Management and Security Server Deployment Guide

MSS Deployment Guide

Table of contents

Management and Security Server Deployment Guide

Management and Security Server (MSS) provides an administrator the means to centrally secure, manage, and monitor users' access to host applications.

MSS contains a new architecture that simplifies deployment, tightens security, improves scaling and high availability, and eases ongoing maintenance.

This guide is intended to walk you through the steps of planning, installing, and configuring your product.

See the MSS Release Notes (page for a list of new features.

In this guide

Legal Notice

Plan

Planning for Deployment

The following steps will help you plan your deployment.

- Determine which installation type (page 4) meets your needs.
- Familiarize yourself with what a standard deployment (page 5) consists of and how many nodes you'll need.
- Learn about the Cluster Certificate (page 7) and Cluster DNS name (page 6) .

Choose an Installation Type

There are two options when considering how to deploy your product. Unless you have specific needs that require using the Linux installer, the virtual software appliance is the suggested default approach.

Installation Type	Description	When to choose	System Requirements
Virtual Appliance	The virtual appliance is a pre-configured virtual machine that contains everything you need to run the system. Deploy the appliance into your virtualization environment using an OVF file, and create as many appliances as needed to meet the demands of your environment.	The simplest and recommended installation type. Choose this option if you: Prefer easy, one-click software updates. Aim to minimize maintenance efforts. Have a virtualization platform that supports OVF (Open Virtualization Format) files.	View System Requirements (page 10)

Installation Type	Description	When to choose	System Requirements
Linux Installer	The Linux installer is a `.sh` script that installs all of the software needed on an existing Linux server of your choice, whether it's virtual or physical.	Choose this option if you: Require a specific Linux distribution. Require more control over the operating system, server configuration, and system updates. Use a cloud provider that doesn't allow BIOS boot time access; thus the appliance cannot be used.	View System Requirements (page 17)

Standard Deployment

We recommend the following default deployment as a starting point:

An external load balancer

Three cluster nodes
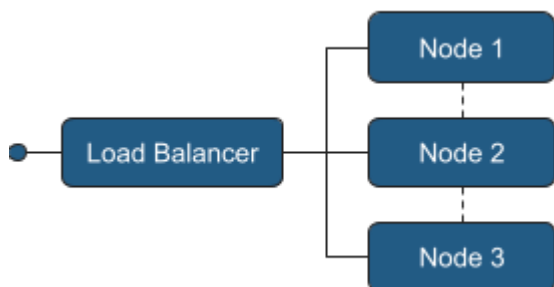
This deployment provides

- **Load Balancing** - User requests are distributed across nodes for performance and availability.
- **High availability** - Ability for one node to go offline without significantly impacting users.
- **Scalability** - Additional capacity may be added as needed.

Requirements - What you provide

Servers / virtual machines that meet the system requirements: Appliance (page 10) or Linux installer (page 17) .

An odd number of nodes is always required.

A load balancer (page 6) with a DNS hostname (page 6) for the cluster.

A certificate key pair (page 7) for securing access to the cluster.

Additional information

Load balancer

An external load balancer is optional but recommended. The specifics about which load balancer to use and the exact configuration are beyond the scope of this documentation.

The load balancer should be configured:

to direct traffic to all available nodes

with the cluster certificate

to use `/ping` as a health endpoint for each node in your cluster

The load balancer does not need to be configured for session affinity or stickiness. Session affinity is automatically handled inside the cluster.

Requests to any node in the cluster are automatically load balanced by the system to nodes across the cluster. This provides a basic level of load balancing regardless of the presence of an external load balancer. An alternative is to use DNS round-robin load balancing, in which the cluster DNS hostname resolves to each node in the cluster.

Cluster DNS name

A DNS hostname is provided by you and will be used when accessing the cluster. This DNS hostname is configured on the cluster as part of the setup process.

The cluster DNS name should resolve to the address of your external load balancer.

If not using an external load balancer, the cluster DNS name should resolve to the IP addresses of each node in the cluster.

Cluster certificate

A certificate key pair is provided by you and is used to secure all communication to the cluster. A self-signed certificate is generated and can be used for accessing the cluster initially, but for a production deployment, we recommend that you provide your own cluster certificate.

The cluster certificate key pair you provide must be in the PEM format.

The certificate should contain the hostname of your load balancer, both as the common name and as a DNS Subject Alternative Name (SAN) entry.

If not using an external load balancer, the certificate should contain a DNS SAN entry for each node in the cluster.

The certificate will additionally be served up by each node in the cluster if accessed directly. If not already present, an additional SAN entry for each node node should be added if direct node access is desired.

Information to gather

While provisioning servers, gather the following information for use in the installation process:

Static IP address

Fully qualified domain name (FQDN) of each node

When using the Appliance, you also need the following network related information:

Network mask — *if you used a static IP address during installation*

Default gateway

DNS Server(s)

Fault Tolerance and Scaling

Maintaining a quorum

To ensure that both service and cluster level operations run smoothly, a **quorum of cluster nodes** must be running at **all times**. A quorum means that more than half (50% + 1) of the nodes need to be running and communicating with each other at any given moment.

Your cluster should always be designed and built to contain an **odd number of nodes**, which helps to maintain a quorum in both normal and adverse networking conditions. Keep this in mind when planning your deployment and looking ahead to maintaining your cluster.

# of nodes in cluster	# of nodes required for a quorum
3	2
5	3
n	(n / 2) + 1

Failure handling

Services. The health of all services in the system is monitored.

- If a service is found to be unhealthy, the system will automatically attempt to self-heal, generally by restarting the process.
- Service interruptions may occur depending on the type of failure.
- Events regarding detected failures can be viewed in the Cluster Management dashboard.

Nodes. When a cluster *node* becomes unavailable for any reason, whether planned or unplanned:

- The cluster will generally move the services that had been running on that node onto other nodes.
- It may take five minutes or more for a node to be recognized as unavailable. This delay is designed to prevent unwarranted service disruptions that could be triggered by temporary conditions, such as intermittent network issues.
- Instructions are provided (in Cluster Management - Nodes help) for gracefully shutting down or rebooting a node. These should be used any time a node is shut down or rebooted.

Scaling

If you need additional capacity beyond what the standard deployment provides, choose from these two options:

- **Vertical** scaling - Add more memory and CPUs to your nodes. This is the suggested method for scaling as it does not require the additional complexity of managing more nodes.
- **Horizontal** scaling - Add more nodes to your cluster. While this allows you to scale as much as needed, it involves managing additional nodes.

> 🔥 **Important**
>
> You must always have an odd number of nodes in your cluster.

Headroom

When building a fault-tolerant cluster, each node must reserve a minimum level of free compute resources so that it can take on additional load when needed.

- When scaling *vertically*, we recommend doubling the required system requirements.

- When scaling *horizontally*, this has been factored in to the system requirements.

Next step

Once you've developed a plan for your deployment, you're ready to .

Install or Update

Install or Update

Check the system requirements, installation instructions, or steps to upgrade your deployment type.

Virtual Appliance

Virtual Appliance - System Requirements

Virtualization platform

The appliance is installed using an OVF (Open Virtualization Format) file and therefore requires a virtualization platform that supports OVF, such as VMWare ESXi.

> 💡 **Note**
>
> The appliance is not supported in public cloud environments.

