

Reflection for the Web Deployment Guide

Reflection for the Web Deployment Guide

13.3.1

Table of contents

Reflection for the Web Deployment Guide	3
In this Guide	3
Legal Notice	3
Plan	4
Planning for Deployment	4
Installation Types	4
Standard Deployment	6
Fault Tolerance and Scaling	9
Next step	10
Install or Update	11
Install or Update	11
Virtual Appliance	11
Linux Installer	18
Uninstalling	23
Next Step after Installing	24
Configure	25
Configure Your Cluster	25
Clustering	26
Next Steps After Configuring Your Cluster	28
Apply	29
Apply your Product Configuration	29
Using Reflection for the Web	29
Technical References	34
Resources	41

Reflection for the Web Deployment Guide

Reflection for the Web is a web application that enables users to access IBM, UNIX, Unisys, OpenVMS, and HP data via a browser. Administrators use the Management and Security Server (MSS) Administrative Console to create, secure, assign, and centrally manage Reflection for the Web sessions.

Beginning with version 13.3, Reflection for the Web (RWeb) 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 or updating, and then configuring your product.

In this Guide

- [Plan for deployment \(page 4\)](#)
- [Install or Update Reflection for the Web \(page 11\)](#)
- [Configure your cluster \(page 25\)](#)
- [Apply your product configuration \(page 29\)](#)

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 \(page 8\)](#) and [Cluster certificate \(page 8\)](#).

Installation Types

There are two options when considering how to deploy Reflection for the Web. 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	<p>The virtual appliance is a pre-configured virtual machine that contains everything you need to run the system.</p> <p>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.</p>	<p>The simplest and recommended installation type. Choose this option if you:</p> <ul style="list-style-type: none"> Prefer easy, one-click software updates. Aim to minimize maintenance efforts. Have a virtualization platform that supports OVF (Open Virtualization Format) files. 	<p>View System Requirements (page 11)</p>


Installation Type	Description	When to choose	System Requirements
Linux Installer	The Linux installer is a <code>.sh</code> script that installs all of the software needed on an existing Linux server of your choice, whether it's virtual or physical.	<p>Choose this option if you:</p> <ul style="list-style-type: none"> 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 18)

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 \(page 11\)](#) or [Linux installer \(page 18\)](#) .

An odd number of nodes – always required.

A [load balancer \(page 7\)](#) with a [DNS hostname \(page 8\)](#) for the cluster.

A [certificate key pair \(page 8\)](#) for securing access to the cluster.

Load balancer

An external load balancer is provided by you. 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 the health endpoint for each node in your cluster

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

Requests to any node in the cluster will be automatically load-balanced by the system to nodes across the cluster. This action 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 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

If 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.

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 \(page 11\)](#) .

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 virtual 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


Other Reflection for the Web requirements

Check the requirements for the Reflection for the Web clients and SDK.

- **Terminal sessions (client)**

When clients are centrally managed by MSS, *only Windows* is supported.

An MSI file (Windows only) is provided to install the RWeb Launcher on Windows workstations, from which the RWeb client is launched.

 **Note**

You will use the MSS Assigned Sessions list and the Reflection for the Web Launcher to launch Reflection for the Web sessions.

• **RWeb SDK**

The RWeb SDK is supported on any platform.

Installing the Appliance

 **Note**

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 the [system requirements \(page 11\)](#) are met, 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:

Fully Qualified Domain Name (FQDN)

IP

Netmask

Gateway

DNS Server

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

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 RWeb, 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`.

[Next step \(page 24\)](#)

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 [Installing the Appliance \(page 14\)](#) 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 appliance registration. Product upgrades require a new activation key.

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 these 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 same 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 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=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

The Linux servers are provided by you. The Linux installer installs the appliance administration console, along with Reflection for the Web.

Supported operating systems

These versions or greater of the following operating systems are supported.

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

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 disable swapping.

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 (sl) 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

Other Reflection for the Web requirements

Check the requirements for the Reflection for the Web clients and SDK.

Terminal session (client)

When clients are centrally managed by MSS, *only Windows* is supported.

An MSI file (Windows only) is provided to install the RWeb Launcher on Windows workstations, from which the RWeb client is launched.

Note

You will use the MSS Assigned Sessions list and the Reflection for the Web Launcher to launch Reflection for the Web sessions.

RWeb SDK

The RWeb SDK is supported on any platform.

