tabs to indent or ui startup will fail. If you used a different path than the default location, enter that path instead.

4. Restart the morpheus-ui to apply the `application.yml` config.

```
sudo morpheus-ctl restart morpheus-ui
```

Terraform is now installed and configured, and Terraform apps can be provisioned from Morpheus.

Creating Terraform App Blueprints

In order to provision Terraform apps, Terraform App Blueprints must be created first.

Important: In Morpheus version 4.2.0, VMware and AWS Cloud types are supported for Terraform App provisioning targets. Additional clouds will be available in later releases.

1. Navigate to *Provisioning -> Blueprints*
2. Select + *ADD*
3. Name the Blueprint and select *Terraform* type.

Note: In order to see the Terraform Blueprint type option, the Morpheus user must have Role permissions for *Provisioning: Blueprints - Terraform* set to *Full*.

4. Select *NEXT*

5. Configure the following:

NAME Name of the

DESCRIPTION Description for you App Blueprints shown in the Apps list (optional)

CATEGORY App Category (optional)

IMAGE Add reference image/picture for your App Blueprint (optional)

CONFIG TYPE (select Terraform, Terraform.json, or Git Repository)

Terraform (.tf)

CONFIG Paste in the .tf contents in the config section. Variables will be presented as input fields during App provisioning, or auto-populated with matching values if contained in a selected TFVAR Secret file added to the Cypher service.

Terraform JSON (.tf.json) Paste in .tf.json contents in the config section. Variables will be presented as input fields during App provisioning, or auto-populated with matching values if contained in a selected TFVAR Secret file added to the Cypher service.

Git Repository

SCM Integration Select a Github SCM integration that has been added in *Administration - Integrations*. If using a Git Repository integration from *Administration - Integrations* this field can be skipped.

Repository Select repository from selected SCM integration, or Git Repository integration from *Administration - Integrations* if no SCM/Github Integration is selected.

BRANCH OR TAG i.e. master (default)

WORKING PATH Enter the repo path for the .tf files (s). `./` is default.

CONFIG .tf files found in the working path will populate in the CONFIG section.

Note: If no files are found please ensure your Github or Git integration is configured properly (Private repos need to have a key pair added to Morpheus, the keypair selected on the integration in Morpheus, and the keypair's public key added to the GitHub users SSH keys in github or to the git repo).

TFVAR SECRET Select a tfvars secret for .tf variables. Tfvars secrets can be added in *Services -> Cypher* using the tfvars/name mountpoint. This allows sensitive data and passwords to be encrypted and securely used with Terraform Blueprints.

OPTIONS example `-var 'instanceName=sampleTfApp'`

6. Select *SAVE*

Your Terraform App is ready to be provisioned from *Provisioning -> Apps*.

Provisioning Terraform Apps

Note: An existing Terraform App Blueprints must be added to *Provisioning -> Blueprints* before it can be provisioned.

Note: In order to provision Terraform Apps in *Provisioning -> Apps*, the Morpheus user must have Role permissions for *Provisioning: Blueprints - Terraform* set to *Provision* or *Full*.

1. Navigate to *Provisioning -> Apps*
2. Select + *ADD*
3. Choose an existing Terraform App Blueprint
4. Select *NEXT*
5. Enter a NAME for the App and select the Group, Default Cloud and Environment (optional)
6. Select *NEXT*
7. Populate any required variables in the *Terraform Variables* section. ..TIP:: If the tf CONFIG data needs to be edited, select the *RAW* section, edit, and then select the *BUILDER* section again. The CONFIG changes from the RAW edit will be updated in the CONFIG section.
8. Select *COMPLETE*

The Terraform App will begin to provision.

Once provisioning is completed, note the TERRAFORM tab in the App details page (*Provisioning -> Apps -> select the App*). This section contains State and Plan output:

Operations

Provisioning

Infrastructure

Backups

Logs

Monitoring

Services

Administration

Instances

Apps

Templates

Automation

Virtual Images

Library

Migrations

Deployments

Apps > terraform-vmw-sample

terraform-vmw-sample

Template: Terraform: VMware

EDIT

ACTIONS +

DELETE

STATUS

HEALTH

100.000%

AVAILABILITY

2MS

RESPONSE TIME

0%

MAX CPU

27%

MEMORY

17%

STORAGE

▼ INFO

Group: All Clouds Demo

Cloud: VMware vCenter

Date Created: 04/19/2018 02:29 PM

Created By: Jeff Wheeler

Price: \$36.7318 / Month

INSTANCES

SECURITY GROUPS

ENVIRONMENT

LOGS

MONITORING

TERRAFORM

State

```
1 data.vsphere_datacenter.dc:
2   id = datacenter-2
3   name = labs-denver
4 data.vsphere_datastore.datastore:
5   id = datastore-204
6   datacenter_id = datacenter-2
7   name = labs-demo-qnap-240
8 data.vsphere_network.network:
9   id = network-51
10  datacenter_id = datacenter-2
11  name = VM Network
12  type = Network
13 data.vsphere_resource_pool.pool:
14  id = resgroup-158
15  datacenter_id = datacenter-2
16  name = labs-den-demo-cluster/Resources
17 data.vsphere_virtual_machine.template:
18  id = 421f5ca5-8f19-d108-b9f0-c977ae44176b
19  alternate_guest_name =
20  datacenter_id = datacenter-2
21  disks.# = 1
```

Plan

```
1
2 No changes. Infrastructure is up-to-date.
3
4 This means that Terraform did not detect any differences between your
5 configuration and real physical resources that exist. As a result, no
6 actions need to be performed.
```

556

Chapter 1. v5.2.0 Highlights

vRealize Orchestrator

The vRealize Orchestrator (vRO) Integration provided for Morpheus enables users to easily trigger existing workflows that may already exist in vRealize Orchestrator. Not only can the user trigger these workflows, but they can also be chained easily into non-vRO workflows and process both output and input parameters of a workflow.

Adding the Integration

Setting up the vRO integration involves some steps which vary depending on the authentication model being used.

When using OAUTH, the Client ID must be gathered first. This can be found by browsing a file on the actual vRA server using SSH. On the vRA server, run the following command: `grep -i cafe_cli= /etc/vcac/solution-users.properties | sed -e 's/cafe_cli=/'`

Secondly, you will need the username, password, and host API URL. Typically, the API URL is run on port 8283. A sample API URL may look like the following example: `https://vrahost.com:8283/`

Be sure to fill in the tenant token as the domain or tenant ID, for example: `vsphere.local`, with a username of `administrator@vsphere.local`.

Note: At times, this can vary depending on how authentication and role assignments for the user have been set up for vRO.

vRA auth uses vRA identity Bearer tokens for API consumption. The only real difference in field requirements when using this authentication mode is that the *Client ID* is no longer needed.

Using vRealize Orchestrator

One of the first things Morpheus does when it is tied into a vRO integration is sync all available workflows by category. These workflows become available when creating a new Morpheus task in *Provisioning -> Automation*. Morpheus allows a user to map these vRO workflows into the task engine. The task engine allows users to design workflows that chain tasks in order or operate at different phases of a provisioning request. For more information on tasks, please read the Automation documentation.

Creating a task for VRO is simple.

First, go to *Provisioning -> Automation* and create a new task. Choose a task type of *vRealize Orchestrator Workflow*. A dropdown will appear allowing one to first select the active vRO Integration you would like to use. Once that is selected, a list of workflows becomes available.

Note: The next part is where things can get a bit tricky. The parameter body (expected in JSON) format can be a bit difficult to track down. One way is to use the Network Chrome inspector when kicking off a sample workflow from the vRO HTML5 client and grabbing the parameter JSON. Another is to query the API yourself and look at the samples from historical run history.

An example payload for the *SSH / Run SSH Command* Workflow would look like this:

```
{
  "parameters": [
    {
      "name": "hostNameOrIP",
      "type": "string",
```

(continues on next page)

(continued from previous page)

```
        "value": {
            "string": {
                "value": "x.x.x.x"
            }
        }
    },
    {
        "name": "port",
        "type": "number",
        "value": {
            "number": {
                "value": 22
            }
        }
    },
    {
        "name": "cmd",
        "type": "string",
        "value": {
            "string": {
                "value": "echo \"Hello <%=instance.name%>\""
            }
        }
    },
    {
        "name": "encoding",
        "type": "string",
        "value": {
            "string": {
                "value": ""
            }
        }
    },
    {
        "name": "username",
        "type": "string",
        "value": {
            "string": {
                "value": "myuser"
            }
        }
    },
    {
        "name": "passwordAuthentication",
        "type": "boolean",
        "value": {
            "boolean": {
                "value": true
            }
        }
    },
    {
        "name": "password",
        "type": "string",
        "value": {
            "string": {
                "value": "password"
            }
        }
    }
}
```

(continues on next page)

(continued from previous page)

```

        }
    },
    {
        "name": "path",
        "type": "string",
        "value": {
            "string": {
                "value": "\\var\\lib\\vco\\app-server\\conf\\vco_key"
            }
        }
    },
    {
        "name": "passphrase",
        "type": "string",
        "value": {
            "string": {
                "value": ""
            }
        }
    }
]
}

```

Note that all Morpheus variables can be injected into the parameter body. In the above example we inject the instance name into the sample command with `<%=instance.name%>`.

Adding this task to a workflow allows the result parameters to be referenced in subsequent tasks called throughout the workflow. For example, a local script task type could reference the output text of the above ssh command by injecting the following results map: `echo "results.vro: <%=results.vro.find{it.name == 'outputText'}?.value?.string?.value%>"`

There are very powerful options available for chaining results and injecting variables relevant to the instance being provisioned or even custom inputs from an operational workflow. Please reference the rest of the Automation documentation for examples.

Backups

Commvault

Adding Commvault Integration

1. Navigate to *Backups -> Services*
2. Select + *ADD*
3. Select Commvault
4. Fill in the following:

Name Name of the Integration in Morpheus

Enabled Enable the Commvault integration

Host IP or Hostname of the Commvault server.

Port Port number configured to access the Commvault server

Username Admin Username for Commvault

Password Password for Username provided (encrypted in Morpheus).

Visibility

Sets Multi-Tenant Visibility

Private Only Available to the Tenant the Integration is added by

Public Available to Sub-Tenants (master tenant option only)

5. *SAVE*

Veeam

Adding Veeam Integration

1. Navigate to *Backups -> Services*

2. Select + *ADD*

3. Select Veeam

4. Fill in the following:

Name Name of the Integration in Morpheus

Enabled Enable the Veeam integration

Host IP or Hostname of the Veeam server, must be HTTPS for VEEAM 10

Port Port number configured to access the Veeam server, must be 9398 for VEEAM 10

Username Admin Username for Veeam

Password Password for Username provided (encrypted in Morpheus).

Visibility

Sets Multi-Tenant Visibility

Private Only Available to the Tenant the Integration is added by

Public Available to Sub-Tenants (master tenant option only)

5. Click *SAVE*

Note: Veeam Backup Enterprise Manager must be installed on the Veeam server in order to successfully integrate Morpheus with Veeam.

Important: Once Veeam service has been integrated with Morpheus, Veeam server(s) will be available to select as the backup provider for VMware, Hyper-V, and vCloud Director cloud integrations (Infrastructure > Clouds > Edit a compatible Cloud). To enable Veeam backups, select the appropriate Veeam server as the “backup provider” for your cloud integrations as needed. Failure to do so will result in blank *Backup Repositories* and *Backup Job Templates* options when configuring Veeam Backups during provisioning.

Rubrik

Adding Rubrik Integration

Note: The Rubrik backup service is currently only supported on the VMware cloud type.

1. Navigate to *Backups -> Services*
2. Select + *ADD*
3. Select Rubrik
4. Fill in the following:

Name Name of the Integration in Morpheus

Enabled Enable the Integration

Host IP or Hostname of the Rubrik api server.

Username Admin Username for Rubrik

Password Password for Username provided (encrypted in Morpheus).

Visibility

Sets Multi-Tenant Visibility

Private Only Available to the Tenant the Integration is added by

Public Available to Sub-Tenants (master tenant option only)

5. *SAVE*

Zerto

Adding Zerto Integration

1. Navigate to *Backups -> Integrations*
2. Select + *ADD*
3. Select Zerto
4. Fill in the following:

Name Name of the Integration in Morpheus

Enabled Enable the Integration

API URL

API URL for Zerto Virtual Manager Example `API URL: https://zvm_IP:9669`

Username Admin Username for Zerto

Password Password for Username provided (encrypted in Morpheus).

Visibility

Sets Multi-Tenant Visibility

Private Only Available to the Tenant the Integration is added by

Public Available to Sub-Tenants (master tenant option only)

5. *SAVE*

Avamar

IMPORTANT: Avamar API must be installed on Avamar server (not installed by default)

Adding Avamar Integration

1. Navigate to *Backups -> Services*
2. Select + *ADD*
3. Select Avamar
4. Fill in the following:

Name Name of the Integration in Morpheus

Enabled Enable the Integration

Host IP or Hostname of the Avamar api server.

Port Port number configured to access the Avamar server

Username Admin Username for Avamar

Password Password for Username provided (encrypted in Morpheus).

Tenant Avamar Tenant/Domain to scope Integration to

Hypervisor Avamar Hypervisor to scope Integration to

Visibility

Sets Multi-Tenant Visibility

Private Only Available to the Tenant the Integration is added by

Public Available to Sub-Tenants (master tenant option only)

5. *SAVE*

Build

Jenkins

The Morpheus Jenkins Integration is easy to add and will allow you to see all jobs, builds, statuses of those builds, commits notes, and links to artifacts.

Adding Jenkins Integration

NEW JENKINS INTEGRATION

×

NAME

Jenkins

☒ ENABLED

JENKINS URL

http://10.30.20.87:8080/

USERNAME

Ryan

PASSWORD

••••••••

SAVE CHANGES

1. Navigate to Administration -> Integrations
2. Select + *NEW INTEGRATION*
3. Select Jenkins
4. Fill in the following:

Name Name of the Integration in Morpheus

Enabled Enable the Integration. Uncheck to disable the Jenkins Integration sync Job.

Jenkins URL Jenkins URL or IP address. ex: `https://jenkins.morpheus.com`

Username Jenkins service account username

Password Jenkins service account password

5. *SAVE CHANGES*

Important: By default Jenkins is configured to run on port 8080. If this has been modified you will need to append the alternate port to the the `Jenkins URL`

Viewing Jobs in Jenkins Integration

In the Morpheus Jenkins integration you can view all of your jobs.

The screenshot shows the Morpheus Jenkins integration interface. The top navigation bar includes tabs for Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Services, and Administration. The 'Administration' tab is active, and the 'Integrations' sub-tab is selected. The main content area displays the 'Labs Jenkins' integration. Below the integration name, there are 'EDIT' and 'DELETE' buttons. The 'Host' is listed as 'https://jenkins.' and the 'Last Update' is '07/27/2019 02:40 AM'. The 'JOBS' tab is selected, showing a list of jobs with columns for NAME, LAST BUILD, QUEUED, and HEALTH. The jobs are listed in descending order of last build time.

NAME	LAST BUILD	QUEUED	HEALTH
[REDACTED]	03/08/2017 06:42 PM	No	ok
[REDACTED]	07/31/2019 05:12 PM	No	ok
[REDACTED]	11/14/2018 08:27 AM	No	error
[REDACTED]	08/31/2018 09:42 PM	No	ok
[REDACTED]	06/27/2017 12:36 PM	No	ok
[REDACTED]	03/10/2017 11:32 AM	No	ok
[REDACTED]	05/09/2018 04:12 PM	No	ok
[REDACTED]	06/08/2015 02:55 PM	No	ok
[REDACTED]	09/07/2018 06:24 PM	No	error
[REDACTED]	05/22/2018 03:33 PM	No	ok
[REDACTED]	07/23/2019 05:31 PM	No	ok
[REDACTED]	04/23/2019 03:17 PM	No	ok
[REDACTED]	04/23/2019 03:17 PM	No	ok
[REDACTED]	11/12/2018 03:04 PM	No	unknown
[REDACTED]	01/23/2017 01:55 PM	No	ok
[REDACTED]	06/14/2016 08:33 PM	No	error
[REDACTED]		No	unknown
[REDACTED]	03/18/2015 09:33 AM	No	ok
[REDACTED]	07/31/2019 02:26 PM	No	ok
[REDACTED]	06/27/2019 12:06 PM	No	ok

Viewing Builds and Build Statuses

In the Morpheus Jenkins integration you can view recent builds with ID, Status, Date, Duration, Artifacts, Commit Notes and Run By user data. Artifacts will automatically link to the Artifact url in Jenkins, and the urls can be used in Morpheus Deployments (dependent on Jenkins configuration).

OperationsProvisioningInfrastructureBackupsLogsMonitoringServicesAdministration

TenantsPlans & PricingRolesUsersIntegrationsPoliciesProvisioningMonitoringBackupsLogsSettings

Integrations > Labs Jenkins

Labs Jenkins

Jenkins

Host: Last Update: 07/27/2019 02:40 AM

JOBS

BUILDS

BUILDS

Search Select

ID	JOB	STATUS	DATE	DURATION	ARTIFACT	COMMIT	RUN BY
#871		Ok	07/31/2019 08:30 PM	12 minutes 11 seconds			
#870		Ok	07/31/2019 07:27 PM	12 minutes 27 seconds			
#869		Ok	07/31/2019 05:56 PM	12 minutes 8 seconds			
#868		Ok	07/31/2019 05:31 PM	11 minutes 59 seconds			
#867		Ok	07/31/2019 05:19 PM	11 minutes 57 seconds			
#985		Ok	07/31/2019 05:12 PM	7 minutes 26 seconds			
#886		Ok	07/31/2019 05:07 PM	12 minutes			
#865		Ok	07/31/2019 04:55 PM	11 minutes 51 seconds			
#864		Ok	07/31/2019 04:43 PM	11 minutes 59 seconds			
#863		Ok	07/31/2019 04:31 PM	12 minutes 2 seconds			

Clouds

AWS

Overview

AWS is the Amazon public cloud, offering a full range of services and features across the globe in various datacenters. AWS provides businesses with a flexible, highly scalable, and low-cost way to deliver a variety of services using open standard technologies as well as proprietary solutions. This section of documentation will help you get Morpheus and AWS connected to utilize the features below:

Features

- Instance, Service, Infrastructure Provisioning & Synchronization
- EKS Cluster Creation & Synchronization
- Morpheus Kubernetes, Docker & KVM Cluster Creation
- ELB Classic Load Balancer Creation & Synchronization
- ELB Application Load Balancer (ALB) Creation & Synchronization
- Security Group Creation & Synchronization
- Security Group Rule Creation & Synchronization
- Network Synchronization
- VPC Creation & Synchronization
- CloudFormation Provisioning & Resource Synchronization
- Terraform Provisioning & Resource Synchronization
- Pricing & Costing Synchronization
- MetaData Tag Creation & Synchronization
- S3 Bucket Creation & Synchronization
- Route53 Automation & Synchronization
- IAM Profile Synchronization and Assignment
- RDS Support
- Backups / Snapshots
- Migrations
- Auto Scaling
- Remote Console (SSH & RDP)
- Lifecycle Management and Resize
- Restore from Snapshots
- Elastic IP Assignment
- Network Pools
- Enhanced Invoice Costing

Requirements

AWS IAM Security Credentials Access Key Secret Key Sufficient User Privileges (see [MinimumIAMPolicies](#) section for more info)

Security Group Configuration for Agent Install, Script Execution, and Remote Console Access

- Typical Inbound ports open from Morpheus Appliance: 22, 5985, 3389 (22 & 3389 required for Console. 22 & 5985 required for agent-less comms)
- Typical Outbound to Morpheus Appliance: 443 (Required for Agent install & comms)

Note: These are required for Morpheus agent install, communication, and remote console access for windows and linux. Other configurations, such as docker instances, will need the appropriate ports opened as well. Cloud-init Agent Install mode does not require incoming access for port 22.

Network(s) IP assignment required for Agent install, Script Execution, and Console if the Morpheus Appliance is not able to communicate with AWS instances private ip's.

Note: Each AWS Cloud in Morpheus is scoped to an AWS Region and VPC. Multiple AWS Clouds can be added and even grouped if different region and VPC combinations are needed. It's also recommended you verify Security Groups are properly configured in all regions Morpheus Clouds will scope to.

Adding an AWS Cloud

1. Navigate to *Infrastructure -> Clouds*
2. Select + *Create Cloud*
3. Select AWS
4. Enter the following:

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When "AUTOMATICALLY POWER ON VMS" is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

REGION Select AWS Region for the Cloud

ACCESS KEY Access Key ID from AWS IAM User Security Credentials.

SECRET KEY Secret Access Key associated with the Access Key ID.

USE HOST IAM CREDENTIALS Check to use use Host IAM Credentials

ROLE ARN Supports security token service (STS) to AssumeRole by entering an AWS Role ARN

INVENTORY

Basic Morpheus will sync information on all EC2 Instances in the selected VPC the IAM user has access to, including Name, IP Addresses, Platform Type, Power Status, and overall resources sizing for Storage, CPU and RAM, every 5 minutes. Inventoried EC2 Instances will appear as Unmanaged VM's.

Full In addition to the information synced from Basic Inventory level, Morpheus will gather Resource Utilization metrics for Memory, Storage and CPU utilization per VM.

Off Existing EC2 Instances will not be inventoried

Note: Cloud Watch must be configured in AWS for Morpheus to collect Memory and Storage utilization metrics on inventoried EC2 instances.

USE VPC Specify if the target account is using EC2-VPC or EC2-Classic Platform. In almost all cases, VPC should be selected, and then select the target VPC from the synced available VPC's list, or *All VPC's*.

5. The AWS cloud is ready to be added to a group and saved. Additional configuration options available:

IMAGE TRANSFER STORE S3 bucket for Image transfers, required for migrations into AWS.

EBS ENCRYPTION Enable or disable encryption of EBS Volumes

COSTING KEY For Gov Cloud pricing only, key for standard managing cost account

COSTING SECRET For Gov Cloud pricing only, secret for standard managing cost account

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME_ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Enhanced Invoice Costing Configuration

As of version 4.2.3, AWS cloud integrations in Morpheus sync billing data solely through the use of AWS Costing & Utilization Reports (CUR). In prior versions users could instead sync billing data through AWS Cost Explorer if desired. Version 4.2.3 also simplified the way CUR reports can be selected or created in order to sync costing data. The section below discusses setting up enhanced costing through CUR reports both in 4.2.3 and versions prior. Keep in mind you must go through this process in version 4.2.3 and higher in order for Morpheus to aggregate billing data.

Note: Even with a costing report configured in the Cloud integration as described below, the COSTING value must also be set to “Costing and Reservations” in order for enhanced invoice data to be brought into Morpheus. Confirm this setting by editing the Amazon Cloud integration, and checking the COSTING value in the Advanced Options panel before continuing.

v4.2.3 and Above

In Morpheus 4.2.3+, edit the Amazon cloud integration or create a new Amazon Cloud to get started. On the Create/Edit Cloud modal, open the advanced options section. The relevant fields for configuring invoice costing are shown below:

COSTING REPORT	<div>New Report</div>
COSTING REPORT NAME	<div>test_report</div>
COSTING FOLDER	<div>reports</div>
COSTING BUCKET	<div>New Report Bucket</div>
COSTING BUCKET NAME	<div>new_bucket</div>
COSTING BUCKET REGION	<div>US West (N. California)</div>
COSTING KEY	<div></div>
COSTING SECRET	<div></div>
LINKED ACCOUNT ID	<div></div>

In the example case above, a new report and a new S3 bucket are created but Morpheus will also sync in buckets and reports that meet the required parameters if they already exist. For reports to be synced they must meet the requirements listed below:

- Hourly time granularity
- Include resource IDs
- GZIP compression
- CSV format

If you don't currently have a report meeting those criteria, you can create one by selecting "New Report" from the REPORT NAME dropdown menu. A new S3 bucket can be created in similar fashion if needed. You may also want to review the section below on configuration for Morpheus 4.2.2 and below to note policies that will be applied to your selected bucket and Cost Explorer permissions required for the AWS cloud user associated with the Morpheus Cloud integration.

In the end, the following fields must be filled in order to complete the process:

- **COSTING BUCKET:** The S3 bucket name
- **COSTING REGION:** The region the bucket was created in
- **COSTING FOLDER:** This is the report path prefix if you configured one earlier
- **COSTING REPORT NAME:** The name given to your CUR report
- **COSTING KEY:** If the IAM user for this AWS cloud integration does not have access to the S3 bucket with the CUR data, enter the AWS Key ID for an IAM user with access
- **COSTING SECRET:** If the IAM user for this AWS cloud integration does not have access to the S3 bucket with the CUR data, enter the AWS Secret Key for the IAM account whose Key ID you entered in the previous field
- **LINKED ACCOUNT ID:** If the IAM user for this AWS cloud integration does not have access to the S3 bucket with the CUR data, enter the AWS account number that the IAM user from the above step resides in

Note: If the AWS cloud account is a GovCloud account, enter the COSTING KEY, COSTING SECRET, and LINKED ACCOUNT ID for the master commercial account your GovCloud account is associated with.

v4.2.2 and Below

Begin by logging into the [AWS Billing Console](#), then click *Create report*.

The screenshot shows the AWS Cost and Usage Reports console. The top navigation bar includes the AWS logo, 'Services', 'Resource Groups', and a star icon. The left sidebar lists navigation options: Home, Cost Management, Cost Explorer, Budgets, Budgets Reports, Savings Plans, **Cost & Usage Reports** (highlighted), Cost Categories, Cost allocation tags, Billing, Bills, Orders and invoices, Credits, Purchase orders (preview), Preferences, Billing preferences, Payment methods, Consolidated billing, and Tax settings. The main content area is titled 'AWS Cost and Usage Reports' and contains the following elements:

- A paragraph: 'AWS Cost and Usage reports provide access to detailed data, enabling you to better analyze and understand your AWS usage amounts underlying those costs. You can customize the content and delivery of your reports and manage them'.
- A blue button labeled 'Create report'.
- A section titled 'Other Reports' with a dropdown arrow.
- A paragraph: 'You will be able to see your spend data within 24 hours of enabling reports.'
- A blue button labeled 'Enable Reports'.
- Three report options, each with an icon and description:
 - Bar chart icon:** 'Analyze your cost and usage using AWS Cost Explorer. AWS Cost Explorer lets you dive deeper into your cost and usage data to identify trends, pinpoint cost drivers, and optimize your spend.' (Note: the text in the image is partially cut off).
 - Line graph icon:** 'Monitor your Reserved Instance (RI) utilization using AWS Cost Explorer. This report allows you to visualize your RI utilization, providing insight into opportunities for increasing your utilization.' (Note: the text in the image is partially cut off).
 - Document icon:** [AWS Usage Report](#). 'A CSV report that provides usage data for select AWS services.'

Include a name for your report and mark the box to “Include resource IDs”. Morpheus uses these resource IDs to map costs to various resources. Click *Next*.

The screenshot shows the AWS Cost and Usage Reports console. The top navigation bar includes the AWS logo, 'Services', 'Resource Groups', and a star icon. The breadcrumb trail is 'AWS Cost and Usage Reports > Create report'. On the left, a sidebar shows the steps: 'Step 1: Report content' (active), 'Step 2: Delivery options', and 'Step 3: Review'. The main content area is titled 'Report content'. It features a 'Report name - required' field with the value 'Morpheus'. Below this, the 'Report includes' section lists several categories: Account identifiers, Invoice and Bill Information, Usage Amount and Unit, Rates and Cost, Product Attributes (e.g., instance type, operating system, and region), Pricing Attributes (e.g., offer types, and lease lengths), and Reservation identifiers and related details (for reserved instances only). The 'Additional report details' section has a checked checkbox for 'Include resource IDs' with an information icon. The 'Data refresh settings' section has a checked checkbox for 'Automatically refresh your Cost & Usage Report when charges are detected for previous months with closed bills.' At the bottom right, there are 'Cancel' and 'Next' buttons.

Step 1
Report content

Step 2
Delivery options

Step 3
Review

Report content

Report name - required

Morpheus

Report includes

- Account identifiers
- Invoice and Bill Information
- Usage Amount and Unit
- Rates and Cost
- Product Attributes (e.g., instance type, operating system, and region)
- Pricing Attributes (e.g., offer types, and lease lengths)
- Reservation identifiers and related details (for reserved instances only)

Additional report details

☒ Include resource IDs ⓘ

Data refresh settings ⓘ

☒ Automatically refresh your Cost & Usage Report when charges are detected for previous months with closed bills.

Cancel Next

On the following page, begin by identifying an S3 bucket to house reports. Click *Configure* near the top of the page and select an existing bucket or create a new one.

Step 1 of 2: Configure S3 Bucket

In order to receive AWS Cost & Usage Reports, you must have an Amazon S3 bucket created and configured with the appropriate access permissions. You can add an existing bucket or create a new one.

Select existing bucket

S3 bucket name

Choose an existing S3 Bucket ▼

Create a bucket

S3 bucket name

Enter a Bucket name

Region

US East (N. Virginia) ▼

OR

If you create a bucket for AWS Cost and Usage reports, default settings and permissions are applied. After the bucket is created, you can update the settings and permissions on the Amazon S3 console.

Some Regions are disabled by default. To create a bucket in one of these Regions, you must enable the Region first. [Learn more](#)

Cancel
Next

After identifying the bucket, you must mark the box to accept the default policy being applied to the bucket. Click *Save*.

Step 2 of 2: Verify policy

The following default policy will be applied to your bucket:

```
{
  "Version": "2008-10-17",
  "Id": "arn:aws:s3:::bucket-name",
  "Statement": [
    {
      "Sid": "AllowBillingReportsAccess",
      "Effect": "Allow",
      "Principal": {
        "Service": "billingreports.amazonaws.com"
      },
      "Action": "*"
    }
  ]
}
```

☒ I have confirmed that this policy is correct

Note: Saving the report in the existing S3 bucket will overwrite the current policy.

Cancel
Previous
Save

The default policy applied to the bucket is below:

```

{
  "Version": "2008-10-17",
  "Id": "SomeID",
  "Statement": [
    {
      "Sid": "SomeStmtID",
      "Effect": "Allow",
      "Principal": {
        "Service": "billingreports.amazonaws.com"
      },
      "Action": [
        "s3:GetBucketAcl",
        "s3:GetBucketPolicy"
      ],
      "Resource": "arn:aws:s3:::bucket-name"
    },
    {
      "Sid": "SomeStmtID",
      "Effect": "Allow",
      "Principal": {
        "Service": "billingreports.amazonaws.com"
      },
      "Action": [
        "s3:PutObject"
      ],
      "Resource": "arn:aws:s3:::bucket-name/*"
    }
  ]
}

```

After choosing a bucket, accepting the default policy, and saving the change, you're brought back to the report delivery page. By default, CUR reports are saved to a folder at the path `my-report-name/date-folder`. If this bucket already contains CUR reports, you may want to specify a prefix path in the "Report path prefix" field. Outside of this field, use the default values as shown in the screenshot below, then click *Next*.

In addition, the AWS cloud user associated with the integration in Morpheus needs IAM policy permission to access Cost Explorer. Attach a policy like the one below to this cloud user:

```
{
  "Version": "2012-10-17",
  "Id": "SomeID",
  "Statement": [
    {
      "Sid": "SomeStmtID",
      "Effect": "Allow",
      "Action": [
        "ce:DescribeReportDefinitions",
        "ce:DescribeCostCategoryDefinition",
        "ce:ListCostCategoryDefinitions"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Note: If the Cost Explorer permissions are granted at the master account level, the user will see all costs for each member account; if granted at the member account, only the costs for that member account are available.

With the AWS console configuration steps complete, we can move back into Morpheus. Keep in mind it is only necessary to set up one AWS cloud for Costing since we process all records in the CUR report.

Once back in Morpheus, add or edit the relevant AWS cloud integration (Infrastructure > Clouds > + *ADD* OR click the pencil icon in the row for the chosen AWS integration). Expand the Advanced Options drawer and complete the following fields:

- **COSTING BUCKET:** The S3 bucket name
- **COSTING REGION:** The region the bucket was created in
- **COSTING FOLDER:** This is the report path prefix if you configured one earlier
- **COSTING REPORT NAME:** The name given to your CUR report
- **COSTING KEY:** If the IAM user for this AWS cloud integration does not have access to the S3 bucket with the CUR data, enter the AWS Key ID for an IAM user with access
- **COSTING SECRET:** If the IAM user for this AWS cloud integration does not have access to the S3 bucket with the CUR data, enter the AWS Secret Key for the IAM account whose Key ID you entered in the previous field
- **LINKED ACCOUNT ID:** If the IAM user for this AWS cloud integration does not have access to the S3 bucket with the CUR data, enter the AWS account number that the IAM user from the above step resides in

Note: If the AWS cloud account is a GovCloud account, enter the COSTING KEY, COSTING SECRET, and LINKED ACCOUNT ID for the master commercial account your GovCloud account is associated with.

COSTING BUCKET	<input type="text"/>
COSTING REGION	<input type="text" value="us-east-2"/>
COSTING FOLDER	<input type="text"/>
COSTING REPORT NAME	<input type="text" value="Morpheus"/>
COSTING KEY	<input type="text"/>
COSTING SECRET	<input type="text"/>
LINKED ACCOUNT ID	<input type="text"/>

Save changes to your cloud integration.

Important: It may take as long as one hour for Morpheus to process the next CUR report.

AWS Reserved Instances and Savings Plans

Amazon AWS public cloud offers Reserved Instances (RI) and Savings Plans, which allow organizations with consistent use patterns to reduce cloud spend significantly. Morpheus analyzes AWS cloud usage and spend, which allows it to make intelligent recommendations that can lead to significant savings. This data can be reviewed in the Reservation Recommendations and Savings Plan Recommendations tables on any AWS Cloud detail page (Infrastructure > Clouds > Selected Amazon Cloud).

Savings Plans potentially offer greater than 70% savings in exchange for a commitment to consistent usage levels for a 1- or 3-year term. Morpheus provides Savings Plan guidance based on learned analytics; allowing you to analyze Savings Plans based on different term commitments and upfront costs to choose the best savings plan.

SAVINGS PLAN RECOMMENDATIONS

EC2 Instance

3-Year

No upfront

30 Days

ITEM		CURRENT SPEND		RECOMMENDATION					TERM COST		
SERVICE	TYPE	AVG / HOUR	TOTAL / MONTH	COMMITMENT	EST. ON-DEMAND	TOTAL / MONTH	SAVINGS	SAVINGS %	ROI (MONTHS)	UPFRONT	TOTAL
EC2	T2	\$2.65	\$1,936.24	\$1.25	\$0.25	\$1,095.28	\$840.96	43.4%	15.64	\$0.00	\$32,928.84
EC2	M4	\$0.35	\$255.87	\$0.17	\$0.02	\$134.20	\$121.67	47.6%	17.12	\$0.00	\$4,362.48
EC2	M4	\$1.15	\$839.54	\$0.75	\$0.25	\$728.72	\$110.81	13.2%	4.75	\$0.00	\$19,710.00
TOTAL			\$3,031.65			\$1,958.21	\$1,073.44			\$0.00	\$57,001.32

Reserved Instances (RI) provide a discounted hourly rate and optional capacity reservation for EC2 instances. AWS billing automatically applies your RI-discounted rate when the attributes of EC2 instance usage match attributes of an active RI. Morpheus provides RI guidance based on learned analytics.

RESERVATION RECOMMENDATIONS

Amazon Elastic Compute Cloud - Compute

1-Year

No upfront

ITEM		CURRENT USAGE			RECOMMENDATION				TERM COST	
SERVICE	TYPE	AVG	TOTAL / MONTH	RESERVED	TOTAL / MONTH	SAVINGS	SAVINGS %	ROI (MONTHS)	UPFRONT	TOTAL
EC2	T2.NANO	16.9	\$71.55	17	\$44.68	\$26.87	38%	4.51	\$0.00	\$536.11
EC2	M4.LARGE	3	\$256.23	3	\$186.37	\$69.86	27%	3.27	\$0.00	\$2,236.43
EC2	T2.NANO	387.97	\$1,954.22	346	\$1,262.90	\$468.72	24%	2.88	\$0.00	\$15,154.80
EC2	T2.NANO	1.99	\$9.87	2	\$6.86	\$3.01	30%	3.66	\$0.00	\$82.34
TOTAL			\$2,291.87		\$1,500.81	\$568.47			\$0.00	\$18,009.68

Minimum AWS IAM Policies

Below are the AWS IAM Permissions covering the minimum access for Morpheus applying to all resources and services.

See http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies.html for more information.

Morpheus Sample AWS IAM Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "autoscaling:DescribeAutoScalingGroups",
        "ce:*",
        "cloudwatch:GetMetricStatistics",
        "ec2:AllocateAddress",
        "ec2:AssignPrivateIpAddresses",
```

(continues on next page)

(continued from previous page)

```
"ec2:AssociateAddress",
"ec2:AttachInternetGateway",
"ec2:AttachNetworkInterface",
"ec2:AttachVolume",
"ec2:AuthorizeSecurityGroupEgress",
"ec2:AuthorizeSecurityGroupIngress",
"ec2:CancelExportTask",
"ec2:CancelImportTask",
"ec2:CopyImage",
"ec2:CopySnapshot",
"ec2>CreateEgressOnlyInternetGateway",
"ec2>CreateImage",
"ec2>CreateInstanceExportTask",
"ec2>CreateInternetGateway",
"ec2>CreateKeyPair",
"ec2>CreateNatGateway",
"ec2>CreateNetworkAcl",
"ec2>CreateNetworkAclEntry",
"ec2>CreateNetworkInterface",
"ec2>CreateSecurityGroup",
"ec2>CreateSnapshot",
"ec2>CreateTags",
"ec2>CreateVolume",
"ec2>DeleteEgressOnlyInternetGateway",
"ec2>DeleteInternetGateway",
"ec2>DeleteKeyPair",
"ec2>DeleteNatGateway",
"ec2>DeleteNetworkAcl",
"ec2>DeleteNetworkAclEntry",
"ec2>DeleteNetworkInterface",
"ec2>DeleteSecurityGroup",
"ec2>DeleteSnapshot",
"ec2>DeleteTags",
"ec2>DeleteVolume",
"ec2:DeregisterImage",
"ec2:DescribeAccountAttributes",
"ec2:DescribeAddresses",
"ec2:DescribeAvailabilityZones",
"ec2:DescribeClassicLinkInstances",
"ec2:DescribeConversionTasks",
"ec2:DescribeEgressOnlyInternetGateways",
"ec2:DescribeExportTasks",
"ec2:DescribeImageAttribute",
"ec2:DescribeImages",
"ec2:DescribeImportImageTasks",
"ec2:DescribeImportSnapshotTasks",
"ec2:DescribeInstances",
"ec2:DescribeInstanceStatus",
"ec2:DescribeInternetGateways",
"ec2:DescribeKeyPairs",
"ec2:DescribeNatGateways",
"ec2:DescribeNetworkAcls",
"ec2:DescribeNetworkInterfaceAttribute",
"ec2:DescribeNetworkInterfaces",
"ec2:DescribeRegions",
"ec2:DescribeSecurityGroupReferences",
"ec2:DescribeSecurityGroups",
```

(continues on next page)

(continued from previous page)

```

"ec2:DescribeSnapshotAttribute",
"ec2:DescribeSnapshots",
"ec2:DescribeStaleSecurityGroups",
"ec2:DescribeSubnets",
"ec2:DescribeTags",
"ec2:DescribeVolumeAttribute",
"ec2:DescribeVolumes",
"ec2:DescribeVolumeStatus",
"ec2:DescribeVpcAttribute",
"ec2:DescribeVpcClassicLink",
"ec2:DescribeVpcClassicLinkDnsSupport",
"ec2:DescribeVpcEndpoints",
"ec2:DescribeVpcEndpointServices",
"ec2:DescribeVpcPeeringConnections",
"ec2:DescribeVpcs",
"ec2:DetachInternetGateway",
"ec2:DetachNetworkInterface",
"ec2:DetachVolume",
"ec2:DisassociateAddress",
"ec2:ImportImage",
"ec2:ImportInstance",
"ec2:ImportKeyPair",
"ec2:ImportSnapshot",
"ec2:ImportVolume",
"ec2:ModifyImageAttribute",
"ec2:ModifyInstanceAttribute",
"ec2:ModifyNetworkInterfaceAttribute",
"ec2:ModifySnapshotAttribute",
"ec2:ModifyVolumeAttribute",
"ec2:RebootInstances",
"ec2:RegisterImage",
"ec2:ReleaseAddress",
"ec2:ReplaceNetworkAclAssociation",
"ec2:ReplaceNetworkAclEntry",
"ec2:ResetImageAttribute",
"ec2:ResetInstanceAttribute",
"ec2:ResetNetworkInterfaceAttribute",
"ec2:ResetSnapshotAttribute",
"ec2:RevokeSecurityGroupEgress",
"ec2:RevokeSecurityGroupIngress",
"ec2:RunInstances",
"ec2:StartInstances",
"ec2:StopInstances",
"ec2:TerminateInstances",
"ec2:UnassignPrivateIpAddresses",
"ec2:UpdateSecurityGroupRuleDescriptionsEgress",
"eks:*",
"iam:ListGroups",
"iam:ListInstanceProfiles",
"iam:ListRoles",
"rds:AddRoleToDBCluster",
"rds:AddTagsToResource",
"rds:ApplyPendingMaintenanceAction",
"rds:AuthorizeDBSecurityGroupIngress",
"rds:CopyDBClusterSnapshot",
"rds:CopyDBParameterGroup",
"rds:CopyDBSnapshot",

```

(continues on next page)

(continued from previous page)

```
"rds:CreateDBCluster",
"rds:CreateDBClusterSnapshot",
"rds:CreateDBInstance",
"rds:CreateDBInstanceReadReplica",
"rds:CreateDBSecurityGroup",
"rds:CreateDBSnapshot",
"rds>DeleteDBCluster",
"rds>DeleteDBInstance",
"rds>DeleteDBSecurityGroup",
"rds>DeleteDBSnapshot",
"rds:DescribeAccountAttributes",
"rds:DescribeCertificates",
"rds:DescribeDBClusterParameterGroups",
"rds:DescribeDBClusterParameters",
"rds:DescribeDBClusters",
"rds:DescribeDBClusterSnapshotAttributes",
"rds:DescribeDBClusterSnapshots",
"rds:DescribeDBEngineVersions",
"rds:DescribeDBInstances",
"rds:DescribeDBLogFiles",
"rds:DescribeDBParameterGroups",
"rds:DescribeDBParameters",
"rds:DescribeDBSecurityGroups",
"rds:DescribeDBSnapshotAttributes",
"rds:DescribeDBSnapshots",
"rds:DescribeDBSubnetGroups",
"rds:DescribeEngineDefaultClusterParameters",
"rds:DescribeEngineDefaultParameters",
"rds:DescribeEventCategories",
"rds:DescribeEvents",
"rds:DescribeOptionGroupOptions",
"rds:DescribeOptionGroups",
"rds:DescribeOrderableDBInstanceOptions",
"rds:ListTagsForResource",
"rds:ModifyDBCluster",
"rds:ModifyDBClusterParameterGroup",
"rds:ModifyDBClusterSnapshotAttribute",
"rds:ModifyDBInstance",
"rds:ModifyDBParameterGroup",
"rds:ModifyDBSnapshotAttribute",
"rds:PromoteReadReplica",
"rds:RebootDBInstance",
"rds:RemoveTagsFromResource",
"rds:RestoreDBClusterFromSnapshot",
"rds:RestoreDBClusterToPointInTime",
"rds:RestoreDBInstanceFromDBSnapshot",
"rds:RestoreDBInstanceToPointInTime",
"rds:RevokeDBSecurityGroupIngress",
"route53:GetHostedZone",
"route53:ListHostedZones",
"route53:ListResourceRecordSets",
"s3:AbortMultipartUpload",
"s3:CreateBucket",
"s3>DeleteBucket",
"s3>DeleteObject",
"s3>DeleteObjectVersion",
"s3:GetBucketLocation",
```

(continues on next page)

(continued from previous page)

```

        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:ListBucketMultipartUploads",
        "s3:ListBucketVersions",
        "s3:ListMultipartUploadParts",
        "s3:PutObject"
    ],
    "Resource": "*"
}
]
}
```

Resource Filter

If you need to limit actions based on filters you have to pull out the action and put it in a resource based policy since not all the actions support resource filters.

See <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-supported-iam-actions-resources.html> for more info on limiting resources by filter.

Resource filter example:

```

{
  "Effect": "Allow",
  "Action": [
    "ec2:StopInstances",
    "ec2:StartInstances"
  ],
  "Resource": *
},
{
  "Effect": "Allow",
  "Action": "ec2:TerminateInstances",
  "Resource": "arn:aws:ec2:us-east-1:123456789012:instance/*",
  "Condition": {
    "StringEquals": {
      "ec2:ResourceTag/purpose": "test"
    }
  }
}
}
```

Azure (Public)

Overview

Morpheus offers a complete Integration with Microsoft Azure including the following:

- Virtual Machine Sync, Create, Delete, Manage, RBAC, Tenant Permissions, Policies
- Resource Group Sync, Create, Delete, RBAC, Tenant Permissions
- Network Sync, Create, Delete, RBAC, Tenant Permissions

- Subnet Sync, Create, Delete, RBAC, Tenant Permissions
- Security Group Sync, Create, Delete, Tenant Permissions
- Security Group Rule Sync, Create, Delete, Tenant Permissions
- ARM Blueprints, Spec Templates, Deployment Logs Sync, Git/GitHub Integration
- MSSQL Service Sync, Create, Delete, Manage, RBAC, Tenant Permissions
- AKS Sync, Sync, Create, Delete, Manage, RBAC, Tenant Permissions
- Backup Create, Delete, Manage, RBAC, Policies
- Storage Sync, Create, Delete, Manage, Browse, RBAC, Tenant Permissions, Policies
- Marketplace Sync
- Private Image Sync & Upload
- Azure Marketplace Custom Library Item Support
- Remote Console (SSH & RDP)
- Lifecycle Management
- Availability Set Support
- Scale Set Sync, Create, Assign, Manage, Delete
- Azure Load Balancer Create, Assign, Manage, Delete, RBAC, Tenant Permissions
- Docker (VM) Cluster Sync, Create, Delete, Manage, RBAC, Tenant Permissions
- Kubernetes (VM) Cluster Sync, Create, Delete, Manage, RBAC, Tenant Permissions
- Service Plan Sync, Tenant Permissions, RBAC
- Pricing Sync RBAC, Tenant Permissions, Markup
- Costing Sync, Reporting, Invoicing
- Reservations Sync, Guidance Recommendations
- Azure Stack Support
- Tag Bi-Directional Sync, Creation, Deletion Policy Enforcement
- Cost Estimator
- Azure US Gov Support
- Azure China Support
- Azure Germany Support
- CSP Account Support

Requirements

Morpheus Azure Integration requires Owner or Contributor access to subscription via App Registration. Adding an Azure Cloud or Clouds to Morpheus will require the following:

- Azure Subscription ID
- Directory (tenant) ID
- Application (client) ID
- Application (client) Secret
- Application (client) must be Owner or Contributor of Subscription

CSP Accounts require the additional following input:

- CSP Directory (tenant) ID
- CSP Application (client) ID
- CSP Application (client) SECRET

Credentials & Permissions

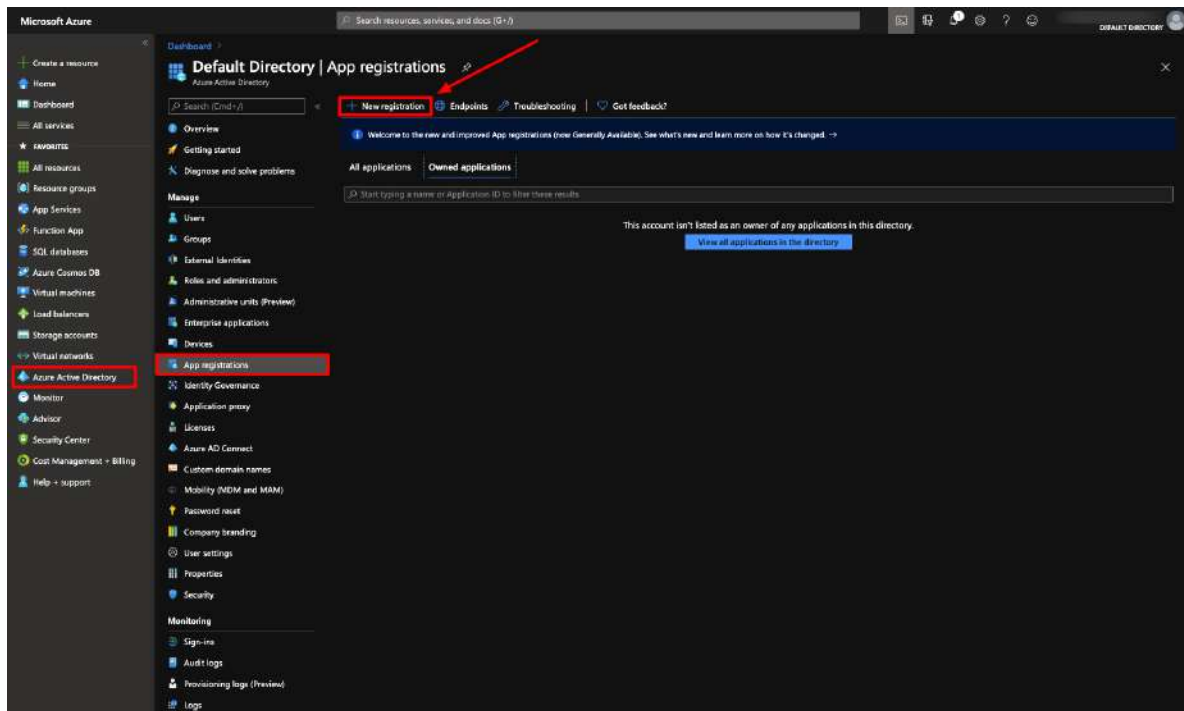
Morpheus authenticates with Azure via an App Registration with an Owner or Contributor Role on a Subscription. Use the steps below to create and collect the required credentials and assign the required permissions to integrate Azure with Morpheus.

Warning: Using an App Registration (service principal) that has selective resource permissions and is not an Owner or Contributor of the Subscription is not supported and will cause failures/issues. Please confirm the App Registration you use to integrate Azure with Morpheus has Owner or Contributor permissions on the specified Subscription before contacting support.

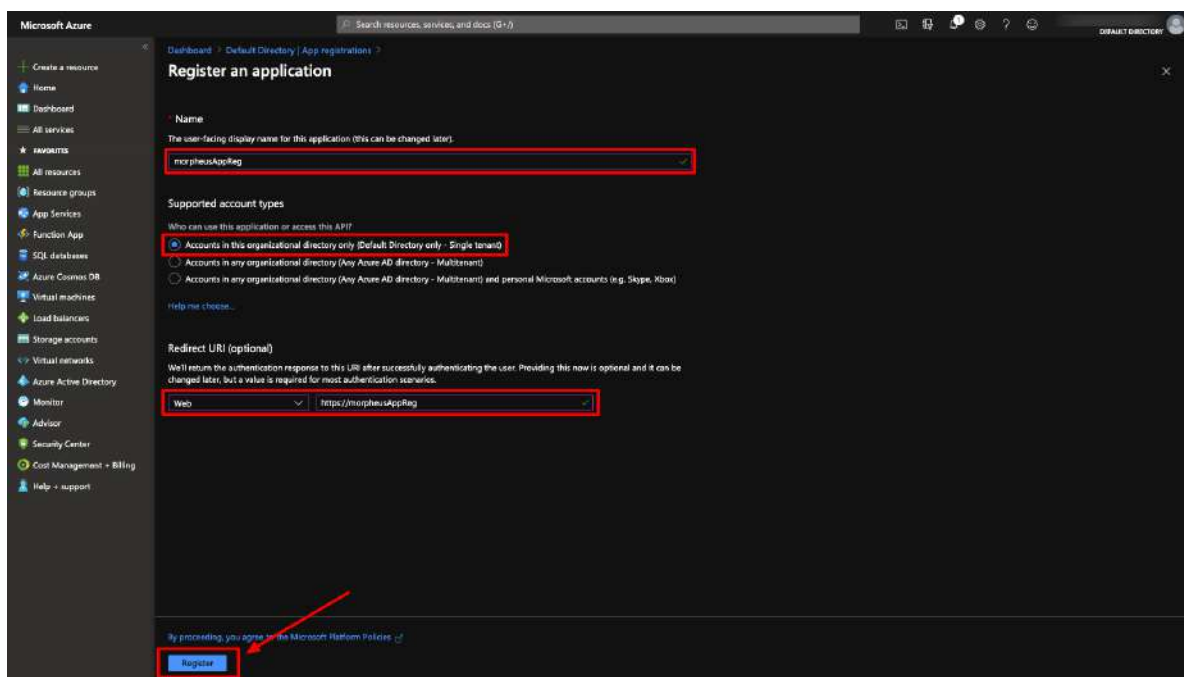
Create an App Registration

If you do not have an existing Azure Active Directory App Registration, or you wish to use an new one for Morpheus, you will need to create one.

1. Log into the Azure portal
2. Select “Azure Active Directory”
3. Select “App Registrations”
4. Select “New Registration”



- Next, give app a name, specify Web app / API for the type (default) and enter any url for the Sign-on URL:
- Click Create and your new App Registration will be created.

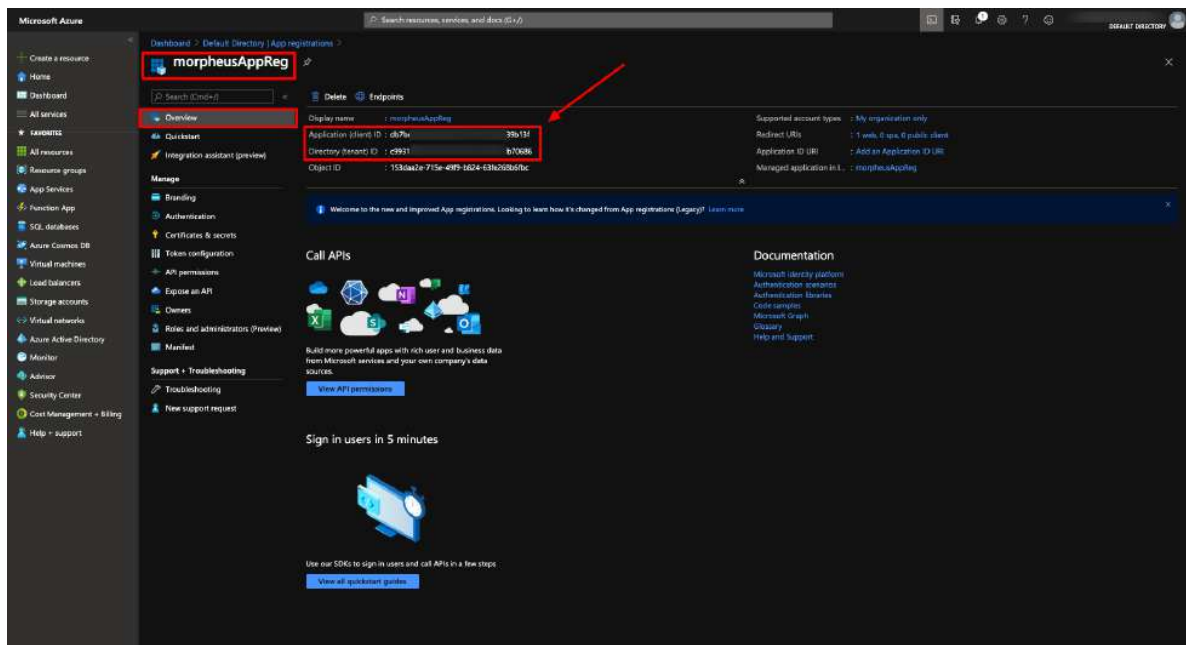


Now that we have (or already had) our App Registration, we will gather the credentials required for the Morpheus Azure integration.