Minimum CPU and memory requirements

Each virtual appliance VM requires:

8 CPU Cores

16 GB RAM

100 GB disk space (SSD)

Fast storage

To ensure optimal performance and reliability, the use of a solid-state drive (SSD) or other fast storage solutions is required. Not using SSD based storage may lead to inconsistent behaviors and errors.

Fixed IP address

A fixed, non-changing IP address is required for each node. DHCP (Dynamic Host Configuration Protocol) is supported but the IP must be reserved and cannot change.

Network ports

The following ports must be exposed and available between all nodes:

Port	Purpose
6443	Kubernetes API Server
8472	Virtual LAN
10250	Kubernetes metrics
2379-2380	etcd

The following port must be exposed internally for administrator access:

Port	Purpose
9443	Appliance Administration Console

The following ports must be exposed for outside access:

Port	Purpose
443	Product access
3000	Security Proxy Server *
8001	AJP **

* The Security Proxy port use is optional.

** The AJP port is used when optionally integrated with Microsoft's IIS web server.

Supported web browsers

The following web browsers are supported:

Google Chrome (recommended)

Mozilla Firefox (recommended)

Microsoft Edge

Installing the Appliance

License entitlement file

A license entitlement file (activation file) is required to install products in the appliance and is available from the product download site. Make sure you download the current version of the activation file for your product.

Installation steps

To install the appliance, first make sure your system meets the , and then perform the following steps:

1. Download and unzip the Appliance ZIP file.

   The ZIP file contains the files needed to install the appliance. All files must reside in the same directory during deployment.

2. Gather the necessary information from your network administrator.

   If using a Static IP Address, gather the following:

Fully Qualified Domain Name (FQDN)

IP

Netmask

Gateway

DNS Server

   If using DHCP IP Address (with a fixed IP), note:

Fully Qualified Domain Name (FQDN)

3. Import the OVF file into your virtualization system to create a new appliance template. Create a new VM instance from the new template.

4. Start the VM. Read and accept the license file.

5. Start to configure the appliance by specifying a password for the root user on the appliance.

6. Click Next to configure the hostname and network options.

   • Specify a fully qualified DNS hostname for the appliance; then select whether to use a Static IP address or DHCP. Click Next.

   • If you use a static IP address, you must specify the IP address assigned to your virtual machine, netmask, the gateway, and DNS server(s) that you gathered.

   • If you are using DHCP, the IP address **must be fixed**; it cannot change over time.

7. Click Next and wait for the initialization to complete.

   During initialization, progress messages appear on a console screen. Initialization takes 5-15 minutes. When a login prompt displays, initialization is complete.

8. After a login prompt is displayed, the appliance console is accessed using the supplied URL. For example: `https://hostname:9443` .

Accessing the appliance console

The appliance console provides a comprehensive set of capabilities, including configuring clusters, adding/removing programs, system configuration, and support and maintenance tasks.

- The Appliance Console is accessed on port 9443, for example: `https://hostname:9443`
- The root account is used to access the console by default. Use the password specified during appliance configuration.
- Log in to the console and browse around to explore the different options and capabilities.

Registering the appliance

Before installing your product, register the appliance. Registration enables you to receive online updates, which reduces the overhead of managing security patches and bug fixes.

1. Log in to the Appliance Administration console using the root account at `https://hostname:9443` .
2. Click Online Update.
3. The Registration dialog should display. If not, click Register.
4. Select Micro Focus Customer Center as the service type.
5. Specify the following information about the account for this appliance:

- Email address of the account in the Customer Center.
- Activation key. To obtain the key:

  a. Log in to Software Licenses and Downloads (SLD) portal (page .

  b. Click the Activations tab.

  c. Locate your product.

  d. Click `Download <Appliance Update Channel Activation Key.txt>`

  e. Open the file to view the activation key.

6. Select an option to share information with Open Text:

Hardware profile

Optional information

7. Click Register.

8. Wait while the appliance registers with the service, then click OK.

   You can now view a list of the needed and installed updates. You can use manual or automatic options to update the appliance.

Installing your product into the appliance

A license entitlement file (activation file) is required to install products in the appliance.

1. Download the current version of the activation file for each product from the Downloads site (where you downloaded this appliance).

2. Log in to the appliance console using the root account at `https://hostname:9443`.

3. Click Products.

4. Click Choose Files and browse to the activation file for each product you want to install.

> 💡 **Note**
>
> At least one activation file for a product, such as Management and Security Server, must be included in the selection.

5. Click Install.

While it may take several minutes for your product to start up and become accessible, you can monitor the cluster status in the appliance console Cluster view. Proceed when the status is `Ready`.

Updating the Appliance

What's required for updates?

- An activation key to register the update channel.
- Each node in the cluster must be in a ready state before you attempt to update.

Registering the appliance for software updates

To receive online updates, which reduce the overhead of managing security patches and bug fixes, register the appliance. See for instructions.

Manage appliance software updates

Software updates are delivered through two mechanisms in the appliance:

- Online Update: Delivers security updates to the OS and installed products. This should be used regularly to keep your system up to date.
- Upgrade: Delivers more significant upgrades to the installed products. Product upgrades require a new activation key and should only be done after proper planning.

> **✏ Notes**
>
> To supervise system changes, we recommend manually updating your appliance and not using the automatic scheduling feature.
>
> Both Online Updates and Upgrades occasionally require rebooting the appliance. A "Reboot Needed" option is displayed in the upper right corner of the Appliance Administration console when this is called for.

Preparing to update

- Be prepared to supply the email address and activation key that were used during initial registration.
- To ensure easy recovery in case of errors, take a snapshot of the current configuration before updating.

> **⚠ Caution**
>
> During the update process, the cluster will be unavailable for end users. Plan your maintenance window accordingly.

Installing updates and upgrades

To install either online updates or upgrades, first ensure all nodes in the deployment are in a `Ready` state using the Cluster view in the Appliance Administration console. Then perform the following steps.

1. In the Cluster view, click Cordon All Nodes. The status of each node will change to `Ready/SchedulingDisabled`. This step can be safely skipped if this button is not available; it is only available in more recent versions of the appliance.

2. On each node in the cluster, update one node at a time by repeating the following steps:

   a. In the Appliance Administration console, click Online Update.

   b. Click Update Now to install the updates.

   c. After the updates are installed, click Close.

   d. If the Reboot button is highlighted, click it to restart the appliance.

   e. If the Upgrade button shows a badge indicating an Upgrade is available, and you wish to upgrade product versions at this time, click Upgrade. If not, skip this step.

   Click Start then review the license.

   Register using the email address used during appliance registration and the new activation key.

   Click OK on the Update Now dialog. Wait while the upgrade is performed.

f. Click Reboot. After the reboot finishes, in the Appliance Administration console click Cluster. Under Cluster Status, wait until the updated node shows a status of `Ready`. It can take up to 15 minutes for the node to become ready.

Throughout the cluster update process, it is normal to see warnings and errors in the lower sections of the Cluster view. These will clear once the entire update process is complete.

g. Click Online Update again to check for and install any new updates that are available. If updates were installed and a reboot was required, wait for the node to show `Ready` again in the Cluster View.

h. Move on to the next node in the cluster.

