Upgrade your NetApp SolidFire all-flash storage system version Element 12.3.x

Element Software

NetApp
November 19, 2025

This PDF was generated from https://docs.netapp.com/us-en/element-software-123/upgrade/concept_element_upgrade_overview.html on November 19, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Upgrade your NetApp SolidFire all-flash storage system version Element 12.3.x . . . 1
- Upgrade sequence overview . . . 1
  - System upgrade sequence . . . 1
- System upgrade procedures . . . 2
  - Update management services . . . 2
  - Upgrade to the latest HealthTools . . . 5
  - Run Element storage health checks prior to upgrading storage . . . 7
  - Upgrade Element software . . . 15
  - Upgrade storage firmware . . . 32
  - Upgrade a management node . . . 41
  - Upgrade the Element Plug-in for vCenter Server . . . 55
- Upgrade your vSphere components for a NetApp SolidFire storage system with the Element Plug-in for vCenter Server . . . 62
  - Find more information . . . 62

Upgrade your NetApp SolidFire all-flash storage system version Element 12.3.x

Upgrade sequence overview

You can keep your SolidFire Element storage system up-to-date after deployment by sequentially upgrading all NetApp storage components.

These components include management services, HealthTools, NetApp Hybrid Cloud Control, Element software, management node, and (depending on your installation) the Element Plug-in for vCenter Server.

Beginning November 2023, you cannot start a component upgrade using NetApp Hybrid Cloud Control or REST API because the signing key certificates (private and public) expired on November 5, 2023. You can resolve this issue by following the workaround documented in the Knowledge Base article SolidFire and HCI upgrades unable to start due to upgrade package upload error.

The system upgrade sequence content describes the tasks that are needed to complete a SolidFire all-flash storage system upgrade. Ideally these procedures are performed as part of the larger upgrade sequence and not in isolation. If a component-based upgrade or update is needed, see the procedure prerequisites to ensure additional complexities are addressed.

The vSphere upgrade sequence including Element Plug-in for vCenter Server content describes additional pre- and post-upgrade steps required to re-install the Element Plug-in for vCenter Server.

What you'll need

- You are running management node 11.3 or later. Newer versions of the management node have a modular architecture that provides individual services.

  To check the version, log in to your management node and view the Element version number in the login banner. If you do not have 11.3, see Upgrade your management node.

- You have upgraded your management services to at least version 2.1.326.

  Upgrades using NetApp Hybrid Cloud Control are not available in earlier service bundle versions.

- You have ensured that the system time on all nodes is synced and that NTP is correctly configured for the storage cluster and nodes. Each node must be configured with a DNS nameserver in the per-node web UI (`https://[IP address]:442`) with no unresolved cluster faults related to time skew.

System upgrade sequence

You can use the following sequence to upgrade your NetApp SolidFire all-flash storage system for Element 12.3.x.

Steps

1. Update management services from Hybrid Cloud Control.

If you are updating management services to version 2.16 or later and you are running a management node 11.3 to 11.8, you will need to increase your management node VM's RAM prior to updating management services.

You must update to the latest management services bundle before upgrading your Element software.

2. (Optional) Upgrade to the latest HealthTools.

Upgrading HealthTools is only required if the management node and Element software you are running is 11.1 or earlier. HealthTools are not required for performing Element upgrades using NetApp Hybrid Cloud Control.

3. Run Element storage health checks prior to upgrading storage.
4. Upgrade your Element software and storage firmware.
5. (Optional) Upgrade your Element storage firmware only.

You might perform this task when a new storage firmware upgrade becomes available outside of a major release.

6. (Optional) Upgrade your management node.

Upgrading the management node operating system is no longer required to upgrade Element software on the storage cluster. If the management node is version 11.3 or higher, you can simply upgrade the management services to the latest version to perform Element upgrades using NetApp Hybrid Cloud Control. Follow the management node upgrade procedure for your scenario if you would like to upgrade the management node operating system for other reasons, such as security remediation.

7. Upgrade your Element Plug-in for vCenter Server.

Find more information

- NetApp Element Plug-in for vCenter Server
- SolidFire and Element Software Documentation

System upgrade procedures

Update management services

You can update your management services to the latest bundle version after you have installed management node 11.3 or later.

Beginning with the Element 11.3 management node release, the management node design has been changed based on a new modular architecture that provides individual services. These modular services provide central and extended management functionality for a SolidFire all-flash storage system. Management services include system telemetry, logging, and update services, the QoSSIOC service for Element Plug-in for vCenter Server, NetApp Hybrid Cloud Control, and more.

About this task

- You must upgrade to the latest management services bundle before upgrading your Element software.

- Management services 2.22.7 includes Element Plug-in for vCenter Server 5.0 which contains the remote plug-in. If you use the Element plug-in, you should upgrade to management services 2.22.7 or later to comply with the VMware directive that removes support for local plug-ins. Learn more.
- For the latest management services release notes describing major services, new features, bug fixes, and workarounds for each service bundle, see the management services release notes

What you'll need

Beginning with management services 2.20.69, you must accept and save the End User License Agreement (EULA) before using the NetApp Hybrid Cloud Control UI or API to upgrade management services:

1. Open the IP address of the management node in a web browser:

```
https://<ManagementNodeIP>
```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Update options

You can update management services using the NetApp Hybrid Cloud Control UI or the management node REST API:

- Update management services using Hybrid Cloud Control (Recommended method)
- Update management services using the management node API

Update management services using Hybrid Cloud Control

You can update your NetApp management services using NetApp Hybrid Cloud Control.

Management service bundles provide enhanced functionality and fixes to your installation outside of major releases.

Before you begin

- You are running management node 11.3 or later.
- If you are updating management services to version 2.16 or later and you are running a management node 11.3 to 11.8, you will need to increase your management node VM's RAM prior to updating management services:
  a. Power off the management node VM.
  b. Change the RAM of the management node VM from 12GB to 24GB RAM.
  c. Power on the management node VM.
- Your cluster version is running NetApp Element software 11.3 or later.
- You have upgraded your management services to at least version 2.1.326. NetApp Hybrid Cloud Control

upgrades are not available in earlier service bundles.

> For a list of available services for each service bundle version, see the Management Services Release Notes.

Steps

1. Open the IP address of the management node in a web browser:

   ```
   https://<ManagementNodeIP>
   ```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the Upgrades page, select the **Management Services** tab.
5. Follow the instructions on the page to download and save a management services upgrade package to your computer.
6. Select **Browse** to locate the package you saved and upload it.

   After you upload the package, the upgrade starts automatically.

   After the upgrade begins, you can see the upgrade status on this page. During the upgrade, you might lose connection with NetApp Hybrid Cloud Control and have to log back in to see the results of the upgrade.

Update management services using the management node API

Users should ideally perform management services updates from NetApp Hybrid Cloud Control. You can however manually upload, extract, and deploy a service bundle update for management services to the management node using the REST API. You can run each command from the REST API UI for the management node.

Before you begin

- You have deployed a NetApp Element software management node 11.3 or later.
- If you are updating management services to version 2.16 or later and you are running a management node 11.3 to 11.8, you will need to increase your management node VM's RAM prior to updating management services:
  a. Power off the management node VM.
  b. Change the RAM of the management node VM from 12GB to 24GB RAM.
  c. Power on the management node VM.
- Your cluster version is running NetApp Element software 11.3 or later.
- You have upgraded your management services to at least version 2.1.326. NetApp Hybrid Cloud Control upgrades are not available in earlier service bundles.

> For a list of available services for each service bundle version, see the Management Services Release Notes.

Steps

1. Open the REST API UI on the management node: `https://<ManagementNodeIP>/mnode`
2. Select **Authorize** and complete the following:
   a. Enter the cluster user name and password.
   b. Enter the client ID as `mnode-client` if the value is not already populated.
   c. Select **Authorize** to begin a session.
   d. Close the window.
3. Upload and extract the service bundle on the management node using this command: `PUT /services/upload`
4. Deploy the management services on the management node: `PUT /services/deploy`
5. Monitor the status of the update: `GET /services/update/status`

   A successful update returns a result similar to the following example:

```
{
"current_version": "2.10.29",
"details": "Updated to version 2.17.52",
"status": "success"
}
```

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Upgrade to the latest HealthTools

Before you begin an Element storage upgrade from 11.1 or earlier, you should upgrade your HealthTools suite. Upgrading HealthTools is only required if the management node and Element software you are running is 11.1 or earlier. HealthTools are not required for performing Element upgrades using NetApp Hybrid Cloud Control.

> Element software 12.3.2 is the final version that you can upgrade to using NetApp HealthTools. If you are running Element software 11.3 or later, you should use NetApp Hybrid Cloud Control to upgrade Element software. You can upgrade Element versions 11.1 or earlier using NetApp HealthTools.

What you'll need

- You are running management node 11.0, 11.1 or later.
- You have upgraded your management services to at least version 2.1.326.

  NetApp Hybrid Cloud Control upgrades are not available in earlier service bundle versions.
- You have downloaded the latest version of HealthTools and copied the installation file to the management node.

You can check the locally installed version of HealthTools by running the `sfupdate-healthtools -v` command.

- To use HealthTools with dark sites, you need to do these additional steps:
  - Download a JSON file from the NetApp Support Site on a computer that is not the management node and rename it to `metadata.json`.
  - Have the management node up and running at the dark site.

About this task

The commands in the HealthTools suite require escalated privileges to run. Either preface commands with `sudo` or escalate your user to root privileges.

The HealthTools version you use might be more up to date than the sample input and response below.

Steps

1. Run the `sfupdate-healthtools <path to install file>` command to install the new HealthTools software.

   Sample input:

   ```
   sfupdate-healthtools /tmp/solidfire-healthtools-2020.03.01.09.tgz
   ```

   Sample response:

   ```
   Checking key signature for file /tmp/solidfirehealthtools-
   2020.03.01.09/components.tgz
   installing command sfupdate-healthtools
   Restarting on version 2020.03.01.09
   sfupdate-healthtools /sf/bin/sfupdate-healthtools -r 2020.03.01.09
   installing command sfupgradecheck
   installing command sfinstall
   installing command sfresetupgrade
   ```

2. Run the `sfupdate-healthtools -v` command to verify the installed version has been upgraded.

   Sample response:

   ```
   Currently installed version of HealthTools:
   2020.03.01.09
   ```

Find more information

- SolidFire and Element Software Documentation

- NetApp Element Plug-in for vCenter Server

Run Element storage health checks prior to upgrading storage

You must run health checks prior to upgrading Element storage to ensure all storage nodes in your cluster are ready for the next Element storage upgrade.

What you'll need

- **Management services**: You have updated to the latest management services bundle (2.10.27 or later).

  > You must upgrade to the latest management services bundle before upgrading your Element software.

- **Management node**:You are running management node 11.3 or later.
- **Element software**: Your cluster version is running NetApp Element software 11.3 or later.
- **End User License Agreement (EULA)**: Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to run Element storage health checks:
  1. Open the IP address of the management node in a web browser:

     ```
     https://<ManagementNodeIP>
     ```

  2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
  3. Select **Upgrade** near the top right of the interface.
  4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Health check options

You can run health checks using the NetApp Hybrid Cloud Control UI, the NetApp Hybrid Cloud Control API, or the HealthTools suite:

- Use NetApp Hybrid Cloud Control to run Element storage health checks prior to upgrading storage (Preferred method)
- Use API to run Element storage health checks prior to upgrading storage
- Use HealthTools to run Element storage health checks prior to upgrading storage

You can also find out more about storage health checks that are run by the service:

- Storage health checks made by the service

Use NetApp Hybrid Cloud Control to run Element storage health checks prior to upgrading storage

Using NetApp Hybrid Cloud Control, you can verify that a storage cluster is ready to be upgraded.

Steps

1. Open the IP address of the management node in a web browser:

```
https://<ManagementNodeIP>
```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select the **Storage** tab.
5. Select the health check for the cluster you want to check for upgrade readiness.
6. On the **Storage Health Check** page, select **Run Health Check**.
7. If there are issues, do the following:
   a. Go to the specific KB article listed for each issue or perform the specified remedy.
   b. If a KB is specified, complete the process described in the relevant KB article.
   c. After you have resolved cluster issues, select **Re-Run Health Check**.

After the health check completes without errors, the storage cluster is ready to upgrade. See storage node upgrade instructions to proceed.

Use API to run Element storage health checks prior to upgrading storage

You can use REST API to verify that a storage cluster is ready to be upgraded. The health check verifies that there are no obstacles to upgrading, such as pending nodes, disk space issues, and cluster faults.

Steps

1. Locate the storage cluster ID:
   a. Open the management node REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/mnode
      ```

   b. Select **Authorize** and complete the following:
      i. Enter the cluster user name and password.
      ii. Enter the client ID as `mnode-client` if the value is not already populated.
      iii. Select **Authorize** to begin a session.
      iv. Close the authorization window.
   c. From the REST API UI, select `GET /assets`.
   d. Select **Try it out**.
   e. Select **Execute**.
   f. From the response, copy the `"id"` from the `"storage"` section of the cluster you intend to check for upgrade readiness.

      Do not use the `"parent"` value in this section because this is the management node's ID, not the storage cluster's ID.

```
"config": {},
"credentialid": "12bbb2b2-f1be-123b-1234-12c3d4bc123e",
"host_name": "SF_DEMO",
"id": "12cc3a45-e6e7-8d91-a2bb-0bdb3456b789",
"ip": "10.123.12.12",
"parent": "d123ec42-456e-8912-ad3e-4bd56f4a789a",
"sshcredentialid": null,
"ssl_certificate": null
```

2. Run health checks on the storage cluster:
   a. Open the storage REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/storage/1/
      ```

   b. Select **Authorize** and complete the following:
      i. Enter the cluster user name and password.
      ii. Enter the client ID as `mnode-client` if the value is not already populated.
      iii. Select **Authorize** to begin a session.
      iv. Close the authorization window.
   c. Select **POST /health-checks**.
   d. Select **Try it out**.
   e. In the parameter field, enter the storage cluster ID obtained in Step 1.

