

Host Access for the Cloud Deployment Guide

Table of contents

Host Access for the Cloud Deployment Guide	4
In this guide:	4
Legal Notice	4
Plan	5
Planning for Deployment	5
Installation Types	5
Standard Deployment	6
Fault Tolerance and Scaling	8
Next Steps	11
Install or Update	12
Install or Update	12
Virtual Appliance	12
Linux Installer	19
Uninstalling	23
Next Steps After Installing	24
Configure	25
Configure Your Deployment	25
Configure Your Cluster	25
Clustering	26
Next Steps After Configuring Your Cluster	28
Apply	29
Apply	29
Create Host Sessions	29
Providing Access to Host Sessions	30
Managing User Preferences	30
Customizing Host Sessions	31
Logging	33
Authentication and Authorization	36
Develop	39
Additional Features	45

Host Access for the Cloud Deployment Guide

Host Access for the Cloud now 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, then configuring Host Access for the Cloud.

In this guide:

[Plan for deployment](#)

[Install or update](#)

[Configure your deployment](#)

[Apply](#)

Legal Notice

© 2024 Open Text

The only warranties for products and services of Open Text and its affiliates and licensors (“Open Text”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Plan

Planning for Deployment

The following steps will help you plan your deployment.

- Determine which installation type meets your needs.
- Familiarize yourself with what a standard deployment consists of and how many nodes you'll need.
- Become familiar with fault tolerance and scaling.
- Learn about the [Cluster DNS name](#) and [Cluster certificate](#).

Installation Types

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


Installation Type	Description	System Requirements
Virtual Appliance	The virtual appliance is a pre-configured virtual machine that contains everything you need to run the system. You deploy the appliance into your virtualization environment using an OVF file, creating as many appliances as needed to meet the demands of your environment. Software updates are handled via an easy to use integrated tool that doesn't require reinstalling the appliance itself.	View System Requirements

Installation Type	Description	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 that's a virtual or physical server. It should be used when your organization requires a specific Linux distribution or more control of the server environment.	View System Requirements

Standard Deployment

We recommend the following default deployment as a starting point:

- An external load balancer (optional)
- Three cluster nodes



This deployment provides

- Load Balancing - Where user requests are distributed across nodes for performance and availability.
- High availability - The ability for one node to go offline without significantly impacting users.
- Scalability - Where additional capacity may be added as needed.

Requirements - What you provide

- Servers or virtual machines that meet the system requirements: [Appliance](#) or [Linux installer](#)
- An odd number of nodes is always required.
- A [load balancer](#) with a [DNS hostname](#) for the cluster.
- A [certificate key pair](#) for securing access to the cluster.

Additional information

Load balancer

An external load balancer is optional but recommended. The specifics on 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

You provide a DNS hostname which is used to access 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

You provide a certificate key pair to use to secure all communication to the cluster. A self-signed certificate is generated which you can use to access the cluster initially, but for 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 you are not using an external load balancer, the certificate should contain a DNS SAN entry for each node in the cluster.
- If accessed directly, the certificate is served up by each node in the cluster. If not already present, an additional DNS SAN entry for each 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 and fully qualified domain name of each node

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

If you chose to use a static IP address during installation, will you need the network mask.

Default gateway

DNS servers

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**. This helps 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.

Number of nodes in cluster	Number of nodes required for a quorum
3	2
5	3

Number of nodes in cluster	Number of nodes required for a quorum
n	$(n / 2) + 1$

Failure handling

The health of all services in the system are monitored.

- If a service is found to be unhealthy, the system automatically attempts to self-heal, generally by restarting the process.
- It may take five minutes or more for a service to be recognized as unhealthy and restarted, potentially on a different node when appropriate.
- Service interruptions may occur depending on the type of failure.
- The Cluster Management dashboard provides a view of events regarding detected failures.

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

- The cluster generally moves 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.
- There are instructions for gracefully shutting down or rebooting a node. These should be used any time a node is shut down or rebooted.

Scaling

A standard deployment of three nodes supports:

- 6000 concurrent host sessions
- The failure of 1 node

If capacity is needed beyond what the standard deployment provides, there are two options: vertical scaling and horizontal scaling. Vertical scaling is recommended in most cases as it does not involve managing additional nodes.

Vertical scaling

To scale vertically, you add more memory and CPU cores to each of your existing nodes. This adds capacity to handle additional requests on the existing nodes, i.e. more concurrent host connections, but it does not increase the number of nodes that can fail without a system disruption.

Host Session Capacity

For each additional 6000 host sessions needed, the following should be added to *each* node in the cluster:

- 2 additional CPU Cores
- 4 GB RAM

Max concurrent host sessions	Required CPU Cores	Required Memory
6000	Base requirement - 8 Cores	Base requirement - 16 GB
12,000	10 Cores	20 GB
18,000	12 Cores	24 GB

Horizontal scaling

To scale horizontally you add more nodes to your cluster. This adds both additional capacity as well as increased resilience for nodes to fail, but involves managing additional nodes and the complexity that comes with that.

Important

You must always have an odd number of nodes in your cluster

Host Session Capacity

When scaling horizontally, *each* node added to the cluster adds capacity for approximately an additional 2000 host sessions.

Max concurrent host sessions	Required Number of Nodes*
6000	Base requirement - 3 nodes
10,000	5 nodes
14,000	7 nodes

* Assuming base system requirements for CPU and Memory

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, these resources are factored into the system requirements.

Next Steps

Once you've developed a plan for your deployment, the next step is to move on to [installing](#).

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	HACloud
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 system requirements, and then perform the following steps:

1. Download and unzip the Appliance ZIP file.

The ZIP file contains the files necessary 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)

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).
- 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 will take 5-15 minutes. Once a login prompt is displayed, 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.

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 Micro Focus download 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(s) for the product you want to install.