3. Once all of the nodes in the cluster have been updated and are ready, in the Cluster view click Scale Cluster on any node.

4. Wait for the cluster to return to a healthy state with all nodes showing `Ready,SchedulingDisabled`, or `Ready` if Cordon was not available. This process can take up to 15 minutes.

5. In the Cluster view, click Uncordon All Nodes. This step can be safely skipped if this button is not available.

6. Once all nodes return to a `Ready` status, the cluster is ready for use.

Using the Subscription Management Tool (SMT) to manage appliance updates

You can use the Micro Focus Subscription Management Tool (SMT), version 2.0, to provide appliance updates on SLES 15 SP5 or OpenSUSE Leap 15 SP5 platforms.

- Learn about SMT (page
- Installing the SMT server (page
- Creating a certificate (page

SMT 2.0 does not automatically create TLS certificates to be used by Apache. You can create certificates manually before configuring SMT.

For example:

```
openssl req -x509 -newkey rsa:4096 -keyout /etc/ssl/servercerts/serverkey.pem -out /etc/ssl/servercerts/servercert.pem -sha256 -days 3650 -nodes -subj "/C=US/ST=WA/L=Tacoma/O=OpenText/OU=CompanySectionName/CN=sles15.microfocus.com"
```

Replace `CN` with your own value.

After successfully installing the SMT server locally and creating the certificates:

1. In appliance console, click Online Update.

2. Select Local SMT as the service type.

3. Specify the fully qualified SMT hostname, for example, `smt.microfocus.com`.

4. Click Register. It will take a few minutes for the updates to become available.

Linux Installer

Linux Installer - System Requirements

The Linux servers are provided by you. The Linux installer installs the appliance administration console, along with your product.

Supported Operating Systems

The following operating systems are supported. These versions or greater.

SUSE Linux Enterprise Server 15 SP5

OpenSUSE Leap 15.5

Red Hat Linux 9

Rocky Linux 9

Oracle Linux 9

AlmaLinux 9

Minimum CPU and memory requirements

The following minimum resources are required for each node. These requirements assume that no other production software is installed on the node. If additional software will be run on the node, more resources need to be added to accommodate the other software accordingly.

8 CPU Cores

16 GB RAM - with swap space disabled

100 GB disk space (SSD) with 80 GB delegated to `/var/opt` and 20 GB for `/opt`

Fast storage

To ensure optimal performance and reliability, the use of a solid-state drive (SSD) or other fast storage solutions is required. Not using SSD-based storage may lead to inconsistent behaviors and errors.

Disable swapping

For optimal performance and reliability, swap must be turned off on every node. Please refer to the specific documentation of your Linux distribution for guidelines on how to accomplish this.

Fixed IP address

A fixed, non-changing IP address is required for each node. DHCP (Dynamic Host Configuration Protocol) is supported but the IP must be reserved and cannot change.

Network ports

The following ports must be exposed and available between all nodes:

Port	Purpose
6443	Kubernetes API Server
8472	Virtual LAN
10250	Kubernetes metrics
2379-2380	etcd

The following ports must be exposed for outside access:

Port	Purpose
443	Product access
3000	Security Proxy Server *
8001	AJP **

* The Security Proxy port use is optional.

** The AJP port is used when optionally integrated with Microsoft's IIS web server.

Additional firewall rules

The following source IP ranges must be added to the trusted zones list:

Source IP Range	Purpose
10.42.0.0/16	Pod communication
10.43.0.0/16	Service communication

Required Third Party Packages

Several packages are installed automatically during product installation. Certain platforms require platform-specific repositories. Ensure the following repositories are configured on these platforms:

For Red Hat, the `epel-release` (Extra Packages for Enterprise Linux) repository is required.

For OpenSUSE, the SUSE Linux Enterprise (sle) repository is required.

The following packages are automatically installed or updated as needed. Some of these packages are platform-specific and are installed only on the applicable platform: `bash`, `curl`, `grep`, `gawk`, `wget`, `jq`, `haveged`, `zip`, `bind-utils`, `sysstat`, `strongswan`, `apparmor-parser`, `util-linux`, `iscsi-initiator-utils` or `open-iscsi`, `nfs-utils` or `nfs-common`, `supportutils` or `sos`

Supported web browsers

The following web browsers are supported:

Google Chrome (recommended)

Mozilla Firefox (recommended)

Microsoft Edge

Installing using the Linux Installer

Installation Steps

To install your product, first make sure your system meets the system requirements (page 17) , and then perform the following steps:

1. From the Downloads site, download the Linux installer script (`install*.sh`) for MSS.

2. Enable execute permissions for the installer: `chmod 744 install*.sh`

3. Ensure that an operating system firewall is not blocking any required ports and that masquerading is enabled.

4. With elevated privileges (for example, `sudo`), run the Linux install script (`.sh`) to install the product.

5. A PGP key is used to verify that the file you are downloading has not been manipulated by a third party. If the displayed signing information represents a known and trusted entity, such as Micro Focus, then enter `y` to install the public key and continue.

   Alternatively, refer to this Knowledge Base article (page to download the key separately and verify the file.

6. When the installation completes, a verification tool is automatically executed.

7. If verification *succeeds* then the services will start automatically; continue with the next step (page 22) .
   If verification *fails*, see Troubleshooting the Linux installation (page 19) .

Troubleshooting the Linux installation

Symptom: "Permission denied" messages with references to "zgrep" in the output.
Possible fix: Check that the AppArmor profile for `zgrep` is not too restrictive for the verification process.

Once the issues are addressed, run `sudo cspctl start` to start the system. Then run `sudo cspctl enable` to have the system start automatically after the server restarts.

If issues remain, please contact Customer Support for assistance.

Upgrading using the Linux installer

When upgrading, it is important to remove any activation files from MSS associated with previous versions of Host Access for the Cloud. Leaving obsolete activation files in place may result in limited access to sessions.

What's required before upgrading?

Administrative privileges for the operating system.

The cluster will be unavailable for end users during the upgrade process. We recommend planning a maintenance window accordingly.

Each node in the cluster must be in a `Ready` state before you attempt to upgrade.

Upgrade steps

To upgrade your product, first ensure all nodes in the deployment are in a `Ready` state by running: `cspctl status` with elevated privileges. Then perform the following steps:

1. From the Micro Focus download site, download the Linux installer script (`install*.sh`) for your product.

2. Enable execute permissions for the installer:

   ```
   chmod 744 install*.sh
   ```

3. On any any node in the cluster run: `cspctl cluster cordon`. The `STATUS` of each node will change to `Ready/SchedulingDisabled`. This command can be safely skipped if it is not available; it is only available in more recent versions of the cspctl.

4. On each node in the cluster, update one node at a time by repeating the following steps:

   a. Copy the installer to the node, run the Linux install script (.sh) with elevated privileges, (for example, sudo), to upgrade the product.

   b. After the upgrade is complete, the verification tool automatically runs.

   If verification succeeds, the services will automatically start.

   If verification fails, review the troubleshooting steps (page 19) .

   c. After the CSP service starts, wait until the updated node shows a `STATUS` of `Ready/SchedulingDisabled`.

   Throughout the cluster upgrade process, it is normal to see warnings and errors in output of `cspctl status`. These will clear once the entire upgrade process is complete.