Installing using the Linux Installer

To install your product, be sure the [system requirements \(page 18\)](#) are met before you proceed.

Installation steps

1. From the Downloads site, download the Linux installer script (`install*.sh`) for Reflection for the Web.
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 automatically start, and you can move on to the [next step \(page 22\)](#).

If verification *fails*, see [Troubleshooting the Linux installation \(page 21\)](#).

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 and `sudo cspctl verify` runs without errors, run `sudo cspctl start` to start the system. Then run `sudo cspctl enable` to have the system start automatically after server restarts.

If issues remain, please contact Customer Support for assistance.

Next step (page 24)

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 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 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 \(page 21\)](#).

- c. After the CSP service starts, wait until the updated node show 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 `cordun` 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 used to install Reflection for the Web.

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-rweb.sh`.

The uninstall process takes a while to complete.

Next Step after Installing

After you install Reflection for the Web and create your cluster, it's time to [Configure your Deployment \(page 25\)](#) .

Configure

Configure Your Cluster

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. Although these settings can be set at a later time, we recommend setting them during initial configuration.

Accessing the MSS Administrative Console

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).

To access the MSS Admin Console:

1. Log in to `https://hostname/adminconsole`.
2. The Admin Console's default password is `admin`.

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

A single certificate, once configured, is automatically shared across all nodes in a cluster.

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.

Next: [Clustering \(page 26\)](#)

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 \(page 6\)](#).

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 \(page 27\)](#) or [Linux installer \(page 27\)](#)

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.

When clustering is complete, the Cluster Status will display a list of all nodes in the cluster with a status of "Ready." **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 step (page 28)

Next Steps After Configuring Your Cluster

You now have a cluster ready for use. You are ready to [apply your configuration \(page 29\)](#) and use Reflection for the Web.

Apply

Apply your Product Configuration

After you install Reflection for the Web and configure your deployment, you are ready to use Reflection for the Web to create and configure host sessions.

- [Using Reflection for the Web \(page 29\)](#)
- [Technical References \(page 34\)](#)
- [Resources \(page 41\)](#)

Using Reflection for the Web

Using Reflection for the Web

If you are new to Reflection for the Web, or need a little refresher, be sure to see [Reflection for the Web Overview \(page 37\)](#) for more detail about how Reflection for the Web works.

If you are already familiar with Reflection for the Web, continue to use your product as you did previously.

- [Create a Reflection for the Web session \(page 29\)](#)
- [Launch a session \(page 30\)](#)
- [Configure and assign a session \(page 30\)](#)
- [Distribute the Reflection for the Web Launcher Installer \(page 31\)](#)

Create a Reflection for the Web session

In the MSS Administrative Console, add a Reflection for the Web session and follow the prompts to download the Reflection for the Web Launcher Installer. Follow these steps.

1. On the MSS Administrative Console - **Manage Sessions** panel, click + ADD.
2. Select your Product – **Reflection for the Web**.
3. Select the Session type, such as IBM 3270.
4. Enter a unique Session name.

5. When prompted to install the Reflection for the Web Launcher, click LAUNCHER.
 - a. Click DOWNLOAD to download the Reflection for the Web Launcher, packaged as `RWebLauncher.msi`. When prompted, click **Save File**.
 - b. From your Downloads location, click and open `RWebLauncher.msi`.
 - c. Proceed through the RWebLauncher Setup wizard and click Finish.

At this point, the Reflection for the Web Launcher is installed, and you are able to add and launch sessions without being interrupted by dialog prompts to install the Launcher.

Note

This is a one-time task. Once the Launcher is installed, you can skip step 5 when creating sessions.

6. Unfortunately, the session you began adding above was not saved.
Return to **Manage Sessions** to repeat the steps to add and the session.
7. Next: [Launch a session \(page 30\)](#) .

Launch a session

You can launch a session by using either option:

- **Assigned Sessions** list. Click the name of the session, and click LAUNCH.
- **MSS Administrative Console**. In the Manage Sessions list, click a session, and click LAUNCH.

Next: [Configure and assign a session \(page 30\)](#)

Configure and assign a session

Once the Reflection for the Web Launcher is installed on the administrator's workstation, you can launch a session, configure it, and assign it to authorized users.

1. On the MSS Administrative Console - **Manage Sessions** panel, click a Reflection for the Web session name.
2. Click **LAUNCH**. You will be prompted to use Zulu Platform x32 Architecture to open the session.
3. **Configure the session settings**, such as Profiling (Administration > User Interface Profiler).
4. **Save the session**. When you click Save/Exit, the session is added to the Manage Sessions list in the MSS Administrative Console.
5. **Assign the session**. In the MSS Administrative Console, click Assign Access.

