

**F5[®] BIG-IQ[®] Centralized Management:
Upgrading a 5.2 or 5.3 DCD Cluster to
Version 5.4 With Minimal Downtime**

Version 5.4



Table of Contents

Data Collection Device Cluster Upgrade Overview.....	5
What is a minimal downtime upgrade?	5
Decide which upgrade guide to use.....	5
Upgrading from version 5.0.....	5
Upgrading from version 5.1.....	5
Upgrading from version 5.2 or 5.3.....	6
How do I upgrade my data collection device cluster to the latest version with minimal downtime?.....	6
Data collection device cluster update requirements.....	8
Prepare for Data Collection Device Upgrade with Minimal Downtime.....	9
Prepare the data collection device cluster for upgrade with minimal downtime.....	9
Check the statistics retention policy.....	9
Check data collection device health.....	9
Confirm disk space is sufficient for minimal downtime upgrade	10
Confirm that the BIG-IP DCD configuration is correct.....	10
Confirm correct minimum master devices setting for version 5.2 or 5.3.....	11
Stop snapshot creation.....	11
Upgrade the Data Collection Devices in Your Cluster with Minimal Downtime.....	13
Prerequisites for upgrading the data collection device to the latest version.....	13
Upgrade the data collection devices to the latest version.....	13
Upload the latest version of the BIG-IQ software image.....	13
Stop sending alerts to the DCD.....	14
Deactivate data collection services.....	14
Upgrade the data collection devices to the latest version.....	15
Confirm that the upgraded DCD is back in service.....	16
Reactivate data collection services after upgrade.....	17
Bring the data collection device back online.....	17
Upgrading BIG-IQ Centralized Management with Data Collection Devices to the Latest Version.....	19
What you need to do before you upgrade BIG-IQ from version 5.x to 5.4.....	19
Remove the secondary BIG-IQ from the HA pair.....	19
Upload the BIG-IQ software image to the secondary BIG-IQ system.....	20
Install the latest version of BIG-IQ software on the HA peer.....	20
Upload the latest version of the BIG-IQ software image.....	21
Upgrade the primary BIG-IQ system to the latest version	22
Completing Post-Upgrade Processes.....	25
Recheck data collection device health.....	25
Reestablish the HA configuration after upgrading to BIG-IQ version 5.4.....	25
Restart snapshot creation.....	26
Re-discover devices and re-import services.....	26
Troubleshooting Minimal Downtime Upgrade Issues.....	29

Table of Contents

What do I do if there is not enough disk space?.....	29
Data collection device cluster status is yellow.....	30
Connectivity Issues.....	30
Index created in higher version.....	30
Data collection device cluster status is red.....	31
Statistics replicas are not enabled.....	31
Generic Failure.....	31
Data collection device cluster is offline.....	32
Election of new master node failed.....	32
Legal Notices.....	35
Legal notices.....	35

Data Collection Device Cluster Upgrade Overview

What is a minimal downtime upgrade?

The minimal downtime upgrade process significantly reduces the amount of time your cluster is offline while you are upgrading to the latest version. The degree to which your cluster can maintain complete online function during upgrade depends on which BIG-IP services and what kind of data you collect.

- Operators using the BIG-IQ to review alerts and events may notice slightly slower performance during the upgrade. Data may take a bit longer to refresh and searches may take a little more time to execute.
- You should not lose any ASM events and FPS alerts during the upgrade.
- You should maintain almost all of your AFM events during the upgrade.
- Statistics data is not collected during the upgrade.
- There is a brief (approximately 5 seconds) downtime when the master node is upgraded.
- There is another short (less than a minute) downtime when the secondary BIG-IQ is paired with the primary.

Decide which upgrade guide to use

There are a number of different upgrade guides you can use to upgrade the BIG-IQ[®] Centralized Management data collection device (DCD) cluster to the latest version. To decide which guide to use, you need to consider a number of factors. First, consider which version you are upgrading from. Then use the appropriate table to determine which guide is most appropriate for your business needs.

Upgrading from version 5.0

To upgrade from version 5.0, use the guide on support.f5.com titled: *F5 BIG-IQ Centralized Management: Upgrading Logging Nodes to Version 5.1*. When a task instructs you to download and install a v5.1.0 file, download and install a v5.4.0 file instead.

Additionally, after the upgrade is complete, you need to use the 5.4.0 user interface to configure an external storage location for the snapshots created for this cluster. To configure the external storage, you will need the IP address, storage file path, domain name, user name, and password for the location you plan to configure. To access the External Storage screen, click **System > BIG-IQ DATA COLLECTION > Big-IQ Data Collection Cluster > Configuration > External Storage & Snapshots** and then click **Configure**. For additional details, refer to *Define external storage snapshots location* on support.f5.com.

Upgrading from version 5.1

Before you upgrade BIG-IQ from version 5.1, you need to answer a couple of questions.

- Do you need to preserve the data on your logging node cluster?
- If so, can you allow your logging node cluster to go offline for an extended period during the upgrade?

Depending on your answers, select one of the following scenarios for further instructions:

I do not need to preserve my data

Use the guide titled: *F5 BIG-IQ Centralized Management: Upgrading Logging Nodes to Version 5.2 Without Preserving Existing Data*. When a task instructs you to download and install a v5.2.0 file, download and install a v5.4.0 file instead.

I need to preserve my data and I can take my logging node cluster offline during upgrade

Use the guide titled: *F5 BIG-IQ Centralized Management: Upgrading Logging Nodes to Version 5.2 While Preserving Existing Data*. When a task instructs you to download and install a v5.2.0 file, download and install a v5.4.0 file instead.

I need to preserve my data, I have at least 3 logging nodes, and I need to minimize logging cluster downtime during upgrade

Use the guide titled: *F5[®] BIG-IQ[®] Centralized Management: Upgrading a Logging Node Cluster from 5.1 to Version 5.4 with Minimal Downtime*.

Regardless of which guide you use, after the upgrade is complete, you need to use the 5.4.0 user interface to configure an external storage location for the snapshots created for this cluster. To configure the external storage, you will need the IP address, storage file path, domain name, user name, and password for the location you plan to configure. To access the External Storage screen, click **System > BIG-IQ DATA COLLECTION > Big-IQ Data Collection Cluster > Configuration > External Storage & Snapshots** and then click **Configure**. For additional detail, refer to *Define external storage snapshots location* on support.f5.com.

Upgrading from version 5.2 or 5.3

To decide which guide to use when you upgrade from BIG-IQ[®] version 5.2 or 5.3, the main question to answer is whether you can allow your data collection device (DCD) cluster to go offline for an extended period during the upgrade. If you can afford the downtime, the workflow is much more straightforward and takes significantly less effort on your part.

I can take my DCD cluster offline during upgrade

Use the guide titled: *F5[®] BIG-IQ[®] Centralized Management: Upgrading a Data Collection Cluster to Version 5.4*

I need to preserve my data, I have at least 3 DCDs, and I need to minimize logging cluster downtime during upgrade

Use the guide titled: *F5[®] BIG-IQ[®] Centralized Management: Upgrading a Data Collection Device Cluster from 5.2 or 5.3 to Version 5.4 with Minimal Downtime*

How do I upgrade my data collection device cluster to the latest version with minimal downtime?

A *data collection device (DCD) cluster* is made up of all of your DCDs, the BIG-IQ[®] Centralized Management device you use to manage them, and an optional BIG-IQ peer device. To avoid taking the DCD cluster offline for an extended time during upgrade, you require a specific workflow.

The following diagram illustrates the task sequence for the minimal downtime workflow.