   It can take up to 15 minutes for the node to become `Ready`.

   d. Move on to the next node.

5. Once all nodes in the cluster have been updated and are ready, run: `cspctl cluster scale` on any node.

6. Wait for the cluster to return to a healthy state with all nodes showing `Ready,SchedulingDisabled`, or `Ready` if cordon was not used. This process can take up to 15 minutes.

7. On any node run: `cspctl cluster uncordon` to allow all nodes to become schedulable again. This command can be safely skipped if it is not available.

8. Once all nodes return to a `Ready` status, the cluster is ready for use.

Uninstalling

The method of uninstalling depends on the deployment method you used to install your product.

> 💡 **Note**
>
> Before uninstalling, always remove the node from the cluster:
>
> Open the MSS Administrative Console (`https://hostname/adminconsole`).
>
> Click Cluster Management > Nodes.
>
> Next to the node you want to remove, click ⋯ Delete.

Virtual Appliance method

To uninstall a product:

1. Open the Appliance Administration Console (`https://hostname:9443`) > Products.

2. Next to the product you wish to uninstall, click ⋯ Uninstall.

   This process takes a while to complete.

Linux installer

To uninstall, run `sudo /opt/opentext/csp/uninstall-mss.sh` .

The uninstall process takes a while to complete.

Add-On Products

To uninstall an MSS Add-on product:

1. Open the MSS Administrative Console > Configure Settings - Product Activation.

2. Click and Remove the product.

Next Step After Installing

Once your product is installed, you are ready to configure your deployment (page 23) .

Configure

Configure Your Deployment

After installing, you have a cluster of one, a single node. The next steps are to configure key cluster settings and then add more nodes to your cluster. These settings can be set at a later time, but we recommend setting them during initial configuration.

Configure Your Cluster

The MSS Administrative (Admin) Console is a central location for system and product configuration. First, use the MSS Admin Console to access Cluster Management, where you will set key cluster settings. Later, use the MSS Admin Console to further configure your product(s).

Access the MSS Admin Console

To access the MSS Admin Console:

1. Log in to `https://hostname/adminconsole` .

2. The Admin Console's default password is `admin` .

3. Once signed in, various views can be loaded using the drop-down menu.

Set the cluster DNS name

1. Register a name in your DNS system that points to your load balancer. If not using a load balancer, the name should resolve to all nodes in your cluster.

2. Log in to the MSS Admin Console at `https://hostname/adminconsole` .

3. From the drop-down menu, click Cluster Management.

4. Click Settings.

5. Set the Cluster DNS Name and click Apply.

6. Use this hostname for accessing all services in the cluster.

Set the cluster certificate

1. Log in to the MSS Admin Console at `https://hostname/adminconsole` .

2. From the drop-down menu, click Cluster Management.

3. Click Settings.

4. Expand the Certificate and Private Key panels and import the certificate and key pair.

5. Click Apply. The cluster certificate will be applied to all cluster endpoints. This may take a few minutes.

Clustering

Your initial installation is a cluster of one, a single node. Now repeat the installation process to create three or more nodes, always ending with an odd number of nodes, as described in the standard deployment (page 5) .

Once you have a set of nodes, the next step is to cluster them together.

> ⚠️ **Caution**
>
> Before proceeding with clustering be aware that:
>
> - Before adding a node to a cluster, all nodes must be in a healthy state.
> - The node that joins a cluster loses its own application data, such as configured sessions. The data present on the node that you are joining is inherited.
> - Before joining a node to a cluster, the node must have the same products installed as those nodes that already participate in the cluster.
> - Removing a node from a cluster results in its data being lost.
> - Removing one healthy node from the cluster results in data loss and the need for a reset, but the remaining node will remain functional.
> - When later replacing a node in a cluster, always remove the existing node before adding a new node.

Follow the clustering steps for your method of deployment: Appliance (page 25) or Linux installer (page 25)

Clustering when using the appliance

To join a new appliance to an existing appliance cluster:

1. Log in to the Appliance Administration console using the root account at `https://hostname:9443` .

2. Click Cluster.

3. Specify the DNS hostname or IP address of the remote appliance to which you are clustering.

4. Specify `root` as the username and enter the password for the root user on the remote appliance.

5. Click Join Cluster.

   The Cluster Status will display a list of all nodes in the cluster with a status of `Ready` when clustering is complete. **The process takes 5-15 minutes to complete**.

Clustering when using Linux installations

To join a new Linux node to an existing cluster:

1. On a node that exists in a cluster, note the following:

   • The hostname or IP address of the host

   • The cluster join token, which is obtained by executing: `sudo cspctl cluster token`

2. On the node that is joining the cluster, execute the following:

   ```
   sudo cspctl cluster join -s <hostname> -t <token>
   ```

   Note that the `hostname` and `token` values were obtained from the existing node in the cluster you are joining (step 1). **The process takes 5-15 minutes to complete**.

To remove a node

To remove a node from a cluster, please refer to the **Cluster Management** help:

1. In the MSS Administrative Console, click Cluster Management from the drop-down menu.

2. Click Nodes, and then open Help (?).

3. Click Nodes and follow the steps to delete a node.

Next Steps After Configuring Your Cluster

You now have a cluster ready for use. You are ready to .

Apply

Apply your Product Configuration

After you install MSS and configure your deployment, you are ready to use MSS.

Information for MSS is always available from the MSS Documentation (page site or from the online Support Resources (page .

Using the MSS components and Add-on Products

Follow the steps for using each component or entitled product .

Using Metering

Use the Metering Server to monitor session activity and to control concurrent access to specific hosts. Metering Reports are available as clients use the metered sessions.

The Metering Server is installed with Management and Security Server (MSS). No separate license is required.

Configure Metering

Metering is configured with the Metering Console.

1. From the MSS Administrative Console drop-down menu, click Cluster Management.

2. Click Services.

3. Next to the `mss-metering` service, hover over and click the link to the **Metering Console**.

   When prompted for a password, enter the MSS administrative password. (If you changed the password, then enter the new one.)

4. Use the Metering Console to configure license pools and server settings and to run reports. Open Help for assistance.

5. Enable the **clients** that are to be metered.

   Refer to your emulator's product documentation to enable metering for that client.

View Metering Reports

After metering is configured and users begin to launch client sessions, you can monitor activity by viewing reports.

1. Open Cluster Management > Services.

2. Next to the `mss-metering` service, hover and click the link to **Metering Reports**.

3. When prompted for a password, enter the MSS Administrator password.

> ○ **Note**
>
> You can set a password for non-administrators to view Metering Reports.

4. Several reports are available:

activity by user, machine, IP address, and other attributes

concurrent usage (to comply with your license)

host connections

How Metering Works

When the configuration is complete, here's how the Metering Server communicates with the metered client.

1. A user starts a client session and initiates a host connection.

2. The session requests a license from the Metering Server, and once granted, the host connection proceeds, and the Metering Server begins to record product usage.

3. The session sends updates to the Metering Server at regular intervals until the user closes the session.

4. The metering data is available for the administrator to generate reports.

Using Security Proxy Server

When you use the Security Proxy Server, data sent between the client session and the Security Proxy is TLS-encrypted. The host is protected from direct user contact.