Note

At least one activation file for a product, such as HACloud, RWeb, or Management and Security Server (for use with Desktop products), 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 its status in the appliance console cluster view.

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 reduces the overhead of managing security patches and bug fixes, register the appliance:

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](#).
 - 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.

Manage appliance software updates

Software updates are delivered through two mechanisms in the appliance:

- Online Update: Delivers regular security patches to the OS and installed products.
- Upgrade: Delivers more significant upgrades to the OS and installed products.

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.

Some product updates involve an OS upgrade. This ensures the availability of OS security updates. In this case, both the Online Update and Upgrade tools are used when updating.

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 update your appliance, first ensure all nodes in the deployment are in a `Ready` state by clicking `Cluster` in the Appliance Administration console. Then perform these steps.

1. 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 `Upgrade` button shows a badge indicating an `Upgrade` is available, click `Upgrade`.

Click `Start`. Then review the license.

Register using the same email address and activation key as used during initial registration.

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

- e. Click `Reboot`. After 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.

- f. 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.
 - g. Move on to the next node in the cluster.
2. Once all of the nodes in the cluster have been updated, enable SSH on any node in the cluster.
 - a. In the Appliance Administration console, click System Services.
 - b. Click SSH.
 - c. Click Action > Start.
 3. On a node in the cluster that has SSH enabled, ssh into the appliance shell (command line interface) as the root user and run: `cspectl cluster scale`
 4. In the Appliance Administration console, click Cluster.

Wait for the cluster to return to a healthy state with all nodes showing `Ready`. This process can take up to 15 minutes.

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 all about SMT](#)
- [Installing the SMT server](#)
- [Creating a certificate](#)

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=smt.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

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)

Note

Be sure that a large part of the disk space is available to `/var/opt/opentext/csp`.

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(Product specific):

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

The following third-party Linux packages are required for installation: `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`

These packages will be installed automatically for you during product installation. However:

- In the case of Red Hat, you must add the EPEL (Extra Packages for Enterprise Linux) repository in order to gain access to these packages. Specifically, the `epel-release` repository is required.
- For OpenSUSE, add the SUSE Linux Enterprise (sle) repository to retrieve the `supportutils` package.

Some of these packages are platform-specific and are not required on all platforms.

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, and 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. 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. Host Access for the Cloud uses a PGP key to verify 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](#) to download the key separately and verify the file.
6. When the install completes, a verification tool is automatically executed.
7. If verification succeeds, then the services automatically start and you can move on to the next steps. If verification fails, see [Troubleshooting](#)

Troubleshooting

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 start the system automatically after server restarts.

If issues remain please contact Micro Focus 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. Then perform these 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 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](#).

- c. After the CSP service starts, wait until the updated node shows a status of `Ready`. The cluster status can be checked by running the following command with elevated privileges: `cspctl status`

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 in the cluster.

4. After all of the nodes in the cluster have been updated, on any node in the cluster with elevated privileges, run: `cspctl cluster scale`
5. Wait for the cluster to return to a healthy state with all nodes showing `Ready` using: `cspctl status`. This process can take up to 15 minutes.

Uninstalling


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

Note


Before uninstalling, always remove the node from the cluster:

Open the MSS Administrative Console (<https://hostname:9443>)

Click Cluster Management > Nodes

Next to the node you want to remove, click  , then Delete

Virtual Appliance method

1. Open the Appliance Administration Console (<https://hostname:9443>) > Products.
2. Next to the product you want to uninstall, click  Uninstall.

This process takes a while to complete.

Linux

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

The uninstall process takes a while to complete.

Next Steps After Installing

Once you've installed your product, the next is to move on to [configuring your deployment](#).

Configure

Configure Your Deployment

After installing, you have a cluster of one, a single node. The next steps are to configure key cluster settings 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](#)

[Clustering](#)

Configure Your Cluster

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

Access the Admin Console

- Access the MSS Admin Console here: `https://hostname/adminconsole`.
- The Admin Console's default password is `admin`.
- Once signed in, you can navigate to various views using the drop-down menu in the upper left.

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 pull-down menu, select Cluster Management.
4. Click Settings.
5. Set the Cluster DNS Name and click Apply.
6. You should now 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, select 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](#) for guidance.

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 health 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.

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 will take 5-15 minutes to complete.**

To remove a node:

To remove a node from a cluster, please refer to the Cluster Management Help.

Next Steps After Configuring Your Cluster

You now have a cluster ready for use, the next step is to proceed to [apply your product configuration](#).

Apply

Apply

After installing, deployment, and configuration is complete, you can now create and customize host sessions, extend those sessions using the tools provided, and set user preferences.

Create Host Sessions

Host Access for the Cloud supports IBM 3270, 5250 and VT hosts as well as UTS, T27 and ALC host types.

Your users gain access to the host through sessions that you create and configure. Sessions are created by an administrator in the MSS Administrative Console. When you launch a session from the Administrative Console, the web client Connection panel opens in a separate browser window. You configure connection options from this panel. Options vary depending on your host type.

To create a session:

1. Create sessions in the MSS Administrative Console. See [Add a Session](#) in the MSS documentation.
2. In the HACloud web client, on the **Create New Session** dialog box, select the type of host you are connecting to from the drop down list.
3. On the Connection panel, identify the name of the host to which you want to connect. You can use a full host name or its IP address.
4. Type the number of the port you want to use.
5. Complete the information needed for the host connection.
6. Save your connection settings.
7. After you complete configuring and testing your session, click Exit to return to the MSS Administrative Console.
8. Using the [Assign Access](#) view in the MSS Administrative Console, assign the session you created to your users.
9. Once sessions have been assigned to your users, you can let them know [how to access their sessions](#).

Providing Access to Host Sessions

Your end users access sessions either through a session server or through the Assigned Sessions list portal. With both options, once authenticated, your users are presented with a list of sessions that they can access and successfully launch.