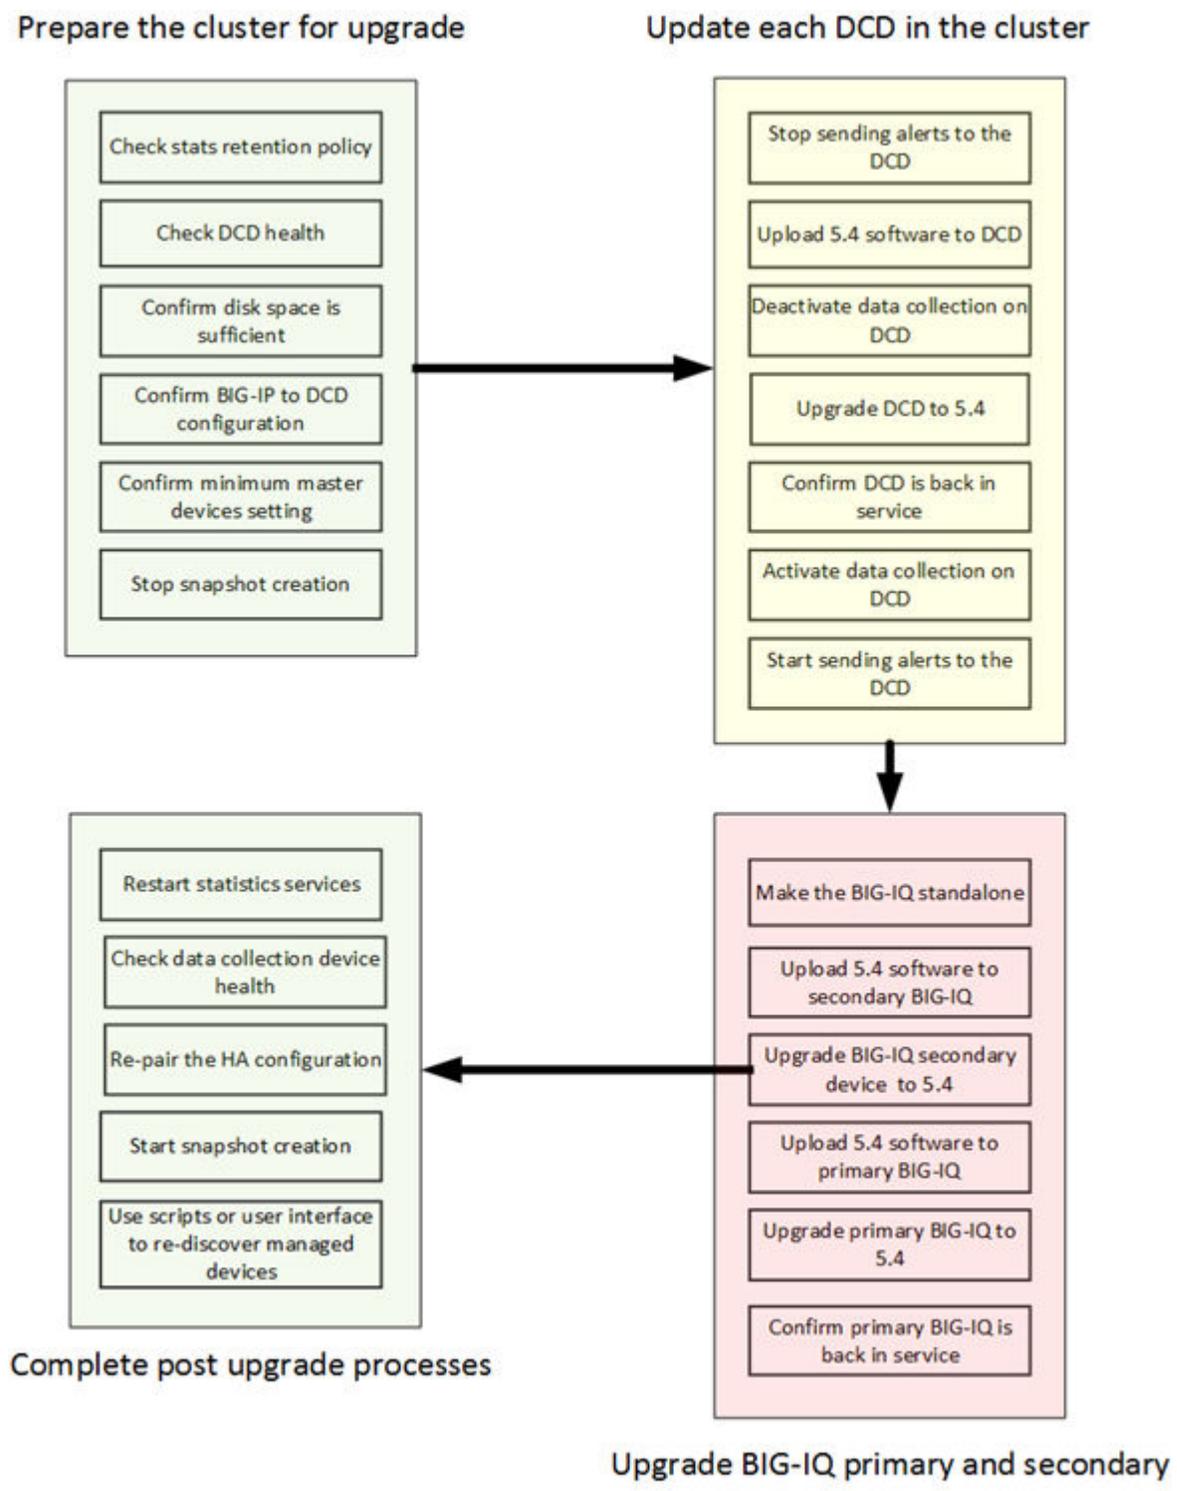


Figure 1: DCD cluster minimal downtime upgrade

Note: To upgrade the DCD cluster successfully with minimal downtime, it is essential that you perform the tasks in the sequence detailed here. You cannot perform these tasks in parallel.

Data collection device cluster update requirements

There are a number of items that you need to have ready access to, before you can successfully complete the data collection device cluster update. It might be helpful to use the following table to record the values, so you have them at hand when needed.

Device	Information needed	Information Recorded
Primary BIG-IQ Centralized Management device	Management IP address admin user name and password root access password	
Secondary BIG-IQ device	Management IP address root access password admin user name and password	
Data collection devices	Management IP address root access password	

Prepare for Data Collection Device Upgrade with Minimal Downtime

Prepare the data collection device cluster for upgrade with minimal downtime

If you choose the minimal downtime method, you can upgrade your data collection device cluster minimizing the time that the cluster is offline.

Important: *You cannot perform the minimal downtime upgrade unless you have at least 3 DCDs in your cluster.*

Check the statistics retention policy

Before you begin an update with minimal downtime, you must confirm that your statistics data is replicated to multiple DCDs. Perform this task on the primary BIG-IQ® system.

1. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.
The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.
2. Click the **Settings** button.
The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
3. On the left, click **Statistics Collection**.
The Statistics Collection Status screen displays the percentage of available disk space currently consumed by statistics data for each container.
4. To change the retention settings for your statistics data, click **Configure**.
The Statistics Retention Policy screen opens.
5. Expand Advanced Settings, and then select the **Enable Replicas** check box.
Replicas are copies of a data set that are available to the DCD cluster when one or more devices within that cluster become unavailable. You must allow the system to create these copies to upgrade your DCD cluster with minimal downtime and lost data.
6. Click **Save & Close**.
7. Return to the BIG-IQ Data Collection Devices screen as in step 2. If you just changed the replication policy, the cluster health status icon should be yellow.
Until the data replicas are distributed to the cluster, the status does not become green.
8. Wait for the cluster health status icon to turn to green, so you know that the data is safely distributed to multiple DCDs.

Check data collection device health

You can use the BIG-IQ Data Collection screen to review the overall health and status of the data collection devices you've configured. You can use the data displayed on this screen both before and after an upgrade to verify that your data collection device cluster configuration is as you expect it to be.

Note: *Perform this task only on the primary BIG-IQ system.*

1. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.
The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.
2. Click the **Settings** button.
The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
3. Check the cluster health status. If the cluster health is not green, then resolve those issues and then repeat the process until the cluster is operating normally.

Confirm disk space is sufficient for minimal downtime upgrade

As part of preparing to upgrade your data collection device with minimal downtime, you must confirm that there is sufficient disk size in the cluster so that when you take a DCD offline, there is room for its data on other devices in the cluster. If the amount of free space in the cluster is less than the amount of data on any one node, then there is insufficient space to upgrade without downtime. If this is the case, you need to either add DCDs or increase storage space on the existing DCDs.