The Security Proxy is a MSS add-on product that can be used by Reflection Desktop and Reflection for the Web. The Security Proxy is automatically installed when you install the appliance or use the Linux installer, but it must be

- **activated** so it can be managed by MSS
- **configured** to trust MSS

Refer to Security Proxy Server (page (in the *MSS product help*) for detailed instructions:

Enabling the Security Proxy Server

Configuring the Security Proxy

Advanced Configuration

Setting the Logging Level

Using FIPS-Approved Mode

Using Terminal ID Manager

The Terminal ID Manager lets you centrally manage and assign terminal and device IDs to emulator sessions. You can pool terminal IDs, track ID usage, and manage inactivity timeout values for specific users, thus conserving terminal ID resources and significantly reducing operating expenses.

The Terminal ID Manager Add-On requires a separate license and an activation file.

Terminal ID Manager is automatically installed, and it needs to be activated and configured.

Refer to the Terminal ID Manager Guide (page to complete the configuration.

Using Automated Sign-On for Mainframe

Automated Sign-On for Mainframe is an add-on product that enables you to configure user access to z/OS mainframe applications using a single login. This add-on product requires a separate license.

Prerequisites and System Requirements

Before installing or configuring Automated Sign-On for Mainframe, the following requirements must be met:

Management and Security Server (the Administrative Server) is installed.

Terminal emulation software, such as Reflection Desktop, is installed on the client and administrator's workstations.

The Automated Sign-On for Mainframe Add-On activation file is available (after purchase).

z/OS with DCAS is installed on the mainframe.

LDAP directory is used for user authorization.

Installing Automated Sign-On for Mainframe

The Automated Sign-On for Mainframe Add-on product is installed with an activation file. Follow these steps.

1. After purchasing Automated Sign-On for Mainframe Add-On, you will receive information about downloading the product activation file: `activation.automated_signon_for_mainframe-<version>.jaw`

2. Download the activation file and note the location.

3. Open the MSS Administrative Console and click **Configure Settings - Product Activation**.

4. Click **ACTIVATE NEW** and browse to `activation.automated_signon_for_mainframe-<version>.jaw` .

5. Click the file. The **Automated Sign-On for Mainframe Add-On** is installed and added to the **Product** list.

6. Restart your browser to ensure that the MSS Administrative Console is fully updated with the new set of activation files. You do not need to restart the Administrative Server (MSS Server) service.

Next step

After Automated Sign-on for Mainframe Add-On is activated, settings must be configured on different systems:

z/OS

Management and Security Server

your emulator

Refer to the Automated Sign-on for Mainframe - Administrator Guide (page for detailed steps.

Using MSS Automated Sign-On for Host Access

MSS Automated Sign-On for Host Access is an add-on product that enables you to configure user access to host applications using a single login. This add-on product requires a separate license.

> ♀ **Note**
>
> If you are using a *z/OS* system, see *Setting Up Automated Sign-On for Mainframe (page 30)*

Prerequisites and System Requirements

Before installing or configuring MSS Automated Sign-On for Host Access, the following requirements must be met:

- Management and Security Server (the Administrative Server) is installed.
- Terminal emulation software, such as Reflection Desktop, is installed on the client and administrator's workstations.
- The MSS Automated Sign-On for Host Access Add-On activation file is available (after purchase).
- LDAP directory is used for user authorization.

Installing MSS Automated Sign-On for Host Access

The MSS Automated Sign-On for Host Access Add-on product is installed with an activation file. Follow these steps.

1. After purchasing MSS Automated Sign-On for Host Access Add-On, you will receive information about downloading the product activation file: `activation.automated_signon_for_hostaccess-<version>.jaw`

2. Download the activation file and note the location.

3. In the Management and Security Server, open the MSS Administrative Console and click **Configure Settings - Product Activation**.

4. Click **ACTIVATE NEW** and browse to `activation.automated_signon_for_hostaccess-<version>.jaw`.

5. Click the file. The **MSS Automated Sign-On for Host Access Add-On** is installed and added to the **Product** list.

6. Restart your browser to ensure that the MSS Administrative Console is fully updated with the new set of activation files. You do not need to restart the MSS Server service.

Next step

After Automated Sign-on for Host Access Add-On is activated, settings must be configured on different systems:

your host

Management and Security Server (MSS)

your emulator

Refer to the technical reference, **Configuring MSS Automated Sign-on for Host Access**, in the *MSS Administrator Guide (page* for details.

Using Micro Focus Advanced Authentication Add-On

This Add-on product enables MSS to use the Micro Focus Advanced Authentication product.

The Advanced Authentication product enables strong multi-factor authentication using a variety of authentication methods, including biometrics, one-time passwords, and smartphone authentication.

Prerequisites and System Requirements

Before installing and configuring the Micro Focus Advanced Authentication Add-On, verify that:

Management and Security Server (MSS) is installed.

Micro Focus Advanced Authentication Add-On is licensed.

The Micro Focus Advanced Authentication server is installed on a separate machine.

> 🔥 **Tip**
>
> Note the server name (or IP address) and the server's port number.

Installing and Configuring Micro Focus Advanced Authentication Add-On

Three basic steps are required to install and configure the Micro Focus Advanced Authentication Add-On. Before starting, make sure you've met the .

Step 1: Install Micro Focus Advanced Authentication Add-On

The Advanced Authentication Add-On is installed with an activation file, as follows.

1. After purchasing Micro Focus Advanced Authentication Add-On, you will receive information about downloading the product activation file: `activation.advanced_authentication-<version>.jaw`

2. Download the activation file and note the location.

3. In the Management and Security Server, open the MSS Administrative Console and click **Configure Settings - Product Activation**.

4. Click **ACTIVATE NEW** and browse to `activation.advanced_authentication-<version>.jaw` .

5. Click the file. The **Advanced Authentication Add-On** is installed and added to the **Product** list.

6. Restart your browser to ensure that the MSS Administrative Console is fully updated with the new set of activation files. You do not need to restart the Administrative Server (MSS Server) service.

Step 2: Set up Advanced Authentication in the MSS Administrative Console

In the MSS Administrative Console:

1. Open **Configure Settings - Authentication & Authorization**, and click **Micro Focus Advanced Authentication**.

2. Open **Help** [?] and follow the steps to configure Advanced Authentication.

Step 3: Configure authentication methods

To configure Advanced Authentication methods, such as Voice, refer to your Micro Focus Advanced Authentication server documentation (page .

Technical References

Technical References

Advanced Settings

Information for your product is always available from the installed documentation or from online resources (page . See the list (on the right) for the settings documented on this page.

Adjusting product settings

You may occasionally need to change or add properties to your product services. Properties are set in the Cluster Management console.
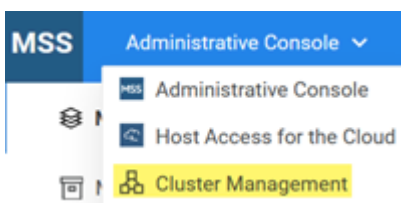
Follow these steps:

1. Log in to the MSS Admin Console at https://hostname/adminconsole, and click Cluster Management from the drop-down menu.

2. Click Services.

3. Click the service of interest, and click ⋯ Edit Properties.

4. Add or edit the key and value accordingly.

5. After all properties are adjusted, redeploy the service.

> **🔥 Important**
>
> Be aware that redeploying services may affect end users who are accessing the service.