Using a load balancer for high availability and scalability is highly recommended. See [Planning for Deployment](#) for more information.

Session Servers

Users most commonly access their sessions by going to the session servers, typically through a load balancer.

End user access on a session server is available at `https://<clusterdns>/webclient`

Assigned Sessions List

Using the Assigned Sessions list, users can launch all of their sessions from a consolidated HTML-based portal. After a user authenticates, they see their list of assigned sessions.

The Assigned Sessions list is available at `https://<clusterdns>/sessions/`

Managing User Preferences

As an administrator you can choose what options users can configure for their sessions. These options are set on a per session basis and all users who have access to a particular session can configure their own session instance.

1. From the left navigation panel, choose **User Preference Rules**.
2. Select which options you want to allow your users to configure.
3. Click Save.

Each user's configurations are specific to their instance of the session and will not conflict with those of other users.

There is a **Restore Defaults** option available on the various settings and display panels. As an administrator, this option restores the web client back to its default settings. For end users this option will restore the values set by the administrator when the session was created.

Warning

When the authentication method is set to None, be aware that all users share the same settings. During session configuration, it is best to not allow users to modify their session settings (User Preference Rules), because they can overwrite each other's choices. To work around this constraint, it is possible to [provide user identities in other ways](#).

Customizing Host Sessions

Customizing Host Sessions

You can use these features to customize sessions for your end users:

- **Plus** - Enable custom controls to provide a more efficient work flow and a more modern and friendly interface. See [Use Plus](#) to customize screens.

Using this option, you can add tool tips to fields, replace old-style numbered lists with more modern drop-down lists, add buttons to the host interface and program them to start macros or perform other actions, and replace manual date entry with a graphical calendar date-picker.

- **Server-side Events** - Supply procedural Java code that extends and improves the presentation of host data.

Using server side events, you can define specific events and suspend the host application, replacing or interrupting it with code that you have supplied to the session, as well as extend error handling options. For example, you can add an event that recognizes when an error occurs and then implements the code to intercept the error, take control, and correct the error. See [Server-side events](#).

- **Advanced** - Only use as directed by Micro Focus Technical Support.

These options are configured on the Customization panel.

1. Click Settings on the toolbar to open the left navigation panel.
2. Click Customization.

Use Plus to customize screens


Note

The Plus feature requires archive files (`.rdar`) produced by Micro Focus Screen Designer version 9.5 or higher. The Screen Designer is available in Micro Focus Rumba Desktop 9.5. Reflection Desktop 16.1 includes a limited version of the Screen Designer. To get access to more controls and full use of Plus and the Screen Designer, you can purchase and install the Micro Focus Reflection Desktop Plus add-on.

1. On the **Customization** panel in the web client, click **Enable Plus**.
2. Select the Plus archive file you want to use from the drop-down list or upload a file from a different location. Plus archive files are identified by a `.rdar` file extension.

Archive files are the output of a Screen Designer project and are used to provide the custom control criteria.

If you are updating the Plus archive (`.rdar`) file associated with your Plus enabled session, you must first delete the folder containing the old `.rdar` file from the session server. After you delete the folder, you can open your Plus enabled session and the new `rdar` file will be downloaded to the session server.

3. Verify that the number of milliseconds for the host settle delay time is accurate. This is the time that the server waits for a synchronous connection before deciding that the host has finished sending data.
4. When you return to your session, Plus is available. Click  on the toolbar to turn off the custom controls.

When you enable Plus for a session, all end users of that session see the Plus icon on the toolbar and any controls made available through the Screen Designer customization file.

More information

[Customizing Host Sessions](#)

Use Server-side Events

Using server-side events, you can supply procedural Java code that can extend and improve the presentation of host data.

The Customization panel tells the web client where to find the event after you configure it. See [Using the Java SDK](#) for instructions on using the SDK and the samples available to you.

1. Open the **Customization** panel in the web client.
2. Under **Server Side Events**, type the fully-qualified class name to the event.
3. Launch the session and test the event.

Access [API documentation and event samples](#).

More information

[Customizing Host Sessions](#)


[Using the Java SDK](#)

Logging

Viewing and downloading log files is accomplished on the Cluster Management console.

Locating log files

To view or download log files:

1. Log in to the MSS Admin Console at `https://hostname/adminconsole`.
2. From the drop-down menu, click Cluster Management.
3. Click Services, and find the `hacloud-session-server` service.
4. Click that service to view the application instances (pods) across the cluster.
5. Click  and then click either
 - View Recent Logs - to view the last 500 entries logged in a browser window.
 - Download Logs - to download the entirety of all the log files
 - Redeploy - to redeploy the service instance. This will affect any users connected to the service.

Setting logging levels

There are various types of logging levels you can use to produce different types of information. You configure logging levels by editing the service's properties in the [Cluster Management console](#).

Use the following format to set logging levels:

Key	Value
<code>logging.level.logger</code>	<code>log level</code>

Where `logger` is the name of the logger to adjust and `log level` is one of the following:

- **Trace** - designates finer-grained informational events than Debug
- **Debug** - designates fine-grained informational events that are most useful to debug an application.
- **Info** - designates informational messages that highlight the progress of the application at coarse-grained level.
- **Warn** - designates potentially harmful situations.
- **Error** - designates error events that might still allow the application to continue running.
- **Fatal** - designates very severe error events that will presumably lead the application to terminate.

Enabling Web client-to-session server logging

While the browser provides a basic mechanism for logging to its JavaScript console, the Web Client extends this capability and, with some configuration, you can log events to the service for viewing by an administrator.


By default, nothing is logged to the session server. **You must set the log level**, following the instructions below, in order to enable this feature.

The available log levels are: `debug`, `info`, `warn`, `error`, or `off`. The default log level is `off`.