Copy Directory (tenant) and Application (client) IDs

The App Registration Directory (tenant) and Application (client) ID are required for the Morpheus Azure integration. Both can be found in the overview section of the App Registration.

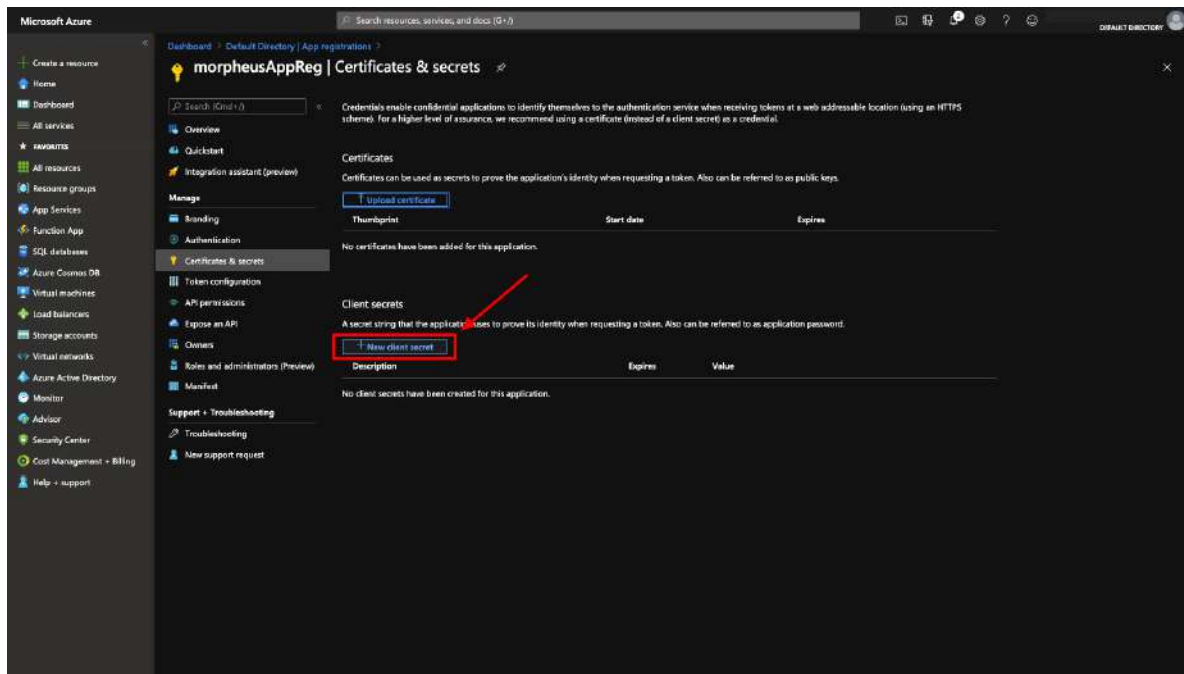
1. Go to the Overview section of your App Registration
2. Copy the Directory (tenant) ID
3. Store/Paste for use as the Tenant ID when Adding your Azure cloud in Morpheus
4. Copy the Application (client) ID
5. Store/Paste for use as the Client ID when Adding your Azure cloud in Morpheus



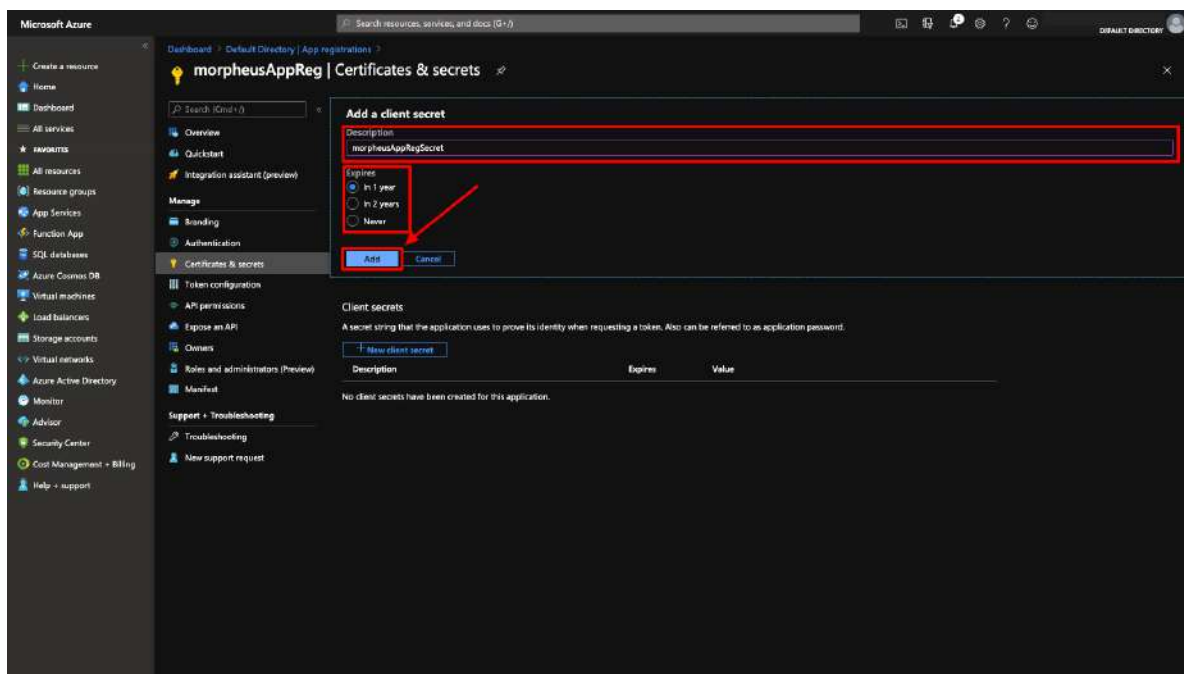
Generate a Client Secret

While still in your App Registration:

1. Select Certificates & secrets in the Manage Section
2. Select + New client secret

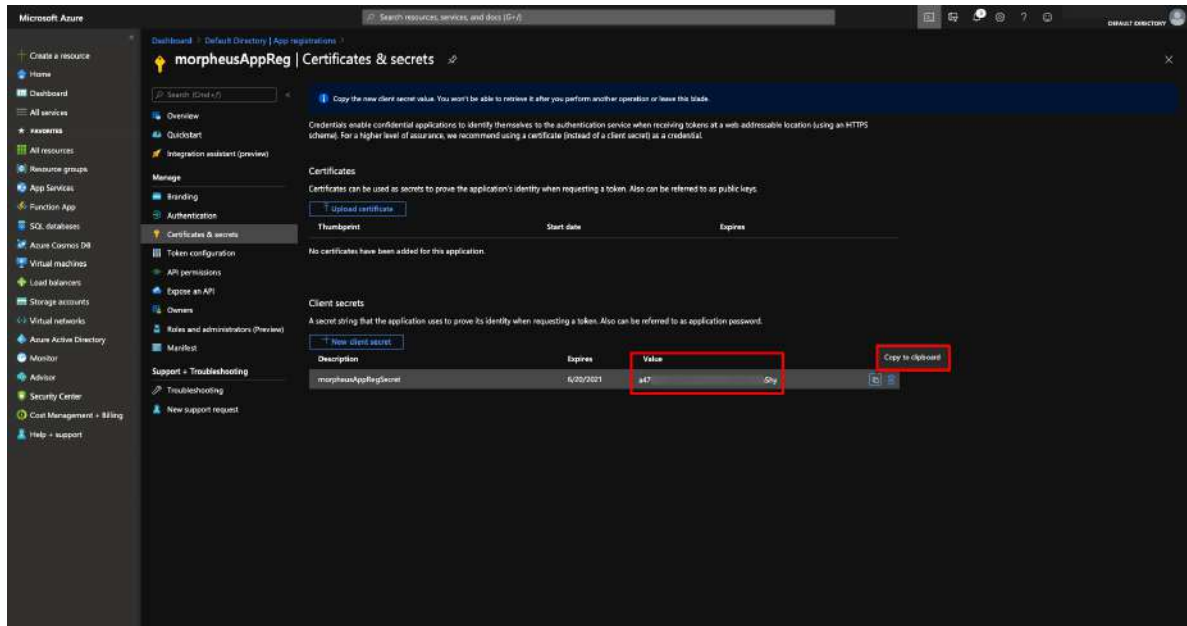


3. The “Add a client secret” modal will come up
4. Add a description to help identify the secret in the future
5. Select a duration
6. Select *Add*



7. Copy the newly generated Client Secret Value. It is important to copy the Client Secret Value now as it will not be displayed/available

Important: Copy the key value before continuing as it will not be displayed/available again.



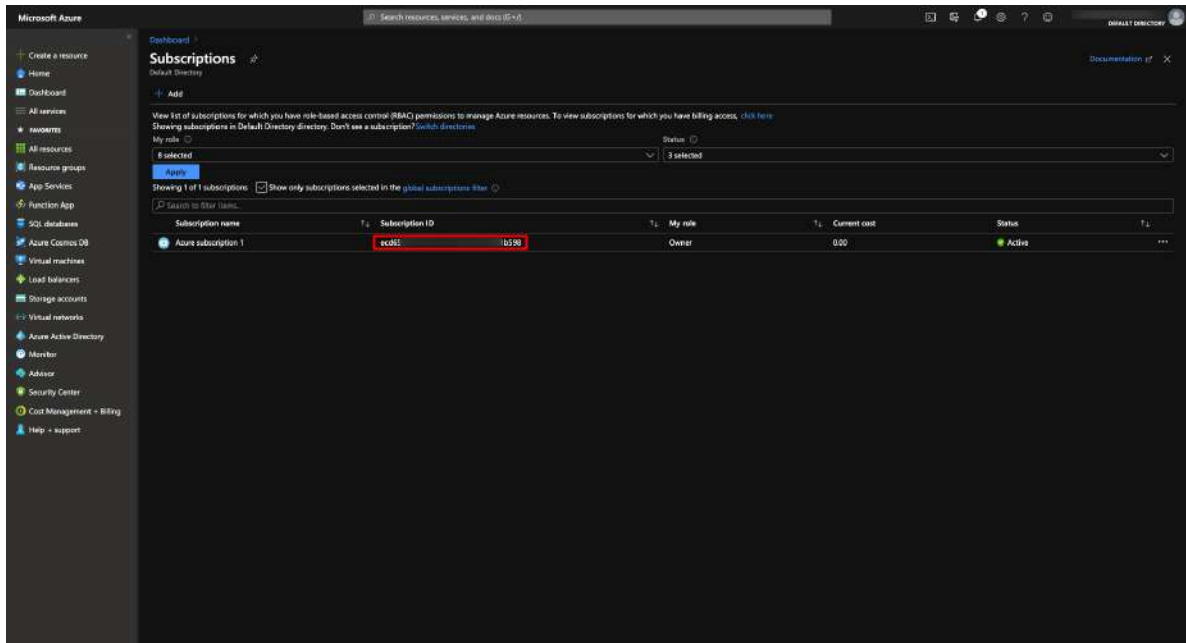
8. Store/Paste for use as the Client Secret when Adding your Azure cloud in Morpheus

You now have 3 or the 4 credentials required for Morpheus Azure cloud integration. The last credential required is the Azure Subscription ID.

Subscription ID

To get the Azure Subscription ID:

1. Navigate to the main Subscriptions section. One way is to search for “Subscriptions” and select Subscriptions in the search results
2. In the main “Subscriptions” section, copy the Subscription ID

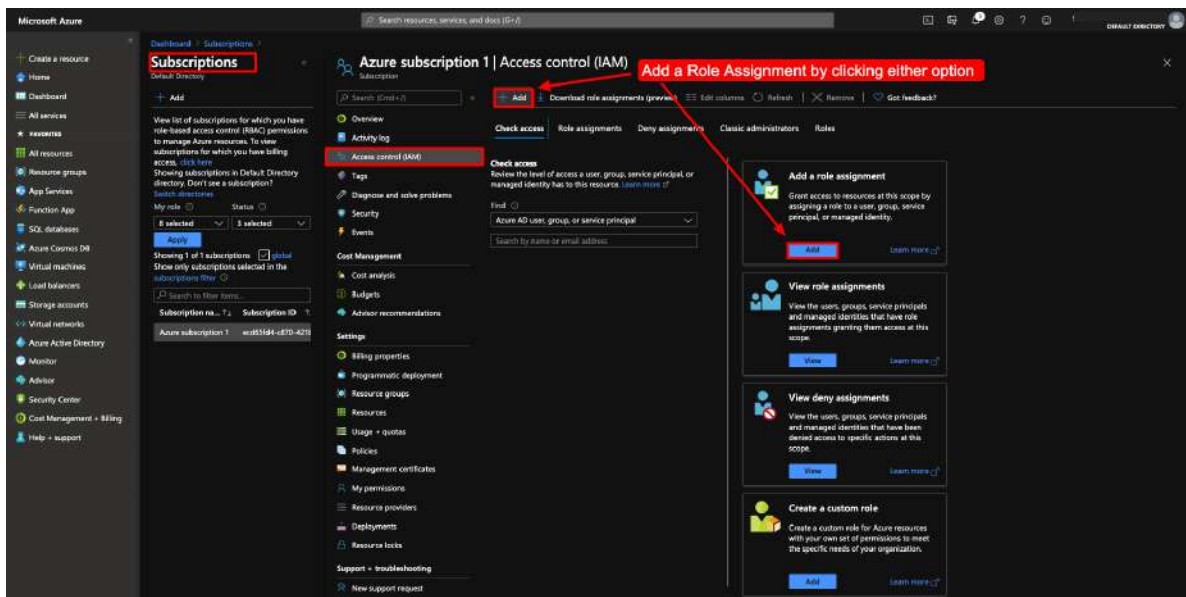


3. Store/Paste for use as the Subscription ID when Adding your Azure cloud in Morpheus

Make App Registration owner or contributor of Subscription

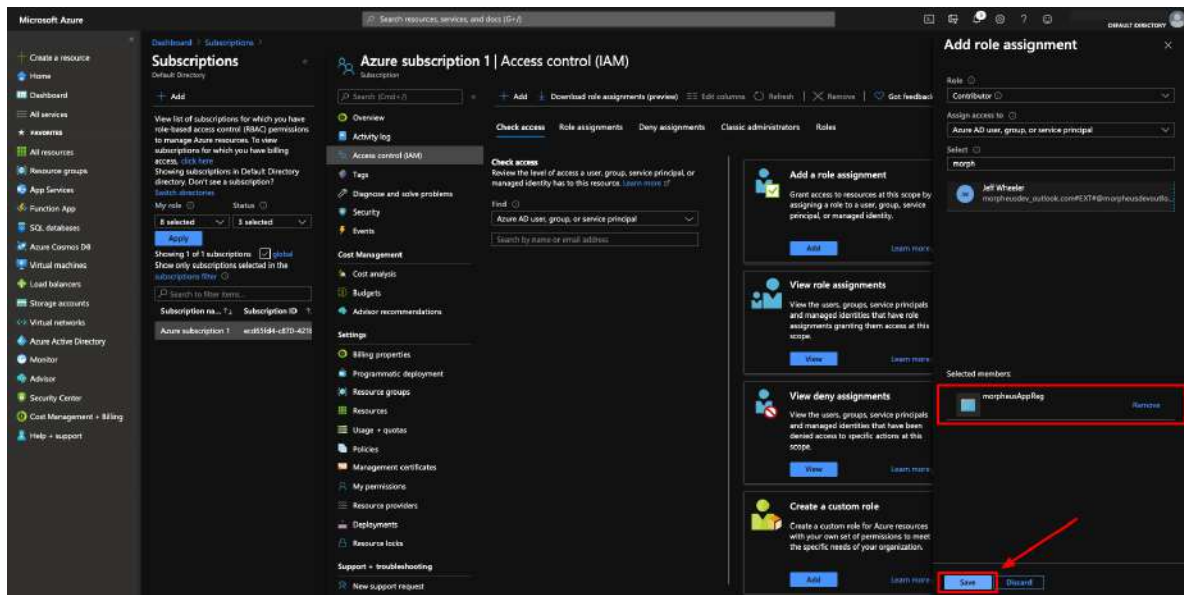
The App Registration created/used needs to be an owner of the Azure Subscription used for the Morpheus cloud integration. If lesser permissions are given or permissions are assigned at individual resource levels, Morpheus will not be able to properly inventory/sync, create and/or remove resources.

1. In the main “Subscriptions” section in Azure, select the Subscription
2. In the Subscription pane, select “Access Control (IAM)”
3. Either Click “+ Add”, and the “Add Role Assignment”, or simply select “Add a role assignment”



4. In the right pane, select “Owner” or “Contributor” Role type

5. Search for the name of the App Registration used for the Morpheus integration
6. Select the App Registration in the search results
7. Select “Save”

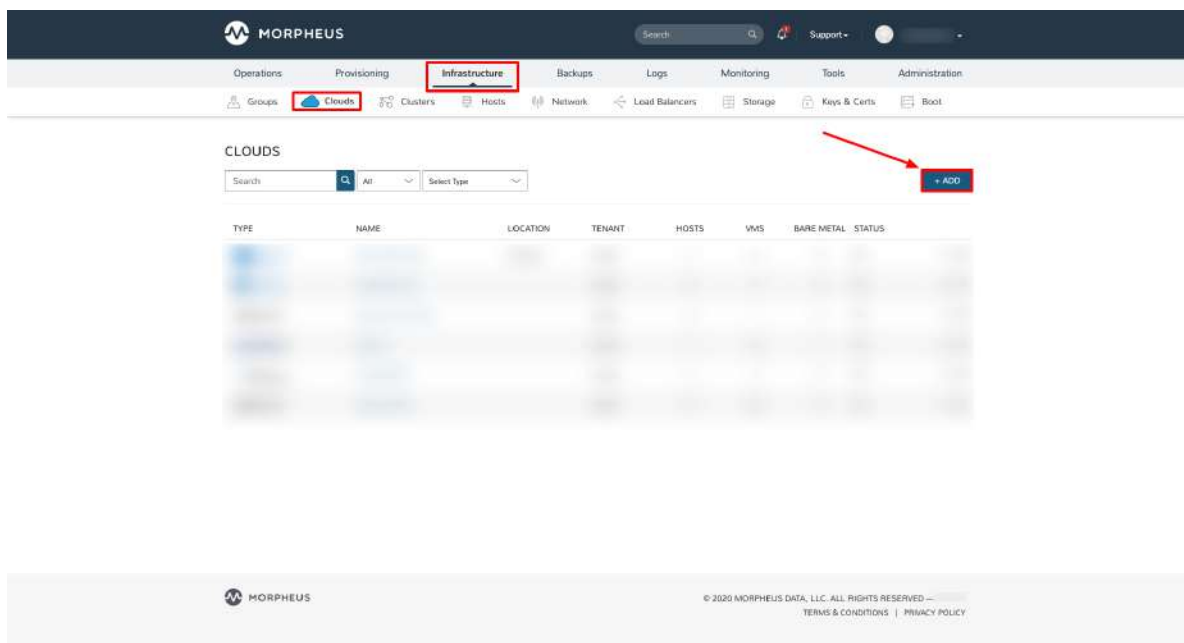


You now have the required Credentials and permissions to add an Azure Cloud Integration(s) into Morpheus.

Add an Azure Cloud Integration

To add a new Azure Cloud integration into Morpheus using the credentials created/collected from the previous section, perform the following:

1. In Morpheus, navigate to Infrastructure -> Clouds and select + ADD



2. Select “AZURE (PUBLIC)” from the Cloud Types list and click *NEXT*

CREATE CLOUD [X]

CLOUD CONFIGURE GROUP REVIEW

Search [Globe Icon] [Lock Icon]

	ALIBABA CLOUD Alibaba is a high performance , highly scalable global public cloud with large Chinese customers.
	AMAZON Amazon cloud
	AZURE (PUBLIC) Azure
	AZURE STACK (PRIVATE) Azure (Private)
	CLOUD FOUNDRY Cloud Foundry
	DELL Dell Server zone - manually managed servers
	DIGITALOCEAN Digital Ocean

PREVIOUS **NEXT**

3. Populate the Following

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provision-

ing.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

CLOUD TYPE

- Standard (Azure Cloud)
- US Gov (Azure US Government)
- German (Azure German Cloud)
- China (Azure China Cloud)

SUBSCRIPTION ID The target Azure Subscription ID obtained from the previous section

TENANT ID The Directory (tenant) ID obtained from the previous section

CLIENT ID The Application (client) ID obtained from the previous section

CLIENT SECRET The Application (client) Secret obtained from the previous section

LOCATION Once valid credentials are populate above and Morpheus is able to successfully authenticate with Azure, the available locations/regions will populate.

RESOURCE GROUP

- Select “All” to scope the Cloud to all available Resource Groups in the specified location/region.
- Select a single Resource Group to limit Morpheus resource creation, selection and discovery to just this Resource Group.

INVENTORY EXISTING INSTANCES Check to enable discovery/inventory of existing VM’s in the scoped Region and Resource Group(s)

INVENTORY LEVEL

Basic Morpheus will sync information on all resources in the selected Resource Group(s), including Name, IP Addresses, Platform Type, Power Status, and overall resources sizing for Storage, CPU and RAM, every 5 minutes. Inventoried VM’s will appear as Unmanaged VM’s.

Full (API Heavy) In addition to the information synced from Basic Inventory level, Morpheus will gather Resource Utilization metrics for Memory, Storage and CPU utilization per VM when available.

Off Existing VM’s will not be inventoried

ACCOUNT TYPE Standard, EA or CSP

Note: For CSP Accounts, also enter CSP TENANT ID, CSP CLIENT ID and CSP CLIENT SECRET in the Advanced Options section.

CREATE CLOUD

CLOUD

CONFIGURE

GROUP

REVIEW

NAME

newAzureCloud

CODE

codeIsUsefulForPoliciesAndAPI

LOCATION

West Central US

Details

CLOUD TYPE

Global

SUBSCRIPTION ID

ecdfb598

TENANT ID

c990686

CLIENT ID

db713f

CLIENT SECRET

.....

LOCATION

West Central US

RESOURCE GROUP

All

☒ INVENTORY EXISTING INSTANCES

INVENTORY LEVEL

Full (API Heavy)

ACCOUNT TYPE

Standard

► Advanced Options

► Provisioning Command

PREVIOUS

NEXT

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

4. Once done configuring the Cloud, select *NEXT*. NOTE all specified values except the Subscription ID can be changes after the Cloud is created.
5. Next select an existing Group to add the Azure Cloud to, or create a new Group, then select *NEXT*

CREATE CLOUD ✕

CLOUD > CONFIGURE > **GROUP** > REVIEW

☐ USE EXISTING ☒ CREATE NEW

Configuration

NAME

CODE

LOCATION

► Advanced Options

PREVIOUS NEXT

6. Review the configuration and then select *COMPLETE*

CREATE CLOUD

CLOUD

CONFIGURE

GROUP

REVIEW

Name: newAzureCloud

Location: West Central US

Domain: localdomain

Group: Azure

Visibility: private

Scale Priority: 1

Type: Azure (Public)

PREVIOUS

COMPLETE

Your new Azure Cloud integration will be created and begin to sync.

Note: The initial sync of an Azure Cloud can take some time due to Marketplace data sync.

MORPHEUS

Search

Support

Operations

Provisioning

Infrastructure

Backups

Logs

Monitoring

Tools

Administration

Group

Clouds

Clusters

Hosts

Network

Load Balancers

Storage

Keys & Certs

Boot

CLOUDS

Search

All

Select Type

+ ADD

TYPE	NAME	LOCATION	TENANT	HOSTS	VMS	BAFE METAL	STATUS
Microsoft Azure	newAzureCloud	West Central US	master	0	0	0	OK

MORPHEUS

© 2020 MORPHEUS DATA, LLC. ALL RIGHTS RESERVED — V. 4.2.2

TERMS & CONDITIONS | PRIVACY POLICY

Azure Stack

Overview

Azure Stack is Microsoft's Azure Cloud for on-premises environments. Azure Stack contains the core Azure services, allowing organizations to take advantage of Azure's offerings with the security, compliance, and financial benefits of hosting it in their own data-centers.

- Virtual Machine Provisioning
- Backups / Snapshots
- Resource Group Sync & Selection
- Network Sync & Selection
- Security Group Sync & Selection
- Storage Account Sync & Selection
- Marketplace Search and Provisioning
- Remote Console
- Periodic Synchronization
- Lifecycle Management and Resize
- Availability Set Support
- Azure Load Balancers
- Azure Storage
- Docker Host Provisioning & Management
- Service Plan Sync
- Pricing Sync with markup options
- Cost Estimator

Combine these features with public Azure and Morpheus can provide a single pane of glass and self service portal for managing instances scattered across both Azure offerings.

Requirements

Azure Stack Accessibility

By default, the Azure Stack management url's are not accessible from an external network. Port mappings and DNS must be configured for communication between the Morpheus Appliance and Azure Stack.

Important: In order to communicate with Azure Stack, Morpheus must be able to reach the internal Azure Stack network. The Azure Stack Portal needs to be exposed to the Morpheus Appliances' network with corresponding entries added to DNS.

One option to expose the Internal Azure Stack network to the Morpheus Appliances network is to use the `Expose-AzureStackPortal.ps1` powershell script from <https://gallery.technet.microsoft.com/scriptcenter/Expose-the-Azure-Stack-7ef68b19>. An Azure Stack Port Mapping Tool is also available.

Below is a sample output from the script for reference:

```
[Admin Portal] Created port mappings on 10.30.23.120 to 192.168.102.8
[Admin Portal] Ports: 13011 30015 13001 13010 13021 13020 443 13003 12646 12647 12648 ↵
↵12649 12650 12495 13026 12499
[Admin Portal] DNS: 10.30.23.120 - adminportal.local.azurestack.external ↵
↵adminmanagement.local.azurestack.external

[Tenant Portal] Created port mappings on 10.30.23.121 to 192.168.102.10
[Tenant Portal] Ports: 13011 30015 13001 13010 13021 13020 443 13003 12646 12647 ↵
↵12648 12649 12650 12495 13026 12499
[Tenant Portal] DNS: 10.30.23.121 - portal.local.azurestack.external management.local.
↵azurestack.external

[Blob Storage] Created port mappings on 10.30.23.122 to 192.168.102.4
[Blob Storage] Ports: 80 443
[Blob Storage] DNS: 10.30.23.122 *.blob.local.azurestack.external

VERBOSE: DNS delegation/forwarding is optional, change the DNS records on MAS-DC01 ↵
↵manually (dnsmgmt.msc from Host).
[DNS Delegation] Created port mappings on 10.30.23.120 to 192.168.200.224
[DNS Delegation] Ports: 53 (TCP/UDP)
[DNS Delegation] DNS: local.azurestack.external NS 10.30.23.120
[DNS Delegation] Change records on MAS-DC01 manually `if` you plan to use DNS ↵
↵forwarding.
[DNS Delegation] Change records back to the original internal IPs before running this ↵
↵script again.

VERBOSE: App Service detected and external IPs specified, creating mappings.
[App Service API] Created port mappings on 10.30.23.123 to 192.168.102.17
[App Service API] Ports: 443
[App Service API] DNS: 10.30.23.123 api.appservice.local.azurestack.external
[App Service Apps] Created port mappings on 10.30.23.124 to 192.168.102.15
[App Service Apps] Ports: 80 443 21 990
[App Service Apps] DNS: 10.30.23.124 *.appservice.local.azurestack.external
```

Azure Stack Resources

The following resources need to be created and configured inside Azure Stack for successful provisioning:

- Resource Group(s)
- Virtual Network(s)
- Storage Account(s)
- Network Security Group(s)
 - Inbound ports open from Morpheus Appliance: 22, 5985, 3389
 - Outbound ports open to Morpheus Appliance: 80, 443

Note: Proper Network and Network Security Group configuration is required for Morpheus agent install, communication, and remote console access. Other configurations, such as docker instances, will need the appropriate ports opened as well.

Required Credentials & Permissions

Credentials to integrate Morpheus with Azure Stack are located in both the public Azure Portal and the Private Azure Stack Portal. The Azure Active Directory Application used must be an owner of the Azure Stack subscription.

Azure Portal:

- Azure Active Directory Application Credentials
- Directory ID
- Management URL
- Identity Resource URL
- Application ID
- Key Value

Azure Stack Portal:

- Azure Stack Subscription ID
- Active Directory App from Azure portal added as owner of the Azure Stack Subscription in Azure Stack.

Adding an Azure Stack Cloud

Configure

1. In the Morpheus UI, navigate to Infrastructure -> Clouds and Select + *CREATE CLOUD*
2. Select *AZURE STACK (PRIVATE)* from the Clouds list and select *NEXT*
3. In the Configure section, enter:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

IDENTITY URL <https://login.microsoftonline.com>

MANAGEMENT URL* Azure AD Azure Stack Administrator app or Microsoft Azure Stack Administrator app url. Example: <https://adminmanagement.local.azurestack.external/>

IDENTITY RESOURCE URL Azure AD Azure Stack Administrator App ID URI Example: <https://adminmanagement.xxxxxxx.onmicrosoft.com/4a80e607-4259-4ac6-83e2-2fabeaf2eh83>

BASE DOMAIN This should match the base domain in your Management url. Example: local.azurestack.external

SUBSCRIPTION ID Subscription ID from Azure Stack portal (this is different from the Subscription ID in you Azure portal used when configuring Azure Stack)

TENANT ID This is the Directory ID from the Azure AD directory

CLIENT ID Application ID of Azure AD app with Azure Stack permissions granted, and has been added as an owner of the Azure Stack subscription (in the Azure Stack portal).

CLIENT SECRET Key Value of Application ID used above

Note: Once all credentials are entered and validated, the Location and Resource Group fields will populate.

Location Select an Azure Stack region for the cloud to scope to. This typically will be “local”.

Resource Group Select All or a single Resource Group to scope the cloud to. Selecting a single Resource Group will only sync resources in that Resource Group and disable Resource Group selection during provisioning. All will sync all resources and allow specifying the Resource Group during provisioning.

Inventory Existing Instances If enabled, existing Virtual Machines will be inventoried and appear as unmanaged Virtual Machines in Morpheus.

4. The Azure Stack cloud is ready to be added to a group and saved. Additional configuration options available:

Note: All fields and options can be edited after the Cloud is created.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

5. Once all options are configured, select NEXT to add the cloud to a Group.

A Group must be specified or created for the new Cloud to be added to. Clouds can be added to additional Groups or removed from Groups after being created.

USE EXISTING Add the new Cloud to an exiting Group in Morpheus .

CREATE NEW Creates a new Group in Morpheus and adds the Cloud to the Group.

6. Confirm all settings are correct and select COMPLETE. The Azure Stack Cloud will be added, and Morpheus will perform the initial cloud sync of:
 - Virtual Machines (if Inventory Existing Instances is enabled)
 - Networks
 - Virtual Images/Blueprints

- Network Security Groups
- Storage Accounts
- Marketplace Catalog
- Availability Sets

Tip: Synced Networks can be configured or deactivated from the Networks section in this Clouds detail page, or in the *Infrastructure -> Networks* section.

Cloud Foundry

Configuration

Adding PCF Cloud From *Infrastructure -> Clouds*

1. Navigate to *Infrastructure -> Clouds*
2. Select + *ADD*
3. Select **CLOUD FOUNDRY** from the Clouds list
4. Select *NEXT*
5. Populate the following:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

API URL Cloud Foundry API Url

CLIENT ID Typically `cf`

CLIENT SECRET Typically blank

USERNAME Enter Username. If using an API Key, enter `apikey` for username, and the API Key as the password.

PASSWORD Enter Password. If using an API Key, the API Key as the password.

ORGANIZATION Select Organization. Dropdown populates upon successful authorization.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

6. Select *NEXT*
7. Select an existing or create a new Group to add the Cloud to. The Cloud can be added to additional Groups in a Groups *Clouds* tab.
8. Select *NEXT*
9. Review and then Select *COMPLETE*

Adding PCF Cloud From *Infrastructure -> Groups*

1. Navigate to *Infrastructure -> Groups*
2. Select a Group
3. Select the *CLOUDS* tab
4. Scroll down to CLOUD FOUNDRY and select + *ADD*
5. Populate the following:

Name Name of the Cloud in Morpheus

Location Description field for adding notes on the cloud, such as location.

Visibility For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT Select a Tenant if Visibility is set to Private to assign to Cloud to that Tenant. Multiple Tenants can be added by editing the cloud after creation.

API URL Cloud Foundry API Url. Example `https://api.cf.morpheusdata.com`

CLIENT ID Typically `cf`

CLIENT SECRET Typically blank

USERNAME Enter Username. If using an API Key, enter `apikey` for username, and the API Key as the password.

PASSWORD Enter Password. If using an API Key, the API Key as the password.

ORGANIZATION Select Organization. Dropdown populates upon successful authorization.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

6. Select *NEXT*
7. Review and then Select *COMPLETE*

Adding Spaces

Cloud Foundry Spaces are referred to as Resource Pools in Morpheus. You can add a new Space by:

1. Navigating to the Cloud and selecting the Resources tab.
2. Then, click :guilabel: '+ Add Resource'.
3. Give the Resource a Name
4. Expand the Managers, Developers, and Auditors section to add specific Cloud Foundry users to the roles. When adding a user to these sections, use their Cloud Foundry email addresses.

Provisioning

Morpheus automatically seeds MySQL, Redis and RabbitMQ PCF Instance Types, as well as a generic Cloud Foundry Instance Type that will create a shell app used in conjunction with deployments. PCF Marketplace items can also be added to the Provisioning Library in the Cloud detail view Marketplace tab. The Marketplace item will be added to the selected Instance Type and available when selecting the Cloud Foundry Cloud during Instance or App Template creation.

Deployments

The Cloud Foundry App Instance Type is used in conjunction with deployments. Users do not have to pick deployment when creating a Cloud Foundry App Instance Type, but then Instance will only be a shell of a Cloud Foundry Application.

A deployment in Morpheus can either point to a git hub repository or contain the actual manifest.yml and associated artifacts required for a Cloud Foundry deployment. During the deployment, Morpheus will gather up the files required. Therefore, if the deployment points to a git hub repository, Morpheus will fetch the files from git hub. Once the files are obtained, Morpheus will deploy the artifacts in a similar fashion to the Cloud Foundry cli. This includes parsing the manifest to obtain the parameters to create or update the Cloud Foundry application. Morpheus will ignore certain fields such as memory and disk size because they are dictated by the selected plan. Other fields are utilized such as routes. After parsing the manifest.yml file (including overwriting certain fields), Morpheus is ready to update or create the App in Cloud Foundry.

After the App is configured, the artifacts references in the Morpheus deployment are uploaded to Cloud Foundry for the App. Note that when paths are referenced in the manifest.yml file, the paths continue to be relative to the manifest. So, a jar file under build/libs would need to be found under the build/libs directory.

If Cloud Foundry services are specified in the manifest, they must already exist within Cloud Foundry. Morpheus App templates can be utilized to wire up Cloud Foundry services created by Morpheus. In this case, Morpheus will add all of the included service names defined in the App template to the manifest.yml services section. Therefore, multiple services can be used and wired up by Morpheus."

Example

To better understand how Morpheus parses the manifest.yml file, let's take a closer look at the Cloud Foundry 'spring-music' project. The project can be found here (<https://github.com/cloudfoundry-samples/spring-music>).

The project contains the required manifest.yml file as well as the source code and build.gradle file to define how the project is to be built. After downloading the project to your local machine, build the project to generate the jar.

Now, let's take a look at the manifest.yml file:

```
---
applications:
- name: spring-music
  memory: 1G
  random-route: true
  path: build/libs/spring-music.jar
```

Using the Cloud Foundry docs (<https://docs.cloudfoundry.org/devguide/deploy-apps/manifest.html>), we can gain a better understanding of how this file is utilized by Cloud Foundry.

- The `-name` parameter defines the name that will be given to the application in Cloud Foundry. Morpheus will overwrite this value with the name given to the Instance being created in Morpheus.
- The `-memory` parameter (as well as the `disk_quota` parameter if specified) will be overwritten by Morpheus based on the plan specified for the Instance.
- The `-path` parameter defines, where relative to the manifest.yml file, your Cloud Foundry application can be found.
- The `-random-route` parameter, as well as all other parameters described in the Cloud Foundry documentation will simply be passed through to Cloud Foundry.

Adding Marketplace Items

1. Navigate to Infrastructure -> Clouds and select your Cloud Foundry Cloud
2. Select the MARKETPLACE tab
3. Select + *ADD MARKETPLACE ITEM*
4. Select the Morpheus Instance Type to add the Marketplace Item to.
5. Enter version
6. Search for and select Marketplace Item
7. Select *SAVE CHANGES*

A Node Type and layout will be created in the Provisioning -> Library section and the layout will be automatically added to the Instance Type selected when adding the Marketplace Item.

Provisioning Instances

Seeded and Marketplace Items

Morpheus automatically seeds MySQL, Redis and RabbitMQ PCF Instance Types, and PCF Marketplace items can also be easily added to the Provisioning Library in the Cloud detail view Marketplace tab. The Marketplace item will be added to the selected Instance Type and available when selecting the Cloud Foundry Cloud during Instance or App Template creation.

1. Navigate to `Provisioning -> Instances` and select an Instance Type with a Cloud Foundry layout (MySQL, Redis and RabbitMQ plus Marketplace additions)
2. Select *NEXT*
3. Select a Group and PCF Cloud
4. Add an Instance Name
5. Optionally select an Environment Tag and/or add a custom Tag
6. Select *NEXT*
7. Select Version and Instance Configuration for a Cloud Foundry layout, ex: *Cloud Foundry MySQL*
8. Select a Plan and available options for the Plan, or use the custom Plan
9. Select a Space to add the Instance to
10. Optionally configure advanced options
11. Select *NEXT*
12. Optionally configure Automation options
13. Select *NEXT*
14. Select *COMPLETE*

Note: Compute, Memory, and CPU stats will be pulled, and a Cloud Foundry monitoring health check will be automatically configured for the instance.

Cloud Foundry App Instance Type

Important: Add Deployments in Provisioning -> Deployments to be used when provisioning a Cloud Foundry App Instance Type.

Note: Minimal options are outlined below.

1. Navigate to `Provisioning -> Instances` and select the *Cloud Foundry App* Instance Type
2. Select *NEXT*
3. Select a Group and PCF Cloud
4. Add an Instance Name
5. Optionally select an Environment Tag and/or add a custom Tag

6. Select *NEXT*
7. Select a Plan and available options for the Plan, or use the custom Plan
8. Select a Space to add the Instance to
9. Select *NEXT*
10. In the Deployments section, select a Deployment and Version to be deployed. These can be git repos or files added in Provisioning -> Deployments

Important: If services are specified in a git repo manifest, Morpheus assumes they already exist in the PCF cloud with matching names.

11. Select *NEXT*
12. Select *COMPLETE*

This will quickly create the Cloud Foundry Application, and then the deployment will follow which may take longer depending on the app configuration. The location will be updated with the route once it is configured.

Note: Compute, Memory, and CPU stats will be pulled, and a Cloud Foundry monitoring health check will be automatically configured for the instance.

Digital Ocean

Add a Digital Ocean Cloud

DigitalOcean Cloud Integration Detail fields:

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

Username DigitalOcean Username

API Key Personal access tokens/Key from the DigitalOcean *API* -> *Tokens/Keys* section.

Data Center Select DigitalOcean DataCenter Region

The Cloud can now be added to a Group or configured with additional Advanced options.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin* -> *settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.

- **Cloud Init / Unattend** (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

ESXi

The ESXi Cloud type enables managing and provisioning to ESXi hosts, even without the ESXi API enabled.

Important: The VMware ESXi integration is for adding a single ESXi / vSphere Hypervisor host. If you have vCenter please use the VMWare vCenter cloud type for full vSphere integration features.

To get started with VMware ESXi, simply add a VMware ESXi Cloud in either the *Infrastructure -> Clouds* or *Infrastructure -> Groups* section.

1. Select + Create Cloud Button
2. Select ESXi from the Add Cloud modal
3. Select NEXT
4. Provide the following information.

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

- ESXi Host name or IP address
- Username (This is normally root)
- Password

Note: If you receive the message “Error! Invalid cloud config” Please ensure you have ssh enabled on the ESXi host.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Important: ESXi provisioning require a vmx file, which is not included in an OVF/OVA export from vCenter. A proper vmx file must be included when adding a vmdk/ovf/ova image to Virtual Images in Morpheus for successful provisioning.

Google Cloud Platform (GCP)

Integration Features

- Provisioning Virtual Machines
- Network tagging
- Private and Local Images
- Google VM Snapshots
- Brownfield Inventory
- Costing
- Right-sizing
- Shared Network Support

Requirements for Integration with Morpheus

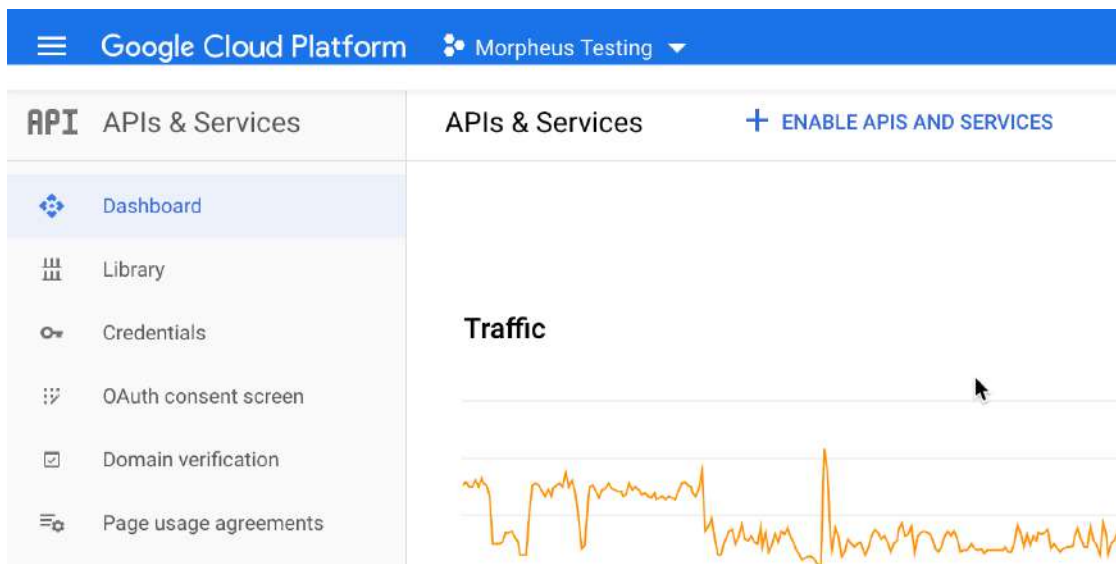
To integrate Morpheus with Google Cloud Platform, you will need the following:

- The Compute Engine API enabled in GCP “APIs & Services”
- Credentials for an IAM service account with Owner or Compute Admin role permissions
- The Project ID, private key, and client email for the service account

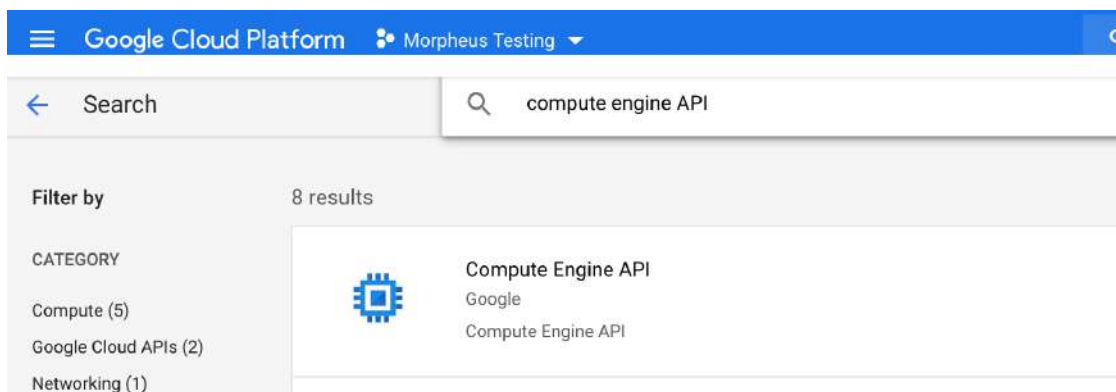
This integration guide goes through the process of configuring your account and obtaining the information necessary to integrate with Morpheus.

Enabling the Compute Engine API

1. Log into the Google Cloud Platform web console
2. Hover over the “APIs & Services” menu and click on Dashboard
3. Click + *ENABLE APIS AND SERVICES*



4. In the search bar, search for “Compute Engine API”

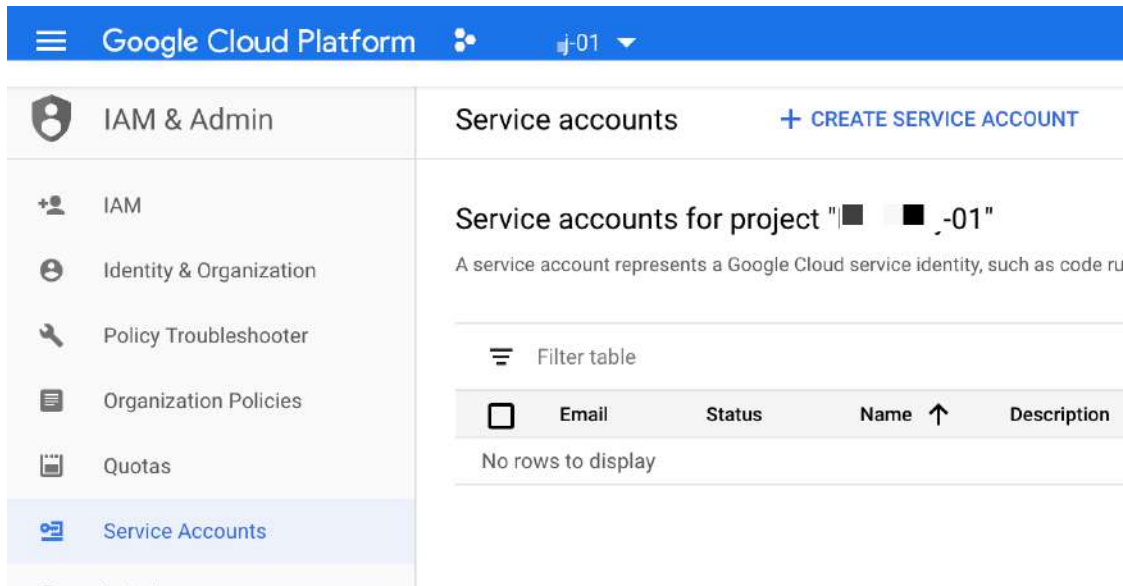


5. Select “Compute Engine API” and click *ENABLE*. It may take a few moments for the API to be fully enabled

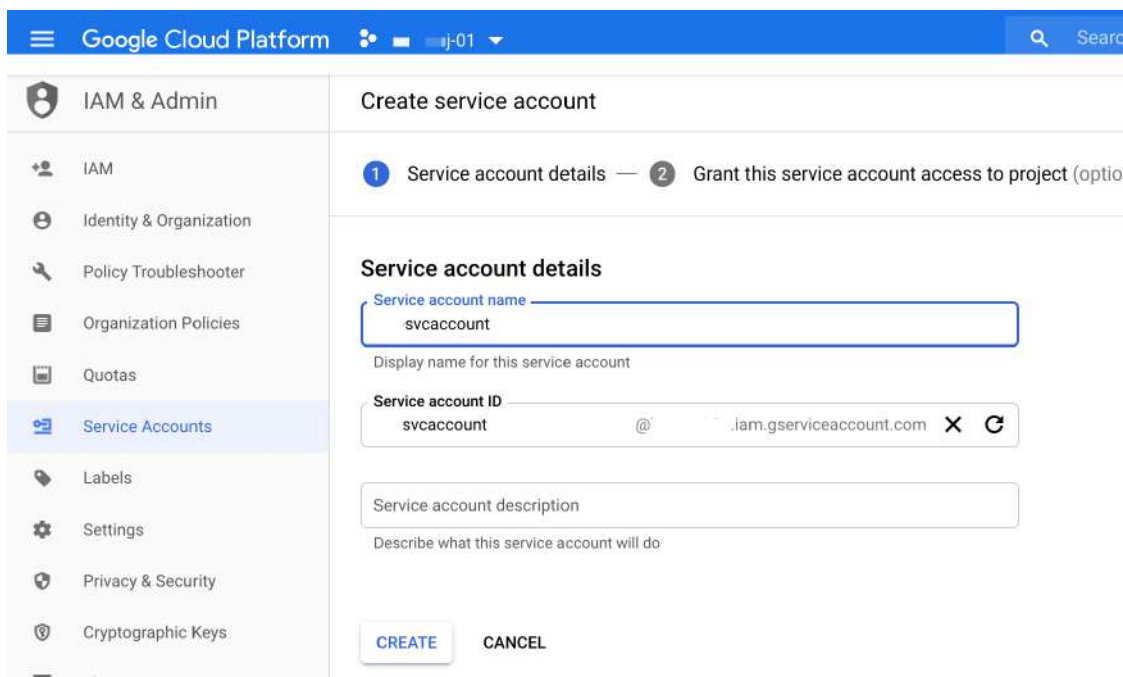
Note: If the button is labeled MANAGE rather than ENABLE, the API is already enabled.

Creating a Service Account

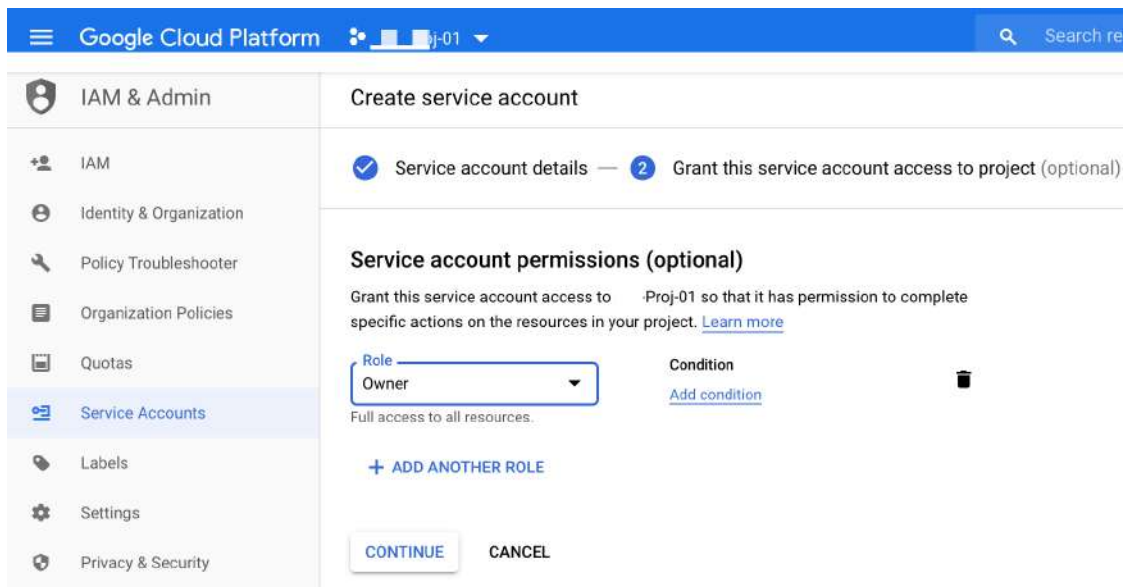
1. From the landing page of the GCP web console, hover over the “IAM & Admin” menu and click on “Service Accounts”
2. Click + *CREATE SERVICE ACCOUNT*



3. Enter at least a name for your new service and click CREATE



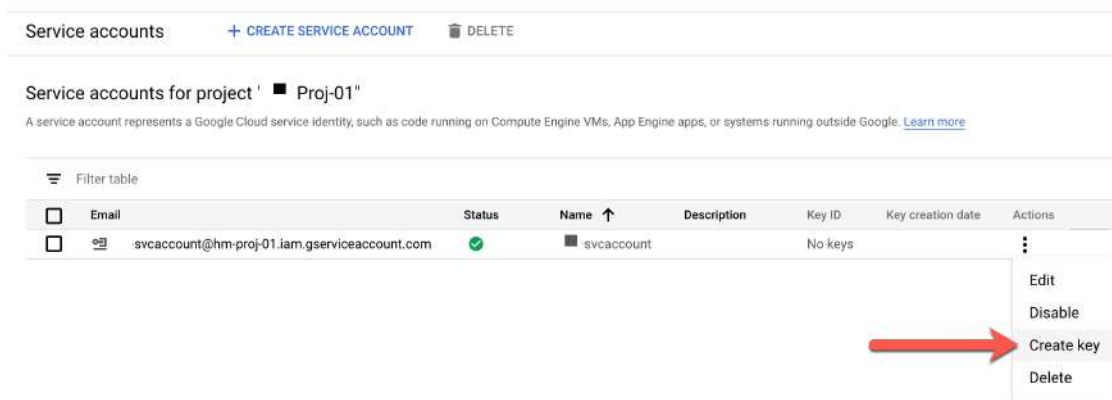
- After creating the service account, you'll be prompted to set a role for the account. In order to fully integrate with Morpheus, you must use an account in the Owner role or the Compute Admin role
- Click CONTINUE



- Following creation of the service account, you'll be taken back to the list of existing service accounts

Generating Keys and Integrating with Morpheus

- From the list of service accounts, click the ellipsis button (...)
- Click "Create Key"



- Select JSON format and click CREATE
- A JSON-formatted document will be downloaded, this document contains the Project ID, private key, and client email values needed to complete the integration process in the next step

Add a GCP Cloud

Note: The JSON-formatted document downloaded when creating a key for your service account contains all of the required values for completing the integration. Consult the above section on generating keys if needed.