      ```
      {
        "config": {},
        "storageId": "123a45b6-1a2b-12a3-1234-1a2b34c567d8"
      }
      ```

   f. Select **Execute** to run a health check on the specified storage cluster.

      The response should indicate state as `initializing`:

```
{
  "_links": {
    "collection": "https://10.117.149.231/storage/1/health-checks",
    "log": "https://10.117.149.231/storage/1/health-checks/358f073f-
896e-4751-ab7b-ccbb5f61f9fc/log",
    "self": "https://10.117.149.231/storage/1/health-checks/358f073f-
896e-4751-ab7b-ccbb5f61f9fc"
  },
  "config": {},
  "dateCompleted": null,
  "dateCreated": "2020-02-21T22:11:15.476937+00:00",
  "healthCheckId": "358f073f-896e-4751-ab7b-ccbb5f61f9fc",
  "state": "initializing",
  "status": null,
  "storageId": "c6d124b2-396a-4417-8a47-df10d647f4ab",
  "taskId": "73f4df64-bda5-42c1-9074-b4e7843dbb77"
}
```

g. Copy the `healthCheckID` that is part of response.

3. Verify the results of the health checks:
   a. Select **GET /health-checks/{healthCheckId}**.
   b. Select **Try it out**.
   c. Enter the health check ID in the parameter field.
   d. Select **Execute**.
   e. Scroll to the bottom of the response body.

      If all health checks are successful, the return is similar to the following example:

```
"message": "All checks completed successfully.",
"percent": 100,
"timestamp": "2020-03-06T00:03:16.321621Z"
```

4. If the `message` return indicates that there were problems regarding cluster health, do the following:
   a. Select **GET /health-checks/{healthCheckId}/log**
   b. Select **Try it out**.
   c. Enter the health check ID in the parameter field.
   d. Select **Execute**.
   e. Review any specific errors and obtain their associated KB article links.
   f. Go to the specific KB article listed for each issue or perform the specified remedy.
   g. If a KB is specified, complete the process described in the relevant KB article.
   h. After you have resolved cluster issues, run **GET /health-checks/{healthCheckId}/log** again.

Use HealthTools to run Element storage health checks prior to upgrading storage

You can verify that the storage cluster is ready to be upgraded by using the `sfupgradecheck` command. This command verifies information such as pending nodes, disk space, and cluster faults.

If your management node is at a dark site without external connectivity, the upgrade readiness check needs the `metadata.json` file you downloaded during HealthTools upgrades to run successfully.

About this task

This procedure describes how to address upgrade checks that yield one of the following results:

- Running the `sfupgradecheck` command runs successfully. Your cluster is upgrade ready.
- Checks within the `sfupgradecheck` tool fail with an error message. Your cluster is not upgrade ready and additional steps are required.
- Your upgrade check fails with an error message that HealthTools is out-of-date.
- Your upgrade check fails because your management node is on a dark site.

Steps

1. Run the `sfupgradecheck` command:

```
sfupgradecheck -u <cluster-user-name> MVIP
```

> For passwords that contain special characters, add a backslash (`\`) before each special character. For example, `mypass!@1` should be entered as `mypass\!\@`.

Sample input command with sample output in which no errors appear and you are ready to upgrade:

```
sfupgradecheck -u admin 10.117.78.244
```

```
check_pending_nodes:
Test Description: Verify no pending nodes in cluster
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltOQAQ/pendingnodes
check_cluster_faults:
Test Description: Report any cluster faults
check_root_disk_space:
Test Description: Verify node root directory has at least 12 GBs of
available disk space
Passed node IDs: 1, 2, 3
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltTQAQ/
SolidFire-Disk-space-error
check_mnode_connectivity:
Test Description: Verify storage nodes can communicate with management
node
Passed node IDs: 1, 2, 3
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltYQAQ/mNodeconnecti
vity
check_files:
Test Description: Verify options file exists
Passed node IDs: 1, 2, 3
check_cores:
Test Description: Verify no core or dump files exists
Passed node IDs: 1, 2, 3
check_upload_speed:
Test Description: Measure the upload speed between the storage node and
the
management node
Node ID: 1 Upload speed: 90063.90 KBs/sec
Node ID: 3 Upload speed: 106511.44 KBs/sec
Node ID: 2 Upload speed: 85038.75 KBs/sec
```

2. If there are errors, additional actions are required. See the following sub-sections for details.

Your cluster is not upgrade ready

If you see an error message related to one of the health checks, follow these steps:

1. Review the `sfupgradecheck` error message.

   Sample response:

```
The following tests failed:
check_root_disk_space:
Test Description: Verify node root directory has at least 12 GBs of
available disk space
Severity: ERROR
Failed node IDs: 2
Remedy: Remove unneeded files from root drive
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltTQAQ/SolidFire-
Disk-space-error
check_pending_nodes:
Test Description: Verify no pending nodes in cluster
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltOQAQ/pendingnodes
check_cluster_faults:
Test Description: Report any cluster faults
check_root_disk_space:
Test Description: Verify node root directory has at least 12 GBs of
available disk space
Passed node IDs: 1, 3
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltTQAQ/SolidFire-
Disk-space-error
check_mnode_connectivity:
Test Description: Verify storage nodes can communicate with management
node
Passed node IDs: 1, 2, 3
More information:
https://kb.netapp.com/support/s/article/ka11A0000008ltYQAQ/mNodeconnectivi
ty
check_files:
Test Description: Verify options file exists
Passed node IDs: 1, 2, 3
check_cores:
Test Description: Verify no core or dump files exists
Passed node IDs: 1, 2, 3
check_upload_speed:
Test Description: Measure the upload speed between the storage node and
the management node
Node ID: 1 Upload speed: 86518.82 KBs/sec
Node ID: 3 Upload speed: 84112.79 KBs/sec
Node ID: 2 Upload speed: 93498.94 KBs/sec
```

In this example, node 1 is low on disk space. You can find more information in the knowledge base (KB) article listed in the error message.

HealthTools is out of date

If you see an error message indicating that HealthTools is not the latest version, follow these instructions:

1. Review the error message and note that the upgrade check fails.

   Sample response:

   ```
   sfupgradecheck failed: HealthTools is out of date:
   installed version: 2018.02.01.200
   latest version: 2020.03.01.09.
   The latest version of the HealthTools can be downloaded from:
   https://mysupport.netapp.com/NOW/cgi-bin/software/
   Or rerun with the -n option
   ```

2. Follow the instructions described in the response.

Your management node is on a dark site

1. Review the message and note that the upgrade check fails:

   Sample response:

   ```
   sfupgradecheck failed: Unable to verify latest available version of
   healthtools.
   ```

2. Download a JSON file from the NetApp Support Site on a computer that is not the management node and rename it to `metadata.json`.
3. Run the following command:

   ```
   sfupgradecheck -l --metadata=<path-to-metadata-json>
   ```

4. For details, see additional HealthTools upgrades information for dark sites.
5. Verify that the HealthTools suite is up-to-date by running the following command:

   ```
   sfupgradecheck -u <cluster-user-name> -p <cluster-password> MVIP
   ```

Storage health checks made by the service

Storage health checks make the following checks per cluster.

Check Name	Node/Cluster	Description
check_async_results	Cluster	Verifies that the number of asynchronous results in the database is below a threshold number.
check_cluster_faults	Cluster	Verifies that there are no upgrade blocking cluster faults (as defined in Element source).
check_upload_speed	Node	Measures the upload speed between the storage node and the management node.
connection_speed_check	Node	Verifies that nodes have connectivity to the management node serving upgrade packages and estimates connection speed.
check_cores	Node	Checks for kernel crash dump and core files on the node. The check fails for any crashes in a recent time period (threshold 7 days).
check_root_disk_space	Node	Verifies the root file system has sufficient free space to perform an upgrade.
check_var_log_disk_space	Node	Verifies that `/var/log` free space meets some percentage free threshold. If it does not, the check will rotate and purge older logs in order to fall under threshold. The check fails if it is unsuccessful at creating sufficient free space.
check_pending_nodes	Cluster	Verifies that there are no pending nodes on the cluster.

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Upgrade Element software

To upgrade NetApp Element software, you can use the NetApp Hybrid Cloud Control UI, REST API, or the HealthTools suite of tools. Certain operations are suppressed during an Element software upgrade, such as adding and removing nodes, adding and removing drives, and commands associated with initiators, volume access groups, and virtual networks, among others.

What you'll need

- **Admin privileges**: You have storage cluster administrator permissions to perform the upgrade.

- **Valid upgrade path**: You have checked upgrade path information for the Element version you are upgrading to and verified that the upgrade path is valid.
  NetApp KB: Upgrade matrix for storage clusters running NetApp Element Software
- **System time sync**: You have ensured that the system time on all nodes is synced and that NTP is correctly configured for the storage cluster and nodes. Each node must be configured with a DNS nameserver in the per-node web UI (`https://[IP address]:442`) with no unresolved cluster faults related to time skew.
- **System ports**: If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See Network ports for more information.
- **Management node**: For NetApp Hybrid Cloud Control UI and API, the management node in your environment is running version 11.3.
- **Management services**: You have updated your management services bundle to the latest version.

> You must upgrade to the latest management services bundle before upgrading your Element software to version 12.3.x. If you are updating your Element software to version 12.3.x, you need management services 2.14.60 or later to proceed.

- **Cluster health**: You have verified that the cluster is ready to be upgraded. See Run Element storage health checks prior to upgrading storage.
- **Updated BMC for H610S nodes**: You have upgraded the BMC version for your H610S nodes. See the release notes and upgrade instructions.
- **End User License Agreement (EULA)**: Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to upgrade Element software:
  1. Open the IP address of the management node in a web browser:

     ```
     https://<ManagementNodeIP>
     ```

  2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
  3. Select **Upgrade** near the top right of the interface.
  4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Upgrade options

Choose one of the following Element software upgrade options:

- Use NetApp Hybrid Cloud Control UI to upgrade Element storage
- Use NetApp Hybrid Cloud Control API to upgrade Element storage
- Upgrade Element software at connected sites using HealthTools
- Upgrade Element software at dark sites using HealthTools

> If you are upgrading an H610S series node to Element 12.3.x and the node is running a version of Element earlier than 11.8, you will need to perform additional upgrade steps (phase 2) for each storage node. If you are running Element 11.8 or later, the additional upgrade steps (phase 2) are not required.

Use NetApp Hybrid Cloud Control UI to upgrade Element storage

Using the NetApp Hybrid Cloud Control UI, you can upgrade a storage cluster.

For potential issues while upgrading storage clusters using NetApp Hybrid Cloud Control and their workarounds, see the KB article.

The upgrade process takes approximately 30 minutes per node for non-H610S platforms.

Steps

1. Open the IP address of the management node in a web browser:

   ```
   https://<ManagementNodeIP>
   ```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select **Storage**.

   The **Storage** tab lists the storage clusters that are part of your installation. If a cluster is inaccessible by NetApp Hybrid Cloud Control, it will not be displayed on the **Upgrades** page.

5. Choose from the following options and perform the set of steps that are applicable to your cluster:

Option	Steps
All clusters running Element 11.8 and later	1. Select **Browse** to upload the upgrade package that you downloaded. 2. Wait for the upload to complete. A progress bar shows the status of the upload. The file upload will be lost if you navigate away from the browser window. An on-screen message is displayed after the file is successfully uploaded and validated. Validation might take several minutes. If you navigate away from the browser window at this stage, the file upload is preserved. 3. Select **Begin Upgrade**. The **Upgrade Status** changes during the upgrade to reflect the status of the process. It also changes in response to actions you take, such as pausing the upgrade, or if the upgrade returns an error. See Upgrade status changes. While the upgrade is in progress, you can leave the page and come back to it later to continue monitoring the progress. The page does not dynamically update status and current version if the cluster row is collapsed. The cluster row must be expanded to update the table or you can refresh the page. You can download logs after the upgrade is complete.

Option	Steps
You are upgrading an H610S cluster running Element version earlier than 11.8.	1. Select the drop-down arrow next to the cluster you are upgrading, and select from the upgrade versions available. 2. Select **Begin Upgrade**. After the upgrade is complete, the UI prompts you to perform phase 2 of the process. 3. Complete the additional steps required (phase 2) in the KB article, and acknowledge in the UI that you have completed phase 2. You can download logs after the upgrade is complete. For information about the various upgrade status changes, see Upgrade status changes.

Upgrade status changes

Here are the different states that the **Upgrade Status** column in the UI shows before, during, and after the upgrade process:

Upgrade state	Description
Up to Date	The cluster was upgraded to the latest Element version available.
Versions Available	Newer versions of Element and/or storage firmware are available for upgrade.
In Progress	The upgrade is in progress. A progress bar shows the upgrade status. On-screen messages also show node-level faults and display the node ID of each node in the cluster as the upgrade progresses. You can monitor the status of each node using the Element UI or the NetApp Element plug-in for vCenter Server UI.
Upgrade Pausing	You can choose to pause the upgrade. Depending on the state of the upgrade process, the pause operation can succeed or fail. You will see a UI prompt asking you to confirm the pause operation. To ensure that the cluster is in a safe spot before pausing an upgrade, it can take up to two hours for the upgrade operation to be completely paused. To resume the upgrade, select **Resume**.
Paused	You paused the upgrade. Select **Resume** to resume the process.