Adjusting the logging level for all Web client users

To adjust the logging level for all Web clients, add the following property to the service

Key	Value
logging.level.com.microfocus.zfe.webclient.core.handler.ClientLoggingHandler-webclient	log_level

 **note**

Use caution when increasing the logging level for all Web Client users in a production environment due to a potential increase in network traffic.

Adjusting the logging level for an individual user

There are two options for adjusting the logging level for individual users:


- To temporarily adjust the logging level for a particular user's Web client instance without requiring a session server restart, instruct the user to add the following URL parameter when loading the Web client in their browser:

- https://<clusterdns>/webclient?log=<log_level> -

- To adjust the logging level for an individual user without requiring them to make changes, follow these directions, [Adjusting Advanced Product Settings](#), using these values:

Key	Value
logging.level.com.microfocus.zfe.webclient.core.handler.ClientLoggingHandler-webclient- <i>username</i>	log_level

and where `username` is the user name of the person whose logging levels you are adjusting.

 **note**

Logging based on a username requires an authentication mode that involves usernames.

Authentication and Authorization

Authentication and Authorization

In HACloud, authentication and authorization are provided by the Host Access Management and Security Server (MSS) and are configured using the Administrative Console.

Authentication validates a user's identity based on some credentials, such as a username/password combination or a client certificate. Authorization is then used to determine which sessions each user can access.

HACloud supports the following authentication methods: None, LDAP, Single Sign-on through IIS, Windows Authentication via Kerberos, X.509 Client Certificates, SiteMinder, OpenID Connect, and SAML.

For general information on choosing and configuring authentication and authorization methods, see [Authentication and Authorization](#) in the MSS documentation.

Note

Some authentication methods require HACloud specific configuration. See the topics under this heading in the table of contents for more information.

Enabling OAuth

OAuth must be enabled when configuring these authentication methods:

Windows Authentication - Kerberos

X.509

OpenID Connect

Steps to enable and configure OAuth

1. [Enable and configure OAuth in MSS](#)
2. [Configure OAuth on Session Server](#)

Configure OAuth on Session Server

1. In the Administrative Console, open Cluster Management from the top left drop-down.
2. On the Services page, locate `hacloud-session-server`. Click `...` and then Edit Properties.
3. Find the entry with `SPRING_PROFILES_ACTIVE` as its Key value.

4. In the corresponding Values field add `,oauth`. Click OK.
5. Click `...` on the `hacloud-session-server` service and click Redeploy All.
6. Select your preferred authentication method:
 - [Windows Authentication - Kerberos](#)
 - [X.509 Configuration](#)
 - [OpenID Connect](#)

Single Sign-on through IIS

See [Single Sign-on through IIS](#) in the MSS Administrative Console documentation if you need further information.

This option uses Microsoft IIS web server.

OpenID Connect

OpenID Connect (OIDC) is an open standard security protocol that delegates user authentication through a third-party identity provider.

In Host Access for the Cloud, end users can interact with their host sessions on the session server without being prompted for credentials, as long as they are authenticated by the third-party identity provider.

Steps to configure OIDC

1. [OAuth](#) must first be enabled before configuring OIDC.
2. [Configure OIDC in the MSS Administrative Console](#).

Windows Authentication - Kerberos

Kerberos is an authentication protocol that uses cryptographic tickets to avoid transmitting plain text passwords. Clients obtain ticket-granting tickets from the Kerberos Key Distribution Center (KDC) and present those tickets as their network credentials to gain access to services.

In Host Access for the Cloud, Kerberos allows end users to access their host sessions on the session server without being prompted for credentials.

 **Note**

Kerberos authentication to AS/400 hosts is also supported, however that functionality is not yet integrated with Kerberos for authenticating end users accessing the session server.

Steps to configure Kerberos

1. OAuth must first be enabled before configuring Kerberos.
2. Configure Kerberos in the MSS Administrative Console.

Configure your browser for Kerberos

In order to sign in using Kerberos, your browser must be configured correctly for Windows Authentication via Kerberos and your machine must be a member of the proper domain (Kerberos realm). Please consult the help for your specific browser for instructions on how to enable Kerberos.

Launch sessions

HACloud sessions need no additional configuration to launch and authenticate using Kerberos, as long as your browser has been configured correctly for Windows Authentication / Kerberos. Just navigate to <https://cluster-dns.mydomain.com> and you'll be automatically logged into the HACloud session server.

SAML

SAML (Security Assertion Markup Language) is an XML-based open standard format that exchanges authentication and authorization data between an identity provider, (the server that issues SAML assertions and performs authentication), and a service provider, (the web server from which you access information or services). MSS acts as the service provider.

See [SAML Configuration Steps](#) in the MSS Administrative Console documentation if you need further information.

 **Note**

The SameSite flag must be adjusted when using SAML behind a load balancer. See [Setting the SameSite Attribute](#).

Troubleshooting the configuration

If you have problems authenticating or see session timeout errors, see [Troubleshooting SAML Setup](#) in the MSS Administrator Guide.

X.509 Authentication

X.509 client authentication allows clients to authenticate to servers with certificates rather than with a user name and password by leveraging the X.509 public key infrastructure (PKI) standard.

MSS has additional information on [X.509 configuration](#).

Enabling X.509 client authentication

- When the user accesses the web client using TLS the browser sends a certificate to the session server identifying the end user and completing the TLS handshake.
- The session server refers to its truststore to check the client's certificate and verify its trust.
- Once the TLS negotiation is complete (the session server trusts the end user), the session server sends the end user's public certificate to MSS for further validation.
- MSS also verifies that it trusts the end users certificate using its trust store.
- When MSS finishes the validation, the end user will have successfully authenticated.

The client's full certificate chain needs to be present in the session server and MSS truststores or alternatively signed by a Certificate Authority that is present in the truststores.