1. Navigate to Infrastructure > Clouds
2. Select + *CREATE CLOUD*, select Google Cloud, and then click *NEXT*.
3. Enter the following into the Create Cloud modal:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

PROJECT ID Google Cloud Project ID

PRIVATE KEY The service account private key. Paste in the entire value between (but not including) the quotation marks in your downloaded JSON document, formatted like the following example: —BEGIN PRIVATE KEY—(your_key)—END PRIVATE KEY—.

CLIENT EMAIL The service account client email, ex: *morpheus@morpheus.iam.gserviceaccount.com*

REGION Regions will auto-populate upon successful authentication with the above credentials. If no regions are found, double check your entered credentials and try again. Select the appropriate region for this Cloud

INVENTORY EXISTING INSTANCES If checked, existing Google Instances will be inventoried and appear as unmanaged virtual machines in Morpheus.

Note: Morpheus scopes Clouds to single regions. Multiple clouds can be added for multi-region support, and then optionally added to the same group.

If advanced options are not needed, click *NEXT* to advance to the Group selection page. Otherwise, continue on with this guide and review advanced or provisioning options.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

4. After reviewing all options, click *NEXT* to advance to the Group selection page. Following Group selection, click *COMPLETE* to finish the integration process. If you've opted to inventory existing Instances, they will be viewable in Morpheus shortly. At this point, you are ready to provision new resources in Google Cloud Platform as needed!

Important: If you experience difficulties adding a GCP Cloud, review the above guide and ensure you've met all requirements for completing the integration. For example, if the Compute Engine API is not enabled, Morpheus will not accept credentials entered on the Create Cloud modal. If you repeatedly run into problems completing the integration process, review the above guide in its entirety and double check that each step is completed and your account meets all configuration requirements.

Windows Images

Morpheus can add custom metadata that will be injected into the unattend conf by GCP during provisioning. This is required for customizations including setting the Windows Administrator password during provisioning. GCP Windows Images must be syspreped using the `GCESysprep` command prior to image creation, and must have platform/os set on the Virtual Image record in Morpheus after image sync for successful customization and Agent Installation.

GCP Windows Requirements

- GCP Windows Images must be syspreped using the `GCESysprep` command prior to Image creation in GCP. Refer to Google's "creating-windows-os-image" doc.
- Once the Image is synced into Morpheus, the Platform (Windows, Windows 2016 etc) must be set on the Morpheus Virtual Image record, otherwise linux is assumed and the metadata will not be generated correctly.
- The Global Windows "Administrator" password must be set in Morpheus under `/admin/provisioning/settings -> Windows Settings -> Administrator Password`, or Administrator and password defined on the Morpheus Virtual Image record.
- Be aware the unattend configuration during startup after sysprep delays causes a reboot and a prolonged finalization process during provisioning, and console/rdp may not be available during this time as windows is configuring.

Note: Some Google provided Windows Images have slow startups that cause the Morpheus Agent service to not start within the default 30 second service startup timeframe, including after initial reboot after sysprep/unattend configuration. This can be adjusted by running `New-ItemProperty -Path "HKLM:\SYSTEM\CurrentControlSet\Control\" -Name "ServicesPipeTimeout" -PropertyType DWORD -Value 180000` in powershell on the Windows Image.

Important: Failure to use a GCP Windows Image that has not been sysprepped using `GCESysprep` will cause Agent Installation, Automation, and Console issues as Morpheus will not be able to set user credentials and authenticate.

Huawei Cloud

Features

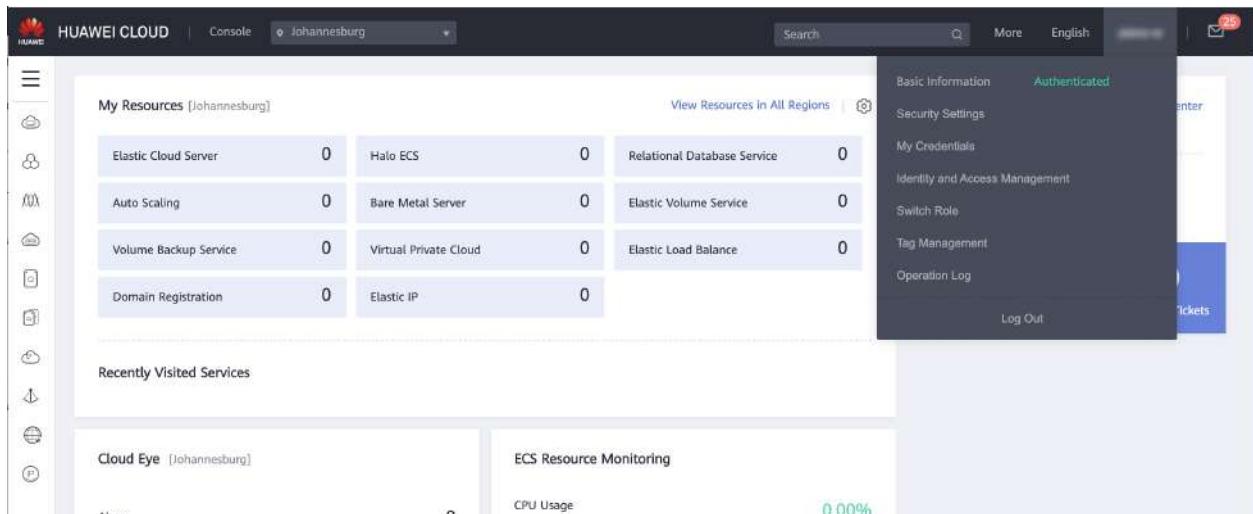
- Virtual machine provisioning
- Backups
- Brownfield VM management and migration
- Hypervisor remote console
- Cloud sync
- Lifecycle management and resizing
- Network security group creation
- Network security group management
- Router and network creation
- Load balancer services
- Docker host management and configuration
- Floating IP assignment
- Huawei OBS buckets (create, manage, delete, and discovery)
- Huawei SFS (create, manage, and delete)

Integrate Huawei Cloud with Morpheus

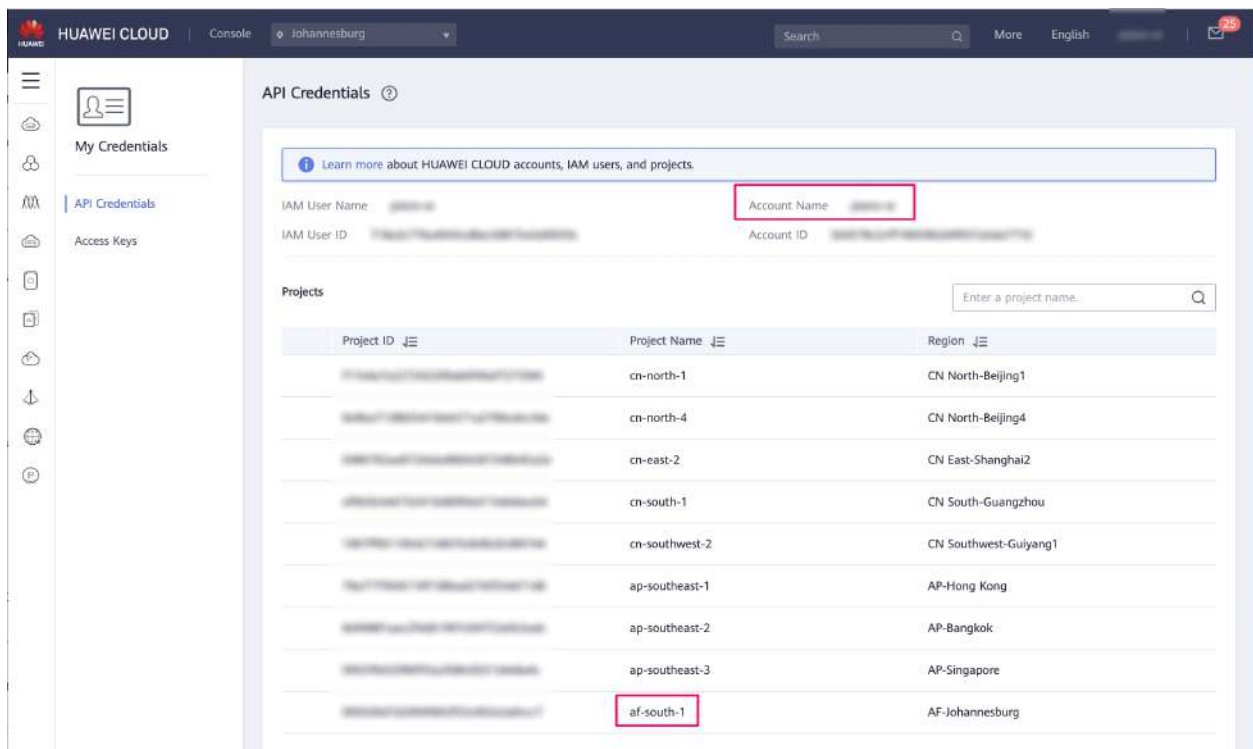
To integrate Huawei Cloud with Morpheus, we'll gather the following pieces of information:

- Account Name
- Identity (IAM) API URL
- Project
- Username
- Password

Begin by logging into your [Huawei Cloud console](#). If you're not currently logged in, you will be prompted to do so. Once on the console page, hover over your username in the upper-right corner of the application window and select "My Credentials".



From the credentials page, we can gather the Account Name and the Project Name, record them for later when we provide the integration information to Morpheus.



To gather the API endpoint URL, take a look at the complete list of [endpoints](#). If a specific endpoint exists for your region, use it. In any other case use the endpoint for all regions. It will be formatted like this: <https://iam.myhuaweicloud.com/v3>.

Identity and Access Management (IAM)

Region Name	Region	Endpoint	Protocol Type
ALL	ALL	iam.myhuaweicloud.com	HTTPS
CN East-Shanghai1	cn-east-3	iam.cn-east-3.myhuaweicloud.com	HTTPS
CN East-Shanghai2	cn-east-2	iam.cn-east-2.myhuaweicloud.com	HTTPS
CN North-Beijing1	cn-north-1	iam.cn-north-1.myhuaweicloud.com	HTTPS
CN North-Beijing2	cn-north-2	iam.cn-north-2.myhuaweicloud.com	HTTPS
CN North-Beijing4	cn-north-4	iam.cn-north-4.myhuaweicloud.com	HTTPS
CN South-Guangzhou	cn-south-1	iam.cn-south-1.myhuaweicloud.com	HTTPS
CN South-Shenzhen	cn-south-2	iam.cn-south-2.myhuaweicloud.com	HTTPS
CN Southwest-Guian g1	cn-southwest-2	iam.cn-southwest-2.myhuaweicloud.com	HTTPS

With this information gathered, and presuming you know the credentials for the service account you wish to use, we can move back into Morpheus-UI.

Navigate to Infrastructure > Clouds and click + *ADD*. Scroll to Huawei Cloud and click *NEXT*. The information we've gathered will be plugged into the CREATE CLOUD modal. The DOMAIN ID field will accept the Account Name field we gathered. Your completed CREATE CLOUD modal will look similar to the one pictured below:

Details

IDENTITY API URL	<input type="text" value="https://iam.myhuaweicloud.com/v3"/>
DOMAIN ID	<input type="text" value="my_account_name"/> <small>This pertains to the OpenStack V3 API and should be ignored when using V2. This is the Domain ID (Not to be confused with Domain Name). The Domain ID can be found via the CLI by typing <code>openstack domain list</code>.</small>
PROJECT	<input type="text" value="af-south-1"/>
REGION	<input type="text"/>
USERNAME	<input type="text" value="my_username"/>
PASSWORD	<input type="password" value="....."/>
IMAGE FORMAT	<input type="text" value="QCOW2"/>

After clicking *NEXT*, add this new Cloud to a Group or create a new Group. On finalizing the wizard, Huawei Cloud will be integrated into Morpheus and ready for provisioning. If you opted to inventory existing workloads, those will be onboarded shortly.

Add/Edit Huawei Cloud Modal Fields

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

IDENTITY API URL The v2 or v3 identity endpoint. See the integration steps above for more detail

DOMAIN ID The DOMAIN ID field takes the “Account Name” as shown on the [Basic Information page](#) of the account. See the integration steps above for more detail

PROJECT The target project name. See the integration steps above for more detail

USERNAME The service account username. See the integration steps above for more detail

PASSWORD The integration service account password. See the integration steps above for more detail

IMAGE FORMAT Select QCOW2, RAW or VMDK image type

Inventory Existing Instances Select for Morpheus to discover and sync existing VMs

Enable Hypervisor Console Hypervisor console support for openstack currently only supports novnc. Be sure the novnc proxy is configured properly in your openstack environment.

Tip: When using the RAW image format, you can bypass the image conversion service within the cloud leading to quicker performance. Other image formats are converted to RAW format and back when performing various actions. Using the RAW format from the start will bypass these conversion steps.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Huawei Scalable File Service (SFS)

The Morpheus integration with Huawei Cloud includes the capability to work with Huawei Scalable File Service (SFS). SFS is shared file storage hosted on Huawei Cloud. By integrating Morpheus with Huawei Cloud you can discover, create, manage, and delete SFS servers, as well as view and work with the file shares and files contained therein.

SFS Server Discovery and Management

On integrating Huawei Cloud with Morpheus, SFS servers and file shares are discovered automatically after a short time. The server(s) can be viewed in Infrastructure > Storage > Servers. By viewing the server detail page and clicking *EDIT*, the storage server can be scoped as needed. Administrators can choose to scope to other Huawei Cloud integrations (if more than one relevant integration currently exists), select from synced availability zones, and scope the storage server to specific Tenants if desired.

EDIT STORAGE SERVER

NAME

Labs Huawei SFS

DESCRIPTION

☒ ENABLED

TYPE

Huawei SFS

CLOUD

QA Huawei

AVAILABILITY
ZONE

Select

Permissions

VISIBILITY

Private

TENANTS

Search

SAVE CHANGES

Additionally, Huawei SFS servers can be created from the storage server list page (Infrastructure > Storage > Servers) directly in Morpheus. Click + *ADD* to begin and set the storage server type value to “Huawei SFS”. Just like with existing synced SFS servers, those created from Morpheus can be scoped as needed.

ADD STORAGE SERVER

X

NAME

New Huawei SFS

DESCRIPTION

☒ ENABLED

TYPE

Huawei SFS

CLOUD

QA Huawei

AVAILABILITY
ZONE

Select

Permissions

VISIBILITY

Private

TENANTS

Search

SAVE CHANGES

SFS File Share Discovery and Management

Discovered file shares will appear among other file shares synced with Morpheus in Infrastructure > Storage > File Shares. Depending on the number of cloud integrations in your Morpheus appliance and the number of cloud integrations available to your user account, this list may be quite large. Using the search bar on this page, we can narrow down the list to file shares displayed to those whose names match the search terms.

NAME	PROVIDER TYPE	SHARE PATH	SOURCE	BACKUP	DEPLOYMENTS
qa-huawei-sfs	Huawei SFS Share	/mnt/qa-huawei-sfs	Labs Huawei SFS	No	No
qa-huawei-sfs	Huawei SFS Share	/mnt/qa-huawei-sfs	Labs Huawei SFS	No	No
qa-huawei-sfs	Huawei SFS Share	/mnt/qa-huawei-sfs	Labs Huawei SFS	No	No
qa-huawei-sfs	Huawei SFS Share	/mnt/qa-huawei-sfs	Labs Huawei SFS	No	No
qa-huawei-sfs	Huawei SFS Share	/mnt/qa-huawei-sfs	Labs Huawei SFS	No	No
qa-huawei-sfs	Huawei SFS Share	/mnt/qa-huawei-sfs	Labs Huawei SFS	No	No

We can drill into individual file shares by clicking on their hyperlinked name in the list of all integrated file shares. From the file share detail page, a list of files will appear on the files tab. Begin the process of adding a new file by clicking + ADD. The Access tab on the file shares detail page allows users to view and manage ACL rules.

Note: A “Failed to load files from storage provider” is present when the Morpheus appliance doesn’t have access to the file share.

New Huawei SFS file shares can be created directly in Morpheus. From the file shares list page, get started by clicking + ADD. Select the type “Huawei SFS Share”. Set the storage service field to a pre-existing Huawei SFS server. A friendly name for the file share in Morpheus and selecting from synced availability zones are required fields.

NEW FILE SHARE



NAME Huawei File Share

STORAGE SERVICE Labs Huawei SFS

AVAILABILITY
ZONE ap-southeast-1a

SHARE SIZE 0 GB

☒ ACTIVE☐ DEFAULT BACKUP TARGET☐ DEFAULT DEPLOYMENT ARCHIVE TARGET☐ DEFAULT VIRTUAL IMAGE STORE

Retention

RETENTION
POLICY None

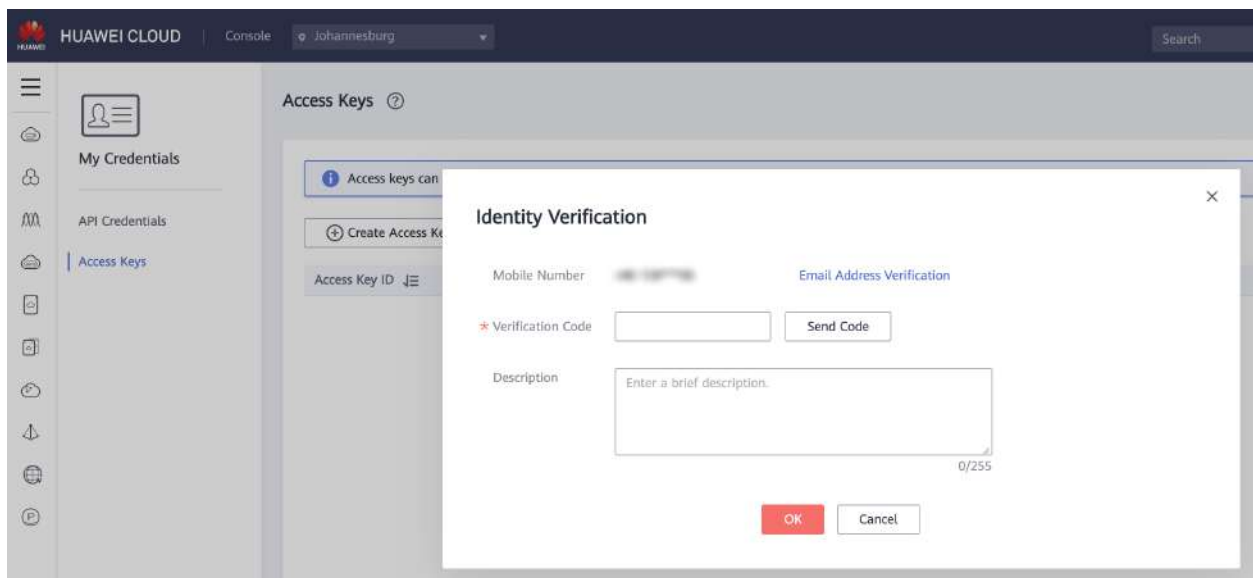
SAVE CHANGES

Huawei Object Storage Service (OBS)

The Morpheus integration with Huawei Cloud also supports Object Storage Service (OBS). Morpheus will automatically onboard existing OBS servers and buckets shortly after completing the cloud integration. Before you can add a new OBS server from Morpheus, you must know or generate a key and secret value from the Huawei console and must provide a Huawei OBS API endpoint.

Generate a Key and Secret

From the Huawei web console, log into the account used to integrate Huawei Cloud with Morpheus. Hover over your account name in the upper-right corner of the application window and click “My Credentials”. Select “Access Keys” from the left-hand sidebar. To create a new key, click + *Create Access Key*. Complete the two-factor authentication steps in the box that appears.



Once the key is generated, download or record the key and store it in a safe location. The key will not be viewable or available for download again after this point.

Create OBS Server in Morpheus

With the key and secret value in hand from the previous section, navigate to Infrastructure > Storage > Servers. Click + *ADD*. On changing the server type to Huawei OBS, you will see the fields for the access key and the secret key. OBS API endpoints can be found in [Huawei endpoint documentation](#). Include those three values in the Create Server modal along with a friendly name for use in Morpheus UI. Just like with SFS objects, we can choose to scope the server to all or specific Tenants at this time.

ADD STORAGE SERVER ×

NAME

DESCRIPTION

☒ **ENABLED**

TYPE

Huawei OBS ▼

ACCESS KEY

SECRET KEY

ENDPOINT

Permissions

VISIBILITY

Private ▼

TENANTS

Search

SAVE CHANGES

Create Huawei OBS Bucket

With an OBS server onboarded or created in Morpheus, you're able to create and manage Huawei OBS buckets as needed. To create a new bucket, navigate to Infrastructure > Storage > Buckets. Click + *ADD* and select "Huawei OBS Bucket". The following fields are required when creating a Huawei OBS bucket:

- **NAME:** A friendly name for use in Morpheus UI
- **STORAGE SERVICE:** Choose the OBS server to associate the new bucket with
- **BUCKET NAME:** The name of the bucket in Huawei Cloud, this must be unique
- **STORAGE CLASS:** If needed, view the [discussion of storage classes](#) in Huawei support documentation
- **BUCKET ACL:** Public Read, Public Read/Write, or Private
- **BUCKET POLICY:** Public Read, Public Read/Write, or Private
- **STORAGE QUOTA:** Set to 0 for no quota

Once finished, click *SAVE CHANGES*

NEW BUCKET

X

NAME

My OBS Bucket

STORAGE SERVICE

Labs Huawei OBS

BUCKET NAME

my-obs-bucket

STORAGE CLASS

Standard

BUCKET ACL

Private

BUCKET POLICY

Private

STORAGE QUOTA

0

GB

☐ CREATE BUCKET

☒ ACTIVE

☐ DEFAULT BACKUP TARGET

☐ DEFAULT DEPLOYMENT ARCHIVE TARGET

☐ DEFAULT VIRTUAL IMAGE STORE

Retention

RETENTION POLICY

None

SAVE CHANGES

Hyper-V

Hyper-V is the virtualized server computing environment introduced by Microsoft. Hyper-V is consumed by Morpheus as a private cloud offering and is a common hypervisor technology in data centers. Morpheus provides an avenue to aggregate Hyper-V resources together to allow efficient and seamless deployment of applications as a virtual machine (VM) or Docker host in the world of Hyper-V.

Features

- Virtual Machine Provisioning
- Containers
- Backups / Snapshots
- Resources Groups
- Migrations
- Auto Scaling
- Load Balancing
- Remote Console
- Periodic Synchronization
- Veeam Integration
- Lifecycle Management and Resize
- Unique Kerberos Authentication

Morpheus can provide a single pane of glass and self-service portal for managing instances scattered across both Hyper-V and public cloud offerings like Azure.

Getting Started

To get started this a few prerequisites must first be met. The Hyper-V host must be installed with its firewall enabled and it can either be joined to a domain or standalone. The Hyper-V host must also have the external network of Hyper-V configured and it can share this network with the management operating system. This document covers Hyper-V 2008 and Hyper-V 2012.

A user account that is part of the local administrators group on the Hyper-V host is also required.

Understand WinRM

Morpheus uses WinRM to communicate to the Hyper-V host for deployment of the Morpheus agent. The Morpheus agent allows for the host dashboard to be populated with information in the form of graphs that cover CPU, Network, Storage, and memory consumption. Furthermore, this agent provides logging and monitoring capabilities.

If Windows Remote Management (WinRM) is not installed and configured, WinRM scripts do not run and the WinRM command-line tool cannot perform data operations or allow for the Morpheus agent to be installed. WinRM uses Http port 5985 or Https port 5986 for communications.

To better understand all of the default settings of WinRM please refer to the below Microsoft link:

[https://msdn.microsoft.com/en-us/library/aa384372\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/aa384372(v=vs.85).aspx)

Native Authentication

To configure WinRM with default settings (WINRM_NATIVE)

Type the following command at a command prompt:

```
$ winrm quickconfig
```

If you are not running under the local computer Administrator account, you must either select Run as Administrator from the Start menu or use the *Runas* command at a command prompt.

When the tool displays `Make these changes [y/n]?`, type `y`.

If configuration is successful, the following output is displayed:

```
$ WinRM has been updated for remote management.
$ WinRM service type changed to delayed auto start.
$ WinRM service started.
$ Created a WinRM listener on HTTP://* to accept WS-Man requests to any IP on this_
↪machine.
```

Keep the default settings for client and server components of WinRM, or customize them. By default Kerberos is enabled and if domain authentication is not being used we want to disable that. Issue the below commands to setup basic authentication:

```
$ winrm set winrm/config/service/Auth @{Basic="true"}
$ winrm set winrm/config/service @{AllowUnencrypted="true"}
$ winrm set winrm/config/service/Auth @{Kerberos="false"}
```

Domain Authentication

To configure WinRM with Domain Authentication (WINRM_INTERNAL)

Type the following command at a command prompt

```
$ winrm quickconfig
```

If you are not running under the local computer Administrator account, you must either select Run as Administrator from the Start menu or use the *runas* command at a command prompt.

When the tool displays `Make these changes [y/n]?`, type `y`.

If configuration is successful, the following output is displayed:

```
$ WinRM has been updated for remote management.
$ WinRM service type changed to delayed auto start.
$ WinRM service started.
$ Created a WinRM listener on HTTP://* to accept WS-Man requests to any IP on this_
↪machine.
```

Keep the default settings for client and server components of WinRM, or customize them. Issue the below commands to setup domain authentication:

```
$ winrm set winrm/config/service/Auth @{Basic="true"}
$ winrm set winrm/config/service @{AllowUnencrypted="false"}
$ winrm set winrm/config/service/Auth @{Kerberos="true"}
```

Kerberos authentication will also need to be configured on the Morpheus appliance to support Windows domain accounts to access the remote host with WINRM_INTERNAL connection type.

On the Morpheus appliance the krb5-user package must be installed.

For Ubuntu the command is as follows:

```
$ sudo apt-get install krb5-user
```

For Centos the command is as follows:

```
$ sudo yum install krb5-workstation pam_krb5 -y
```

Create a file in /etc called krb5.conf and replace the domain name with the name of the domain to be used. In this case we used Morpheus .com as the domain.

```
[libdefaults]
    default_realm = |morpheus| .COM
    dns_lookup_kdc = true
    verify_ap_req_nofail = false
    default_tgs_enctypes = rc4-hmac
    default_tkt_enctypes = rc4-hmac
[realms]
    |morpheus| .COM = {
        kdc = win-ad.|morpheus| .COM:88
        admin_server = win-ad.|morpheus| .COM:749
    }
[domain_realm]
    .|morpheus| .COM = |morpheus| .COM
    |morpheus| .COM = |morpheus| .COM
[login]
    krb4_convert = true
    krb4_get_tickets = false
```

After creation of the krb5.conf a keytab file is also required. See below on instructions on how to create a keytab file.
<http://www.itadmintools.com/2011/07/creating-kerberos-keytab-files.html>

Adding Hyper-V as a Private Cloud

The Hyper-V host is prepared for Morpheus to communicate with it via WinRM so the Hyper-V private cloud is ready to be configured. Create a group and then create a Morpheus cloud for Hyper-V. Populate the information as shown in Figure 1: specific for the environment being configured.

CREATE CLOUD

×

CLOUD

CONFIGURE

GROUP

REVIEW

NAME

San Mateo Hyper-V

LOCATION

San Mateo, CA

DNS DOMAIN

localdomain

VISIBILITY

Public

▼

SCALE PRIORITY

1

Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

Details

HYPER-V HOST

192.168.163.141

WORKING PATH

D:\Morpheus

VM PATH

D:\VMs

DISK PATH

D:\VirtualDisks

USERNAME

Administrator


PASSWORD

.....

PREV

NEXT


























Note: The working path, vm path, and disk path should be created on the Hyper-V host by the Hyper-V administrator. If these paths are not created they will need to be setup and the Hyper-V settings will need to be adjusted to reference them.



Service Plans

A default set of Service Plans are created in Morpheus for the VMware provisioning engine. These Service Plans can be considered akin to AWS Flavors or Openstack Flavors. They provide a means to set predefined tiers on memory, storage, cores, and cpu. Price tables can also be applied to these so estimated cost per virtual machine can be tracked as well as pricing for customers. By default, these options are fixed sizes but can be configured for dynamic sizing. A service plan can be configured to allow a custom user entry for memory, storage, or cpu. To configure this, simply edit an existing Service Plan tied to Hyper-V or create a new one. These all can be easily managed from the Admin | Service Plans & Pricing section.

SELECT: PROVISION TYPE ▾
+ CREATE SERVICE PLAN

TYPE NAME	CLOUDS	MEMORY	STORAGE	PRICE SETS		
 Memory: 3.5GB, 1 vCPU	Google Labs	3.5GB	0MB	0	/	🗑️
 Memory: 7.2GB, 2 vCPU	Google Labs	7.2GB	0MB	0	/	🗑️
 Memory: 15GB, 4 vCPU	Google Labs	15.0GB	0MB	0	/	🗑️
 Memory: 30GB, 8 vCPU	Google Labs	30.0GB	0MB	0	/	🗑️
 Memory: 60GB, 16 vCPU	Google Labs	60.0GB	0MB	0	/	🗑️
 Memory: 120GB, 32 vCPU	Google Labs	120.0GB	0MB	0	/	🗑️
 Memory: 13GB, 2 vCPU	Google Labs	13.0GB	0MB	0	/	🗑️
 Memory: 26GB, 4 vCPU	Google Labs	26.0GB	0MB	0	/	🗑️
 Memory: 52GB, 8 vCPU	Google Labs	52.0GB	0MB	0	/	🗑️
 Memory: 104GB, 16 vCPU	Google Labs	104.0GB	0MB	0	/	🗑️
 Memory: 208GB, 32 vCPU	Google Labs	208.0GB	0MB	0	/	🗑️
 Memory: 1.7GB, 2 vCPU	Google Labs	1.7GB	0MB	0	/	🗑️
 Memory: 3.5GB, 4 vCPU	Google Labs	3.5GB	0MB	0	/	🗑️
 Memory: 7.2GB, 8 vCPU	Google Labs	7.2GB	0MB	0	/	🗑️
 Memory: 14GB, 16 vCPU	Google Labs	14.0GB	0MB	0	/	🗑️
 Memory: 28GB, 32 vCPU	Google Labs	28.0GB	0MB	0	/	🗑️
 Memory: 512MB Storage: 10GB	San Mateo Hyper-V	512.0MB	10.0GB	0	/	🗑️
 Memory: 1GB Storage: 10GB	San Mateo Hyper-V	1.0GB	10.0GB	0	/	🗑️
 Memory: 2GB Storage: 20GB	San Mateo Hyper-V	2.0GB	20.0GB	0	/	🗑️
 Memory: 4GB Storage: 40GB	San Mateo Hyper-V	4.0GB	40.0GB	0	/	🗑️
 Memory: 8GB Storage: 80GB	San Mateo Hyper-V	8.0GB	80.0GB	0	/	🗑️
 Memory: 16GB Storage: 160GB	San Mateo Hyper-V	16.0GB	160.0GB	0	/	🗑️
 Memory: 24GB Storage: 240GB	San Mateo Hyper-V	24.0GB	240.0GB	0	/	🗑️
 Memory: 32GB Storage: 320GB	San Mateo Hyper-V	32.0GB	320.0GB	0	/	🗑️
 HV_4_240	San Mateo Hyper-V	3.9GB	234.3GB	0	/	🗑️

Docker

So far this document has covered how to add the Hyper-V cloud integration and has enabled users the ability to provision virtual machine-based instances via the Add Instance catalog in Provisioning. Another great feature provided by Morpheus out of the box is the ability to use Docker containers and even support multiple containers per Docker host. To do this a Docker Host must first be provisioned into Hyper-V (multiple are needed when dealing with horizontal scaling scenarios).

To provision a Docker Host simply navigate to the Clusters tab of the Cloud detail page or Infrastructure > Clusters section. From there click + *ADD CLUSTER* to add a Hyper-V Docker Host. A cluster is created when adding Docker hosts, even when only one host is needed.

Morpheus views a Docker host just like any other hypervisor with the caveat being that it is used for running containerized images instead of virtualized ones. Once a Docker Host is successfully provisioned a green checkmark will appear to the right of the host marking it as available for use. In the event of a failure click into the relevant host that failed and an error explaining the failure will be displayed in red at the top.

Some common error scenarios include network connectivity. For a Docker Host to function properly, it must be able to resolve the Morpheus appliance url which can be configured in Admin | Settings. If it is unable to resolve and negotiate with the appliance then the agent installation will fail and provisioning instructions will not be able to be issued to the host.

KVM

Adding VLANs to Morpheus KVM Hosts (CentOS)

Overview

Morpheus KVM is a powerful, cheaper alternative to virtualization when it comes to other hypervisor offerings. It is also very capable of setting up complex shared storage and multiple networks across many hosts. Currently this process is a manual process but will become automated in the coming months. This guide will go over how to configure VLANs on a Morpheus KVM Host.

Getting Started

To get started, the first step is to go ahead and add the KVM host to morpheus and allow morpheus to configure it just like any other kvm host. When provisioning a manual kvm host be sure to enter the proper network interface name for the management network (not the trunk port). For example `eno2` could be a management network while `eno1` could be the trunk port network that the VLAN's are going to be on as in this example.

Setting up a VLAN Interface

Before a VLAN can be used by KVM, an interface definition must first be configured for said vlan. In CentOS this is done by defining a network script in `/etc/sysconfig/network-scripts`.

Note: It is highly recommended that `NM_CONTROLLED` is set to `NO` or NetworkManager is disabled entirely as it tends to get in the way.

If our trunk network is called `eno1` we need to make a new script for each VLAN ID we would like to bridge onto. In our example we are going to look at **VLAN 211**. To do this we need to make a new script called `ifcfg-eno1.211` (note the VLAN Id is a suffix to the script name after a period as this is conventional and required).

```
TYPE=Ethernet
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=none
NAME=enol.211
DEVICE=enol.211
ONBOOT=yes
NM_CONTROLLED=no
VLAN=yes DEVICETYPE=ovs
OVS_BRIDGE=br211
```

There are a few important things to note about this script. Firstly there is a flag called `VLAN=yes` that enables the kernel tagging of the VLAN. Secondly we have defined an `OVS_BRIDGE` name. Morpheus utilizes openvswitch for its networking which is a very powerful tool used even by Openstack's Neutron. It supports not just VLANs but VxLAN interfacing.

The **OVS_BRIDGE** name means we also need to define a bridge port script called `br211` by making a script called `ifcfg-br211`:

```
DEVICE=br211
ONBOOT=yes
DEVICETYPE=ovs
TYPE=OVSBridge
NM_CONTROLLED=no
BOOTPROTO=none
HOTPLUG=no
```

These configurations will enable persistence on these interfaces so that a reboot of the host will retain connectivity to the bridges. Next up, the interfaces need to be brought online. This can be done by restarting all network services but if a typo is made networking could be stuck disabled and access over SSH could be broken. To do this by interface simply run:

```
ifup enol.211
ifup br211
ovs-vsctl
add-br br211
```

Defining a LibVirt Network

Now that the bridge interface is defined properly for OVS, it must be defined in LibVirt so that Morpheus will detect the network and KVM can use it properly. By convention, these resource configurations are stored in `/var/morpheus/kvm/config`.

An XML definition must be created to properly define the network. In this case the network is named `public 185.3.48.0.xml`:

```
<network>
<name>public 185.3.48.0</name>
<forward mode="bridge"/>
<bridge name="br211"/>
<virtualport type="openvswitch"/>
</network>
```

This configuration defines the network name that will be synced into morpheus for selection as well as the type of interface being used (in this case a bridge to the `br211` interface over openvswitch).

Now that this xml specification is defined it must be registered with libvirt via the virsh commands:

```
virsh net-define "public 185.3.48.0.xml"
virsh net-autostart "public 185.3.48.0"
virsh net-start "public 185.3.48.0"
```

Once this is completed, simply refresh the cloud in morpheus and wait for the network to sync into the networks list. Once the network is synced make sure the appropriate settings are applied to it within Morpheus. This includes setting the CIDR, Gateway, Nameservers and if using IP Address Management, the IPAM Pool.

Canonical MAAS

MAAS from Canonical is an open-source tool, server orchestration tool. It's designed to allow administrators to build a data center from on-premises, bare-metal servers where large networks of individual units can be discovered, deployed, and reconfigured.

Integrating MAAS and Morpheus

Integrating MAAS with Morpheus is a simple process requiring the MAAS API URL and API Key. We'll start by gathering what we need from the MAAS UI, then move back into Morpheus to store the required details.

We can gather the API URL by clicking on the username in the upper-right corner of the MAAS UI window. From this preferences page, click on "API keys" as shown in the screenshot:

MAAS

Machines Devices Controllers KVM Images DNS AZs Subnets Settings

Morpheus Log out

My preferences

Details

API keys

SSH keys

SSL keys

Username *

Morpheus

Required. 150 characters or fewer. Letters, digits and @/./+/_ only.

Full name (optional)

Email address *

morpheus@morpheusdata.com

[Change password...](#)

Save

MAAS name: morph MAAS

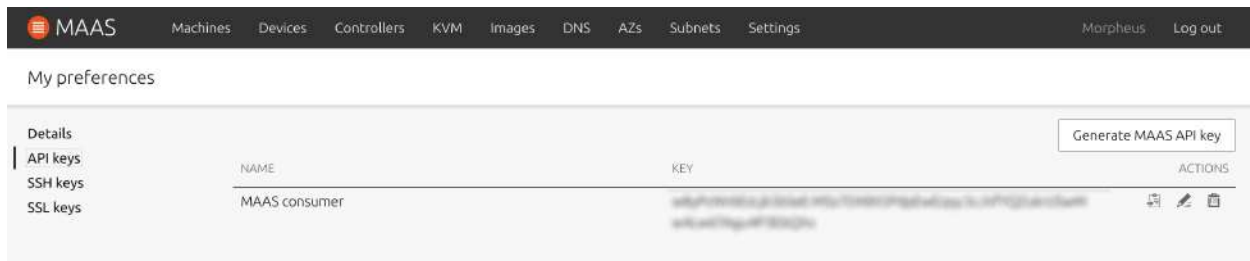
MAAS version: 2.7.3 (8290-g.ebe2b9884)

[View release notes](#) · [View documentation](#) · [Legal information](#) · [Give feedback](#)

© 2020 Canonical Ltd. Ubuntu and Canonical are registered trademarks of Canonical Ltd.

CANONICAL

From the API keys page, select the displayed key and copy it. Alternatively, you can click the copy button in the UI to add the full key to your clipboard. Store this key in an accessible location for a later step.



In addition to the API key, we need the MAAS API URL. This URL is given in the format `http://<maas-hostname-or-ip>:5240/MAAS/api/2.0`. Plug the hostname or host IP address into the example shown in the previous sentence and store the complete API URL for use in the next step.

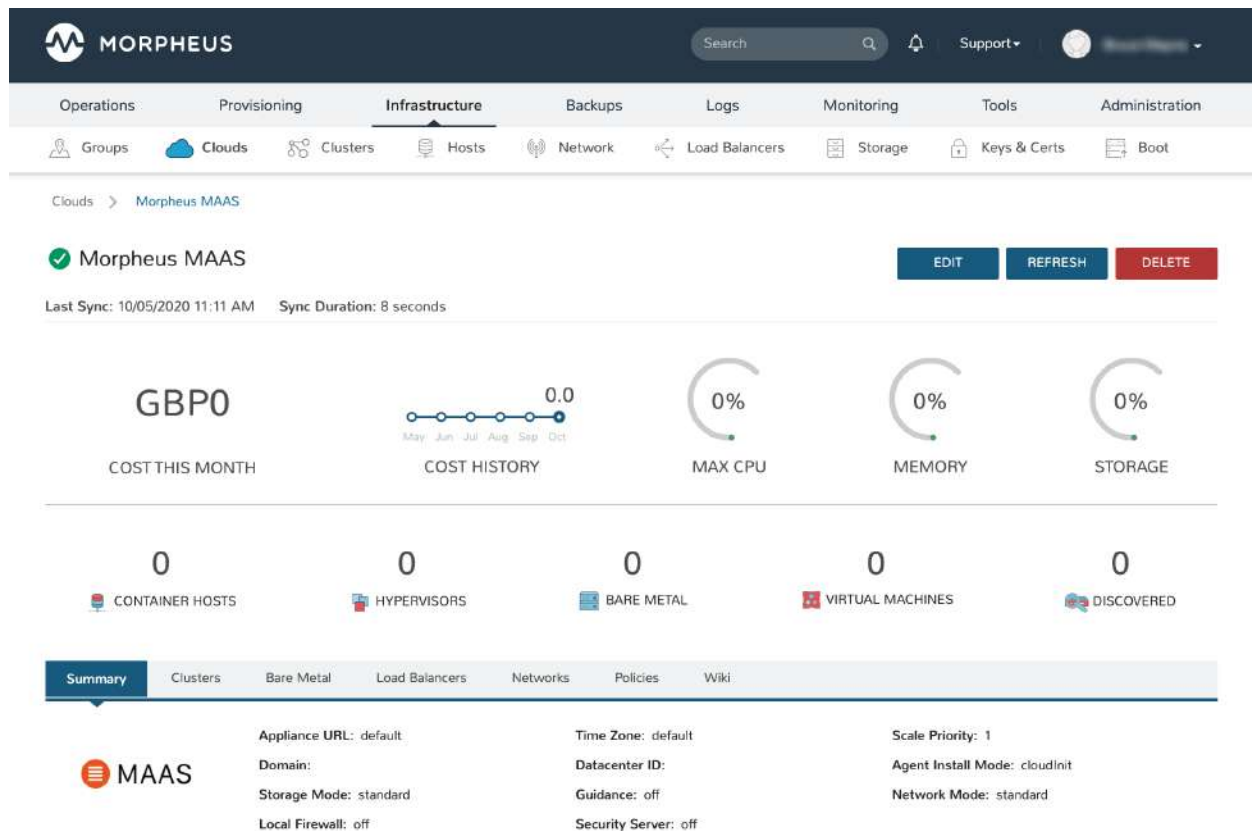
In Morpheus, navigate to the list of integrated Clouds and start a new MAAS Cloud integration:

1. Infrastructure > Clouds
2. Click + *ADD*
3. Select “MAAS”
4. Click *NEXT*

On the “CREATE CLOUD” modal, you must at least give a friendly name for the Cloud in Morpheus, MAAS API URL and API KEY. An example is shown below:

Tip: You’ll know the credentials are entered correctly when your list of MAAS resource pools is synced in as you can see in the example screenshot.

Click *NEXT* and add this new Cloud to an existing Group or create a new Group for it. Once you advanced past the end of the wizard, the Cloud is added and Morpheus begins to inventory (if you’ve opted to inventory when adding the Cloud).



Mac Stadium

Overview

MacStadium is a provider of enterprise-class hosting solutions for Apple Mac infrastructure. It can be used to deploy a hosted private cloud for large-scale CI/CD or even a single Mac mini to test an iOS app. It allows virtualized Mac build machines

Features

- Virtual Machine Provisioning
- Backups / Snapshots
- Resource Groups
- Datastores and DRS Clusters
- Distributed Switches
- Datacenter / Cluster scoping
- Brownfield VM management and migration
- VMware to VMware migrations
- VMDK/OVF image conversion support

- Hypervisor Remote Console
- Periodic Synchronization
- Veeam Backup Integration
- Lifecycle Management and Resize

On top of all these features, Morpheus also adds additional features to VMware that do not exist out of the box to make it easier to manage in multitenant environments as well as hybrid cloud environments:

- Cloud-Init Support
- VHD to VMDK Image Conversion
- QCOW2 to VMDK Image Conversion
- Multitenancy resource allocation
- Virtual Image management (Blueprints)
- Auto-scaling and recovery

Getting Started

To get started with VMware, simply start by adding a Cloud in the Infrastructure -> Clouds section.

The screenshot shows the Morpheus 'CREATE CLOUD' dialog box. The 'CONFIGURE' tab is selected, displaying the following fields and options:

- NAME:** Text input field.
- LOCATION:** Text input field.
- DOMAIN:** Text input field with 'localdomain' entered.
- VISIBILITY:** Dropdown menu set to 'Public'.
- SCALE PRIORITY:** Text input field with '1' entered. A note below states: 'Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.'
- Details section:**
 - API URL:** Text input field with 'https://vcenter.domain.com/sdk' entered.
 - USERNAME:** Text input field.
 - PASSWORD:** Text input field.
 - VDC:** Dropdown menu with 'No datacenters found: check your config'.
 - CLUSTER:** Dropdown menu with 'No clusters found: check your config'.
 - RESOURCE POOL:** Dropdown menu with 'No resource pools found: check your config'.
 - RPC MODE:** Dropdown menu set to 'SSH / WinRM'.
 - Checkboxes:**
 - ☐ HIDE HOST SELECTION FROM USERS
 - ☐ INVENTORY EXISTING INSTANCES
 - ☐ ENABLE VNC

At the bottom of the dialog are 'PREV' and 'NEXT' buttons. The background shows the Morpheus dashboard with a sidebar containing 'Groups', 'Clouds', and 'Hosts' sections, and a main area with a 'CLOUDS' list and a '+ CREATE CLOUD +' button.

To start adding a VMware cloud there will be some things you will need:

vCenter API Url Typically this is the url to the vCenter web client with a /sdk in the path

Username/Password A set of credentials with high level access to VMware (ensure the account has Datacenter level access)

Once these fields are entered, some selections will start pre-populating. A cloud integration is scoped to a specific data center, and can optionally be scoped down to a single cluster or even a single resource pool. If the drop downs do not populate, please verify the api url is resolvable, morpheus has access to vCenter on 443, and the provided credentials are correct and the user has sufficient permissions.

Another cool feature provided with the cloud integration is optional *Resource Pool* scoping. One can choose to allow the cloud to provision into All Resource Pools or a singular Resource Pool. When choosing *All*, these Resource Pools can be managed from a sub-account and visibility perspective via the Cloud Detail page (multi-tenancy).

The VMware cloud integration provides a few additional options including allowing users to make host selections or keeping that aspect hidden such that the best host is automatically chosen for the requested provision.

The *RPC Mode* feature can be configured to allow Morpheus to install its agent on the Guest operating system via either SSH/WinRM or VMware Tools Guest Process feature. The VMware tools Guest Execution API can be tricky so it is recommended to use SSH/WinRM if possible. However, if it is not possible for the Appliance to have outbound access to all networks in which VMs are being provisioned to the SSH/WinRM ports (22, 5985 respectively) then Guest Execution is the only option.

The *Use VNC* console option on the VMware cloud requires special configuration on each ESXI host but allowed hypervisor level remote console support. (See the Advanced Section for details)

When following this add cloud wizard an option will be presented to create a group or add to an existing group. These groups can be given provisioning permission via role based access control. It is normally recommended that groups are organized such that one cloud exists in one group unless the networks are setup such that internal routing is possible between the clouds. This is very useful for bursting, or hybrid cloud configurations.

Windows Provisioning Tips

By default when provisioning windows templates, Morpheus performs guest customizations which initiates a sysprep. This resets the Administrator user and password. Morpheus will set the Administrator password from Administration > Provisioning > Windows Settings > Password.

Users can also set the username on an image as Administrator and enter a different password if unique passwords are required per image.

Guest customizations are required when assigning static IP's manually or using IP pools. They can be disabled per virtual image advanced settings under Provisioning > Virtual Images > Edit Image > Advanced > Uncheck "Force Guest Customization" if using DHCP. However the SID will not be changed from the source template. In addition, new VM's will not be able to join a domain that had already been joined by the source template or any other VM's with that SID.

Existing Instances

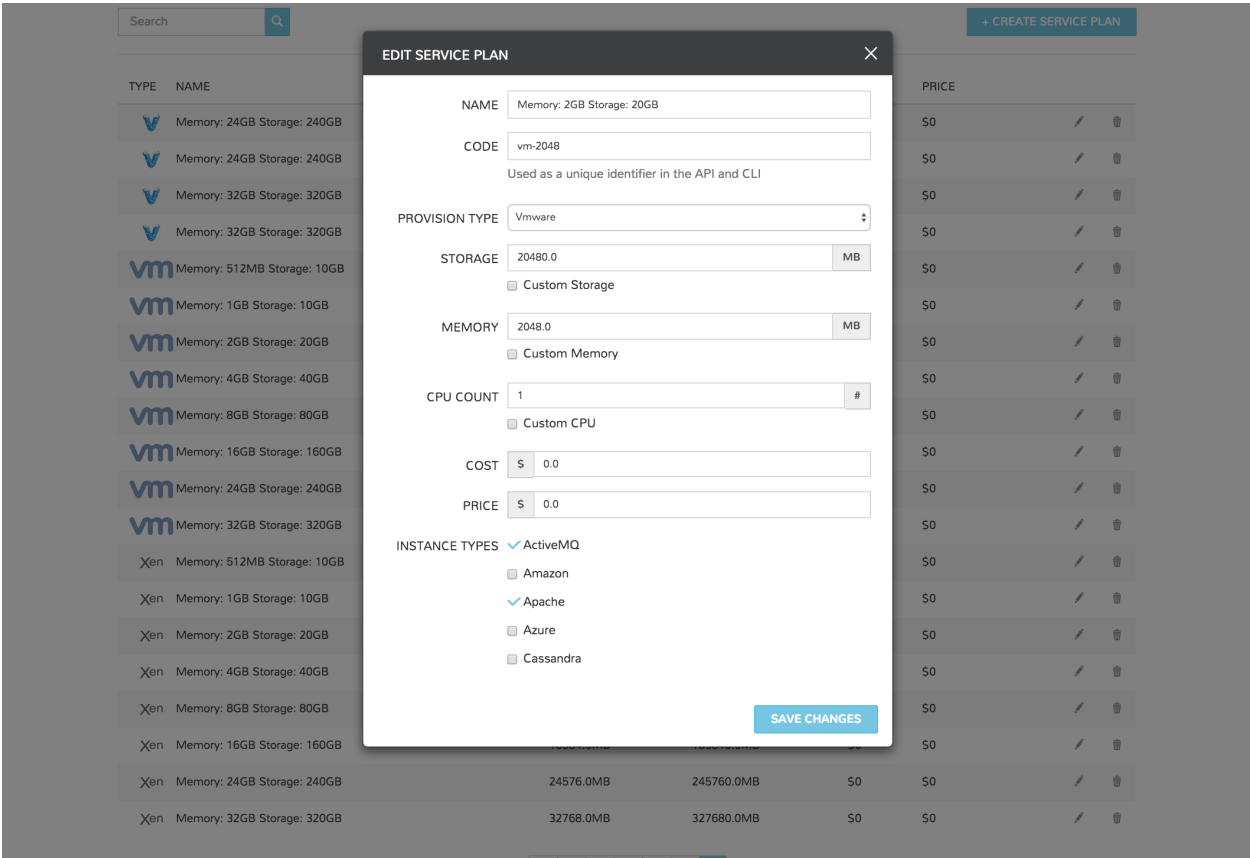
Morpheus provides several features regarding pulling in existing virtual machines and servers in an environment. Most cloud options contain a checkbox titled '*Inventory Existing Instances*'. When this option is selected, all VMs found within the specified scope of the cloud integration will be scanned periodically and Virtual Machines will be synced into Morpheus. By default these virtual machines are considered 'unmanaged' and do not appear in the Provisioning -> Instances area but rather Infrastructure -> Hosts -> Virtual Machines. However, a few features are provided with regards to unmanaged instances. They can be assigned to various accounts if using a multitenant master account, however it may be best suited to instead assign the 'Resource Pool' to an account and optionally move all servers with regards to that pool (more on this later). A server can also be made into a managed server. During this process remote access is requested and an agent install is performed on the guest operating system. This allows for guest operations regarding log acquisition and stats. If the agent install fails,

a server will still be marked as managed and an Instance will be created in *Provisioning*, however certain features will not function. This includes stats collection and logs.

Note: All Cloud data is resynchronized on a 5 minute interval. This includes Datastores, Resource Pools, Networks, Blueprints, and Virtual Machines.

Service Plans

A default set of Service Plans are created in Morpheus for the VMware provisioning engine. These Service Plans can be considered akin to AWS Flavors or Openstack Flavors. They provide a means to set predefined tiers on memory, storage, cores, and cpu. Price tables can also be applied to these so estimated cost per virtual machine can be tracked as well as pricing for customers. By default, these options are fixed sizes but can be configured for dynamic sizing. A service plan can be configured to allow a custom user entry for memory, storage, or cpu. To configure this, simply edit an existing Service Plan tied to VMware or create a new one. These all can be easily managed from the Admin -> Plans & Pricing section.



Virtual Images / Blueprints

Morpheus will automatically take an inventory of all blueprints configured in vCenter and present them as options during provisioning. However, in order for Morpheus to properly provision these virtual machines and provide accurate stats and health of these virtual machines, an agent must be installed during virtual machine startup. This means remote access needs to be granted at the guest operating system level to Morpheus. To properly configure these virtual images, find the relevant images in `Provisioning -> Virtual Images` and edit the entry. On this form, a few options are presented. The first is a check box asking whether or not cloud-init is enabled. If cloud-init is enabled, simply provide the default OS username configured (for Ubuntu the username is *ubuntu* and for CentOS the username is *centos*). For those looking to add cloud-init to existing blueprints Morpheus requires no special configuration and can use the default *cloud.cfg* settings.

A global cloud-init username/password can also be configured per account as well as a keypair via the `Admin->Provisioning` settings section. The great benefit of utilizing cloud-init is default blueprints do not need common credential sets thereby increasing provisioning security.

Windows systems do not typically support cloud-init. So simply turn this checkbox off and provide the *Administrator* credentials. It should be noted that these credentials are encrypted in the database. If using WinRM for the RPC Mode instead of VMware tools, a Local or Domain Administrator account credential set can be provided instead.

Snapshots

Morpheus allows the ability to create a snapshot of a VM in VMware vCenter. From the instance detail page, simply select `Actions -> Create Snapshot` to begin creation of a new Snapshot. Existing snapshots can be viewed in the `BACKUPS` tab on the instance detail page. Snapshots taken in vCenter will sync into Morpheus every five minutes. To revert to a previous snapshot, click on the revert icon located on the right side of the Snapshot. Snapshots can be deleted by clicking on the trash can icon.

Note: Access to Snapshots can be limited or removed entirely for specific user roles as needed. To edit a role's Snapshots permissions, go to `Administration > Roles > (Your selected role) > Snapshots`. Users can be given Full, Read-only, or No access.