Kubernetes dashboard

The Kubernetes dashboard is a web-based interface where you can monitor applications running in a cluster, specify or modify resources, and troubleshoot issues.

To use the Kubernetes dashboard:

1. Log into the MSS Admin Console at `https://hostname/adminconsole` .

2. From the drop-down menu, click Cluster Management.

3. Click Advanced.

4. Slide the button to enable the dashboard.

5. Copy the authentication token, click the URL, and paste the token into the field provided.

6. Click Sign in.

Kubernetes configuration file

The KubeConfig file is available for advanced configuration purposes.

Click **Download KubeConfig File** from the Cluster Management > Advanced page to use with the Kubernetes command line tool, `kubectl` .

You can install the `kubectl` tool from the Kubernetes project at https://kubernetes.io/ (page .

Shell into an application instance (pod) in Kubernetes

1. Log in to the Kubernetes Dashboard (page 35) .

2. Under Workloads, click Pods.

3. Use the Name column to locate the pod of interest, and click ⋮ .

4. Click Exec to use a shell to access the pod's file system.

Add Kubernetes tools to the environment (Linux installation)

> 🔥 **Important**
>
> These steps apply only to a Linux installation. If you installed using the appliance, they do not apply.

By default, Kubernetes tools are not added to the system's environment. To simplify access to these tools, follow these steps:

Log in to the Linux host.

```
cd /opt/opentext/csp/bin
sudo -s
. ./env.sh
```

> 💡 **Note**
>
> The first dot, known as the "dot operator," executes the script in the current shell instead of a subshell, ensuring that the current shell's environment is updated.

Enable the SSH service for the appliance

1. Log into the Appliance Administration console using the root account at `https://hostname:9443`.

2. Click System Services.

3. Select SSH, and then from the Action menu, click Start.

   To automatically start the SSH service after system restarts, click Options, and Set as Automatic.

Migrating data from Legacy Deployments

You can migrate your data from a legacy installation to the new container-based deployment.

Use the new migration tool to export data from your previous installation into a zip file. Then import the data into the new installation.

What's required?

- The existing data must be on a current major release of your product.

- OS administrative privileges to run the migration tool.

- A new single-node installation to import the data.

- To run the migration tool,

Data that is NOT migrated

kerberos settings

metering report data

security proxy configuration

- passwords

  For example the MSS Admin password will remain the same before and after migration.

Run the migration tool

Log in to the MSS Admin Console at `https://hostname/adminconsole` . Click Configure Settings > Migration.

Open the help for more information and detailed migration steps.

Next step

After your data is migrated, you are ready to configure your cluster (page 23) .

Integrating with IIS

Integrating with IIS

A default web server is installed for application access. Integration with IIS is used for the following purposes:

- IIS Single Sign-on authentication.

- SiteMinder authentication.

- Use existing web server certificates on IIS.

- Comply with Common Criteria security requirements.

Using IIS Single Sign-on for authentication

Integration with IIS is a three step process.

Integrate MSS with IIS using the ISAPI redirector (page 38) .

Integrate products with IIS using the IIS reverse proxy (page 42) .

Refer to the *MSS Administrator Guide* to configure MSS for Single Sign-on for IIS authentication (page .

Using IIS for SiteMinder authentication

Integration with IIS is a three step process.

A SiteMinder Web Agent is installed on IIS and is configured to protect web resources. Refer to the SiteMinder documentation for more information.

Integrate with IIS using the IIS reverse proxy (page 42) .

Refer to the *MSS Administrator Guide* to configure MSS for SiteMinder authentication (page .

Using IIS for other purposes

If you require an IIS front-end for other purposes, integrate IIS using the IIS reverse proxy (page 42) .

Integrate MSS with IIS for IIS Single Sign-On

This section describes how to integrate MSS with IIS, for the purpose of using IIS Single Sign-On.

> 💡 **Note**
>
> When integrated with IIS, a common/shared certificate and private key is used to provide security (HTTPS) for the Cluster DNS entrypoint and IIS.

Refer to the Microsoft IIS documentation for instructions on how to install IIS and its features.

Requirements

IIS 8.0 or higher

IIS features that must be enabled:

ISAPI Extensions

ISAPI Filters

Windows authentication

A common/shared certificate and private key pair that will be used by both the Cluster DNS entrypoint and the IIS website.

Check your firewall settings to ensure that requests from the IIS server to the cluster's AJP port are allowed (default is 8001).

DNS resolution must be working properly between IIS and the Cluster DNS name for transparent Single Sign-On to succeed.

IIS integration steps

Download and install the ISAPI redirector

1. Download the ISAPI redirector DLL from the Apache Tomcat website at https://dlcdn.apache.org/tomcat/tomcat-connectors/jk/binaries/windows/ (page

2. Select the ZIP file for x86-64, unless a different platform is required.

3. On the machine where IIS is installed, create a directory that will be used to contain the redirector files.

4. Unzip the redirector files into the directory.

5. Create a copy of the file named `isapi_redirect.dll`. Rename the copied file to `isapi_redirect_sec.dll`.

Create a configuration file for the redirector

1. In the same directory, create a file named `isapi_redirect.properties`.

2. Copy this content to the file:

```
worker_file=workers.properties
worker_mount_file=uriworkermap.properties
log_level=emerg
log_file=iis_redirect.log
extension_uri=/tomcat/isapi_redirect.dll
```

Create a configuration file for the second redirector

1. In the same directory, create a file named `isapi_redirect_sec.properties`.

2. Copy this content to the file:

```
worker_file=workers.properties
worker_mount_file=uriworkermap_sec.properties
log_level=emerg
log_file=iis_redirect_sec.log
extension_uri=/tomcat/isapi_redirect_sec.dll
```

Create a worker file for the redirector

1. In the same directory, create a file named `workers.properties`.

2. Copy this content to the file:

```
ps=\
worker.list=ajp13_worker
worker.ajp13_worker.port=8001
worker.ajp13_worker.host=
worker.ajp13_worker.type=ajp13
worker.ajp13_worker.secret=changeit
worker.ajp13_worker.lbfactor=1
worker.loadbalancer.type=lb
worker.loadbalancer.balanced_workers=ajp13_worker
```

3. Specify the cluster DNS name as the value for the property named `worker.ajp13_worker.host`.

Create a URI mapping file for the first redirector

1. In the same directory, create a file named `uriworkermap.properties`.

2. Copy this content to the file:

```
default.worker=ajp13_worker
/mss|/*=$(default.worker)
/tidm|/*=$(default.worker)
/adminconsole|/*=$(default.worker)
!/adminconsole/plugins|/*=$(default.worker)
/login|/*=$(default.worker)
/sessions|/*=$(default.worker)
```

Create a URI mapping file for the second redirector

1. In the same directory, create a file named `uriworkermap_sec.properties`.

2. Copy this content to the file:

```
default.worker=ajp13_worker
/iisred|/*=$(default.worker)
```

Add the virtual directory to IIS

1. Open the IIS Manager application.

2. In the Connections pane, expand the tree to view the website to integrate.

3. Right-click the website to integrate, and click Add Virtual Directory...

4. Specify an Alias value of `tomcat`.

5. For the Physical path value, browse to the directory that contains the ISAPI redirector DLL files.

6. Click OK to close the dialog.

7. In the Connections pane, right-click the `tomcat` virtual directory, and select Edit Permissions...

8. Click the Security tab, Edit..., and then click Add...

9. In the Enter the object names to select box, add the following local Groups:

IUSR

IIS_IUSRS

Click OK to close each dialog in succession.