The browser determines what client certificate to send using a browser or smart card specific configuration.

Basic steps

1. First, [enable OAuth](#).
2. Configure X.509 in the [MSS Administrative Console](#).
3. Ensure the user certificate is properly installed on the client system.

Develop

Developing with HACloud

Host Access for the Cloud has a collection of APIs and libraries that help you develop efficient client/server and Web applications that integrate host data into various development environments.

You can also extend the web client without affecting the installed files. This ability provides you with a wide range of options to tailor the web client to your own needs.

- [Extending the Web Client](#) you can enhance and broaden the scope of the web client using custom code, such as CSS or JavaScript.
- [Using the Java SDK](#) you can use the provided Java API to enhance the presentation of host data using server side events.
- [Enabling the JavaScript API](#) you can embed the web client in your own web site.
- [Using the Connector for Windows](#) you can interact with host sessions in your .NET application or within Visual Basic for Applications using the API and samples provided.

[HACloud API Documentation](#)


Extending the Web Client

You can update, modify, and customize the presentation of the web client by using your own HTML, CSS, or JavaScript from within the browser.

You can take advantage of extensions to make visual changes to the web client and to customize the application. The web client hosts your custom HTML or CSS code, making it easy to modify and support.

Adding an extension

Before proceeding keep in mind that although Host Access for the Cloud provides the ability to plan and use custom code, the code itself must be supported by the team that produced it.

You enable extensions in the Host Access For The Cloud console. There are detailed instructions available from the Extension page help icon .

warning

During a product upgrade extensions are disabled. This means that, after an upgrade, you must verify that the product is working as expected without extensions, and then re-enable the extensions using the steps to add custom code.

When you add extensions to the web client, the modifications are visible to all your users and apply to all sessions.

To enable extensions

1. From the drop-down menu in the Administrative Console, select Host Access for the Cloud.
2. Click Extensions
3. Ensure extensions are uploaded to enable the Extensions toggle.
4. Toggle the Extensions switch to activate extensions.
5. Redeploy the service.
6. After extensions are enabled, from the same drop-down menu, select Host Access for the Cloud, and then Extensions. From here you can upload, download, or remove your extensions.

Making extensions available without client authentication

Files within the `/client` directory are protected using the level of authentication you selected in MSS.

To share files without requiring authentication:

1. Create a zip file that contains a directory named `public`. This directory will contain your extension code.
2. In this same zip file, add any other client or server extensions under the relevant directories (`/client` or `/server`).
3. Using the Extension page from the Host Access for the Cloud drop-down menu, upload the zip file. Call the URL `/public/*` to access your files.

Extension sample

In this example, once extensions are enabled, you can add some custom CSS and JavaScript code to change the menu label font color and print text to the JavaScript console.

You will create three files; `custom.css`, `custom.js`, and `index.html`


Step 1

Create `index.html`. Place your extension files in the file to create an entry point:

```
<!-- Define the link to the external style sheet -->
<link href="client/custom.css" rel="stylesheet">
<!-- Define the external JavaScript file -->
<script src="client/custom.js"></script>
```

Step 2

Change the default black menu labels to orange:



Create custom.css to change the color to orange:

```
/* Change link text to Orange */  
a span {  
  color: #ff5d28;  
}
```

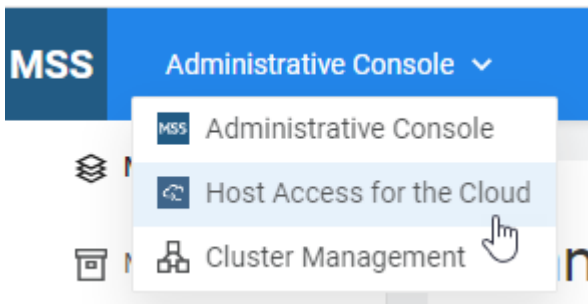
Step 3

Create custom.js to send text to the JavaScript console:

```
//Print message to the JavaScript console  
console.log('Hello World!');
```


Step 4

Create a zip file containing client and server folders, each containing the relevant extension file. Using the Extension page from the Host Access for the Cloud drop-down menu, upload the zip file.



Step 5

Deploy the files. The the results should look like this:



And the "Hello World" text is visible in the JavaScript console:




More Information:

Using the Java SDK

Working with [server side events](#) and the Host Access for the Cloud SDK you can supply procedural Java code that can extend and improve the presentation of host data. To help you create server side events, Host Access for the Cloud has an SDK to provide you with a starting point.

The API documentation, including samples, is available in a ZIP file (`hacloud-extras-<version>.zip`) from the same download location as the Host Access for the Cloud product files. They are also always available online: [Javadocs](#).

1. [Enable extensions](#).

You enable extensions in the Host Access For The Cloud console. There are detailed instructions available from the Extension page help icon .

2. Write the Java code necessary to accomplish the task and compile the code into a Java class within a JAR (Java Archive) file.
3. Copy the JAR file to a folder named `/server` and create a ZIP file containing the `/server` folder at the root and any other extensions you want to use.
4. From the Admin Console drop-down menu, open Host Access for the Cloud.
5. Using the Extension page, upload the ZIP file.
6. Redeploy the service.
7. As you configure the session in the web client, from Settings, click Customization, and type the fully qualified class name to the event.
8. Launch the session and test the event.

Using the JavaScript API

Note

The SameSite flag must be modified when using the JavaScript API. See [Setting the SameSite Attribute](#).

Using JavaScript in a browser you can embed the web client in a web page. Your end users, by accessing a common web page, can interact with the web client and connect to the host application providing the ability to:

- Programatically interact with host sessions.

- Run it “headlessly”, meaning you can access all its functionality without having a visible interface embedded in the web page.

The API documentation, including tutorials and the JavaScript library, is available in a ZIP file (`hacloud-extras-<version>.zip`) from the same download location as the Host Access for the Cloud product files. They are also always available online: [HACloud JavaScript API](#).