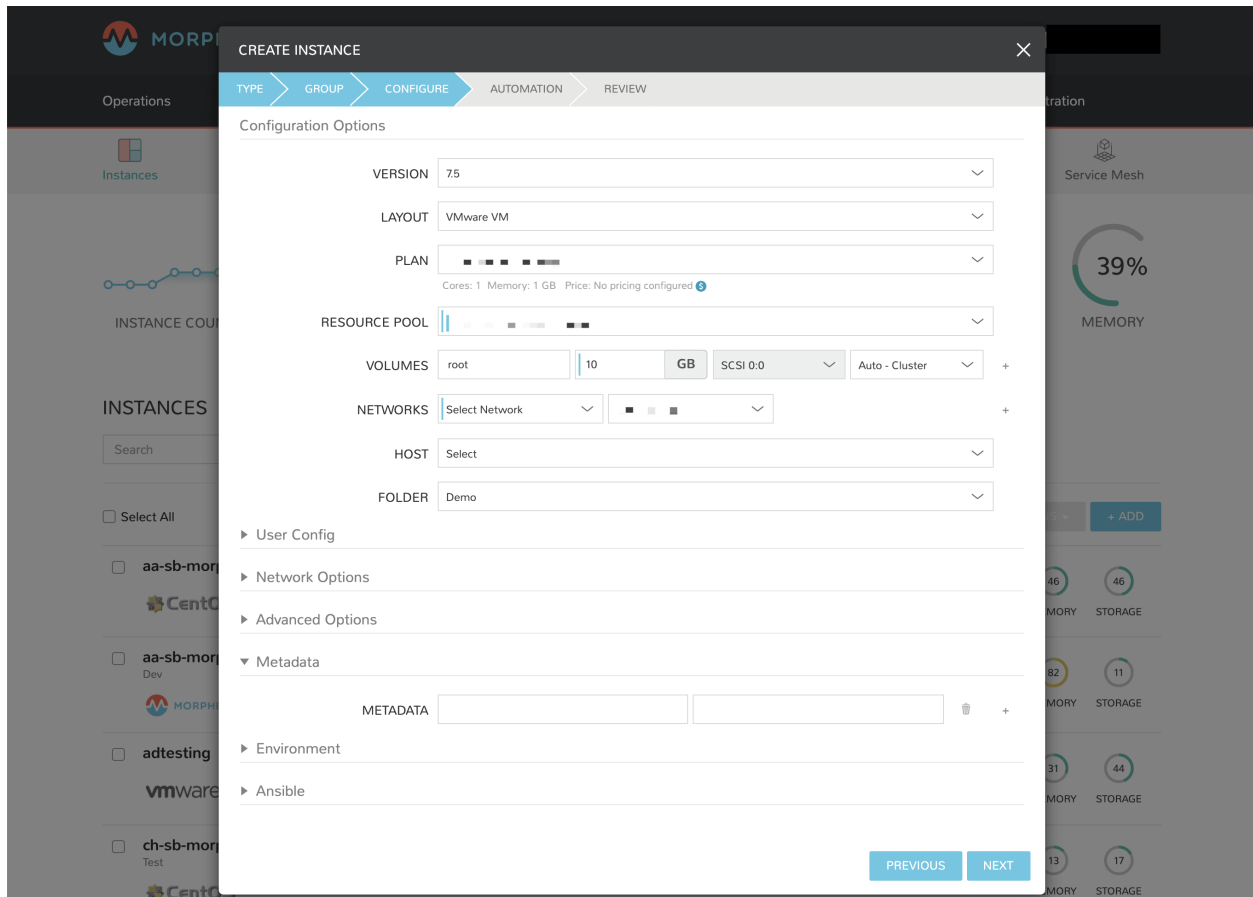
Tagging and Metadata

As of Morpheus version 4.1.0, tagging support is included for vCenter in addition to the other clouds that have already supported it in past versions. Tags will sync to vCenter from Morpheus and existing tags are also inventoried from vCenter into Morpheus.

Note: This feature requires a minimum API version of vCenter 6.5. The API version can be edited by navigating to 'Infrastructure > Clouds' and clicking the edit (pencil) button in the row for the relevant cloud. The field is labeled 'VERSION'.

Tags can be created on-demand when provisioning from the 'CONFIGURE' tab of the 'CREATE INSTANCE' wizard (`Provisioning > Instances`). Within the 'Metadata' drawer, you will see sets of fields to enter key/value pairs. On creation of the instance, this metadata will be synced into vCenter.

'Option Types' from your library can also be exported as metadata for use with vCenter. When adding or editing a new Option Type (`Provisioning > Library > OPTION TYPES`), simply mark the box labeled 'EXPORT AS METADATA'. The 'FIELD NAME' becomes the tag category in VMWare.



Docker

So far this document has covered how to add the VMware cloud integration and has enabled users the ability to provision virtual machine based instances via the *Add Instance* catalog in *Provisioning*. Another great feature provided by Morpheus out of the box is the ability to use Docker containers and even support multiple containers per Docker host. To do this a Docker Host must first be provisioned into VMware (multiple are needed when dealing with horizontal scaling scenarios).

To provision a Docker Host simply navigate to the Clusters tab of the Cloud detail page or Infrastructure > Clusters section. From there, click + *ADD CLUSTER* to add a VMware Docker Host. This host will show up in the Hosts tab next to other ESXi servers that were inventoried by the VMware cloud integration. Morpheus views a Docker host just like any other Hypervisor with the caveat being that it is used for running containerized images instead of virtualized ones. Once a Docker Host is successfully provisioned a green checkmark will appear to the right of the host marking it as available for use. In the event of a failure click into the relevant host that failed and an error explaining the failure will be displayed in red at the top.

Some common error scenarios include network connectivity. For a Docker Host to function properly, it must be able to resolve the Morpheus appliance url which can be configured in Administration > Settings. If it is unable to resolve and negotiate with the appliance then the agent installation will fail and provisioning instructions will not be able to be issued to the host.

Multitenancy

A very common scenario for Managed Service Providers is the need to provide access to VMware resources on a customer by customer basis. With VMware several administrative features have been added to ensure customer resources are properly scoped and isolated. For VMware it is possible to assign specific *Networks*, *Datastores*, and *Resource Pools* to customer accounts or even set the public visibility of certain resources, therefore allowing all sub accounts access to the resource.

LABS VMWARE EDIT DELETE

Location: San Mateo
Cloud Group: Labs VMware
Servers: 4

HOSTS VIRTUAL MACHINES BARE METAL SECURITY GROUPS LOAD BALANCERS NETWORKS DATA STORES **RESOURCE POOLS**

Search

NAME	VISIBILITY	ACCOUNT	ACTIONS
Resources	Private	morpheus-qa	ACTIONS ▾
Brian	Private	morpheus-qa	ACTIONS ▾
Macbook	Private	morpheus-qa	ACTIONS ▾
David	Private	morpheus-qa	ACTIONS ▾
Macbook	Private	morpheus-qa	ACTIONS ▾

© 2016 MORPHEUS DATA, LLC. ALL RIGHTS RESERVED
TERMS AND CONDITIONS | PRIVACY POLICY

Advanced

There are several advanced features provided within Morpheus that can leverage some cool aspects of VMware. One of these features is Remote Console support directly to the hypervisor. To enable this feature a few prerequisites must be met. First, the Morpheus appliance must have network access to the ESXi hosts within VCenter. Secondly, firewall settings need to be adjusted on each ESXi host. This can be done in VSphere under firewall configuration on the host. Simply check the *gdbserver* option, which will open up the necessary ports (starting at 5900 range).

Important: Hypervisor Console for vCenter 6.5 requires Morpheus v3.2.0+

Now that the ESXi hosts are ready to utilize remote console, simply edit the cloud in Morpheus via Infrastructure → Clouds. Check the option that says *Use VNC*. It is important to note that currently this functionality only works for newly provisioned vm's provisioned directly via Morpheus. This should change soon however.

It is also possible to import vm snapshots for backup or conversion purposes from VCenter and also an ESXi host. However, this does require that the ESXi host license has an enterprise level license as it will not allow the appliance

to download a virtual image if it is not a paid VMware license.

Nutanix

Overview

Nutanix simplifies datacenter infrastructure by integrating server and storage resources allowing applications to run at scale. Morpheus provides an avenue to enhance the Nutanix resources to allow efficient and seamless deployment of applications as a virtual machine (VM) or as a container on a Docker host.

Features

- Virtual Machine Provisioning
- Containers
- Backups / Snapshots
- Resources Groups
- Migrations
- Auto Scaling
- Load Balancing
- Remote Console
- Periodic Synchronization
- Lifecycle Management and Resize

Morpheus can provide a single pane of glass and self-service portal for managing multiple Nutanix Clusters and allowing the seamless deployment of applications.

Note: Prism Central is not currently supported as a Cloud endpoint target

Getting Started

To get started this a few prerequisites must first be met. The Nutanix cluster should be provisioned and available on the network. Morpheus will look login to the Nutanix cluster with the Nutanix admin credentials and is typically located at the <https://fqdn:9440> url.

Adding a Nutanix Cloud

The Nutanix cluster should be available and responding to the <https://fqdn:9440> url for authentication by Morpheus .

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected powered on state of managed VM's and power on any managed VM's in the cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When "AUTOMATICALLY POWER ON VMS" is enabled, the power state of managed VM's should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

API URL URL of the Nutanix Prism API, example: <https://10.30.21.220:9440>. Prism Central is not currently supported as a Cloud endpoint target

USERNAME Nutanix admin username

PASSWORD Nutanix admin password

Inventory Existing Instances If enabled, existing Virtual Machines will be inventoried and appear as unmanaged Virtual Machines in Morpheus .

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Service Plans

A default set of Service Plans are created in Morpheus for the VMware provisioning engine. These Service Plans can be considered akin to AWS Flavors or Openstack Flavors. They provide a means to set predefined tiers on memory, storage, cores, and cpu. Price tables can also be applied to these so estimated cost per virtual machine can be tracked as well as pricing for customers. By default, these options are fixed sizes but can be configured for dynamic sizing. A service plan can be configured to allow a custom user entry for memory, storage, or cpu. To configure this, simply edit an existing Service Plan tied to Nutanix or create a new one. These all can be easily managed from the Admin | Service Plans & Pricing section.

Docker

So far this document has covered how to add the Nutanix cloud integration and has enabled users the ability to provision virtual machine based instances via the Add Instance catalog in Provisioning. Another great feature provided by Morpheus out of the box is the ability to use Docker containers and even support multiple containers per Docker host. To do this a Docker Host must first be provisioned into Nutanix (multiple are needed when dealing with horizontal scaling scenarios).

To provision a Docker Host, simply navigate to the Cloud detail page or Infrastructure > Clusters section. From there click + *ADD CLUSTER* to add a Nutanix Docker Host. Morpheus views a Docker host just like any other Hypervisor with the caveat being that it is used for running containerized images instead of virtualized ones. Once a Docker Host is successfully provisioned a green checkmark will appear to the right of the host marking it as available for use. In the event of a failure click into the relevant host that failed and an error explaining the failure will be displayed in red at the top.

Some common error scenarios include network connectivity. For a Docker Host to function properly, it must be able to resolve the Morpheus appliance url which can be configured in Admin Settings. If it is unable to resolve and negotiate with the appliance then the agent installation will fail and provisioning instructions will not be able to be issued to the host.

Snapshots

Morpheus allows the ability to create a snapshot of a Nutanix instance. From the instance detail page, simply select Actions -> Create Snapshot to begin creation of a new Snapshot. Existing snapshots can be viewed in the BACKUPS tab on the instance detail page. Snapshots taken outside Morpheus will be synced every five minutes. To revert to a previous snapshot, click on the revert icon located on the right side of the Snapshot. Snapshots can be deleted by clicking on the trash can icon.

Note: Access to Snapshots can be limited or removed entirely for specific user roles as needed. To edit a role's Snapshots permissions, go to Administration > Roles > (Your selected role) > Snapshots. Users can be given Full, Read-only, or No access.

Openstack

Overview

Openstack is becoming a widely used on-premise infrastructure orchestration platform. It has a wide array of contributors and enterprise sponsorships. There are several variations on Openstack as well. Morpheus supports integration with all the various platform offerings and ranges in support all the way back to Openstack Juno. The complete list of compatible versions is listed in our [Integration Compatibility table](#). It leverages the APIs and provides full functionality as a self service portal in front of Openstack.

Features

- Virtual Machine Provisioning
- Backups and Snapshots
- Security Group Management
- Disk Mode support Local/Image (via Ceph)
- Floating IP Assignment support
- Brownfield VM management and Migration
- Lifecycle Management and Resize
- Docker Host management and configuration
- Manila File Services (SFS)
- Object Storage (OBS)
- Network Lifecycle
- LBaaS/Octavia Load Balancing Services

On top of all these features, Morpheus also adds additional features to Openstack that do not exist out of the box to make it easier to manage in multitenant environments as well as hybrid cloud environments:

- Image to QCOW2 Image Conversion
- QCOW2 to RAW Image Conversion
- Multitenancy resource allocation
- Virtual Image management (Blueprints)
- Auto-scaling and recovery
- Instance Cloning
- Morpheus Kubernetes Cluster Deployment

Tip: To allow Morpheus to list Hypervisor Hosts, ensure the Openstack user used for the Cloud Integration has sufficient privileges for `os_compute_api:os-hypervisors` in `/etc/nova/policy.json` in Openstack.

Getting Started

OpenStack Clouds are very easy to integrate with Morpheus. First, go to the `Infrastructure > Clouds` section and click + *ADD*. Select OpenStack to begin the integration process, most branded flavors of OpenStack will work with this Cloud selection as well.

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

IDENTITY API URL v2.0 or v3 Identity endpoint.

DOMAIN ID For *Default* domains, Default can be used. For other domain the Domain ID must be entered, not the Domain Name.

PROJECT Target project

USERNAME Service Username

PASSWORD Service user password

OS VERSION Select Openstack Version.

IMAGE FORMAT Select QCOW2, RAW or VMDK Image Type

LB TYPE

Select LB Type for Openstack LB syncing and creation

Inventory Existing Instances Select for Morpheus to discover and sync existing VM's

Enable Hypervisor Console Hypervisor console support for openstack currently only supports novnc. Be sure the novnc proxy is configured properly in your openstack environment. When disabled Morpheus will use ssh and rdp for console connection (vm/host credentials required)

Note: The user which is used connect to a project only needs to be a member ('_member_') of the project rather than an admin. Admin will work but it exposes some additional items to the project that an Openstack Admin typically does not want portal users to see.

Most of the information in the dialog can be acquired from the Openstack dashboard, under *Project > Access & Security > API Access*. The API URL that is needed is the one tied to *Identity*. The Domain and Project inputs typically correlate to the multitenant domain setup within Openstack (sometimes just left at default) as well as

the project name given to instances. Morpheus allows multiple integrations to the same Openstack cluster to be scoped to various domains and projects as needed.

The remaining options help Morpheus determine which API capabilities exist in the selected Openstack environment. Hence the need for the Openstack version and image format. If a newer Openstack cluster is being used then exists in the dropdown, simply select the most recent version in the dropdown and this should function sufficiently until the new version is added.

Tip: Some Openstack environments do not support QCOW2 and force RAW image formats (like Metapod). This is due to some network overhead in Ceph created by using QCOW2. Morpheus keeps two copies of Openstack image templates for this exact purpose.

Saving this cloud integration should perform a verification step and close upon successful completion.

Existing Instances

Morpheus provides several features regarding pulling in existing virtual machines and servers in an environment. Most cloud options contain a checkbox titled *'Inventory Existing Instances'*. When this option is selected, all VMs found within the specified scope of the cloud integration will be scanned periodically and Virtual Machines will be synced into Morpheus.

By default these virtual machines are considered 'unmanaged' and do not appear in the Provisioning -> Instances area but rather Infrastructure -> Hosts -> Virtual Machines. However, a few features are provided with regards to unmanaged instances. They can be assigned to various accounts if using a multi-tenant master account, however it may be best suited to instead assign the 'Resource Pool' to an account and optionally move all servers with regards to that pool (more on this later).

A server can also be made into a managed server. During this process remote access is requested and an agent install is performed on the guest operating system. This allows for guest operations regarding log acquisition and stats. If the agent install fails, a server will still be marked as managed and an Instance will be created in *Provisioning*, however certain features will not function. This includes stats collection and logs.

Service Ports

Morpheus consumes the following OpenStack service ports by default as part of its cloud integration. If your OpenStack implementation has been configured to use alternate service ports, these can be overwritten in the Cloud configuration under the Service Endpoints section when adding or editing the Cloud integration.

▼ Service Endpoints

COMPUTE SERVICE	<input type="text" value="http://10.30.21.150:8774"/>
IMAGE SERVICE	<input type="text" value="http://10.30.21.150:9292"/>
STORAGE SERVICE	<input type="text" value="http://10.30.21.150:8776"/>
NETWORK SERVICE	<input type="text" value="http://10.30.21.150:9696"/>
LOAD BALANCER SERVICE	<input type="text"/>
OBJECT STORAGE SERVICE	<input type="text"/>
SHARED FILE SYSTEM SERVICE	<input type="text"/>

Default Service Ports

- Identity: 5000
- Compute: 8774
- Compute_Legacy: 8774 v2
- Image: 9292
- Key Manager: 9311
- Network: 9696
- Volume API v2: 8776 v2
- Volume API v3: 8776 v3
- Manila: 8786

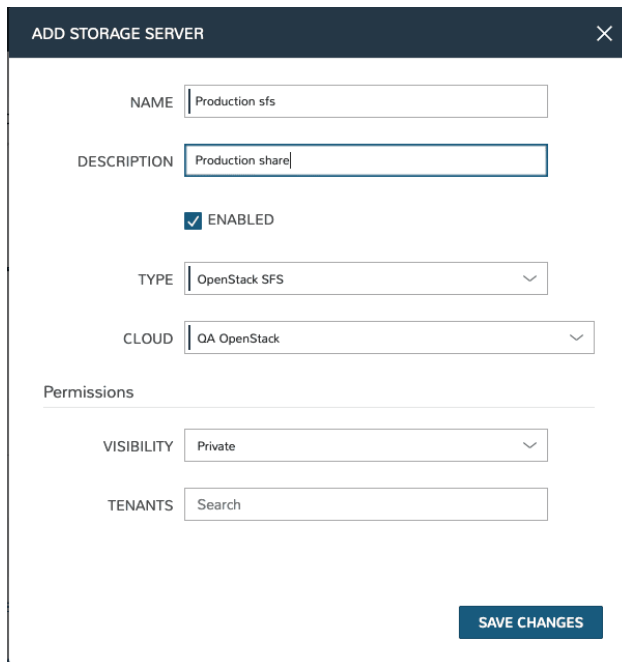
OpenStack Scalable File Service (SFS)

The Morpheus integration with Openstack Cloud includes the capability to work with Openstack Scalable File Service (SFS). SFS is shared file storage hosted on Openstack Cloud. By integrating Morpheus with Openstack you can discover, create, manage, and delete SFS servers, as well as view and work with the file shares and files contained therein.

SFS Server Discovery and Management

On integrating Openstack Cloud with Morpheus, SFS servers and file shares are discovered automatically after a short time. The server(s) can be viewed in Infrastructure > Storage > Servers. By viewing the server detail page and clicking *EDIT*, the storage server can be scoped as needed. Administrators can choose to scope to other Openstack Cloud integrations (if more than one relevant integration currently exists), select from synced availability zones, and scope the storage server to specific Tenants if desired.

Additionally, Openstack SFS servers can be created from the storage server list page (Infrastructure > Storage > Servers) directly in Morpheus. Click *+ADD* to begin and set the storage server type value to “Openstack SFS”. Just like with existing synced SFS servers, those created from Morpheus can be scoped as needed.



ADD STORAGE SERVER

NAME

DESCRIPTION

☒ ENABLED

TYPE

CLOUD

Permissions

VISIBILITY

TENANTS

SAVE CHANGES

SFS File Share Discovery and Management

Discovered file shares will appear among other file shares synced with Morpheus in Infrastructure > Storage > File Shares. Depending on the number of cloud integrations in your Morpheus appliance and the number of cloud integrations available to your user account, this list may be quite large. Using the search bar on this page, we can narrow down the list to only file shares whose names match the search terms.

We can drill into individual file shares by clicking on the hyperlinked name in the list of all integrated file shares. From the file share detail page, a list of files will appear on the files tab. Begin the process of adding a new file by clicking +ADD. The Access tab on the file shares detail page allows users to view and manage ACL rules.

Note: “Failed to load files from storage provider” is present when the Morpheus appliance doesn’t have access to the file share.

New Openstack SFS file shares can also be created directly in Morpheus. From the file shares list page, get started by clicking +ADD. Select the type “Openstack SFS Share”. Set the storage service field to a pre-existing Openstack SFS server. Setting a friendly name for the file share in Morpheus and selecting from synced availability zones is required.

Advanced

There are a few advanced features when it comes to provisioning on top of Openstack. Most of these present themselves in the provisioning wizard. This includes OS Volume Type (Local or Volume) which dictates whether the main OS disk is copied and run off the hypervisor or remotely mounted as a volume via Glacier. Some Openstack setups only configure hypervisors with minimal local disks so volume type is needed.

Another option during provisioning is “Assign Floating IP”. This option does exactly what it says and is similar to the feature on the Openstack instance dashboard itself. It should be noted that this will attempt to acquire a floating IP from the project and, if out of capacity, will attempt to increase capacity to the project if the cloud credentials provided have sufficient administrative privileges to do so.

Docker

So far this document has covered how to add the Openstack cloud integration and has described how to provision virtual machine-based Instances via the *Add Instance* catalog in *Provisioning*. Another great feature provided by Morpheus out of the box is the ability to work with Docker containers and even support multiple containers per Docker host. To do this, a Docker host must first be provisioned into Openstack (multiple hosts are needed when dealing with horizontal scaling scenarios).

To provision a Docker Host, navigate to Infrastructure > Clusters and click + *ADD CLUSTER*. Complete the provisioning wizard including selecting the appropriate Group and Cloud. Alternatively, you can navigate to the Clusters tab for a specific Cloud (Infrastructure > Clouds > Specific Cloud detail page > Clusters tab) and begin the process of provisioning a Docker host to that Cloud from there. Once completed, this host will show up in the Hosts sections (Infrastructure > Hosts OR Infrastructure > Clouds > Specific Cloud detail page > Hosts tab). Morpheus views a Docker host just like any other Hypervisor with the caveat being that it is used for running containerized images instead of virtualized ones.

Once a Docker Host is successfully provisioned, a green checkmark will appear to the right of the host marking it as available for use. In the event of a failure, click into the relevant host that failed and an error explaining the failure will be displayed in red at the top.

Some common error scenarios include network connectivity. For a Docker Host to function properly, it must be able to resolve the Morpheus appliance URL which can be configured in Administration > Settings. If it is unable to resolve and negotiate with the appliance, the Morpheus Agent installation will fail and provisioning instructions will not be able to be issued to the host.

Oracle VM

Add an Oracle VM Cloud

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

API URL Oracle VM API URL. ex: <https://10.20.30.40:7002/ovm/core/wsapi/rest>

USERNAME Oracle VM User

PASSWORD Oracle VM User Password

REPOSITORY Available repositories will auto-populate upon successful authentication with the above credentials. Select appropriate repository for this Cloud.

SERVER POOL Available server pools will auto-populate upon successful authentication with the above credentials. Select appropriate server pool for this Cloud.

Inventory Existing Instances If enabled, existing Virtual Machines will be inventoried and appear as unmanaged Virtual Machines in Morpheus .

The Cloud can now be added to a Group or configured with additional Advanced options.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Oracle Cloud

Add Oracle Public Cloud

Important: A Keypair (both public and private keys) must be added to Morpheus with the Public Key in ssh-rsa format added to Oracle Cloud users keys in Oracle Cloud console for authentication.

Note: Information on uploading the Public Key and generating Tenancy's OCID and User's OCID can be found at <https://docs.cloud.oracle.com/iaas/Content/API/Concepts/apisigningkey.htm>

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

TENANCY OCID The OCID string from *Tenancy Information* section in Oracle Cloud

USER OCID OCID String for the OPC API user

SELECT KEY PAIR Select a keypair added to Morpheus matching the public key added to specified OPC API user

REGION Select the OPC region (populates after successful account authentication)

COMPARTMENT Select Compartment (populates after successful account authentication)

INVENTORY Turn on for Morpheus to discover and sync existing VMs

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Enable Live Costing for Oracle Public Cloud

Morpheus version 4.2.1 and higher support live costing data from the Oracle Cloud metering API. In order to authenticate with this API, edit your existing Oracle Cloud account integration or begin the process of newly integrating an account that wasn't previously consumable in Morpheus (Infrastructure > Clouds > +ADD).

In the advanced options section of the add/edit cloud modal for Oracle Public Cloud, the **COSTING KEY** and **COSTING SECRET** fields must be completed to work with metering API data in Morpheus. Unlike the OCI API authentication used to initially integrate Oracle Cloud, the metering API uses token-based authentication. We must access a Client ID and Client Secret value from the Oracle Public Cloud console to complete these fields.



The screenshot shows a section titled "Advanced Options" with a downward arrow icon. Below this title are two input fields. The first field is labeled "COSTING KEY" and the second field is labeled "COSTING SECRET". Both fields are empty text boxes.

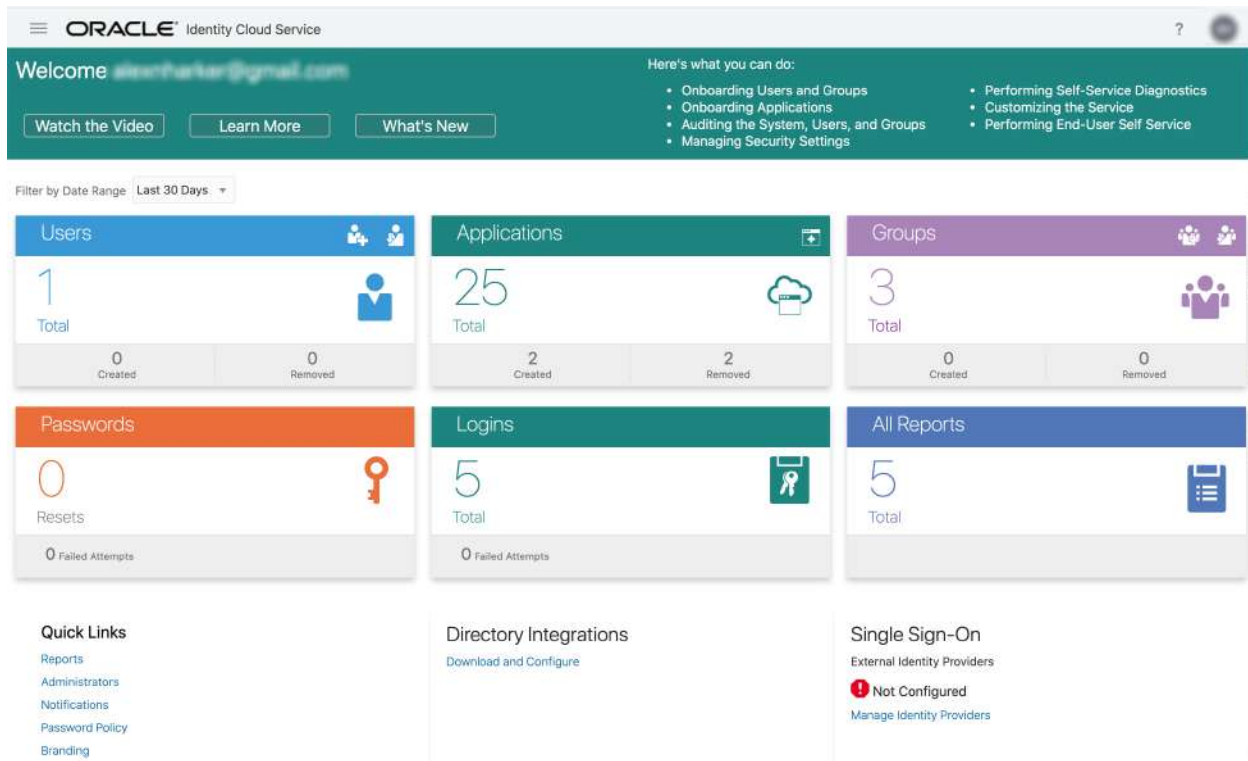
Navigate to Oracle cloud sign in page, the URL for which is similar to the following example:

```
https://idcs-00a0xxxxxxxxxxxxx.identity.oraclecloud.com/ui/v1/signin
```

If you're not redirected to the admin console similar to the one pictured below, log out and replace 'signin' at the end of the URL with 'adminconsole' as in the following example:

```
https:// idcs-00a0xxxxxxxxxxxxx.identity.oraclecloud.com/ui/v1/adminconsole
```

You'll immediately be redirected back to the same signin page but in doing that you should be taken to the admin console after authenticating your session once again.



ORACLE Identity Cloud Service

Welcome alex.harker@gmail.com

Watch the Video Learn More What's New

Here's what you can do:

- Onboarding Users and Groups
- Onboarding Applications
- Auditing the System, Users, and Groups
- Managing Security Settings
- Performing Self-Service Diagnostics
- Customizing the Service
- Performing End-User Self Service

Filter by Date Range: Last 30 Days

Users	Applications	Groups
1 Total	25 Total	3 Total
0 Created	2 Created	0 Created
0 Removed	2 Removed	0 Removed

Passwords	Logins	All Reports
0 Resets	5 Total	5 Total
0 Failed Attempts	0 Failed Attempts	

Quick Links

- Reports
- Administrators
- Notifications
- Password Policy
- Branding

Directory Integrations

Download and Configure

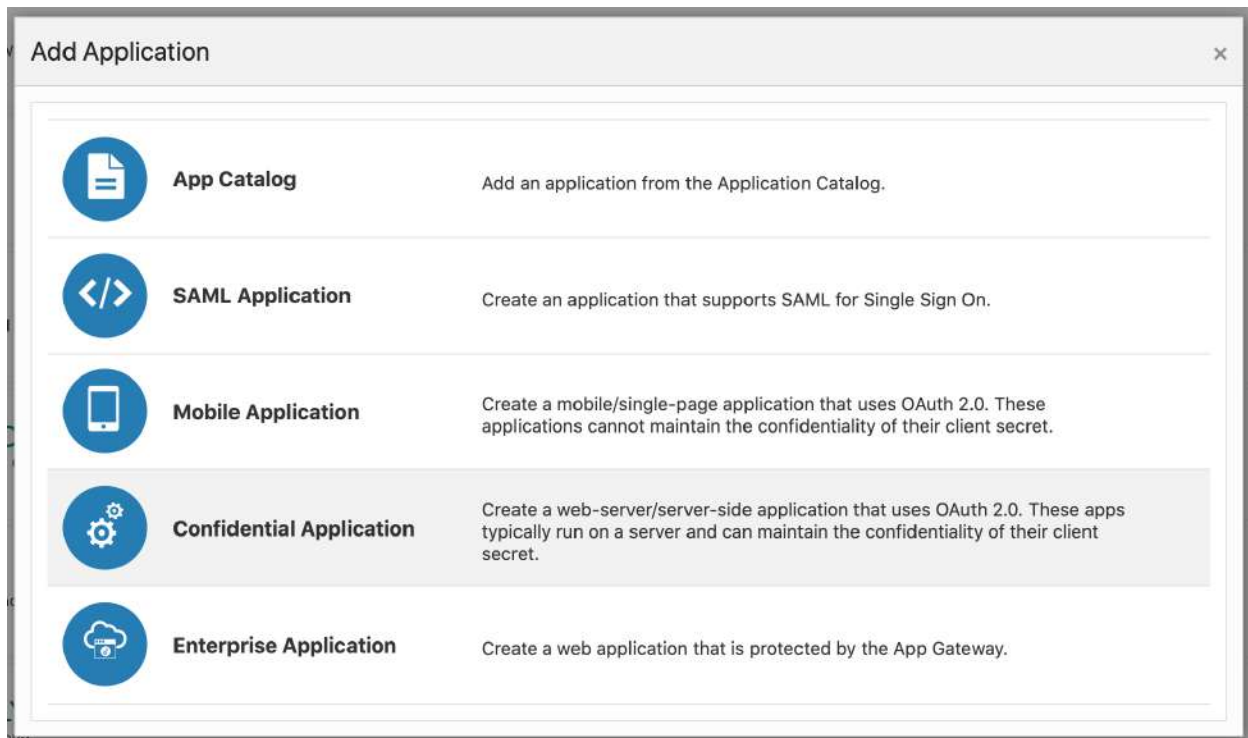
Single Sign-On

External Identity Providers

Not Configured

Manage Identity Providers

Create a new application and select the type “Confidential Application”.



Add Application

- App Catalog**: Add an application from the Application Catalog.
- SAML Application**: Create an application that supports SAML for Single Sign On.
- Mobile Application**: Create a mobile/single-page application that uses OAuth 2.0. These applications cannot maintain the confidentiality of their client secret.
- Confidential Application**: Create a web-server/server-side application that uses OAuth 2.0. These apps typically run on a server and can maintain the confidentiality of their client secret.
- Enterprise Application**: Create a web application that is protected by the App Gateway.

On the Details tab, enter a “Name” value and click “Next”.

Add Confidential Application

Cancel 1 Details 2 Client 3 Resources 4 Authorization Next >

App Details

Name morpheus

Description

Application Icon

On the Client tab, choose to “Configure this application as a client now” to reveal additional fields. Then, in the Authorization section, mark the boxes for “Client Credentials” and “JWT Assertion”.

☒ Configure this application as a client now ☐ Skip for later

Authorization

Allowed Grant Types ☐ Resource Owner ☒ Client Credentials ☒ JWT Assertion ☐ SAML2 Assertion ☐ Refresh Token ☐ Authoriza

☐ Device Code

Allow non-HTTPS URLs ☐

Redirect URL

Logout URL

Post Logout Redirect URL

Security ☐ Trusted Client ☐ Certificate

Allowed Operations ☐ Introspect ☐ On behalf Of

Bypass Consent ☐

In the Token Issuance Policy section, click the “+Add Scope” button. Click the right-facing arrow button in the row for “CloudPortalResourceApp”. Mark the box to give read access for metering and click “Add”.

Select Scope

< Back

CloudPortalResourceApp

☐ urn:opc:resource:consumer:cp:itas:myservices::read

☐ urn:opc:resource:consumer:cp:itas:myservices::all

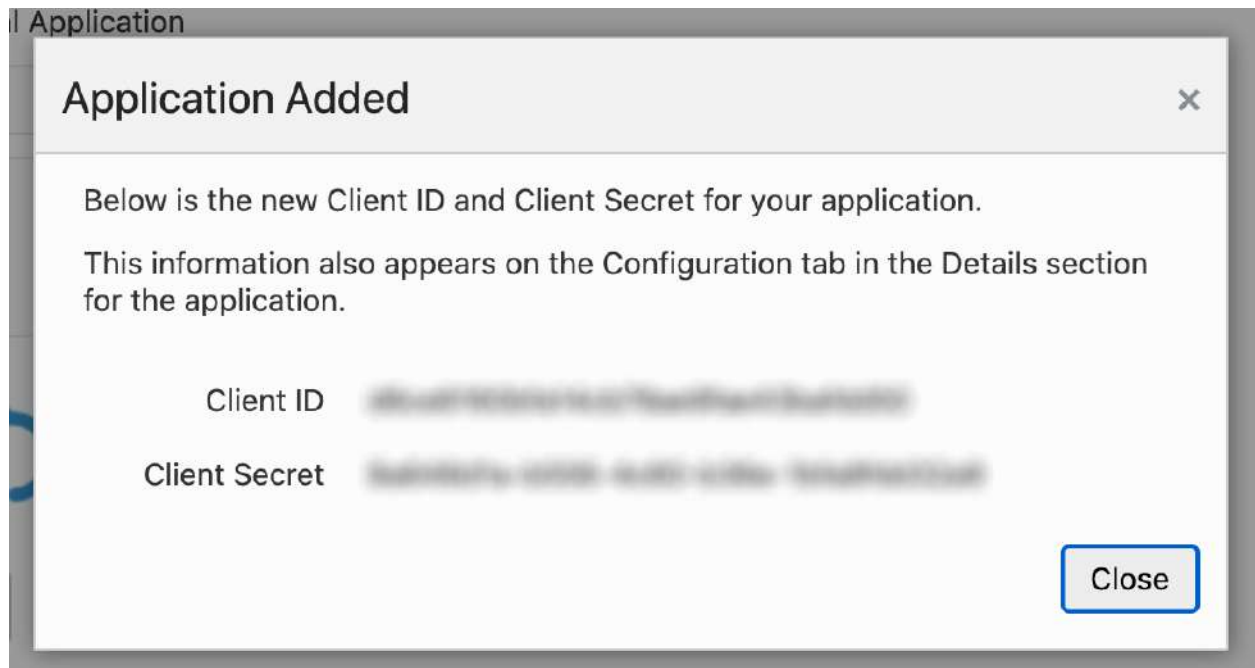
☒ urn:opc:resource:consumer:cp:itas:metering::read

☐ urn:opc:resource:consumer:cp:itas:metering::all

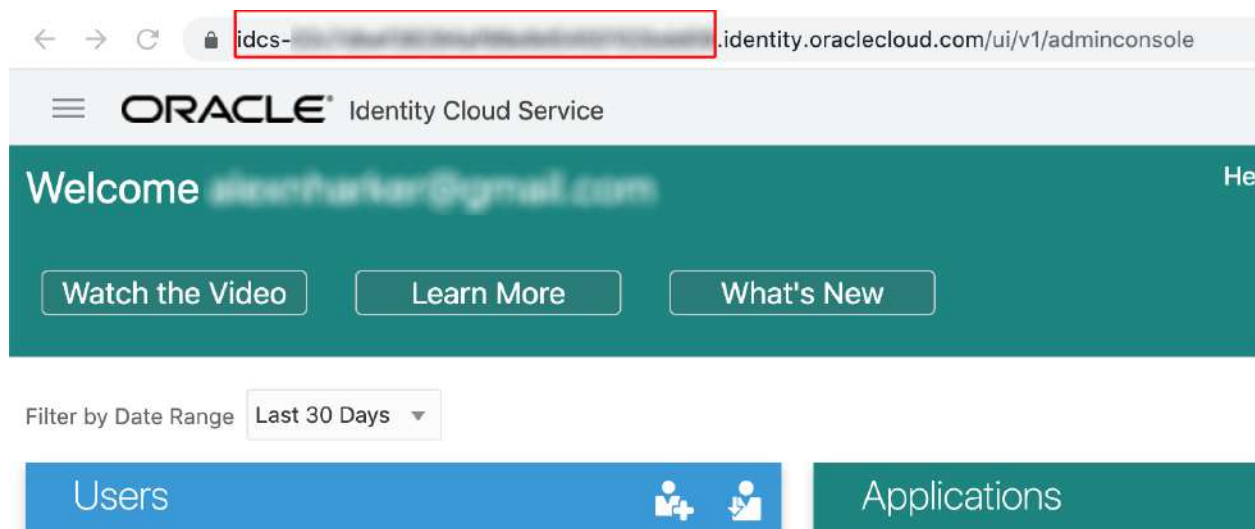
Add

Click “Next” until the “Finish” button is shown, then click “Finish”

The Client ID and Client Secret value will be shown at this point. If these values need to be referenced in the future, simply edit the application and go to the Configuration tab. The Client ID and Client Secret are shown in the General Information section.



Back in Morpheus, enter these values in the COSTING KEY and COSTING SECRET fields of the add/edit cloud modal for your Oracle Public Cloud integration. You also need to fill in the IDENTITY SERVICE value. This value can be found in the URL for your Oracle admin console as shown in the image below. It will be in a format `idcs-xxxxxx`.



Save changes to the Cloud.

SCVMM

Requirements

- Access to SCVMM host on port 5985 for Agent installation
- Morpheus Agent installation (installed on the target SCVMM host via port 5985 and WinRM)
- User with administrator privileges

Agent Requirement

SCVMM and Hyper-V integrations utilize the Morpheus Agent for communication with the Morpheus appliance, making the Morpheus Agent required. This also means SCVMM and Hyper-V Clouds can only point to one Morpheus Appliance at any given time. If another Morpheus Appliance adds an SCVMM or Hyper-V Cloud that is already managed by another Morpheus Appliance, the Morpheus Agent `appliance_url` will be updated to point to the new Morpheus appliance `url`, and the previous Morpheus Appliance will no longer be able to communicate with the SCVMM Cloud or Hyper-V Cloud until the Agent configuration is updated to point to the previous appliance again. In Morpheus version 4.2.1 and higher, multiple Morpheus clouds can be created by integrating with the same SCVMM host. This allows users to create separate Clouds with are scoped to different SCVMM Cloud, Host and/or Cluster combinations.

Note: Morpheus only supports integration with standalone SCVMM installations and not high-availability cluster installation at this time.

Add a SCVMM Cloud

1. Navigate to `Infrastructure > Clouds`
2. Select + *CREATE CLOUD*, select SCVMM, and then click *Next*.
3. Enter the following into the Create Cloud modal:

Note: You will need to open port 5985 in order for Morpheus to communicate to SCVMM. You will also want to make sure the SCVMM Controller has WinRM enabled.

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

SCVMM HOST IP address or URL of SCVMM host server

USERNAME SCVMM Username, for example: svc.scvmm

PASSWORD SCVMM user Password

CLOUD To scope the SCVMM Integration to a single Cloud, select it from the Cloud dropdown, which populates after establishing communication and authorization over port 5985 using the supplied username and password. To scope to all Clouds, leave the dropdown selection as `Select Cloud`

HOST GROUP To scope the SCVMM Integration to a single host group, select a host group from the dropdown list. To scope to all host groups, select `All Hosts`

CLUSTER To scope the SCVMM Integration to a single cluster, select a cluster from the dropdown list. To scope to all host groups, select `All`

LIBRARY SHARE Select a Library Share to be used with the cloud integration

SHARED CONTROLLER

When creating additional Morpheus clouds that point to an SCVMM host already integrated with this appliance, select the appropriate shared controller value from the dropdown.

Important: Only set `SHARED CONTROLLER` on additional Morpheus clouds and not on the Primary Morpheus SCVMM cloud. Failure to set the `SHARED CONTROLLER` on secondary Morpheus clouds pointed to the same SCVMM cluster will cause agent comm issues resulting in provisioning failures.

WORKING PATH Path for Morpheus to write to, for example `c:\cloud`

DISK PATH Path for Virtual Disks, for example `c:\virtualdisks`

HIDE HOST SELECTION FROM USERS Prevents host selection from appearing in provisioning wizards

INVENTORY EXISTING INSTANCES Enable for Morpheus to automatically discover existing VMs in the scoped resources

ENABLE HYPERVISOR CONSOLE Enable to use VNC Hypervisor Console for Morpheus console connection as opposed to the default SSH and RDP console connection methods. Requires resolution of all Hyper-V host names and access over port 443 from the Morpheus appliance to Hyper-V hosts.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in `admin -> settings`) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud.
ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

4. After clicking *NEXT*, the new Cloud can be added to a Group or configured with additional advanced options.

Softlayer

Add a Softlayer Cloud

Cloud Configuration

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

Details

Username Softlayer Username

API Key Softlayer User API Key, accessible in the Softlayer Portal under `Account -> Users -> View API Key`

Datacenter Datacenters will auto-populate upon successful authentication with the above credentials. Select appropriate Datacenter for this Cloud.

Object Store Select the destination Object Store

Inventory Existing Instances If enabled, existing Softlayer Instances will be inventoried and appear as unmanaged Virtual Machines in Morpheus .

The Cloud can now be added to a Group or configured with additional Advanced options.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME_ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER_ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK_MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL_FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY_SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST_PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE_MODE Single Disk, LVM or Clustered

BACKUP_PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION_PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS_INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE_REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG_MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE_MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT_INSTALL_MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API_PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL_AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

UCS Manager

Overview

The Morpheus UCS Manager Integration enables UCS M B and C Chassis Inventory, VM and Container Host Bare Metal Provisioning, PXE boot with IPMI, Storage Profile, SAN Connection Profile, Server Pool, BIOS Profile, Boot Profile, Maintenance Profile, UUID Pool and Disk Group Profile sync.

Adding UCS Manager Cloud

1. Navigate to *Infrastructure -> Clouds*
2. Select + *ADD*
3. Select **UCS MANAGER** from the Clouds list
4. Populate the following:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

UCS MANAGER IP or hostname of UCS Manager

USERNAME UCS Manager User

PASSWORD UCS Manager Password

ORGANIZATION

- EXISTING (select)
- **NEW (create)**
 - **ORG NAME** Enter name for the new Organization

SERVER PREFIX String provisioned servers will be prefixed with

DATA DISK MODE

- LVM data disk
- Single Disk

DATA VOLUME Defaults to `/dev/sdb` * Check to enable SOFTWARE RAID

NET INTERFACE Defaults to `eth0`

5. Select *NEXT*
6. Select an existing or create a new Group to add the Cloud to. The Cloud can be added to additional Groups in a Groups *Clouds* tab.
7. Select *NEXT*
8. Review and then Select *COMPLETE*

UpCloud

Overview

UpCloud is a cloud hosting provider that offers both Linux and Windows virtual machines on their MAXIOPS infrastructure which is billed as I.A.A.S (infrastructure-as-a-service). They have datacenters based in the UK, USA, Germany, Netherlands, Singapore and Finland. Servers can be created a lightning fast 45 seconds with their faster than SSD technology.

Features

- Virtual Machine Provisioning
- Containers
- Backups / Snapshots
- Migrations
- Auto Scaling
- Load Balancing
- Remote Console
- Periodic Synchronization
- Lifecycle Management and Resize
- Inventory
- Cloudinit

Requirements

An UpCloud User with API, Server and Storage permissions is required.

To enable API access for a Main Account UpCloud User:

1. Login to UpCloud
2. Select *My Account* -> *User Accounts*
3. Select *Change* on the target user
4. Check the box for *API connections: Allow API connections from*
5. Under *Access Permissions* -> Allow access to individual servers, check the box for *User has control access to all servers.*
6. Under *Access Permissions* -> Allow control access to individual storages, check the box for *User has control access to all storages*
7. Save

To Enable API, API, Server and Storage permissions for a SubAccount User:

When creating or editing a Sub Account UpCloud user:

1. Check the box for *API connections: Allow API connections from*
2. Under *Access Permissions* -> Allow access to individual servers, check the box for *User has control access to all servers.*
3. Under *Access Permissions* -> Allow control access to individual storages, check the box for *User has control access to all storages*
4. Save

Adding an UpCloud Cloud

Configure

1. Navigate to *Infrastructure* -> *Clouds*
2. Select *+ Create Cloud Button*
3. Select UpCloud from the Add Cloud modal
4. Select *NEXT*
5. Enter the following:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

USERNAME UpCloud User Account Username

PASSWORD UpCloud User Account Password

ZONE Select UpCloud Datacenter to scope cloud to

INVENTORY

- *Off*: Existing UpCloud Servers will not be inventoried in Morpheus
- *Basic*: Existing Servers are Inventoried with Power state, Memory and Cores statistics synced.
- *Full*: Existing Servers are Inventoried with Power state, Memory and Cores statistics, plus IP Addresses, Storage Info, and Console VNC Information.

Note: Full Inventory level recommended. Basic Inventory level can reduce Cloud Sync times when inventorying Datacenters with large amounts of servers. Credentials need to be added by editing the Virtual Machine in order to connect.

The Cloud can now be added to a Group or configured with additional Advanced options.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Group

A Group must be specified or created for the new Cloud to be added to. Clouds can be added to additional Groups or removed from Groups after being created.

- **USE EXISTING:** Add the new Cloud to an exiting Group in Morpheus .
- **CREATE NEW:** Creates a new Group in Morpheus and adds the Cloud to the Group.

Review

Confirm all settings are correct and select *COMPLETE*.

The UpCloud Cloud will be added, and Morpheus will perform the initial cloud sync of:

- UpCloud Servers will added as Virtual Machines (if Inventory is enabled)
- UpCloud Templates (My Templates) will sync and be added to ``Provisioning -> Virtual Images`.

Note: The Console tab will only appear for Inventoried Servers if Inventory Level is set to *Full*

Provisioning to UpCloud

Instances and Apps can be created using the private Images synced from UpCloud or from the Morpheus provided Image Catalog.

Provision a synced Image

Images synced from UpCloud can be provisioned by using:

- The *UPCLOUD* Instance Type and selecting the Image from the Image dropdown in the configure section when provisioning and Instance, App, or creating an App Blueprint.
- Creating custom Library Instance Types and selecting a synced Image when creating a Node Type for the custom Instance Type.

Important: Synced images should be configured prior to provisioning by editing the Image in the *Provisioning -> Virtual Images* section.

Provision a Morpheus provided UpCloud Image

Morpheus provides a number of pre-configured Images that are available in the default Morpheus Catalog when provisioning and Instance, App, or creating an App Blueprint. UpCloud Images are included in the following Instance Types in the default Morpheus catalog.

- ACTIVEMQ
- APACHE
- CASSANDRA
- DEBIAN
- ELASTICSEARCH
- GRAILS
- JAVA
- MONGO
- MYSQL
- NGINX

- PHP
- RABBITMQ
- REDIS
- OMCAT
- UBUNTU
- WINDOWS
- GRAILS

vCloud Director

Configuration

Add vCD Cloud From Infrastructure > Clouds

1. Navigate to Infrastructure -> Clouds
2. Select + *ADD*
3. Select **VCLLOUD DIRECTOR** from the Clouds list
4. Select *NEXT*
5. Populate the following:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

API URL

vCloud Director API Url Example: `https://org.vcd.company.com`

USERNAME vCD Organization Administrator User

..NOTE:: User must have an Organizational Administrator Role in the selected Origination for successful provisioning

PASSWORD vCD Organization Administrator User password

ORGANIZATION Select Organization. Dropdown populates upon successful authorization.

VDC Select VDC. Dropdown populates upon successful authorization.

Inventory Existing Instances If enabled, existing Virtual Machines will be inventoried and appear as unmanaged Virtual Machines in Morpheus .

NOTE: Multiple Organizations/VDC's can be added by creating additional Clouds in Morpheus.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus

will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

6. Select *NEXT*
7. Select an existing or create a new Group to add the Cloud to. The Cloud can be added to additional Groups in a Groups *Clouds* tab.
8. Select *NEXT*
9. Review and then Select *COMPLETE*

Add vCD Cloud From *Infrastructure -> Groups*

1. Navigate to *Infrastructure -> Groups*
2. Select a Group
3. Select the *CLOUDS* tab
4. Scroll down to **VCLLOUD DIRECTOR** and select + *ADD*
5. Populate the following:

Name Name of the Cloud in Morpheus

Location Description field for adding notes on the cloud, such as location.

Visibility For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

API URL

vCloud Director API Url Example: `https://org.vcd.company.com`

USERNAME vCD Organization Administrator User

NOTE:: User must have an Organizational Administrator Role in the selected Origination for successful provisioning

PASSWORD vCD Organization Administrator User password

ORGANIZATION Select Organization. Dropdown populates upon successful authorization.

VDC Select VDC. Dropdown populates upon successful authorization.

Inventory Existing Instances If enabled, existing Virtual Machines will be inventoried and appear as unmanaged Virtual Machines in Morpheus .

NOTE: Multiple Organizations/VDC's can be added by creating additional Clouds in Morpheus.

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud.
ex: Cherwell

AGENT INSTALL MODE

- SSH / WINRM: Morpheus will use SSH or WINRM for Agent install.
- Cloud Init / Unattend (when available): (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

6. Select *NEXT*

7. Review and then Select *COMPLETE*

How to create vCloud Director templates for Morpheus

To create a Windows Template

Create a new machine in VMware vCenter and install a base version of your preferred Windows build.

1. Apply any service packs / updates to the operating system.
2. Set the Network location to Private the below PowerShell will set the location.

```
Get-NetConnectionProfile | Set-NetconnectionProfile -NetworkCategory private
```

3. Configure WinRM to allow remote management and open the firewall.

- To do this, under local computer Administrator, open a command prompt and run `winrm quickconfig`

4. Install VMware tools
5. Install .Net at least 4.5
6. Enable remote PowerShell this can be done in PowerShell.

```
Enable-PSremoting
```

7. Shutdown the virtual machine and convert to a template.

Note: Do not run sysprep

To create a Linux Centos template

Create a new machine in VMware vCenter and install a base version of your preferred Linux distro build. If you are using cloud init as part of your image you will need to ensure your virtual machine has a cdrom.

1. Before installing the operating system setup a single `ext` or `xfs` partition without a swap disk (This is so that growpart can extend the disk. growpart currently does not support lvm)
2. Install the distro and apply any updates to the operating system and security updates
3. Install cloud-init using command `yum install cloud-init`
4. Install cloud-utils-growpart using command `yum install cloud-utils-growpart`
5. Install vmware tools
6. Install git by running `yum install git`
7. `epel-release`
8. `selinux` set to permissive (enforced can cause problems with cloud-init)

To create a Linux Ubuntu template

Create a new machine in VMware vCenter and install a base version of your preferred Linux distro build. If you are using cloud init as part of your image you will need to ensure your virtual machine has a cdrom.

1. Before installing the operating system setup a single `ext` partition without a swap disk (This is so that growpart can extend the disk. growpart currently does not support lvm)
2. Install the distro and apply any updates to the operating system and security updates
3. Ensure you have set a root password
4. Install cloud-init by running `sudo apt install cloud-init`
5. Install cloud-utils-growpart `sudo apt install cloud-utils`
6. Install desired hypervisor drivers (Virtio, Open-VM Tools)
7. Install git by running `sudo apt install git`
8. As Debian 9 includes network manager ensure this is disabled. Change the below file

```
/etc/NetworkManager/NetworkManager.conf
```

to the following:

```
managed=false
```

We also recommend disabling network manager and setting the network adapter to `eth0` rather than the automatically assigned name. <https://support.morpheusdata.com/hc/en-us/articles/115002881228-Creating-a-CentOS-7-Morpheus-VMware-Image>

To import your template into vCloud director you will need to login as either an administrator or organisation administrator.

Once logged into vCloud director you will then need select `Manage Organizations` and then select your organization.

From within the organisation click on `Catalogues` > select an existing catalogue or create a new catalogue.

Note: Please note once you connect Morpheus to your vCD environment, it will create a catalogue called Auto Morpheus. This is a working catalogue and is ignored by Morpheus when searching for images, so any images in the catalogue will not be synced into Morpheus

Open the catalogue and select the import template from vCenter and then browse the data stores for your templates. Select your template and the type in a new name and description then check the copy template into vCloud director.

Once you click ok the import process will begin. When the import has completed the template will appear in Morpheus within `Provisioning > Virtual Images`

If the image does not appear within the virtual images you may need to use the filters to filter the virtual images by the vmware (vmdk / ovf / ova) type.

You may also need to refresh the cloud. To do this go to `Infrastructure > Clouds > select the vCloud Director cloud > select Refresh`.

VMware vCenter

Overview

VMware is a very common cloud integration choice supported by Morpheus . They have provided a top notch virtualization solution and one might argue pioneered the virtualization space altogether. As such, many companies utilize this technology and all the features that come with it, so Morpheus covers a broad feature set in vCenter.

