Planning and Implementing an F5® BIG-IQ® Centralized Management Deployment

Version 5.2

Table of Contents

Planning and Implementing a Centralized Management Deployment

What elements make up a Centralized Management deployment?

A F5® BIG-IQ® Centralized Management solution involves a number of components. This diagram illustrates those that make up the BIG-IQ® Centralized Management solution.

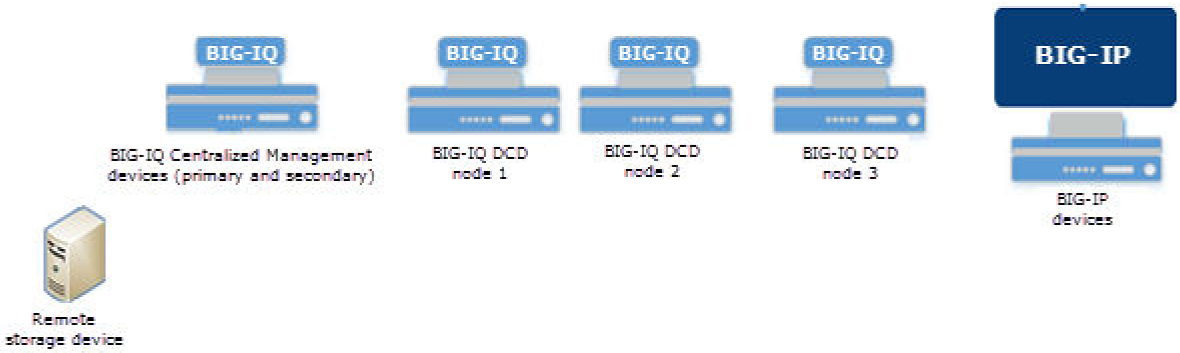

Figure 1: Centralized Management network topology

BIG-IQ Centralized Management device

F5® BIG-IQ® Centralized Management is a platform that helps you manage BIG-IP® devices, and all of their services (such as LTM®, AFM™, or ASM®), from one location. That means you and your co-workers don't have to log in to individual BIG-IP systems to get your job done. Using BIG-IQ Centralized Management, you can centrally manage your BIG-IP devices, performing operations such as backups, licensing, monitoring, and configuration management. And because access to each area of BIG-IQ is role based, you can limit access to users maximizing work flows while minimizing errors and potential security issues.

BIG-IQ data collection device

A *BIG-IQ data collection device* (DCD) is a specially-provisioned BIG-IQ® system that you use to manage and store alerts, events, and statistical data from one or more BIG-IP® systems.

Configuration tasks on the BIG-IP system determine when and how alerts or events are triggered on the client. The alerts or events are sent to a BIG-IQ data collection device, and the BIG-IQ system retrieves them for your analysis. When you opt to collect statistical data from the BIG-IP devices, the DCD periodically retrieves those statistics from your devices, and then processes and stores that data.

The group of data collection devices and BIG-IQ systems that work together to store and manage your data are referred to as the *data collection cluster*. The individual data collection devices are generally referred to as *nodes*.

BIG-IP device

A BIG-IP device runs a number of licensed components designed around application availability, access control, and security solutions. These components run on top ofF5® TMOS®. This custom operating system is an event driven operating system designed specifically to inspect network and application traffic and make real-time decisions based on the configurations you provide. The BIG-IP software runs on both hardware and virtualized environments.

Remote storage device

The remote storage device is necessary only when your deployment includes a data collection device (DCD) and you plan to store backups of your events, alerts, and statistical data for disaster recovery requirements.

Network Requirements for a BIG-IQ Centralized Management Deployment

Before you deploy a Centralized Management solution

Before you begin to deploy a BIG-IQ® system, you should complete these preparations.

- Determine the deployment scenario that works best for your needs.
- Create the interfaces, communications, and networks needed to support your deployment scenario
- Configure your network (including switches and firewalls) to permit BIG-IQ network traffic to flow based on the deployment scenario you choose.
- Assemble the passwords, IP addresses, and licensing information needed for the BIG-IQ cluster components.

Planning for a Centralized Management deployment

To successfully deploy a BIG-IQ® Centralized Management solution, you may need to coordinate with several people in your company.

If you use BIG-IQ virtual editions, you might need to coordinate with the people who manage your virtual environment, so they can provision the virtual machines with the required amount of CPUs, memory, and network interfaces. Further, you’ll need to coordinate with the people who manage the storage for the virtual machines to make sure each virtual machine is provisioned with the necessary storage to support the BIG-IQ environment. You also might need to provide the virtual environment team a copy of the BIG-IQ virtual machine image (available from *https://downloads.f5.com*), depending how they operate.

If you use BIG-IQ 7000 devices in your network, you need to coordinate with the people who manage the data center where the BIG-IQ devices are housed to make arrangements for the devices to be racked, powered on, and connected to your network.

There are also several tasks to coordinate with your networking team:

- IP address allocation for the BIG-IQ nodes, depending on your deployment model.
- Creation of networks, VLANs, and so on dependent on your deployment model.
- Any routing configuration required to ensure traffic passes between the BIG-IQ nodes and the BIG-IP devices.
- Additional networking configuration required to support the BIG-IQ system's operation.

Finally, you may need to coordinate with your network firewall administrators, depending on the network configuration at your company. The BIG-IQ software needs to communicate between BIG-IQ nodes and BIG-IP systems; and, if there are firewalls in the network path, firewall rules probably need to be configured to permit that traffic. For additional detail about required network ports and protocols, refer to *Open ports required for data collection device cluster deployment*.

Determining the network configuration needed for your deployment

There are three common deployment scenarios for the BIG-IQ® system. The scenario most appropriate for you depends on what you want to do.

Table 1: BIG-IQ deployment options

Which deployment type should you choose?	What functions does your deployment need to perform?	Which hardware components and networks do you need?
Simple management and configuration.	Manage and configure the BIG-IP® devices. For example, taking backups, licensing virtual editions, and configuring local traffic and security policies.	All you need is one or more BIG-IQ system and the BIG-IP devices you want to manage. This configuration uses a single management network.
Advanced management and configuration.	Manage and configure the BIG-IP devices. For example, taking backups, licensing virtual editions, and configuring local traffic and security policies. Collect and view Local Traffic, DNS, and Device statistical data from the BIG-IP devices. Collect, manage, and view events and alerts from BIG-IP devices provisioned with the APM®, FPS, or ASM® components.	You need BIG-IQ systems, data collection devices, and an external storage device. This configuration needs a single, management network and an internal BIG-IQ cluster network.
Large-scale, distributed management and configuration	Manage and configure the BIG-IP devices. For example, taking backups, licensing virtual editions, and configuring local traffic and security policies. Collect and view Local Traffic, DNS, and Device statistical data from the BIG-IP devices. Collect, manage, and view events and alerts from BIG-IP devices provisioned with the APM, FPS, or ASM components. Separate network traffic to support large, distributed deployments of the F5 BIG-IQ Centralized Management solution for improved performance, security, and interactions in multiple data center environments.	You need BIG-IQ systems, data collection devices, and an external storage device. This configuration needs an external network, a management network, and an internal network.

Network environment for simple management and configuration

To deploy this configuration, all you need is one or more BIG-IQ® systems and the BIG-IP® devices you want to manage. The number of BIG-IQ systems you need depends on how much redundancy your business requires. A second system provides high availability failover capability. Or, for disaster recovery capability, you could operate multiple data centers, each with its own set of BIG-IQ systems. (For additional detail, refer to *Managing Disaster Recovery Scenarios*.)

The simple management and configuration uses a single management network. Traffic on the management network is used to do the following:

- Provide communication between the BIG-IQ system and DCD nodes.
- Enable bidirectional traffic between the BIG-IQ systems and the BIG-IP devices.
- If you use a secondary high availability BIG-IQ system, enable traffic between the BIG-IQ systems. This traffic keeps the state information synchronized on your BIG-IQ systems.
- Provide access the BIG-IQ user interface. You can also use it to access the BIG-IQ system using SSH if you need to run manual commands.