Using the Connector for Windows

Host Access for the Cloud Connector for Windows allows you to interact with host sessions in your .NET application or within Visual Basic for Applications.

The installers, API documentation, and examples are available in a ZIP file (`hacloud-extras-<version>.zip`) from the same download location as the Host Access for the Cloud product files. The documentation is also always available online: [HACloud Connector for Windows](#).

Here are a few things to keep in mind as you prepare to install:

- Two install platforms are available: a 32-bit version and a 64-bit version. Depending on which one you install, the default base install path will be `C:\Program Files (x86)\Microsoft Focus\HACloud\Connector for Windows` or `C:\Program Files\Microsoft Focus\HACloud\Connector for Windows`
- The installation platform you choose also determines the solution platform in which you can develop. For example: If you have installed the 32-bit version of Microsoft Office® and want to use Visual Basic for Applications with the connector, then you must install the 32-bit version of the Host Access for the Cloud Connector for Windows.
- .NET 4.6.2 or higher is required.
- The Connector for Windows supports these authentication methods: LDAP, Kerberos, and None. Authentication is configured in the MSS Administrative Console.
- To use the Windows connector with Kerberos:

You must first be authenticated to the Windows Domain.

Because the MSS server must be trusted for integrated authentication, the Cluster DNS URL should be added to the Intranet Zone site list in the Internet Options for the Windows Operating System. (See Microsoft documentation for instructions).

When connecting a client application to a session server in the cluster, the DNS endpoint must include the HACloud webclient context. For example, `https://clusterdns/webclient`.

Examples and connector documentation

Documentation is available to reference from your IDE. There are also samples to help you take advantage of the connector.

1. Navigate to the install directory. In a default install, either `C:\Program Files (x86)\Micro Focus\HACloud\Connector for Windows` Or `C:\Program Files\Micro Focus\HACloud\Connector for Windows` depending on your platform.
2. In the Connector for Windows directory you will find:
 - `MicroFocus.ZFE.Connector.dll` - a .NET Framework assembly to reference in your C# or .NET project.
 - `MicroFocus.ZFE.Connector.tlb` - a Type Library to use in your COM or Visual Basic for Applications project.
 - `\help` - this directory contains information which will aid in using the connector.
 - `\samples` - this directory contains the code samples that provide a starting point for developing your own applications.

Using the connector with Microsoft Visual Studio

If you are using Microsoft Visual Studio to develop applications, keep these things in mind:

- When using Microsoft Visual Studio with Connector for Windows, make sure your solution platform is set to either x86 or x64, depending on your installation. Because of the native components used within the Connector for Windows SDK, the Any CPU platform is not supported. Use the Configuration Manager for your Visual Studio Solution to create a platform for x86 or x64.
- When adding a reference to the Connector for Windows library, Visual Studio may set the Copy Local reference property to True. This should be set to False so that the library and its dependencies are executed from the SDK install directory.

Additional Features

Using Metering

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

Metering is automatically set globally for all emulation sessions created by the session server.

 **Note**

In the event that all licenses are in use and you attempt to make a connection, the session will be disconnected. To determine whether the host has disconnected or the metering service has stopped the connection, see the [log file](#).

Set Up the Terminal ID Manager

The Management and Security Server provides a Terminal ID Manager to 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.

[Setting up the Terminal ID Manager](#) requires that this feature is first enabled in MSS. The Terminal ID Manager Guide has complete instructions on how to configure the Terminal ID Manager.

After Terminal ID Manager has been configured in MSS, each HACloud session can then be configured to use the Terminal ID Manager in Connection Settings.

See the MSS [Terminal ID Manager Guide](#) for instructions on configuring Terminal ID Manager. Once configured, each HACloud session can then be configured to use Terminal ID Manager in the web client under Connection Settings.

 **Tip**

If MSS and Host Access for the Cloud are installed on the same machine, no additional configuration is needed.

Set Up Automated Single Sign-On for Mainframe

[Automated Sign-On for Mainframe - Administrator Guide](#) has additional information on configuring this option.

Automated Sign-On for Mainframe is an add-on product to Management and Security Server that enables an end user to authenticate to a terminal emulation client and be automatically logged on to a host application on the z/OS mainframe.

1. Install and configure the Automated Sign-On for Mainframe add-on for Management and Security Server. You can find complete instructions in the [Automated Sign-On for Mainframe - Administrator Guide](#).

2. After the Management and Security Server setup is complete, open the Administrative Console to add sessions and map users to those sessions. During that process, you can complete the additional configuration needed to implement automated sign-on.
3. A Host Access for the Cloud macro sends the user's mainframe username and pass ticket to the host application. The user is then automatically logged in. To help users create the macro there is a macro API containing an AutoSignon object and a sample macro in the Users Guide.

Set Up Kerberos for AS/400 Single Sign-on

Kerberos is an authentication protocol that uses cryptographic tickets to avoid transmitting plain text passwords. Client services obtain ticket-granting tickets from the Kerberos Key Distribution Center (KDC) and present those tickets as their network credentials to gain access to services.

Note

Kerberos authentication for end users accessing the session server is also supported, however that functionality is not yet integrated with Kerberos for AS/400 authentication. This feature allows for automated sign on from the session server to an AS/400 host. MSS must be configured with an authentication method that results in a user principle that is resolvable in the Kerberos Active Directory domain, for example LDAP, SAML, or Siteminder. A Windows Active Directory Server is required.

By using Kerberos, after an initial domain sign-on, users do not have to enter their credentials when accessing AS/400 sessions in Host Access for the Cloud.

An overview of enabling and using this feature can be found in the MSS Administrative Console > Host Access for the Cloud panel documentation.

Technical References

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.
- There are detailed migration steps, including a complete list of what data can be migrated to the new installation. Open MSS Administrative Console > Configure Settings > Migration.