Features

- Virtual Machine Provisioning
- Backups / Snapshots
- Resource Groups
- Datastores and DRS Clusters
- Distributed Switches
- Datacenter / Cluster scoping
- Brownfield VM management and migration
- VMware to VMware migrations
- VMDK/OVF image conversion support
- Hypervisor Remote Console
- Periodic Synchronization
- Veeam Backup Integration
- Lifecycle Management and Resize
- Metadata tag sync

On top of all these features, Morpheus also adds additional features to VMware that do not exist out of the box to make it easier to manage in multitenant environments as well as hybrid cloud environments:

- Cloud-Init Support
- VHD to VMDK Image Conversion

- QCOW2 to VMDK Image Conversion
- Multitenancy resource allocation
- Virtual Image management (Blueprints)
- Auto-scaling and recovery

Getting Started

To get started with VMware, simply start by adding a Cloud in the Infrastructure -> Clouds section.

To start adding a VMware cloud there will be some things you will need:

vCenter API Url Typically this is the url to the vCenter web client with a /sdk in the path

Username/Password A set of credentials with high level access to VMware (ensure the account has Datacenter level access)

Once these fields are entered, some selections will start pre-populating. A cloud integration is scoped to a specific data center, and can optionally be scoped down to a single cluster or even a single resource pool. If the drop downs do not populate, please verify the api url is resolvable, morpheus has access to vCenter on 443, and the provided credentials are correct and the user has sufficient permissions.

Another cool feature provided with the cloud integration is optional *Resource Pool* scoping. One can choose to allow the cloud to provision into All Resource Pools or a singular Resource Pool. When choosing *All*, these Resource Pools can be managed from a sub-account and visibility perspective via the Cloud Detail page (multi-tenancy).

The VMware cloud integration provides a few additional options including allowing users to make host selections or keeping that aspect hidden such that the best host is automatically chosen for the requested provision.

The *RPC Mode* feature can be configured to allow Morpheus to install its agent on the Guest operating system via either SSH/WinRM or VMware Tools Guest Process feature. The VMware tools Guest Execution API can be tricky so it is recommended to use SSH/WinRM if possible. However, if it is not possible for the Appliance to have outbound access to all networks in which VMs are being provisioned to the SSH/WinRM ports (22, 5985 respectively) then Guest Execution is the only option.

The *Use VNC* console option on the VMware cloud requires special configuration on each ESXI host but allowed hypervisor level remote console support. (See the Advanced Section for details)

When following this add cloud wizard an option will be presented to create a group or add to an existing group. These groups can be given provisioning permission via role based access control. It is normally recommended that groups are organized such that one cloud exists in one group unless the networks are setup such that internal routing is possible between the clouds. This is very useful for bursting, or hybrid cloud configurations.

Windows Provisioning Tips

By default when provisioning windows templates, Morpheus performs guest customizations which initiates a sysprep. This resets the Administrator user and password. Morpheus will set the Administrator password from Administration > Provisioning > Windows Settings > Password.

Users can also set the username on an image as Administrator and enter a different password if unique passwords are required per image.

Guest customizations are required when assigning static IP's manually or using IP pools. They can be disabled per virtual image advanced settings under Provisioning > Virtual Images > Edit Image > Advanced > Uncheck "Force Guest Customization" if using DHCP. However the SID will not be changed from the source template. In addition, new VM's will not be able to join a domain that had already been joined by the source template or any other VM's with that SID.

Existing Instances

Morpheus provides several features regarding pulling in existing virtual machines and servers in an environment. Most cloud options contain a checkbox titled '*Inventory Existing Instances*'. When this option is selected, all VMs found within the specified scope of the cloud integration will be scanned periodically and Virtual Machines will be synced into Morpheus. By default these virtual machines are considered 'unmanaged' and do not appear in the Provisioning -> Instances area but rather Infrastructure -> Hosts -> Virtual Machines. However, a few features are provided with regards to unmanaged instances. They can be assigned to various accounts if using a multitenant master account, however it may be best suited to instead assign the 'Resource Pool' to an account and optionally move all servers with regards to that pool (more on this later). A server can also be made into a managed server. During this process remote access is requested and an agent install is performed on the guest operating system. This allows for guest operations regarding log acquisition and stats. If the agent install fails, a server will still be marked as managed and an Instance will be created in *Provisioning*, however certain features will not function. This includes stats collection and logs.

Note: All Cloud data is resynchronized on a 5 minute interval. This includes Datastores, Resource Pools, Networks, Blueprints, and Virtual Machines.

Service Plans

A default set of Service Plans are created in Morpheus for the VMware provisioning engine. These Service Plans can be considered akin to AWS Flavors or Openstack Flavors. They provide a means to set predefined tiers on memory, storage, cores, and cpu. Price tables can also be applied to these so estimated cost per virtual machine can be tracked as well as pricing for customers. By default, these options are fixed sizes but can be configured for dynamic sizing. A service plan can be configured to allow a custom user entry for memory, storage, or cpu. To configure this, simply edit an existing Service Plan tied to VMware or create a new one. These all can be easily managed from the Admin -> Plans & Pricing section.

Virtual Images / Blueprints

Morpheus will automatically take an inventory of all blueprints configured in vCenter and present them as options during provisioning. However, in order for Morpheus to properly provision these virtual machines and provide accurate stats and health of these virtual machines, an agent must be installed during virtual machine startup. This means remote access needs to be granted at the guest operating system level to Morpheus. To properly configure these virtual images, find the relevant images in Provisioning -> Virtual Images and edit the entry. On this form, a few options are presented. The first is a check box asking whether or not cloud-init is enabled. If cloud-init is enabled, simply provide the default OS username configured (for Ubuntu the username is *ubuntu* and for CentOS the username is *centos*). For those looking to add cloud-init to existing blueprints Morpheus requires no special configuration and can use the default *cloud.cfg* settings.

A global cloud-init username/password can also be configured per account as well as a keypair via the Admin->Provisioning settings section. The great benefit of utilizing cloud-init is default blueprints do not need common credential sets thereby increasing provisioning security.

Windows systems do not typically support cloud-init. So simply turn this checkbox off and provide the *Administrator* credentials. It should be noted that these credentials are encrypted in the database. If using WinRM for the RPC Mode instead of VMware tools, a Local or Domain Administrator account credential set can be provided instead.

Snapshots

Morpheus allows the ability to create a snapshot of a VM in VMware vCenter. From the instance detail page, simply select **Actions** -> **Create Snapshot** to begin creation of a new Snapshot. Existing snapshots can be viewed in the **BACKUPS** tab on the instance detail page. Snapshots taken in vCenter will sync into Morpheus every five minutes. To revert to a previous snapshot, click on the revert icon located on the right side of the Snapshot. Snapshots can be deleted by clicking on the trash can icon.

Note: Access to Snapshots can be limited or removed entirely for specific user roles as needed. To edit a role's Snapshots permissions, go to **Administration > Roles > (Your selected role) > Snapshots**. Users can be given Full, Read-only, or No access.

Tagging and Metadata

As of Morpheus version 4.1.0, tagging support is included for vCenter in addition to the other clouds that have already supported it in past versions. Tags will sync to vCenter from Morpheus and existing tags are also inventoried from vCenter into Morpheus.

Note: This feature requires a minimum API version of vCenter 6.5. The API version can be edited by navigating to **'Infrastructure > Clouds'** and clicking the edit (pencil) button in the row for the relevant cloud. The field is labeled **'VERSION'**.

Tags can be created on-demand when provisioning from the **'CONFIGURE'** tab of the **'CREATE INSTANCE'** wizard (Provisioning > Instances). Within the **'Metadata'** drawer, you will see sets of fields to enter key/value pairs. On creation of the instance, this metadata will be synced into vCenter.

'Option Types' from your library can also be exported as metadata for use with vCenter. When adding or editing a new Option Type (Provisioning > Library > **OPTION TYPES**), simply mark the box labeled **'EXPORT AS METADATA'**. The **'FIELD NAME'** becomes the tag category in VMWare.

CREATE INSTANCE

TYPE > GROUP > **CONFIGURE** > AUTOMATION > REVIEW

Configuration Options

VERSION: 7.5

LAYOUT: VMware VM

PLAN: [Progress Bar]

Resource Pool: [Progress Bar]

VOLUMES: root | 10 GB | SCSI 0:0 | Auto - Cluster

NETWORKS: Select Network

HOST: Select

FOLDER: Demo

▶ User Config
 ▶ Network Options
 ▶ Advanced Options
 ▼ Metadata
 METADATA: [] []
 ▶ Environment
 ▶ Ansible

PREVIOUS NEXT

Docker

So far this document has covered how to add the VMware cloud integration and has enabled users the ability to provision virtual machine based instances via the *Add Instance* catalog in *Provisioning*. Another great feature provided by Morpheus out of the box is the ability to use Docker containers and even support multiple containers per Docker host. To do this a Docker Host must first be provisioned into VMware (multiple are needed when dealing with horizontal scaling scenarios).

To provision a Docker Host simply navigate to the Clusters tab of the Cloud detail page or Infrastructure > Clusters section. From there, click + **ADD CLUSTER** to add a VMware Docker Host. This host will show up in the Hosts tab next to other ESXi servers that were inventoried by the VMware cloud integration. Morpheus views a Docker host just like any other Hypervisor with the caveat being that it is used for running containerized images instead of virtualized ones. Once a Docker Host is successfully provisioned a green checkmark will appear to the right of the host marking it as available for use. In the event of a failure click into the relevant host that failed and an error explaining the failure will be displayed in red at the top.

Some common error scenarios include network connectivity. For a Docker Host to function properly, it must be able to resolve the Morpheus appliance url which can be configured in Administration > Settings. If it is unable to resolve and negotiate with the appliance then the agent installation will fail and provisioning instructions will not be able to be issued to the host.

Multitenancy

A very common scenario for Managed Service Providers is the need to provide access to VMware resources on a customer by customer basis. With VMware several administrative features have been added to ensure customer resources are properly scoped and isolated. For VMware it is possible to assign specific *Networks*, *Datastores*, and *Resource Pools* to customer accounts or even set the public visibility of certain resources, therefore allowing all sub accounts access to the resource.

MORPHEUS Bertram Labs

Dashboard Provisioning **Infrastructure** Backups Logs Monitoring Reports Admin

Groups Clouds Hosts Load Balancers Storage Key Pairs Certificates Boot Security Groups

Clouds > Labs VMware

LABS VMWARE [EDIT](#) [DELETE](#)

Location: San Mateo
Cloud Group: Labs VMware
Servers: 4

HOSTS VIRTUAL MACHINES BARE METAL SECURITY GROUPS LOAD BALANCERS NETWORKS DATA STORES **RESOURCE POOLS**

Search

NAME	VISIBILITY	ACCOUNT	ACTIONS
Resources	Private	morpheus-qa	ACTIONS ▾
Brian	Private	morpheus-qa	ACTIONS ▾
Macbook	Private	morpheus-qa	ACTIONS ▾
David	Private	morpheus-qa	ACTIONS ▾
Macbook	Private	morpheus-qa	ACTIONS ▾

MORPHEUS © 2016 MORPHEUS DATA, LLC. ALL RIGHTS RESERVED
TERMS AND CONDITIONS | PRIVACY POLICY

Advanced

There are several advanced features provided within Morpheus that can leverage some cool aspects of VMware. One of these features is Remote Console support directly to the hypervisor. To enable this feature a few prerequisites must be met. First, the Morpheus appliance must have network access to the ESXi hosts within VCenter. Secondly, firewall settings need to be adjusted on each ESXi host. This can be done in VSphere under firewall configuration on the host. Simply check the *gdbserver* option, which will open up the necessary ports (starting at 5900 range).

Important: Hypervisor Console for vCenter 6.5 requires Morpheus v3.2.0+

Now that the ESXi hosts are ready to utilize remote console, simply edit the cloud in Morpheus via Infrastructure -> Clouds. Check the option that says *Use VNC*. It is important to note that currently this functionality only works for newly provisioned vm's provisioned directly via Morpheus . This should change soon however.

It is also possible to import vm snapshots for backup or conversion purposes from VCenter and also an ESXi host. However, this does require that the ESXi host license has an enterprise level license as it will not allow the appliance

to download a virtual image if it is not a paid VMware license.

VMware Permissions

Usage

vCenter

- Non-Propagating

Datacenter

- Non-Propagating

Cluster

- Non-Propagating

Host

- Non-Propagating

Datastore/Datastore Cluster

- Propagating

Privileges

Datastore/Datastore Cluster

- Allocate Space
- Browse Datastore
- Low Level file Operations
- Remove File
- Update virtual machine files
- Update virtual machine metadata

Distributed Switch

- Port configuration operation
- Port setting operation

Global

- Log Event
- Manage custom attributes
- Set custom attribute

Network

- Assign Network
- Configure
- Remove

Resource

- Apply recommendation
- Assign vApp to resource pool
- Assign virtual machine to resource pool
- Migrate powered off virtual machine
- Migrate powered on virtual machine

Scheduled task

- Create tasks
- Modify task
- Remove task
- Run task

Tasks

- Create task
- Update task

Virtual Machine

- Configuration (all)
- Guest Operations (all)
- Interaction (all)
- Inventory (all)
- Provisioning (all)
- Service configuration (all)
- Snapshot management (all)
- vSphere Replication (all)

vApp

- Clone
- Export
- Import

vSphere Tagging

- Assign or Unassign vSphere Tag
- Create vSphere Tag
- Create vSphere Tag Category
- Delete vSphere Tag
- Delete vSphere Tag Category
- Edit vSphere Tag
- Edit vSphere Tag Category
- Modify UsedBy Field For Category
- Modify UsedBy Field For Tag

- `privilege.InventoryService.Tagging.CreateScope.label`
- `privilege.InventoryService.Tagging.DeleteScope.label`

Creating a Morpheus VMware Image

Overview

Morpheus comes out of the box with a default set of blueprints for use in many modern deployment scenarios. These consist mostly of base operating system images with a few additional adjustments. These adjustments typically include the addition of cloud-init (which is highly recommended to be used in most environments, but not mandatory). However, in many on-premise deployments there are custom image requirements as well as networking requirements. This guide will go over how to create a VMware Images for use within Morpheus.

Supported Versions

2008R2,2012,2012R2,2016,2019

Creating a Windows Image

Create a new machine in VMware vCenter and install a base version of your preferred Windows build. The smaller the VMDK drive, typically the faster you can clone and deploy. Utilizing Morpheus, provisioning and post deploy scripts can expand drives to desired sizing.

1. Ensure VMtools is installed on the operating system.
2. Apply any service packs / updates to the operating system.
3. Configure WinRM to allow remote management and open the firewall. This is optional if using VMtools RPC mode for agent install and Morpheus Agent for guest exec. To enable this, under local computer Administrator, open a command prompt and run

```
winrm quickconfig
```

4. Install .Net at least 4.5
5. Ensure Windows Firewall will allow WinRM connections.
6. Shutdown the virtual machine and convert to a template.

Note: WinRM is not required and is used as a fallback when using vmtools guest exec and customizations

Note: Morpheus will sysprep images based on the “Force Guest Customizations” flag under the Virtual Image’s settings when using DHCP. Ensure a sysprep has not been performed on the template if this flag is enabled or if using Static IPs/IP Pools when provisioning, which will always use Guest Customizations and trigger a sysprep.

Creating a CentOS/RHEL Image

Create a new virtual machine in VMware vCenter and install a base version of your preferred Linux distro build. If you are using cloud init as part of your image you will need to ensure your virtual machine has a cdrom.

1. Before installing the operating system setup a single `ext` or `xfs` partition without a swap disk (This is so that `growpart` can extend the disk. `growpart` currently does not support `lvm`)
2. Install the distro and apply any updates to the operating system and security updates
3. Install cloud-init using command `yum install cloud-init`
4. Install cloud-utils-growpart using command `yum install cloud-utils-growpart`
5. Install open-vm-tools using command `yum install open-vm-tools`
6. Install git by running `yum install git`
7. Install epel-release repo using command `yum install epel-release`
8. selinux set to permissive (enforced can cause problems with cloud-init) `sudo vi /etc/selinux/config`

Cloud-Init

To get started with a base CentOS image we first install cloud-init. This is a relatively simple process using yum:

```
yum -y install epel-release
yum -y install git wget ntp curl cloud-init dracut-modules-growroot
rpm -qa kernel | sed 's/^kernel-//' | xargs -I {} dracut -f /boot/initramfs-{}.img {}
```

There are two parts to this yum installation. We are first ensuring some core dependencies are installed for automation as well as cloud-init. git for example is installed for use by ansible playbook automation down the line and is therefore optional if not using ansible. The dracut-modules-growroot is responsible for resizing the root partition upon first boot to match the virtual disk size that was potentially adjusted during provisioning.

A great benefit to using cloud-init is credentials don't have to be locked into the blueprint. It is advisable, within Morpheus, to configure the default cloud-init user that gets created when the vm boots automatically by cloud-init. This is located in the *Administration -> Provisioning -> Cloud-Init Settings* section.

Network Interfaces

A slightly annoying change with CentOS 7 is that the network interfaces have changed naming convention. You may notice when running `ifconfig` that the primary network interface is set to something like `ens2344` or some other random number. This naming is dynamic typically by hardware id and we don't want this to fluctuate when provisioning the blueprint in various VMware environments. Fortunately, there is a way to turn this functionality off and restore the interface back to `eth0`.

Firstly we need to adjust our bootloader to disable interface naming like this.

```
sed -i -e 's/quiet/quiet net.ifnames=0 biosdevname=0/' /etc/default/grub
grub2-mkconfig -o /boot/grub2/grub.cfg
```

The above command adds a few arguments to the kernel args list (namely `net.ifnames=0` and `biosdevname=0`). It may be useful to view the `/etc/default/grub` file and ensure these settings were indeed applied.

The next step is to adjust the network-scripts in CentOS. we need to ensure we have a file called `/etc/sysconfig/network-scripts/ifcfg-eth0`

Below is a script that we run on our packer builds to prepare the machines network configuration files.

```
export iface_file=$(basename "$(find /etc/sysconfig/network-scripts/ -name 'ifcfg*' -
↪not -name 'ifcfg-lo' | head -n 1)")
export iface_name=${iface_file:6}
echo $iface_file
echo $iface_name
sudo mv /etc/sysconfig/network-scripts/$iface_file /etc/sysconfig/network-scripts/
↪ifcfg-eth0
sudo sed -i -e "s/$iface_name/eth0/" /etc/sysconfig/network-scripts/ifcfg-eth0
sudo bash -c 'echo NM_CONTROLLED="no" >> /etc/sysconfig/network-scripts/ifcfg-eth0'
```

This script tries to ensure there is a new ifcfg-eth0 config created to replace the old ens config file. Please do verify this config exists after running. If it does not you will have to be sure to build one on your own.

```
TYPE=Ethernet
DEVICE=eth0
NAME=eth0
ONBOOT=yes
NM_CONTROLLED="no"
BOOTPROTO="dhcp"
DEFROUTE=yes
```

Creating an Ubuntu Image

Create a new machine in VMware vCenter and install a base version of your preferred Linux distro build. If you are using cloud init as part of your image you will need to ensure your virtual machine has a cdrom.

1. Before installing the operating system setup a single `ext` partition without a swap disk (This is so that growpart can extend the disk. growpart currently does not support lvm)
2. Install the distro and apply any updates to the operating system and security updates
3. Ensure you have set a root password
4. Install cloud-init by running `sudo apt install cloud-init`
5. Install cloud-utils-growpart `sudo apt install cloud-utils`
6. Install desired hypervisor drivers (Virtio, Open-VM Tools)
7. Install git by running `sudo apt install git`
8. As Debian 9 includes network manager ensure this is disabled, set `/etc/NetworkManager/NetworkManager.conf` to `managed=false`

We also recommend disabling network manager and setting the network adapter to eth0 rather than the automatically assigned name as described in the CentOS/RHEL section above.

Gotyas

SELinux can cause issues with cloud-init when in enforced mode. It may be advisable to set this to permissive unless it is mandatory within your organization to use an enforced SELinux configuration. If that is the case please see the documentation for the `cloud_init_t` security policies.

Network Manager will also prevent the required restart of the Network Service when assigning static IP's. Disable Network Manager when possible or Static IP assignment may not work until the Network Service is restarted manually.

A Note on Proxies

Proxy configurations are known to vary in some organizations and makes building a base blueprint a little more difficult. In order to fully configure proxies a few environment variables must be set in the `/etc/environment` file (This can be done automatically in a default user-data script for cloud-init as well in edit cloud).

```
http_proxy="http://myproxyaddress:8080"
https_proxy="http://myproxyaddress:8080"
ftp_proxy="http://myproxyaddress:8080"
no_proxy=127.0.0.1,localhost,applianceUrl
https_no_proxy=127.0.0.1,localhost,applianceUrl
```

Important: It is very important to properly set the `no_proxy` list (`applianceUrl`) should be replaced with the actual appliance url. In future releases, morpheus plans to automatically take care of this.

Note: If using cloud-init agent install mode these settings need to be set in the custom Cloud-Init User data section of “Edit Cloud” or “Edit Virtual Image”

Important: If using this virtual machine as a docker host, proxy settings must also be configured in the docker config. See Docker guides for instructions on how to properly set this. If necessary this can be wrapped in a task automation workflow for your own use.

Introduction

This guide is designed to help you get started and quickly get the most out of Morpheus with VMWare. By the end, you will integrate your first cloud, configure networking, prepare and consume images, provision instances, and get started with automation. We will briefly discuss installation and account setup but will provide links to additional resources for those very first steps. For the most part, this guide assumes you are able to get Morpheus installed and are ready to move forward from that point. There is a lot more to see and do in Morpheus that is beyond the scope of this guide. For more, consult the complete Morpheus documentation or take part in our user community forum.

Installation & Setup

In the simplest configuration, Morpheus needs one appliance server which will contain all the components necessary to orchestrate virtual machines and containers. Full requirements, including storage and networking considerations, can be found in Morpheus documentation [here](#). In order to provision any new instances, hosts, or applications, (or convert any discovered resources to managed resources) you will need a valid license. If you don’t have one, you can request a lab license for free at [Morpheus Hub](#). Once obtained, the license can be applied in Administration > Settings > LICENSE.

Groups

Groups in Morpheus define which resources a user has access to. Clouds are added to groups and a user can only access clouds that are in the groups to which their roles give them access. More information on Morpheus groups is [here](#). A deep dive into groups goes beyond the scope of this guide but it's often useful to create a group that contains all clouds for testing purposes. We will create that group now so that we can add our first cloud into this group in the next section.

Navigate to *Infrastructure > Groups*. Here we will see a list of all configured groups but, of course, this will be empty immediately after installation. Click “+CREATE”. Give your group a name, such as “All Clouds”. The “CODE” field is used when calling Morpheus through Morpheus API or Morpheus CLI. It's useful in most cases to have an “All Clouds” group for testing purposes so this will likely help you down the road.

NEW GROUP

Configuration

NAME

CODE

LOCATION

► Advanced Options

SAVE CHANGES

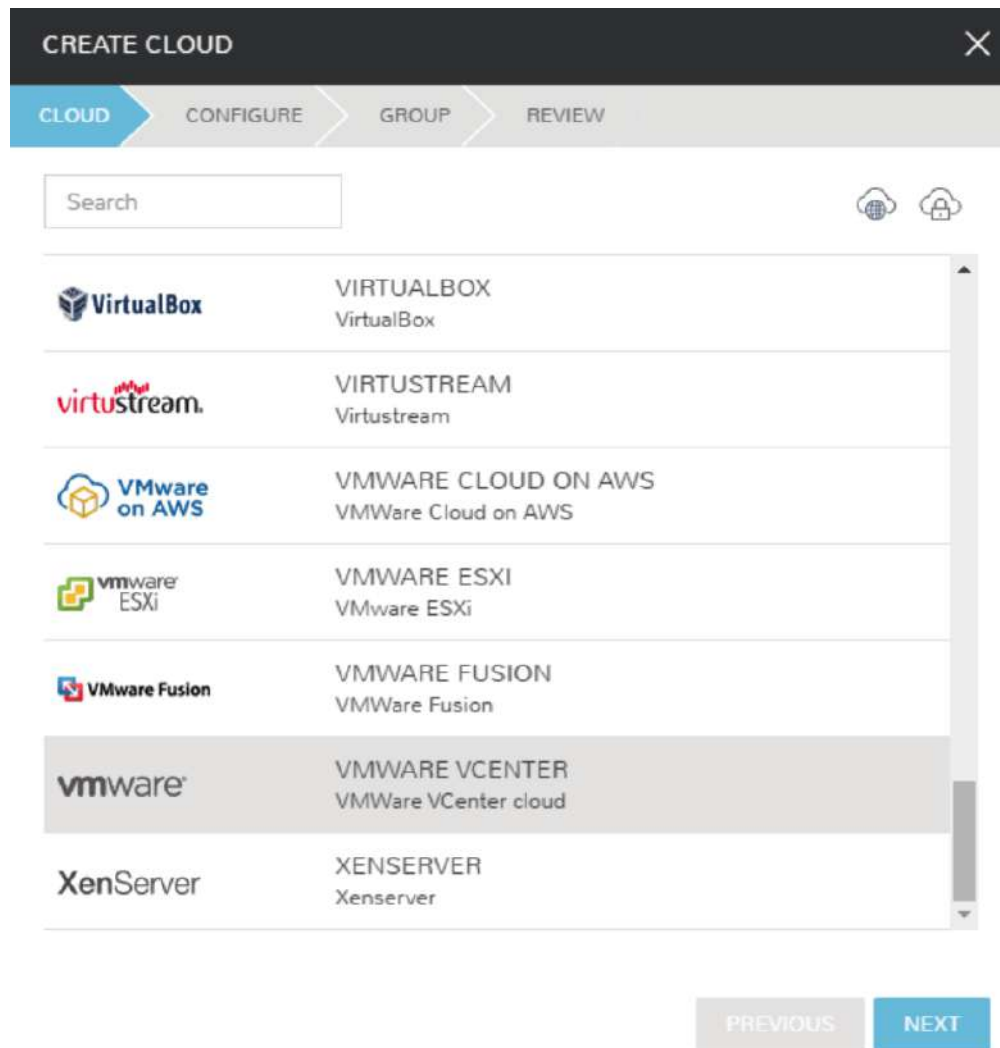
Click “SAVE CHANGES”. Your group is now ready to accept clouds.

Integrating Your First Cloud

Clouds in Morpheus consist of any consumable endpoint whether that be On-Prem, Public clouds, or even bare metal. In this guide, we will focus on integrating and working with VMWare vCenter.

To get started, we will navigate to *Infrastructure > Clouds*. This is the cloud detail page which lists all configured clouds. It will be empty if you've just completed installation and setup of Morpheus but soon we will see our integrated vCenter cloud here.

Click the “+ADD” button to pop the “CREATE CLOUD” wizard. Select “VMWARE VCENTER” and click the “NEXT” button.



On the “CONFIGURE” tab, we’re asked to set the initial connection strings into vSphere. The **API URL** should be in the following format: <https://<URL>/sdk>. The **USERNAME** should be in [user@domain](#) format.

CREATE CLOUD [X]

CLOUD > **CONFIGURE** > GROUP > REVIEW

NAME: Lab vCenter

CODE: labvc

LOCATION: Denver, CO

Details

API URL: https://192.168.0.110/sdk

USERNAME: administrator@vsphere.local

PASSWORD:

VERSION: 6.7+ ▼

VDC: Home ▼

CLUSTER: All ▼

Morpheus allows vCenter clouds to be scoped to the **VDC** and **CLUSTER** or even the specific **RESOURCE POOL** if you choose. Once you’ve entered your URL and credentials, these dropdown menus will become populated.

The **RPC MODE** setting determines how Morpheus will connect to VMs and make configuration and scripting calls if **Morpheus Agent** is not installed. In a VMware environment we have the additional option to select VMware Tools if WinRM/SSH are not available.

Additionally, we can opt to **INVENTORY EXISTING INSTANCES** to begin polling VMs for statistics and rightsizing recommendations as well as **ENABLE HYPERVISOR CONSOLE** to use native vSphere console with port 443 connectivity between Morpheus and ESXi hosts.

To move on, expand the “Advanced Options” section.

Within the “Advanced Options” drawer are additional configurations to consider for your first cloud. Some of these won’t be usable until they reference additional configured integrations. Common settings to consider are **DOMAIN**, **STORAGE TYPE**, **APPLIANCE URL** (overrides the Morpheus URL for external systems), **GUIDANCE** (setting “Manual” will make recommendations for rightsizing), and **AGENT INSTALL MODE**.

▼ Advanced Options

☐ ENABLE DISK TYPE SELECTION☐ ENABLE NETWORK INTERFACE TYPE SELECTIONSTORAGE TYPE DOMAIN SCALE PRIORITY

Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL TIME ZONE

Once you're satisfied with your selections, click "NEXT"

We have now arrived at the "GROUP" tab. In this case, we will mark the radio button to "USE EXISTING" groups if you wish to use the group we configured earlier.

CREATE CLOUD [X]

CLOUD > CONFIGURE > **GROUP** > REVIEW

☒ USE EXISTING ☐ CREATE NEW

GROUP

PREVIOUS NEXT

Once you've selected the group, click "NEXT"

On the final tab of the "CREATE CLOUD" wizard, you'll confirm your selections and click "COMPLETE". The new cloud is now listed on the cloud detail page. After a short time, Morpheus will provide summary information and statistics on existing virtual machines, networks, and other resources available in the cloud.

Viewing Cloud Inventory

Now that we've integrated our first VMware cloud, we can stop for a moment to review what Morpheus gives us from the cloud detail page. We can see that Morpheus gives us estimated costs and cost histories, metrics on used resources, and also lists out resource counts in various categories including container hosts, hypervisors, and virtual machines. We can drill into these categories to see lists of resources in the various categories individual resources within them by clicking on the category tabs. We can link to the detail page for any specific resource by clicking on it from its resource category list.

Configuring Resource Pools

With our VMware cloud configured, Morpheus will automatically sync in available resource pools and data stores.

For resource pools, once Morpheus has had time to ingest them, then will be visible from the cloud detail page. Navigate to *Infrastructure > Clouds > (your VMware cloud) > RESOURCES tab*. In here, we are able to see and control access to the various resource pools that have been configured in vCenter. For example, we can restrict access to a specific resource pool within Morpheus completely by clicking on the "ACTIONS" button, then clicking "Edit". If we unmark the "ACTIVE" button and then click "SAVE CHANGES" we will see that the resource pool is now grayed out in the list. The resources contained in that pool will not be accessible for provisioning within Morpheus.

POOLS

<input type="text" value="Search"/>				
NAME	DESCRIPTION	VISIBILITY	DEFAULT	TENANT
Demo-vSAN		Private		morpheus
Demo		Private		morpheus
Pied Piper		Private		morpheus
Prod		Private		morpheus
Sandbox		Private		morpheus

Often our clients will want to make specific blocks of resources available to their own customers. This can be easily and conveniently controlled through the same "EDIT RESOURCE POOL" dialog box we were just working in. If we expand the "Group Access" drawer, we are able to give or remove access to each pool to any group we'd like. We can also choose to make some or all of our resource pools available to every group. Specific resource pools can also be defined as the default for each group if needed.

EDIT RESOURCE POOL

×

NAME
Demo-vSAN

☐ MOVE SERVERS

☐ ACTIVE

☐ DEFAULT

▼ Group Access

GROUP	ACCESS	DEFAULT
all	<input checked="" type="checkbox"/>	
AA	<input type="checkbox"/>	<input type="checkbox"/>
AA-DB	<input type="checkbox"/>	<input type="checkbox"/>

Additionally, we may choose to allow only certain service plans to be provisioned into a specific pool of resources. For example, perhaps a specific cluster is my SQL cluster and only specific services plans should be consumable within it. We can control that through this same dialog box.

Configuring Data Stores

To take a look at data stores, we'll move from the "RESOURCES" tab to the "DATA STORES" tab on our cloud detail page.

Morpheus gives the user similar control with data stores to what we saw with our resources pools earlier. Just like with resource pools, we can disable access within Morpheus completely by clicking on "ACTIONS" and then "Edit". If we unmark the "ACTIVE" checkbox and click "SAVE CHANGES", you will see that specific data store has been grayed out.

SUMMARY	CLUSTERS	HOSTS	VMS	CONTAINERS	LOAD BALANCERS	NETWORKS	DATA STORES	RESOURCES	POLICIES	WIKI
<div> <div>Search</div> <div>Q</div> <div>ACTIONS +</div> </div>										
<input type="checkbox"/>	NAME	TYPE CAPACITY ONLINE VISIBILITY TENANT								
<input type="checkbox"/>	ds-65-root	Vmfs	461.1GiB	Yes	Private	morpheus	ACTIONS ▼			
<div> <div>morpheuscamusara</div> <div>morpheus Morpheus Shaltar EMC IT SAN AM troublemaker Camset Anich rtuler KDE I neronha</div> </div>										

Just like with resource pools, we are also able to scope data stores to specific groups. This ensures that the members of each group are only able to consume the data stores they should have access to.

EDIT DATA STORE

×

NAME

ds-65-root

ACTIVE
☐

▼ Group Access

GROUP	ACCESS
all	<input checked="" type="checkbox"/>
AA	<input type="checkbox"/>
AA-DB	<input type="checkbox"/>

Configuring Network for Provisioning

When configuring networking, we can set global defaults by going to *Infrastructure > Network > NETWORKS* tab. Here we can add or configure networks from all clouds integrated into Morpheus. Depending on the number of clouds Morpheus has ingested, this list may be quite large and may also be paginated across a large number of pages. In such a case, it may be easier to view or configure networks from the specific cloud detail page so that networks from other clouds are not shown.

NETWORKS

NETWORKS

NETWORK GROUPS

ROUTERS

IP POOLS

DOMAINS

PROXIES

SECURITY GROUPS

INTEGRATIONS

Search

+

ADD

NAME	TYPE	CLOUD	CIDR	POOL	DHCP	VISIBILITY	TENANTS	
10.30.20.0	OracleVM Network	Morpheus Oracle VM	10.30.20.0/22	✓	Public	morpheus		ACTIONS ▼
172.31.0.0/20 (subnet-63dff13b)	Amazon Subnet	ah-only	172.31.0.0/20	✓	Private	morpheus		ACTIONS ▼
172.31.16.0/20 (subnet-22110ed6)	Amazon Subnet	ah-only	172.31.16.0/20	✓	Private	morpheus		ACTIONS ▼

Still in *Infrastructure > Network*, make note of the “INTEGRATIONS” tab. It’s here that we can set up any integrations that may be relevant, such as IPAM integrations. Generally speaking, when adding IPAM integrations, we simply need to name our new integration, give the API URL, and provide credentials. There’s more information in the [IPAM integration](#) section of Morpheus Docs.

ADD IPAM INTEGRATION ✕

NAME

☒ **ENABLED**

URL

https://x.x.x.x/wapi/v2.2.1

USERNAME

PASSWORD

THROTTLE RATE

0

ms

In *Infrastructure > Networking* we can also set up IP address pools from the IP Pools tab. These pools can be manually defined, known as a Morpheus-type IP pool, or they can come from any IPAM integrations you’ve configured. As instances are provisioned, Morpheus will assign IP addresses from the pool chosen during provisioning. When the instance is later dissolved, Morpheus will automatically release the IP address to be used by another instance when needed. When adding or editing a network, we can opt to scope the network to one of these configured IP address pools.

CREATE NETWORK POOL ✕

NAME

POOL TYPE

Morpheus

▼

IP Ranges

STARTING ADDRESS

ENDING ADDRESS

192.168.0.2

-

192.168.0.255

+

SAVE CHANGES

Since this guide is focused on working within a VMware cloud that we integrated at the start, we will take a look at our network configurations on the cloud detail page as well. Navigate to *Infrastructure > Clouds > (your VMware cloud) > NETWORKS tab*. Just as with resource pools and data stores, we have the ability to make certain networks inactive in Morpheus, or scope them to be usable only for certain groups or tenants.

SUMMARY	CLUSTERS	HOSTS	VMS	CONTAINERS	LOAD BALANCERS	NETWORKS	DATA STORES	RESOURCES	POLICIES	WIKI
---------	----------	-------	-----	------------	----------------	----------	-------------	-----------	----------	------

NETWORKS						
<input type="text" value="Search"/>						<input type="button" value="ACTIONS"/>
<input type="checkbox"/> NAME	TYPE	CIDR	POOL	DHCP VISIBILITY	TENANT	
<input type="checkbox"/> default	KVM Host Bridge	192.168.122.1/24		✓ Private	morpheus	ACTIONS ▾
<input type="checkbox"/> docker bridge	Docker Bridge			✓ Private	morpheus	ACTIONS ▾
<input type="checkbox"/> Garf-Network	Overlay	0.0.0.0/10	Vanishing-IP-Test	Private	morpheus	ACTIONS ▾

Prepping an Image

As we'll discuss and try out in the next section, Morpheus comes out of the box with a default set of blueprints that are relevant to many modern deployment scenarios. For the most part, these are base operating system images with a few additional adjustments. However, in many on-premise deployments, there are often custom image and networking requirements. We will work with images included in Morpheus by default in this guide but it's important to discuss how to prep custom images as well.

Creating a Windows Image

The following versions of Windows Server are supported:

- 2008 R2
- 2012
- 2012 R2
- 2016
- 2019

To start, create a new Windows machine in vCenter using a base version of your selected Windows build.

Note: It's recommended to make the VMDK drive as small as possible for your purposes as this generally speeds cloning and deploy times. Morpheus provisioning and post-deploy scripts allow to to expand the drive to any size that you need.

Once the machine is created, ensure VMtools is installed on the operating system. Then, apply all updates and service packs. Next, configure WinRM and open the firewall:

```
winrm quickconfig
```

Note: WinRM configuration is optional if using VMtools RPC mode for agent install and Morpheus Agent for guest exec.

Next, we'll install .NET 4.5 or higher. Ensure Windows Firewall will allow WinRM connections and shut down the virtual instance. Finally, convert it to a template.

Note: Morpheus will Sysprep images based on the "Force Guest Customizations" flag under VM settings when using DHCP. If this flag is enabled or if using static IP addresses or IP pools when provisioning, ensure a Sysprep has not been performed. In such cases, guest customization will always be performed and a Sysprep will be triggered.

Creating a CentOS/RHEL Image

Create a new machine in vCenter and install a base version of your preferred Linux distro.

Note: If you are using cloud-init as part of your image, you will need to ensure your virtual machine has a cdrom.

Before installing the operating system, set up a single ext or xfs partition without a swap disk. Next, install the distro applying any updates to the operating system or security updates. Once the operating system is running and updated, install the following:

```
yum install cloud-init
yum install cloud-utils-growpart
yum install open-vm-tools
yum install git
yum install epel-release
```

Set selinux to permissive as the enforced setting can cause problems with cloud-init:

```
sudo vi /etc/selinux/config
```

Cloud-Init

We'll get started by installing cloud-init using the following command:

```
yum -y install epel-release
yum -y install git wget ntp curl cloud-init dracut-modules-growroot
rpm -qa kernel | sed 's/^kernel-//' | xargs -I {} dracut -f /boot/initramfs-{}.img {}
```

Note: The above command will install some core dependencies for cloud-init and automation later as you work with your provisioned instances. For example, we install Git here as it is used for Ansible automation. If you had no plans to use Ansible, this installation could be skipped. The dracut-modules-growroot is responsible for resizing the root partition upon initial boot which was potentially adjusted during provisioning.

One key benefit of using cloud-init is that we don't have to lock credentials into the blueprint. We recommend configuring a default cloud-init user that will get created automatically when the VM is booted by cloud-init. We can define that default user in *Administration > Provisioning > Cloud-Init*.

Network Interfaces

As of CentOS 7, network interface naming conventions have changed. You can check this by running *ifconfig* and noting that the primary network interface has some value similar to "ens2344". The naming is dynamic and typically set based on hardware ID. We don't want this to fluctuate when provisioning this blueprint in our VMware environments. To accomplish this end, we will rename the interface back to "eth0" using the steps below.

First, adjust the bootloader to disable interface naming:

```
sed -i -e 's/quiet/quiet net.ifnames=0 biosdevname=0/' /etc/default/grub
grub2-mkconfig -o /boot/grub2/grub.cfg
```

The next step is to adjust network scripts in CentOS. Start by confirming the presence of a file called */etc/sysconfig/network-scripts/ifcfg-eth0*. Once confirmed, run the following script:

```
export iface_file=$(basename "$(find /etc/sysconfig/network-scripts/ -name 'ifcfg*' -
↪not -name 'ifcfg-lo' | head -n 1)")
export iface_name=${iface_file:6}
```

(continues on next page)

(continued from previous page)

```
echo $iface_file
echo $iface_name
sudo mv /etc/sysconfig/network-scripts/$iface_file /etc/sysconfig/network-scripts/
↪ifcfg-eth0
sudo sed -i -e "s/$iface_name/eth0/" /etc/sysconfig/network-scripts/ifcfg-eth0
sudo bash -c 'echo NM_CONTROLLED="\no\" >> /etc/sysconfig/network-scripts/ifcfg-eth0'
```

This script tries to confirm there is a new *ifcfg-eth0* config created to replace the old config file. Confirm this config exists after running and if not you will have to build your own:

```
TYPE=Ethernet
DEVICE=eth0
NAME=eth0
ONBOOT=yes
NM_CONTROLLED="no"
BOOTPROTO="dhcp"
DEFROUTE=yes
```

For more on CentOS/RHEL image prep, including additional configurations for specific scenarios, take a look at the [VMware image prep](#) page in Morpheus Docs.

Creating an Ubuntu Image

Create a new machine in vCenter and install a base version of your preferred Linux distro.

Note: If you are using cloud-init as part of your image, you will need to ensure your virtual machine has a cdrom.

Before installing the operating system, set up a single ext partition without a swap disk. Install the distro and apply any operating system and security updates. Ensure you’ve set a root password.

Install cloud-init and cloud-utils-growpart:

```
sudo apt install cloud-init
sudo apt install cloud-utils
```

Install desired hypervisor drivers, such as Virto or Open-VM Tools

Install Git:

```
sudo apt install git
```

Since Debian 9 includes network manager, ensure this is disabled. You can do this by editing the configuration file at */etc/NetworkManager/NetworkManager.conf*. Within that file, update the “managed” flag to false:

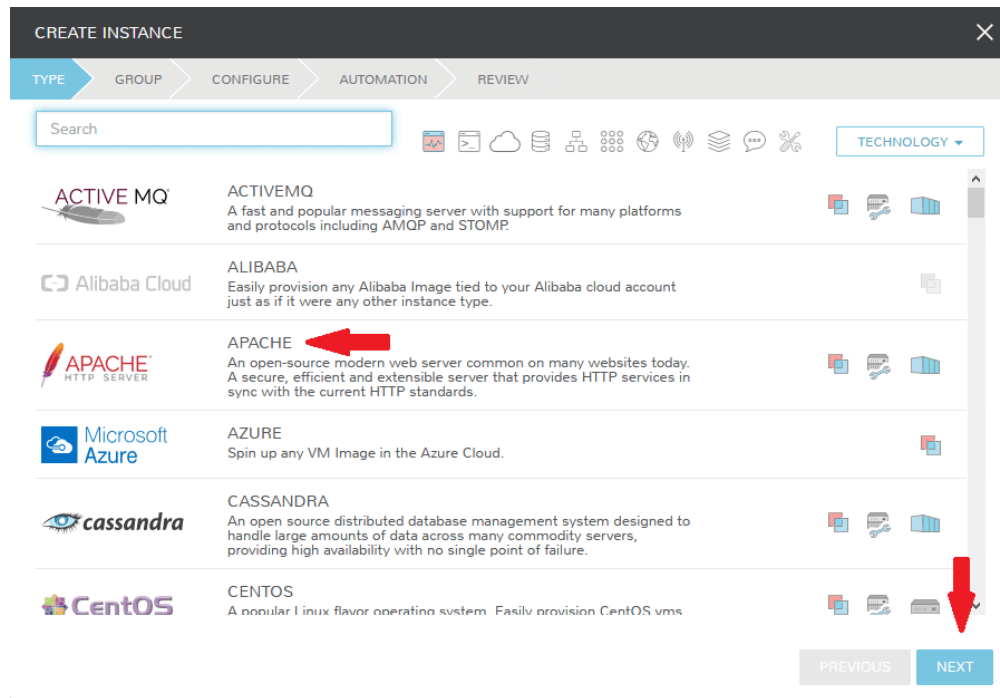
```
managed=false
```

We also recommend setting the network adapter to “eth0”. This process is described above in the “Network Interfaces” section of the CentOS image prep guide above.

Provisioning Your First Instance

At this point, we are ready to provision our first image. As a first instance, we'll provision an Apache web server to our vCenter cloud.

Navigate to *Provisioning > Instances*. If any instances are currently provisioned, we will see them listed here. To start a new instance we click the “+ADD” button to pop the “CREATE INSTANCE” wizard. We'll scroll down to and select the Apache instance type and click “NEXT”.



First, we'll specify the group to provision into which determines the clouds available. If you've followed this guide to this point, you should at least have a group that houses all of your clouds which you can select here. This will allow us to select the vCenter cloud from the “CLOUD” dropdown menu. Provide a unique name to this instance and then click “NEXT”

From the “CONFIGURE” tab, we're presented with a number of options. The options are cloud and layout-specific, more generalized information on creating instances and available options is [here](#). For our purposes, we'll select the following options:

- **LAYOUT:** Includes options such as the base OS, custom layouts will also be here when available
- **PLAN:** Select the resource plan for your instance. Some plans have minimum resource limits, Morpheus will only show plans at or above these limits. User-defined plans can also be created in *Administration > Plans & Pricing*.
- **VOLUMES and DATASTORES:** The minimum disk space is set by the plan, this value may be locked if you've selected a custom plan that defines the volume size
- **NETWORKS:** Select a network, note that IP pools must be linked with the networks defined in VMware in order to assign static IP addresses

Under the “User Config” drawer, mark the box to “CREATE YOUR USER”. Click “NEXT”.

CREATE INSTANCE

TYPE > GROUP > **CONFIGURE** > AUTOMATION > REVIEW

Configuration Options

LAYOUT ESXi Apache on Ubuntu 14.04

PLAN 1 Core, 512MB Memory
Cores: 1 Memory: 512 MB

VOLUMES root 10 GB Auto - Datastore

NETWORKS VM Network DHCP

▼ User Config

☒ **CREATE YOUR USER**

USER GROUP Select

► Network Options

► Advanced Options

► Metadata

► Environment

PREVIOUS NEXT

Note: “CREATE YOUR USER” will seed a user account into the VM with credentials set in your Morpheus user account settings. If you’ve not yet defined these credentials, you can do so by clicking on your username in the upper-right corner of the application window and selecting “USER SETTINGS”.

For now, we’ll simply click “NEXT” to move through the “AUTOMATION” tab but feel free to stop and take a look at the available selections here. There is more information later in this guide on automation and even more beyond that in the rest of Morpheus docs.

Review the settings for your first instance and click “COMPLETE”.

CREATE INSTANCE

TYPE

GROUP

CONFIGURE

AUTOMATION

REVIEW

apachetest

admin (esxtank)

Summary

INSTANCE OPTIONS

NAME: apachetest

GROUP: admin

CLOUD: esxtank

TYPE: Apache

PLAN: 1 Core, 512MB Memory

Cores: 1 Memory: 512MB

VERSION: 2.4

LAYOUT: ESXi Apache on Ubuntu 14.04

VOLUMES

ROOT: 10 GB Auto - Datastore SCSI 0

NETWORKS

VM NETWORK: DHCP

Options

PREVIOUS

COMPLETE

We are now dropped back onto the instances list page. We can see a new entry in the list at this point with a status indicator that the new machine is being launched (rocket icon in the status field). We can double click on the instance in the list to move to the instance detail page. For now we will see a progress bar indicating that the instance is being created and is starting up. The exact amount of time this process will take depends on your environment and selections made when provisioning the instance. Initially, Morpheus will guess as to how long this will take and the progress bar may not be accurate. Over time, Morpheus will learn how long these processes take and progress bar accuracy will improve. For more detailed information on the status of various provisioning processes, we can scroll down and select the “HISTORY” tab. The “STATUS” icon will change from the blue rocket to a green play button when the instance is fully ready. Furthermore, we can click on the hyperlinked IP address in the “VMS” section of this page to view a default page in a web browser to confirm success.

INFO

Group: admin

Created By: Nick Celebic

Cores: 1

Source Image: Morpheus Apache 2.4 on Ubuntu 14.04.3

Cloud: esxtank

Layout: ESXi Apache on Ubuntu 14.04

Memory: 512.0MiB

Date Created: 11/13/2019 01:44 PM

Version: 2.4

Total Storage: 10.0GiB

VMS

	STATUS	NAME	TYPE	CLOUD	LOCATION	COMPUTE	MEMORY	STORAGE	ACTIONS
<input type="checkbox"/>		apachetest	Apache 2.4	esxtank		<div>0</div>	<div>0</div>	<div>0</div>	

SUMMARY

WIKI

DEPLOY

STORAGE

NETWORK

LOGS

BACKUPS

ENVIRONMENT

SCALE

HISTORY


	NAME	DESCRIPTION	CREATED BY	START DATE	ETA/DURATION	STATUS	ERROR
	apachetest	Provision	Nick Celebic	11/13/2019 01:44 PM	00:18:44	<div></div>	
		Prepare Resources		11/13/2019 01:44 PM	00:18:44	Complete	
		Prepare Image		11/13/2019 01:44 PM	00:00:01	Complete	
		Configure Instance		11/13/2019 01:44 PM	260ms	Complete	
		Deploy Instance		11/13/2019 01:44 PM	00:18:44	<div></div>	

Creating Your First Library Item

In the prior section, we manually provisioned our first instance. However, Morpheus allows you to build a catalog of custom provisionable items to simplify and speed provisioning in the future. In this section, we'll build a catalog item and show how that can translate into quick instance provisioning after configuration.

Note: Before starting this process, it's important to decide which virtual image you plan to use. If you're not using a Morpheus-provided image, you'll want to ensure it's uploaded. You will not be able to complete this section without selecting an available image. In this example we will use Morpheus Redis 3.0 on Ubuntu 14.04.3 v2.

Navigate to *Provisioning > Library > NODE TYPES* and click "+ADD".



In this example, I am going to set the following options in the “NEW NODE TYPE” wizard:

- **NAME**
- **SHORT NAME**
- **VERSION:** 1 (In this particular case, the version is not important)
- **TECHNOLOGY:** VMware
- **VM IMAGE:** Morpheus Redis 3.0 on Ubuntu 14.04.3 v2

Note: Within the “VMware VM Options” section you should add anything that will always be used for this node, regardless of the specific deployment details. This can include LDAP Authentication, bash scripts that should run on installation, among other things.

NAME

customnodetype

SHORT NAME

customnodetype

The short name is a name with no spaces used for display in your container list.

VERSION

1

TECHNOLOGY

VMware

ENVIRONMENT VARIABLES

Name

Value

⚙️

+

VMware VM Options

VM IMAGE

Morpheus Redis 3.0 on Ubuntu 14.04.3 v2

LOG FOLDER

CONFIG FOLDER

DEPLOY FOLDER

(Optional) If using deployment services, this mount point will be replaced with the contents of said deployments.

EXTRA OPTIONS

Name

Value

🗑️

+

SERVICE PORTS

port

⬆️⬆️⬆️

name

No LB

⌵

+

SCRIPTS

Search


FILE TEMPLATES

Search

▶ Advanced Options

SAVE CHANGES

With the new node created, we'll now add a new instance type which will be accessible from the provisioning wizard once created. Move from the "NODE TYPES" tab to the "INSTANCE TYPES" tab and click "+ADD".



In the “NEW INSTANCE TYPE” wizard, I’ll simply enter a **NAME** and **CODE** value. Click “SAVE CHANGES”.

Now that we’ve created a new instance type, access it by clicking on the name in the list of custom instances you’ve created. In my case, I’ve given the name “NewInstanceType”.

LIBRARY

INSTANCE TYPES LAYOUTS NODE TYPES OPTION TYPES OPTION LISTS FILE TEMPLATES SCRIPTS SPEC TEMPLATES CLUSTER LAYOUTS

Search SELECT TECHNOLOGY SELECT CATEGORY + ADD

	NAME	TECHNOLOGY	CATEGORY	FEATURED	
	NewInstanceType		Web		ACTIONS ▾
	ActiveMQ	Mixed	Messaging		ACTIONS ▾
	Alibaba	Alibaba	Cloud		ACTIONS ▾
	Amazon Api		Apps		ACTIONS ▾
	AmazonMQ		Messaging		ACTIONS ▾

Once we’ve opened the new instance type, by default, we should be on the “LAYOUTS” tab. Click “+ADD LAYOUT”.

I’ve set the following fields on my example layout:

- **NAME**
- **VERSION:** This is the version number of the layout itself, which is labeled 1.0 in the example
- **TECHNOLOGY:** VMware

- **Nodes:** Select the node we created earlier, if desired you can specify multiple nodes

Click “SAVE CHANGES”.

NEW LAYOUT

NAME

customlayout

VERSION

1.0

DESCRIPTION

☒ CREATABLE

TECHNOLOGY

VMware

MINIMUM MEMORY

0

MB

This will override any memory requirement set on the virtual image

WORKFLOW

☐ SUPPORTS CONVERT TO MANAGED

ENVIRONMENT VARIABLES

Name

Value

+

Option Types

Search option types

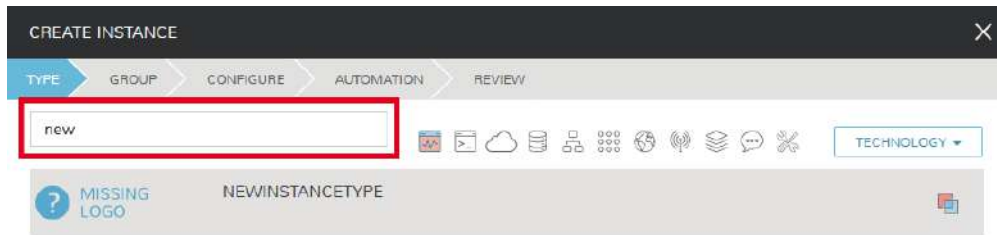
Nodes

customnodetypevm

customnodetypevm (1)

SAVE CHANGES

At this point we’ve completed the setup work and can now provision the instance we’ve created to our specifications. Navigate to *Provisioning > Instances* and click “+ADD”. From the search bar we can search for the new instance type we’ve created. In the example case, we called it “newinstancetype”. Click “NEXT”.