Note: *The number of devices of each type that will best meet your company's needs depends on a number of factors. Refer to the Data collection device sizing guidelines for additional detail.*

The figure illustrates the network topology required for a simple management and configuration deployment.

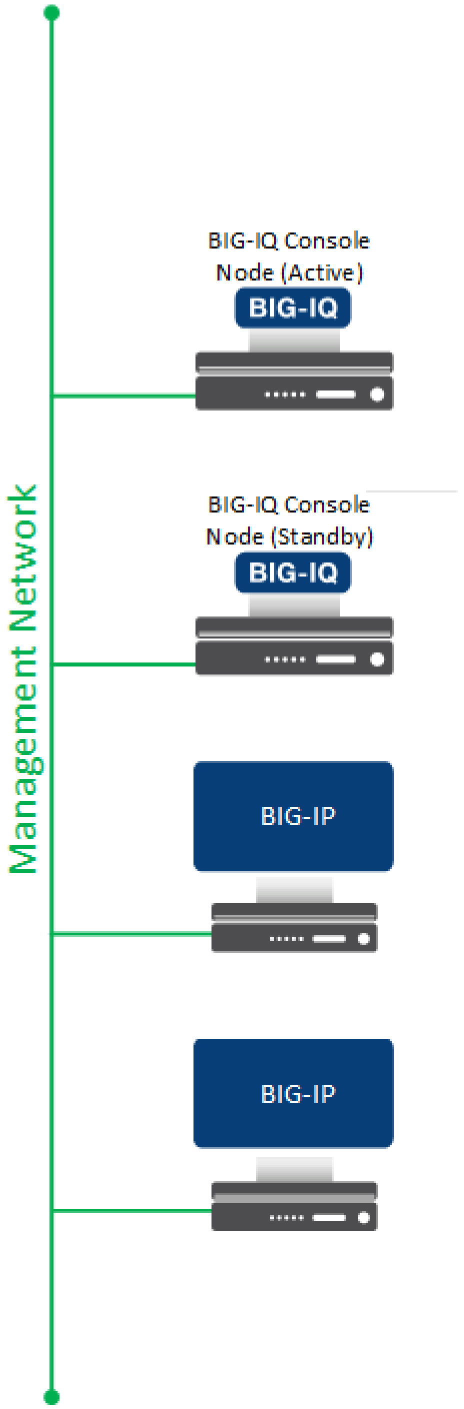

Figure 2: Centralized management network topology

Use the table to record the IP addresses for the devices in the BIG-IQ deployment.

Device Type	Management IP address(es)
Primary BIG-IQ system	
Secondary BIG-IQ system	
BIG-IP devices	

Network environment for advanced management and configuration

To deploy this configuration, you need BIG-IQ® systems, Data Collection Devices, and an external storage device. This configuration needs a single management network, an internal BIG-IQ cluster network, and an optional external storage device for backing up alert, event, and statistical data..

Note: *The number of devices of each type that will best meet your company's needs depends on a number of factors. Refer to the Data collection device sizing guidelines for additional detail.*

Traffic on the management network is used to do the following:

- Provide communication between the BIG-IQ system and DCD nodes.
- Enable bidirectional traffic between the BIG-IQ systems and the BIG-IP devices.
- If you use a secondary high availability BIG-IQ system, enable traffic between the BIG-IQ systems. This traffic keeps the state information synchronized on your BIG-IQ systems.
- Provide access the BIG-IQ user interface. You can also use it to access the BIG-IQ system using SSH if you need to run manual commands.

The internal network is used to replicate data to maintain the BIG-IQ Centralized Management cluster.

Note: *It is best practice to isolate the traffic between BIG-IQ cluster nodes for performance and improved security.*

This figure illustrates the network topology required for an advanced management and configuration deployment.

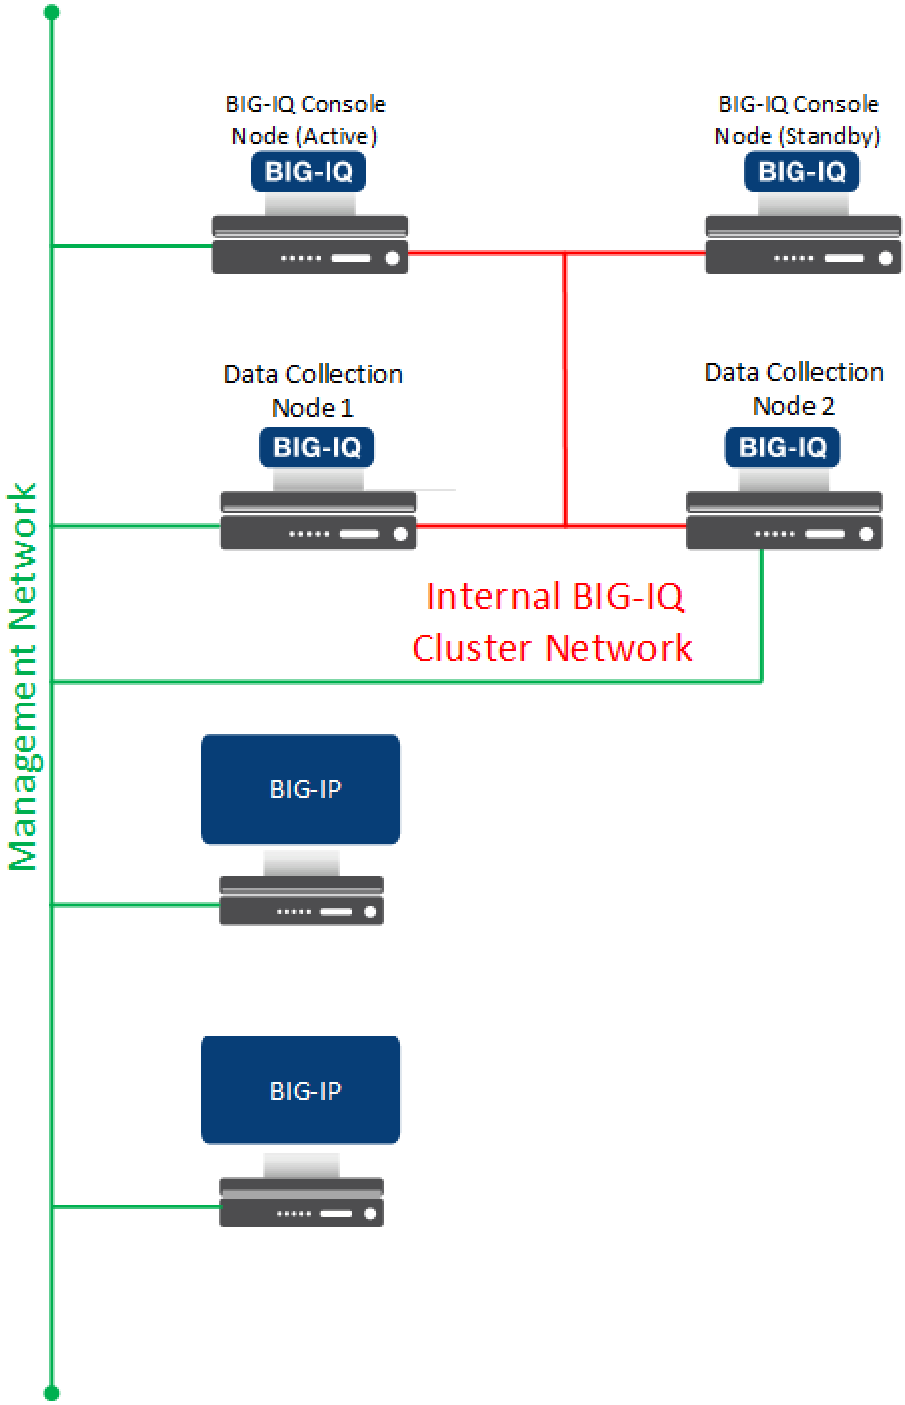

Figure 3: Centralized management and enhanced monitoring network topology

Use the table to record the IP addresses for the devices in the BIG-IQ deployment.