Enable execution of IIS Handler Mappings

In the IIS Manager application:

1. In the Connections pane, select the tomcat virtual directory.

2. In the tomcat Home view, double-click Handler Mappings.

3. In the Actions view, click Edit Feature Permissions...

4. Select the Execute checkbox and then OK.

Configure Windows authentication for the second ISAPI redirector

In the IIS Manager application:

1. In the Connections pane, right-click the tomcat virtual directory and select Switch to Content View.

2. In the tomcat Content view, right-click `isapi_redirect_sec.dll` and select Switch to Features View.

3. In the `isapi_redirect_sec.dll` Home view, double-click Authentication.

4. Disable Anonymous Authentication.

5. Enable Windows Authentication.

Configure the ISAPI Filters for the website

In the IIS Manager application:

1. In the Connections pane, under the Sites node, click the website to integrate.

2. In the website's Home view, double-click ISAPI Filters.

3. In the Actions pane, click Add..., and specify the following values:

Name: `isapi_redirect`

Executable: Browse and select the file named `isapi_redirect.dll`

Click OK to close the dialog.

4. In the Actions pane, click Add..., and specify the following values:

Name: `isapi_redirect_sec`

Executable: Browse and select the file named `isapi_redirect_sec.dll`

Click OK to close the dialog.

5. In the Actions pane, click View Ordered List...

6. Ensure that `isapi_redirect` is at the top of the list, and that `isapi_redirect_sec` is second in the list.

Enable ISAPI extensions for IIS

In the IIS Manager application:

1. In the Connections pane, select the top-most node, for the IIS server.

2. In the server's Home page, double-click ISAPI and CGI Restrictions.

3. Add the `isapi_redirect.dll` and `isapi_redirect_sec.dll` files, and for each select the checkbox to Allow extension path to execute.

4. Restart IIS.

> 🔥 **Important**
>
> Integration is not complete. The MSS Admin Console plug-ins and products will not work until an IIS Reverse Proxy is added -- the next step.

Next Step: Configure the IIS Reverse Proxy

Next, configure the IIS Reverse Proxy (page 42) for products and MSS Admin Console plug-ins.

Troubleshooting

Ensure firewalls are not interfering with connections between IIS and MSS.

You may need to restart IIS or redeploy MSS.

Ensure the user is authenticated with IIS, prior to accessing MSS or any other applications.

Inspect the IIS logs and Windows Event logs for any information regarding issues.

Enable logging for the Tomcat ISAPI redirector and inspect the logs.

If an HTTP 500 error is encountered, launch a browser from the IIS host to obtain richer information from IIS about the failure.

Integrate with the IIS Reverse Proxy

Integration with an IIS reverse proxy is used for the following purposes:

Enable IIS access to products like Host Access for the Cloud, Reflection for the Web, or Metering.

Use SiteMinder for authentication.

Comply with Common Criteria security requirements.

> ⓘ **Note**
>
> When integrated with IIS, a common/shared certificate and private key is used to provide security (HTTPS) for the Cluster DNS entrypoint and IIS.

Refer to the Microsoft IIS help documentation for instructions on how to install IIS and its features.

Requirements

IIS 8.0 or higher

The IIS URL Rewrite module must be installed. See IIS URL Rewrite (page for information on how to install this.

IIS Application Request Routing (ARR) 3.0 or later is required. See IIS Application Request Routing (page for information on how to install this.

**Additional Requirement when using Host Access for the Cloud**

The IIS WebSockets protocol must be enabled. See IIS 8.0 WebSocket Protocol Support (page for information on how to enable this protocol.

Configure the IIS Reverse Proxy

1. Launch the Internet Information Services (IIS) Manager, navigate to the website you want to use, and open the URL Rewrite feature.
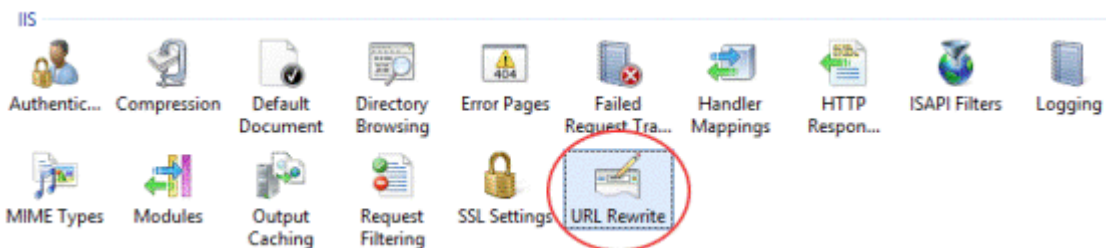
2. Choose the Add Rule(s) action and add the Reverse Proxy rule.
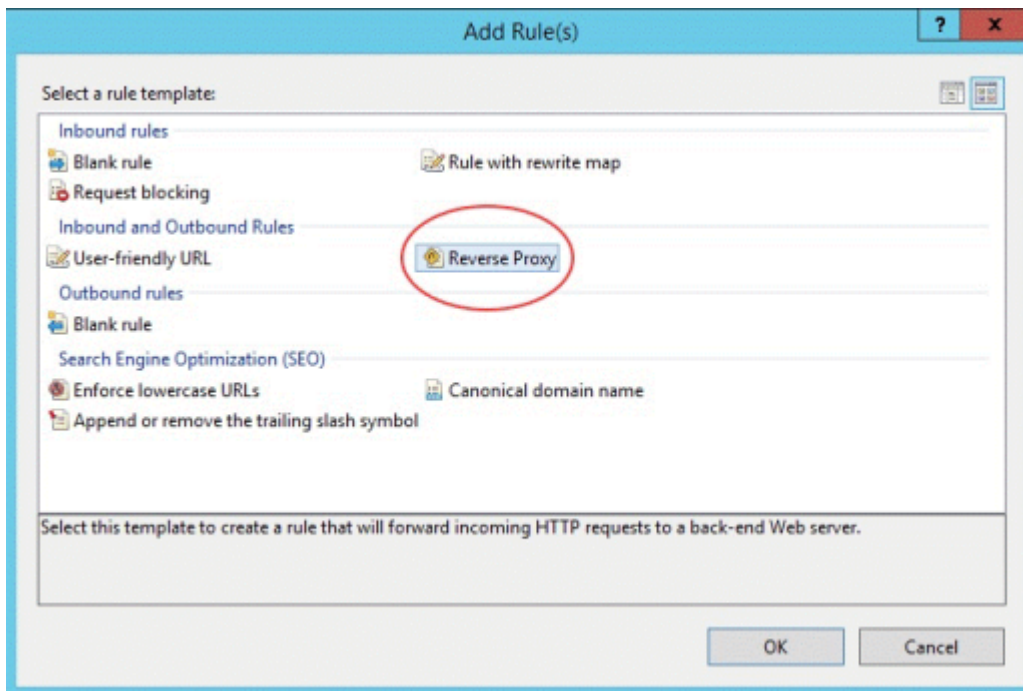
3. For the inbound rule, enter the Cluster DNS name, and de-select Enable SSL Offloading.