6. If LDAP authorization is enabled, you can search for a particular user or group. Select a user or group and check the session(s) you want to assign.

If LDAP is not enabled, you can assign sessions to All Users.

7. Click APPLY.

In the MSS Administrative Console, click **Currently Assigned** to see that the session is assigned to the user or group.

8. In the **Manage Sessions** list, click the session name and scroll to see the direct URL for the direct link to the session.

9. Deploy the session and—the Reflection for the Web Launcher—by choosing a distribution option.

10. Next: [Distribute the Reflection for the Web Launcher Installer \(page 31\)](#) to end users.

Note

After the Reflection for the Web Launcher is installed on the users' workstations, you can return to the Administrative Console to refine session settings. Refer to the [MSS Deployment Guide \(page 31\)](#) to set up Metering, the Security Proxy, Terminal ID Manager, and other features.

Distribute the RWeb Launcher

Distribute the Reflection for the Web Launcher Installer

The Reflection for the Web Launcher is installed on the administrator's workstation. Now you need to distribute the Reflection for the Web Launcher Installer to users' workstations so they can launch their assigned Reflection for the Web sessions.

The **Reflection for the Web Launcher Installer**, packaged as `RWebLauncher.msi`, can be distributed either by [using a software deployment system \(page 32\)](#), such as Microsoft Group Policies, or by [enabling individual downloads \(page 32\)](#).

Note

With either method, be sure to apply the [security updates \(page 32\)](#) as they become available.

- [Using a software deployment system \(page 32\)](#)
- [Enabling individual downloads \(page 32\)](#)
- [JRE security updates \(page 32\)](#)
- [Troubleshooting the Reflection for the Web Launcher \(page 33\)](#)

Using a software deployment system

Deploy the Reflection for the Web Launcher Installer (`RWebLauncher.msi`) file using the system of your choice, such as Microsoft Group Policies.

For reference, the Reflection for the Web Launcher Installer, packaged as `RWebLauncher.msi` , can be accessed from the MSS server: `MSS\server\web\webapps\rwebclient\ex\RWebLauncher.msi` .

Note

If you are using a web proxy, use the provided PowerShell script to generate a transform (`*.mst`) file that customizes `RWebLauncher.msi` with your preferred networking settings. Then, deploy the transform file (`RWebLauncherEx.mst`) along with the Reflection for the Web Launcher (`RWebLauncher.msi`), and install them together on the client machines. See [When using a web proxy \(page 35\)](#) .

Enabling individual downloads

Rather than using a software deployment system, you can let users download and install the Reflection for the Web Launcher the first time they open a Reflection for the Web session.

After you assign sessions to users, you need to distribute the session URL so an authorized user can access their **Assigned Sessions** list and launch a Reflection for the Web session directly.

When a user clicks a session link—before the Reflection for the Web Launcher is installed— they will be prompted to download and install the Launcher.

1. Provide each user with this link to their Assigned Sessions list:

`http://hostname[:port]/sessions` , where `<hostname>` is the name of the host where MSS is installed.

The list contains the sessions that were assigned to that authorized user.

2. The user clicks a link to launch a Reflection for the Web session.

When a Reflection for the Web session is launched, a transient dialog appears that provides the user with the ability to download the MSI installer directly. The dialog is automatically dismissed after a few seconds.

3. After the Reflection for the Web Launcher is installed on the users' workstations, their sessions launch directly, and the Download button can be ignored.

JRE security updates

JRE security updates are provided for the Reflection for the Web Launcher (in the `RWebLauncher.msi` file).

When product updates are available, the Reflection for the Web automated installer updates the Reflection for the Web Launcher file on the MSS server: `MSS\server\web\webapps\lweb-client\ex\RWebLauncher.msi`.

Important

Make sure the end-user workstations run the updated `.msi` file.

Troubleshooting the Reflection for the Web Launcher

If you encounter issues with the Reflection for the Web Launcher, troubleshoot as follows.

- [Error: "The system administrator has set policies to prevent this installation."](#) (page 33)
- [Use the Java Console](#) (page 33)

Error: "The system administrator has set policies to prevent this installation."

A Group Policy can be enabled that prohibits a Standard User in Windows from running the `"msiexec"` application.