CREATE INSTANCE

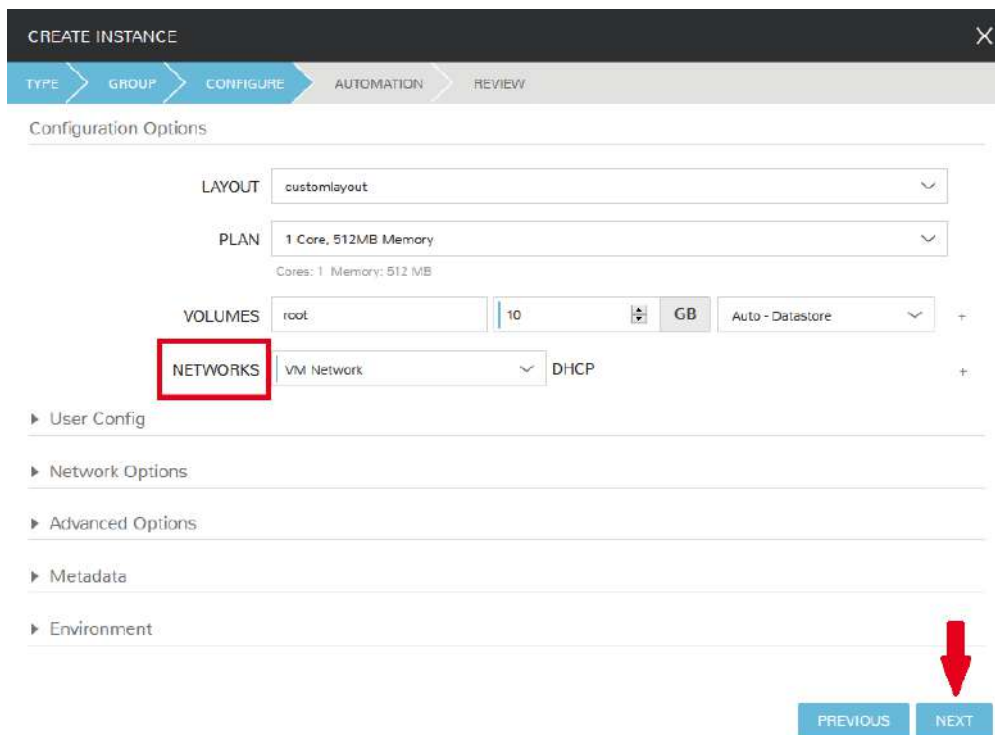
TYPE GROUP CONFIGURE AUTOMATION REVIEW

new

MISSING LOGO NEWINSTANCETYPE



As before, we can select a group and cloud to provision this new instance. Click “NEXT”. On the “CONFIGURE” tab, make note that the layout and plan are already selected because they were configured as part of creating the new instance type. Select a network and click “NEXT”. Once again we will also click “NEXT” through the “AUTOMATION” tab. Finally, click “COMPLETE”.



CREATE INSTANCE

TYPE GROUP CONFIGURE AUTOMATION REVIEW

Configuration Options

LAYOUT customlayout

PLAN 1 Core, 512MB Memory

Cores: 1 Memory: 512 MB

VOLUMES root 10 GB Auto - Datastore

NETWORKS VM Network DHCP

User Config

Network Options

Advanced Options

Metadata

Environment

PREVIOUS NEXT

As before when we manually provisioned an instance, Morpheus will now begin to spin up the new VM. How long this will take depends on your environment but Morpheus will predict how long this process will take and represent that on a progress bar. Over time, Morpheus begins to learn how long these processes take and becomes more accurate

in predicting spin-up time.

Once the provisioning process has completed, open the instance detail page in Morpheus and click on the “CONSOLE” tab. You’ll be logged in with your user account and are then able to confirm the machine is ready and available.

The screenshot shows the Morpheus instance detail page for an instance named 'newinstance'. At the top, there are buttons for 'EDIT', 'ACTIONS', and 'DELETE'. Below this, a row of metrics includes STATUS (green play button), HEALTH (blue question mark), LAST BACKUP (grey minus), AVAILABILITY (100.00%), RESPONSE TIME (N/A), MAX CPU (1%), MEMORY (39%), and STORAGE (27%). The 'INFO' section displays details: Group: admin, Created By: Nick Celebic, Cloud: esxtank, Date Created: 11/14/2019 12:21 PM, Cores: 1, Layout: customlayout, Memory: 512.0MB, Version: 1.0, Source Image: Morpheus Redis 3.0 on Ubuntu 14.04.3 v2, Provision Time: 14 minutes 29 seconds, and Total Storage: 10.0GiB. The 'VMS' section contains a table with columns: STATUS, NAME, TYPE, CLOUD, LOCATION, COMPUTE, MEMORY, STORAGE, and ACTIONS. The table has one row for 'newinstance' with a green play button icon. Below the table, there are tabs for SUMMARY, WIKI, STORAGE, NETWORK, LOGS, BACKUPS, ENVIRONMENT, HISTORY, and CONSOLE. The 'CONSOLE' tab is selected, showing a terminal window with the text: (Connected), Paste Text Here, SEND CTRL+ALT+DELETE, CUSTOMNODETYPE_14 - ..., and a terminal prompt: Welcome to Ubuntu 14.04.3 LTS (GNU/Linux 3.19.0-80-generic x86_64) * Documentation: https://help.ubuntu.com/ Last login: Thu Nov 14 17:55:47 2019 from 192.168.88.61 ncelebic@newinstance:~\$

Automation and Configuration Management

Morpheus automation is composed of Tasks and Workflows. A task could be a script added directly, scripts or blueprints pulled from the Morpheus Library, playbooks, recipes, or a number of other things. The complete list of task types can be found in the [Automation section](#) of Morpheus docs. Tasks can be executed individually but they are often combined into workflows. We can opt to run a workflow at provision time or they can be executed on existing instances through the Actions menu.

In this guide we will set up an Ansible integration, create a task, add the task to a workflow, and run the workflow against a new and existing instance. If you’ve worked through this guide to this point, you should already have an Apache instance running. If you don’t yet have that, provision one before continuing with this guide and ensure it’s reachable on port 80.

The screenshot shows the Morpheus Administration page. The top navigation bar includes Operations, Provisioning, Infrastructure, Backups, Logs, Monitoring, Tools, and Administration. The 'Administration' tab is selected. Below the navigation bar, there are icons for Tenants, Plans & Pricing, Roles, Users, Integrations, Policies, Provisioning, Monitoring, Backups, Logs, and Settings. The 'Integrations' section is active, showing a table with columns: NAME and a '+ NEW INTEGRATION' button. A dropdown menu is open, showing a list of integration options: automation, Chef, Puppet, Ansible, Ansible Tower, vRealize Orchestrator, Salt Master, dns, and Microsoft DNS. Red arrows point to the '+ NEW INTEGRATION' button and the 'Ansible' option in the dropdown.

We'll first set up the Ansible integration, you can integrate with the sample repository referenced here or integrate with your own. Go to 'Administration > Integrations'. Click "+NEW INTEGRATION" and select Ansible from the dropdown menu. Fill in the following details:

- **NAME**
- **ANSIBLE GIT URL:** <https://github.com/ncelebic/morpheus-ansible-example>, or enter the URL for your own Ansible git repository
- **PLAYBOOKS PATH**
- **ROLES PATH**
- Mark the box to "USE MORPHEUS AGENT COMMAND BUS"

Note: If your git repository requires authentication, you should create a keypair and use the following URL format: [git@github.com:ncelebic/morpheus-ansible-example.git](https://github.com:ncelebic/morpheus-ansible-example.git).

The screenshot shows the 'NEW ANSIBLE INTEGRATION' form. The following fields and options are highlighted with red boxes:

- NAME:** NewAnsibleIntegration
- ENABLED:** ☒ ENABLED
- ANSIBLE GIT URL:** <https://github.com/ncelebic/morpheus-ansible-example>
- KEY PAIR:** (Dropdown menu)
- PLAYBOOKS PATH:** /
- ROLES PATH:** /roles
- GROUP VARIABLES PATH:** (Empty field)
- HOST VARIABLES PATH:** (Empty field)
- USE ANSIBLE GALAXY:** ☐
- ENABLE VERBOSE LOGGING:** ☐
- USE MORPHEUS AGENT COMMAND BUS:** ☒

A red arrow points down to the **SAVE CHANGES** button at the bottom right of the form.

Click "SAVE CHANGES". You'll now see our new Ansible integration listed among any other configured integrations.

tions. If we click on this new integration to view detail, a green checkmark icon indicates the git repository has been fully synced.

With the Ansible integration set up, we can now create a task that includes our playbook. Go to *Provisioning > Automation*, click “+ADD”. We’ll first set our “TYPE” value to Ansible Playbook so that the correct set of fields appear in the “NEW TASK” wizard. Set the following options:

- **NAME**
- **ANSIBLE REPO:** Here we will choose the Ansible integration that we just created
- **PLAYBOOK:** In our example case, enter ‘playbook.yml’

The screenshot shows the 'NEW TASK' form. The fields are: NAME (New Ansible Task), CODE, TYPE (Ansible Playbook), ANSIBLE REPO (NewAnsibleIntegration), GIT REF, PLAYBOOK (playbook.yml), TAGS, SKIP TAGS, and COMMAND OPTIONS. Below these are 'Target Options' and 'EXECUTE TARGET' (Resource). A red arrow points to the 'SAVE CHANGES' button at the bottom right.

Click “SAVE CHANGES” to save our new task. We can test the new task on our Apache VM now by going to *Provisioning > Instances* and clicking into our VM. From the “ACTIONS” menu select “Run Task”. From the “TASK” dropdown menu, select the task we just added and click “EXECUTE”.

The screenshot shows the 'EXECUTE TASK?' dialog box. It asks 'Are you sure you would like to perform this task operation?'. There is a 'TASK' dropdown menu showing 'New Ansible Task'. A red arrow points to the 'EXECUTE' button.

To see the progress of the task, click on the “HISTORY” tab and click on the (i) button to the right of each entry in the list. In this case, we can also see the results of the task by clicking on the link in the “LOCATION” column of the “VMS” section.

Now that our task is created, we can put it into a workflow. Back in *Provisioning > Automation* we will click on the “WORKFLOWS” tab. Click “+ADD” and select Provisioning Workflow. We’ll give the new workflow a name and expand the Post Provision section. As we begin to type in the name of the task we’ve created, it should appear as a selection. Click “SAVE CHANGES”.

NEW WORKFLOW ✕

NAME

DESCRIPTION

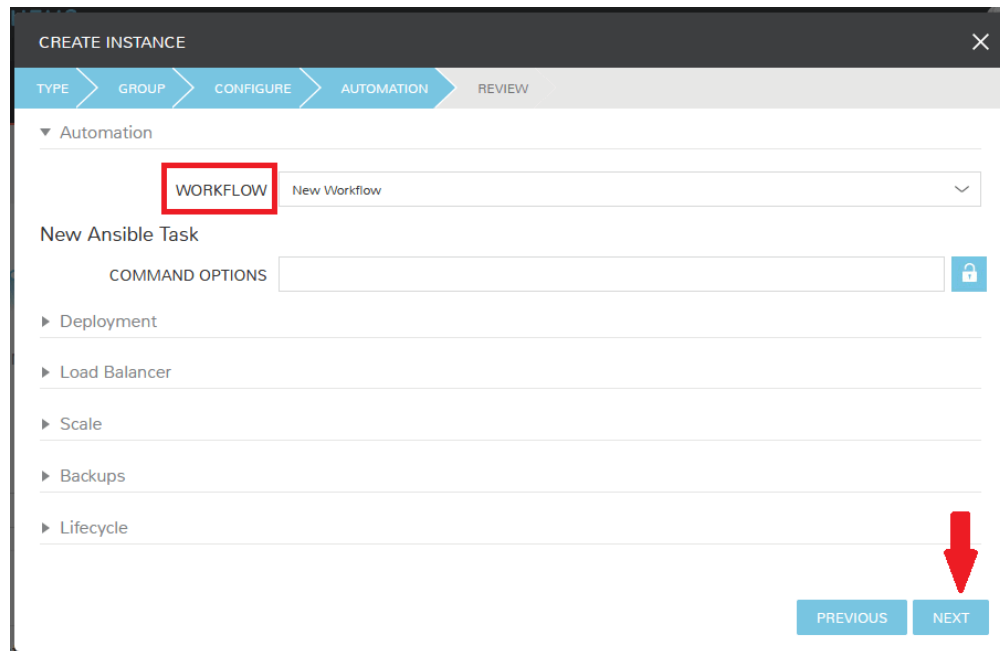
PLATFORM ▼

Tasks

- ▶ Pre Provision
- ▼ Provision
-
- ▼ Post Provision
 -
 - ▶ **New Ansible Task**
- ▶ Stop Service
- ▶ Pre Deploy
- ▶ Deploy
- ▶ Reconfigure
- ▶ Teardown

SAVE CHANGES

Now that we have a workflow, return to *Provisioning > Instances* and begin to provision another Apache instance. More detailed instructions on provisioning a new Apache instance are included earlier in this guide if needed. Now, when you reach the “AUTOMATION” section of the “CREATE INSTANCE” wizard, we have a workflow to select. From the “WORKFLOW” dropdown menu, select the workflow we just created and complete provisioning of the new instance.



As the instance is provisioning, we can go to the “HISTORY” tab and see Morpheus executing the tasks that were contained in our workflow.

This is just one example of using Morpheus to automate the process of configuring and instance to your needs. There are a number of other automation types that can be built into your workflows as well. For further information, take a look at the [automation integrations](#) guide in Morpheus docs.

Conclusion

At this point you should be up and running in Morpheus, ready to consume VMware. This guide only scratches the surface, there is a lot more to see and do in Morpheus. Take a look at the rest of [Morpheus Docs](#) for more information on supported integrations and other things possible.

VMware Fusion

Add a VMware Fusion Cloud

1. Navigate to Infrastructure -> Clouds
2. Select + *CREATE CLOUD*, select VMware Fusion, and then click *Next*.
3. Enter the following into the Create Cloud modal:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

VMWARE FUSION HOST IP or URL of VMware Fusion Host

WORKING PATH Existing folder Morpheus will write to on Host

USERNAME Host Username

PASSWORD Host Password

BRIDGE NAME Will auto-populate upon successful authentication with the Fusion Host (E.X. ‘EN0: ETHERNET’)

4. The Cloud can now be added to a Group or configured with additional Advanced options.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM’s if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Xen Server

Add a Xen Server Cloud

1. Navigate to *Infrastructure -> Clouds*
2. Select + *CREATE CLOUD*, select Xen, and then click *Next*.
3. Enter the following into the Create Cloud modal:

NAME Name of the Cloud in Morpheus

CODE Unique code used for api/cli, automation and policies.

LOCATION Description field for adding notes on the cloud, such as location.

VISIBILITY For setting cloud permissions in a multi-tenant environment. Not applicable in single tenant environments.

TENANT If Visibility is set to Private, select the Tenant the Cloud resources will assigned to.

ENABLED When disabled, automatic Cloud sync is paused and the Cloud will not be selectable for provisioning.

AUTOMATICALLY POWER ON VMS When enabled, Morpheus will maintain the expected power state of managed VMs. Morpheus will power on any managed VMs in the Cloud that have been shut down for unknown reasons (not powered off by Morpheus) to ensure availability of services.

Note: When “AUTOMATICALLY POWER ON VMS” is enabled, the power state of managed VMs should be maintained in Morpheus. This setting is not applicable to discovered/unmanaged resources.

API URL IP or URL of Xen Host. ex: *xenserver.domain.com*

CUSTOM PORT Port for non standard xen server clouds

USERNAME Xen Host Username

PASSWORD Xen Host Password

Inventory Existing Instances If enabled, existing Virtual Machines will be inventoried and appear as unmanaged Virtual Machines in Morpheus .

4. The Cloud can now be added to a Group or configured with additional Advanced options.

Advanced Options

DOMAIN Specify a default domain for instances provisioned to this Cloud.

SCALE PRIORITY Only affects Docker Provisioning. Specifies the priority with which an instance will scale into the cloud. A lower priority number means this cloud integration will take scale precedence over other cloud integrations in the group.

APPLIANCE URL Alternate Appliance url for scenarios when the default Appliance URL (configured in *admin -> settings*) is not reachable or resolvable for Instances provisioned in this cloud. The Appliance URL is used for Agent install and reporting.

TIME ZONE Configures the time zone on provisioned VM's if necessary.

DATACENTER ID Used for differentiating pricing among multiple datacenters. Leave blank unless prices are properly configured.

NETWORK MODE Unmanaged or select a Network Integration (NSX, ACI etc)

LOCAL FIREWALL On or Off. Enable to managed Host and VM firewall/IP Table rules (linux only)

SECURITY SERVER Security Server setting is for Security Service Integrations such as ACI

TRUST PROVIDER Select Internal (Morpheus) or an existing Trust Provider Integration

STORAGE MODE Single Disk, LVM or Clustered

BACKUP PROVIDER Select Internal Backups (Morpheus) or a Backup Integration

REPLICATION PROVIDER Sets the default Replication Provider for the Cloud. Select an existing Replication Provider Integration

GUIDANCE Enable Guidance recommendations on cloud resources.

COSTING Enable for Morpheus to sync Costing data from the Cloud provider, when available. If your organization utilizes reserved instances and you want to pull in related pricing data, select *Costing and Reservations*. If this is not relevant, select *Costing* to save money on additional calls to the AWS Cost Explorer API or similar service for other clouds.

DNS INTEGRATION Records for instances provisioned in this cloud will be added to selected DNS integration.

SERVICE REGISTRY Services for instances provisioned in this cloud will be added to selected Service Registry integration.

CONFIG MANAGEMENT Select a Chef, Salt, Ansible or Puppet integration to be used with this Cloud.

CMDB Select CMDB Integration to automatically update selected CMDB.

CHANGE MANAGEMENT Select an existing Change Management Integration to set on the Cloud. ex: Cherwell

AGENT INSTALL MODE

- **SSH / WINRM:** Morpheus will use SSH or WINRM for Agent install.
- **Cloud Init / Unattend (when available):** (DEFAULT) Morpheus will utilize Cloud-Init or Cloudbase-Init for agent install when provisioning images with Cloud-Init/Cloudbase-Init installed. Morpheus will fall back on SSH or WINRM if cloud-init is not installed on the provisioned image. Morpheus will also add Agent installation to Windows unattend.xml data when performing Guest Customizations or utilizing syspreped images.

API PROXY Set a proxy for outbound communication from the Morpheus Appliance to the Cloud endpoints. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

INSTALL AGENT Enable to have Agent Installation on by default for all provisioning into this Cloud. Disable for Agent Installation to be off by default for all provisioning into this Cloud.

Provisioning Options

PROXY Set a proxy for inbound communication from Instances to the Morpheus Appliance. Proxies can be added in the *Infrastructure -> Networks -> Proxies* tab.

Bypass Proxy for Appliance URL Enable to bypass proxy settings (if added) for Morpheus Agent communication to the Appliance URL.

NO PROXY Include a list of IP addresses or name servers to exclude from proxy traversal

USER DATA (LINUX) Add cloud-init user data. Morpheus 4.1.0 and earlier assumes bash syntax. Morpheus 4.1.1 and later supports all User Data formats. Refer to <https://cloudinit.readthedocs.io/en/latest/topics/format.html> for more information.

Creating a CentOS 7 Morpheus Image

Overview

Morpheus comes out of the box with a default set of blueprints for use in many modern deployment scenarios. These consist mostly of base operating system images with a few additional adjustments. These adjustments typically include the addition of cloud-init (which is highly recommended to be used in most environments, but not mandatory). However, in many on-premise deployments there are custom image requirements as well as networking requirements. This guide will go over how to create a base CentOS 7 Image for use within Morpheus.

Creating a CentOS 7 Morpheus VMware Image

VMWare

When running in VMWare it is highly recommended that VMWare Tools be installed. Without it, Morpheus will have difficulty assessing the host ip address and performing some additional automation tasks for the operating system.

Cloud-Init

To get started with a base CentOS image we first install cloud-init. This is a relatively simple process using yum:

```
yum -y install epel-release
yum -y install git wget ntp curl cloud-init dracut-modules-growroot
rpm -qa kernel | sed 's/^kernel-//' | xargs -I {} dracut -f /boot/initramfs-{}.img {}
```

There are two parts to this yum installation. We are first ensuring some core dependencies are installed for automation as well as cloud-init. git for example is installed for use by ansible playbook automation down the line and is therefore optional if not using ansible. The dracut-modules-growroot is responsible for resizing the root partition upon first boot to match the virtual disk size that was potentially adjusted during provisioning.

A great benefit to using cloud-init is credentials don't have to be locked into the blueprint. It is advisable, within Morpheus, to configure the default cloud-init user that gets created when the vm boots automatically by cloud-init. This is located in the *Administration -> Provisioning -> Cloud-Init Settings* section.

Network Interfaces

A slightly annoying change with CentOS 7 is that the network interfaces have changed naming convention. You may notice when running ifconfig that the primary network interface is set to something like ens2344 or some other random number. This naming is dynamic typically by hardware id and we don't want this to fluctuate when provisioning the blueprint in various VMWare environments. Fortunately, there is a way to turn this functionality off and restore the interface back to eth0.

Firstly we need to adjust our bootloader to disable interface naming like this.

```
sed -i -e 's/quiet/quiet net.ifnames=0 biosdevname=0/' /etc/default/grub
grub2-mkconfig -o /boot/grub2/grub.cfg
```

The above command adds a few arguments to the kernel args list (namely `net.ifnames=0` and `biosdevname=0`). It may be useful to view the `/etc/default/grub` file and ensure these settings were indeed applied.

The next step is to adjust the network-scripts in CentOS. we need to ensure we have a file called `/etc/sysconfig/network-scripts/ifcfg-eth0`

Below is a script that we run on our packer builds to prepare the machines network configuration files.

```
export iface_file=$(basename "$(find /etc/sysconfig/network-scripts/ -name 'ifcfg*' -
↪not -name 'ifcfg-lo' | head -n 1)")
export iface_name=${iface_file:6}
echo $iface_file
echo $iface_name
sudo mv /etc/sysconfig/network-scripts/$iface_file /etc/sysconfig/network-scripts/
↪ifcfg-eth0
sudo sed -i -e "s/$iface_name/eth0/" /etc/sysconfig/network-scripts/ifcfg-eth0
sudo bash -c 'echo NM_CONTROLLED="no" >> /etc/sysconfig/network-scripts/ifcfg-eth0'
```

This script tries to ensure there is a new ifcfg-eth0 config created to replace the old ens config file. Please do verify this config exists after running. If it does not you will have to be sure to build one on your own.

```
TYPE=Ethernet
DEVICE=eth0
NAME=eth0
ONBOOT=yes
NM_CONTROLLED="no"
BOOTPROTO="dhcp"
DEFROUTE=yes
```

Gotyas

SELinux can cause issues with cloud-init when in enforced mode. It may be advisable to set this to permissive unless it is mandatory within your organization to use an enforced SELinux configuration. If that is the case please see the documentation for the cloud_init_t security policies.

Network Manager will also prevent the required restart of the Network Service when assigning static IP's. Disable Network Manager when possible or Static IP assignment may not work until the Network Service is restarted manually.

A Note on Proxies

Proxy configurations are known to vary in some organizations and makes building a base blueprint a little more difficult. In order to fully configure proxies a few environment variables must be set in the */etc/environment* file (This can be done automatically in a default user-data script for cloud-init as well in edit cloud).

```
http_proxy="http://myproxyaddress:8080"
https_proxy="http://myproxyaddress:8080"
ftp_proxy="http://myproxyaddress:8080"
no_proxy=127.0.0.1,localhost,applianceUrl
https_no_proxy=127.0.0.1,localhost,applianceUrl
```

Important: It is very important to properly set the no_proxy list (applianceUrl) should be replaced with the actual appliance url. In future releases, morpheus plans to automatically take care of this.

Note: If using cloud-init agent install mode these settings need to be set in the custom Cloud-Init User data section of "Edit Cloud" or "Edit Virtual Image"

Important: If using this virtual machine as a docker host, proxy settings must also be configured in the docker config. See Docker guides for instructions on how to properly set this. If necessary this can be wrapped in a task automation workflow for your own use.

Morpheus Cloud Capability Coverage

Table 13: Morpheus Cloud Capability Coverage

Cloud Integration	Ubuntu	CentOS	Debian	Linux Guest Cust	Cloud Init
Amazon	Yes	Yes	Yes	N/A	yes
Alibaba	Yes	Yes	Yes	N/A	Yes
Azure	Yes	Yes	Yes	Yes	Cloud Limited
Digital Ocean	Yes	Yes	Yes	N/A	Yes
Google Cloud	Yes	Yes	Yes	N/A	Yes
Huawei	Yes	Yes	Yes	N/A	Yes
Hyperv	Yes	Yes	Yes	N/A	Yes
IBM Cloud / Softlayer	Yes	Yes	Yes	N/A	N/A
KVM	Yes	Yes	Yes	N/A	Yes
Nutanix	Yes	Yes	Yes	N/A	Yes
Openstack	Yes	Yes	Yes	N/A	Yes
Oracle Cloud	Yes	Yes	Yes	N/A	Yes
OVF	Yes	Yes	Yes	N/A	Yes
OTC	Yes	Yes	Yes	N/A	Yes
SCVMM	Yes	Yes	Yes	NO	Yes
Upcloud	Yes	Yes	Yes	N/A	Yes
VCD	Yes	Yes	Yes	No	Yes
Vmware	Yes	Yes	Yes	Yes	Yes
Xen	Yes	Yes	Yes	No	Yes

Containers

Docker Registry

Overview

Without any additional configuration Morpheus can provision images from Docker's public hub at <https://hub.docker.com/> using their public api at <https://index.docker.io/v1/>

However, many organizations maintain private Docker registries for security measures. Additional public and private Docker registries can be added to Morpheus.

Adding a Docker Registry Integration

1. Navigate to *Administration -> Integrations*
2. Click "New Integration"
3. Select the *Docker Repository* Type
4. Add the following:
 - Name** Name for the Registry in Morpheus
 - Repository url** Docker Registry url or IP address
 - Username** Username if private registry

Password Password if private registry

5. Save Changes

Note: You must either have signed certificates for your registry or configure your docker host(s) to accept insecure registries

Provisioning an Instance from Docker Registry

Docker images from the Integrated Registry can be provisioned using the generic *Docker* Instance Type, or by adding images to Node Types for custom Library Instance Types.

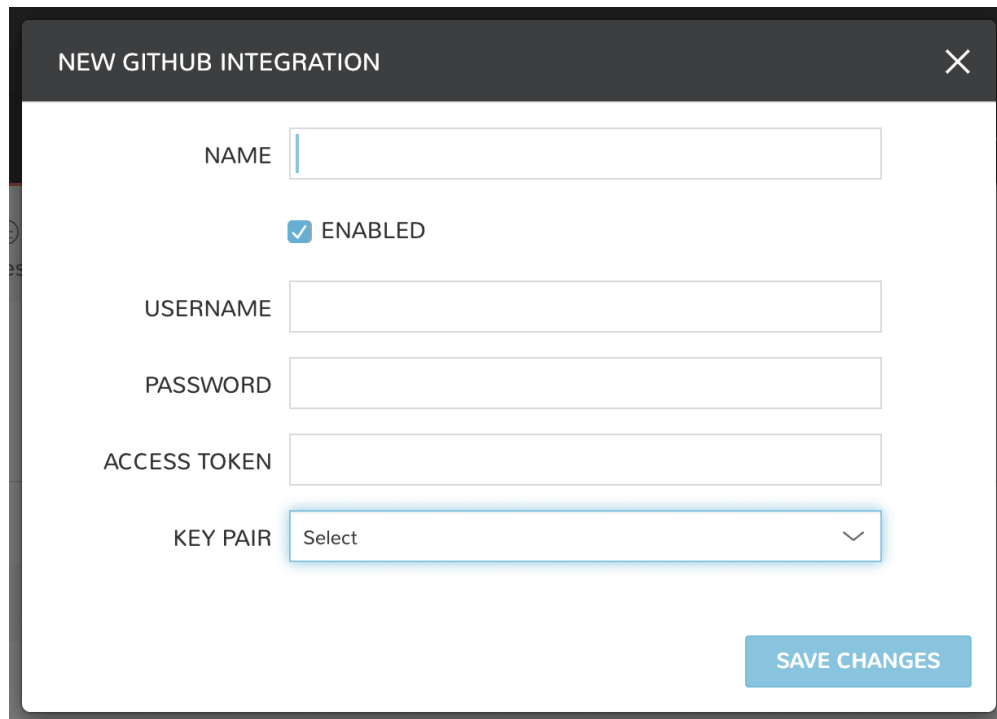
Deployment

Git

Authentication

Add a private Github or Git repository. Authentication can be handled by supplying any one of the following:

- Username and password
- Access token
- Key pair



NEW GITHUB INTEGRATION

NAME

☒ ENABLED

USERNAME

PASSWORD

ACCESS TOKEN

KEY PAIR

SAVE CHANGES

Note: Git and Github integrations can be authenticated over HTTPS with a username and password or with an access token. They are authenticated over SSH by providing a key pair. In previous versions of Morpheus, Git (not Github)

integrations could only be authenticated over SSH.

Key pairs are stored in Morpheus and selected from a dropdown menu when needed. To add a key pair to Morpheus:

1. Generate an SSH key pair, or use an existing SSH key pair.
2. Navigate to Infrastructure -> Keys & Certs
3. Select + *ADD*
4. Enter both the Public and Private keys
5. Click *SAVE CHANGES*

DNS

AWS Route53

Overview

Morpheus integrates directly with Amazon Route 53 to automatically create DNS entries for Instances provisioned to a configured Cloud or Group. Morpheus also syncs in Route 53 Domains for easy selection while provisioning, or setting as the default Domain on a Cloud or Network.

Add Route 53 Integration

Route 53 can be added in the *Administration* or *Infrastructure* sections:

1. In Administration -> Integrations, select + *New Integration*
2. In Infrastructure -> Networks -> Services, select *Add Service*
3. Provide the following:

TYPE Route 53

NAME Name for the Integration in Morpheus

REGION AWS Region for the Integration

ACCESS KEY AWS User IAM Access Key

SECRET KEY AWS User IAM Secret Key

4. Once saved the Integration will be added and visible in both Administration -> Integrations and Infrastructure -> Networks -> Services

Note: All fields can be edited after saving.

Domains

Once the integration is added, Route 53 Domains will sync and listed under Infrastructure -> Networks -> Domains.

Note: Default Domains can be set on Networks and Clouds, and can be selected when provisioning. Additional configuration options are available by editing a domain in *Networks -> Domains*

Configuring Route 53 with Clouds and Groups

DNS Integrations are available in the *DNS Integration* dropdown in Cloud and Group settings.

Morpheus will register Instances with the DNS provider when provisioned into a Cloud or Group with a DNS Integration added.

Add DNS Integration to a Cloud

1. In Infrastructure -> Clouds edit the target Cloud.
2. Expand the *Advanced Options* section.
3. In the *DNS Integration* dropdown, select an available DNS Integration.
4. Save Changes

Add DNS Integration to a Group

1. In Infrastructure -> Groups select the target Group.
2. Select the *Edit* button for the Group
3. Expand the *Advanced Options* section.
4. In the *DNS Integration* dropdown, select an available DNS Integration.
5. Save Changes

Note: Instances provisioned into a Cloud or Group with a DNS Integration added will be registered as instance-name.domain with the DNS Provider during provisioning, and de-registered at teardown.

Microsoft DNS

Overview

Morpheus integrates directly with Microsoft DNS to automatically create DNS entries for Instances provisioned to a configured Cloud or Group. Morpheus also syncs in Microsoft DNS Domains for easy selection while provisioning, or setting as the default Domain on a Cloud or Network.

Add Microsoft DNS Integration

Important: The Morpheus Microsoft DNS integration works over http/5985. If you have turned off the http listener on 5985 and only enabled https/5986 it will fail.

Microsoft DNS can be added in the *Administration* or *Infrastructure* sections:

1. In *Administration* -> *Integrations*, select *+ New Integration*
2. In *Infrastructure* -> *Networks* -> *Services*, select *Add Service*
3. Provide the following:
 - TYPE** Microsoft DNS
 - NAME** Name for the Integration in Morpheus
 - DNS SERVER** IP or resolvable hostname of DNS server
 - USERNAME** DNS provider username
 - PASSWORD** DNS provider user password
 - ZONE** (Optional) Enter a dns zone to limit scope
 - CREATE POINTERS** Enabled to create A records during provisioning
4. Once saved the Integration will be added and visible in both *Administration* -> *Integrations* and *Infrastructure* -> *Networks* -> *Services*

Note: All fields can be edited after saving.

Domains

Once the integration is added, Microsoft DNS Domains will sync and listed under *Infrastructure* -> *Networks* -> *Domains*.

Note: Default Domains can be set on *Networks* and *Clouds*, and can be selected when provisioning. Additional configuration options are available by editing a domain in *Networks* -> *Domains*

Configuring Microsoft DNS with Clouds and Groups

DNS Integrations are available in the *DNS Integration* dropdown in *Cloud* and *Group* settings. Morpheus will register Instances with the DNS provider when provisioned into a *Cloud* or *Group* with a DNS Integration added.

To take full advantage of the Morpheus Microsoft DNS integration, a service account in the Admins group is not required. However, an account must have the following minimum access to use all features:

- Read, Create, and Delete rights on objects
- Belongs to the local group “**WinRMRemoteWMIUsers__**”
- WinRM Quickconfig must be run on the DNS server
- CIMv2 needs access according to instructions in our [KnowledgeBase](#)

Add DNS Integration to a Cloud

1. In *Infrastructure* -> *Clouds* edit the target Cloud.
2. Expand the *Advanced Options* section.
3. In the *DNS Integration* dropdown, select an available DNS Integration.
4. Save Changes

Add DNS Integration to a Group

1. In *Infrastructure* -> *Groups* select the target Group.
2. Select the *Edit* button for the Group
3. Expand the *Advanced Options* section.
4. In the *DNS Integration* dropdown, select an available DNS Integration.
5. Save Changes

Note: Instances provisioned into a Cloud or Group with a DNS Integration added will be registered as instance-name.domain with the DNS Provider during provisioning, and de-registered at teardown.

Power DNS

Overview

Morpheus integrates directly with Power DNS to automatically create DNS entries for Instances provisioned to a configured Cloud or Group. Morpheus also syncs in Power DNS Domains for easy selection while provisioning, or setting as the default Domain on a Cloud or Network.

Add Power DNS Integration

Power DNS can be added in the *Administration* or *Infrastructure* sections:

1. In *Administration* -> *Integrations*, select *+ New Integration*
2. In *Infrastructure* -> *Networks* -> *Services*, select *Add Service*
3. Provide the following:

TYPE Power DNS

NAME Name for the Integration in Morpheus

API HOST URL of Power DNS API. Example: `http://10.30.20.10:8081`

Token Power DNS API Token

Version Power DNS API Version

4. Once saved the Integration will be added and visible in both *Administration* -> *Integrations* and *Infrastructure* -> *Networks* -> *Services*

Note: All fields can be edited after saving.

Domains

Once the integration is added, Power DNS Domains will sync and listed under Infrastructure -> Networks -> Domains.

Note: Default Domains can be set on Networks and Clouds, and can be selected when provisioning. Additional configuration options are available by editing a domain in *Networks -> Domains*

Configuring Power DNS with Clouds and Groups

DNS Integrations are available in the *DNS Integration* dropdown in Cloud and Group settings.

Morpheus will register Instances with the DNS provider when provisioned into a Cloud or Group with a DNS Integration added.

Add DNS Integration to a Cloud

1. In Infrastructure -> Clouds edit the target Cloud.
2. Expand the *Advanced Options* section.
3. In the *DNS Integration* dropdown, select an available DNS Integration.
4. Save Changes

Add DNS Integration to a Group

1. In Infrastructure -> Groups select the target Group.
2. Select the *Edit* button for the Group
3. Expand the *Advanced Options* section.
4. In the *DNS Integration* dropdown, select an available DNS Integration.
5. Save Changes

Note: Instances provisioned into a Cloud or Group with a DNS Integration added will be registered as instance-name.domain with the DNS Provider during provisioning, and de-registered at teardown.

Identity Management

Active Directory

Overview

Active Directory is Microsoft's primary authentication service widely used in Enterprise organizations and even via Microsoft's cloud services. While Active Directory also supports LDAP protocol support (which Morpheus can integrate with as well), the main Active Directory integration can also be utilized. It is even possible to map Active Directory groups to equivalent Roles within Morpheus. Morpheus will connect over port 389 for non-secure LDAP and port 636 for secure LDAP.

Note: To use Active Directory, a valid / trusted SSL certificate must be in place on the Active Directory services (self signed will not work).

Adding an Active Directory Integration

1. Navigate to `Administration -> Tenants`
2. Select a Tenant
3. Select *IDENTITY SOURCES*
4. Select + *IDENTITY SOURCE*
5. Choose "Active Directory"
6. Populate the following:

Name Unique name for authentication type.

AD Server Hostname or IP address of AD Server.

Domain Domain name of AD Domain.

Binding Username Service account username for bind user.

Binding Password Password for bind service account.

Required Group The AD group users must be in to have access (optional)

Default Role The default role a user is assigned if no group is listed under AD user that maps under Role Mappings section.

Service Account Holder This is the admin account type in Morpheus and an AD group can be created and populated to a user that this role should be assigned. Roles are assigned dynamically based on group membership.

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

1. Select *SAVE CHANGES*.

Now allowed AD users can login to Morpheus via their Active Directory credentials and a User will be automatically generated to Morpheus with matching metadata and mapped Role permissions.

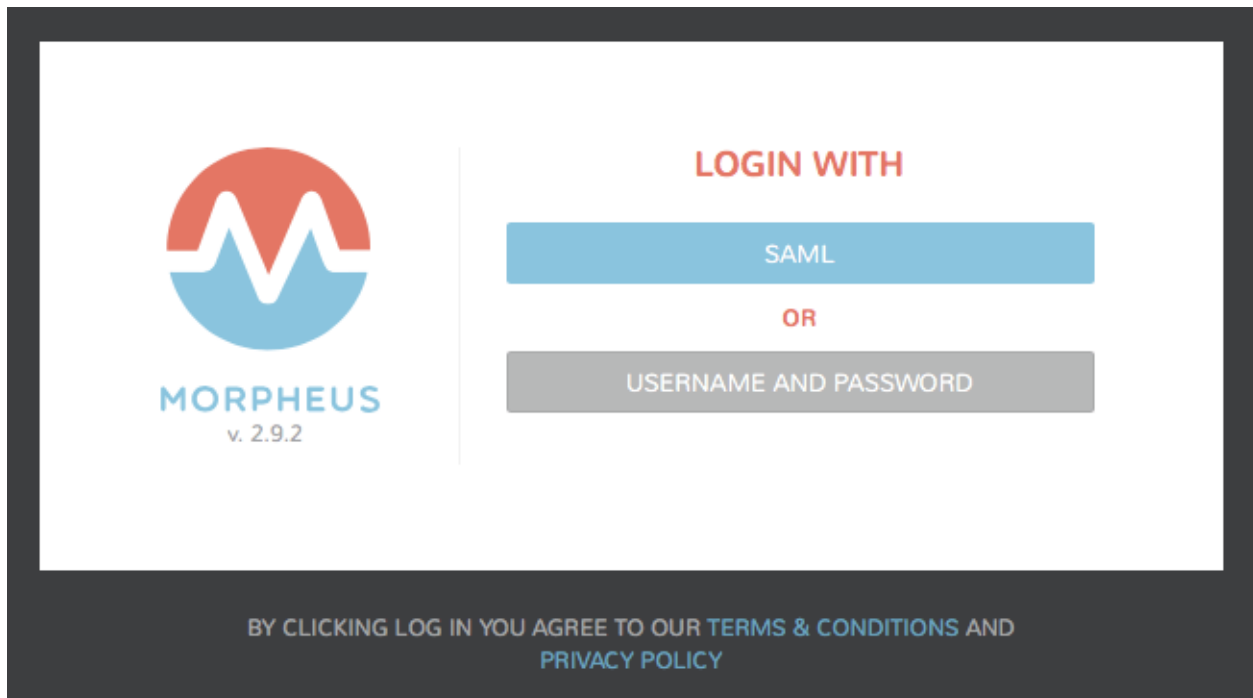
Note: Only the username is required with password, not the `username@domain`.

Note: Sub-tenant Morpheus API authentication for Active Directory generated users is not currently supported.

SAML Integration

Overview

The Morpheus SAML identity source integration allows customers to add user SSO to Morpheus , authenticated by external login SAML providers.



Adding a SAML Integration

To add a SAML integration:

1. Navigate to Administration -> Tenants
2. Select a tenant.
3. Select IDENTITY SOURCES in the Tenant detail page
4. Select + *ADD IDENTITY SOURCE*.
5. Select SAML (external login) from the TYPE field
6. Add a Name and optional Description for the SAML integration

NEW IDENTITY SOURCE

Identity Source

TYPE: SAML

NAME:

DESCRIPTION:

SAML Configuration

LOGIN REDIRECT URL:

☐ Do not include SAMLRequest parameter

LOGOUT POST URL:

SIGNING PUBLIC KEY:

☒ Do not validate SAMLResponse signatures

ADVANCED VALIDATION OPTIONS: [Show/Hide](#)

Role Attribute Value: [Show/Hide](#)

Role Mappings

DEFAULT ROLE: Developer

ROLE ATTRIBUTE NAME: memberOf

REQUIRED ROLE ATTRIBUTE VALUE:

DEVELOPER:

DIPESH-008USER:

IMAGE DEVELOPER:

LIPSCOMBE USER:

MORPHEUS ADMIN:

MORPHEUS SALES:

SAVE CHANGES

There are 3 sections with fields that need to be populated depending on the desired configuration:

- SAML Configuration
- Role Mappings
- User Attribute Names

SAML Configuration

LOGIN REDIRECT URL This is the SAML endpoint Morpheus will redirect to when a user signs into Morpheus via SAML.

LOGOUT POST URL The url morpheus will post to when a SAML user log out of Morpheus to log out of the SAML provider as well.

SIGNING PUBLIC KEY Add the X.509 Certificate public key from the SAML provider.

Role Mappings

DEFAULT ROLE Role a saml user will be assigned by default when no role is mapped

ROLE ATTRIBUTE NAME The name of the attribute field that will map to morpheus roles, such a MemberOf

REQUIRED ROLE ATTRIBUTE VALUE Role attribute value that a user must be assigned/a member of to be authorized, such as group or role in the SAML SP.

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

The rest of the Role Mapping Fields will be the existing Roles in morpheus with a Role Attribute Value field.

User Attribute Names

GIVEN NAME ATTRIBUTE NAME SAML SP field value to map to Morpheus user First Name

SURNAME ATTRIBUTE NAME SAML SP field value to map to Morpheus user Last Name

EMAIL ATTRIBUTE NAME SAML SP field value to map to Morpheus user email address

EDIT IDENTITY SOURCE



Identity Source

TYPE	SAML [Beta]
ACTIVE	Yes
NAME	SAML
DESCRIPTION	onelogin SAML

SAML Configuration

LOGIN REDIRECT URL	https://morpheusdata-dev.onelogin.com/trust/saml2/http- <input type="checkbox"/> Do not include SAMLRequest parameter
LOGOUT POST URL	https://morpheusdata-dev.onelogin.com/trust/saml2/http-
SIGNING PUBLIC KEY	MIIEFzCCAv+gAwIBAgIUayYdMuoXBTGcalAARanxhRJwwtQwDQYJKoZIhvcNAQEF

Role Mappings

DEFAULT ROLE	System Admin
ROLE ATTRIBUTE NAME	MemberOf
REQUIRED ROLE ATTRIBUTE VALUE	dev
LEGACY ACCOUNT ADMIN	Role Attribute Value

User Attribute Names

Show/Hide

GIVEN NAME ATTRIBUTE NAME	firstName
SURNAME ATTRIBUTE NAME	lastName
EMAIL ATTRIBUTE NAME	Email

DEACTIVATE

DELETE

SAVE CHANGES

Once populated, select SAVE CHANGES and the SAML identity source integration will be added.

In the Identity Sources section, important information for configuration of the SAML integration is provided. Use the SP ENTITY ID and SP ACS URL for configuration on the external login SAML provider side.

- SP ENTITY ID
- SP ACS URL*
- IDP LOGIN REDIRECT URL
- IDP LOGOUT POST URL
- SP METADATA

Master Account

<input type="text" value="Search"/>		+ ADD IDENTITY SOURCE	
TYPE	NAME	DETAILS	ACTIVE
SAML [Beta]	SAML <i>onelogin SAML</i>	SP ENTITY ID: https://someip.com/saml/CDWPjmZt SP ACS URL: https://someip.com/saml/CDWPjmZt IDP LOGIN REDIRECT URL: https://someip.com/saml/CDWPjmZt IDP LOGOUT POST URL: https://someip.com/saml/CDWPjmZt SP METADATA: VIEW	Yes

Sample Metadata code output:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?><EntityDescriptor entityID=
  ↪ "https://someip.com/saml/CDWPjmZt" xmlns="urn:oasis:names:tc:SAML:2.0:metadata">
  ↪ <SPSSODescriptor AuthnRequestsSigned="false" WantAssertionsSigned="true"
  ↪ protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol"><NameIDFormat>
  ↪ urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified</NameIDFormat>
  ↪ <AssertionConsumerService index="0" isDefault="true" Binding=
  ↪ "urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Location="https://someip.com/
  ↪ externalLogin/callback/CDWPjmZt" /></SPSSODescriptor></EntityDescriptor>
```

Note: Different SAML providers will have different field names and requirements. A onelogin SAML Test Connector (IdP w/attr) was used for the example integration this article.

Onelogin SAML SSO

For Onelogin SAML integration, the following fields are mapped:

- LOGIN REDIRECT URL : SAML 2.0 Endpoint (HTTP)
- LOGOUT POST URL : SLO Endpoint (HTTP)
- SIGNING PUBLIC KEY : X.509 Certificate
- SP ENTITY ID: ACS (Consumer) URL Validator
- SP ACS URL: ACS (Consumer) URL

Azure Active Directory SSO (SAML)

Azure Active Directory Single Sign-on can be added as a Identity Source in Morpheus using the SAML Identity Source Type. The Azure AD SSO configuration is slightly different than other SAML providers, and this guide will assist in adding a Azure AD SSO Identity Source.

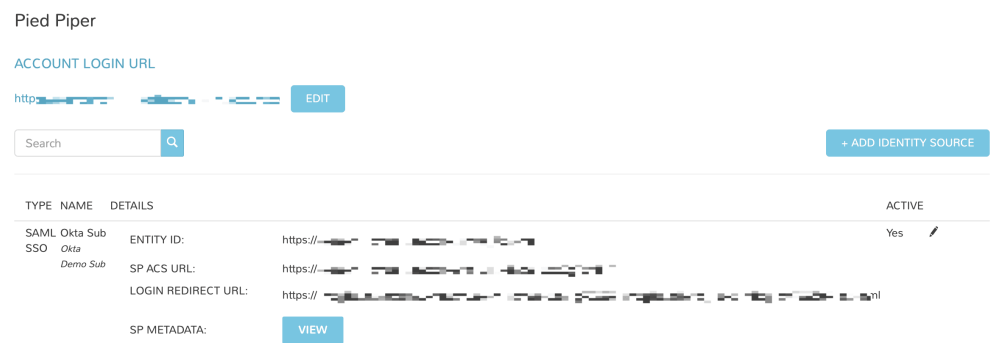
Create a Azure AD SAML Integration

Azure requires inputting the *Identifier (Entity ID)* and *Reply URL (Assertion Consumer Service URL)* in the Azure SSO configuration before it provides the Endpoints and Certificate necessary to add the Integration into Morpheus. In order to get the *Identifier (Entity ID)* and *Reply URL (Assertion Consumer Service URL)* to input into Azure SSO config, we need to create a base SAML Integration in Morpheus first.

To add a base SAML integration:

- 1. Navigate to Administration -> Tenants
- 2. Select a tenant.
- 3. Select IDENTITY SOURCES in the Tenant detail page
- 4. Select + ADD IDENTITY SOURCE.
- 5. Select SAML SSO from the TYPE field
- 6. Add a Name, optional Description and any value in the LOGIN REDIRECT URL field. Since we do not have the LOGIN REDIRECT URL from Azure yet, type any text such as test into the LOGIN REDIRECT URL field so the Identity Source Integration can be saved and the *Identifier (Entity ID)* and *Reply URL (Assertion Consumer Service URL)* generated. We will edit the Integration with the proper LOGIN REDIRECT URL after configuring SSO in Azure.
- 7. Select SAVE CHANGES.

Upon save, the *Entity ID (Identifier (Entity ID))* and *SP ACS URL (Reply URL (Assertion Consumer Service URL))* will be provide in the Identity Source list view. Copy these for use in Azure SSO config.



Configure Azure SSO

This guide assumes an Azure AD Application has already been created in Azure with a subscription level high enough to configure SSO in the application. Please refer to Azure documentation if this has not already been configured.

- Next, in the Azure Active Directory Application details page, select `Single sign-on`, then enter the following:
 - Single Sign-on Mode dropdown** Select `SAML-based Sign-on`
 - Identifier (Entity ID)** Enter the `Entity ID` URL from the Morpheus Identity Source Integration above.
 - Reply URL (Assertion Consumer Service URL)** Enter the `SP ACS URL` from the Morpheus Identity Source Integration above.
- Save and click the `Test SAML Settings` button. Azure will confirm connection with Morpheus
- In Azure's `User Attributes & Claims` settings (step 2), select `Add a group claim` with value `user.groups [SecurityGroup]`

User Attributes & Claims config

Table 14: Required Claim

Claim name	Value
Unique User Identifier (Name ID)	<code>user.userprincipalname [nameid-format:emailAddress]</code>

Table 15: Additional Claims

Claim name	Value
<code>http://schemas.microsoft.com/ws/2008/06/identity/claims/groups</code>	<code>user.groups [SecurityGroup]</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress</code>	<code>user.mail</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname</code>	<code>user.givenname</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name</code>	<code>user.userprincipalname</code>
<code>http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname</code>	<code>user.surname</code>

- Copy or keep available for reference the the Claim Names/Namespace URLs for entering Role Attribute Values in the Morpheus Identity Source Integration.
- In Azure SSO config, if one has not been generated, select `Create new certificate` to generate a new SAML Signing Certificate.
- Enter a valid email address to receive certificate expiration notifications (these are not Morpheus-generated email).
- In Azure SSO config, select `Configure {AD App Name}`
- In the `Configure sign-on` pane, copy the following:
 - SAML Single Sign-On Service URL** This will be used for the LOGIN REDIRECT URL in the Morpheus Identity Source Integration settings
 - Sign-Out URL** This will be used for the LOGOUT REDIRECT URL in the Morpheus Identity Source Integration settings

- Click on the **SAML XML Metadata** link, open the xml file, and copy the key between the **<X509Certificate>** and **</X509Certificate>**. This will be used for the *Public Key* value in the SAML RESPONSE section of the Morpheus Identity Source Integration settings

Example Key (this key is an example and is not valid):