Data that is NOT migrated

kerberos settings

metering report data

security proxy configuration

- passwords

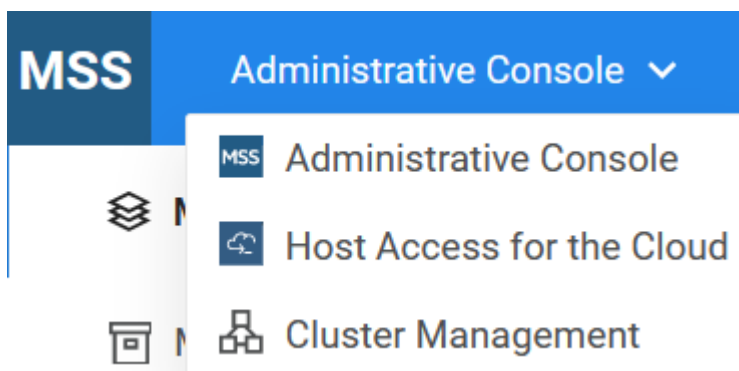
For example the MSS Admin password will remain the same before and after migration.

Advanced Product Settings

Information for HACloud is always available from the installed documentation or from [online resources](#).

You may occasionally have to change or add properties to your product services.


Properties are set in the Cluster Management console.



Follow these steps:

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

Click Services.

Click the service of interest, and click  Edit Properties.

Add or edit the key and value accordingly.

After all properties are adjusted, redeploy the service.

Important

Be careful when redeploying services as this will 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. See [instructions on how to use the Kubernetes dashboard](#).

Setting the SameSite Attribute

To help prevent cross-site request forgery attacks, the default SameSite attribute on the session server cookie has been updated from **None** (less restrictive) to **Lax** (more restrictive).


With the attribute set to Lax, the session server cookie will not be sent on cross-site requests as is often the case with the JavaScript SDK and in SAML authentication.

This change affects two areas of HACloud:

- The JavaScript SDK and
- SAML authentication behind a load balancer

In these cases you'll need to adjust the attribute value to **None**.



1. Open the Administrative Console, and launch Cluster Management.



2. Click Services and locate the `hacloud-session-server` service.
3. Open the Edit Properties option from the `⋮` menu.
4. Add a new property:
Name - `samesite.cookie.attribute`
Value - `none`
5. Click OK.
6. From the same `⋮` menu associated with the `hacloud-session-server` service, choose Redeploy All. Click Yes when prompted. Wait for the session servers to redeploy.

Setting the SameSite attribute for a multi-node cluster

Additional configuration is needed when using the JavaScript SDK with a multi-node (multiple session servers) cluster. The SameSite attribute must also be configured for the session server load balancer cookie.

1. Follow step 1 above.
2. Click Advanced from the left pane.
3. Enable the Kubernetes Dashboard. This may take several seconds.
4. Copy the Authentication Token using the copy icon .
5. Launch the Kubernetes Dashboard using the Kubernetes Dashboard URL.
6. Paste the Authentication Token you copied to the clipboard, and click Sign In.
7. Under Service in the left pane, click Services.
8. Locate (you may need to scroll) and open the `hacloud-session-server-lb` service.
9. On the `hacloud-session-server-lb` service page, click  to edit the service.
10. In the Edit Resource dialog box, locate the annotations section near the top of the file. Change the `samesite` annotation, `traefik.ingress.kubernetes.io/service.sticky.cookie.samesite: none`, from `lax` to `none`.

```
l annotations:
  meta.helm.sh/release-name: hacloud
  meta.helm.sh/release-namespace: hacloud
  traefik.ingress.kubernetes.io/service.sticky.cookie: 'true'
  traefik.ingress.kubernetes.io/service.sticky.cookie.httponly: 'true'
  traefik.ingress.kubernetes.io/service.sticky.cookie.name: session-server-lb
  traefik.ingress.kubernetes.io/service.sticky.cookie.samesite: lax
  traefik.ingress.kubernetes.io/service.sticky.cookie.secure: 'true'
```

11. Click Update. On the `hacloud-session-server-lb` service page, verify that the `traefik.ingress.kubernetes.io/service.sticky.cookie.samesite` annotation shows the new value.
12. Sign out and close the Kubernetes Dashboard and the Administrative Console.

Modifying the Size Limit on Upload File Transfer Operations

There is a 50MB file size limit on file transfer upload operations.

To modify the file size limit, you need to [adjust your service settings](#) and redeploy the session server.

Both keys need to be reset.

For example:

Key	Value
<code>spring.servlet.multipart.maxfilesize</code>	100MB
<code>spring.servlet.multipart.maxrequestsize</code>	100MB


Configuring User Names when Using Anonymous Access Control

Users need access to their macros, user configurations, and other personalized settings whether they are authenticated through Management and Security Server or not. These settings are collectively referred to as User Preferences.

When MSS is configured for authentication, using LDAP or SAML for example, a username is determined when a user logs in. The user's settings are centrally saved in MSS using that username for all future logins.

However when the MSS Authentication Method is set to None, also known as anonymous mode, no unique username is available for the system to identify that particular user when they return in the future. In this configuration, all users share the same settings. If one user changes a setting, that setting will be changed for all other users.

Because that may not always be the desired behavior, Host Access for the Cloud supports a number of ways that, as an administrator, you can configure a unique identifier for each user so their customized settings can be stored and retrieved during future logins.



note

These configuration modifications do not alter the security considerations of using Management and Security Server in anonymous mode.

Configuration options

There are four different configuration options you can choose from when configuring user name identifiers.

To implement any of these options, you need to [adjust your service settings](#) and redeploy the `hacloud-session-server` service.