Device Type	**Management IP addresses**	**Internal Network IP addresses**
Primary BIG-IQ system		
Secondary BIG-IQ system		

Device Type	Management IP addresses	Internal Network IP addresses
Data Collection Device management IP addresses		
BIG-IP devices		
Remote storage device		

Network environment for large-scale, distributed management and configuration

To deploy this configuration, you need BIG-IQ® systems, Data Collection Devices, and an external storage device. This configuration needs an external network, a management network, an internal network, and an optional external storage device for backing up alert, event, and statistical data.

The external network routes traffic between the BIG-IQ Centralized Management cluster and the BIG-IP® devices.

The internal network is used to replicate data to maintain the BIG-IQ Centralized Management cluster.

***Note:** It is best practice to isolate the traffic between BIG-IQ cluster nodes for performance and improved security.*

Traffic on the management network is used to do the following:

- Provide communication between the BIG-IQ system and DCD nodes.
- Enable bidirectional traffic between the BIG-IQ systems and the BIG-IP devices.
- If you use a secondary high availability BIG-IQ system, enable traffic between the BIG-IQ systems. This traffic keeps the state information synchronized on your BIG-IQ systems.
- Provide access the BIG-IQ user interface. You can also use it to access the BIG-IQ system using SSH if you need to run manual commands.

This figure illustrates the network topology required for this deployment.

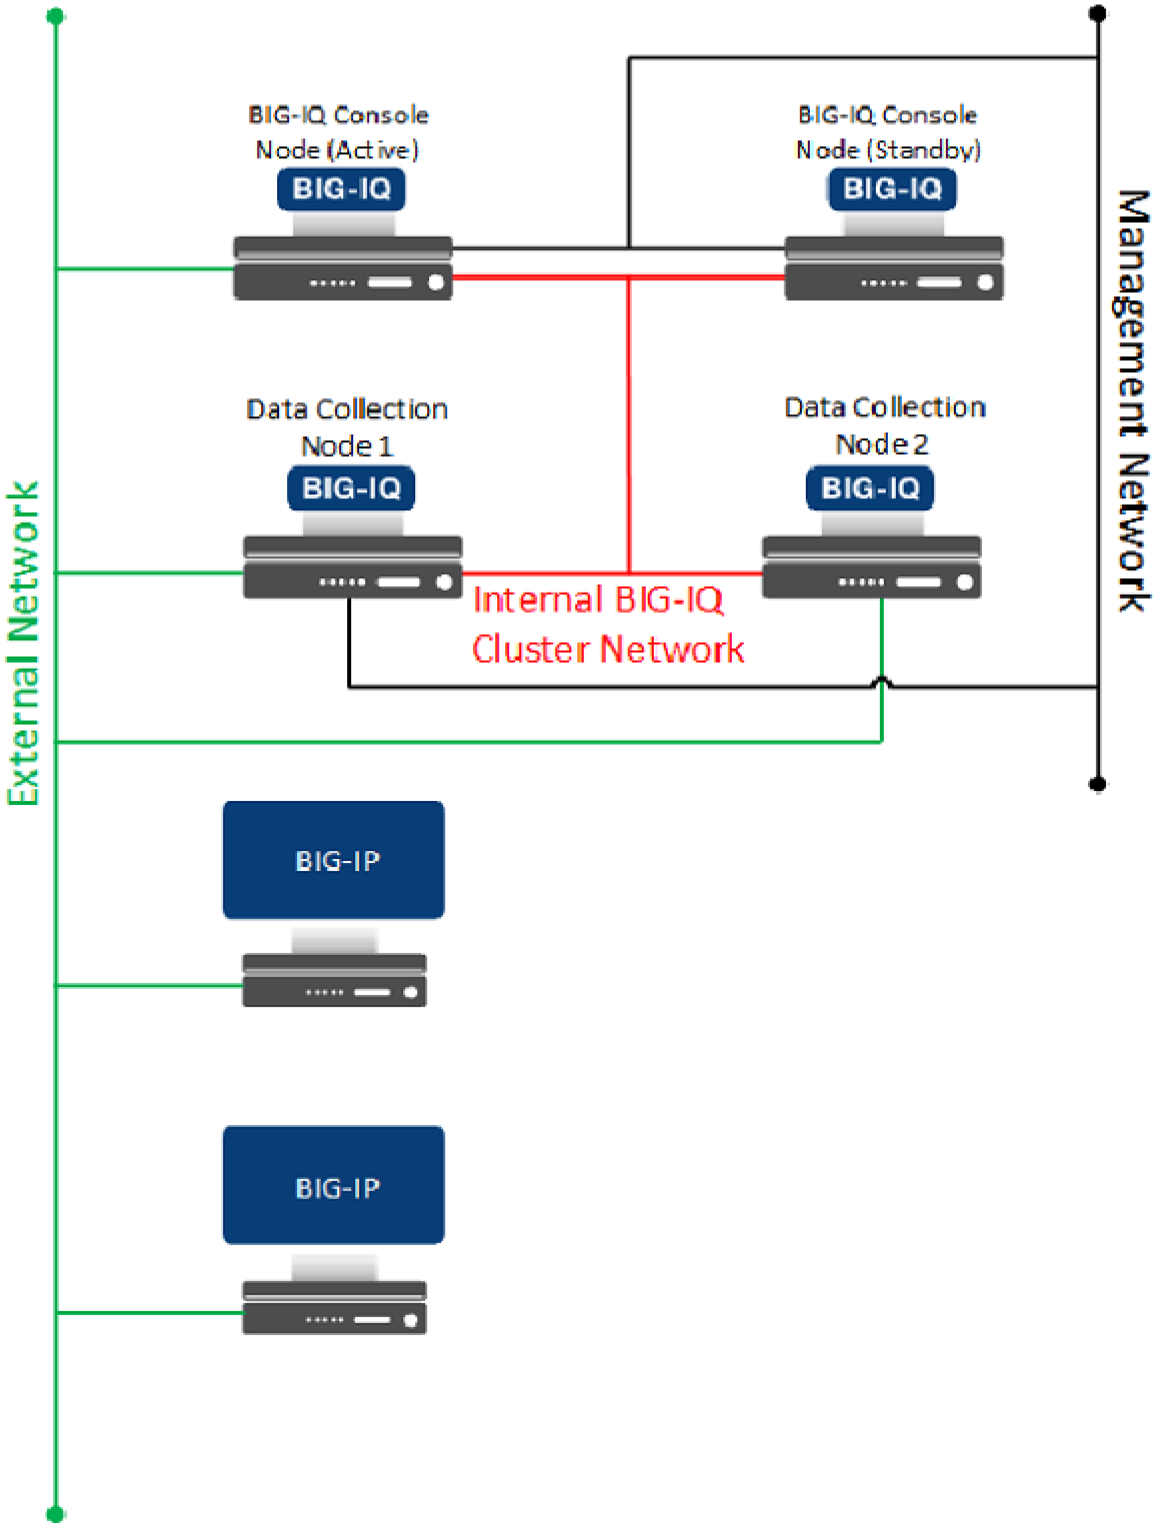

Figure 4: Centralized management, enhanced monitoring, and improved performance network topology

Use the table to record the IP addresses for the devices in the BIG-IQ deployment.

Device Type	Management IP addresses	Internal Network IP addresses	External Network IP addresses
Primary BIG-IQ system			
Secondary BIG-IQ system			
Data Collection Device management IP addresses			
BIG-IP devices			
Remote storage device			

Open ports required for data collection device cluster deployment

The BIG-IQ® system must have bidirectional communication with the devices in your network to successfully manage them. The ports described in the table must be open to allow for this required two-way communication. You might have to contact a firewall or network administrator to verify that these ports are open, or to have them opened if they are not.