Upgrade state	Description
Error	An error has occurred during the upgrade. You can download the error log and send it to NetApp Support. After you resolve the error, you can return to the page, and select **Resume**. When you resume the upgrade, the progress bar goes backwards for a few minutes while the system runs the health check and checks the current state of the upgrade.
Complete with Follow-up	Only for H610S nodes upgrading from Element version earlier than 11.8. After phase 1 of the upgrade process is complete, this state prompts you to perform phase 2 of the upgrade (see the KB article). After you complete phase 2 and acknowledge that you have completed it, the status changes to **Up to Date**.

Use NetApp Hybrid Cloud Control API to upgrade Element storage

You can use APIs to upgrade storage nodes in a cluster to the latest Element software version. You can use an automation tool of your choice to run the APIs. The API workflow documented here uses the REST API UI available on the management node as an example.

Steps

1. Download the storage upgrade package to a device that is accessible to the management node.

   Go to the Element software downloads page and download the latest storage node image.

2. Upload the storage upgrade package to the management node:
   a. Open the management node REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/package-repository/1/
      ```

   b. Select **Authorize** and complete the following:
      i. Enter the cluster user name and password.
      ii. Enter the client ID as `mnode-client`.
      iii. Select **Authorize** to begin a session.
      iv. Close the authorization window.
   c. From the REST API UI, select **POST /packages**.
   d. Select **Try it out**.
   e. Select **Browse** and select the upgrade package.
   f. Select **Execute** to initiate the upload.
   g. From the response, copy and save the package ID (`"id"`) for use in a later step.
3. Verify the status of the upload.
   a. From the REST API UI, select **GET /packages/{id}/status**.
   b. Select **Try it out**.

c. Enter the package ID you copied in the previous step in **id**.
d. Select **Execute** to initiate the status request.

   The response indicates `state` as `SUCCESS` when complete.

4. Locate the storage cluster ID:
   a. Open the management node REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/inventory/1/
      ```

   b. Select **Authorize** and complete the following:
      i. Enter the cluster user name and password.
      ii. Enter the client ID as `mnode-client`.
      iii. Select **Authorize** to begin a session.
      iv. Close the authorization window.
   c. From the REST API UI, select **GET /installations**.
   d. Select **Try it out**.
   e. Select **Execute**.
   f. From the response, copy the installation asset ID (`"id"`).
   g. From the REST API UI, select **GET /installations/{id}**.
   h. Select **Try it out**.
   i. Paste the installation asset ID into the **id** field.
   j. Select **Execute**.
   k. From the response, copy and save the storage cluster ID (`"id"`) of the cluster you intend to upgrade for use in a later step.
5. Run the storage upgrade:
   a. Open the storage REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/storage/1/
      ```

   b. Select **Authorize** and complete the following:
      i. Enter the cluster user name and password.
      ii. Enter the client ID as `mnode-client`.
      iii. Select **Authorize** to begin a session.
      iv. Close the authorization window.
   c. Select **POST /upgrades**.
   d. Select **Try it out**.
   e. Enter the upgrade package ID in the parameter field.

f. Enter the storage cluster ID in the parameter field.

The payload should look similar to the following example:

```
{
  "config": {},
  "packageId": "884f14a4-5a2a-11e9-9088-6c0b84e211c4",
  "storageId": "884f14a4-5a2a-11e9-9088-6c0b84e211c4"
}
```

g. Select **Execute** to initiate the upgrade.

The response should indicate the state as `initializing`:

```
{
  "_links": {
    "collection": "https://localhost:442/storage/upgrades",
    "self": "https://localhost:442/storage/upgrades/3fa85f64-1111-
4562-b3fc-2c963f66abc1",
    "log": https://localhost:442/storage/upgrades/3fa85f64-1111-4562-
b3fc-2c963f66abc1/log
  },
  "storageId": "114f14a4-1a1a-11e9-9088-6c0b84e200b4",
  "upgradeId": "334f14a4-1a1a-11e9-1055`-6c0b84e2001b4",
  "packageId": "774f14a4-1a1a-11e9-8888-6c0b84e200b4",
  "config": {},
  "state": "initializing",
  "status": {
    "availableActions": [
      "string"
    ],
    "message": "string",
    "nodeDetails": [
      {
        "message": "string",
        "step": "NodePreStart",
        "nodeID": 0,
        "numAttempt": 0
      }
    ],
    "percent": 0,
    "step": "ClusterPreStart",
    "timestamp": "2020-04-21T22:10:57.057Z",
    "failedHealthChecks": [
      {
```

```
        "checkID": 0,
        "name": "string",
        "displayName": "string",
        "passed": true,
        "kb": "string",
        "description": "string",
        "remedy": "string",
        "severity": "string",
        "data": {},
        "nodeID": 0
      }
    ]
  },
  "taskId": "123f14a4-1a1a-11e9-7777-6c0b84e123b2",
  "dateCompleted": "2020-04-21T22:10:57.057Z",
  "dateCreated": "2020-04-21T22:10:57.057Z"
}
```


h. Copy the upgrade ID (`"upgradeId"`) that is part of the response.

6. Verify the upgrade progress and results:
   a. Select **GET /upgrades/{upgradeId}**.
   b. Select **Try it out**.
   c. Enter the upgrade ID from the previous step in **upgradeId**.
   d. Select **Execute**.
   e. Do one of the following if there are problems or special requirements during the upgrade:

Option	Steps
You need to correct cluster health issues due to `failedHealthChecks` message in the response body.	1. Go to the specific KB article listed for each issue or perform the specified remedy. 2. If a KB is specified, complete the process described in the relevant KB article. 3. After you have resolved cluster issues, reauthenticate if needed and select **PUT /upgrades/{upgradeId}**. 4. Select **Try it out**. 5. Enter the upgrade ID from the previous step in **upgradeId**. 6. Enter `"action":"resume"` in the request body. `{ "action": "resume" }` 7. Select **Execute**.
You need to pause the upgrade because the maintenance window is closing or for another reason.	1. Reauthenticate if needed and select **PUT /upgrades/{upgradeId}**. 2. Select **Try it out**. 3. Enter the upgrade ID from the previous step in **upgradeId**. 4. Enter `"action":"pause"` in the request body. `{ "action": "pause" }` 5. Select **Execute**.

Option	Steps
If you are upgrading an H610S cluster running an Element version earlier than 11.8, you see the state `finishedNeedsAck` in the response body. You need to perform additional upgrade steps (phase 2) for each H610S storage node.	1. See Upgrading H610S storage nodes to Element 12.3.x (phase 2) and complete the process for each node. 2. Reauthenticate if needed and select **PUT /upgrades/{upgradeId}**. 3. Select **Try it out**. 4. Enter the upgrade ID from the previous step in **upgradeId**. 5. Enter `"action":"acknowledge"` in the request body. `{ "action": "acknowledge" }` 6. Select **Execute**.

f. Run the **GET /upgrades/{upgradeId}** API multiple times, as needed, until the process is complete.

During the upgrade, the `status` indicates `running` if no errors are encountered. As each node is upgraded, the `step` value changes to `NodeFinished`.

The upgrade has finished successfully when the `percent` value is `100` and the `state` indicates `finished`.

What happens if an upgrade fails using NetApp Hybrid Cloud Control

If a drive or node fails during an upgrade, the Element UI will show cluster faults. The upgrade process does not proceed to the next node, and waits for the cluster faults to resolve. The progress bar in the UI shows that the upgrade is waiting for the cluster faults to resolve. At this stage, selecting **Pause** in the UI will not work, because the upgrade waits for the cluster to be healthy. You will need to engage NetApp Support to assist with the failure investigation.

NetApp Hybrid Cloud Control has a pre-set three-hour waiting period, during which one of the following scenarios can happen:

- The cluster faults get resolved within the three-hour window, and upgrade resumes. You do not need to take any action in this scenario.
- The problem persists after three hours, and the upgrade status shows **Error** with a red banner. You can resume the upgrade by selecting **Resume** after the problem is resolved.
- NetApp Support has determined that the upgrade needs to be temporarily aborted to take corrective action before the three-hour window. Support will use the API to abort the upgrade.

Aborting the cluster upgrade while a node is being updated might result in the drives being ungracefully removed from the node. If the drives are ungracefully removed, adding the drives back during an upgrade will require manual intervention by NetApp Support. The node might be taking longer to do firmware updates or post update syncing activities. If the upgrade progress seems stalled, contact NetApp Support for assistance.

Upgrade Element software at connected sites using HealthTools

Steps

1. Download the storage upgrade package; go to the Element software downloads page and download the latest storage node image to a device that not the management node.

   You need the latest version of HealthTools to upgrade Element storage software.

2. Copy the ISO file to the management node in an accessible location like /tmp.

   When you upload the ISO file, make sure that the name of the file does not change, otherwise later steps will fail.

3. **Optional**: Download the ISO from the management node to the cluster nodes before the upgrade.

   This step reduces the upgrade time by pre-staging the ISO on the storage nodes and running additional internal checks to ensure that the cluster is in a good state to be upgraded. Performing this operation will not put the cluster into "upgrade" mode or restrict any of the cluster operations.

   ```
   sfinstall <MVIP> -u <cluster_username> <path-toinstall-file-ISO> --stage
   ```

   Omit the password from the command line to allow `sfinstall` to prompt for the information. For passwords that contain special characters, add a backslash (`\`) before each special character. For example, `mypass!@1` should be entered as `mypass\!\@`.

   Example
   See the following sample input:

   ```
   sfinstall 10.117.0.244 -u admin /tmp/solidfire-rtfisodium-11.0.0.345.iso
   --stage
   ```

   The output for the sample shows that `sfinstall` attempts to verify if a newer version of `sfinstall` is available:

```
sfinstall 10.117.0.244 -u admin
/tmp/solidfire-rtfisodium-11.0.0.345.iso 2018-10-01 16:52:15:
Newer version of sfinstall available.
This version: 2018.09.01.130, latest version: 2018.06.05.901.
The latest version of the HealthTools can be downloaded from:
https:// mysupport.netapp.com/NOW/cgi-bin/software/
or rerun with --skip-version-check
```

See the following sample excerpt from a successful pre-stage operation:

When staging completes, the message will display `Storage Node Upgrade Staging Successful` after the upgrade event.

```
flabv0004 ~ # sfinstall -u admin
10.117.0.87 solidfire-rtfi-sodium-patch3-11.3.0.14171.iso --stage
2019-04-03 13:19:58: sfinstall Release Version: 2019.01.01.49 Management
Node Platform:
Ember Revision: 26b042c3e15a Build date: 2019-03-12 18:45
2019-04-03 13:19:58: Checking connectivity to MVIP 10.117.0.87
2019-04-03 13:19:58: Checking connectivity to node 10.117.0.86
2019-04-03 13:19:58: Checking connectivity to node 10.117.0.87
...
2019-04-03 13:19:58: Successfully connected to cluster and all nodes
...
2019-04-03 13:20:00: Do you want to continue? ['Yes', 'No']: Yes
...
2019-04-03 13:20:55: Staging install pack on cluster nodes
2019-04-03 13:20:55: newVersion: 11.3.0.14171
2019-04-03 13:21:01: nodeToStage: nlabp2814, nlabp2815, nlabp2816,
nlabp2813
2019-04-03 13:21:02: Staging Node nlabp2815 mip=[10.117.0.87] nodeID=[2]
(1 of 4 nodes)
2019-04-03 13:21:02: Node Upgrade serving image at
http://10.117.0.204/rtfi/solidfire-rtfisodium-
patch3-11.3.0.14171/filesystem.squashfs
...
2019-04-03 13:25:40: Staging finished. Repeat the upgrade command
without the --stage option to start the upgrade.
```

The staged ISOs will be automatically deleted after the upgrade completes. However, if the upgrade has not started and needs to be rescheduled, ISOs can be manually de-staged using the command:

```
sfinstall <MVIP> -u <cluster_username> --destage
```

After the upgrade has started, the de-stage option is no longer available.

4. Start the upgrade with the `sfinstall` command and the path to the ISO file:

   `sfinstall <MVIP> -u <cluster_username> <path-toinstall-file-ISO>`

   Example

   See the following sample input command:

   ```
   sfinstall 10.117.0.244 -u admin /tmp/solidfire-rtfi-sodium-
   11.0.0.345.iso
   ```

   The output for the sample shows that `sfinstall` attempts to verify if a newer version of `sfinstall` is available:

   ```
   sfinstall 10.117.0.244 -u admin /tmp/solidfire-rtfi-sodium-
   11.0.0.345.iso
   2018-10-01 16:52:15: Newer version of sfinstall available.
   This version: 2018.09.01.130, latest version: 2018.06.05.901.
   The latest version of the HealthTools can be downloaded from:
   https://mysupport.netapp.com/NOW/cgi-bin/software/ or rerun with --skip
   -version-check
   ```

   See the following sample excerpt from a successful upgrade. Upgrade events can be used to monitor the progress of the upgrade.