Important: *If your cluster has multiple zones, you must perform the disk space check for each zone.*

1. Use SSH to log in to a device in the cluster.
You must log in as `root` to perform this procedure.
2. Determine the storage space requirement for your DCD cluster using the following command:
`curl localhost:9200/_cat/allocation?v`

```
shards disk.indices disk.used disk.avail disk.total disk.percent host ip node
57 397.5mb 2gb 7.8gb 9.8gb 20 10.10.10.5 10.10.10.5 8637c04c-1b83-4795-b1f0-347ac733fd10
56 471.7mb 2.2gb 7.5gb 9.8gb 23 10.10.10.3 10.10.10.3 9d718ba7-5bb9-4866-9aa3-4677a1f60e46
56 393mb 2.1gb 7.7gb 9.8gb 21 10.10.10.2 10.10.10.2 8c4e58b4-a005-404f-9a53-6e318ec0e381
57 444.2mb 2gb 7.8gb 9.8gb 20 10.10.10.10 10.10.10.10 11ac40f9-5b13-4f9a-a739-0351858ba571
```

3. Analyze the storage space requirement for your DCD to determine if there is sufficient disk space.
In the previous example, there is plenty of space. The DCD consuming the most data is only consuming 2.2 GB, and each of the other DCDs has almost 8 GB free. So when that DCD goes offline to upgrade, the system can move the 2.2 GB of data to the remaining 15.5 GB of free space. If these numbers were reversed, so that the DCD consuming the most storage had 7.8 GB of data, and the remaining DCDs only had 6.3 GB free, there would be insufficient space to move the data when that DCD went offline.

If there is sufficient space, you can proceed. Otherwise, you need to either add DCDs, or add DCD storage space.

Confirm that the BIG-IP DCD configuration is correct

When preparing to upgrade your data collection device with minimal downtime, you must confirm that the BIG-IP® device to DCD configuration is correct, so that when a particular DCD is upgraded, data that was being routed to it is automatically routed to another DCD in the cluster. There are two settings of particular concern:

- Confirm that data sent from the BIG-IP devices is not being sent to just one DCD. Each BIG-IP device must be configured to send data to multiple DCDs.
- Confirm that the BIG-IP devices are configured with appropriate monitors that allow for traffic to switch to a different DCD when one DCD is taken offline.

1. Analyze the data routing configuration for all of the BIG-IP devices that send data to your DCD cluster.

2. If you find a BIG-IP device that is configured to send data to only one DCD, change that configuration before proceeding.
Refer to the BIG-IP documentation on support.f5.com for details about how to configure the BIG-IP to DCD routing.
3. Analyze the monitor configuration for all of the BIG-IP devices that send data to your DCD cluster. Make sure that each device is configured to send data to an alternate DCD if one DCD goes offline.
4. If you find a BIG-IP device that is not configured with appropriate monitors, change that configuration before proceeding.
Refer to the BIG-IP documentation on support.f5.com for details about how to configure BIG-IP monitors correctly.

Confirm correct minimum master devices setting for version 5.2 or 5.3

In preparing to upgrade your data collection device with minimal downtime, you must confirm that the minimum master devices setting is correctly configured so that when a DCD is upgraded, its data can fail over to another DCD in the cluster.

Note: Perform this task only on the primary BIG-IQ system.

1. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.
The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.
2. Click the **Settings** button.
The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
3. Confirm that the value of the **Minimum Master Eligible Devices** setting is correct.
 - If you are upgrading a single zone DCD cluster, confirm that the number of devices is less than or equal to the number of DCDs in the cluster.
 - If you are upgrading a multiple zone DCD cluster, confirm that the number of devices is less or equal to than the number of DCDs in each zone.
4. If the value is too high, click **Update**, and type in a number that is less than the number of DCDs in the cluster.
5. Click the **Save & Close** button at the bottom of the screen.

Stop snapshot creation

Because of the mixed software version environment that occurs during the upgrade process, if snapshot schedules are configured for the cluster, you should stop creating snapshots before you begin the upgrade to prevent possible issues with your data.

1. Use SSH to log in to the primary BIG-IQ system for this cluster.
You must log in as `root` to perform this procedure.
2. Retrieve the list of scheduled snapshots using the following command: `restcurl cm/shared/esmgmt/es-snapshot-task | grep task-scheduler`

```
config # restcurl cm/shared/esmgmt/es-snapshot-task | grep task-scheduler
"link": "https://localhost/mgmt/shared/task-scheduler/scheduler/0fdf50ec-8a17-3da9-b717-c63637ccc68a"
"link": "https://localhost/mgmt/shared/task-scheduler/scheduler/0af33352-2f33-32b3-85cb-1281bb88c249"
"link": "https://localhost/mgmt/shared/task-scheduler/scheduler/2ad770a8-bdb0-3383-99a9-300846eb0972"
```

In the example here, there are 3 snapshots scheduled.

Prepare for Data Collection Device Upgrade with Minimal Downtime

3. Stop each of the schedules using the following command: `restcurl -X PATCH -d '{"status":"DISABLED"}' shared/task-scheduler/scheduler/<SNAPSHOT ID>`

```
#restcurl -X PATCH -d '{"status":"DISABLED"}'  
shared/task-scheduler/scheduler/0af33352-2f33-32b3-85cb-1281bb88c249  
{ "id": "0af33352-2f33-32b3-85cb-1281bb88c249", "status":"DISABLED", ...}
```

After you run the command for each scheduled snapshot, no more snapshots are created.

Upgrade the Data Collection Devices in Your Cluster with Minimal Downtime

Prerequisites for upgrading the data collection device to the latest version

Before upgrading the F5® BIG-IQ® Centralized Management data collection device, perform these tasks.

Tasks	Additional information
Reactivate the BIG-IQ system license.	You must do this for each data collection device in the cluster. For specific instructions about how to reactivate a license, refer to the <i>F5® BIG-IQ® Central Management: Licensing and Initial Setup</i> guide.
Decide which disk volume to install the upgrade on. At least two volumes are required to upgrade BIG-IQ.	If the machine you are upgrading does not have two volumes, you must create one using a <code>CLI</code> command. The upgrade task provides options for either case.

Upgrade the data collection devices to the latest version

After you prepare the data collection device (DCD) cluster for upgrade, use these procedures to upgrade the data collection devices in your cluster. Perform the following tasks for each DCD in the cluster.

1. Upload the latest version of the BIG-IQ version software image.
2. From the BIG-IP device(s), take the DCD out of the alert-traffic flow.
3. Deactivate data collection services.
4. Upgrade the DCD to the latest version.
5. Confirm that the DCD is back in service
6. Reactivate data collection services.
7. From the BIG-IP device(s), bring the DCD back online.

Important: To successfully upgrade your DCD cluster with minimal downtime, it is critical that you perform this sequence of tasks to each DCD in the cluster serially, not in parallel.

Upload the latest version of the BIG-IQ software image

Before you can upload the software image to your data collection device (DCD), you must have first downloaded it from the F5 Downloads site.

You upload the latest BIG-IQ® version of the software image to your DCD to make it available for this upgrade.

1. At the top of the screen, click **System**.
2. On the left, click **THIS DEVICE > SOFTWARE MANAGEMENT > Available Images**.
3. Click the **Upload Image** button.
4. Click the **Choose File** button and go to the location to which you downloaded the image, and click the **Open** button to upload it to BIG-IQ.

5. Click the **Upload** button.

The screen refreshes to display the progress of the upload.

When the image is finished uploading, it shows in the Available Images list.

Stop sending alerts to the DCD

To upgrade your DCD cluster with minimum downtime, before you begin the process of updating your DCDs, you must complete the upgrade preparation detailed in *Prepare for Data Collection Device Upgrade With Minimal Downtime*.

This is the first task you must perform on each DCD in the cluster to make sure that you do not lose any data during the upgrade. When you take the DCD you are about to upgrade offline, the BIG-IP® device routes the data to another DCD in the cluster.

Log in to each BIG-IP device that sends data to your DCD cluster, and revise the configuration so that the data it sends to the cluster does not include the particular DCD you are preparing to upgrade.

Note: The process for stopping data collection for one DCD in the cluster depends on the kind of data you collect.