```
MIIC8ECCAdigAwIBAgIQEOZX1N5wY9Dc6Ow1sKEMzANBgkqhkiG9w0BAQsFADA0MTIwMAYDVQQDEylNaWNyb3NvZnQ
→V6GcBpRkoxJd0DLbhubwd0kp65LD9IIh5PUY2ohBHvrFAy3SZ04mXoH7LWvY3oNrxaNAksbYF6phOkONf/
→XeTdZor14xdGnTuD9zRqPsJHHisyfFBUG/CxYxzO6w9fAPzJGLzc0Y7o5lMW2OjINaQI4R/
→pqp3qw+nYf7DXSzY6tf1Sspk64jffZDt1jSVjD7upMITKPeOCRmeBUCnebJzwXqFBO7l4Vf5gloEJyftT7Wpr4VVmo
→pH6xzQVRz0GZQpol9ViQJJbJJqhLm4LjWT9VU2lYqdi0NdgtK7QthZo4J0ZFdUG6qfFTfPKqVn0AEHxiM4JWxfigz
→y56+ksYSRP87XdOcVvTftHYmQnDOF0qKrpqMK7LtmsEwqc7rKX7nTCenZnBEOCFDBVH4QEzMrAznpEdJnQs9nJZB
→sec
```

9. Save the SSO config in Azure AD app and return to Morpheus

Edit the existing Azure AD SAML Integration

Now that we have the required information, we can finalize the Azure AD SAML Integration in Morpheus

1. Edit the existing Azure AD SAML Integration created in the first step and populate the following:

LOGIN REDIRECT URL Add the SAML Single Sign-On Service URL copied from Azure SSO config.

LOGOUT REDIRECT URL Add the Sign-Out URL copied from Azure SSO config.

SAML RESPONSE Set to “Validate Assertion Signature”, then in the “Public Key” field enter the Public Key value we discussed in the last section

GIVEN NAME ATTRIBUTE NAME (May have to click “show” to see hidden SAML Assertion Attribute Names fields)

Enter the `givenname` Namespace url from Azure SSO config: <http://schemas.xmlsoap.org/ws/2005/05/identity/claims>

SURNAME ATTRIBUTE NAME Enter the `emailaddress` Namespace url from Azure SSO config: <http://schemas.xmlsoap.org/ws/2005/05/identity/claims>

EMAIL ATTRIBUTE NAME (May need to scroll down within the SAML Assertion Attribute Names section see this field)

Enter the `surname` Namespace url from Azure SSO config: <http://schemas.xmlsoap.org/ws/2005/05/identity/claims>

Configure Role Mappings

Role mappings will map Azure AD Groups to Morpheus Roles. Azure AD users will be assigned Roles in Morpheus upon signing in based on their Group Membership in Azure AD.

Important: Use an Azure Groups Object ID, not Group name, when entering Role Mappings. Example: 7626a4a2-b388-4d9b-a228-72ce9a33bd4b

DEFAULT ROLE Role a Azure AD user will be assigned by default upon signing in to Morpheus using this Identity Source.

REQUIRED AZURE AD GROUP OBJECT ID Object ID of Azure AD Group a user must be a member of to be authorized to sign in to Morpheus. Users not belonging to this Group will not be authorized to login to Morpheus. This field is optional, and if left blank, any user from the Azure AD App will be able to sign in to Morpheus and will be assigned the Default Role if no Role Mappings match AD Group membership.

GROUP ASSERTION ATTRIBUTE NAME Enter `http://schemas.microsoft.com/ws/2008/06/identity/claims/groups` for Azure AD SSO

Additional Role Mappings The existing Roles in Morpheus will be listed. To map a Morpheus Role to an Azure AD Group, enter the Object ID of the desired Azure AD Group in the *Role Attribute Value* field for the corresponding Morpheus Role.

Important: Use an Azure Groups Object ID, not Group name, when entering Role Mappings. Example: 7626a4a2-b388-4d9b-a228-72ce9a33bd4b

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

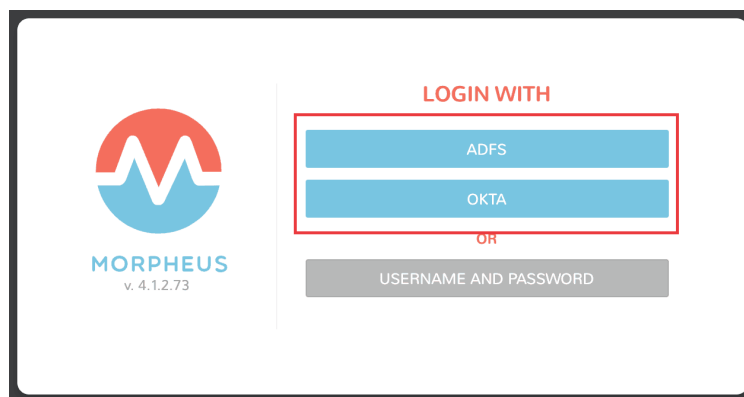
Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

Once populated, select *SAVE CHANGES* and the SAML identity source integration will be added. The Identity Source can be edited anytime to deactivate or change Role Mappings or other values.

Note: If Role mappings are edited after Azure AD SSO users have signed into Morpheus, currently logged in users will need to log out of Morpheus for the new Role mappings to take effect, when applicable.

Signing In to Morpheus

When there is an active SAML/Azure AD SSO Identity Source Integration, a new button will appear on the Morpheus login page with the name of the Identity Source Integration as the button title. Example: *ADFS*. Another button titled “USERNAME AND PASSWORD” is also added for Morpheus account authentication outside of an Identity Source.



- **SAML/Azure AD SSO users can log into Morpheus by clicking the SAML button** This will redirect the User to Azure AD app sign in url. If they are currently signed into Azure and authorized, the user will be instantly signed into Morpheus.
- Local Morpheus users can select “USERNAME AND PASSWORD” to sign in with their local credentials as before.

Note: If no local users other than the System Admin have been created, “USERNAME AND PASSWORD” option will not be displayed, only the SAML option.

OneLogin

Adding OneLogin Identity Source Integration

1. Navigate to Administration -> Tenants
2. Select the Tenant to add the Identity Source Integration
3. Select *IDENTITY SOURCES*
4. Select + *IDENTITY SOURCE*
5. Enter the following:

TYPE OneLogin

NAME

Name of the Identity Source Integration in Morpheus

DESCRIPTION Optional Description of the Identity Source

ONELOGIN SUBDOMAIN

example: morpheus-dev

Warning: Please verify the subdomain carefully. An invalid subdomain will cause authentication attempts by OneLogin users to fail.

ONELOGIN REGION Specify US or EU region

API CLIENT SECRET OneLogin API Client Secret from the Settings - API section in OneLogin portal

API CLIENT ID OneLogin API Client ID from the Settings - API section in OneLogin portal

REQUIRED ROLE Enter a role if OneLogin users logging into morpheus must have at least this OneLogin role to gain access to Morpheus.

DEFAULT ROLE The default Morpheus Role applied to users created from OneLogin Integration if no other role mapping is specified below

ROLE MAPPINGS Existing Morpheus Roles will be listed with fields to enter OneLogin Roles to map to. Users with OneLogin roles matching the role mappings will be assigned the appropriate Role(s) in Morpheus when signing in.

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (Administration > Users > Selected user).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

6. Select *SAVE CHANGES* and the OneLogin Integration will be added.

Users can now login to Morpheus with OneLogin credentials. The first Login will create a user in Morpheus matching the Username, email and Password from OneLogin. If a **REQUIRED ROLE** is specified in the Identity Source settings, only users with that Role in OneLogin will be able to login to Morpheus.

Important: OneLogin users will not authenticate in Morpheus if there is an existing Morpheus User with matching username or email address.

Okta

Overview

Morpheus allows users to integrate an Okta deployment for user management and authentication. In Morpheus, identity sources are added on a per-Tenant basis and Morpheus allows you to map Okta user groups to Morpheus user groups. User accounts are automatically created with matching metadata and role permissions when users are authenticated.

Adding an Okta Integration

1. Navigate to *Administration* -> *Tenants*
2. Select a Tenant
3. Select *IDENTITY SOURCES*
4. Select + *IDENTITY SOURCE*
5. Choose TYPE: "Okta"
6. Populate the following, then select *SAVE CHANGES*:

Name Unique name for authentication type

Description A description for your new Okta Identity Source

Okta URL Your Okta URL

Administrator API Token Your Okta Administrator API Token

Required Group The Okta group that users must be in to have access (optional)

Default Role The default role a user is assigned if no group is listed under an Okta user that maps within the Morpheus Role Mappings section

ENABLE ROLE MAPPING PERMISSION When selected, Tenant users with appropriate rights to view and edit Roles will have the ability to set role mapping for the Identity Source integration. This allows the Tenant user to edit only the role mappings without viewing or potentially editing the Identity Source configuration.

MANUAL ROLE ASSIGNMENT When selected, administrators can manually edit Roles for users created through this identity source integration from the user detail page (*Administration > Users > Selected user*).

Note: For more on Identity Source role mapping permissions, see the [associated guide](#) in our KnowledgeBase.

Now, allowed Okta users can log into Morpheus via their Okta credentials and a user will be automatically generated within Morpheus with matching metadata and mapped Role permissions.

Note: If you've created multi-tenant roles, these will also appear here and can be mapped to Okta user groups allowing you to map users to equivalent user groups in Morpheus.

ITSM

ServiceNow

Overview

IT Service Management (ITSM) is an important area of focus for many organizations. Organizations invested in ServiceNow as an ITSM provider will find that Morpheus integrates tightly with some of the most-used features. After integrating ServiceNow with Morpheus, both environments can be used interchangeably and the results are synced to both places. This guide walks administrators through the process of integrating ServiceNow with Morpheus and how Morpheus can be used to effectively leverage the best of ServiceNow.

Tip: The ServiceNow integration guide is also available as a [PDF download](#), which includes additional example use cases and screenshots.

Add ServiceNow Integration

1. Navigate to Administration > Integrations
2. Select + *NEW INTEGRATION*
3. Select "ServiceNow" from the dropdown list
4. Add the following:

NAME A friendly name to describe the ServiceNow integration in Morpheus.

ENABLED Check "Enabled" to allow consumption of this ServiceNow integration in Morpheus.

HOST URL of the ServiceNow instance (ex: <https://your.instance.service-now.com>), keep in mind you can create multiple ServiceNow integrations in Morpheus if needed.

USER/PASSWORD A user in ServiceNow that is able to access the REST interface and create/update/delete incidents, requests, requested items, item options, catalog items, workflows, etc. The list of necessary roles includes `x_moda_morpheus_ca.integration` (available if the Morpheus ServiceNow plugin is installed from the ServiceNow Store), `catalog_admin`, `itil`, `rest_service`, and `import_transformer`.

CMDB CUSTOM MAPPING If needed, administrators can opt to populate a specific field in the ServiceNow table and such mapping is identified here with a JSON code snippet. Below is an example that populates the `object_id` field in the CM database with the Morpheus instance name:


```
{
  "object_id": "<%=instance.name%>";
  "SN_field_id2": "<%=morph.varname2%>";
  "SN_field_id3": "<%=morph.varname3%>"
}
```

CMDB BUSINESS OBJECT Allows the user to define the table CMDB records are written to if they prefer this over Morpheus default. By default, Morpheus writes to the `cmdb_ci_vm_instance` table.

5. Save Changes

ServiceNow Configuration Management Database (CMDB)

ServiceNow CMDB is central to maintaining a complete record of IT infrastructure at many organizations. The Morpheus ServiceNow integration can create and update configuration item (CIs) records as new services are provisioned or existing services are reconfigured. Once a ServiceNow integration is set as the CMDB for a Cloud or Group, CI records are created and managed by Morpheus.

Setting a CMDB on a Group

When adding or editing a Morpheus Group, any active ServiceNow integration can be set as the CMDB.

1. Navigate to Infrastructure > Groups
2. Select an existing Group name from the list
3. Click *EDIT*
4. Under “Advanced Options”, select an active ServiceNow integration from the CMDB dropdown menu

This setting is also available when creating a Group. Rather than selecting an existing Group in step two above, click + *CREATE* to make a new Group.

EDIT GROUP ✕

Configuration

NAME

CODE

LOCATION

▼ Advanced Options

DNS SERVICE

CMDB

SERVICE REGISTRY

CONFIG MANAGEMENT

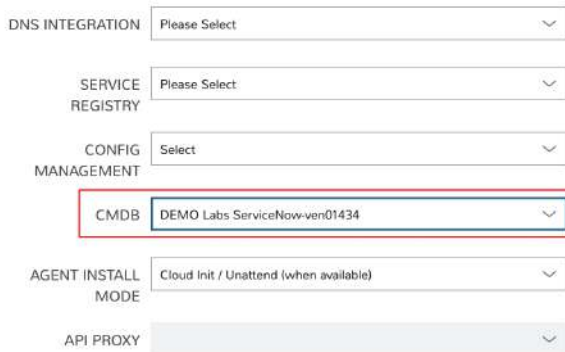
SAVE CHANGES

Setting a CMDB on a Cloud

When adding or editing a Morpheus Cloud, any active ServiceNow integration can be set as the CMDB.

1. Navigate to Infrastructure > Clouds
2. Select an existing Cloud name from the list
3. Click *EDIT*
4. Under “Advanced Options”, select an active ServiceNow integration from the CMDB dropdown menu

This setting is also available when creating a Cloud. Rather than selecting an existing Cloud in step two above, click + *ADD* to make a new Cloud.



The screenshot shows a configuration form for a Morpheus Cloud. The form includes several dropdown menus: 'DNS INTEGRATION' (Please Select), 'SERVICE REGISTRY' (Please Select), 'CONFIG MANAGEMENT' (Select), 'CMDB' (DEMO Labs ServiceNow-ver01434), 'AGENT INSTALL MODE' (Cloud Init / Unattend (when available)), and 'API PROXY'. The 'CMDB' dropdown is highlighted with a red rectangle, indicating it is the focus of the instruction.

Provisioning and CI Records

With a ServiceNow instance integrated with Morpheus and the instance set as the CMDB for a Morpheus Group or Cloud, we will see CI records created as new resources are provisioned to the Cloud or Group in Morpheus. After the provisioning process has completed, a CI record should exist with a name value equal to the Instance name in Morpheus.

Provisioned and active Instances in Morpheus will have CI records with an “On” state in ServiceNow. After they are deleted in Morpheus, the state value will be rolled to “Terminated” in ServiceNow as expected.

Morpheus will also populate a number of additional fields in ServiceNow including IP address, FQDN and more. Custom views can be created in ServiceNow to expose these fields.

ServiceNow Approval Policies

Morpheus offers its own approval engine out of the box, but some organizations prefer ServiceNow to be their final approval authority. With a ServiceNow instance integrated with Morpheus, administrators can create provision approval policies and tie them to an active ServiceNow integration. With the policy in place, any new provisioning within the policy scope (Global, Group, Cloud, User, or Role) is sent to ServiceNow for approval before provisioning will go ahead in Morpheus. Approvals are synced between the two applications every minute.

Add ServiceNow Provision Approval Policy to a Cloud

Note: Any Instance provisioned into a Cloud with an approval policy enabled will not proceed without the required approval.

To add a ServiceNow Approval policy to a Cloud:

1. Navigate to `Infrastructure > Clouds`
2. Select a Cloud by clicking on the desired Cloud name link
3. Select the **POLICIES** tab
4. Click + *ADD POLICY*
5. Select `Provision Approval` from the type dropdown
6. Optionally enter a description for the Policy
7. Configure the following:

APPROVAL INTEGRATION Select the ServiceNow Integration already configured in `Administration > Integrations` to use for the approval policy.

WORKFLOW Select the ServiceNow workflow for the approval in ServiceNow (if desired). These workflows are configured and synced in from the ServiceNow Integration.

TENANTS (if applicable) Only required for multi-tenant permission scoping. For the policy to apply to a Subtenant, type the name of the tenant(s) and select the Tenant(s) from the typeahead list.

8. Save Changes

Add ServiceNow Provision Approval Policy to a Group

Note: Any Instance provisioned into a Group with an approval policy enabled will not proceed without the required approval.

To add a ServiceNow Approval policy to a Group:

1. Navigate to `Infrastructure > Groups`
2. Select a Group by clicking on the Group name
3. Select the **POLICIES** tab
4. Click + *ADD POLICY*
5. Select `Provision Approval`
6. Optionally enter a description for the Policy
7. Configure the following:

APPROVAL INTEGRATION Select the ServiceNow Integration already configured in `Administration > Integrations` to use for the approval policy.

WORKFLOW Select the ServiceNow workflow for the approval in ServiceNow (if desired). These workflows are configured and synced in from the ServiceNow Integration.

TENANTS (if applicable) Only required for multi-tenant permission scoping. For the policy to apply to a Subtenant, type the name of the tenant(s) and select the Tenant(s) from the typeahead list.

8. Save Changes

Using ServiceNow Approval Policies

Any Instance provisioned into a Cloud or Group with an approval policy enabled will be in a PENDING state until the request is approved.

Instances pending a ServiceNow approval will show “Waiting for Approval” with the Requested Item number and Request number, ex: Waiting for Approval [RITM0010002 – REQ0010002].

ServiceNow approval requests are displayed in *Operations > Approvals*. Instances pending a ServiceNow approval must be approved in ServiceNow for provisioning to initiate. Approval requests from a ServiceNow approval policy cannot be approved in Morpheus, only approvals originating from Morpheus.

ServiceNow approval requests are displayed in Morpheus under *Operations > Approvals*. Pending ServiceNow approval requests can be cancelled in Morpheus by selecting the request and then selecting **ACTIONS > Cancel**.

Once a pending ServiceNow approval request is approved in ServiceNow, the Instance(s) will begin to provision in Morpheus within one minute of being approved in ServiceNow.

ServiceNow Monitoring Integration Settings

Note: A ServiceNow integration must be already configured in *Administration > Integrations* to enable ServiceNow monitoring.

The ServiceNow monitoring integration is enabled and configured in *Administration > Settings > Monitoring*. As long as the “Enabled” switch is activated, Morpheus will report monitoring data to ServiceNow. Configuration selections are described below:

Enabled Enables the ServiceNow monitoring integration

Integration Select from an existing ServiceNow integration in *Administration > Integrations*

New Incident Action The ServiceNow action to take when a Morpheus incident is created

Close Incident Action The Service Now action to take when a Morpheus incident is closed

Incident Severity Mapping

Morpheus Severity	ServiceNow Impact
Info	Low/Medium/High
Warning	Low/Medium/High
Critical	Low/Medium/High

Once finished working with configuration, click **APPLY**

ServiceNow Service Catalog Integration

In addition to integrating with key ServiceNow features, Morpheus offers a free plugin directly from the ServiceNow Store. At the time of this writing, the plugin supports ServiceNow releases New York, Orlando, and Paris. Once the plugin is installed, Morpheus Instance Types, Blueprints, and Self-Service Catalog Items can be presented as provisioning options in the ServiceNow catalog for ordering. The following is a guide to installing the Morpheus ServiceNow application.

Important: A valid SSL Certificate is required on the Morpheus Appliance for the ServiceNow plugin to be able to communicate with the appliance.

ServiceNow Configuration

1. Install the Morpheus plugin from the ServiceNow store
2. Navigate to Morpheus Catalog > Properties
3. Set the following properties:

Morpheus Appliance Endpoint The full URL to your Morpheus appliance

Username Username of the Morpheus administrator user

Password Password of the Morpheus administrator user

MID Server If desired, specify the name of a configured MID server to use

Adding to ServiceNow Catalog

Once the ServiceNow plugin is installed and configured, items can be added to the ServiceNow catalog from back in Morpheus. Follow the guide below to expose Morpheus Clouds, Library Items, and Blueprints to users in the ServiceNow catalog.

1. Navigate to *Administration > Integrations*
2. Select the relevant ServiceNow integration
3. From the Instances tab we can + *ADD CLOUD* or + *ADD LIBRARY ITEM*
4. From the Blueprints tab we can + *ADD BLUEPRINT*
5. From the Catalog Items tab, we can + *ADD CATALOG ITEM*
6. Back in ServiceNow, access the Morpheus plugin from the Service Catalog
7. Exposed Morpheus Library Items, Catalog Items, and Blueprints are visible here for ServiceNow users with sufficient role permissions

Integrations > DEMO Labs ServiceNow-ven01434

✓ DEMO Labs ServiceNow-ven01434 EDIT DELETE









service**now**

Host: https://ven01434.service-now.com Last Update: 12/01/2020 02:08 PM

Instances Blueprints Catalog Items




EXPOSED CLOUDS

First, select which clouds you would like to have exposed to ServiceNow for provisioning. + ADD CLOUD

TYPE	NAME	GROUP	LOCATION	
 amazon web services™	AWS	ServiceNow	US West (N. California)	
 Microsoft Azure	Azure US	ServiceNow	West US	
 Google Cloud Platform	Google Labs	ServiceNow		
 vmware™	VMware Demo Cluster	ServiceNow		

EXPOSED LIBRARIES

Next, select which library items you would like to have exposed to ServiceNow as Catalog Items. + ADD LIBRARY ITEM

TYPE	CLOUD	NAME	VERSION	
 APACHE®		Amazon Apache on Ubuntu 14.04	2.4	

Cherwell

Add Cherwell Integration

1. Navigate to Administration -> Integrations
2. Select + NEW INTEGRATION
3. Select Cherwell from the dropdown.
4. Add the following:

NAME Name of the Integration in Morpheus.

ENABLE Leave checked to enable the Integration.

HOST Url of the Cherwell Instance

USER Enter in username

PASSWORD Above Cherwell user's password

CLIENT KEY Provide your Cherwell client key

CREATED BY USER This is the full name of a user in the Cherwell system. When a new change management record is created in the Cherwell system, this user will be added to the record as the user that created it.

START DAYS FROM NOW Number of days from now to set proposed start date

END DAYS FROM NOW Number of days from now to set proposed end date

CUSTOM MAPPING This is an optional json object that allows the custom setting of the Cherwell fields on the Change Request object.

Note: The keys in the map correspond to the name of the field on the Change Request in Cherwell that you would like to set (see <https://bertram.d.pr/1Ziuhy> for a reference). In addition, the value in the map corresponds to the value you wish to use. Within the value, Morpheus variables may be used. Here is an example for setting the Description is:

```
{
  "Description": "Created from Morpheus by ${instance.createdByUsername}
  ↳ in ${zone.name}"
}
```

5. Save Changes

Remedy

PreRequisites

The user used for this integration need to be an Administrator in Remedy or have all the permissions to the form that is outlined in the table below.

API Endpoint	Action	BMC Form
/api/arsys/v1/entry/CTM:People	GET	CTM:People

Table 16 – continued from previous page

API Endpoint	Action	BMC Form
/api/arsys/v1/entry/COM:Company?q=%27Status%27=%22Enabled%22&fields=values(Company)	GET	COM:Company
/api/arsys/v1/entry/User	GET	User
/api/arsys/v1/entry/Group	GET	Group
/api/arsys/v1/entry/CHG:Infrastructure%20Change	POST	CHG:Infrastructure
/api/arsys/v1/entry/CHG:Infrastructure%20Change	PUT	CHG:Infrastructure
/api/arsys/v1/entry/CHG:Infrastructure%20Change	GET	CHG:Infrastructure
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_DiskDrive	POST	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_DiskDrive	PATCH	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_DiskDrive	GET	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_DiskDrive	DELETE	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_IPEndpoint	POST	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_IPEndpoint	PATCH	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_IPEndpoint	GET	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_IPEndpoint	DELETE	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Memory	POST	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Memory	PATCH	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Memory	GET	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Memory	DELETE	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Processor	POST	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Processor	PATCH	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Processor	GET	BMC.CORE:BMC
/api/cmdb/v1.0/instances/BMC.ASSET/BMC.CORE/BMC_Processor	DELETE	BMC.CORE:BMC
/api/arsys/v1/entry/AST:ComputerSystem	GET	AST:ComputeSystem
/api/arsys/v1/entry/AST:ComputerSystem	PUT	AST:ComputeSystem
/api/arsys/v1/entry/AST:ComputerSystem	POST	AST:ComputeSystem
/api/arsys/v1/entry/AST:IPEndpoint	GET	AST:IPEndpoint
/api/arsys/v1/entry/AST:IPEndpoint	PUT	AST:IPEndpoint
/api/arsys/v1/entry/AST:IPEndpoint	POST	AST:IPEndpoint
/api/arsys/v1/entry/AST:DiskDrive	GET	AST:DiskDrive
/api/arsys/v1/entry/AST:DiskDrive	PUT	AST:DiskDrive
/api/arsys/v1/entry/AST:DiskDrive	POST	AST:DiskDrive
/api/arsys/v1/entry/AST:Processor	GET	AST:Processor
/api/arsys/v1/entry/AST:Processor	PUT	AST:Processor
/api/arsys/v1/entry/AST:Processor	POST	AST:Processor
/api/arsys/v1/entry/AST:Memory	GET	AST:Memory
/api/arsys/v1/entry/AST:Memory	PUT	AST:Memory
/api/arsys/v1/entry/AST:Memory	POST	AST:Memory
/api/jwt/login	POST	

Add Remedy Integration

1. Navigate to Administration -> Integrations

2. Select + NEW INTEGRATION

3. Select Remedy from the dropdown.

4. Add the following:

NAME Name of the Integration in Morpheus.

ENABLE Leave checked to enable the Integration.

REMEDY HOST Url of the Remedy Instance. e.g: <http://xx.xx.xx.xx:8008>

USER Enter in username

PASSWORD Above Remedy user's password

COMPANY The dropdown will populate with values as soon as the auth using the above creds are successful

APPROVAL USER Full name of the user as it appear in Remedy. E.g: userid 'anish' would have full name as "Anish Abraham"

5. Save Changes

Keys and Certificates

Venafi

Overview

Morpheus integrates with Venafi to sync and request SSL certificates

Add Venafi

1. Navigate to Administration > Integrations

2. Select + *NEW INTEGRATION*

3. Enter in the following:

- Name
- Venafi Host
- Username
- Password

4. Click *SAVE CHANGES*

Link Venafi To Cloud

To add Venafi as the *Trust Provider* for a cloud

1. Navigate to Infrastructure > Clouds
2. Select Cloud
3. Select *EDIT*
4. Under *Advanced Options* select the Venafi integration from the *TRUST PROVIDER* dropdown
5. Select *SAVE CHANGE*

Load Balancers

AzureLB

Add Azure Load Balancer

1. Navigate to Infrastructure -> Load Balancers
2. Select + *ADD*
3. Select *Azure Load Balancer*
4. Fill in the following:
 - CLOUD** Select the Cloud the Load Balancer will be available for
 - NAME** Name of the Load Balancer in Morpheus
 - DESCRIPTION** Identifying information displayed on the Load Balancer list page.
 - VISIBILITY** Define Multi-Tenant permissions
 - RESOURCE GROUP** Select the Resource Group the Load Balancer will be linked to
5. Save changes

F5 Load Balancers

Add F5 Load Balancer

To add a F5 Load Balancer Integration:

1. Navigate to *Infrastructure -> Load Balancers*
2. Select + *ADD*
3. Select *F5 BigIP*
4. Fill in the following:
 - GROUP** Select the Group the Load Balancer will be available for
 - CLOUD** Select the Cloud the Load Balancer will be available for
 - NAME** Name of the Load Balancer in Morpheus

DESCRIPTION Identifying information displayed on the Load Balancer list page.

VISIBILITY Define Multi-Tenant permissions

API HOST IP or resolvable hostname url.

API PORT Typically 8443

USERNAME API user

PASSWORD API user password

MANAGEMENT URL Example: `https://10.30.20.31:8443/xui/`

Advanced Options (optional)

- VIRTUAL NAME
- POOL NAME
- SERVER NAME

5. Save Changes

Virtual Servers

Instances attached to an F5 will be listed in the Virtual servers tab. Virtual servers can also be manually added in this section.

Add Virtual Server

1. Navigate to *Infrastructure -> Load Balancers*
2. Select F5 Integration name to drill into the detail page
3. Select + *ADD* in the VIRTUAL SERVERS tab
4. Fill in the following:
 - **NAME** Name of the Virtual Server in Morpheus
 - **DESCRIPTION** Description of the Virtual Server in Morpheus
 - **Enabled** Uncheck to keep the configuration but disable F5 availability in Morpheus
 - **VIP TYPE**
 - Standard
 - Forwarding (Layer 2)
 - Forwarding (IP)
 - Performance (HTTP)
 - Performance (Layer 4)
 - Stateless
 - Reject
 - DHCP
 - Internal
 - Message Routing

- **VIP HOSTNAME** Enter Hostname of the VIP (optional)
- **VIP ADDRESS** Enter IP address for the VIP
- **VIP PORT** Enter port used for the VIP
- **SOURCE ADDRESS** Enter Virtual Server source address
- **PROTOCOL** tcp, udp, or sctp
- **PROFILES** Search for and select from available PROFILES
- **POLICIES** Search for and select from available POLICIES
- **IRULES** Search for and select from available RULE SCRIPTS
- **PERSISTENCE**
 - cookie
 - dest-addr
 - global-settings
 - hash
 - msrdp
 - sip
 - source-addr
 - ssl
 - universal
- **DEFAULT POOL** Select from available POOLS

5. Select *SAVE CHANGES*

Policies

Policies will be synced and listed in the Policies tab. These policies will be available options when creating Virtual Servers.

Pools

Create Pool

NAME Name of the POOL in Morpheus

DESCRIPTION Description of the POOL in Morpheus

BALANCE MODE

- Round Robin
- Least Connections

SERVICE PORT Specify SERVICE PORT for the POOL

MEMBERS Search for and select from available NODES

MONITORS Search for and select from available Monitors

Profiles

SSL Profiles are synced and will be created when an SSL Certificate is assigned in the Load balancer section when provisioning or editing a Load balancer on an Instance.

Monitors

Create Monitor

NAME Name of the MONITOR in Morpheus

DESCRIPTION Description of the MONITOR in Morpheus

PARENT MONITOR Select from available MONITORS

DESTINATION Specify Destination, such a *:443. Default is *:*

INTERVAL Specify Monitor Interval. Default is 5

TIMEOUT Specify Monitor Timeout. Default is 15

MONITOR CONFIG Enter monitor config.

Nodes

Create Node

NAME Name of the NODE in Morpheus

DESCRIPTION Description of the NODE in Morpheus

ADDRESS Enter node address

MONITOR Select from available MONITORS

SERVICE PORT Specify SERVICE PORT for the NODE

Rule Scripts

Rule Scripts will be synced and listed in the RULE SCRIPTS tab. These rules will be available options when creating Virtual Servers.

Citrix NetScaler



Add NetScaler Integration

To add a NetScaler Load Balancer Integration:

1. Navigate to *Infrastructure -> Load Balancers*
2. Select + *ADD*
3. Select *Citrix NetScaler*
4. Fill in the following:

GROUP * Select the Group the Load Balancer will be available for.

CLOUD * Select the Cloud the Load Balancer will be available for.

NAME * Name of the Load Balancer in Morpheus.

DESCRIPTION Identifying information displayed on the Load Balancer list page.

VISIBILITY

Define Tenant Visibility

- Public: Available to all Tenants.
- Private: Only available to specified Tenant.

Tenant If Visibility is set to private, define the Tenant the Load Balancer will be available in.

API URL *

URL of the NetScaler API

- Example: <http://10.30.21.55>

API PORT *

NetScaler API port

- Example: 80

USERNAME * NetScaler service account username

PASSWORD * NetScaler service account password

VIRTUAL NAME

Naming Pattern for new NetScaler Virtual Servers

- If blank, defaults to `morph_lb_${loadBalancer.id}`

SERVICE NAME

Naming Pattern for new NetScaler Services

- If blank, defaults to `morph_service_${container.id}`

SERVER NAME

Naming Pattern for new NetScaler Servers

- If blank, defaults to `morph_server_${server.id}`

Add Load Balancer to Instance

Load Balancers can be added to Instances during Provisioning or to existing Instances. For Load Balancer settings to appear during provisioning, or for the scale tab to be available on an Instance, the instances Node Type must have a LB port defined.

Note: For Load Balancer settings to appear during provisioning, or for the scale tab to be available on an Instance, the instances Node Type must have a LB port defined.

Add Load Balancer during Provisioning

In the Instance Provisioning wizard, Load Balancers can be configured in the Automation -> Load Balancer section.

1. Navigate to *Provisioning* -> *Instances*.
2. Select + *ADD*.
3. Select an Instance Type that supports scaling. (ENABLE SCALING (HORIZONTAL) flagged on Instance Type configuration)
4. Proceed with Instance configuration to the Automation section.
5. Fill in the following:

VIP ADDRESS

Define IP Address for the Virtual Server

- Example: 10.30.23.191

VIP PORT

Define port for the Virtual Server

- Example: 80

VIP HOSTNAME

Define hostname that will resolve to the VIP IP.

- Example: jwDemoHaApp59.den.example.com

VIRTUAL SERVICE NAME Define name for the Virtual Service. Defaults to `${instance.name}`

BALANCE MODE

Specify balance mode for the VIP

- Least Connections
- Round Robin

STICKY MODE

Specify sticky session options for the VIP

- Source IP
- Cookie

SHARED VIP ADDRESS Select if VIP is shared, then enter DIRECT VIP ADDRESS

SSL CERT

SSL Certificate that will be applied to the VIP.

- No SSL
- Select existing Certificate from Infrastructure -> Keys & Certs or from a Trust Provider Integration.

USE EXTERNAL ADDRESS FOR BACKEND NODES

- Select if traffic from LB to Backend Nodes needs to be sent to the External Addresses, or leave deselected to use Internal Addresses for Backed Nodes.
6. Optionally configure auto-scaling configuration in the `Scale` section
 7. Select *NEXT* and provision the Instance.

After all nodes in the Instance are provisioned, the LB configuration will be added to the Instance and Virtual Servers, Services and Servers will be created and configured on the NetScaler. The Load Balancer settings and status will be visible in the Instance details page **LOAD BALANCER** section, with additional details, links, and configurations options available in the **SCALE** tab.

Logs

LogRhythm

Adding LogRhythm Integration

1. Navigate to Administration -> Logs
2. Expand the LogRhythm section
3. Enable the integration
4. Fill in the following:
 - Enabled** Enable the LogRhythm integration
 - Host** IP or Hostname of the LogRhythm server.
 - Port** Port configured to access the LogRhythm server .
5. *SAVE*

Splunk

Overview

The Morpheus Splunk Integration allows forwarding logs from managed Linux hosts and vm's to a target Splunk listener by changing the rsyslogd config on linux vm's to point to Splunk forwarders. The logs will be forwarded from the clients, not from the Morpheus Appliance.

Adding Splunk Integration

1. Add a syslog listener configuration in Splunk.
2. Navigate to Administration > Settings > Logs
3. Expand the Splunk section
4. Enable the integration
5. Fill in the following:

Enabled Enable the Splunk integration

Host IP or Hostname of the Splunk server.

Port Port configured to access the Splunk server.

6. *SAVE*

Once added, syslogs from managed Linux hosts and vm's will be forwards from the clients to the target Splunk listener.

Syslog

Adding Syslog Integration

1. Navigate to Administration -> Logs
2. Expand the Morpheus logging section
3. Add the Syslog forwarding rules
4. Select *QUICK ADD*

Monitoring

ServiceNow Monitoring Integration

Note: A ServiceNow Integration must be already configured in Administration -> Integrations to enable the ServiceNow Monitoring Integration. Refer to the [ServiceNow](#) configuration guide for more information.

Enabled Enables the ServiceNow Monitoring Integration

Integration Select from a ServiceNow Integration added in Administration -> Integrations

New Incident Action The Service Now action to take when a Morpheus incident is created.

Close Incident Action The Service Now action to take when a Morpheus incident is closed.

Incident Severity Mapping

Morpheus Severity	ServiceNow Impact
Info	Low/Medium/High
Warning	Low/Medium/High
Critical	Low/Medium/High

AppDynamics

AppDynamics is a very powerful performance and application monitoring tool. It features advanced correlation features and profiling capabilities for a very wide range of application platforms including native Docker support. Due to the level of capabilities of AppDynamics there are more required settings to integrate it with Morpheus.

Configuring The AppDynamics Integration

1. Navigate to Administration > Monitoring
2. Expand the AppDynamics section
3. Toggle the Enable slider
4. Fill out desired fields
5. Save

Once saved, all hosts will automatically be configured to install the AppDynamics agent.

AppDynamics is capable of being run as a paid SaaS based service as well as an on premise installation and Morpheus supports both configurations. Most input fields related to connecting to AppDynamics provide helpful tips as to what information exactly needs provided and where to acquire it.

NewRelic

Configuring The NewRelic Integration

1. Navigate to Administration > Monitoring
2. Expand the NewRelics section
3. Toggle the Enable slider
4. Enter License Key to be used when installing the New Relic agent in order for the agent to report data to your New Relic account

Note: The License Key is the 40-character hexadecimal string that New Relic provides when you sign up for your account.

Networking

Infoblox

Features

- Network Pools synchronization
- DNS Zone & Zone record synchronization
- Host Record synchronization
- Total & Free IP status bar for networks
- Network Grid and List view with IP Status and records, date and user tracking

- Automatic and manual IP Reservations, DNS A/PTR record creation and deletion
- Use script variables like `<%= variableX %>` for evaluation of the key data in extended attributes

Adding Infoblox Integration

Note: Making full use of the Morpheus Infoblox integration requires credentials for an Infoblox user account with API access granted, access to list the pools and zones you wish to work with, and rights to create and destroy records. See Infoblox documentation for more information on user rights administration in that product.

1. Navigate to Infrastructure - Network - Integrations
2. Select + *ADD* -> IPAM -> Infoblox
3. Enter the following:

ADD IPAM INTEGRATION

NAME

☒ ENABLED

URL

https://x.x.x.x/wapi/v2.2.1

USERNAME

PASSWORD

THROTTLE RATE

0

ms

☒ DISABLE SSL SNI VERIFICATION

☐ INVENTORY EXISTING

NETWORK FILTER

ZONE FILTER

TENANT MATCH ATTRIBUTE

IP MODE

Static IPs

EXTRA ATTRIBUTES

Accepts a JSON input of custom attributes that can be saved on Host Record in Infoblox. These Must be first defined as extra attributes in Infoblox and values can be injected for the user creating the record and the date of assignment. The available injectable attributes are: `userId` , `username` , and `dateCreated` . They can be injected with `<%= %>` .

SAVE CHANGES

NAME Name of the Integration in Morpheus

Enabled Deselect to disable the Integration

URL Infoblox wapi url. Example: <https://x.x.x.x/wapi/v2.2.1>

USERNAME Infoblox user username

PASSWORD Infoblox user password

THROTTLE RATE In milliseconds (ms)

DISABLE SSL SNI VERIFICATION Leave selected to disable SSL SNI Verification

INVENTORY EXISTING Mark this option to inventory existing network pools from Infoblox

NETWORK FILTER Filter which networks are synced into Morpheus. Example: Network Filter: [network_view=default&*Building=work]

ZONE FILTER Filter terms for Zone Records

TENANT MATCH ATTRIBUTE This can be set to the name of the extended attribute in Infoblox where Morpheus will check for the id of a morpheus tenant. This allows for setting the tenant's Morpheus id to an extended attribute field on a network view or network in Infoblox, and when the network or view is discovered by morpheus, it will be auto assigned to the right tenant.

IP MODE Static IPs or DHCP Reservations

EXTRA ATTRIBUTES Accepts a JSON input of custom attributes that can be saved on host records in Infoblox

4. Select *SAVE CHANGES*

Upon save the Infoblox IPAM integration will be created and the following will sync:

- Infoblox networks will be synced in and populate in the *Infrastructure - Network - IP Pools* tab and in the Infoblox detail page under the *NETWORK POOLS* tab
- Host Records will sync and populate in the Network Pool detail view (select an IP Pool name to view)
- DNS Zones will sync and populate under *Infrastructure - Network - Domains* and in the Infoblox detail page under the *HOSTS* tab
- DNS Zone Records will sync and populate

Adding IP Pools to Networks

Morpheus can automatically assign the next available Infoblox IP in an IP/Network Pool and create the corresponding DNS records, as well as remove the records upon teardown. To enable this, add an Infoblox IP/Network Pool to the *Network Pool* section on a Network(s).

1. Navigate to *Infrastructure > Network > Networks*
2. Select a Network name and click *EDIT*
3. In the *NETWORK POOL* section, search for and select the name of the IP/Network Pool.
 - Gateway, DNS and CIDR must be populated for static/pool IP assignment
 - Select *Allow IP Override* to allow selecting between DHCP, Static entry and Pool Selection at provision time (if desired)
 - Deselect DHCP server if a DHCP server will not be used on the network (only static and/or IP Pool IP assignment)

4. Select *SAVE CHANGES*

Creating Host Records

1. Select a Network Pool from *Infrastructure > Network > IP Pools* or *Infrastructure > Network > Services > Infoblox*
2. Select + *ADD*
3. Enter the following

CREATE HOST RECORD

HOSTNAME

sample

IP ADDRESS

10.30.23.88

DOMAIN

infoblox.den.bertramlabs.com

☒ Create DNS Records

SAVE CHANGES

HOSTNAME Hostname for the record

IP ADDRESS IP address for the Host Record

DOMAIN Select an Infoblox Zone

Create DNS Records Select to create DNS A and PTR Records in Infoblox

4. Select *SAVE CHANGES*

Creating Zone Records

1. Select a Domain from *Infrastructure > Network > Domains* or *Infrastructure > Network > Services > Infoblox > Zones*
2. Select + *ADD*
3. Enter the following

CREATE ZONE RECORD

×

NAME

sample

TYPE

A

▼

CONTENT

10.30.22.89

TTL

86400

SAVE CHANGES

NAME Name for the record, such as Hostname

Type A, AAAA, CNAME, MX, NS, PTR, SOA, or TXT

CONTENT Content of the record, such as IP or A Record

TTL Time To Live value

4. Select *SAVE CHANGES*

phpIPAM

Configuration

1. Within phpIPAM dashboard, enable api in Administration > phpIPAM settings > feature settings. Toggle API switch to on and save.
2. Go to Admin > API > create API key.
3. Create unique App ID.
4. Enable read/write/admin access under **App Permissions**.
5. Under **App Security** select none.

Add phpIPAM integration to Morpheus

1. Navigate to Infrastructure > Network > Integrations
2. Select + *ADD* -> IPAM -> phpIPAM
3. Enter the following:
 - Name
 - **URL** Add /api/ to end of URL ex. http://10.30.20.196/api/
 - **App ID** From phpIPAM API Key
 - Username
 - Password
 - Enable or Disable SSL SNI Verification
 - Enter Network Filter
4. Select *SAVE IPAM INTEGRATION*

NSX-V

Overview

- SUMMARY
- TRANSPORT ZONES
- SWITCHES
- LOGICAL SWITCHES
- FIREWALL
- LOGICAL ROUTERS
- EDGE GATEWAYS

Add NSX-V Integration

1. Navigate to Infrastructure > Network > Integrations
2. Select Select + *ADD* -> VMWare NSX-V
3. Enter the following:
 - NAME** Name for the NSX Integration in Morpheus
 - API HOST** URL of NSX Manager
 - USERNAME** NSX Manager Admin Username
 - PASSWORD** NSX Manager Admin password
 - VMWARE CLOUD** Select the existing VMware cloud associated with this NSX integration
4. Select *ADD NETWORK INTEGRATION*

Once the NSX Integration is added Morpheus will sync in existing Transport Zones, Logical Switches, and Edge Gateways. New Transport Zones, Logical Switches, and Edge Gateways can be now be created.

Switches

Morpheus version 4.1.2 adds SWITCHES tab to view switches associated with the selected NSX integration. Information displayed on each switch include the following:


- NAME
- UPLINK PORT
- TYPE
- SWITCH ID
- DESCRIPTION

✓ Labs NSX

EDITDELETEACTIONS ▾

Host: https://10.30.23.6 Last Update: 01/15/2020 10:19 AM

SUMMARYTRANSPORT ZONES**SWITCHES**LOGICAL SWITCHESFIREWALLLOGICAL ROUTERSEDGE GATEWAYS

Search 

NAME	UPLINK PORT	TYPE	SWITCH ID	DESCRIPTION
Testing-dvSwitch	Uplink 1	NSX VDS	dvs-15	

Create NSX Logical Switch and Edge Gateway

Important: Prior to creating a Logical Switch and Edge Gateway, associated External VMware Networks must be configured in Morpheus. Navigate to *INFRASTRUCTURE* -> *NETWORK* and edit any Distributed Switch Groups that will be used and populate the Gateway, DNS and CIDR

1. Navigate to *INFRASTRUCTURE* -> *NETWORK*
2. Select the *SERVICES* tab
3. Select the name of NSX Integration
4. Select the *LOGICAL SWITCHES* tab
5. Select + *CREATE NSX LOGICAL SWITCH*
6. Populate the following for the Logical Switch and Edge Gateway Configurations:

Logical Switch Configuration:

NAME Name of the Logical Switch

DESCRIPTION Description of the Logical Switch

TRANSPORT ZONE Select an existing Transport Zone

CIDR Add the CIDR for the Logical Switch. Example: 10.30.28.0/24

TENANT NAME Enter Tenant name for the Logical Switch (Optional)

Edge Gateway Configuration:

HOSTNAME Enter Hostname of the Edge Gateway

SIZE Select Size of the Edge Gateway

EXTERNAL NETWORK Select the External Network for the Edge Gateway.

Important: The Gateway, DNS and CIDR must be populated on an external network for it to be selectable when creating an Edge Gateway.

IP ADDRESS Populate IP address to be assigned to the Edge Gateway

DATA STORE Select the Datastore for the Gateway

RESOURCE POOL Select the Resource Pool for the Gateway

FOLDER Select a Folder for the Edge Gateway (optional)

USERNAME Enter a Username for the Edge Gateway

PASSWORD Enter a Password for the Edge Gateway


Note: Password length must be at-least 12 characters and at-max 255 characters. It must contain mix of alphabets with both upper case and lower case, numbers and at-least one special character. Password must not contain username as substring. Character must not consecutively repeat 3 or more times.

7. Select + *CREATE*

Morpheus version 4.1.2 also extends the details we can see on existing Edge Gateways. First, to view the list of Edge Gateways, navigate to your selected NSX integration, and click on the **EDGE GATEWAYS** tab. Here you will see a list of existing Edge Gateways, including their **NAME** and **DESCRIPTION** values. To see the enhanced details view for your Edge Gateways, click on the name of a selected Edge Gateway.

labs-edge-1 EDIT DELETE

SUMMARY **FIREWALL** **DHCP** **ROUTING**

 **Name:** labs-edge-1 **Type:** NSX Edge Gateway **Provider Id:** edge-8
Date Created: 2019-11-18 21:58:09.0

▼ **INTERFACES**

POSITION	NAME	TYPE	LINK	NETWORK	IP ADDRESS	SUBNET	ENABLED
0	external0	uplink	vLAN02	vLAN02	10.30.23.240	10.30.23.240/22	Yes
1	internal1	internal	vxw-dvs-15-virtualwire-14-sid-5000-labs-net-1	vxw-dvs-15-virtualwire-14-sid-5000-labs-net-1	172.16.21.1	172.16.21.1/24	Yes

▼ **FIREWALL**

Enabled: Yes Version: 20 Default Policy: deny

The new Edge Gateway detail view includes the following tabs:

- **SUMMARY:** Includes general configuration details for the selected Edge Gateway
- **FIREWALL:** Includes firewall configuration detail and includes the ability to create rules
- **DHCP:** Includes details on IP pools
- **ROUTING:** Includes details on configured routes and includes the ability to create routes

Firewall

Morpheus version 4.1.2 adds a FIREWALL tab which allows you to view existing firewall rules as well as create new rules and groups. From the rules summary list, the following fields are displayed for each rule:




- NAME
- TYPE
- POLICY
- DIRECTION
- SOURCE
- DESTINATION
- APPLICATION

RULES

Search

Q

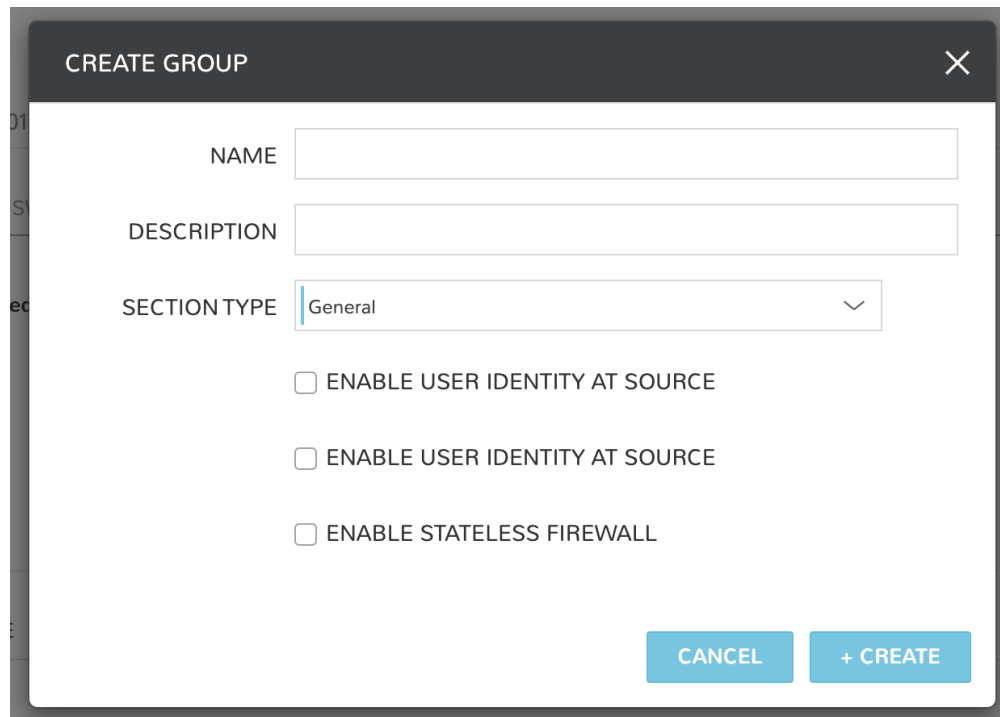
ACTIONS ▾

NAME	TYPE	POLICY	DIRECTION	SOURCE	DESTINATION	PROTOCOL	
BungeTest							
bw-test-section-1							
bw-test-section-2							
▼ Default Section Layer3							
Default Rule NDP	custom	accept	inout	Any	Any	Any	
Default Rule DHCP	custom	accept	inout	Any	Any	Any	
Default Rule	custom	accept	inout	Any	Any	Any	
TestSection							

Morpheus also allows you to create new firewall groups and new firewall rules.

To create a new group:

1. Click on the **ACTIONS** button from within the list of firewall rules
2. Click “Create Group”



CREATE GROUP

NAME

DESCRIPTION

SECTION TYPE

☐ ENABLE USER IDENTITY AT SOURCE

☐ ENABLE USER IDENTITY AT SOURCE

☐ ENABLE STATELESS FIREWALL

CANCEL + CREATE

To create a new rule:

1. Click on the *ACTIONS* button from within the list of firewall rules
2. Click “Create Rule”

NAME	
DESCRIPTION	
DIRECTION	ingress
SOURCE TYPE	Any
SOURCE	
DESTINATION TYPE	Any
DESTINATION	
SERVICE TYPE	Custom Rule
PORT RANGE	

CANCEL + CREATE

accept inout Any Any

Logical Routers

Morpheus version 4.1.2 adds a Logical Routers section to the NSX integration, including the ability to view and create new logical routers. From the LOGICAL ROUTERS tab, a list of logical routers associated with your selected integration is shown. Values displayed for each logical router include the following:

- STATUS
- NAME
- DESCRIPTION

To create a new logical router:

1. Navigate to the LOGICAL ROUTERS tab for the chosen integration
2. Click on + *CREATE NSX LOGICAL ROUTER*
3. Complete the presented modal
4. Click *ADD NETWORK ROUTER*

ADD NETWORK ROUTER

NAME

DESCRIPTION

☒ ENABLED

HOSTNAME

TENANT NAME

SIZE compact

DATA STORE Select

RESOURCE POOL Select

FOLDER Select

MANAGEMENT Select

NSX-T

Overview

VMware NSX-T offers network virtualization allowing for creation and management of software-based virtual networks in an efficient and programmatic way. Morpheus offers a full-featured integration with NSX-T, exposing its networking abstractions in the following sections of the Morpheus NSX-T integration:

- SUMMARY
- TRANSPORT ZONES
- SEGMENTS
- FIREWALL
- TIER-1 GATEWAYS
- TIER-0 GATEWAYS

This guide goes through the process of integrating an existing NSX-T installation with Morpheus and working with the associated objects synced in with the integration. For more on installing NSX-T and an overview of its concepts, please review the [NSX-T overview documentation](#) provided by VMware.

Add NSX-T Integration to Morpheus

1. Navigate to Infrastructure > Network > Integrations
2. Select Select + *ADD* > VMWare NSX-T
3. Enter the following:
 - **NAME:** Name for the NSX Integration in Morpheus
 - **API HOST:** URL of the NSX Manager (ex. <https://x.x.x.x/api>)
 - **USERNAME:** NSX Manager Admin Username
 - **PASSWORD:** NSX Manager Admin password
 - **VMWARE CLOUD:** Select the existing VMware cloud associated with this NSX integration
4. Select *ADD NETWORK INTEGRATION*

Once the NSX Integration is added Morpheus will sync in existing Transport Zones, Segments, firewall groups and rules, and Gateways. We can also manage these synced items from within Morpheus UI, including the ability to create, edit, and delete them.

Summary View

The SUMMARY tab contains the default view when accessing an NSX-T integration. From the summary view we can see the health status of the NSX-T server, and details about interfaces and group status.

Transport Zones

Access Transport Zones by navigating to Infrastructure > Networks > Integrations > (Your NSX-T Integration) > Transport Zones tab. We can delete Transport Zones by clicking on the trash can icon to the far right of each list item. The default view lists each Transport Zone and provides the following information about them:

- **NAME:** The given name for the Transport Zone
- **DESCRIPTION:** A given description value (if available)
- **TRAFFIC TYPE:** “Overlay” or “VLAN”
- **N-VDS NAME:** The name of the NSX-managed virtual distributed switch
- **STATUS:** An icon indicating the current status of the Transport Zone
- **HOST MEMBERSHIP CRITERIA:** “Standard” or “Enhanced Datapath”

Creating NSX-T Transport Zones

1. Navigate to *Infrastructure -> Network*
2. Select the *Integrations* tab
3. Select the name of NSX-T integration
4. Select the *Transport Zones* tab
5. Select + *CREATE NSX-T TRANSPORT ZONE*
6. After completing the required fields and any desired optional fields, click + *CREATE*

Segments

Access Segments by navigating to Infrastructure > Networks > Integrations > (Your NSX-T Integration) > Segments tab. We can delete Segments by clicking on the trash can icon to the far right of each list item or edit them by clicking on the pencil icon. The default view lists each Segment and provides the following information about them:

- **STATUS:** An icon indicating the current status of the Transport Zone
- **NAME:** The given name for the Segment
- **TRAFFIC TYPE:** “Overlay” or “VLAN”
- **N-VDS NAME:** The name of the NSX-managed virtual distributed switch
- **STATUS:** An icon indicating the current status of the Transport Zone
- **HOST MEMBERSHIP CRITERIA:** “Standard” or “Enhanced Datapath”

Creating NSX-T Segments

1. Navigate to *Infrastructure -> Network*
2. Select the *Integrations* tab
3. Select the name of NSX-T integration
4. Select the *Segments* tab
5. Select + *CREATE NSX-T SEGMENT*
6. Complete the fields in the CREATE NETWORK modal
7. Click *SAVE CHANGES*

Note: NSX-T Segments can be scoped to specific Groups and Tenants when creating or editing the Segment.

Firewall

Access firewalls by navigating to Infrastructure > Networks > Integrations > (Your NSX-T Integration) > Firewall tab. We can delete firewall groups by clicking on the trash can icon at the end of each row. Additionally each group can be expanded (when applicable) to reveal the firewall rules within the group. Individual rules can be edited or deleted by clicking on pencil or trash can icon at the end of the row. The default view lists each Segment and provides the following information about them:

- **NAME:** The name of the rule or group within Morpheus
- **CATEGORY:** “Ethernet”, “Emergency”, “Infrastructure”, “Environment”, or “Application”
- **ENABLED:** Applies only to rules, the rule is enabled when the check mark is present
- **POLICY:** Applies only to rules, “Allow”, “Drop”, or “Reject”
- **DIRECTION:** Applies only to rules, “In”, “Out”, or “In-Out”
- **SOURCE:** Applies only to rules, “Any”, by default
- **DESTINATION:** Applies only to rules, “Any”, by default
- **APPLICATION:** Applies only to rules, “Any”, by default

Creating NSX-T Firewall Groups

1. Navigate to *Infrastructure -> Network*
2. Select the *Integrations* tab
3. Select the name of NSX-T integration
4. Select the *Firewall* tab
5. Select *ACTIONS*
6. Select *Create Group*
7. Complete the fields in the CREATE GROUP modal:
 - **NAME:** The name of the rule or group within Morpheus
 - **DESCRIPTION:** An optional description value for the group
 - **CATEGORY:** “Ethernet”, “Emergency”, “Infrastructure”, “Environment”, or “Application”
8. Click *SAVE CHANGES*

Creating NSX-T Firewall Rules

1. Navigate to *Infrastructure -> Network*
2. Select the *Integrations* tab
3. Select the name of NSX-T integration
4. Select the *Firewall* tab
5. Select *ACTIONS*
6. Select *Create Rule*
7. Complete the fields in the CREATE RULE modal:
 - **NAME:** The name of the rule or group within Morpheus
 - **DESCRIPTION:** An optional description value for the rules
 - **ENABLED:** Rule is enforced when checked
 - **DIRECTION:** “In”, “Out”, or “In-Out”
 - **SOURCES:** “Any”, by default
 - **DESTINATIONS:** “Any”, by default
 - **SERVICES:** “Any”, by default
 - **PROFILES:** “Any”, by default
 - **SCOPES:** “Any”, by default
 - **POLICY:** “Allow”, “Drop”, or “Reject”
8. Click + *CREATE*

Tier-1 Gateways

Access Tier-1 Gateways by navigating to Infrastructure > Networks > Integrations > (Your NSX-T Integration) > Tier-1 Gateways tab. We can edit a Gateway by clicking the pencil icon in each row or delete the Gateway by clicking on the trash can icon. The default page for Tier-1 Gateways displays the following information on each:

- **STATUS:** An icon indicating the status of each gateway
- **NAME:** The given name of the gateway
- **DESCRIPTION:** An optional description value for the gateway

Creating Tier-1 Gateways

1. Navigate to *Infrastructure -> Network*
2. Select the *Integrations* tab
3. Select the name of NSX-T integration
4. Select the *Tier-1 Gateways* tab
5. Select + *CREATE NSX-T TIER-1 GATEWAY*
6. Complete the fields in the ADD NETWORK ROUTER modal:
 - **GROUP:** If desired, scope the Tier-1 Gateway to a Morpheus Group
 - **NAME:** The name of the Tier-1 Gateway within Morpheus
 - **ENABLED:** Tier-1 Gateway is available for use when checked
 - **TIER-0 Gateway:** Select an existing and enabled Tier-0 Gateway
 - **EDGE CLUSTER:** Select an existing Edge Cluster
7. Make selections as needed in the “Route Advertisement” section
8. Click *ADD NETWORK ROUTER*

Tier-0 Gateways

Access Tier-0 Gateways by navigating to Infrastructure > Networks > Integrations > (Your NSX-T Integration) > Tier-0 Gateways tab. We can edit a Gateway by clicking the pencil icon in each row or delete the Gateway by clicking on the trash can icon. The default page for Tier-0 Gateways displays the following information on each:

- **STATUS:** An icon indicating the status of each gateway
- **NAME:** The given name of the gateway
- **DESCRIPTION:** An optional description value for the gateway

Creating Tier-0 Gateways

1. Navigate to *Infrastructure* -> *Network*
2. Select the *Integrations* tab
3. Select the name of NSX-T integration
4. Select the *Tier-0 Gateways* tab
5. Select + *CREATE NSX-T TIER-0 GATEWAY*
6. Complete the fields in the *ADD NETWORK ROUTER* modal:
 - **GROUP:** If desired, scope the Tier-0 Gateway to a Morpheus Group
 - **NAME:** The name of the Tier-0 Gateway within Morpheus
 - **ENABLED:** Tier-1 Gateway is available for use when checked
 - **HA MODE:** “Active Active” or “Active Standby”
 - **EDGE CLUSTER:** Select an existing Edge Cluster
7. Make selections as needed in the routing and BGP sections
8. Click *ADD NETWORK ROUTER*

Cisco ACI

Overview

The screenshot displays the Cisco ACI Sandbox interface. At the top, there's a header with the Cisco ACI logo and the text "ACI Sandbox". Below this, a navigation bar contains tabs: SUMMARY, TENANTS, APPLICATION PROFILES (which is selected), ENDPOINT GROUPS, CONTEXTS, BRIDGE DOMAINS, FILTERS, and CONTRACTS. The main content area is titled "APPLICATION PROFILES" and includes a search bar and a "+ CREATE ACI APP PROFILE" button. A table lists existing application profiles with columns for NAME, DESCRIPTION, and DN. Each row also has a "Remove" link.

NAME	DESCRIPTION	DN	
A1_AP1		uni/tn-A1/ap-A1_AP1	Remove
A2_AP1		uni/tn-A2/ap-A2_AP1	Remove
access		uni/tn-infra/ap-access	Remove
apBankPac		uni/tn-BankPac/ap-apBankPac	Remove
appProfile		uni/tn-zxdf/ap-appProfile	Remove

Add ACI as a network and security integration. Inventory your existing ACI configurations. Create networks, bridge domains, application profiles, tenants, endpoint groups, contexts, filters and contracts. Provision instances into new endpoint groups and define security groups that apply contracts on provision.

From Morpheus below can be created:

- Tenants
- ANP's
- EPG's

- Contexts
- Bridge Domains
- Filters
- Contracts

Note: Morpheus to ACI Sync Job Schedule: Every 5 minutes

Note: Morpheus connects to ACI APIC over port 443

Add Network Integration

1. Navigate to Infrastructure -> Networks -> Integrations
2. Select **+ADD** -> *Networking* -> Cisco ACI
3. Populate the following:

NAME

ACI Integration Name/Label in Morpheus This is unique to Morpheus and not part of authentication

URL ACI fabric url, eg `https://apicdc.company.com`

USERNAME ACI aaaUser name attribute

PASSWORD ACI aaaUser pwd attribute

TENANT Populates upon authentication, tenant selection not required

4. Select **ADD NETWORK INTEGRATION**

Configure Cloud Network Mode

For your ACI Integration to be available during provisioning, ACI needs to be defined on a Cloud or multiple Clouds
NETWORK MODE advanced options.

1. Select an existing VMware vCenter Cloud
2. Select *EDIT*
3. Expand the *Advanced Options* section
4. Select ACI Integration in **NETWORK MODE** dropdown
5. Select *SAVE*

Instance Provisioning

▼ ACI Options

ENDPOINT GROUP

APP PROFILE

▼ ACI Security

CONSUMES

PROVIDES

Once ACI is integration to a cloud, it can be used during instance provisioning:

1. From the EPG drop down, either an existing EPG can be selected or a new one can be created. It is the same for ANP, either create a new one or choose an existing.
2. Under ACI security consumes and provides, contracts can be searched when you enter a name. When the provisioning wizard is completed, it will provision the instance and apply the ACI options and Security. This can be viewed under the instance page, or via REST API and CLI.

Blueprint Configuration

▼ ACI Options

ENDPOINT GROUP

APP PROFILE

▼ ACI Security

CONSUMES

PROVIDES

- In a Blueprint, you can define the ANP and EPG of each Tier
- Variables can be used for EPG and ANP names.
- This could be useful to create blueprints for dev testing to isolate from prod networks.
- This can be hybrid based on the VMM domains in APIC.

Bluecat

Overview

Morpheus integrates with Bluecat IPAM to scope pools to networks for static IP assignment from Bluecat to your Morpheus Instances.

Adding Bluecat to Morpheus

1. Navigate to `Infrastructure > Network > Integrations`
2. Click + *ADD*
3. Select *Bluecat*
4. Enter in the following information

Name Name of the Bluecat Integration in Morpheus

Enabled Uncheck to disable sync with the Bluecat endpoint

URL URL of the Bluecat server, ex: `http://10.30.20.10`

Username Username of Bluecat API User. API and root level propagating read access required, read/write access required for target Networks and Domains.

Password Bluecat User password

Network Filter Optionally enter the id of a config, block or network, or comma separated combination of configs, blocks and/or networks.

5. Click *SAVE CHANGES*

The Bluecat Integration will be saved, IP pools will sync in and populate under `Infrastructure > Network > IP Pools`, and Domain will populate in `Infrastructure > Network > Domains`. Pools and Domains can also be found in the Bluecat Integration details page, which can be accessed by clicking on the name of the added Bluecat Integration in `Infrastructure > Network > Services`.

Important: *Quick Deployments* must be enabled in Bluecat for Morpheus to create instantly available DNS records when using Bluecat DNS.

Adding IP Pools to Networks

Morpheus can automatically assign the next available Bluecat IP in an IP/Network Pool and create the corresponding DNS records, as well as remove the records upon teardown. To enable this, add an Bluecat IP/Network Pool to the *Network Pool* section on a Network(s).

1. Navigate to `Infrastructure - Network- Networks`
2. Select a Network name and *EDIT*, or select *ACTIONS - Edit*
3. In the *NETWORK POOL* section, search for and select the name of the IP/Network Pool.
 - Gateway, DNS and CIDR must be populated for static/pool IP assignment
 - Select *Allow IP Override* to allow selecting between DHCP, Static entry and Pool Selection at provision time

- Deselect DHCP server if a DHCP server will not be used on the network (only static and/or IP Pool IP assignment)
4. Select *SAVE CHANGES*

SolarWinds IPAM

Features

- Automate static IP assignment across environments using Solarwinds IPAM
- Network pool sync: Network pools can be set on networks in Morpheus for automated IP allocation and record creation
- Optional Network Pool allocation and record sync. `Inventory Existing` option syncs all individual IP records and corresponding status. Inventory is not required for provisioning
- Grid and list displays with IP record overlays and color coding for static, available, reserved and transient status
- Manual IP Host record creation from Network Pool detail pages

Adding SolarWinds to Morpheus

1. Navigate to `Infrastructure > Network > Integrations`
2. Click + *ADD*
3. Select *SolarWinds*
4. Enter in the following information

Name Name of the SolarWinds Integration in Morpheus

Enabled Deselect to disable sync with the SolarWinds endpoint

URL URL of the SolarWinds server, ex: `http://10.30.20.10:17778`

Username Username of SolarWinds API User. API and root level propagating read access required, read/write access required for target networks and domains.

Password SolarWinds User password

5. Click *SAVE CHANGES*

Consuming SolarWinds in Morpheus

On saving your new integration, SolarWinds networks will be synced and can be viewed by navigating to `Infrastructure > Network > IP POOLS`. They're also viewable from the detail section of the SolarWinds integration at `Infrastructure > Network > INTEGRATIONS > (your SolarWinds integration) > NETWORK POOLS`.

Host records can also be viewed here by clicking on the name of a SolarWinds network.

Note: Morpheus SolarWinds integration does not support zone record syncing despite the presence of the ZONES tab on the integration detail page. This is a UI feature carried over from other networking integrations and is not supported at this time.

Morpheus can automatically assign the next available SolarWinds IP in an IP/Network Pool and create the corresponding DNS records. Morpheus will also remove the records upon teardown. To enable this, add an SolarWinds IP/Network Pool to the *Network Pool* section on a Network(s).

- ### 1.3. Hide Blueprint fields

- Select **ALLOW IP OVERRIDE**, if desired, to allow selecting between DHCP, static IP address entry, and pool address selection at provision time
 - Deselect DHCP server if a DHCP server will not be used on the network (only static and/or IP Pool IP assignment)
4. Select *SAVE CHANGES*

Creating Host Records

1. Select a Network Pool from Infrastructure > Network > IP POOLS
2. Click + *ADD*
3. Enter the following
HOSTNAME Hostname for the record
IP ADDRESS IP address for the Host Record
4. Select *SAVE CHANGES*



CREATE HOST RECORD

HOSTNAME

IP ADDRESS

SAVE CHANGES

Unisys Stealth

Introduction

Unisys Stealth is a network security tool for safeguarding sensitive information across shared networks. By creating pre-defined communities of interest (COIs) and curating user access to these groups, the need to create separate networks for handling restricted data is reduced.

Morpheus includes a full integration with Stealth allowing administrators to create and manage COIs, work with configurations and roles, and provision new endpoints into COIs.

Stealth Concepts

- **Communities of Interest (COIs):** A collection of endpoints cryptographically separated so they can only communicate to each other
- **Endpoints:** Any system or device with a Stealth agent
- **Configuration:** Tells the endpoints which authorization methods and services to use for obtaining COI membership
- **Role:** Defines COI membership. Users and groups are assigned to a role which grants access to that role's COIs
- **Filters:** Customizes Stealth communication. More specifically, filters constrain Stealth communication to specific addresses, ports, or protocols and allow Stealth endpoints to communicate with non-Stealth endpoints

Integrating Stealth with Morpheus

1. Navigate to Infrastructure > Network > Integrations
2. Click + *ADD*
3. Complete the following fields:
 - **NAME:** A name for this Stealth integration in Morpheus
 - **API HOST:** The address for the server hosting Stealth (ex: <https://x.x.x.x:8448>)
 - **USERNAME:** A username for the portaladmin, or the user who logs into the web console
 - **PASSWORD:** A password for the account
 - **MANAGER USERNAME:** A username for the manager account
 - **MANAGER PASSWORD:** A password for the manager account
4. Click *ADD SECURITY INTEGRATION*

ADD SECURITY INTEGRATION X

NAME

API HOST

USERNAME

PASSWORD

MANAGER USERNAME

MANAGER PASSWORD

► Advanced Options

ADD SECURITY INTEGRATION

Summary View

The default view when accessing a Stealth integration in Morpheus is the Summary view. In addition to basic information about the Stealth server itself, we can see system status, license and service information.

Networks > Security Services > Labs Stealth

✓ Labs Stealth EDIT DELETE ACTIONS

Last Update: 06/10/2020 09:47 AM

Summary Endpoints Configurations Roles COIs Filters

STEALTH Name: Labs Stealth Type: Stealth Date Created: 2020-05-08 18:50:43.0

▼ STEALTH SYSTEM STATUS

STATUS	NAME	LOCATION	ROLES
!			License, Authorization
!			Management, License, Authorization
✓			License, Authorization

▼ STEALTH LICENSE

STATUS	TYPE	USED	AVAILABLE	MAX
✓	Client	0	25	25
✓	Server	2	23	25
✓	Mobile	0	25	25
✓	VGD	0	25	25
✓	AWSMarketplace	0	0	0

Endpoints

Endpoints are any system or device with a Stealth agent. Stealth endpoints can be provisioned in Morpheus in the same way other cloud resources are provisioned. With Stealth integrated, workloads provisioned on the appointed networks are assigned Stealth configuration and a Stealth role during the provisioning process. Based on a user's assigned roles and COIs assigned to those roles, only workloads on the appointed COIs will be visible to the user. Additionally, workloads can only see other workloads within their COI.

Note: Machines on the same network which are not Stealth-managed will be able to see and communicate with each other but will not be able to see workloads which are assigned to a Stealth COI.

Endpoints View

The endpoints view will display all available Stealth-managed resources. Endpoints are not created here but will be synced into this view as they are created (through Morpheus provisioning or outside creation). Stealth-managed endpoints can be deleted by clicking the trash can icon at the end of each row in this view.

The following fields are exposed in the endpoints list view:

- **DISPLAY NAME**
- **NAME**
- **DESCRIPTION**

Configurations

Configurations in Stealth are the top-level construct and house COIs, roles, groups, users and endpoints. Your Stealth integration will include at least one configuration but often they will include more.

Configurations are primarily created and managed from the Stealth console but we can opt to remove them from Morpheus by clicking the trash can icon at the end of each row on the configurations list page. Configurations are selected along with a Stealth role at provision time in Morpheus.

Configurations View

The following fields are exposed in the configurations list view:

- **NAME:** The name of the configuration
- **DESCRIPTION:** A description value for the configuration

Roles

Users are placed into roles which are associated with COIs. A user's role(s) determine which COIs he or she will be able to see and access. Roles are synced into Morpheus once the integration process is complete and can be viewed in the Roles list view. Roles can also be created from the Morpheus integration as described later in this section.

Roles View

The following fields are exposed in the roles list view:

- **NAME:** The name of the role
- **DESCRIPTION:** A description value for the role

Note: More detail on each item in the roles list can be revealed by clicking on the (i) icon in each row, including the COIs associated with the role.

Creating Stealth Roles

1. Navigate to Infrastructure > Network > Integrations > (Your Stealth integration) > Roles
2. Click + *CREATE ROLE*
3. Complete the following fields:
 - **NAME:** The name for the new role
 - **DESCRIPTION:** A description value for the new role
 - **CONFIGURATION:** Select an existing Stealth configuration to associate with the role
 - **ROLE TYPE:** Identifies how the role is used. Can be Global (for roles used to isolate endpoints and users), Service (for roles used by endpoints in service mode to access an authorization service) or Work-Group (for roles used by endpoints in normal operation)
 - **FILTER SET:** Choose a filter set to apply to the role to allow or deny clear text communication with non-Stealth-managed endpoints
 - **COIs:** Select the COIs to be associated with the role
 - **PROVISION CHANGES:**
4. Click *ADD ROLE*

ADD ROLE

NAME

DESCRIPTION

CONFIGURATION

Management

ROLE TYPE

Global

FILTER SET

Labs Networking

COIS

Select

+

☐ PROVISION CHANGES

ADD ROLE

COIs (Communities of Interest)

COIs exist within configurations and create a logical separation between endpoints in separate COIs. Communication between endpoints in the COI is encrypted and those outside the COI are unable to see or access endpoints despite being on the same network.

On completing the integration, Morpheus will sync in existing COIs. COIs can also be created from Morpheus UI which is described later in this section. COIs are deleted by clicking on the trash can icon at the end of each row in the COIs list view.

COIs View

The following fields are exposed in the roles list view:

- **NAME:** The name of the COI
- **DESCRIPTION:** A description value for the COI

Creating Stealth COIs

1. Navigate to Infrastructure > Network > Integrations > (Your Stealth integration) > COIs
2. Click + *CREATE COI*
3. Complete the following fields:
 - **NAME:** The name for the new COI
 - **DESCRIPTION:** A description value for the new COI
 - **TYPE:** Workgroup or Service
 - **DIRECTION:** Default (enables COI to accept inbound and initiate outbound tunnels), Initiate (restricts the COI to only initiate outbound tunnels), or Accept (restricts the COI to only accept inbound tunnels)
4. Click *CREATE COI*

CREATE COI

NAME

DESCRIPTION

TYPE

DIRECTION

CREATE COI

Filters

Filters customize Stealth communication. More specifically, filters constrain Stealth communication to specific addresses, ports, or protocols and allow Stealth endpoints to communicate with non-Stealth endpoints.

Filters are synced into Morpheus when integrating with Stealth and are viewable from the filters list view. They are created and managed from within the Stealth console itself.

When accessing the filters list view, all filter sets are displayed. Each filter set can be expanded to view the individual filters within. Information on each filter is displayed once the filter set has been expanded to reveal the individual filters.

Provisioning with Stealth

In order to provision new Stealth-managed endpoints, Stealth must be integrated with Morpheus as described above. In addition, Stealth must be selected as the Security Server for the cloud you're provisioning into. Security servers can be selected at the time a new Cloud integration is created or by editing an existing Cloud integration.

Choosing a Cloud Security Server

Assuming the Cloud is already integrated with Morpheus, use the steps below to set the security server and activate Stealth prompts at provision time. The steps to set the security server during the time the cloud is initially integrated with Morpheus is very similar.

1. Navigate to Infrastructure > Clouds > (Your Selected Cloud)
2. Click *EDIT*
3. Click on Advanced Options to reveal additional selections
4. In the dropdown for SECURITY SERVER, choose an existing Stealth integration

Provisioning to a Stealth-enabled Cloud

Once we have selected our Stealth integration as the security server for at least one Cloud in Morpheus, new Instances (endpoints) can be provisioned and managed by Stealth.

1. Navigate to Provisioning > Instances
2. Click + *ADD*
3. Select the Instance Type, Cloud, and Group making sure to choose a Cloud that has been set up for an existing Stealth integration
4. On the Configure tab of the provisioning wizard, choose a Stealth Configuration and a Stealth Role according to the needs of the machine(s) being provisioning
5. Once the provisioning process is complete, the new Stealth-managed endpoints will be available and restricted based on the Stealth implementation

CREATE INSTANCE

TYPE

GROUP

CONFIGURE

AUTOMATION

REVIEW

Configuration Options

VERSION

7.5

LAYOUT

VMware VM

PLAN

1 CPU, 512MB Memory

Cores: 1 Memory: 512 MB Price: \$33.50 / Month

RESOURCE POOL

VOLUMES

root

10

GB

SCSI 0:0

Auto - Datastore

+

NETWORKS

VMXNET 3

Network Default

+

HOST

Select

FOLDER

Tests

STEALTH CONFIGURATION

Select

STEALTH ROLE

Select

▶ User Config

▶ Advanced Options

PREVIOUS

NEXT

Service Discovery

Consul

Morpheus can integrate with Consul to automatically install the Consul Agent in Client Mode on Instances and configure communication with the Consul host.

Add Consul Integration

1. Navigate to *Administration -> Integrations* and select + *New Integration*
2. Select Integration Type *Consul Service Registry*
3. Populate the following fields:
 - Name** Name of the Consul Integration in Morpheus
 - Enabled** Enabled by default
 - Consul Host** IP or Url of the Consul Host
 - Consul Http Port** Http port of the Consul Host
 - Username** Consul Host User
 - Password** Consul Host User Password
 - Datacenter ID** Validator key for the organization
4. Save Changes

The added Consul Integration is now available for use in Morpheus , but must be scoped to a Cloud or Group to automatically install the Consul Agent while provisioning.

Scope Consul Integration to a Cloud

1. Navigate to *Infrastructure -> Clouds*
2. Edit the target Cloud
3. Expand the *Advanced Options* section
4. In the *Service Registry* dropdown, select the Consul Integration.
5. Save Changes

Scope Consul Integration to a Group

1. Navigate to *Infrastructure -> Groups*
2. Edit the target Group
3. Expand the *Advanced Options* section
4. In the *Service Registry* dropdown, select the Consul Integration.
5. Save Changes

And that's it. After your integration is set up, all containers deployed within the Group or Cloud integrated will provision with the Consul Agent in Client Mode, gossiping to your Consul Server!

Storage

3Par

Adding 3Par Storage Server

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the SERVERS tab, Click the + *ADD* button.
4. From the ADD STORAGE SERVER wizard input the following:
NAME Name of the Storage Server in Morpheus
TYPE Select *3Par*
URL URL Of 3Par Server Example : *https://192.168.190.201:8008*
USERNAME Add your administrative user account.
PASSWORD Add your administrative password.
5. Select *SAVE CHANGES*

The 3Par Storage Server will be added and displayed in the Buckets tab.
Buckets, Files Shares and Storage Groups will be synced in.

AzureStorage

To Add Azure Storage

1. Navigate to *Infrastructure -> Storage* Hola
2. Select + *ADD*
3. From the New Storage Provider Wizard input the following:
Name Name of the storage provider.
Provider Type Azure
Storage Account Add Storage Account
Storage Key Add Storage Key
Share Name Add Share Name
Targets
 - Default Backup Target
 - Default Deployment Archive Target
 - Default Virtual Image Store
4. *Save Changes*

Dell ECS

Overview

Morpheus integrates with DELL EMC ECS via the ECS api. This allows Morpheus to talk directly to the ECS services.

When you add a ECS Server, Morpheus will sync in the following.

- Storage Groups
- Buckets
- File shares

Users will be able to create the following items within ECS without direct access to the ECS console.

- Buckets
- File shares

Storage Servers

The first step in the Dell EMC ECS integration is to add a Dell EMC ECS Storage Server. Once added, Buckets, Files Shares and Storage Groups will be synced in and can be access and managed in Morpheus.

Adding Dell EMC ECS Storage Server

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the SERVERS tab, Click the + *ADD* button.
4. From the ADD STORAGE SERVER wizard input the following:

NAME Name of the Storage Server in Morpheus

TYPE Select *Dell EMC ECS*

URL URL Of DELL EMC ECS Server Example : *https://192.168.190.200:4443*

Tip: The port 4443 is the api port for ECS api. This may be different depending on your configuration

USERNAME Add your administrative user account.

PASSWORD Add your administrative password.

S3 SERVICE URL (Optional) Add your S3 service url Example: *http://192.168.190.220:9020*

Note: S3 SERVICE URL is not required if you are not planning on using ECS S3.

5. Select *SAVE CHANGES*

The Dell EMC ECS Storage Server will be added and displayed in the Buckets tab.

Buckets, Files Shares and Storage Groups will be synced in.

Buckets

- **Buckets** will be listed in *Infrastructure - Storage - Buckets*
 - Buckets can be created and deleted with *Infrastructure - Storage* Role Permissions
 - Buckets can be browsed with *Infrastructure: Storage Browser* Role permissions
 - File and folders can be uploaded, downloaded and deleted with Full *Infrastructure: Storage Browser* Role permissions.

Adding Dell EMC ECS Buckets

Note: A Dell ECS Storage Server must be configured in *Infrastructure - Storage - Servers* prior to adding a Dell ECS Bucket.

To Add a Dell ECS Storage Bucket:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the BUCKETS tab, Click the + *ADD* button.
4. Select *Dell EMC ECS Bucket* from the dropdown list
5. From the NEW BUCKET Wizard input the following:

NAME Name of the Bucket in Morpheus.

STORAGE SERVICE Select existing Dell EMC ECS Storage Server (configured in *Infrastructure - Storage - Servers*)

BUCKET NAME Enter a name for the new Dell ECS bucket.

USER Your Dell EMC ECS S3 user account

SECRET KEY

Your Dell EMC ECS S3 Secret Example: jW+pFyAPtSS5FuEqKwt44xlpM/2

NAMESPACE Select Dell EMC ECS Namespace for the Bucket

STORAGE GROUP Select a Dell EMC ECS Storage Group

Default Backup Target Sets this bucket as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this Bucket will be presented.

Archive Snapshots Enabled to export VM snapshots to this Bucket when creating VMware Backups, after which the snapshot will be removed from the target hypervisor.

Default Deployment Archive Target Sets this Bucket as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this bucket as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the Bucket will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the bucket.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup Bucket.

BACKUP BUCKET Search for and select the Bucket the files will be backed up to.

DELETE OLD FILES

This option will delete files from this bucket after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the Bucket.

6. Select *SAVE CHANGES*

The Bucket will be created and displayed in the Buckets tab.

- To browse, upload, download, or delete files from this Bucket, select the name of the Bucket.
- To edit the Bucket, select the edit icon or select the name of the Bucket and select *ACTIONS - EDIT*.

Warning: Repointing a bucket that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a Bucket, select the trash icon or select the name of the Bucket and select *DELETE*.

Warning: When deleting a Bucket, all Deployment Versions and Backups associated with the Bucket will be deleted.

Add Dell EMC ECS File Shares

To Add a Dell EMC ECS File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *Dell EMC ECS Share* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

STORAGE SERVICE Select existing Dell EMC ECS Storage Server (configured in *Infrastructure - Storage - Servers*)

SHARE PATH

Enter Dell EMC ECS Share Path Example: `ecs-file-share-1`

USER Dell EMC ECS User

SECRET KEY Dell EMC ECS Secret key

Volume Size Specify volume size for the File Share (in MB)

Allowed IP's

Specify IP Addresses to limit accessibility to the File Share

Leave blank for open access Click the + symbol to the right of the first ALLOWED IPS field to add multiple IP's

NAMESPACE Select Dell EMC ECS Namespace (synced)

STORAGE GROUP Select Dell EMC ECS Storage Group (synced)

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

Isilon

Add Dell EMC Isilon Storage Server

Important: Enable insecure mode on the NFS settings. This allows non-root ports to be used. Setting the insecure/privileged mode will require a restart of the Isilon nodes.

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the SERVERS tab, Click the + *ADD* button.
4. From the ADD STORAGE SERVER wizard input the following:

NAME Name of the Storage Server in Morpheus

TYPE Select *Dell EMC Isilon*

URL URL Of Dell EMC Isilon Server Example : *https://192.168.190.202:8080*

USERNAME Add your administrative user account.

PASSWORD Add your administrative password.

PROVISION USER Select Provision User

PROVISION GROUP Select Provision Group

ROOT PATH

Enter Root Path Example : ``

5. Select *SAVE CHANGES*

The Dell EMC Isilon Storage Server will be added and displayed in the Buckets tab.

Buckets, Files Shares and Storage Groups will be synced in.

Add Dell EMC Isilon File Share

To Add a Dell EMC Isilon File Share:

1. Select the Infrastructure link in the navigation bar.
2. Select the Storage link in the sub navigation bar.
3. In the FILE SHARES tab, Click the + *ADD* button.
4. Select *Dell EMC Isilon Share* from the dropdown list
5. From the NEW FILE SHARE Wizard input the following:

NAME Name of the File Share in Morpheus.

STORAGE SERVICE Select existing Dell EMC Isilon Storage Server (configured in *Infrastructure - Storage - Servers*)

SHARE PATH

Enter Dell EMC Isilon Share Path Example: *ecs-file-share-1*

Volume Size Specify volume size for the File Share (in MB)

Allowed IP's

Specify IP Addresses to limit accessibility to the File Share

Leave blank for open access Click the + symbol to the right of the first ALLOWED IPS field to add multiple IP's

NAMESPACE Select Dell EMC Isilon Namespace (syncd)

STORAGE GROUP Select Dell EMC Isilon Storage Group (syncd)

Default Backup Target Sets this File Share as the default backup target when creating Backups. If selected the option to update existing Backup configuration to use this File Share will be presented.

Archive Snapshots Enabled to export VM snapshots to this File Share when creating VMware Backups, after which the snapshot will be removed from the source Cloud.

Default Deployment Archive Target Sets this File Share as the default storage target when uploading Deployment files in the *Deployments* section.

Default Virtual Image Store Sets this File Share as the default storage target when uploading Virtual Images from the *Virtual Images* section, importing Images from Instance Actions, creating Images with the *Image Builder* and when creating new images from *Migrations*.

RETENTION POLICY

None Files in the File Share will not be automatically deleted or backed up.

Backup Old Files

This option will backup files after a set amount of time and remove them from the File Share.

DAYS OLD Files older than the set number of days will be automatically backed up to the selected Backup File Share.

BACKUP File Share Search for and select the File Share the files will be backed up to.

DELETE OLD FILES

This option will delete files from this File Share after a set amount of days.

DAYS OLD Files older than the set number of days will be automatically deleted from the File Share.

6. Select *SAVE CHANGES*

The File Share will be created and displayed in the File Shares tab.

- To browse, upload, download, or delete files from this File Share, select the name of the File Share.
- To edit the File Share, select the edit icon or select the name of the File Share and select *ACTIONS - EDIT*.

Warning: Repointing a File Share that is in use may cause loss of file references. Ensure data is mirrored first.

- To delete a File Share, select the trash icon or select the name of the File Share and select *DELETE*.

Warning: When deleting a File Share, all Deployment Versions and Backups associated with the File Share will be deleted.

Supported Integration Versions

Morpheus supports an extensive range of software integrations and versions past and present. Current iterations of Amazon AWS, Microsoft Azure, Google Cloud Platform, Digital Ocean, HPE OneView, OpenTelekom Cloud, IBM Bluemix, Softlayer and UpCloud are all supported.

In addition, Morpheus is verified to work with, but not limited to:

Integrations

Note: Current iterations of Amazon AWS, Microsoft Azure, Google Cloud Platform, Digital Ocean, HPE OneView, OpenTelekom Cloud, IBM Bluemix, Softlayer and UpCloud are all supported.

Integration	Supported Version(s)	Notes
Ansible	2.7.x	
Ansible Tower	3.3.x	
App Dynamics	4.5.x	
Azure Stack	2002 back to 1908	2019-03-01-hybrid api-profile version used which is supported in 1908 and later Azure Stack versions
Cisco ACI	3.10	
Commvault	v11 sp 19	
Jenkins	< 2.176.1	
Kubernetes	1.x	
Kubernetes	Major:"1", Minor:"17", GitVersion:"v1.17.3"	
Microsoft Hyper-V	2012R2, 2016, 2019	
Nutanix AHV	5.0 - 5.10 Note: Prism Central is not a supported endpoint	In 5.5 - 5.7 if Prism Central is managing Prism Element, image creation will not function due to PC Image Management.
Openstack	Juno, Kilo, Liberty, Mitaka, Newton, Ocata, Pike, Queens, Rocky, Stein, Train	
Rubrik	4.2	
SCVMM	2016, 2019	
ServiceNow	Kingston, London, Madrid, New York, Orlando, and Paris	
Splunk	7.10	
Terraform	v0.11.x, v0.12.18+	v0.13.x has not been validated
vCloud Director	8.20, 9.1, 9.5, 9.7*, 10.0	*vCD 9.7 supported on API v31
Veeam	9.5u3, 9.5u4, 10	
VMware ESXi	5.5, 6.0, 6.5, 6.7, 7	
VMware Fusion	8, 9, 10+	
VMware NSX	-V, -T v3	
VMware vCenter	5.5, 6.0, 6.5, 6.7, 7	
XenServer	7.x	

Note: Non-listed versions may be compatible but are not verified.

Note: Non-listed versions may be compatible but are not verified.

If you have any specific requirements please contact support@morpheusdata.com

1.3.14 v5.2.0 Release Notes

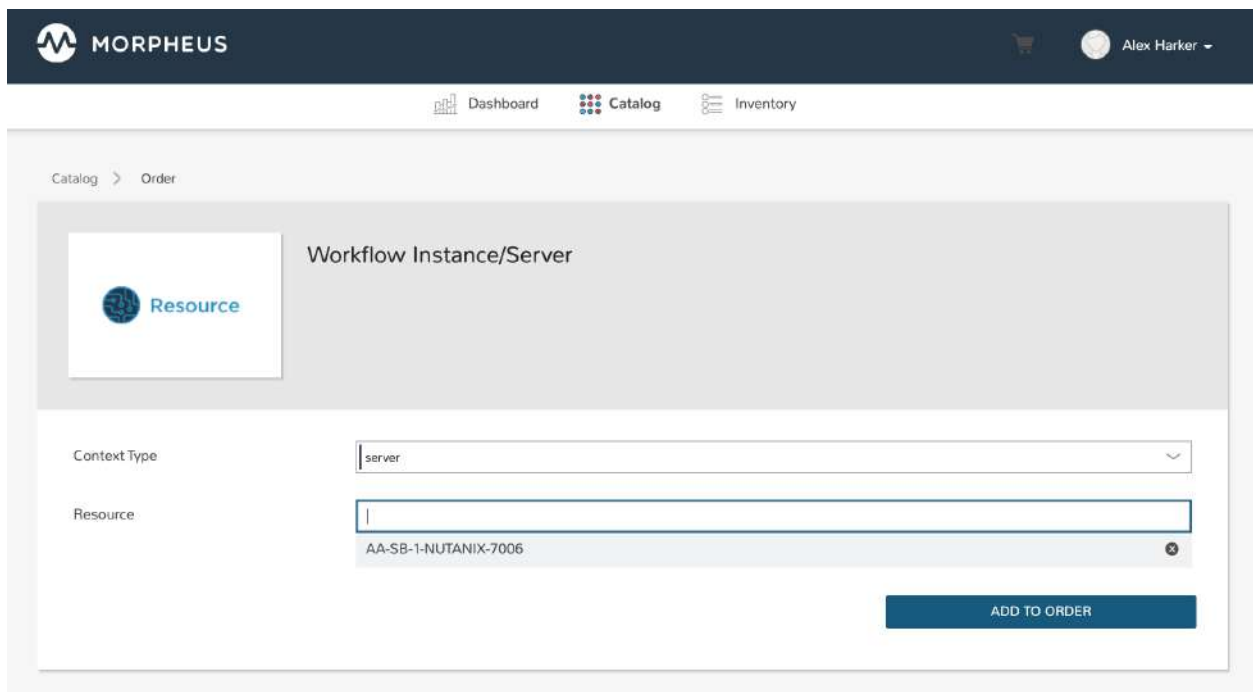
Note: This list includes improvements added in 4.2.4 which were not part of the 5.0.0 beta. Users upgrading from 4.x.x may also want to review the [5.0.0 change list](#) to get a complete picture of the changes.

v5.2.0 Highlights

Service Catalog Persona Improvements

The Morpheus version 5.0.0 beta introduced Personas, which are a new approach for optimizing and simplifying self-service for targeted audiences. The first Persona to ship is [Service Catalog](#), which sees additional improvements in the 5.2.0 LTS release.

- Make Morpheus Operational Workflows available for order from the Service Catalog and run them against selected targets
- With added API/CLI support, work with Personas, create and manage Catalog Items, and make selections from the catalog through Morpheus API and CLI tools
- Inventory list view now includes much greater detail about each inventory item
- Categorize items under selected headers for enhanced discoverability as the catalog grows
- With the added quantity selector, order additional copies of items in your cart without creating duplicate selections



ServiceNow Integration Improvements

Morpheus 5.2.0 brings improvements to the ServiceNow integration, including upgrades to incident surfacing, integration with service catalog items, and more.

- “Morpheus Incident” alerts are now more insightful including links to the related Morpheus incident or check, severity information, and other details about the failing check
- Provision Service Catalog Items through the Morpheus ServiceNow plugin
- Added the capability to identify a MID server once on the properties page rather than setting it individually for each call
- Pricing data is now available to the ServiceNow plugin when ordering Service Catalog items. This is made available on the XML as a monthly price, users would have to modify the form UI to surface this information

Hide Blueprint fields

Administrators have been able to lock Blueprint fields to restrict editing during App provisioning for a long time. Fields can now be hidden from view completely by toggling the lock/unlock icon to the hidden setting. When users provision Apps based on this Blueprint, they will not see hidden fields at all. This gives administrators additional flexibility to mask unneeded complexity from users.

EDIT BLUEPRINT

STRUCTURE

A Hide

App

APACHE

Group: All Clouds, Cloud: QA Amazon

ubuntu

Group: All Clouds, Cloud: QA VMware

MySQL

Group: All Clouds, Cloud: QA Amazon

NGINX

Group: All Clouds, Cloud: QA Docker

Builder

Raw

Preview

CONFIGURATION

☐ ALLOW EXISTING INSTANCES

▼ Instance Info

NAME

apache-amazon

DESCRIPTION

Description

▼ Configuration Options

LAYOUT

Amazon Apache on Ubuntu 14.04

PLAN

Amazon T2 Micro - 1 Core, 1GB Memory

Cores: 1 Memory: 1 GB Price: \$13.7016 / Month

RESOURCE POOL

labs

VOLUMES

root

20

GB

gp2

data

10

GB

gp2

NETWORKS

labc web 1a / subnet...

DHCP

All New Features

- Amazon: **Amazon AWS Cloud Integration Improvements**
 - Deploying MySQL or SQL Server with Amazon RDS now automatically creates the corresponding check and Instance status is reported
 - Hong Kong region (ap-east-1) support added
- Amazon EKS: Support added for version 1.7.x
- Azure AKS: Support added for version 1.7.11
- Azure Public Cloud: **Azure Cloud Integration Improvements**
 - Option to enable Azure Guest OS Diagnostics when provisioning Instance or App
 - Added option to enable Azure Boot Diagnostics when provisioning an Instance or App
 - Set disk encryption (user or platform-managed) and an encryption set (if user-managed) for an Azure Cloud integration (Add/Edit Cloud modal)
- Azure Stack: Azure Stack: Added support for ARM templates
- BlueCat: Support added for Bluecat 9.x +
- Blueprints: Added ability to set fields as hidden when creating a Blueprint. These fields will not be visible when provisioning an App from the Blueprint. Previously fields could be locked but not hidden.
- Clouds: Cloud sync enhancements including variable sync schedules that auto-adjust per cloud, resulting in optimized sync times and reduction in sync overlaps and record lock conflicts
- KVM: **KVM Improvements**
 - Console access is now available for VMs on the KVM server which were not provisioned by Morpheus
- Networks: Visibility and Tenant permissions added for IP Pools (select Permissions under the “MORE” menus on the IP Pools)
- NSX-V: Create and manage DHCP Pools for Edge Gateway routers
- NSX-T: **NSX-T Integration Improvements**
 - Visibility and Tenant permissions added for Transport Zones (select Permissions under the “MORE” menus on the Transport Zone tab)
 - Visibility and Tenant permissions added for Edge Clusters (select Permissions under the “MORE” menus on the Edge Clusters tab)
 - Create and manage NAT rules for T-0 and T-1 routers (see NAT tab on the detail page for a T-0 or T-1 router)
 - Role permissions added to control access to the T-0 and T-1 routers tabs for an NSX-T network integration
 - Interfaces tab for T-1 routers renamed to Service Interfaces for clarity
- OpenStack: Backup process improved to handle longer running jobs for backing up large instances
- Openstack: `Service Endpoints` section added to Cloud config for manually overriding OpenStack API endpoints
- Policies: Load balancer pricing is factored when enforcing budget policies during provisioning and reconfiguration
- Pricing: **Load Balancer Price Tracking**
 - Load balancer support in Price Plans, Price Sets, and Prices (Administration > Plans & Pricing)

- Load balancer price data sync for Azure and Amazon
 - Automatically apply Price Plans to load balancers based on Plan configuration
 - Usage and Billing data for load balancers
- Provisioning: Set a value to be prepended to all environment variables loaded as part of Instance or App provisioning
- Proxies: Global proxy setting now applies to all Morpheus functionality, including local integrations such as Ansible and Terraform
- Reports: “Invoice Details” report added to list of available report types. For a selected Cloud, group invoices by up to two additional parameters (Region, Cloud, Plan, Tag or Tenant)
- Security Scanning: **Run SCAP Scans to Confirm Security Compliance**
 - Create Jobs to run SCAP scans against any group of Instances or Servers either on-demand or on recurring schedules
 - View complete SCAP evaluation reports on your systems from inside the Morpheus UI
 - View and track security scan run histories
- Roles: **Role Permission Changes**
 - Network integration firewall permissions (Infrastructure > Network > Integrations > Selected integration > Firewalls) now have their own setting (Infrastructure: Network Firewalls). Previously they were inherited from the “Network: Integrations” permission
 - Role permissions added to control access to the T-0 and T-1 routers tabs for an NSX-T network integration
 - Security: Scanning **Feature Access Permission added**
 - Determines access to the Security Packages tab on the Jobs list page (Provisioning > Jobs), Security Scanning type Jobs, and Security Subtab inside the Software tab on a server detail page where the results of security scans are viewed
 - Allows access to view, create, and run security scans on existing systems, as well as view the results of previously-run scans
 - This permission is recommended for those responsible for security compliance of existing systems
- Security Scanning: Windows support added for SCAP security scans
- Service Catalog: **Service Catalog Persona Improvements**
 - Operational Workflows can be made available as Service Catalog Items and ordered by Service Catalog Persona users
 - Catalog Items can be categorized under specific headers for easier discoverability as the catalog grows
 - Quantity selector added for items in cart to avoid the need for adding duplicate items
- ServiceNow: **ServiceNow Integration Improvements**
 - “Morpheus Incident” alerts are now more insightful including links to the related Morpheus incident or check, severity information, and other details about the failing check
 - Provision Service Catalog Items through the Morpheus ServiceNow plugin
 - Inventory list view now includes much greater detail about each inventory item
 - Added the capability to identify a MID server once on the properties page rather than setting it individually for each call

- Pricing data is now available to the ServiceNow plugin when ordering Service Catalog items. This is made available on the XML as a monthly price, users would have to modify the form UI to surface this information
- Tasks: Tasks now have a detail page with a Summary tab showing the script and a Workflows tab listing the Workflows in which the Task is used
- Tenants: Metadata, specifically an account number, account name, and customer number, can now be tracked for Tenants
- UI: **Interface and Usability Improvements**
 - Administrators can now determine the required length and complexity of user passwords (Administration > Settings > Appliance > User Management Settings)
 - When applying state to Terraform and CloudFormation Apps, a friendly progress bar is displayed to indicate the change
 - Icons added for AWS services (such as in Service Catalog), including AWS App Mesh, AWS SQS, and AWS SDB
 - MySQL tmp file location moved from `/tmp` to `/var/run/morpheus/mysqld`
 - Session expiration times can now be configured (Administration > Settings > Appliance), if desired a window can also be displayed at a specified time to warn about the impending logout
 - All navigations bars with potential for high tab counts now handle this scenario gracefully
 - Visibility column added to Catalog Item list (Tools > Self Service) to conveniently indicate whether an item is shared with Tenants
 - Friendly error messages are surfaced if there is a problem creating the items checked out in a Service Catalog cart, the Instance was simply not created and log access was needed to see what went wrong
 - CenturyLink Edge Cloud type renamed to Lumen Edge
- Workflows: “Configuration” phase added to Provisioning Workflows. Tasks in this phase are run prior to the initial provision.

Morpheus API & CLI Improvements

- Billing: Optional parameters added to support pagination of large returns
- Deployments: **Deployments API/CLI Improvements**
 - Support for adding files to a Deployment version
 - Support for managing Instance deploys (appDeploys). This used to only provide endpoints for a specific instance to deploy and list deploys. Now it has full CRUD, and list shows account wide deploys. See *morpheus deploys*.
- Instances: Support added for filtering by `expireDate` and `shutdownDate`
- Instances: Search by tag names and values
- Personas: **Personas and Service Catalog Persona Functionality Added**
 - Set the default Persona for a user
 - Create catalog items for Service Catalog Persona users
 - Set Role permissions regarding Persona and Catalog Item type access
 - Browse the catalog, add items to cart, and checkout as a Service Catalog Persona user

- Search: Global search added similar to the global search bar that has existed in the UI
- Tenants: Account (Tenant) metadata field support added (customerNumber, accountNumber, and accountName)
- Virtual Images: Associated volumes are returned with maxStorage viewable for each

Fixes

- ACI: Fixed invalid display error when creating ACI Application Profile ⁺
- ACI: Fixed network deletion issue caused by illegal characters in CIDR ⁺
- ACI: Fixed potential issues preventing deletion of Cisco ACI Integrations ⁺
- Agent: Auto-recovery settings now enabled for Morpheus Windows Agent service. ⁺
- Amazon: ALB's: Fix for adding ALBs in a subtenant ⁺
- Amazon: Fixed Security Groups stat always showing 0 on the Resources tab of the Cloud summary page ⁺
- Amazon: Fixed running state sync for Amazon Instance/Container status when Morpheus Agent is not installed on associated VM ⁻²
- Amazon: Fixed S3 Bucket create and delete not utilizing AWS Cloud API Proxy config
- Ansible Tower: Fixed invalid Ansible Tower integration link in cloud details pages ⁺
- Ansible: Ansible integrations now utilize Global Proxy config for repo connections
- API/CLI: Fixed config property of Azure image type missing in GET and POST (CLI) ⁺
- API/CLI: Fixed Task creation when using a repository source ⁺
- API/CLI: Validation and response added when passing invalid value for POST /api/roles:roleType ⁺
- Apps: Fixed inconsistent app, node and execution statuses during App provisioning when a Workflow Task fails during the Provision phase ⁺
- Apps: Fixed issue with Zone selection and Instance Configurations ⁻²
- Apps: Updated the NAME property for VM and Container lists on App Detail views to match Instance Detail views ⁺
- Archives: Fixed timeout issue with archive files > 1GB caused by legacy Archives path set in default Nginx config
- Azure: Fix for automated Active Directory domain joins due to -NewName ⁺
- Azure: Fixed long running provision timeouts for ARM Instance Spec Templates ⁺
- Cloud Formation: Fixed issue creating Lambda resources from CF Blueprints. (Note: Lambda resource objects will be added in future release) ⁺
- Commvault: Fixed issue with subtenant Commvault Backup Job completion when Backup and Backup Job names use custom values
- Git: Fixed issue deleting Git integrations with existing file content associations ⁺
- Github: Github integrations now utilize Global Proxy config for Github connections
- Health: Fixed display of Memory: System Swap and Memory: Free Swap values in the Appliance Health section. ⁺

- Hosts: The Remove Infrastructure and Preserve Volumes checkboxes are now present and functional when performing bulk VM delegations (Infrastructure > Hosts > Virtual Machines) ⁺
- Identity Sources: SAML: Fixed issue with checkbox rendering in Firefox browsers
- Networks: If a user has only read-level permission for the “Infrastructure: Network Routers” feature, the + *CREATE NEIGHBOR* button on the BGP tab of the Router Detail page is now hidden ⁺
- Networks: If a user has only read-level permission for the “Infrastructure: Network Routers” feature, that user no longer has the ability to edit or delete router firewall rules ⁺
- NSX-V: Create and manage DHCP Pools for Edge Gateway routers ⁺
- Nutanix: Fixed issue provisioning custom images stored in Amazon S3 buckets ⁺
- Nutanix: Removed root disk storage container selection during provisioning as Nutanix requires root disk must be created on the same storage container as the template ⁺
- Openstack Clouds: Fixed security groups scoped to “All” Clouds which previously were not displayed during provisioning ⁺
- OTC: Fixed issue with long running Instance backups not exporting ⁺
- Policies: Fixed issue where VM tags were allowed to be changed to values not compliant with an active, strictly-enforced tag policy. ⁺
- Policies: Fixed issue with expiration policies not removing resources which are in a failed state ⁺
- Policies: Updated email notification Instance links to redirect to subtenant logins ⁺
- PowerDNS: Fixed display issue with PowerDNS record “Content” fields ⁺
- Provisioning: Fixed sudoer permissions for users created during provisioning when the associated Linux username contains a . ⁺
- PXE Boot: Fixed editing of discovered MAC Addresses ⁺
- RDS: Fixed issue with editing Power Schedules for AWS RDS Instances ⁺
- Reconfigure: Fixed page error when decimal is specified in a disk size during reconfigure ⁺
- Reports: Fixed issue with Instance Inventory Summary Report potentially showing old resource values on re-configured Instances ⁺
- Role: Fixed issues with Persona permissions not copying to Multitenant roles, the default Persona is copied now as well ⁻²
- SCVMM: Adding a disk, resizing a data disk, or removing a data disk during reconfigure will no longer trigger a restart ⁺
- SCVMM: Fixed adding disks during reconfigure of Generation 2 virtual machines ⁺
- SCVMM: Fixed issue where selected SCVMM Cloud was not being passed SCVMM VM config ⁺
- SCVMM: Fixed issue with optical drive being removed during provisioning of Generation 2 virtual machines ⁺
- SCVMM: Fixed Instance reconfigure startup memory and fixed memory allocation ⁺
- SCVMM: Fixed startup memory and fixed memory allocations when dynamic memory is enabled ⁺
- Security: XSS vulnerability removed ⁺
- Tags: Fixed error when trying to create a tag without a value ⁺
- Tenants: Fixed Tenant deletion issue related to existing network_security_server records
- Terraform: Resolved issue where *NEXT* button would become re-enabled on App provisioning prior to completion of validations over 35 seconds ⁺

- vCloud Director: Fixed issue with frequent usage record restarts ⁺
- VMware: Fixed datastore cluster references for datastores shared across multiple clusters ⁺
- VMware: Fixed issue with high-resolution hypervisor consoles showing blank on initial uncompressed connection ⁺
- VMware: Fixed issue with Subtenant setting VMware Folder Group Access permissions ⁺
- Workflows: Fixed issue with Reboot tasks potentially causing Instance state to show as Running when a Provision phase task has failed ⁺

Note: ⁺ indicates items also released in v4.2.4

Note: ⁻² indicates items included in v5.2.0-2

1.3.15 v5.2.0 Compatibility & Breaking Changes

When installing and upgrading to Morpheus v5.2.0, refer to the following to ensure compatibility.

Breaking Changes

- 4.2.1+: Appliance: OS: Ubuntu 14.04 has reached its end of life (EOL) and is no longer supported as a Morpheus Appliance Host Operating System. Any Morpheus Appliance running on 14.04 must be upgraded to 16.04, 18.04 or 20.04 BEFORE upgrading to 4.2.1+. Upgrades on 14.04 will not succeed
- 4.2.1+: Clouds: VirtualBox, VirtuSteam, and MetaCloud Cloud Types are no longer supported or available
- 4.2.1+: Puppet: Morpheus integration now supports version 6+. Puppet versions prior to 6 are no longer supported
- 4.2.1+: Tasks: Python: Virtual environment are now used for Python Tasks. **Note:** `virtualenv` is required on all Appliance App nodes: `pip install virtualenv`
- 4.2.4: For appliances with externalized MySQL databases, due to MySQL deprecation of the “EDT” timezone you may need to update your database timezone to UTC or another compatible value. If this is not done, you will receive errors referencing timezone and Morpheus will not start. Morpheus should handle this change automatically for all-in-one appliances.
- 5.0.0+: When upgrading to 5.0.0+ from 4.x.x, any bearer tokens that have been generated are deleted which requires users to request new bearer tokens

Morpheus Application OS

Morpheus can be installed on the following platforms. Please note the table below is for Morpheus Application OS support, not Morpheus Agent OS Support.

Important: Existing Morpheus Appliances on 14.04 must upgrade to 16.04 or 18.04 PRIOR to upgrading to v4.2+.

Table 17: Supported Appliance Operating Systems

OS	Version(s)	Notes
Amazon Linux	2	
CentOS	7.x, 8.x	
Debian	9, 10	FreeRDP 2.0 is not compatible with Debian 9. Guacd will remain at 1.0.0 for Appliances running on 9.
RHEL	7.x, 8.x	
SUSE Linux Enterprise Server (SLES)	12, 15	
Ubuntu	16.04, 18.04	14.04 is no longer supported for Appliance OS. Existing Appliances on 14.04 must upgrade to 16.04 or 18.04 PRIOR to upgrading to v4.2.1+. Note: 14.04 is still supported by the Morpheus Agent.

Services

v5.2.0 Service Version Changes

- MySQL: Upgraded to 5.7.32 for non-fips versions (5.7.29)
- Nginx: Upgraded to v1.19.3
- RabbitMQ: Upgraded to v3.8.9
- Tomcat: Upgraded to 9.0.39

v5.2.0 Service Version Compatibility

When externalizing MySQL, Elasticsearch and/or RabbitMQ services, the following versions are compatible with version Morpheus v5.2.0

Service	Compatible Branch	Morpheus Installer Version
MySQL	v5.7	v5.7.32
MySQL (FIPS)	v5.7	v.5.7.29
Percona	5.7, WSREP 31	n/a
Elasticsearch	v7.x	v7.8.1
RabbitMQ	v3.5-3.8	v3.8.9
Tomcat		v9.0.39
Nginx		v1.19.3

Important: Elasticsearch 7.x is required for v5.2.0. Refer to [Upgrading](#) section for more information.

Security

Important: Please be aware of the default security enhancements added to v4.1.2+ and assess potential impacts to your environment, including agent installation and frontend load balancers.

CVEs Addressed

- CVE-2017-5929
- CVE-2019-2692
- CVE-2020-2933
- CVE-2020-14338
- CVE-2020-15250

Integrations

Note: Current iterations of Amazon AWS, Microsoft Azure, Google Cloud Platform, Digital Ocean, HPE OneView, OpenTelekom Cloud, IBM Bluemix, Softlayer and UpCloud are all supported.

Integration	Supported Version(s)	Notes
Ansible	2.7.x	
Ansible Tower	3.3.x	
App Dynamics	4.5.x	
Azure Stack	2002 back to 1908	2019-03-01-hybrid api-profile version used which is supported in 1908 and later Azure Stack versions
Cisco ACI	3.10	
Commvault	v11 sp 19	
Jenkins	< 2.176.1	
Kubernetes	1.x	
Kubernetes	Major:"1", Minor:"17", GitVersion:"v1.17.3"	
Microsoft Hyper-V	2012R2, 2016, 2019	
Nutanix AHV	5.0 - 5.10 Note: Prism Central is not a supported endpoint	In 5.5 - 5.7 if Prism Central is managing Prism Element, image creation will not function due to PC Image Management.
Openstack	Juno, Kilo, Liberty, Mitaka, Newton, Ocata, Pike, Queens, Rocky, Stein, Train	
Rubrik	4.2	
SCVMM	2016, 2019	
ServiceNow	Kingston, London, Madrid, New York, Orlando, and Paris	
Splunk	7.10	
Terraform	v0.11.x, v0.12.18+	v0.13.x has not been validated
vCloud Director	8.20, 9.1, 9.5, 9.7*, 10.0	*vCD 9.7 supported on API v31
Veeam	9.5u3, 9.5u4, 10	
VMware ESXi	5.5, 6.0, 6.5, 6.7, 7	
VMware Fusion	8, 9, 10+	
VMware NSX	-V, -T v3	
VMware vCenter	5.5, 6.0, 6.5, 6.7, 7	
XenServer	7.x	

Note: Non-listed versions may be compatible but are not verified.