   ```
   # sfinstall 10.117.0.161 -u admin solidfire-rtfi-sodium-11.0.0.761.iso
   2018-10-11 18:28
   Checking connectivity to MVIP 10.117.0.161
   Checking connectivity to node 10.117.0.23
   Checking connectivity to node 10.117.0.24
   ...
   Successfully connected to cluster and all nodes
   ################################################################
   You are about to start a new upgrade
   10.117.0.161
   10.3.0.161
   solidfire-rtfi-sodium-11.0.0.761.iso
   Nodes:
   10.117.0.23 nlabp1023 SF3010 10.3.0.161
   10.117.0.24 nlabp1025 SF3010 10.3.0.161
   10.117.0.26 nlabp1027 SF3010 10.3.0.161
   10.117.0.28 nlabp1028 SF3010 10.3.0.161
   ################################################################
   Do you want to continue? ['Yes', 'No']: yes
   ...
   ```

```
Watching for new network faults. Existing fault IDs are set([]).
Checking for legacy network interface names that need renaming
Upgrading from 10.3.0.161 to 11.0.0.761 upgrade method=rtfi
Waiting 300 seconds for cluster faults to clear
Waiting for caches to fall below threshold
...
Installing mip=[10.117.0.23] nodeID=[1] (1 of 4 nodes)
Starting to move primaries.
Loading volume list
Moving primary slice=[7] away from mip[10.117.0.23] nodeID[1] ssid[11]
to new ssid[15]
Moving primary slice=[12] away from mip[10.117.0.23] nodeID[1] ssid[11]
to new ssid[15]
...
Installing mip=[10.117.114.24] nodeID=[2] (2 of 4 nodes)
Starting to move primaries.
Loading volume list
Moving primary slice=[5] away from mip[10.117.114.24] nodeID[2] ssid[7]
to new ssid[11]
...
Install of solidfire-rtfi-sodium-11.0.0.761 complete.
Removing old software
No staged builds present on nodeID=[1]
No staged builds present on nodeID=[2]
...
Starting light cluster block service check
```

If you are upgrading an H610S series node to Element 12.3.x and the node is running a version of Element earlier than 11.8, you will need to perform additional upgrade steps (phase 2) for each storage node. If you are running Element 11.8 or later, the additional upgrade steps (phase 2) are not required.

Upgrade Element software at dark sites using HealthTools

You can use the HealthTools suite of tools to update NetApp Element software at a dark site that has no external connectivity.

What you'll need

1. Go to the Element software downloads page.
2. Select the correct software release and download the latest storage node image to a computer that is not the management node.

   You need the latest version of HealthTools to upgrade Element storage software.

3. Download this JSON file (https://library.netapp.com/ecm/ecm_get_file/ECMLP2840740) from the NetApp Support Site on a computer that is not the management node and rename it to `metadata.json`.

4. Copy the ISO file to the management node in an accessible location like `/tmp`.

You can do this by using, for example, SCP. When you upload the ISO file, make sure that the name of the file does not change, otherwise later steps will fail.

Steps

1. Run the `sfupdate-healthtools` command:

```
sfupdate-healthtools <path-to-healthtools-package>
```

2. Check the installed version:

```
sfupdate-healthtools -v
```

3. Check the latest version against the metadata JSON file:

```
sfupdate-healthtools -l --metadata=<path-to-metadata-json>
```

4. Ensure that the cluster is ready:

```
sudo sfupgradecheck -u <cluster_username> -p <cluster_password> MVIP
--metadata=<path-to-metadata-json>
```

5. Run the `sfinstall` command with the path to the ISO file and the metadata JSON file:

```
sfinstall -u <cluster_username> <MVIP> <path-toinstall-file-ISO>
--metadata=<path-to-metadata-json-file>
```

See the following sample input command:

```
sfinstall -u admin 10.117.78.244 /tmp/solidfire-rtfi-11.3.0.345.iso
--metadata=/tmp/metadata.json
```

Optional You can add the `--stage` flag to the `sfinstall` command to pre-stage the upgrade in advance.

If you are upgrading an H610S series node to Element 12.3.x and the node is running a version of Element earlier than 11.8, you will need to perform additional upgrade steps (phase 2) for each storage node. If you are running Element 11.8 or later, the additional upgrade steps (phase 2) are not required.

What happens if an upgrade fails using HealthTools

If the software upgrade fails, you can pause the upgrade.

You should pause an upgrade only with Ctrl-C. This enables the system to clean itself up.

When `sfinstall` waits for cluster faults to clear and if any failure causes the faults to remain, `sfinstall` will not proceed to the next node.

Steps

1. You should stop `sfinstall` with Ctrl+C.
2. Contact NetApp Support to assist with the failure investigation.
3. Resume the upgrade with the same `sfinstall` command.
4. When an upgrade is paused by using Ctrl+C, if the upgrade is currently upgrading a node, choose one of these options:
   - **Wait**: Allow the currently upgrading node to finish before resetting the cluster constants.
   - **Continue**: Continue the upgrade, which cancels the pause.
   - **Abort**: Reset the cluster constants and abort the upgrade immediately.

Aborting the cluster upgrade while a node is being updated might result in the drives being ungracefully removed from the node. If the drives are ungracefully removed, adding the drives back during an upgrade will require manual intervention by NetApp Support. The node might be taking longer to do firmware updates or post update syncing activities. If the upgrade progress seems stalled, contact NetApp Support for assistance.

Upgrading H610S storage nodes to Element 12.3.x (phase 2)

If you are upgrading an H610S series node to Element 12.3.x and the node is running a version of Element earlier than 11.8, the upgrade process involves two phases.

Phase 1, which is performed first, follows the same steps as the standard upgrade to Element 12.3.x process. It installs Element Software and all 5 firmware updates in a rolling fashion across the cluster one node at a time. Due to the firmware payload, the process is estimated to take approximately 1.5 to 2 hours per H610S node, including a single cold-boot cycle at the end of the upgrade for each node.

Phase 2 involves completing steps to perform a complete node shutdown and power disconnect for each H610S node that are described in a required KB. This phase is estimated to take approximately one hour per H610S node.

After you complete phase 1, four of the five firmware updates are activated during the cold boot on each H610S node; however, the Complex Programmable Logic Device (CPLD) firmware requires a complete power disconnect and reconnect to fully install. The CPLD firmware update protects against NVDIMM failures and metadata drive eviction during future reboots or power cycles. This power reset is estimated to take approximately one hour per H610S node. It requires shutting down the node, removing power cables or disconnecting power via a smart PDU, waiting approximately 3 minutes, and reconnecting power.

Before you begin

- You have completed phase 1 of the H610S upgrade process and have upgraded your storage nodes using

one the standard Element storage upgrade procedures.

Phase 2 requires on-site personnel.

Steps

1. (Phase 2) Complete the power reset process required for each H610S node in the cluster:

If the cluster also has non-H610S nodes, these non-H610S nodes are exempt from phase 2 and do not need to be shut down or have their power disconnected.

a. Contact NetApp Support for assistance and to schedule this upgrade.
b. Follow the phase 2 upgrade procedure in this KB that is required to complete an upgrade for each H610S node.

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Upgrade storage firmware

Starting with Element 12.0 and management services version 2.14, you can perform firmware-only upgrades on your storage nodes using the NetApp Hybrid Cloud Control UI and REST API. This procedure does not upgrade Element software and enables you to upgrade storage firmware outside of a major Element release.

What you'll need

- **Admin privileges**: You have storage cluster administrator permissions to perform the upgrade.
- **System time sync**: You have ensured that the system time on all nodes is synced and that NTP is correctly configured for the storage cluster and nodes. Each node must be configured with a DNS nameserver in the per-node web UI (`https://[IP address]:442`) with no unresolved cluster faults related to time skew.
- **System ports**: If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See Network ports for more information.
- **Management node**: For NetApp Hybrid Cloud Control UI and API, the management node in your environment is running version 11.3.
- **Management services**: You have updated your management services bundle to the latest version.

For H610S storage nodes running Element software version 12.0, you should apply D-patch SUST-909 before you upgrade to storage firmware bundle 2.27. Contact NetApp Support to obtain the D-patch before you upgrade. See Storage Firmware Bundle 2.27 Release Notes.

You must upgrade to the latest management services bundle before upgrading the firmware on your storage nodes. If you are updating your Element software to version 12.2 or later, you need management services 2.14.60 or later to proceed.

- **Cluster health**: You have run health checks. See Run Element storage health checks prior to upgrading storage.

- **End User License Agreement (EULA)**: Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to upgrade storage firmware:
  1. Open the IP address of the management node in a web browser:

     ```
     https://<ManagementNodeIP>
     ```

  2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
  3. Select **Upgrade** near the top right of the interface.
  4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.
- **Updated BMC for H610S nodes**: You have upgraded the BMC version for your H610S nodes. See release notes and upgrade instructions.

For a complete matrix of firmware and driver firmware for your hardware, see supported storage firmware versions for SolidFire storage nodes

Upgrade options

Choose one of the following storage firmware upgrade options:

- Use NetApp Hybrid Cloud Control UI to upgrade storage firmware
- Use NetApp Hybrid Cloud Control API to upgrade storage firmware

Use NetApp Hybrid Cloud Control UI to upgrade storage firmware

You can use the NetApp Hybrid Cloud Control UI to upgrade the firmware of the storage nodes in your cluster.

What you'll need

- If your management node is not connected to the internet, you have downloaded the storage firmware bundle.

For potential issues while upgrading storage clusters using NetApp Hybrid Cloud Control and their workarounds, see the KB article.

The upgrade process takes approximately 30 minutes per storage node. If you are upgrading an Element storage cluster to storage firmware newer than version 2.76, individual storage nodes will only reboot during the upgrade if new firmware was written to the node.

Steps

1. Open the IP address of the management node in a web browser:

   ```
   https://<ManagementNodeIP>
   ```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select **Storage**.

The **Storage** tab lists the storage clusters that are part of your installation. If a cluster is inaccessible by NetApp Hybrid Cloud Control, it will not be displayed on the **Upgrades** page. If you have clusters running Element 12.0 or later, you will see the current firmware bundle version listed for these clusters. If the nodes in a single cluster have different firmware versions on them or as the upgrade progresses, you will see **Multiple** in the **Current Firmware Bundle Version** column. You can select **Multiple** to navigate to the **Nodes** page to compare firmware versions. If all your clusters are running Element versions earlier than 12.0, you will not see any information about firmware bundle version numbers.

If the cluster is up to date and/or no upgrade packages are available, the **Element** and **Firmware Only** tabs are not displayed. These tabs are also not displayed when an upgrade is in progress. If the **Element** tab is displayed, but not the **Firmware Only** tab, no firmware packages are available.

5. Select the drop-down arrow next to the cluster you are upgrading.
6. Select **Browse** to upload the upgrade package that you downloaded.
7. Wait for the upload to complete. A progress bar shows the status of the upload.

   The file upload will be lost if you navigate away from the browser window.

   An on-screen message is displayed after the file is successfully uploaded and validated. Validation might take several minutes. If you navigate away from the browser window at this stage, the file upload is preserved.

8. Select **Firmware Only**, and select from the upgrade versions available.
9. Select **Begin Upgrade**.

The **Upgrade Status** changes during the upgrade to reflect the status of the process. It also changes in response to actions you take, such as pausing the upgrade, or if the upgrade returns an error. See Upgrade status changes.

While the upgrade is in progress, you can leave the page and come back to it later to continue monitoring the progress. The page does not dynamically update status and current version if the cluster row is collapsed. The cluster row must be expanded to update the table or you can refresh the page.

You can download logs after the upgrade is complete.

Upgrade status changes

Here are the different states that the **Upgrade Status** column in the UI shows before, during, and after the upgrade process:

Upgrade state	Description
Up to Date	The cluster was upgraded to the latest Element version available or the firmware was upgraded to the latest version.

Upgrade state	Description
Unable to Detect	This status is displayed when the storage service API returns an upgrade status that is not in the enumerated list of possible upgrade statuses.
Versions Available	Newer versions of Element and/or storage firmware are available for upgrade.
In Progress	The upgrade is in progress. A progress bar shows the upgrade status. On-screen messages also show node-level faults and display the node ID of each node in the cluster as the upgrade progresses. You can monitor the status of each node using the Element UI or the NetApp Element plug-in for vCenter Server UI.
Upgrade Pausing	You can choose to pause the upgrade. Depending on the state of the upgrade process, the pause operation can succeed or fail. You will see a UI prompt asking you to confirm the pause operation. To ensure that the cluster is in a safe spot before pausing an upgrade, it can take up to two hours for the upgrade operation to be completely paused. To resume the upgrade, select **Resume**.
Paused	You paused the upgrade. Select **Resume** to resume the process.
Error	An error has occurred during the upgrade. You can download the error log and send it to NetApp Support. After you resolve the error, you can return to the page, and select **Resume**. When you resume the upgrade, the progress bar goes backwards for a few minutes while the system runs the health check and checks the current state of the upgrade.

What happens if an upgrade fails using NetApp Hybrid Cloud Control

If a drive or node fails during an upgrade, the Element UI will show cluster faults. The upgrade process does not proceed to the next node, and waits for the cluster faults to resolve. The progress bar in the UI shows that the upgrade is waiting for the cluster faults to resolve. At this stage, selecting **Pause** in the UI will not work, because the upgrade waits for the cluster to be healthy. You will need to engage NetApp Support to assist with the failure investigation.

NetApp Hybrid Cloud Control has a pre-set three-hour waiting period, during which one of the following scenarios can happen:

- The cluster faults get resolved within the three-hour window, and upgrade resumes. You do not need to take any action in this scenario.
- The problem persists after three hours, and the upgrade status shows **Error** with a red banner. You can resume the upgrade by selecting **Resume** after the problem is resolved.
- NetApp Support has determined that the upgrade needs to be temporarily aborted to take corrective action before the three-hour window. Support will use the API to abort the upgrade.

Aborting the cluster upgrade while a node is being updated might result in the drives being ungracefully removed from the node. If the drives are ungracefully removed, adding the drives back during an upgrade will require manual intervention by NetApp Support. The node might be taking longer to do firmware updates or post update syncing activities. If the upgrade progress seems stalled, contact NetApp Support for assistance.

Use NetApp Hybrid Cloud Control API to upgrade storage firmware

You can use APIs to upgrade storage nodes in a cluster to the latest Element software version. You can use an automation tool of your choice to run the APIs. The API workflow documented here uses the REST API UI available on the management node as an example.