Table 2: Ports required for data collection device cluster deployment

Source IP Address	Destination IP Address	Destination Port	Protocol	Is port Configurable?	Is the Protocol Configurable?	Purpose	Connection Origination
Management IP address or external self IP address of the BIG-IQ console. *See table note 1.	Management IP address or self IP address of the BIG-IP device. *See table note 1.	443 (SSL) 22 (SSH) *See table note 4.	TCP	No	No	Device-level discovery, device configuration changes, and device operations (backup, licensing, and so on), health checking, and some statistics (For example, Access or ADC object status).	From BIG-IQ console to BIG-IP devices.
Management IP address or external self IP address of the BIG-IQ data collection device. *See table notes 1 and 2.	Management IP address or self IP address of the BIG-IP device. *See table note 1.	443 (SSL)	TCP	No	No	Statistics collection for Local Traffic, Device, and DNS objects.	From BIG-IQ data collection devices to BIG-IP devices.
Management IP address or internal self IP address of the BIG-IQ console. *See table note 1.	Management IP address or internal self IP address of the BIG-IQ console. *See table note 1.	443 (SSL)	TCP	No	No	BIG-IQ cluster synchronization and cluster maintenance.	From the active BIG-IQ console to the standby BIG-IQ console. From the BIG-IQ standby console to the BIG-IQ active console.
Management IP address or internal self IP address of the active BIG-IQ console. *See table note 1.	Management IP address or internal self IP address of the standby BIG-IQ console. *See table note 1.	27017	TCP	No	No	BIG-IQ high availability cluster data replication.	From the active BIG-IQ console to the standby BIG-IQ console. From the BIG-IQ standby console to the BIG-IQ active console.
Management IP address or internal self IP address of the BIG-IQ console and the data collection device. *See table notes 1 and 2.	Management IP address or internal self IP address of the BIG-IQ console and the data collection device. *See table notes 1 and 3.	9300	TCP	Yes	No	Internal node-to-node communication to maintain data consistency and replication across clusters when data collection nodes are used.	Full Mesh That is, all BIG-IQ console and data collection devices can originate a connection for this purpose.
Management IP address or self IP address of the BIG-IP device.	Management IP address or self IP address of the BIG-IQ data collection device.	8514	TCP	No	No	Logging profile communication for Web Application Security. This traffic uses the syslog protocol documented in RFC 5424.	From BIG-IP devices to BIG-IQ data collection devices. When you have multiple data collection devices, you need to make sure data

Source IP Address	Destination IP Address	Destination Port	Protocol	Is port Configurable?	Is the Protocol Configurable?	Purpose	Connection Origination
*See table note 1.	*See table notes 1 and 3.						can pass to all devices in the cluster.
Management IP address or self IP address of the BIG-IP device. *See table note 1.	Management IP address or self IP address of the BIG-IQ data collection device. *See table notes 1 and 3.	8008	TCP	No	No	Logging profile communication for Fraud Protection Service, this traffic uses the syslog protocol documented in RFC 5424.	From BIG-IP devices to BIG-IQ data collection devices. When you have multiple data collection devices, you need to make sure data can pass to all devices in the cluster.
Management IP address or self IP address of the BIG-IP device.	Management IP address or self IP address of the BIG-IQ data collection device. *See table notes 1 and 3.	9997	TCP	No	No	For access to events; this traffic uses the syslog protocol documented in RFC 5424.	From BIG-IP devices to BIG-IQ data collection devices. When you have multiple data collection devices, you need to make sure data can pass to all devices in the cluster.
Client IP address	BIG-IQ Management IP address or self IP address for all BIG-IQ instances in the cluster.	443 (SSL) 22 (SSH)	TCP	No	No	For management access to BIG-IQ GUI or API (port 443) or shell access to BIG-IQ (port 22)	From the client workstation to the BIG-IQ device.

Note: *1: Whether you use the management IP address or the self IP address depends on your network configuration.*

Note: *2: For clusters with multiple data collection devices, traffic must be able to originate from any device in the cluster.*

Note: *3: For clusters with multiple data collection devices, the destination can be any device in the cluster.*

Note: *Port 22 (SSH) is only required for BIG-IP® versions 11.5.0 to 11.6.0*

Passwords required for data collection device cluster deployment

To install and configure a data collection device cluster, you use the default passwords for all of the devices in your cluster. If (as recommended) you intend to schedule regular snapshots of your logging data, you need root access credentials for the machine on which you plan to store these snapshots.

Table 3: Passwords for data collection device cluster deployment

User Name	Default Password	Access Rights/Role
admin	admin	This user type can access all aspects of the BIG-IQ® system from the system's user interface.
root	default	This user has access to all aspects of the BIG-IQ system from the system's console command line.

Licenses required for data collection device cluster deployment

To install and configure a data collection device cluster, you need a license for each device.

Data collection device sizing guidelines

This material is being replaced with a new guide scheduled for inclusion in the next release.

Deploying a Data Collection Device

How do I deploy a data collection device cluster?

To manage the data generated by BIG-IP® devices on BIG-IQ® Centralized Management, you deploy a network of devices called a *data collection device (DCD) cluster*, and then configure that cluster to meet your business needs.

To deploy a DCD cluster, you should:

- Prepare your network environment
- Install the DCDs
- Discover and activate the DCDs
- Define an external location to store snapshots
- Enable data collection for the DCD cluster (or configure a BIG-IP system to send alerts or events to the cluster)
- Configure the BIG-IQ console that manages the DCD cluster for HA, if needed.

License basic setup for a data collection device?

The BIG-IQ® data collection device runs as a virtual machine in supported hypervisors, or on the BIG-IQ 7000 series platform. You license the data collection device using the base registration key you purchased. The *base registration key* is a character string that the F5 license server uses to provide access to data collection device features.

You license data collection device in one of the following ways:

- If the system has access to the internet, you can have the data collection device contact the F5 license server and automatically activate the license.
- If the system is not connected to the internet, you can manually retrieve the activation key from a system that is connected to the internet, and transfer it to the data collection device.
- If your data collection device is in a closed-circuit network (CCN) that does not allow you to export any encrypted information, you must open a case with F5 support.

When you license the data collection device, you:

- Specify a host name for the system.
- Assign a management port IP address.
- Specify the IP address of your DNS server and the name of the DNS search domain.
- Specify the IP address of your Network Time Protocol (NTP) servers and select a time zone.
- Change the administrator’s default admin and root passwords.

Automatically license BIG-IQ and perform initial setup

You must have a base registration key before you can license the BIG-IQ® system. If you do not have a base registration key, contact the F5 Networks sales group (`f5.com`).

If the BIG-IQ® system is connected to the public internet, you can follow these steps to automatically perform the initial license activation and perform the initial setup.

1. Use a browser to log in to BIG-IQ by typing `https://`*`<management_IP_address>`*, where *`<management_IP_address>`* is the address you specified for device management.
2. Click **Activate**.
3. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.

***Important:** If you are setting up a data collection device, you have to use a registration key that supports a data collection device license.*

4. In the **Add-On Keys** field, paste any additional license key you have.
5. To add another additional add-on key, click the + sign and paste the additional key in the new **Add-On Keys** field.
6. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button.
7. In the **Hostname** field, type a fully-qualified domain name (FQDN) for the system.

   You cannot change this name after you add it. The FQDN can consist of letters and numbers, as well as the characters underscore (_), dash (-), or period (.).
8. In the **Management Port IP Address** and **Management Port Route** fields, type the IP address for the management port IP address and route.

   ***Note:** The management port IP address must be in Classless Inter-Domain Routing (CIDR) format. For example: `10.10.10.10/24`.*
9. Specify what you want the BIG-IQ to use for the **Discovery Address**.
   - To use the management port, select **Use Management Address**.
   - To use the internal self IP address, select **Self IP Address**, and type the IP address.