Option	Description
If you collect FPS alerts or IPsec events:	<ol style="list-style-type: none">1. On the BIG-IP device, select Local Traffic > Pools > Pool List.2. Select the pool that sends alerts and events to the DCD.3. Click Members.4. Select the pool member that sends alerts to the target DCD, and disable it.
If you collect ASM event:	<ol style="list-style-type: none">1. On the BIG-IP device, select Local Traffic > Pools > Pool List.2. Select the pool that contains the members that send events to the DCD.3. Click Members.4. Select the pool member that sends alerts to the target DCD, and disable it.
If you collect APM events:	You don't need to disable the DCD from the BIG-IP device. All you need to do is deactivate the service.

Deactivate data collection services

You must perform this task for each DCD in the cluster to make sure that you minimize data loss during the upgrade. When you take the DCD you are about to upgrade off line, the BIG-IP® device routes the data to another DCD in the cluster.

1. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.
The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.
2. Select the device name of the DCD that you want to upgrade first.
The Properties screen for the selected DCD opens.
3. On the left, select **Services**.
4. For each active service, click **Deactivate**.

Important: Before you exit this screen, make a note of the Listener Address for this DCD. You will need it later in the upgrade process.

5. Click **Cancel**.

The screen closes and BIG-IQ deactivates your services.

Upgrade the data collection devices to the latest version

Note: If the machine you are upgrading does not have two volumes, you must create one using a CLI command. The following steps provide options for either case.

Before upgrading the data collection device (DCD), you must have downloaded the latest version of the BIG-IQ software image from the F5 downloads site.

The upgrade process installs the new version of the software on one DCD. You must repeat this process for each DCD in your cluster.

1. At the top of the screen, click **System**.
2. On the left, expand **THIS DEVICE > Software Management**, and then click **Available Images**. The Available Images screen lists the software image files that are available for install on this BIG-IQ system.
3. Select the check box for the image you want to install and click **Install**.
4. From the **Target Volume** list, select the volume you want to install the image on.

Note: If the Target Volume list is empty, the machine you are upgrading does not have a second volume available for you to install the upgrade. You must create one to proceed.

- If you selected a target volume, complete the remaining steps to finish the upgrade.
 - If you do not have a second volume, use the following procedure to install the upgrade for this machine:
 - Use SSH to log in to the machine you want to upgrade. You must log in as `root` to perform this procedure.
 - Navigate to the folder where you install software using the command: `cd /shared/images`
 - Install the upgrade to a new volume by running one of the following commands:
 - If you are upgrading a virtual machine: `tmsh install sys software image <full name of the iso file> volume hd1.2 create-volume reboot`
 - If you are upgrading a BIG-IQ 7000 series platform: `tmsh install sys software image <full name of the iso file> volume md1.2 create-volume reboot`
 - While the system installs the upgrade to the new volume, you can check on the status of the upgrade using the command: `tmsh show sys software status`.
 - When the upgrade finishes installing, the machine reboots. Proceed to the next to the last (login) step in this task.
5. To prompt the data collection device to reboot into the new software installation volume, select the **Reboot into Target Volume** check box.
 6. Click the **Install** button.
A confirmation dialog box appears.
 7. Click the **Continue** button.
 8. Wait while the data collection device loads the new software and reboots.
Depending on your configuration and the number of devices you are managing, this could take up to 15 minutes. During this time, it is important that you not interrupt the installation process by restarting services or the server.
 9. Log back in to the data collection device to confirm that the upgrade succeeded.
 10. If needed, extend the `/var` partition.

The default size of the `/var` file system in a newly installed or upgraded node is 30 GB. This amount of available storage might not be enough to store your data. If you previously extended your disk

space, you'll need to do it again after you upgrade, as explained in *K16103: Extending disk space on BIG-IQ Virtual Edition*.

Important: *Upgrading a node requires at least two volumes, so be sure you can extend the /var file system on both volumes to the same size; otherwise, upgrades might fail.*

Note: *Remember to complete this process for each DCD in the cluster.*

Important: *In the unlikely event that you are unable to log in to a BIG-IQ® 7000 series platform after you upgraded it, refer to SOL40338232: The BIG-IQ system interface might be inaccessible after the BIG-IQ system is upgraded from BIG-IQ Centralized Management version to 5.2.0 at support.f5.com/kb/en-us/solutions/public/k/40/sol40338232.html for more information.*

Confirm that the upgraded DCD is back in service

After you upgrade a DCD as part of the minimum downtime workflow, confirm that it is back in service before proceeding to upgrade the next DCD in the cluster.

1. Log in to the BIG-IQ system primary with your admin user name and password.
2. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.

The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.

3. Note the status icon color for the DCD that you just upgraded. If it is not green yet, wait until it turns to green before proceeding.

When the icon turns green, it indicates that the primary BIG-IQ system and the DCD are communicating.

4. Use SSH to log in to the DCD you just upgraded.
You must log in as `root` to perform this procedure.

Important: *It is critical that you log in to the correct device. Log in to the DCD, not the primary BIG-IQ.*

5. Determine whether the DCD cluster service restarted, using the following command:
`bigstart status elasticsearch`

In this example, the DCD cluster has been up and running for less than an hour.

```
elasticsearch run (pid 5110) .5 hours
```

6. Determine whether this DCD has rejoined the DCD cluster, using the following command:
`curl localhost:9200/_cat/nodes?v`

In the following example, the recently upgraded DCD (10.11.0.96) has rejoined the cluster (three other DCDs, the primary BIG-IQ, and the secondary BIG-IQ).

```
curl localhost:9200/_cat/nodes?v
host      ip          heap.percent ram.percent load node.role master
name
10.11.0.96 10.11.0.96    2           63 0.01 d         m       d6211
10.11.0.97 10.11.0.97    2           63 0.01 d         m       d6212
10.11.0.98 10.11.0.98    2           63 0.01 d         m       d6213
10.11.0.99 10.11.0.99    2           63 0.01 d         m       d6214
10.144.73.126 10.144.73.126 22          45 0.04 -         *       09020
b0c-7021-458f-89f4-5e89faf64670
10.144.73.127 10.144.73.127 22          45 0.05 -         m       09020
c0c-2235-234a-54b2-2f45fab3233b
```

7. If the DCD did not rejoin the cluster, restart the cluster service using this command:
`bigstart restart elasticsearch`
The cluster service restarts, and the DCD rejoins the cluster. Repeat step 7 to confirm.
8. Log in to the BIG-IQ system primary again with your admin user name and password.
9. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.
The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.
10. Note the Data Cluster Status icon color. Do not proceed until it turns green.
When the icon turns green, it indicates that data is again distributed to all members of the DCD cluster.

Note: If the DCD cluster status icon does not turn green, refer to *Troubleshooting Minimal Downtime Upgrade Issues*. There are several common issues that you can resolve by following the recommended actions.

The DCD cluster is now fully restored to service.

Now you need to reactivate services and bring the upgraded DCD back online so that it can start collecting data again.

Reactivate data collection services after upgrade

After you upgrade a DCD and it is again a fully functioning member of the DCD cluster, you need it to start collecting data again.

Note: Perform this task only if you are upgrading from BIG-IQ version 5.2 or 5.3 to version 5.4.

1. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION > BIG-IQ Data Collection Devices**.
The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster.
2. Select the device name of the DCD you just upgraded.
The Properties screen for the selected DCD opens.
3. On the left, select **Services**.
4. For each service that you deactivated previously, click **Activate**.

Important: Use the value from the Listener Address setting that you noted when you deactivated services for this DCD.

5. Click **Cancel**.

The screen closes and BIG-IQ activates your services.

Bring the data collection device back online

After you upgrade a DCD and reactivate data collection services, you need to start sending data to it again.

Log in to each BIG-IP® device that sends data to your DCD cluster and revise the configuration so that the data it sends to the cluster includes the DCD you just upgraded.

Note: The process for starting data collection for one DCD in the cluster depends on the kind of data you collect.