Steps

1. Download the storage firmware upgrade package to a device that is accessible to the management node; go to the Element software downloads page and download the latest storage firmware image.
2. Upload the storage firmware upgrade package to the management node:
   a. Open the management node REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/package-repository/1/
      ```

   b. Select **Authorize** and complete the following:
      i. Enter the cluster user name and password.
      ii. Enter the client ID as `mnode-client`.
      iii. Select **Authorize** to begin a session.
      iv. Close the authorization window.
   c. From the REST API UI, select **POST /packages**.
   d. Select **Try it out**.
   e. Select **Browse** and select the upgrade package.
   f. Select **Execute** to initiate the upload.
   g. From the response, copy and save the package ID (`"id"`) for use in a later step.
3. Verify the status of the upload.
   a. From the REST API UI, select **GET /packages/{id}/status**.
   b. Select **Try it out**.
   c. Enter the firmware package ID you copied in the previous step in **id**.
   d. Select **Execute** to initiate the status request.

      The response indicates `state` as `SUCCESS` when complete.
4. Locate the installation asset ID:
   a. Open the management node REST API UI on the management node:

      ```
      https://<ManagementNodeIP>/inventory/1/
      ```

b. Select **Authorize** and complete the following:

   i. Enter the cluster user name and password.

   ii. Enter the client ID as `mnode-client`.

   iii. Select **Authorize** to begin a session.

   iv. Close the authorization window.

c. From the REST API UI, select **GET /installations**.

d. Select **Try it out**.

e. Select **Execute**.

f. From the response, copy the installation asset ID (`id`).

```
"id": "abcd01e2-xx00-4ccf-11ee-11f111xx9a0b",
"management": {
  "errors": [],
  "inventory": {
    "authoritativeClusterMvip": "10.111.111.111",
    "bundleVersion": "2.14.19",
    "managementIp": "10.111.111.111",
    "version": "1.4.12"
```

g. From the REST API UI, select **GET /installations/{id}**.

h. Select **Try it out**.

i. Paste the installation asset ID into the **id** field.

j. Select **Execute**.

k. From the response, copy and save the storage cluster ID (`"id"`) of the cluster you intend to upgrade for use in a later step.

```
"storage": {
  "errors": [],
  "inventory": {
    "clusters": [
      {
        "clusterUuid": "a1bd1111-4f1e-46zz-ab6f-0a1111b1111x",
        "id": "a1bd1111-4f1e-46zz-ab6f-a1a1a111b012",
```

5. Run the storage firmware upgrade:

   a. Open the storage REST API UI on the management node:

```
https://<ManagementNodeIP>/storage/1/
```

   b. Select **Authorize** and complete the following:

i. Enter the cluster user name and password.

ii. Enter the client ID as `mnode-client`.

iii. Select **Authorize** to begin a session.

iv. Close the window.

c. Select **POST /upgrades**.

d. Select **Try it out**.

e. Enter the upgrade package ID in the parameter field.

f. Enter the storage cluster ID in the parameter field.

g. Select **Execute** to initiate the upgrade.

The response should indicate state as `initializing`:

```
{
  "_links": {
    "collection": "https://localhost:442/storage/upgrades",
    "self": "https://localhost:442/storage/upgrades/3fa85f64-1111-
4562-b3fc-2c963f66abc1",
    "log": https://localhost:442/storage/upgrades/3fa85f64-1111-4562-
b3fc-2c963f66abc1/log
  },
  "storageId": "114f14a4-1a1a-11e9-9088-6c0b84e200b4",
  "upgradeId": "334f14a4-1a1a-11e9-1055-6c0b84e2001b4",
  "packageId": "774f14a4-1a1a-11e9-8888-6c0b84e200b4",
  "config": {},
  "state": "initializing",
  "status": {
    "availableActions": [
      "string"
    ],
    "message": "string",
    "nodeDetails": [
      {
        "message": "string",
        "step": "NodePreStart",
        "nodeID": 0,
        "numAttempt": 0
      }
    ],
    "percent": 0,
    "step": "ClusterPreStart",
    "timestamp": "2020-04-21T22:10:57.057Z",
    "failedHealthChecks": [
      {
        "checkID": 0,
```

```
        "name": "string",
        "displayName": "string",
        "passed": true,
        "kb": "string",
        "description": "string",
        "remedy": "string",
        "severity": "string",
        "data": {},
        "nodeID": 0
      }
    ]
  },
  "taskId": "123f14a4-1a1a-11e9-7777-6c0b84e123b2",
  "dateCompleted": "2020-04-21T22:10:57.057Z",
  "dateCreated": "2020-04-21T22:10:57.057Z"
}
```


h. Copy the upgrade ID (`"upgradeId"`) that is part of the response.

6. Verify the upgrade progress and results:
   a. Select **GET /upgrades/{upgradeId}**.
   b. Select **Try it out**.
   c. Enter the upgrade ID from the previous step in **upgradeId**.
   d. Select **Execute**.
   e. Do one of the following if there are problems or special requirements during the upgrade:

Option	Steps
You need to correct cluster health issues due to `failedHealthChecks` message in the response body.	1. Go to the specific KB article listed for each issue or perform the specified remedy. 2. If a KB is specified, complete the process described in the relevant KB article. 3. After you have resolved cluster issues, reauthenticate if needed and select **PUT /upgrades/{upgradeId}**. 4. Select **Try it out**. 5. Enter the upgrade ID from the previous step in **upgradeId**. 6. Enter `"action":"resume"` in the request body. `{ "action": "resume" }` 7. Select **Execute**.
You need to pause the upgrade because the maintenance window is closing or for another reason.	1. Reauthenticate if needed and select **PUT /upgrades/{upgradeId}**. 2. Select **Try it out**. 3. Enter the upgrade ID from the previous step in **upgradeId**. 4. Enter `"action":"pause"` in the request body. `{ "action": "pause" }` 5. Select **Execute**.

f. Run the **GET /upgrades/{upgradeId}** API multiple times, as needed, until the process is complete.

During the upgrade, the `status` indicates `running` if no errors are encountered. As each node is upgraded, the `step` value changes to `NodeFinished`.

The upgrade has finished successfully when the `percent` value is `100` and the `state` indicates `finished`.

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Upgrade a management node

You can upgrade your management node to management node version 12.3.x from version 11.0 or later.

Upgrading the management node operating system is no longer required to upgrade Element software on the storage cluster. If the management node is version 11.3 or higher, you can simply upgrade the management services to the latest version to perform Element upgrades using NetApp Hybrid Cloud Control. Follow the management node upgrade procedure for your scenario if you would like to upgrade the management node operating system for other reasons, such as security remediation.

The vCenter Plug-in 4.4 or later requires a management node 11.3 or later that is created with modular architecture and provides individual services.

Upgrade options

Choose one of the following management node upgrade options:

- Management node 12.3.2 contains a security mitigation for storage clusters with the Virtual Volumes (VVols) feature enabled. If your storage cluster is already at Element 12.3 and the VVols feature is enabled, you should upgrade to 12.3.2.
- There are no additional functionality changes or bug fixes in management node 12.3.1. If you are already running management node 12.3, you do not need to upgrade it to 12.3.1.

- If you are upgrading from management node 12.2:
  Upgrade a management node to version 12.3.x from 12.2
- If you are upgrading from management node 12.0:
  Upgrade a management node to version 12.3.x from 12.0
- If you are upgrading from management node 11.3, 11.5, 11.7, or 11.8:
  Upgrade a management node to version 12.3.x from 11.3 through 11.8
- If you are upgrading from management node 11.0 or 11.1:
  Upgrade a management node to version 12.3.x from 11.1 or 11.0
- If you are upgrading from a management node version 10.x:
  Migrating from management node version 10.x to 11.x

Choose the following option if you have **sequentially** updated (1) your management services version and (2) your Element storage version and you want to **keep** your existing management node:

If you do not sequentially update your management services followed by Element storage, you cannot reconfigure reauthentication using this procedure. Follow the appropriate upgrade procedure instead.

- If you are keeping existing management node:
  Reconfigure authentication using the management node REST API

Upgrade a management node to version 12.3.x from 12.2

You can perform an in-place upgrade of the management node from version 12.2 to version 12.3.x without needing to provision a new management node virtual machine.

> The Element 12.3.x management node is an optional upgrade. It is not required for existing deployments.

What you'll need

- The RAM of the management node VM is 24GB.
- The management node you are intending to upgrade is version 12.0 and uses IPv4 networking. The management node version 12.3.x does not support IPv6.

> To check the version of your management node, log in to your management node and view the Element version number in the login banner.

- You have updated your management services bundle to the latest version using NetApp Hybrid Cloud Control. You can access NetApp Hybrid Cloud Control from the following IP: `https://<ManagementNodeIP>`
- If you are updating your management node to version 12.3.x, you need management services 2.14.60 or later to proceed.
- You have configured an additional network adapter (if required) using the instructions for configuring an additional storage NIC.

> Persistent volumes might require an additional network adapter if eth0 is not able to be routed to the SVIP. Configure a new network adapter on the iSCSI storage network to allow the configuration of persistent volumes.

- Storage nodes are running Element 11.3 or later.

Steps

1. Log in to the management node virtual machine using SSH or console access.
2. Download the management node ISO for Element software from the NetApp Support Site to the management node virtual machine.

> The name of the ISO is similar to `solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso`

3. Check the integrity of the download by running md5sum on the downloaded file and compare the output to what is available on the NetApp Support Site for Element software, as in the following example:

   ```
   sudo md5sum -b <path to iso>/solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso
   ```

4. Mount the management node ISO image and copy the contents to the file system using the following commands:

   ```
   sudo mkdir -p /upgrade
   ```

```
sudo mount <solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso>
/mnt
```

```
sudo cp -r /mnt/* /upgrade
```

5. Change to the home directory, and unmount the ISO file from `/mnt`:

```
sudo umount /mnt
```

6. Delete the ISO to conserve space on the management node:

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-
XX.X.X.XXXX.iso
```

7. On the management node that you are upgrading, run the following command to upgrade your management node OS version. The script retains all necessary configuration files after the upgrade, such as Active IQ collector and proxy settings.

```
sudo /sf/rtfi/bin/sfrtfi_inplace
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
```

The management node reboots with a new OS after the upgrade process completes.

> After you run the sudo command described in this step, the SSH session is killed. Console access is required for continued monitoring. If no console access is available to you when performing the upgrade, retry the SSH login and verify connectivity after 15 to 30 minutes. Once you log in, you can confirm the new OS version in the SSH banner that indicates that the upgrade was successful.

8. On the management node, run the `redeploy-mnode` script to retain previous management services configuration settings:

> The script retains previous management services configuration, including configuration from the Active IQ collector service, controllers (vCenters), or proxy, depending on your settings.

```
sudo /sf/packages/mnode/redeploy-mnode -mu <mnode user>
```

> If you had previously disabled SSH functionality on the management node, you need to disable SSH again on the recovered management node. SSH capability that provides NetApp Support remote support tunnel (RST) session access is enabled on the management node by default.

Upgrade a management node to version 12.3.x from 12.0

You can perform an in-place upgrade of the management node from version 12.0 to version 12.3.x without needing to provision a new management node virtual machine.

> The Element 12.3.x management node is an optional upgrade. It is not required for existing deployments.

What you'll need

- The management node you are intending to upgrade is version 12.0 and uses IPv4 networking. The management node version 12.3.x does not support IPv6.

  > To check the version of your management node, log in to your management node and view the Element version number in the login banner.

- You have updated your management services bundle to the latest version using NetApp Hybrid Cloud Control. You can access NetApp Hybrid Cloud Control from the following IP: `https://<ManagementNodeIP>`
- If you are updating your management node to version 12.3.x, you need management services 2.14.60 or later to proceed.
- You have configured an additional network adapter (if required) using the instructions for configuring an additional storage NIC.

  > Persistent volumes might require an additional network adapter if eth0 is not able to be routed to the SVIP. Configure a new network adapter on the iSCSI storage network to allow the configuration of persistent volumes.

- Storage nodes are running Element 11.3 or later.

Steps

1. Configure the management node VM RAM:
   a. Power off the management node VM.
   b. Change the RAM of the management node VM from 12GB to 24GB RAM.
   c. Power on the management node VM.
2. Log in to the management node virtual machine using SSH or console access.
3. Download the management node ISO for Element software from the NetApp Support Site to the management node virtual machine.

   > The name of the ISO is similar to `solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso`

4. Check the integrity of the download by running md5sum on the downloaded file and compare the output to what is available on the NetApp Support Site for Element software, as in the following example:

   ```
   sudo md5sum -b <path to iso>/solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso
   ```

5. Mount the management node ISO image and copy the contents to the file system using the following commands:

```
sudo mkdir -p /upgrade
```

```
sudo mount <solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso>
/mnt
```

```
sudo cp -r /mnt/* /upgrade
```

6. Change to the home directory, and unmount the ISO file from `/mnt`:

```
sudo umount /mnt
```

7. Delete the ISO to conserve space on the management node:

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-
XX.X.X.XXXX.iso
```

8. On the management node that you are upgrading, run the following command to upgrade your management node OS version. The script retains all necessary configuration files after the upgrade, such as Active IQ collector and proxy settings.

```
sudo /sf/rtfi/bin/sfrtfi_inplace
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
```

The management node reboots with a new OS after the upgrade process completes.

> After you run the sudo command described in this step, the SSH session is killed. Console access is required for continued monitoring. If no console access is available to you when performing the upgrade, retry the SSH login and verify connectivity after 15 to 30 minutes. Once you log in, you can confirm the new OS version in the SSH banner that indicates that the upgrade was successful.

9. On the management node, run the `redeploy-mnode` script to retain previous management services configuration settings:

> The script retains previous management services configuration, including configuration from the Active IQ collector service, controllers (vCenters), or proxy, depending on your settings.

```
sudo /sf/packages/mnode/redeploy-mnode -mu <mnode user>
```

SSH capability that provides NetApp Support remote support tunnel (RST) session access is disabled by default on management nodes running management services 2.18 and later. If you had previously enabled SSH functionality on the management node, you might need to disable SSH again on the upgraded management node.