     ***Important:** If you are configuring a data collection device, you must use the internal self IP address.*

     ***Note:** The self IP address must be in Classless Inter-Domain Routing (CIDR) format. For example: `10.10.10.10/24`.*
10. In the **DNS Lookup Servers** field, type the IP address of your DNS server.

    You can click the **Test Connection** button to verify that BIG-IQ can reach that IP address.
11. In the **DNS Search Domains** field, type the name of your search domain.

    The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
12. In the **Time Servers** field, type the IP addresses of your Network Time Protocol (NTP) server.

    You can click the **Test Connection** button to verify that BIG-IQ can reach the IP address.
13. From the **Time Zone** list, select your local time zone.
14. In the **Old Password** fields, type the default admin and root passwords, and then type a new password in the **Password** and **Confirm Password** fields.

Manually license BIG-IQ and perform initial setup

You must have a base registration key before you can license the BIG-IQ® system. If you do not have a base registration key, contact the F5 Networks sales group (`f5.com`).

If the BIG-IQ® system is not connected to the public internet, use this procedure to manually activate the license and perform the initial setup.

1. Use a browser to log in to BIG-IQ by typing `https://<management_IP_address>`, where `<management_IP_address>` is the address you specified for device management.
2. Click **Activate**.
3. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.

   ***Important:** If you are setting up a data collection device, you have to use a registration key that supports a data collection device license.*
4. In the **Add-On Keys** field, paste any additional license key you have.

5. For the **Activation Method** setting, select **Manual** and click the **Generate Dossier** button.
   The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.
6. Select and copy the text displayed in the **Device Dossier** field.
7. Click the **Access F5 manual activation web portal** link.
   The Activate F5 Product site opens.
8. Into the **Enter your dossier** field, paste the dossier.

   Alternatively, if you saved the file, click the **Choose File** button and navigate to it.

   After a pause, the screen displays the license key text.
9. Click **Next**.

   If you are setting up this device for the first time, the Accept User Legal Agreement screen opens.
10. To accept the license agreement, select **I have read and agree to the terms of this license**, and click **Next**. button.
    The licensing server creates the license key text.
11. Copy the license key.
12. In the **License Text** field on BIG-IQ, paste the license text.
13. Click the **Activate License** button.
14. In the **Hostname** field, type a fully-qualified domain name (FQDN) for the system.

    You cannot change this name after you add it. The FQDN can consist of letters and numbers, as well as the characters underscore (_), dash (-), or period (.).
15. In the **Management Port IP Address** and **Management Port Route** fields, type the IP address for the management port IP address and route.

    Note: *The management port IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.
16. Specify what you want the BIG-IQ to use for the **Discovery Address**.

    - To use the management port, select **Use Management Address**.
    - To use the internal self IP address, select **Self IP Address**, and type the IP address.

      Important: *If you are configuring a data collection device, you must use the internal self IP address.*

      Note: *The self IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.
17. In the **DNS Lookup Servers** field, type the IP address of your DNS server.

    You can click the **Test Connection** button to verify that BIG-IQ can reach that IP address.
18. In the **DNS Search Domains** field, type the name of your search domain.

    The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
19. In the **Time Servers** field, type the IP addresses of your Network Time Protocol (NTP) server.

    You can click the **Test Connection** button to verify that BIG-IQ can reach the IP address.
20. From the **Time Zone** list, select your local time zone.
21. In the **Old Password** fields, type the default admin and root passwords, and then type a new password in the **Password** and **Confirm Password** fields.

Discover and activate a data collection device

Using BIG-IQ® Centralized Management, you can discover a data collection device, and add it to the Logging Group. The BIG-IQ can then access the data on the discovered data collection device. You can then receive data from multiple BIG-IP® systems. This unified view makes browsing easier, and provides a complete view of application alert or event activity and statistics data.

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click **Add** .
4. On the New BIG-IQ Data Collection Device screen, fill in as appropriate:
   a) In **Management Address**, type the management IP address.
   b) In **Username**, type the user name for an administrator on the data collection device (for example, `admin`.
   c) In **Password**, type the password for an administrator on the data collection device (for example, `admin`.
   d) In **Data Collection IP Address**, type the IP address of the data collection device internal self IP address.
   e) For **Data Collection Port**, the default value is `9300`. The BIG-IQ uses this port for internal polling and communication with the data collection devices.
   f) For **Zone**, select the zone in which the data collection device resides.

      Your Admin sets up the zones so that the nodes in your cluster are distributed equitably for disaster recovery purposes.
5. Click the **Add** button at the bottom of the screen to add the data collection device to the system.

   Note: *This operation might take a minute or two.*

6. Repeat the preceding steps for each data collection device you want to configure.
7. To activate this device for the service you want to monitor, on the BIG-IQ Data Collection Devices screen, in the Services column, click **Add Services**.
   The Services screen for this data collection device opens.
8. For the service you want to add, confirm that the **Listener Address** correctly specifies the external self IP address of the data collection device, and click **Activate**.
   When the service is successfully added, the **Service Status** changes to `Active`.
9. Click **Save & Close**.

Once discovered and activated, this data collection device collects the data generated by the configured BIG-IP systems. Thus, BIG-IQ provides a single view of all alert or event entries and statistics data.

Important: *The* ***Total Document Count*** *is not a report of the number of alerts or events sent to the data collection device. Instead, it is a sum of various document types sent to the data collection device. Events and alerts are included in this list, but this total includes other document types as well.*

Modifying alert log indices for Access

Before you can configure the indices for a data collection device, you must activate services for the components that you want to collect data for.

Alert log indices determine the physical characteristics of what is sent to the data collection device. You can set up index rotation for Access alerts.

Important: *The configuration process for log indices varies significantly depending on the component that generates the data. Use the process that corresponds to the component you are trying to set up.*

1. At the top of the screen, click **System**.

2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. On the left, click **Logging Data Collection** .
5. For Access Policy (APM), click the **Configure** button.
   The Access Indices screen opens.
6. Perform the next two steps for each section on this screen.

   ***Important:** To avoid a mismatch in the reports generated from your logging data, use the same indices values for the **access-event-logs** and **access stats**.*

7. Specify the **Rotation Type**.
   - To chunk your data based on the amount of data:
     1. Select **Size Based**
     2. For the **Max Index Size**, type the size of the chunks you want event log data sent to the DCD.

     ***Note:** For example, if you type `1000`, when the event log data reaches a size of 1 Gig, it will be sent to the DCD.*
   - To chunk your data based on the increments of time:
     1. **Select Time Based**
     2. For the **Rotation Period**, specify a time unit, and type how many of those you want to comprise each data chunk sent to the DCD.

     ***Note:** For example, if you type `.5` and select **Hours**, the event log data will be sent to the DCD every half hour.*
8. For the **Retained Index Count**, type the total number of indexes you want to store on the DCD.

   The maximum amount of data stored on the DCD is determined by this setting. When the amount of data reaches this size, the oldest data is truncated or discarded.
9. Click **Save & Close** to save the indices configuration settings.

Modifying alert log indices for Web Application Security

Before you can configure the indices for a data collection device, you must activate services for the components that you want to collect data for.

Alert log indices determine the physical characteristics of what is sent to the Data Collection Device. Use this task to set up index rotation for Web Application Services alerts.

***Important:** The configuration process for log indices varies significantly depending on the component that generates the data. Use the process that corresponds to the component you are trying to set up.*

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.

The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.

4. On the left, click **Logging Data Collection** .
5. For Web Application Security (ASM), click the **Configure** button.
   The ASM Indices screen opens.
6. Specify the **Rotation Type**.
   - To chunk your data based on the amount of data:
     1. Select **Size Based**
     2. For the **Max Index Size**, type the size of the chunks you want event log data sent to the DCD.

     ***Note:** For example, if you type `1000`, when the event log data reaches a size of 1 Gig, it will be sent to the DCD.*
   - To chunk your data based on the increments of time:
     1. **Select Time Based**
     2. For the **Rotation Period**, specify a time unit, and type how many of those you want to comprise each data chunk sent to the DCD.