Upgrade the Data Collection Devices in Your Cluster with Minimal Downtime

Option	Description
If you collect FPS alerts:	<ol style="list-style-type: none">1. On the BIG-IP device, select Local Traffic > Pools > Pool List.2. Select the pool that sends alerts and events to the DCD.3. Click Members.4. Select the pool member that sends alerts to the target DCD, and enable it.
If you collect ASM events:	<ol style="list-style-type: none">1. On the BIG-IP device, select Local Traffic > Pools > Pool List.2. Select the pool that contains the members that send events to the DCD.3. Click Members.4. Select the pool member that sends alerts to the target DCD, and enable it.
If you collect APM events:	You don't need to enable the DCD from the BIG-IP device. All you need to do is activate the service.

Now that the upgraded DCD has rejoined the cluster, the services are reactivated, and the BIG-IP devices are sending data again, you can update the next DCD in the cluster.

Upgrading BIG-IQ Centralized Management with Data Collection Devices to the Latest Version

What you need to do before you upgrade BIG-IQ from version 5.x to 5.4

Before upgrading F5® BIG-IQ® Centralized Management, perform the following tasks.

Tasks	Additional information
Create a backup of the BIG-IQ system's current compressed user configuration set (UCS) and store it on a remote server.	The UCS file includes: system-specific configuration files; license, user account and password information; and SSL certificates and keys. You can use this backup if you want to restore to the previous version of BIG-IQ.
Decide which disk volume to install the upgrade on. At least two volumes are required to upgrade BIG-IQ.	If the machine you are upgrading does not have two volumes, you must create one using a CLI command. The upgrade task provides options for either case.
Deploy any staged configuration changes to your managed devices.	Deploy configuration changes you have staged for your devices. When addressing configuration conflicts for each BIG-IP device, we recommend you use BIG-IP to override the configuration settings stored on BIG-IQ. Refer to the section titled, <i>Re-discover devices and re-import services from the user interface</i> .
If you are currently using a self-IP address for device discovery, make a note of that IP address.	You'll need to enter it when you perform setup after you upgrade and reboot the BIG-IQ system.

If you're currently running a version of BIG-IQ prior to version 5.0, you must first upgrade to version 5.0 before you can upgrade to version 5.4. For more information, refer to the guide titled, *F5 BIG-IQ Centralized Management: Upgrading BIG-IQ to Version 5.0*.

Remove the secondary BIG-IQ from the HA pair

If the F5® BIG-IQ® Centralized Management system is configured in an HA pair, you must remove the secondary BIG-IQ system before you upgrade it.

1. Log in to the primary BIG-IQ with your admin user name and password.
2. At the top of the screen, click **System**.
3. On the left, click **BIG-IQ HA**.
4. Click **Remove Secondary**.
A dialog box opens, prompting you to confirm that you want to remove the HA secondary device from this group.
5. Click **Remove** to confirm that you want to take the HA secondary device out of the group.
The system logs you out of the BIG-IQ while it removes the secondary device.
6. Log back in to the primary BIG-IQ with your admin user name and password.
For a while, both the primary and the secondary BIG-IQ devices continue to display. After a few minutes, the screen updates to display a single standalone device.

Upload the BIG-IQ software image to the secondary BIG-IQ system

Before you can upload the software image to your secondary BIG-IQ[®] system, you must have first downloaded it from the F5 Downloads site.

Upload the latest version of the BIG-IQ software image to your secondary BIG-IQ system to make it available for this upgrade.

1. At the top of the screen, click **System**.
2. On the left, expand **THIS DEVICE > Software Management**, and then click **Available Images**. The Available Images screen lists the software image files that are available for installation on this BIG-IQ system.
3. Click the **Upload Image** button.
4. Click the **Choose File** button and go to the location to which you downloaded the image, and click the **Open** button to upload it to BIG-IQ.
5. Click the **Upload** button.
The screen refreshes to display the progress of the upload.

When the image is finished uploading, it shows in the Available Images list.

Install the latest version of BIG-IQ software on the HA peer

Note: If the machine you are upgrading does not have two volumes, you must create one using a CLI command. The following steps provide options for either case.

Install the latest version of the BIG-IQ software on the peer BIG-IQ system so it will be running the same version as the primary BIG-IQ system you are upgrading.

1. Log on to the system you are going to establish as the secondary BIG-IQ system's command line as **root** and type the following command: `/usr/bin/clear-rest-storage`.
Although this step is not required, it clears the database storage on the system so the upgrade goes more quickly. After you upgrade both peer systems and re-establish the HA pair, the peer BIG-IQ will synchronize its database with this BIG-IQ system and repopulate the database.
2. Log on to the system you are going to establish as the secondary BIG-IQ system's user interface.
3. If you ran the `clear-rest-storage` command, complete the setup wizard. Otherwise, continue to the next step.
4. At the top of the screen, click **System Management**.
5. On the left, click **BIG-IQ DEVICES > BIG-IQ HA**.
6. On the left, click **Software Version**.
7. Click the **Update** button.
8. From the **Software Image** list, select the image you want to install.
9. From the **Target Volume** list, select the volume you want to install the image on.

Note: If the Target Volume list is empty, the machine you are upgrading does not have a second volume available for you to install the upgrade. You must create one to proceed.

- If you selected a target volume, complete the remaining steps to finish the upgrade.

- If you do not have a second volume, use the following procedure to install the upgrade for this machine:
 - Use SSH to log in to the machine you want to upgrade. You must log in as `root` to perform this procedure.
 - Navigate to the folder where you install software using the command: `cd /shared/images`
 - Install the upgrade to a new volume by running one of the following commands:
 - If you are upgrading a virtual machine: `tmsh install sys software image <full name of the iso file> volume hd1.2 create-volume reboot`
 - If you are upgrading a BIG-IQ 7000 series platform: `tmsh install sys software image <full name of the iso file> volume md1.2 create-volume reboot`
 - While the system installs the upgrade to the new volume, you can check on the status of the upgrade using the command: `tmsh show sys software status`.
 - When the upgrade finishes installing, the machine reboots. Proceed to the next to the last (login) step in this task.
- 10. To prompt BIG-IQ to reboot into the new software installation volume, select the **Reboot into Target Volume** check box.
- 11. Click the **Apply** button.

A popup screen opens, prompting you to confirm the installation.
- 12. Click the **Continue** button.
- 13. Wait while BIG-IQ loads the new software and reboots.

Depending on your configuration and the number of devices you are managing, this could take up to an hour. During this time, it is important that you not interrupt the installation process by restarting services or the server.
- 14. Log in to the BIG-IQ that will be the secondary BIG-IQ system to confirm that the upgrade succeeded and complete the setup wizard.

To complete the setup task, you must supply a master key pass phrase that all devices in the logging node cluster must share. Make note of the phrase you use. Each logging node, as well as the primary and secondary BIG-IQ systems must use this phrase before they can join the cluster.
- 15. If needed, extend the `/var` partition.

The default size of the `/var` file system in a newly installed or upgraded node is 30 GB. This amount of available storage might not be enough to store your data. If you previously extended your disk space, you'll need to do it again after you upgrade, as explained in *K16103: Extending disk space on BIG-IQ Virtual Edition*.

Important: *Upgrading a node requires at least two volumes, so be sure you can extend the `/var` file system on both volumes to the same size; otherwise, upgrades might fail.*

Upload the latest version of the BIG-IQ software image

Before you can upload the software image to the BIG-IQ[®] system, you must have first downloaded it from the F5 Downloads site.

Upload the BIG-IQ version software image to the primary BIG-IQ system to make it available for this upgrade.

1. At the top of the screen, click **System**.
2. On the left, click **THIS DEVICE > SOFTWARE MANAGEMENT > Available Images**.
3. Click the **Upload Image** button.
4. Click the **Choose File** button and go to the location to which you downloaded the image, and click the **Open** button to upload it to BIG-IQ.

5. Click the **Upload** button.

The screen refreshes to display the progress of the upload.

When the image is finished uploading, it shows in the Available Images list.

Upgrade the primary BIG-IQ system to the latest version

Note: If the machine you are upgrading does not have two volumes, you must create one using a CLI command. The following steps provide options for either case.