Upgrade a management node to version 12.3.x from 11.3 through 11.8

You can perform an in-place upgrade of the management node from version 11.3, 11.5, 11.7, or 11.8 to version 12.3.x without needing to provision a new management node virtual machine.

The Element 12.3.x management node is an optional upgrade. It is not required for existing deployments.

What you'll need

- The management node you are intending to upgrade is version 11.3, 11.5, 11.7, or 11.8 and uses IPv4 networking. The management node version 12.3.x does not support IPv6.

  To check the version of your management node, log in to your management node and view the Element version number in the login banner.

- You have updated your management services bundle to the latest version using NetApp Hybrid Cloud Control. You can access NetApp Hybrid Cloud Control from the following IP: `https://<ManagementNodeIP>`
- If you are updating your management node to version 12.3.x, you need management services 2.14.60 or later to proceed.
- You have configured an additional network adapter (if required) using the instructions for configuring an additional storage NIC.

  Persistent volumes might require an additional network adapter if eth0 is not able to be routed to the SVIP. Configure a new network adapter on the iSCSI storage network to allow the configuration of persistent volumes.

- Storage nodes are running Element 11.3 or later.

Steps

1. Configure the management node VM RAM:
   a. Power off the management node VM.
   b. Change the RAM of the management node VM from 12GB to 24GB RAM.
   c. Power on the management node VM.
2. Log in to the management node virtual machine using SSH or console access.
3. Download the management node ISO for Element software from the NetApp Support Site to the management node virtual machine.

   The name of the ISO is similar to `solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso`

4. Check the integrity of the download by running md5sum on the downloaded file and compare the output to what is available on NetApp Support Site for Element software, as in the following example:

```
sudo md5sum -b <path to iso>/solidfire-fdva-<Element release>-patchX-
XX.X.X.XXXX.iso
```

5. Mount the management node ISO image and copy the contents to the file system using the following commands:

```
sudo mkdir -p /upgrade
```

```
sudo mount <solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso>
/mnt
```

```
sudo cp -r /mnt/* /upgrade
```

6. Change to the home directory, and unmount the ISO file from `/mnt`:

```
sudo umount /mnt
```

7. Delete the ISO to conserve space on the management node:

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-
XX.X.X.XXXX.iso
```

8. On the 11.3, 11.5, 11.7, or 11.8 management node, run the following command to upgrade your management node OS version. The script retains all necessary configuration files after the upgrade, such as Active IQ collector and proxy settings.

```
sudo /sf/rtfi/bin/sfrtfi_inplace
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
```

   The management node reboots with a new OS after the upgrade process completes.

   After you run the sudo command described in this step, the SSH session is killed. Console access is required for continued monitoring. If no console access is available to you when performing the upgrade, retry the SSH login and verify connectivity after 15 to 30 minutes. Once you log in, you can confirm the new OS version in the SSH banner that indicates that the upgrade was successful.

9. On the management node, run the `redeploy-mnode` script to retain previous management services configuration settings:

   The script retains previous management services configuration, including configuration from the Active IQ collector service, controllers (vCenters), or proxy, depending on your settings.

```
sudo /sf/packages/mnode/redeploy-mnode -mu <mnode user>
```

SSH capability that provides NetApp Support remote support tunnel (RST) session access is disabled by default on management nodes running management services 2.18 and later. If you had previously enabled SSH functionality on the management node, you might need to disable SSH again on the upgraded management node.

Upgrade a management node to version 12.3.x from 11.1 or 11.0

You can perform an in-place upgrade of the management node from 11.0 or 11.1 to version 12.3.x without needing to provision a new management node virtual machine.

What you'll need

- Storage nodes are running Element 11.3 or later.

  Use the latest HealthTools to upgrade Element software.

- The management node you are intending to upgrade is version 11.0 or 11.1 and uses IPv4 networking. The management node version 12.3.x does not support IPv6.

  To check the version of your management node, log in to your management node and view the Element version number in the login banner.

- For management node 11.0, the VM memory needs to be manually increased to 12GB.
- You have configured an additional network adapter (if required) using the instructions for configuring a storage NIC (eth1) in the management node user guide your product.

  Persistent volumes might require an additional network adapter if eth0 is not able to be routed to the SVIP. Configure a new network adapter on the iSCSI storage network to allow the configuration of persistent volumes.

Steps

1. Configure the management node VM RAM:
   a. Power off the management node VM.
   b. Change the RAM of the management node VM from 12GB to 24GB RAM.
   c. Power on the management node VM.
2. Log in to the management node virtual machine using SSH or console access.
3. Download the management node ISO for Element software from the NetApp Support Site to the management node virtual machine.

   The name of the ISO is similar to `solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso`

4. Check the integrity of the download by running md5sum on the downloaded file and compare the output to what is available on NetApp Support Site for Element software, as in the following example:

```
sudo md5sum -b <path to iso>/solidfire-fdva-<Element release>-patchX-
XX.X.X.XXXX.iso
```

5. Mount the management node ISO image and copy the contents to the file system using the following commands:

```
sudo mkdir -p /upgrade
```

```
sudo mount solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso /mnt
```

```
sudo cp -r /mnt/* /upgrade
```

6. Change to the home directory, and unmount the ISO file from /mnt:

```
sudo umount /mnt
```

7. Delete the ISO to conserve space on the management node:

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-
XX.X.X.XXXX.iso
```

8. Run one of the following scripts with options to upgrade your management node OS version. Only run the script that is appropriate for your version. Each script retains all necessary configuration files after the upgrade, such as Active IQ collector and proxy settings.
   a. On an 11.1 (11.1.0.73) management node, run the following command:

```
sudo /sf/rtfi/bin/sfrtfi_inplace
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
sf_keep_paths="/sf/packages/solidfire-sioc-4.2.3.2288
/sf/packages/solidfire-nma-1.4.10/conf /sf/packages/sioc
/sf/packages/nma"
```

   b. On an 11.1 (11.1.0.72) management node, run the following command:

```
sudo /sf/rtfi/bin/sfrtfi_inplace
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
sf_keep_paths="/sf/packages/solidfire-sioc-4.2.1.2281
/sf/packages/solidfire-nma-1.4.10/conf /sf/packages/sioc
/sf/packages/nma"
```

c. On an 11.0 (11.0.0.781) management node, run the following command:

```
sudo /sf/rtfi/bin/sfrtfi_inplace
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
sf_keep_paths="/sf/packages/solidfire-sioc-4.2.0.2253
/sf/packages/solidfire-nma-1.4.8/conf /sf/packages/sioc
/sf/packages/nma"
```

The management node reboots with a new OS after the upgrade process completes.

> ⓘ After you run the sudo command described in this step, the SSH session is killed. Console access is required for continued monitoring. If no console access is available to you when performing the upgrade, retry the SSH login and verify connectivity after 15 to 30 minutes. Once you log in, you can confirm the new OS version in the SSH banner that indicates that the upgrade was successful.

9. On the 12.3.x management node, run the `upgrade-mnode` script to retain previous configuration settings.

> ⓘ If you are migrating from an 11.0 or 11.1 management node, the script copies the Active IQ collector to the new configuration format.

a. For a single storage cluster managed by an existing management node 11.0 or 11.1 with persistent volumes:

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user> -pv <true -
persistent volume> -pva <persistent volume account name - storage
volume account>
```

b. For a single storage cluster managed by an existing management node 11.0 or 11.1 with no persistent volumes:

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user>
```

c. For multiple storage clusters managed by an existing management node 11.0 or 11.1 with persistent volumes:

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user> -pv <true -
persistent volume> -pva <persistent volume account name - storage
volume account> -pvm <persistent volumes mvip>
```

d. For multiple storage clusters managed by an existing management node 11.0 or 11.1 with no persistent volumes (the `-pvm` flag is to provide one of the cluster's MVIP addresses):

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user> -pvm <mvip for
persistent volumes>
```

10. (For all NetApp SolidFire all-flash storage installations with NetApp Element Plug-in for vCenter Server) Update the vCenter Plug-in on the 12.3.x management node by following the steps in the Upgrade the Element Plug-in for vCenter Server topic.
11. Locate the asset ID for your installation using the management node API:
    a. From a browser, log into the management node REST API UI:
        i. Go to the storage MVIP and log in.
        This action causes certificate to be accepted for the next step.
    b. Open the inventory service REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```

    c. Select **Authorize** and complete the following:
        i. Enter the cluster user name and password.
        ii. Enter the client ID as `mnode-client`.
        iii. Select **Authorize** to begin a session.
        iv. Close the window.
    d. From the REST API UI, select **GET /installations**.
    e. Select **Try it out**.
    f. Select **Execute**.
    g. From the code 200 response body, copy the `id` for the installation.

    Your installation has a base asset configuration that was created during installation or upgrade.

Migrating from management node version 10.x to 11.x

If you have a management node at version 10.x, you cannot upgrade from 10.x to 11.x. You can instead use this migration procedure to copy over the configuration from 10.x to a newly deployed 11.1 management node. If your management node is currently at 11.0 or higher, you should skip this procedure. You need management node 11.0 or 11.1 and the latest HealthTools to upgrade Element software from 10.3 + through 11.x.

Steps

1. From the VMware vSphere interface, deploy the management node 11.1 OVA and power it on.

2. Open the management node VM console, which brings up the terminal user interface (TUI).
3. Use the TUI to create a new administrator ID and assign a password.
4. In the management node TUI, log in to the management node with the new ID and password and validate that it works.
5. From the vCenter or management node TUI, get the management node 11.1 IP address and browse to the IP address on port 9443 to open the management node UI.

   ```
   https://<mNode 11.1 IP address>:9443
   ```

6. In vSphere, select **NetApp Element Configuration** > **mNode Settings**. (In older versions, the top-level menu is **NetApp SolidFire Configuration**.)
7. Select **Actions** > **Clear**.
8. To confirm, select **Yes**. The mNode Status field should report Not Configured.

   When you go to the **mNode Settings** tab for the first time, the mNode Status field might display as **Not Configured** instead of the expected **UP**; you might not be able to choose **Actions** > **Clear**. Refresh the browser. The mNode Status field will eventually display **UP**.

9. Log out of vSphere.
10. In a web browser, open the management node registration utility and select **QoSSIOC Service Management**:

    ```
    https://<mNode 11.1 IP address>:9443
    ```

11. Set the new QoSSIOC password.

    The default password is `solidfire`. This password is required to set the new password.

12. Select the **vCenter Plug-in Registration** tab.
13. Select **Update Plug-in**.
14. Enter required values. When you are finished, select **UPDATE**.
15. Log in to vSphere and select **NetApp Element Configuration** > **mNode Settings**.
16. Select **Actions** > **Configure**.
17. Provide the management node IP address, management node user ID (the user name is `admin`), password that you set on the **QoSSIOC Service Management** tab of the registration utility, and vCenter user ID and password.

    In vSphere, the **mNode Settings** tab should display the mNode status as **UP**, which indicates management node 11.1 is registered to vCenter.

18. From the management node registration utility (`https://<mNode 11.1 IP address>:9443`), restart the SIOC service from **QoSSIOC Service Management**.
19. Wait for one minute and check the **NetApp Element Configuration** > **mNode Settings** tab. This should display the mNode status as **UP**.

If the status is **DOWN**, check the permissions for `/sf/packages/sioc/app.properties`. The file should have read, write, and execute permissions for the file owner. The correct permissions should appear as follows:

```
-rwx------
```

20. After the SIOC process starts and vCenter displays mNode status as **UP**, check the logs for the `sf-hci-nma` service on the management node. There should be no error messages.
21. (For management node 11.1 only) SSH into the management node version 11.1 with root privileges and start the NMA service with the following commands:

```
# systemctl enable /sf/packages/nma/systemd/sf-hci-nma.service
```

```
# systemctl start sf-hci-nma21
```

22. Perform actions from vCenter to remove a drive, add a drive or reboot nodes. This triggers storage alerts, which should be reported in vCenter. If this is working, NMA system alerts are functioning as expected.
23. If ONTAP Select is configured in vCenter, configure ONTAP Select alerts in NMA by copying the `.ots.properties` file from the previous management node to the management node version 11.1 `/sf/packages/nma/conf/.ots.properties` file, and restart the NMA service using the following command:

```
systemctl restart sf-hci-nma
```

24. Verify that ONTAP Select is working by viewing the logs with the following command:

```
journalctl -f | grep -i ots
```

25. Configure Active IQ by doing the following:
    a. SSH in to the management node version 11.1 and go to the `/sf/packages/collector` directory.
    b. Run the following command:

```
sudo ./manage-collector.py --set-username netapp --set-password --set
-mvip <MVIP>
```

    c. Enter the management node UI password when prompted.
    d. Run the following commands:

```
./manage-collector.py --get-all
```

```
sudo systemctl restart sfcollector
```

e. Verify `sfcollector` logs to confirm it is working.

26. In vSphere, the **NetApp Element Configuration** > **mNode Settings** tab should display the mNode status as **UP**.
27. Verify NMA is reporting system alerts and ONTAP Select alerts.
28. If everything is working as expected, shut down and delete management node 10.x VM.

Reconfigure authentication using the management node REST API

You can keep your existing management node if you have sequentially upgraded (1) management services and (2) Element storage. If you have followed a different upgrade order, see the procedures for in-place management node upgrades.

Before you begin

- You have updated your management services to 2.10.29 or later.
- Your storage cluster is running Element 12.0 or later.
- Your management node is 11.3 or later.
- You have sequentially updated your management services followed by upgrading your Element storage. You cannot reconfigure authentication using this procedure unless you have completed upgrades in the sequence described.