When this policy is enabled, the user will see this error message when running the `RWebLauncher.msi` file: "The system administrator has set policies to prevent this installation."

Workaround: Use the "Run As..." feature in Windows, and run the `RWebLauncher.msi` file using the built-in "Administrator" account in Windows.

Use the Java Console


To troubleshoot the Reflection for the Web Launcher, open the **Reflection for the Web Launcher Settings** application to configure Java Console output, logging, caching, etc.

1. From the Start menu, select **Reflection for the Web Launcher Settings**.

The Iced Tea Web Control Panel opens.

2. Once launched, look at the **Debugging** section for options to troubleshoot Web Start and Reflection for the Web.

This screenshot shows the **Debugging Settings** after they have been enabled.



Technical References

Technical References

[Advanced Settings \(page 34\)](#)

[Reflection for the Web Launcher: a Web Start \(JNLP\) solution \(page 35\)](#)

[Reflection for the Web Overview \(page 37\)](#)

[Terms used with Reflection for the Web \(page 39\)](#)

[Migrating Legacy Data \(page 36\)](#)

[When using a web proxy \(page 35\)](#)


[Checklist for Planning \(page 40\)](#)

Advanced Product Settings


Information for your product is always available from the installed documentation or from [online resources \(page .](#)

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. See [instructions on how to use the Kubernetes dashboard \(page .](#)

Reflection for the Web Launcher: A Web Start (JNLP) solution

The **Reflection for the Web Launcher** is the client-side application that uses Web Start (JNLP) to launch Reflection for the Web sessions. This OpenJDK implementation eliminates the need for Oracle's JRE and the Java browser plug-in on either the administrator or client machines.

Note

JRE security updates are provided in Reflection for the Web product updates.

When using a web proxy

You can use your web proxy along with the Reflection for the Web Launcher to launch Reflection for the Web sessions. To configure web proxy settings centrally, you may use the provided PowerShell script.

The PowerShell script

Reflection for the Web installs a PowerShell script to

```
<installation directory>/server/web/webapps/rweb-client/utilities/rweb-launcher-msi-transformer/.
```

The PowerShell script generates a transform (*.mst) file that can be used to customize your installation. The script enables you to manage the settings that would otherwise need to be accessed on each workstation.

While the default network settings for the Reflection for the Web Launcher may work, they may need to be adjusted for a more complex network environment, such as with a web proxy.

Run the PowerShell script to centrally manage your network settings

Follow these steps:

1. Download and unzip this file:

<https://cluster-dns-ingress/rweb-client/utilities/MSITransformer.zip> (page

2. Download the MSI to be modified using this URL:

<https://localhost:8443/rweb-client/ex/RWebLauncher.msi> (page

3. Run this command:

```
.\InstallerTransformer.ps1 -f <path to MSI> -t <path to new Transform file>
```

The Result

The script confirms that the transform was successfully created in this directory: `<installation directory>\server\web\webapps\rweb-client\ex\RWebLauncherEx.mst`

After you deploy the transform file (`RWebLauncherEx.mst`) along with the Reflection for the Web Launcher (`RWebLauncher.msi`), your custom network settings will be deployed to the user's workstation the next time the user opens a Reflection for the Web session.

For deployment options, see [Distribute the Reflection for the Web Launcher Installer \(page 31\)](#) .

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.

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.

How to migrate your data

1. Log in to the MSS Admin Console at <https://hostname/adminconsole>.
2. Click Configure Settings > Migration.
3. Open the help for more information and detailed migration steps.

Next step

After your data is migrated, you are ready to [configure your cluster \(page 25\)](#).

Reflection for the Web Overview

Reflection for the Web provides Java-based applets to deploy web-based terminal emulation sessions to your users. Reflection for the Web's terminal sessions are centrally managed and secured using the Management and Security Server (MSS) Administrative Console.

Your Reflection for the Web license entitles you to

- Centralized management (MSS), which is automatically installed with Reflection for the Web
- Security Proxy
- Terminal ID Manager


Using Reflection for the Web and MSS, you can configure secure web-based terminal emulation sessions that connect to host applications located inside or outside the firewall.

Briefly, here's how it works:

1. An administrator installs Reflection for the Web on a server and either installs or uses an existing installation of Management and Security Server (MSS).
2. The administrator uses the MSS Administrative Console to create, configure, and secure terminal emulation sessions. Optional security settings can be configured on a per-session basis.
3. A user clicks a link to start a terminal session.
4. The Reflection for the Web session is downloaded to the user's workstation.