     ***Note:** For example, if you type `.5` and select **Hours**, the event log data will be sent to the DCD every half hour.*
7. For the **Retained Index Count**, type the total number of indexes you want to store on the DCD.

   The maximum amount of data stored on the DCD is determined by this setting. When the amount of data reaches this size, the oldest data is truncated or discarded.
8. Click **Save & Close** to save the indices configuration settings.

Modifying event log indices for FPS

Before you can configure the indices for a data collection device, you must activate services for the components that you want to collect data for.

Event log indices determine the physical characteristics of what is sent to the data collection device. You can set up index rotation for Fraud Protection Services events.

***Important:** The configuration process for log indices varies significantly depending on the component that generates the data. Use the process that corresponds to the component you are trying to set up.*

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. On the left, click **Logging Data Collection**.
5. for Fraud Protection (FPS), click the **Configure** button .
   The FPS Indices screen opens.
6. Specify the **Rotation Type**.
   - To chunk your data based on the amount of data:

1. Select **Size Based**
2. For the **Max Index Size**, type the size of the chunks you want event log data sent to the DCD.

__Note:__ For example, if you type `1000`, when the event log data reaches a size of 1 Gig, it will be sent to the DCD.

- To chunk your data based on the increments of time:

  1. **Select Time Based**
  2. For the **Rotation Period**, specify a time unit, and type how many of those you want to comprise each data chunk sent to the DCD.

  __Note:__ For example, if you type `.5` and select __Hours__, the event log data will be sent to the DCD every half hour.

7. For the **Retained Index Count**, type the total number of indexes you want to store on the DCD.

   The maximum amount of data stored on the DCD is determined by this setting. When the amount of data reaches this size, the oldest data is truncated or discarded.
8. Click **Save & Close** to save the indices configuration settings.

Configure secure communications for data collection device

You need a signed SSL certificate before you can configure HTTPS communications to a data collection device.

If you want to secure the communications between the BIG-IP® devices and your data collection device cluster using SSL encryption, you must provide a signed SSL certificate to the BIG-IP devices and F5® BIG-IQ® Centralized Management systems. You do this by configuring both the BIG-IP device and the data collection device.

__Note:__ The BIG-IP device that generates Fraud Protection Service alerts must be configured to send its alerts to the data collection device (DCD). This process is documented in a separate guide. The guide F5 Fraud Protection Service: Configuration, Version 13.0 provides complete setup instructions for using FPS on a BIG-IP system. Complete the standard setup as documented in the guide, except when you configure the alert server pool, add your DCDs to an alerts pool using their internal self IP addresses.

1. Use SSH to log in to the data collection device.
2. Replace the content of the `/etc/httpd/conf/ssl.crt/` directory on the data collection device with your signed SSL certificate.
3. Replace the content of the `/etc/httpd/conf/ssl.key/` directory on the data collection device with your signed SSL key.
4. To apply these changes to the data collection device, type: `bigstart restart webd` and then press Enter.
5. Log out of the data collection device.

Add a proxy for secure communication

Before you can perform this task, you must be logged in as Admin, and you must have configured a proxy server that your data collection device cluster can access.

As a security precaution, you may want to configure a proxy to route communications. For example you might use it to route your forwarded alerts or download alert rules from the security operations center. Or you might want to use a proxy to avoid exposing the BIG-IQ® device when you download ASM® signature files.

__Important:__ To use a proxy for Fraud Protection Service, you must configure a proxy on each device (each data collection device and both the primary and the secondary BIG-IQ devices) in the cluster. The proxy *names you specify for each node in the cluster must match exactly, but the IP address and port number for the proxy can be different from device to device.*

1. At the top of the screen, click **System**.
2. On the left, click **PROXIES**.
3. On the Proxies screen, click **Add**.
4. If you are configuring an HA peer group, from **Device**, select the primary BIG-IQ system.
5. For **Name**, type a name for the proxy you want to use.

   Important: *The proxy name must match across all devices in the cluster. The proxy addresses and port can vary.*

6. For **Address**, type the IP address of the proxy server.
7. For **Port**, type the port that you want the proxy server to use.
8. If the proxy server requires authentication, type the **User Name** and **Password** for the proxy.
9. To add another proxy, click the plus sign in the upper right hand corner, and then repeat the preceding 4 steps.
10. If your network configuration includes an HA peer, repeat steps 3 - 9, but this time, in step 4, select the HA secondary system.

    Note: *The proxy name for the HA secondary must match the name used for the primary. The proxy addresses and port can vary.*

11. Click **Save & Close**

You need to add a proxy for each data collection device in the cluster.

Note: *Remember, the proxy name must match across all devices in the cluster. The proxy addresses and port can vary.*

Define external storage snapshots location

Before you can configure the external snapshot storage location, you need the following information for the machine you will use to store your data collection device (DCD) snapshots:

- Storage-machine-IP-address
- Storage-file-path
- User name, password, and (optionally) the domain for the user account configured on the external storage device
- Read/Write permissions for the storage file path

You need snapshots to perform software upgrades and to restore your old data.

When you create DCD snapshots, they need to be stored on a machine other than the DCD. You define the location for the snapshot using the BIG-IQ® Centralized Management device.

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. For **External Storage**, click **Configure**.

The External Storage popup screen opens.

5. In the **User name** and **Password** fields, type the user name and password for the user account configured on the external storage device.
6. For the **Domain**, type in the domain name for the user account configured on the external storage device.
7. For the **Storage Path**, type the path to the external storage location.

   You can specify the device using the IP address or the host name. Additionally, you need to specify the path to the folder on the external storage device. For example:

   ```
   //<storage machine ip-address>/<storage-file-path>
   ```

   Note: *Remember, the folder you specify must have full read, write, and execute permissions.*

8. To test the settings just specified, click **Test**.
   A message displays to tell you whether the test completes successfully. If it does not, correct the settings and permissions until it completes successfully.
9. When the external storage is specified successfully, click **Save**.

The storage location should now be accessible to the all of the devices in the DCD cluster.

Define snapshot schedules

Before you define snapshot schedules, you must have defined the snapshot storage location.

Snapshots of the data sent to your data collection devices are an essential safeguard for your data. If the machine that stores the data fails, the data can be restored using these snapshots. These snapshots are created based on the snapshot schedules you define. F5 recommends that you schedule snapshots at least every 6 hours, and retain at least 4 snapshots.

Note: *You perform this task on the BIG-IQ® Centralized Management device; not on the data collection device (DCD).*

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. To view the list of snapshot schedules for this device,in the External Storage & Snapshots area, for **Snapshot Schedules** click the **View Schedules** button.
   The BIG-IQ Data Collection Snapshot Schedules screen opens.
5. To define a new snapshot schedule for this device, click , click **Create**.
   The New Logging Snapshot Schedules screen opens.
6. For the **Snapshot Name Prefix**, type the string that you want to use to identify the snapshots created by this schedule.

   For example `snapshot_`.
7. In **Snapshots to Keep**, specify the number of snapshots that you want to accumulate before they are deleted for space constraints.

   For example, if you specify `25`, then the system will retain a maximum of 25 snapshots before it starts to delete older snapshots as new snapshots are created. You can save up to 100.
8. Define how you want the snapshots scheduled.

Option	Description
Schedule the interval at which you want to create snapshots:	You schedule the system to take snapshots indefinitely. Snapshots are created at the frequency you specify. **1.** Select **Repeat Interval**. **2.** Specify the **Snapshot Frequency**. **3.** Select a time increment. For example, if you set the frequency to **6** and **Hours**, the first DCD snapshot is taken immediately (on **Save**). Subsequent snapshots are taken every 6 hours.
Schedule specific days on which you want to create snapshots:	You schedule the system to take snapshots on specific days. **1.** Select **Days of the Week**. **2.** For the **Days of the Week** setting, select the days on which you want backups to occur. **3.** For the **Start Date**, select the time (date, hour, minute, and AM or PM) on which you want backups to start.