Steps

1. Open the management node REST API UI on the management node:

   ```
   https://<ManagementNodeIP>/mnode
   ```

2. Select **Authorize** and complete the following:
   a. Enter the cluster user name and password.
   b. Enter the client ID as `mnode-client` if the value is not already populated.
   c. Select **Authorize** to begin a session.
3. From the REST API UI, select **POST /services/reconfigure-auth**.
4. Select **Try it out**.
5. For the **load_images** parameter, select `true`.
6. Select **Execute**.

   The response body indicates that reconfiguration was successful.

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Upgrade the Element Plug-in for vCenter Server

For existing vSphere environments with a registered NetApp Element Plug-in for VMware vCenter Server, you can update your plug-in registration after you first update the management services package that contains the plug-in service.

You can update the plug-in registration on vCenter Server Virtual Appliance (vCSA) or Windows using the registration utility. You must change your registration for the vCenter Plug-in on every vCenter Server where you need to use the plug-in.

Management services 2.22.7 includes Element Plug-in for vCenter Server 5.0 which contains the remote plug-in. If you use the Element plug-in, you should upgrade to management services 2.22.7 or later to comply with the VMware directive that removes support for local plug-ins. Learn more.

Element vCenter Plug-in 5.0 or later

This upgrade procedure covers the following upgrade scenarios:

- You are upgrading to Element Plug-in for vCenter Server 5.3, 5.2, 5.1, or 5.0.
- You are upgrading to an 8.0 or 7.0 HTML5 vSphere Web Client.

Element Plug-in for vCenter 5.0 or later is not compatible with vCenter Server 6.7 and 6.5.

When you upgrade from Element Plug-in for vCenter Server 4.x to 5.x, the clusters already configured with the plug-in are lost because the data cannot be copied from a vCenter instance to a remote plug-in. You must re-add the clusters to the remote plug-in. This is a one-time activity when upgrading from a local plug-in to a remote plug-in.

Element vCenter Plug-in 4.10 or earlier

This upgrade procedure covers the following upgrade scenarios:

- You are upgrading to Element Plug-in for vCenter Server 4.10, 4.9, 4.8, 4.7, 4.6, 4.5, or 4.4.
- You are upgrading to a 7.0, 6.7, or 6.5 HTML5 vSphere Web Client.

  - The plug-in is not compatible with VMware vCenter Server 8.0 for Element Plug-in for VMware vCenter Server 4.x.
  - The plug-in is not compatible with VMware vCenter Server 6.5 for Element Plug-in for VMware vCenter Server 4.6, 4.7, and 4.8.

- You are upgrading to a 6.7 Flash vSphere Web Client.

The plug-in is compatible with vSphere Web Client version 6.7 U2 for Flash, 6.7 U3 (Flash and HTML5), and 7.0 U1. The plug-in is not compatible with version 6.7 U2 build 13007421 of the HTML5 vSphere Web Client and other 6.7 U2 builds released prior to update 2a (build 13643870). For more information about supported vSphere versions, see the release notes for your version of the plug-in.

What you'll need

- **Admin privileges**: You have vCenter Administrator role privileges to install a plug-in.
- **vSphere upgrades**: You have performed any required vCenter upgrades before upgrading the NetApp Element Plug-in for vCenter Server. This procedure assumes that vCenter upgrades have already been completed.
- **vCenter Server**: Your vCenter Plug-in version 4.x or 5.x is registered with a vCenter Server. From the registration utility (`https://<ManagementNodeIP>:9443`), select **Registration Status**, complete the necessary fields, and select **Check Status** to verify that the vCenter Plug-in is already registered and the version number of the current installation.
- **Management services updates**: You have updated your management services bundle to the latest version. Updates to the vCenter plug-in are distributed using management services updates that are released outside of major product releases for NetApp SolidFire all-flash storage.
- **Management node upgrades**:
  - Beginning with Element vCenter plug-in 5.0, you are running a management node that has been upgraded to version 12.3.x or later.
  - For Element vCenter plug-in 4.4 to 4.10, you are running a management node that has been upgraded to version 11.3 or later. vCenter Plug-in 4.4 or later requires an 11.3 or later management node with a modular architecture that provides individual services. Your management node must be powered on with its IP address or DHCP address configured.
- **Element storage upgrades**:
  - Beginning with Element vCenter plug-in 5.0, you have a cluster running NetApp Element software 12.3.x or later.
  - For Element vCenter plug-in 4.10 or earlier, you have a cluster running NetApp Element software 11.3 or later.
- **vSphere Web Client**: You have logged out of the vSphere Web Client before beginning any plug-in upgrade. The web client will not recognize updates made during this process to your plug-in if you do not log out.

Steps

1. Enter the IP address for your management node in a browser, including the TCP port for registration:
   `https://<ManagementNodeIP>:9443`
   The registration utility UI opens to the **Manage QoSSIOC Service Credentials** page for the plug-in.

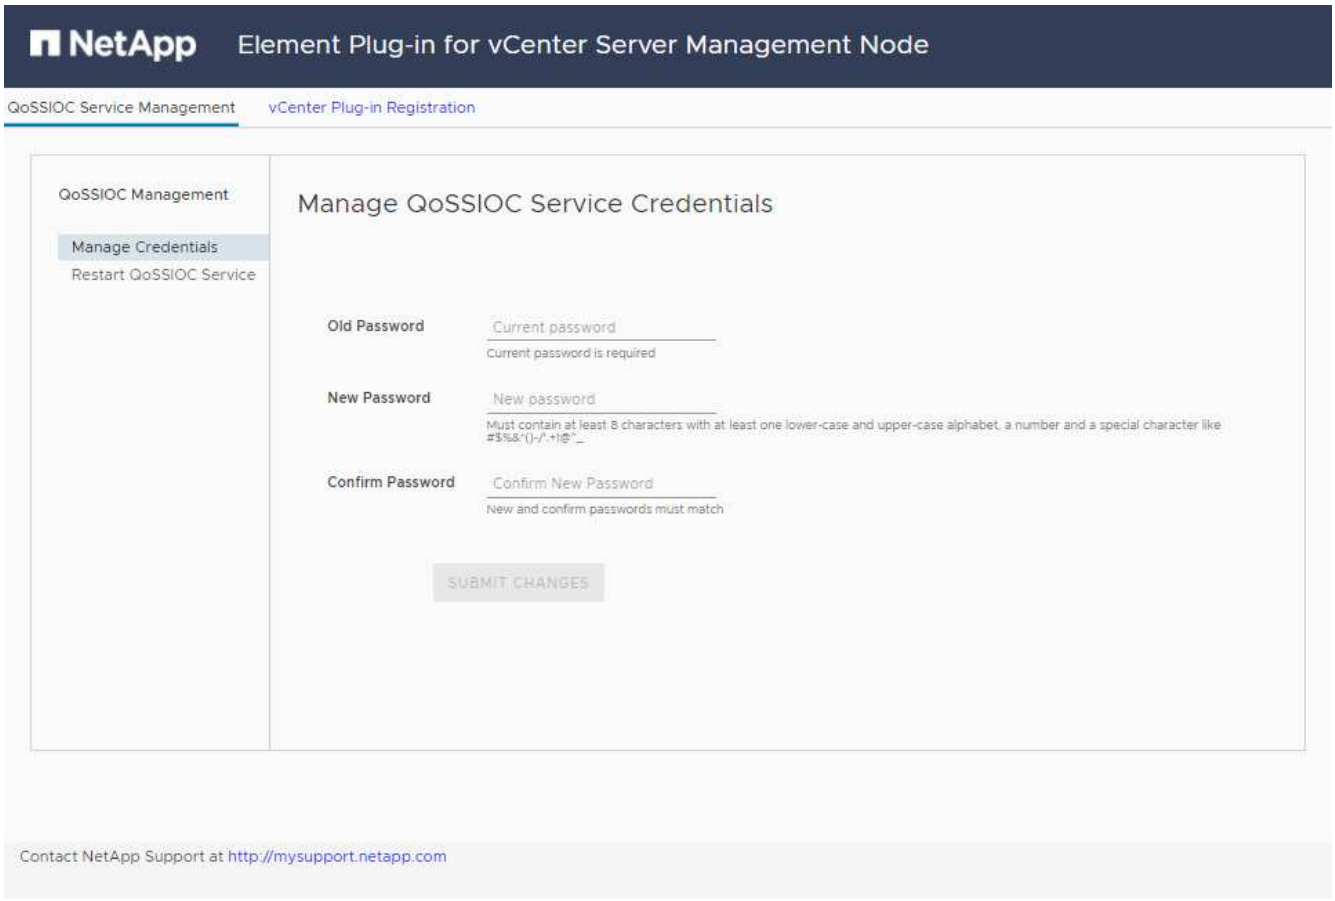

2. Select **vCenter Plug-in Registration**.
   - The vCenter Plug-in Registration page for Element Plug-in for vCenter Server 5.x:

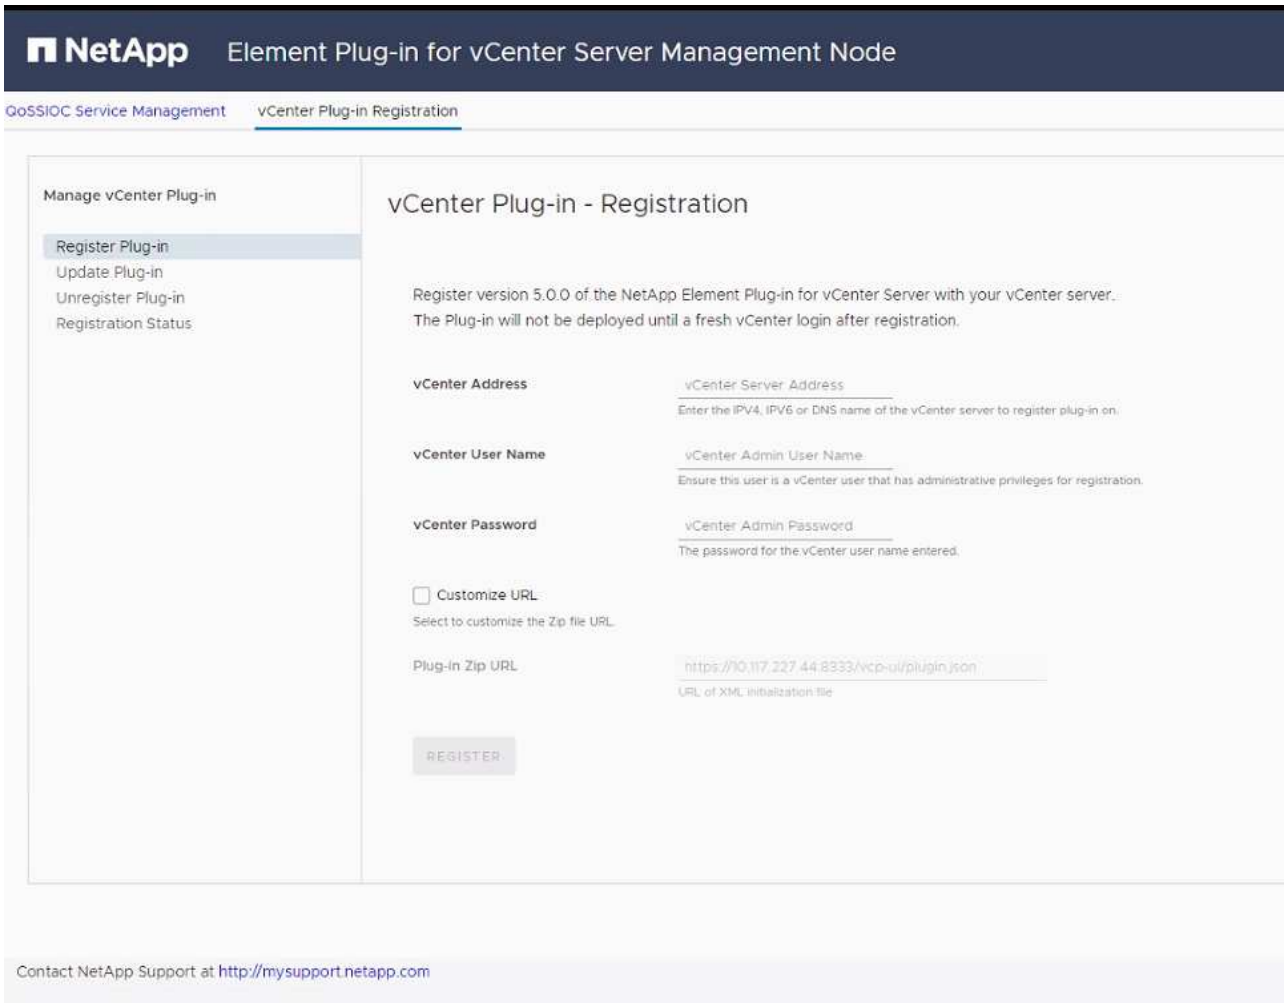

- The vCenter Plug-in Registration page for Element Plug-in for vCenter Server 4.10 or earlier:

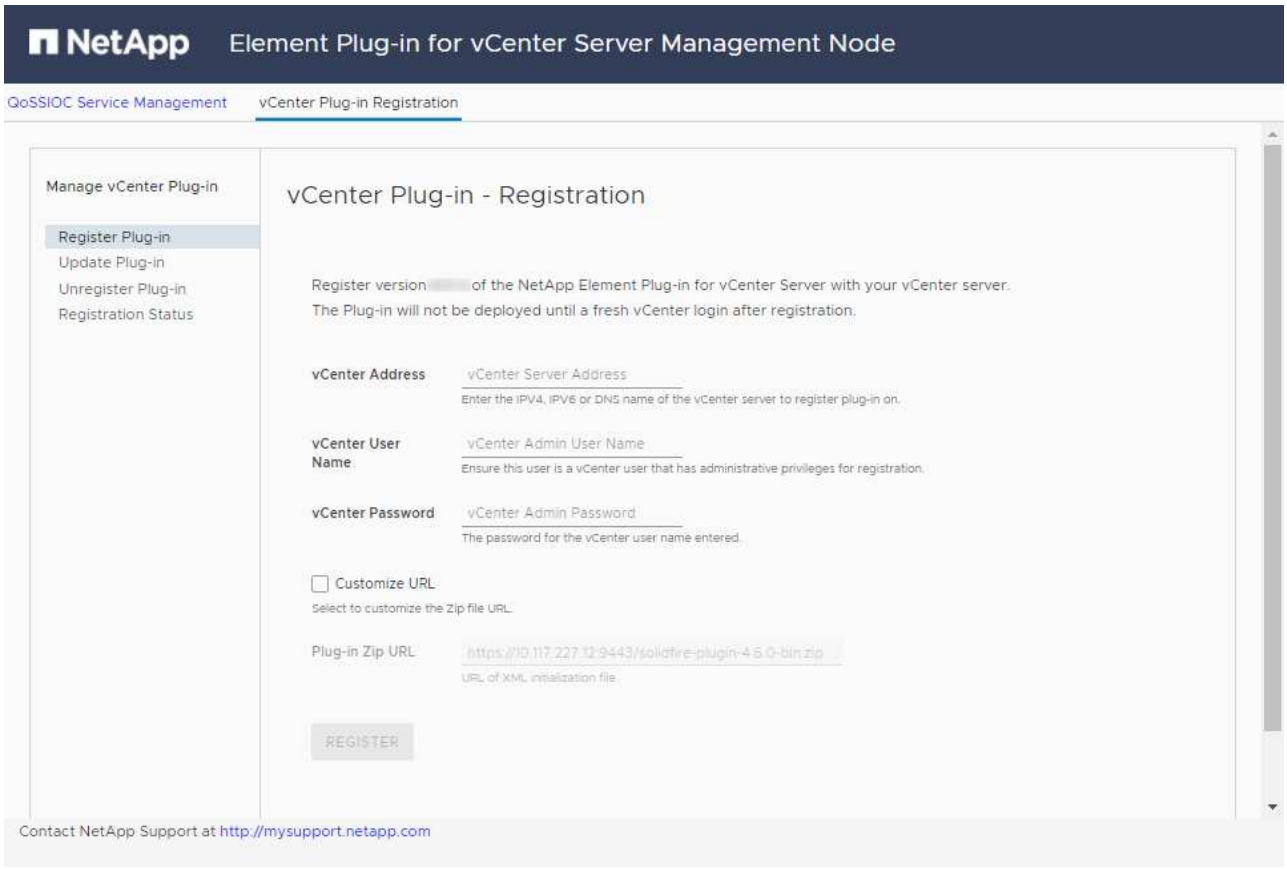

3. Within **Manage vCenter Plug-in**, select **Update Plug-in**.
4. Confirm or update the following information:
   a. The IPv4 address or the FQDN of the vCenter service on which you will register your plug-in.
   b. The vCenter Administrator user name.

      The user name and password credentials you enter must be for a user with vCenter Administrator role privileges.

   c. The vCenter Administrator password.
   d. (For in-house servers/dark sites) Depending on your Element Plug-in for vCenter version, a custom URL for the plug-in JSON file or plug-in ZIP:
      i. Beginning with Element Plug-in for vCenter Server 5.0, a custom URL for the plug-in JSON file.

         You can select **Custom URL** to customize the URL if you are using an HTTP or HTTPS server (dark site) or have modified the JSON file name or network settings. For additional configuration steps if you intend to customize a URL, see Element Plug-in for vCenter Server documentation about modifying vCenter properties for an in-house (dark site) HTTP server.

      ii. For Element Plug-in for vCenter Server 4.10 or earlier, a custom URL for the plug-in ZIP.

You can select **Custom URL** to customize the URL if you are using an HTTP or HTTPS server (dark site) or have modified the ZIP file name or network settings. For additional configuration steps if you intend to customize a URL, see Element Plug-in for vCenter Server documentation about modifying vCenter properties for an in-house (dark site) HTTP server.

5. Select **Update**.

   A banner appears in the registration utility UI when the registration is successful.

6. Log in to the vSphere Web Client as a vCenter Administrator. If you are already logged in to the vSphere Web Client, you must first log out, wait two to three minutes, and then log in again.

   This action creates a new database and completes the installation in the vSphere Web Client.

7. In the vSphere Web Client, look for the following completed tasks in the task monitor to ensure installation has completed: `Download plug-in` and `Deploy plug-in`.
8. Verify that the plug-in extension points appear in the **Shortcuts** tab of the vSphere Web Client and in the side panel.
   - Beginning with Element Plug-in for vCenter Server 5.0, the NetApp Element Remote Plugin extension point appears:

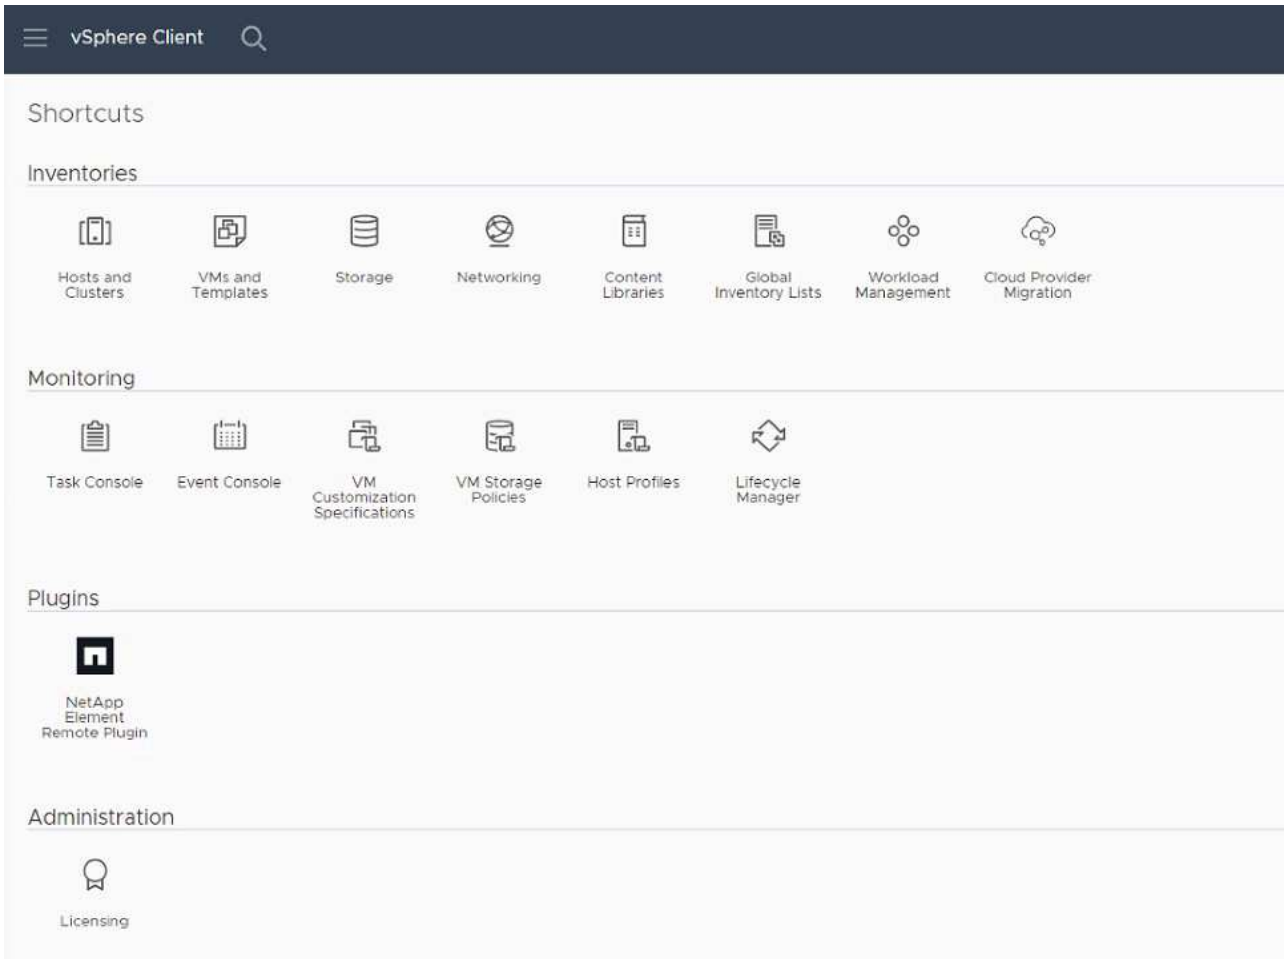

- For Element Plug-in for vCenter Server 4.10 or earlier, the NetApp Element Configuration and Management extension points appear:

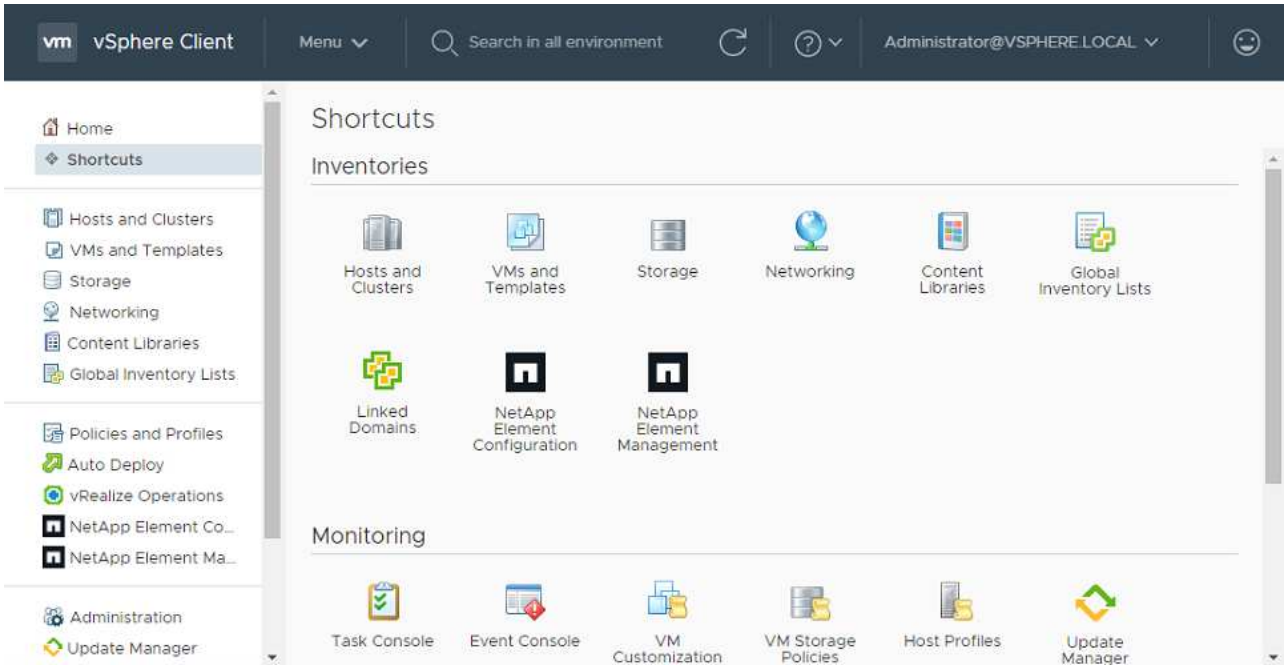

> If the vCenter Plug-in icons are not visible, see Element Plug-in for vCenter Server documentation about troubleshooting the plug-in.
>
> After upgrading to NetApp Element Plug-in for vCenter Server 4.8 or later with VMware vCenter Server 6.7U1, if the storage clusters are not listed or a server error appears in the **Clusters** and **QoSSIOC Settings** sections of the NetApp Element Configuration, see Element Plug-in for vCenter Server documentation about troubleshooting these errors.

9. Verify the version change in the **About** tab in the **NetApp Element Configuration** extension point of the plug-in.

   You should see the following version details or details of a more recent version:

```
NetApp Element Plug-in Version: 5.3
NetApp Element Plug-in Build Number: 9
```

> The vCenter Plug-in contains online Help content. To ensure that your Help contains the latest content, clear your browser cache after upgrading your plug-in.

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Upgrade your vSphere components for a NetApp SolidFire storage system with the Element Plug-in for vCenter Server

When you upgrade the VMware vSphere components of your SolidFire Element storage installation, there are some additional steps you will need to take for systems with Element Plug-in for vCenter Server.

Steps

1. For vCSA upgrades, clear QoSSIOC settings in the plug-in (**NetApp Element Configuration > QoSSIOC Settings**). The **QoSSIOC Status** field displays `Not Configured` after the process is complete.
2. For vCSA and Windows upgrades, unregister the plug-in from the vCenter Server with which it is associated using the registration utility.
3. Upgrade vSphere, including vCenter Server, ESXi, VMs, and other VMware components.

   You should upgrade to NetApp Element Plug-in for vCenter Server 5.0 or later to give you the capability to deploy the plug-in with VMware vCenter 7.0 Update 3 without having to apply a workaround.

   With Element Plug-in for vCenter Server 4.x, when you upgrade to VMware vCenter Server 7.0 Update 3, the plug-in fails to deploy. To resolve this issue using Spring Framework 4, see this KB article.

4. Register the Element Plug-in for vCenter Server again with vCenter.
5. Add clusters using the plug-in.
6. Configure QoSSIOC settings using the plug-in.
7. Enable QoSSIOC for all datastores controlled by the plug-in.

Find more information

- SolidFire and Element Software Documentation
- NetApp Element Plug-in for vCenter Server

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.