- To use an HTTP request header value as the user name

Key	Value
<code>zfe.principal.name.provider</code>	<code>com.microfocus.zfe.webclient.security.mss.HeaderKeyAnonymousPrincip</code>
<code>zfe.principal.name.identifier</code>	<i>the-header-key-to-be-used</i>

- To use an HTTP request cookie value as the user name

Key	Value
<code>zfe.principal.name.provider</code>	<code>com.microfocus.zfe.webclient.security.mss.CookieKeyAnonymousPrincip</code>
<code>zfe.principal.name.identifier</code>	<i>the-cookie-key-to-be-used</i>

- To use an HTTP request URL parameter as the user name

Key	Value
<code>zfe.principal.name.provider</code>	<code>com.microfocus.zfe.webclient.security.mss.UrlParameterAnonymousPrin</code>
<code>zfe.principal.name.identifier</code>	<i>the-url-parameter-key-to-be-used</i>

- To use the client IP address as the user name

Key	Value
<code>zfe.principal.name.provider</code>	<code>com.microfocus.zfe.webclient.security.mss.RemoteAddrAnonymousPrinci</code>

Troubleshooting the configuration

If any of your users experience problems when connecting to a Host Access for the Cloud web application after you have made the configuration changes, check the following:

- Users experience a *503 Service Unavailable* message when connecting to a Host Access for the Cloud web application. First [check the log file](#) for `hacloud-session-server`, then:
- If the log file contains this message: *Unable to create AnonymousPrincipalNameProvider instance for class...*, then the `zfe.principal.name.provider` property is probably mis-typed. Check the spelling and letter case to remedy this issue.
- If the log file contains this message: *zfe.principal.name.identifier is not defined*, then the property is missing. Ensure the property is defined to remedy this issue.
- Users are unable to properly authenticate.

Users should receive an error message indicating the initial HTTP request to the Host Access for the Cloud web application did not contain the required informat

Configuring Cross-Origin Resource Sharing (CORS)

As a security measure, modern web browsers restrict the types of interactions that are permitted between distinct web sites. This can cause problems when attempting cross-site integration, for example when embedding the HACloud web client into another website, such as a portal. CORS is a standard mechanism that you can use to specify that the browser permit access from one site to another site.

You can configure the HACloud service to include the required CORS HTTP header when it responds to the web requests by [adjusting your service settings](#) and redeploying the `hacloud-session-server` service.

1. In the Edit Properties field, add:

Key	Value
<code>CORS_ALLOWED_ORIGINS</code>	<code>https://integration-server1.com</code>

2. Redeploy the `hacloud-session-server` service .

You can set this value to a comma-delimited list of allowed origins or use `*` to allow access from all origins (Allowing this kind of open access may be a security risk). If you use the wild card option(`*`), be aware that web browsers impose additional restrictions, such as limited Cookie access. For more information, see [Cross-Origin Resource Sharing \(CORS\) - HTTP/MDN](#).


Enabling FIPS Level Security

The Federal Information Processing Standards (FIPS) 140-2 validated cryptographic modules are used by the US federal government as a security regulation standard.

Host Access for the Cloud supports this standard and you can easily enable FIPS mode by [adding a property](#) in the Admin Console's Cluster Management interface to the `hacloud-session-server` service using these values:

```
key = com.attachmate.integration.container.FIPS.enabled
```


```
value = true
```

Redeploy the service. Click  associated with the `hacloud-session-server` service, and click Redeploy All.

You can verify that FIPS is enabled by looking at the log file.

Open the Admin Console

Select Host Access for the Cloud


Click  associated with the `hacloud-session-server` service and choose Download Logs.

Verify that FIPS mode is set to true; `FIPS mode: true`.

Using Single Session Mode

You can use single session mode and provide URLs to particular sessions that are launched using the name parameter, (for example a direct link on a company portal page). To enable the launch of a single session use the query parameter `singleSession`. You can use this parameter on its own to just launch the web client in single session mode, for example, `https://<clusterdns>/webclient?singleSession`, or it can be used in conjunction with a named session parameter to launch a particular named session in single session mode: `https://<sessionserver>/webclient?singleSession&name=HumanResources`. The order of the parameters does not matter.

When your users access a single session, they cannot switch between open sessions and cannot open new sessions. A new session will not be launched if the specified session already exists when the user opens the link.



If you want all sessions on the session server to run in single session mode:

1. Follow the steps for [setting properties](#) in the Admin Console.
2. Use the following values to edit the property for the `hacloud-session-server` service:

Key = `webclient.singleSession`

Value = `true`

3. Redeploy the `hacloud-session-server` service.

Known Issues

These issues have been identified in previous releases and are known issues.

Recommended browsers

It is highly recommended that you use Google Chrome or Mozilla Firefox.

Key mapping issues with different browsers

Certain keys on a numeric keypad and some browser-specific keys cannot be mapped. For example, in Chrome, `Ctrl+n` and `Ctrl+w` cannot be mapped.

Host specific issues

The following are issues that are specific to different host types.

Displaying the Euro character

If the EURO character does not display correctly on the terminal screen, talk to your system administrator to make sure the host character set for the session is setup correctly. By default, Host Access for the Cloud uses a character set which does not support the Euro character (€). To display the Euro character, change the character set to one that supports the Euro character.

Issues encountered with VT hosts

Type	Description
Performance issues	Heavy text output, such as form "Is-IR" may cause slow performance Scrolling regions may appear slow or choppy Cursor movement may be slow or choppy

Character sets	Graphical characters and some character sets are not supported. Some non-English characters may cause the terminal display to freeze.
Other VT issues	Insert/delete column (DECIC, DECDC) may fail. VT400 will not recognize DECSCS.

Field outlines in 3270 sessions

The 3270 attributes for field outlines are not fully supported. Host Access for the Cloud currently supports underline and overline; however, left vertical line, right vertical line, and combinations of the four line types are not yet supported.