Before upgrading the primary BIG-IQ in a DCD cluster you must:

1. Prepare the DCD cluster for upgrade. See *Prepare for Data Collection Device Upgrade* on support.f5.com.
2. Update the DCDs in the cluster.
3. Download the BIG-IQ version 5.4 .iso image from the F5 downloads site.

Upgrade the BIG-IQ primary to take advantage of the newest functionality and features.

1. Log in to the primary BIG-IQ system with your admin user name and password.
2. At the top of the screen, click **System**.
3. On the left, expand **THIS DEVICE > Software Management**, then click **Available Images**. The Available Images screen lists the software image files that are available for install on this BIG-IQ system.
4. Select the check box for the image you want to install and click **Install**.
5. From the **Target Volume** list, select the volume you want to install the image on.

Note: If the Target Volume list is empty, the machine you are upgrading does not have a second volume available for you to install the upgrade. You must create one to proceed.

- If you selected a target volume, complete the remaining steps to finish the upgrade.
 - If you do not have a second volume, use the following procedure to install the upgrade for this machine:
 - Use SSH to log in to the machine you want to upgrade. You must log in as `root` to perform this procedure.
 - Navigate to the folder where you install software using the command: `cd /shared/images`
 - Install the upgrade to a new volume by running one of the following commands:
 - If you are upgrading a virtual machine: `tmsh install sys software image <full name of the iso file> volume hd1.2 create-volume reboot`
 - If you are upgrading a BIG-IQ 7000 series platform: `tmsh install sys software image <full name of the iso file> volume md1.2 create-volume reboot`
 - While the system installs the upgrade to the new volume, you can check on the status of the upgrade using the command: `tmsh show sys software status`.
 - When the upgrade finishes installing, the machine reboots. Proceed to the next to the last (login) step in this task.
6. To prompt BIG-IQ to reboot into the new software installation volume, select the **Reboot into Target Volume** check box.
 7. Click the **Install** button.
 8. Click the **Continue** button.
 9. Wait while BIG-IQ loads the new software and reboots.

Depending on your configuration and the number of devices you are managing, this could take up to an hour. During this time, it is important that you not interrupt the installation process by restarting services or the server.

10. Log in to the primary BIG-IQ to confirm that the upgrade succeeded.

11. If needed, extend the `/var` partition.

The default size of the `/var` file system in a newly installed or upgraded node is 30 GB. This amount of available storage might not be enough to store your data. If you previously extended your disk space, you'll need to do it again after you upgrade, as explained in *K16103: Extending disk space on BIG-IQ Virtual Edition*.

Important: *Upgrading a node requires at least two volumes, so be sure you can extend the `/var` file system on both volumes to the same size; otherwise, upgrades might fail.*

Even though you can log in to the BIG-IQ system after the software is installed, the system continues running background database re-indexing processes. For large configurations, re-indexing can take up to an hour. If you perform any searches on objects before re-indexing is complete, BIG-IQ might not return the expected results. During this time, you can continue with the rest of the upgrade process.

Completing Post-Upgrade Processes

Recheck data collection device health

You can use the Summary Status screen to review the overall health and status of the data collection devices you've configured. You can use the data displayed on this screen both before and after an upgrade to verify that your data collection device cluster configuration is as you expect.

1. At the top of the screen, click **System**.
2. At the top of the screen, click **System**, and then, on the left, click **BIG-IQ DATA COLLECTION** and select **BIG-IQ Data Collection Cluster**.
 - Under Summary, you can view information detailing how much data is stored, as well as how the data is stored.
 - Under Configuration, you can access the screens that control DCD cluster performance.
3. Note the status icon color for the cluster. If it is not green yet, wait until it turns to green before proceeding. When the icon turns green it indicates that the primary BIG-IQ DCD cluster is functioning again.
4. Analyze the data collection device cluster details listed in the Summary area, and make sure that the values after upgrade match the values from the health check you did before the upgrade. At a minimum, you should confirm that the following values are correct:
 - Data Node Count

Reestablish the HA configuration after upgrading to BIG-IQ version 5.4

After you upgrade both F5® BIG-IQ® Centralized Management systems in a HA configuration, you can re-associate the secondary system with the primary BIG-IQ system.

1. Log in to primary BIG-IQ system with your administrator user name and password.
2. At the top of the screen, click **System**.
3. On the left, click **BIG-IQ HA**.
4. Click the **Add Secondary** button.
5. In the **IP Address** field, type the discovery address you specified on the BIG-IQ system during setup. This is the same IP address the peers in a high availability confirmation use to communicate.
6. In the **User name** and **Password** fields, type the administrative user name and password for the system.
7. In the **Root Password** field, type the root password for the system.
8. Click the **Add** button to add this device to this high availability configuration.

Even though you can log in to the secondary BIG-IQ after the you reestablish the HA configuration, the system continues some database re-indexing processes in the background. For larger configurations, that can take up to an hour. If you perform any searches on objects before it's done re-indexing, BIG-IQ might not return the expected results. During this time, you can use the primary BIG-IQ.

Next, you should verify that both BIG-IQ systems have the same configuration.

Restart snapshot creation

If you stopped creating snapshots before upgrading, now that the upgrade is complete, you need to restart the snapshots to protect your data.

1. Use SSH to log in to the primary BIG-IQ system for this cluster.
You must log in as `root` to perform this procedure.
2. Retrieve the list of scheduled snapshots by running the following command: `restcurl cm/shared/esmgmt/es-snapshot-task | grep task-scheduler`

```
config # restcurl cm/shared/esmgmt/es-snapshot-task | grep task-scheduler
"link": "https://localhost/mgmt/shared/task-scheduler/scheduler/0fdf50ec-8a17-3da9-b717-c63637ccc68a"
"link": "https://localhost/mgmt/shared/task-scheduler/scheduler/0af33352-2f33-32b3-85cb-1281bb88c249"
"link": "https://localhost/mgmt/shared/task-scheduler/scheduler/2ad770a8-bdb0-3383-99a9-300846eb0972"
```

In this example, there are 3 snapshots scheduled.

3. Start each of the schedules by running the following command: `restcurl -X PATCH -d '{"status":"ENABLED"}' shared/task-scheduler/scheduler/<SNAPSHOT ID>`

```
#restcurl -X PATCH -d '{"status":"ENABLED"}'
shared/task-scheduler/scheduler/0af33352-2f33-32b3-85cb-1281bb88c249
```

After you run the command for each scheduled snapshot, snapshots will again be created following the snapshot schedule.

Re-discover devices and re-import services

After you upgrade F5® BIG-IQ Centralized Management, you must rediscover your managed devices and reimport their services for AFM, ASM, DNS, and LTM so you can start managing those devices with the new features introduced in this release.

1. At the top of the screen, click **Devices**.
2. Select the check box next to the devices you want to rediscover and reimport services for.
3. Click the **More** button and select **Re-discover and Re-import**.
4. In the **Name** and **Description** fields, type a name and an optional description to identify this task.
5. For the Conflict Resolution Policy setting, select an option for how you want to handle any differences found between the configuration on BIG-IQ and the BIG-IP device during the re-discover and re-import task.
 - **Use BIG-IQ** to use the configuration setting stored on BIG-IQ if any differences are found.
 - **Use BIG-IP** to override the configuration settings stored on BIG-IQ with the settings from the BIG-IP device.

Important:

Some new features are introduced with each BIG-IQ release so it's a good idea to use the BIG-IP device's configuration after you upgrade BIG-IQ. This ensures you don't inadvertently overwrite a configuration that wasn't previously supported. BIG-IP devices are re-imported in the order listed, from top to bottom. You can use the arrow keys to change the processing order.

When you select **Use BIG-IP** to resolve conflicts, the BIG-IP device used to resolve those conflicts should appear last in the re-import list. If two or more BIG-IP devices contain the same object with different values, only the value in the last imported BIG-IP is used to resolve the conflict for all the BIG-IP devices.

6. If you want to save a snapshot of the BIG-IP device's configuration before importing their services, select the **Create a snapshot of the current configuration before importing** check box.
7. Click the **Create** button at the bottom of the screen.

Troubleshooting Minimal Downtime Upgrade Issues

What do I do if there is not enough disk space?