9. Click **Save & Close** to save the new schedule.

Overview of configuring the data collection device to BIG-IP device connection

The workflow to configure data to route from the BIG-IP® devices to your data collection device (DCD) cluster depends on the type of data you want to collect.

- To collect statistics data, refer to *Discover and activate a data collection device*.
- To collect Access Policy Manager® data, refer to *Configuring remote logging for Access Policy Manager*.
- To collect Fraud Protection Services data, refer to *Configuring BIG-IP FPS devices to route alerts to a data collection device*.
- To collect Web Application Security data, refer to:
  - *Configuring the BIG-IP logging profile*
  - *Virtual servers that remote logging uses to route event logs*
  - *Assigning the logging profile to a virtual server*

.

Configure remote logging for Access Policy Manager

BIG-IP® devices that you configure for remote logging send Access reporting and SWG log report data to the BIG-IQ® data collection device for storage and management.

1. At the top left of the screen, click **Monitoring** > **DASHBOARDS** > **Access**.
2. Click **Remote Logging Configuration**.
   The Remote Logging Configuration screen opens to display all of the discovered BIG-IP devices that are provisioned with the Access service.
3. Select the BIG-IP devices for which you want to enable remote logging, and then click **Configure**.
   The hostname of the primary data collection device is displayed, and the status changes to let you know whether the enable request was successful.

Configuring BIG-IP FPS devices to route alerts to a data collection device

The BIG-IP® device that generates Fraud Protection Service alerts must be configured to send its alerts to the data collection device (DCD). This process is documented in a separate guide. The guide *F5® Fraud Protection Service: Configuration, Version 13.0* provides complete setup instructions for using FPS on a BIG-IP® system. Complete the standard setup as documented in the guide, except when you configure the alert server pool, add your DCDs to an alerts pool using their internal self IP addresses.

Note: *Although DCDs use their own version of load balancing to level the data stored on each node, it is best practice to configure the BIG-IP pool members with a load balancing method that ensures smooth traffic flow to the DCDs. The load balancing method you configure should:*

- Distribute traffic between the nodes.
- Ensure that, if a DCD goes offline, the BIG-IP device must still be able send traffic to the available DCDs without dropping alerts.

The default port to specify is `8008`, but you can use a different port if your DCD is configured for it. To ensure that alerts are received even if one DCD goes down, specify at least one alternative DCD.

Configure the BIG-IP logging profile

For Web Application Security users, this is the first of three tasks required to route the BIG-IP® event logs to a BIG-IQ® data collection device. You configure the BIG-IP system by creating a logging profile and assigning the logging profile to a virtual server, and then deploying it to the BIG-IP system. The *logging profile* defines the content of the events, and identifies the data collection device to which the events are sent.

1. At the top of the screen, click **Configuration**.
2. On the left, click **SECURITY** > **Shared Security** > **Logging Profiles**.
   The Logging Profiles screen opens to display the logging profiles that have been configured on this device.
3. On the Logging Profiles screen, click **Create**.
   The New Logging Profile screen opens, showing the Properties information.
4. On the Properties screen, edit as appropriate:
   a) In the **Name** field, type a unique name for this new profile. This field is required.
   b) For the **Description**, you can specify an optional description for the logging profile.
   c) For the **Partition**, you can specify the partition to which the logging profile belongs. Only users with access to a partition can view the objects (such as the logging profile) that it contains. If the logging profile resides in the `Common` partition, all users can access it. Although this field is pre-populated with `Common` by default, you can set the partition when creating logging profiles by typing a unique name for the partition.

   Note: *The partition with the name you specify must already exist on the BIG-IP device. No whitespace is allowed in the partition name.*

   d) To specify the devices to which you want to deploy this logging profile, select the devices in the **Available** list, and click the right arrow to add them to the **Selected** list.
5. On the left, click **Application Security**, and then select the **Enabled** check box.

   The screen displays the Application Security settings.

   a) Select the **Remote Storage Enabled** check box.
   The screen displays additional settings, and the **Local Storage** option becomes active.
   b) Clear the **Local Storage** check box.
   c) Specify the appropriate **Logging Format**.

   - If the BIG-IP device runs version 12.0 or later, select **BIG-IQ**.
   - If the BIG-IP device runs a version earlier than 12.0, select **Comma-Separated Values**.
     Several new settings appear.

     - For **Storage Format**, select **User Defined**.
     - In the **Selected** field, paste the following text:

```
unit_hostname="%unit_hostname%",management_ip_address="%management_ip_address%",
http_class_name="%http_class_name%",web_application_name="%http_class_name
%",policy_name="%policy_name%",
policy_apply_date="%policy_apply_date%",violations="%violations%",support_id="%support_id%",
request_status="%request_status%",response_code="%response_code%",ip_client="%ip_client%",
route_domain="%route_domain%",method="%method%",protocol="%protocol
%",query_string="%query_string%",
x_forwarded_for_header_value="%x_forwarded_for_header_value%",sig_ids="%sig_ids
%",sig_names="%sig_names%",
date_time="%date_time%",severity="%severity%",attack_type="%attack_type
%",geo_location="%geo_location%",
ip_address_intelligence="%ip_address_intelligence%",username="%username
%",session_id="%session_id%",
src_port="%src_port%",dest_port="%dest_port%",dest_ip="%dest_ip
%",sub_violations="%sub_violations%",
virus_name="%virus_name%",uri="%uri%",request="%request
%",violation_details="%violation_details%",
header="%headers%",response="%response%
```

Note: *The line breaks in the example above were necessary due to screen width; remove all of them after you paste this data. It must be a single string with no white space.*

d) For **Protocol**, select **TCP**.

e) For the **Server Addresses** settings, specify the address you want to use:

  1. In the **IP Address** field, type the data collection node's management IP address.
  2. Specify the port to use for your data.
     - If you are setting up a logging profile for Web Application Security, type `8514` in the **Port** field.
     - If you are setting up a logging profile for Fraud Protection Service, type `8008` in the **Port** field.
  3. Click the **Add** button to add the address and port to the list of servers.

f) For the **Maximum Entry Length**, select **64k**.

g) In the Storage Filter area, from the **Request Type** list, select **All requests**.

6. If you want to specify Protocol Security options, on the left click **Protocol Security**, then select the **Enabled** check box: the Protocol Security settings display. Edit as appropriate.
7. If you want to specify Network Firewall options, on the left click **Network Firewall**, then select the **Enabled** check box: the Network Firewall settings display. Edit as appropriate.
8. If you want to specify Network Address Translation options, on the left click **Network Address Translation**, then select the **Enabled** check box: the Network Address Translation settings display. Edit as appropriate.
9. If you want to specify DoS Protection options, on the left click **DoS Protection**, then select the **Enabled** check box: the DoS Protection settings display. Edit as appropriate.
10. Click **Save & Close** to save the new profile.

The new logging profile is added to the list of profiles defined on this device.

Before you can begin using this profile, you must assign it to a virtual server and then deploy the virtual server to the BIG-IP device.

Virtual servers that remote logging uses to route alert or event logs

You can either create a new virtual server on the BIG-IP® device that creates the alert or event, or you can use a virtual server that already exists on that device.

Creating a virtual server for remote logging

If the device for which you are configuring remote logging does not have a virtual server, you need to create one.

1. At the top of the screen, click **Configuration**.

2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays a list of virtual servers defined on this device.
4. Click **Create**.
   The Virtual Servers - New Item screen opens.
5. In the **Name** field, type in a name for the virtual server you are creating.
6. From the **Device** list, select the device on which to create the virtual server.
7. In the **Description** field, type in a brief description for the virtual server you are creating.
8. For the **Destination Address**, type the IP address of the destination you want to add to the Destination list.

   The format for an IPv4 address is `I<a>.I<b>.I<c>.I<d>`. For example, `172.16.254.1`.