5. The user connects to and communicates with the host system using the downloaded emulation applet.

The diagram below depicts the interaction between Reflection for the Web, the MSS Administrative Server, and the optional Security Proxy Server to provide enhanced security.



1. Reflection for the Web user connects to the MSS Administrative Server.
2. User authenticates to a directory server (LDAP/Active Directory) or other identity management system – optional.
3. Directory server provides user and group identity – optional.
4. The Administrative Server sends a list of **Assigned Sessions** to the authenticated client. The user clicks a session.
5. When the optional Security Proxy Server is configured for use by a session, emulation applet makes a TLS connection to Security Proxy Server and sends it a signed token.
6. When present, the Security Proxy Server validates session token and establishes a connection to the `host:port` it specifies.
7. When no Security Proxy Server is present or a session is not configured to use it, an authenticated user connects directly to the host.

Administrative Server

The Management and Security Server *Administrative Server* includes the **MSS Administrative Console** and terminal emulation files, which are installed together on a web server.

After you install (or point to an existing) Management and Security Server, you can open the Administrative Console, which is a self-contained web application. Use the Administrative Console to manage and configure web-based terminal sessions. With Reflection for the Web, Java-based applets deploy terminal emulation sessions to your users.

Optional Components

Your Reflection for the Web license entitles you to these optional components in Management and Security Server:

- **Metering Server** monitors the use of terminal sessions.
- **Security Proxy Server** acts as a proxy for terminal sessions, routing encrypted network traffic to and from user workstations.
- **Terminal ID Manager** spools terminal IDs, tracks ID usage, and manages inactivity timeout values for specific users.

Note

Your Reflection for the Web license includes the Security Proxy and Terminal ID Manager, which are Add-On Products to Management and Security Server.

For information about installing, configuring, and using these components, see the [MSS Deployment Guide](#) (page .

Terms used with Reflection for the Web

- Java Cryptography Extension (JCE)

The Java Cryptography Extension (JCE) provides a framework and implementations for encryption, key generation and key agreement, and Message Authentication Code (MAC) algorithms.

- Java Runtime Environment (JRE)

The JRE is a subset of the JDK for end-users. It includes a Java Virtual Machine and a Java interpreter and provides a unified interface to Java programs, regardless of the underlying operating system.

- Java Server Pages (JSP)

A Java technology that helps software developers serve dynamically generated web pages based on HTML, XML, or other document types.

- Java Software Development Kit (JDK)

The JDK (previously called the Java SDK) is the software development environment for writing Java applets or applications; it is a superset of the Java Runtime Environment and the Java Virtual Machine.

- Java Virtual Machine (JVM or VM)

The JVM is the part of Java that interprets Java bytecode. Because the JVM is part of the JDK, it has the same version number. When a browser supports a specific version of the JDK, this includes the JVM.

- OpenJDK

Open Java Development Kit is a free and open-source implementation of the Java Platform, Standard Edition (Java SE). OpenJDK produces a number of components: the virtual machine (HotSpot), the Java Class Library and the Java compiler (javac), and does not include the web-browser plug-in or Web Start.

- Reflection for the Web Launcher

A client-side application that uses a JNLP implementation, along with Management and Security Server (MSS), to launch the emulator sessions. The Reflection for the Web Launcher does not need Oracle's JRE or the Java browser plug-in.

- Reflection for the Web Launcher Installer

The Windows .msi file package that installs the Reflection for the Web Launcher.

- Web Start (JNLP)

The Java Network Launch Protocol (JNLP) enables an application to be launched on a client desktop by using resources that are hosted on a remote web server. A properly configured browser passes JNLP files to a Java Runtime Environment (JRE), which in turn downloads the application onto the user's machine and starts executing it. The Reflection for the Web Launcher installs an OpenJDK JRE with Web Start (JNLP) to launch Reflection for the Web sessions.

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 Reflection for the Web as you did previously.

Create a session and install the RWeb Launcher.


Launch a session.

Configure and assign sessions.

Distribute the Reflection for the Web Launcher.

Resources

Information for Reflection for the Web is always available from the installed product documentation or online resources

- Reflection for the Web Help. In an open session, click  to see the emulation help.
- [Reflection for the Web Reference Guide \(page](#) - includes scripting, HTML examples, and other advanced topics.
- [MSS Product Documentation \(page](#)
- [Support Resources \(page](#)