If there is not enough disk space to install the new software, you need to extend the `/var` partition. The default size of the `/var` file system in a newly installed node is 10 GB. This volume size might be insufficient to store your data. You can see how to extend this file system to a larger size in knowledge article K16103. Refer to: *K16103: Extending disk space on BIG-IQ Virtual Edition at support.f5.com/csp/article/K16103*. Because upgrading a node requires at least two volumes, you must ensure that both volumes can have their `/var` file system extended to the same size, or upgrades might fail.

Symptom

If the message: `UCS restore failure` displays during software installation, it may well be due to insufficient disk space.

To determine if this is the case:

1. Log in to the BIG-IQ system or DCD on which the software is failing to install.
2. Run the command: `tmsh show sys software`.

The system displays `failed (UCS application failed; unknown cause in response)`.

3. Navigate to the `liveinstall.log` file in the `/var/log/` folder. If the issue triggering the error message is insufficient disk space, then the file contains the following error message:

```
info: capture: status 256 returned by command: F5_INSTALL_MODE=install
F5_INSTALL_SESSION_TYPE=hotfix chroot
/mnt/tm_install/9934.NdHXAL /usr/local/bin/im -force /var/local/ucs/config.ucs
info: >++++ result:
info: Extracting manifest: /var/local/ucs/config.ucs
info: /var: Not enough free space
info: 3404179456 bytes required info: 2206740480 bytes available
info: /var/local/ucs/config.ucs: Not enough free disk space to install!
info: Operation aborted.
info: >----
info: Removing boot loader reference

Terminal error: UCS application failed; unknown cause.
*** Live install end at 2017/03/21 17:00:54: failed (return code 2) ***
```

Recommended Actions

1. Switch back to the older (pre-upgrade) volume.
 - a. Log in to the BIG-IQ system or DCD on which the software is failing to install.
 - b. Change to the pre-upgrade volume by running the command `switchboot -b <old-volume-name>`.
 - c. Restart the device by running the command `reboot`.
2. Clean up the audit log entries and snapshot objects.
3. Retry the installation.

Data collection device cluster status is yellow

After upgrading a data collection device (DCD), if the cluster status is yellow (unhealthy) instead of green (healthy), there are a number of potential causes and corresponding corrective actions you can attempt to resolve the issue.

Connectivity Issues

There could be unassigned replica shards in the cluster. Network connectivity issues can cause relocation of shards to a newly upgraded DCD to fail.

Symptoms

1. Log in to the primary BIG-IQ DCD for the cluster.
2. Navigate to the `/var/log/elasticsearch/` directory.
3. Examine the `eslognode.log` file, to see if there is an add/remove/add cycle due to temporary network connectivity issues.

If there are unassigned replica shards, you will find a pattern similar to the following:

```
[2017-05-03 10:49:20,885][INFO ][cluster.service ]
[f992a8aa-49c8-47ba-a59a-2be863f3a042] added {{7c49c2ef-ad2f-41f5-ab15-e01084b20364}
{_xdtat15RGSn-pIC192ORA}{10.10.10.2}{10.10.10.2:9300}{zone=Seattle, master=true}},
reason: zen-disco-join(join from node[{{7c49c2ef-ad2f-41f5-ab15-e01084b20364}
{_xdtat15RGSn-pIC192ORA}{10.10.10.2}{10.10.10.2:9300}{zone=Seattle, master=true}}])
[2017-05-03 10:49:32,172][INFO ][cluster.service ] [f992a8aa-49c8-47ba-a59a-2be863f3a042]
removed {{7c49c2ef-ad2f-41f5-ab15-e01084b20364}{_xdtat15RGSn-pIC192ORA}{10.10.10.2}
{10.10.10.2:9300}{zone=Seattle, master=true}}, reason: zen-disco-node_failed(
{{7c49c2ef-ad2f-41f5-ab15-e01084b20364}{_xdtat15RGSn-pIC192ORA}{10.10.10.2}{10.10.10.2:9300}
{zone=Seattle, master=true}}, reason transport disconnected
[2017-05-03 10:49:36,210][INFO ][cluster.service ] [f992a8aa-49c8-47ba-a59a-2be863f3a042]
added
{{7c49c2ef-ad2f-41f5-ab15-e01084b20364}{_xdtat15RGSn-pIC192ORA}{10.10.10.2}{10.10.10.2:9300}
{zone=Seattle, master=true}}, reason: zen-disco-join(join from node[{{7c49c2ef-ad2f-41f5-
ab15-e01084b20364}
{_xdtat15RGSn-pIC192ORA}{10.10.10.2}{10.10.10.2:9300}{zone=Seattle, master=true}}])
```

Note: The sample log file above illustrates the added/removed/added cycle. Note the presence of the word added on line 2, followed by removed on line 7, and then added again on line 11.

Recommended Actions

1. Log in to the DCD that you just upgraded.
 2. Restart the elastic search service by running the command: `bigstart restart elasticsearch`.
- Shard assignment can now succeed, so the cluster status should change to healthy (green).

Index created in higher version

When a DCD is upgraded, ElasticSearch begins creating indexes on that DCD. Because the upgraded DCD is running a different software version than the other DCDs in the cluster, shards cannot be replicated between the DCDs. This can result in unassigned replica shards in the cluster.

Symptoms

1. Log in to the upgraded DCD.

2. Check for unassigned replica shards by running the following command:

```
curl -s localhost:9200/_cat/shards?v | grep UNASSIGNED
```

3. If there are no unassigned shards, there should be no response to the command. If there are unassigned shards, try the recommended actions.

Recommended Actions

- For a single zone DCD cluster, continue the upgrade process by upgrading another DCD in the cluster. When another DCD in the cluster is running the same version, shards can begin replicating again and the cluster status should become healthy.
- For a multiple zone DCD cluster, continue the upgrade process but upgrade the next DCD in a different zone than the first DCD. When there are DCDs running the same version in different zones, shards can begin replicating again and the cluster status should become healthy.

Data collection device cluster status is red

After upgrading a data collection device (DCD), if the cluster status is red (unhealthy) instead of green (healthy), there are a number of potential causes and corresponding corrective actions you can attempt to resolve the issue.

Statistics replicas are not enabled

If statistics replicas were not enabled before you upgraded, the cluster will not create replicas of your data and the cluster health will be unhealthy.

Recommended Actions

1. Log in to the primary BIG-IQ for the DCD cluster.
2. Navigate to the Statistics Retention Policy screen and expand the Advanced Settings, then select **Enable Replicas**.
3. Login to the most recently upgraded DCD.
4. Change to the pre-upgrade volume by running the command `switchboot -b <old-volume-name>`.
5. Reboot the DCD by running the command `reboot`.
6. Wait for the rebooted DCD to join the cluster.
7. Wait for the cluster status to return to green (indicating that the cluster has successfully replicated your data shards.)
8. Repeat the upgrade for the DCDs in the cluster.

After the upgrade, the cluster status should change to healthy (green).

Generic Failure

Sometimes for no discernable reason, the DCD cluster fails to assign the primary data shards.

Symptoms

There are no symptoms to confirm this other than the cluster status unhealthy (red). However, if you have tried other corrective actions and the problem persists, you can try this remedy to see if it solves the problem.

Recommended Actions

- Log in to the primary BIG-IQ for the DCD cluster.
- Change to the pre-upgrade volume by running the command `switchboot -b <old-volume-name>`.
- Wait for the cluster status to return to green (indicating that the cluster has successfully replicated your data shards.)
- Check to see that replicas exist for each primary shard.