4. Check the outbound rule Rewrite the domain names... and enter the hostname or IP address of the IIS server in the To: box.

5. Click OK to create the new Reverse Proxy Rule.

Configure MSS to use IIS as the front-end

1. Log in to the MSS Admin Console at `https://clusterdns/adminconsole` .

2. From the drop-down menu, click Cluster Management.

3. Under Services, find the service named `mss-mss-server` .

4. Click `...` Edit Properties, and add this key/value pair:

```
management.server.iis.url
```

```
https://<IIS server NetBIOS hostname>/mss
```

Example: `https://iishostname/mss`

> **♀ Note**
>
> When using IIS Single Sign-On authentication, use a hostname without dots in the `management.server.iis.url` property. This will help your browser recognize the IIS host as part of your intranet and enable transparent Integrated Windows Authentication (IWA).
>
> Alternatively, you can use a hostname with dots, like an FQDN. However, in this scenario, you need to configure your browser to trust the site for transparent Integrated Windows Authentication (IWA) to function properly. Refer to your browser's Help documentation for detailed instructions.
>
> This note does not apply to SiteMinder authentication.

5. Click ••• Redeploy All.

> **♀ Note**
>
> Be aware that end users may be affected when redeploying a service.

Additional configuration when using Host Access for the Cloud

1. Log in to the MSS Admin Console at `https://clusterdns/adminconsole` .

2. From the drop-down menu, select Cluster Management.

3. Under Services, find the service named `hacloud-session-server` .

4. Click ••• Edit Properties, and add the following key/value pairs:

```
websocket.compression.enable
```

```
false
```

```
server.compression.enabled
```

```
false
```

```
websocket.allowed.origins
```

```
https://<IIS server name or IP address>
```
Example: `https://iishostname`

```
management.server.iis.url
```

```
https://<IIS server name>/mss
```
Example: `https://iishostname/mss`

> **💡 Note**
>
> When using IIS Single Sign-On authentication, use a hostname without dots in the `management.server.iis.url` property. This will help your browser recognize the IIS host as part of your intranet and enable transparent Integrated Windows Authentication (IWA).
>
> Alternatively, you can use a hostname with dots, like an FQDN. However, in this scenario, you need to configure your browser to trust the site for transparent Integrated Windows Authentication (IWA) to function properly. Refer to your browser's Help documentation for detailed instructions.
>
> This note does not apply to SiteMinder authentication.

5. Click ••• Redeploy All.

> **💡 Note**
>
> Be aware that end users may be affected when redeploying a service.

Configure the Cluster DNS entrypoint to use IIS as the front-end

1. Set the Cluster DNS name (page 23) to the IIS server's hostname or FQDN.

2. Set the Cluster certificate (page 24) to use the same certificate and private key that's used by the IIS website.

3. Add the IIS certificate to the Trusted Root Certificate Authorities in Windows, using the "Manage Computer Certificates" MMC snap-in. This enables IIS to trust the cluster.

After configuring the IIS Reverse Proxy

When using IIS Single Sign-On authentication

Then, refer to the *MSS Administrator Guide* to configure MSS for Single Sign-on for IIS authentication (page .

When using SiteMinder authentication

Then, refer to the *MSS Administrator Guide* to configure MSS for SiteMinder authentication (page .

How Management and Security Server works

This diagram depicts the flow of secure interactions between a client and the host in a typical host session, using Management and Security Server. Note the option to use the Security Proxy Server and other Add-On products.

Host Access Management and Security Server
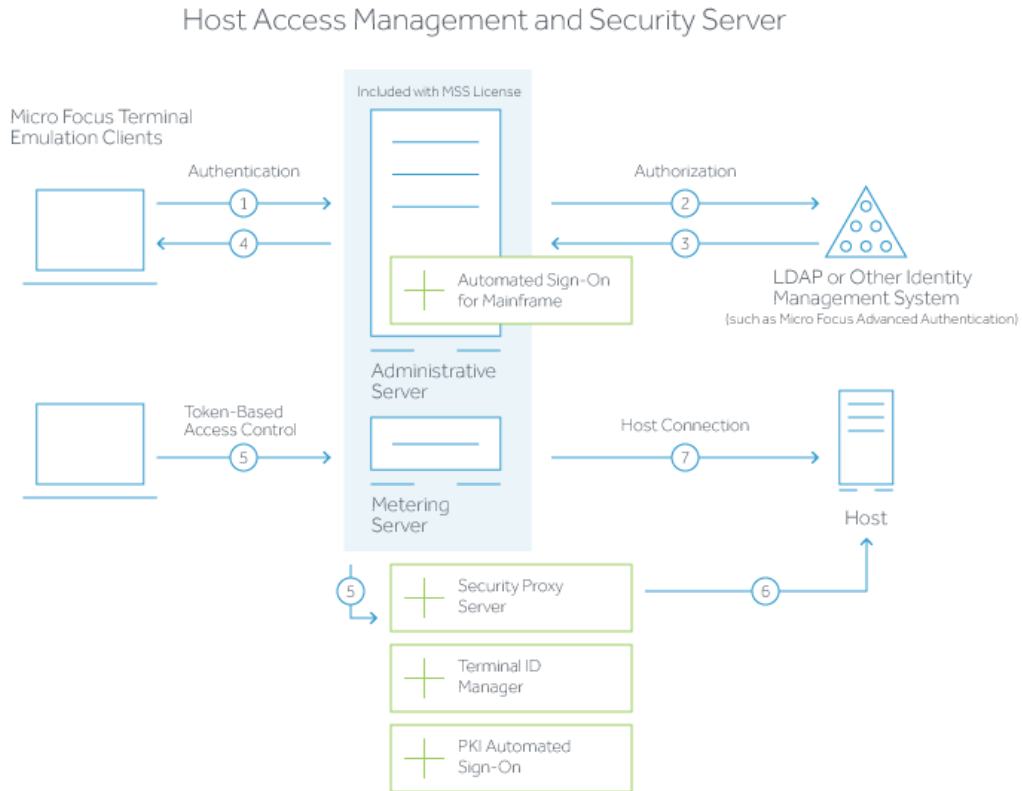
1. User connects to the Administrative Server.

2. User authenticates to a directory server (LDAP/Active Directory) or other identity management system (optional).

3. The directory server provides user and group identity (optional).

4. The Administrative Server sends an emulation session to the authorized client.

5. When the Security Proxy Server is configured for use by a session, the emulation client makes a TLS connection to Security Proxy and sends it a signed token.

6. The Security Proxy Server validates the session token and establishes a connection to the specified host:port.

7. When no Security Proxy is present or a session is not configured to use it, the authorized user connects directly to the host.

Checklist for Planning

As you plan your deployment, consider the workflow required to install and begin using MSS. It may be helpful to check each step as you proceed.

:material-checkbox-blank-outline: Choose a deployment type: virtual appliance or Linux installer.

:material-checkbox-blank-outline: Determine how many nodes you need.

:material-checkbox-blank-outline: Follow the installation steps for your preferred deployment type.

:material-checkbox-blank-outline: Configure your deployment.

:material-checkbox-blank-outline: Continue to use MSS as you did previously.