1. Log in to the BIG-IQ system or DCD on which the software is failing to install.
2. Check for replicas of each primary by running the command `curl -s localhost:9200/_cat/shards?v`. A typical response would look like this:

Note: In the following sample response, note that each primary index (designated with a `p` in the `docs` column) has a corresponding replica (designated with an `r` in the `docs` column), and that the replica exists on a node with an IP address that is different than the node that the primary index is on.

index	shard	prirep	state	docs	store	ip
node						
websafe_2017-07-21t00-00-00-0700	da66-4587-bb70-5d2dbc254a05	2	r	STARTED	0	159b 10.145.193.11 f7b33853-
websafe_2017-07-21t00-00-00-0700	687a7b3b-7dc3-4074-9b22-1e26c6092877	2	p	STARTED	0	159b 10.145.192.202
websafe_2017-07-21t00-00-00-0700	da66-4587-bb70-5d2dbc254a05	1	r	STARTED	0	159b 10.145.193.11 f7b33853-
websafe_2017-07-21t00-00-00-0700	687a7b3b-7dc3-4074-9b22-1e26c6092877	1	p	STARTED	0	159b 10.145.192.202
websafe_2017-07-21t00-00-00-0700	da66-4587-bb70-5d2dbc254a05	3	r	STARTED	0	159b 10.145.193.11 f7b33853-
websafe_2017-07-21t00-00-00-0700	687a7b3b-7dc3-4074-9b22-1e26c6092877	3	p	STARTED	0	159b 10.145.192.202
websafe_2017-07-21t00-00-00-0700	da66-4587-bb70-5d2dbc254a05	4	r	STARTED	0	159b 10.145.193.11 f7b33853-
websafe_2017-07-21t00-00-00-0700	687a7b3b-7dc3-4074-9b22-1e26c6092877	4	p	STARTED	0	159b 10.145.192.202
websafe_2017-07-21t00-00-00-0700	da66-4587-bb70-5d2dbc254a05	0	r	STARTED	0	159b 10.145.193.11 f7b33853-
websafe_2017-07-21t00-00-00-0700	687a7b3b-7dc3-4074-9b22-1e26c6092877	0	p	STARTED	0	159b 10.145.192.202

3. If there are replicas for each primary shard, repeat the upgrade for the DCDs in the cluster.

After you upgrade the DCDs in the cluster, the cluster status should change to healthy (green).

Data collection device cluster is offline

After upgrading a data collection device (DCD), if the cluster status is completely offline, there is one primary potential cause and corresponding corrective action you can attempt to resolve the issue.

Election of new master node failed

When the DCD cluster's master node is rebooted, a new master must be elected. Sometimes that election can fail, which causes the cluster to be offline.

Symptoms

There are a number of different symptoms that can indicate that the master election has failed. Master election failure can only occur when three conditions are met:

- You are upgrading from version 5.2.0 or 5.3.0 to version 5.4.0.
- The DCD that was upgraded and rebooted was the master node before the upgrade.
- Statistics replicas were enabled shortly before the upgrade and reboot.

One symptom of an election failure is when an API call to the cluster responds with a `master not discovered` message.

1. Use SSH to log in to the primary BIG-IQ for the DCD cluster.
2. Check the cluster status by submitting the following API call: `curl http://localhost:9200/_cat/nodes?v.`
3. If the response to the API call is similar to the following, the master election failed.

```
{ "error": { "root_cause":
[{"type": "master_not_discovered_exception", "reason": null}], "type": "master_not_discovered_exception", "reason": null }, "status": 503 } $
```

Another symptom of an election failure is when the ElasticSearch log file contains error messages.

1. Use SSH to log in to the primary BIG-IQ for the DCD cluster.
2. Navigate to the `/var/log/elasticsearch/` folder and open the log file `eslognode.log`.
3. Examine the log file. The following text snippets are examples of the two error messages that signal an election failure. If either of these messages are in the log file, the master election failed.

```
[2017-06-01 16:13:42,792][ERROR][discovery.zen ] [eb1d873a-ffdb-4ae8-ae22-946554970c54]
unexpected failure during [zen-disco-join(elected_as_master,
[4] joins received)] RoutingValidationException[[Index [statistics_t10_device_2017-152-21]:
Shard [2] routing table has wrong number of replicas,
expected [0], got [1], Index [statistics_t10_device_2017-152-21]: Shard [1] routing table
has wrong number of replicas, expected [0], got [1], Index
[statistics_t10_device_2017-152-21]: Shard [4] routing table has wrong number of replicas,
expected [0], got [1], Index [statistics_t10_device_2017-152-21]:
Shard [3] routing table has wrong number of replicas, expected [0], got [1], Index
[statistics_t10_device_2017-152-21]: Shard [0] routing table has
wrong number of replicas, expected [0], got [1]]]
```

```
[2017-06-01 16:14:48,992][INFO ][discovery.zen ] [163e016c-3827-496c-b306-b2972d60c8df]
failed to send join request to master [{eb1d873a-ffdb-4ae8-ae22-946554970c54}
{IyzDax59Swy6pNqg2YU0-Q}{10.145.192.147}{10.145.192.147:9300}{data=false, zone=L,
master=true}], reason [RemoteTransportException[[eb1d873a-ffdb-4ae8-ae22-946554970c54]
[10.145.192.147:9300][internal:discovery/zen/join]]; nested: IllegalStateException[Node
[{eb1d873a-ffdb-4ae8-ae22-946554970c54}{IyzDax59Swy6pNqg2YU0-Q}{10.145.192.147}
{10.145.192.147:9300}{data=false, zone=L, master=true}] not master for join request]; ]
[2017-06-01 16:14:50,124][WARN ][rest.suppressed ] path: /_bulk, params:
{}ClusterBlockException[blocked by: [SERVICE_UNAVAILABLE/2/no master];]
```

Recommended Actions

Repeat these steps for each DCD in the cluster.

1. Use SSH to log in to a DCD in the cluster.
2. Restart the ElasticSearch service by running the command: `bigstart restart elasticsearch.`

Restarting the service for each DCD in the cluster triggers a new master node election. After the last DCD in the cluster is restarted, the cluster status should change to healthy (green).

Legal Notices

Legal notices

Publication Date

This document was published on February 16, 2018.

Publication Number

MAN-0676-01

Copyright

Copyright © 2018, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

For a current list of F5 trademarks and service marks, see <http://www.f5.com/about/guidelines-policies/trademarks>.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at: <https://f5.com/about-us/policies/patents>.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Legal Notices

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Index

B

- BIG-IP logging profile
 - checking configuration 9
- BIG-IQ
 - before you upgrade 13, 19
- BIG-IQ HA pair
 - re-establishing configuration 25
 - upgrading 19
- BIG-IQ latest version
 - about upgrading data collection devices 13
- BIG-IQ upgrade
 - bringing DCD back online 17
 - confirm DCD in service 16
 - taking DCD offline 14

C

- cluster configuration
 - checking data collection device 25
 - checking for data collection device 9
- Connectivity Issues 30

D

- data collection device
 - checking health 25
 - health 25
 - upgrading to latest version 15
- data collection device cluster
 - upgrade without downtime 9
- Data collection device cluster is offline 32
- data collection device upgrade
 - about 6
 - prerequisites 8
- DCD upgrade
 - about 6
- devices
 - re-discover from the user interface after upgrading 26
- disk space
 - confirming sufficient for upgrade 10, 11
- disk space insufficient 29–31

E

- Election of new master node failed 32

G

- Generic Failure 31

H

- HA pair
 - re-establishing BIG-IQ HA configuration 25

I

- Index created in higher version 30

L

- logging profile
 - checking data collection device 9
 - checking health 9

P

- primary BIG-IQ in HA pair
 - upgrading from BIG-IQ version 5.2 to version 5.4 22

R

- rediscover devices
 - from the user interface 26
- reimport services
 - from the user interface 26
- retention policy
 - checking 9

S

- secondary BIG-IQ in an HA pair to latest version
 - upgrading 20
- services
 - re-import from the user interface after upgrading 26
- shards
 - checking policy 9
- snapshots
 - starting for upgrade without down time 26
 - stopping for upgrade without down time 11
- software
 - uploading to BIG-IQ 13, 20, 21
- statistics data
 - managing replicas 9
- Statistics replicas are not enabled 31
- sufficient disk space
 - confirm for upgrade without down time 10
 - confirming for upgrade without down time 10, 11

T

- the latest version
 - before you upgrade to 13

U

- update requirements
 - for data collection device 8
- update requirements worksheet 8
- upgrade prerequisites 13, 19
- upgrade process
 - for data collection device cluster 6

Index

upgrade process (*continued*)
 for DCD cluster 6
 upgrading from BIG-IQ version 5.2 to version 5.4 22
upgrade without downtime
 data collection device cluster 9
upgrade workflow options 5, 6