   The format for an IPv6 address is `I<a>:I<b>:I<c>:I<d>:I<e>:I<f>:I<g>:I<h>.`.

   For example, `2001:db8:85a3:8d3:1319:8a2e:370:7348`.
9. In the **Service Port** field, type a service port number, or select a type from the list.

   When you select a type from the list, the value in the **Service Port** field changes to reflect the associated default, which you can change.
10. Click **Save**.
    The system creates the new virtual server with the settings you specified.
11. Click **Save** to save the assignment. Or, click **Save & Close** to save the assignment and return to the Virtual Servers screen.

A virtual server that can be used to route alert or event data to the logging node is created for the BIG-IP® device.

Before the BIG-IP device can actually use this new virtual server, you must deploy it to the device.

Assign the logging profile to a virtual server

After configuring a logging profile on the BIG-IQ® system, you must assign it to a virtual server and deploy it to the BIG-IP® device from which you want to collect event logs.

1. At the top of the screen, click **Configuration**.
2. On the left, click **SECURITY** > **Shared Security** > **Virtual Servers**.
   The screen displays a list of virtual servers that are configured with devices that have been provisioned and discovered.
3. On the Virtual Servers screen, click the name of the virtual server you want to use.
   The Virtual Servers - Properties screen opens.
4. From the **Log Profiles** list, under **Available**, click a logging profile and move it to the **Selected** list.
5. Click **Save & Close** to save the assignment and return to the Virtual Servers screen.

The virtual server is now associated with the logging profile.

Before the BIG-IP system(s) can start sending alert or event logs to the data collection device, you must deploy the changes you just made to the BIG-IP device.

Managing a Data Collection Device Cluster

Data collection device best practices

There are a number of useful concepts to consider when you manage data collection devices for off-box log storage. This reference material might prove helpful in setting up and maintaining your data collection device (DCD) configuration.

Important: *As part of maintaining a DCD cluster, you might need to remove one or more devices from your DCD cluster. When you remove a DCD from the cluster, BIG-IQ® Centralized Management moves the data to another device in the cluster. Whenever you move data, losing part or all of that data is a risk. Therefore, before you remove a DCD from the cluster, F5 recommends creating a snapshot to back up your logging data.*

Restore data collection device snapshots

You can use the BIG-IQ® user interface to restore data collection device (DCD) snapshots.

Please note:

- The restore operation requires a down time during which no BIG-IQ or DCD work is performed.
- During the restore operation, no data sent to the DCD is retained.
- The restore operation restores only the data from the time before the chosen snapshot was created. Data from the time that the chosen snapshot was created to the current time is not restored.
- Before initiating a snapshot restore, make sure that sufficient disk space is allocated to the /var folder on the device to which you are restoring the snapshot.

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. You have two options for choosing a snapshot and starting the restore, using the settings in the External Storage & Snapshot area near the bottom of the screen.

Option	Description
To restore from the most recent snapshot:	Next to **Last Snapshot/Time**, click **Restore Latest**.
To select the snapshot that you want to restore:	1. Click the **View History** button. 2. Choose the snapshot you wish to restore, and click **Restore**.

Delete a data collection device snapshot

If you determine that there are issues with a specific snapshot, you can delete it so that you cannot accidentally restore to it in the future.

***Note:** You perform this task on the BIG-IQ® Centralized Management device; not on the data collection device (DCD).*

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. Near the bottom of the screen, click the **View History** button.
   The BIG-IQ Data Collection Snapshots screen opens.
5. Browse through the list to find the snapshot you want to delete.
6. Select the check box for the snapshot you want to delete, and click **Delete**.

Check data collection device health

You can use the BIG-IQ® Data Collection Device Settings 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 DCD cluster configuration is as you expect it to be.

***Note:** You perform this task on the BIG-IQ Centralized Management device; not on the data collection device (DCD).*

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Click the **Settings** button.
   The Settings screen opens to display the current state of the DCD cluster defined for this BIG-IQ device.
4. Inspect the DCD cluster details listed in the Summary area.

   #### Data Cluster Status

   Look here for information about the current state of the cluster.

   #### Master Device

   The read-only Master Device field displays the host name of the BIG-IQ device that manages and monitors the health of this DCD cluster.

Devices in Cluster

Displays the total number of devices in the cluster including DCDs, the BIG-IQ Centralized Management devices and the optional peer.

Data Collection Devices in Cluster

Displays the number of DCDs that have been added to the cluster.

Total Document Count

Displays the number of all document types stored on the cluster. Alerts and events are included in this list, but the total includes other types of document as well.

Total Document Size

Displays the amount of disk space consumed by the documents stored for this cluster.

This information provides a fairly detailed overview that describes the DCD cluster you have created to store data. After you complete an upgrade, you can check the health to verify that the cluster restored successfully.

Index rotation policy

The optimum settings used to configure your data collection device (DCD) indices depend on a number of key factors.

- The system provides the ability to dynamically create new indices based on either a specified interval or a specified size. The primary goal to consider when you make these decisions is how to maintain a maximum disk allocation for the DCD data, while maintaining capacity for new data that flows in.
- Secondary considerations include search optimization, and the ability to optimize old indices to reduce their size.
- Generally, the best policy is one that does not create unnecessary indices. The more indices, the lower the overall performance, because your searches have to deal with more shards. For example, if a module knows that it has a low indexing volume (thousands/day) then it makes the most sense to have a large aggregation per rotation (5 days or 30 days). For components like Web Application Security that probably have high indexing volumes, it makes more sense to rotate every 8 hours (which reduces the number of retained indices).
- Index rotation also allows changing sharding and replica counts by changing the template on a given index type. New indices created from that template will contain the new shard and replica count properties.

This table shows the default configuration values for each index running on BIG-IQ® Centralized Management. These values are based on anticipated data ingestion rates and typical usage patterns.

Component	Index Name	Minimum Number of DCDs	Rotation Policy	Retained Index Count	Approximate time window	Size of /var file system
Access	access-event-logs	2	Time/5 days	19	95 days	500 GB
Access	access-stats	2	Time/5 days	19	95 days	500 GB
Web Application Security	asmindex	2	Size/100000 MB	5	N/A	500 GB
FPS	websafe	2	Time/30 days	100	8 years	10 GB

If multiple modules are running on a given DCD or if you have higher inbound data rates, you might have to adjust these values to keep the `/var` file system from filling up. (There is a default alert to warn of this when the file system becomes 80% full.)

The simplest resolution is to revise the retained index count; lowering this value reduces the disk space requirements, but it will also reduce the amount of data available for queries. For details on changing this setting, refer to the modifying indices topic for the component you are configuring.

Managing Disaster Recovery Scenarios

How does a data collection device cluster deal with disaster recovery scenarios?

The BIG-IQ® system uses high availability and zone awareness functions to maintain data collection device operations even when a node or an entire data center goes down. DCDs in each data center are assigned to the appropriate zone. This zone awareness enables the system to manage the distribution of your data and maintain DCD operation in all but the most severe outages.

For a better understanding of how this process works, consider the following example. A hypothetical company named Acme has two data centers; one in Seattle and the other in Boston. Acme wants to ensure data reliability, and has set up these data centers so that if one goes offline, the DCD data it was receiving is routed to the other data center. To achieve this, Acme has an HA pair of BIG-IQ console nodes for viewing and managing the data, and six DCDs divided equally between the two data centers. Two BIG-IP® devices are used to load balance data and configure Fraud Protection Services and Application Security Manager™ settings for both data centers.

The HA pair ensures that one BIG-IQ console node is always available for managing configuration data. The standby console is available for viewing configuration data, and managing data. The DCDs are treated as one large cluster that is split between two sites. Each data node is replicated, so that if one goes down, or even if an entire data center goes down, the data is still available.

Important: *The total number of DCDs should be an even number so that they can be split evenly between the data centers. This configuration makes it easier to configure zones.*

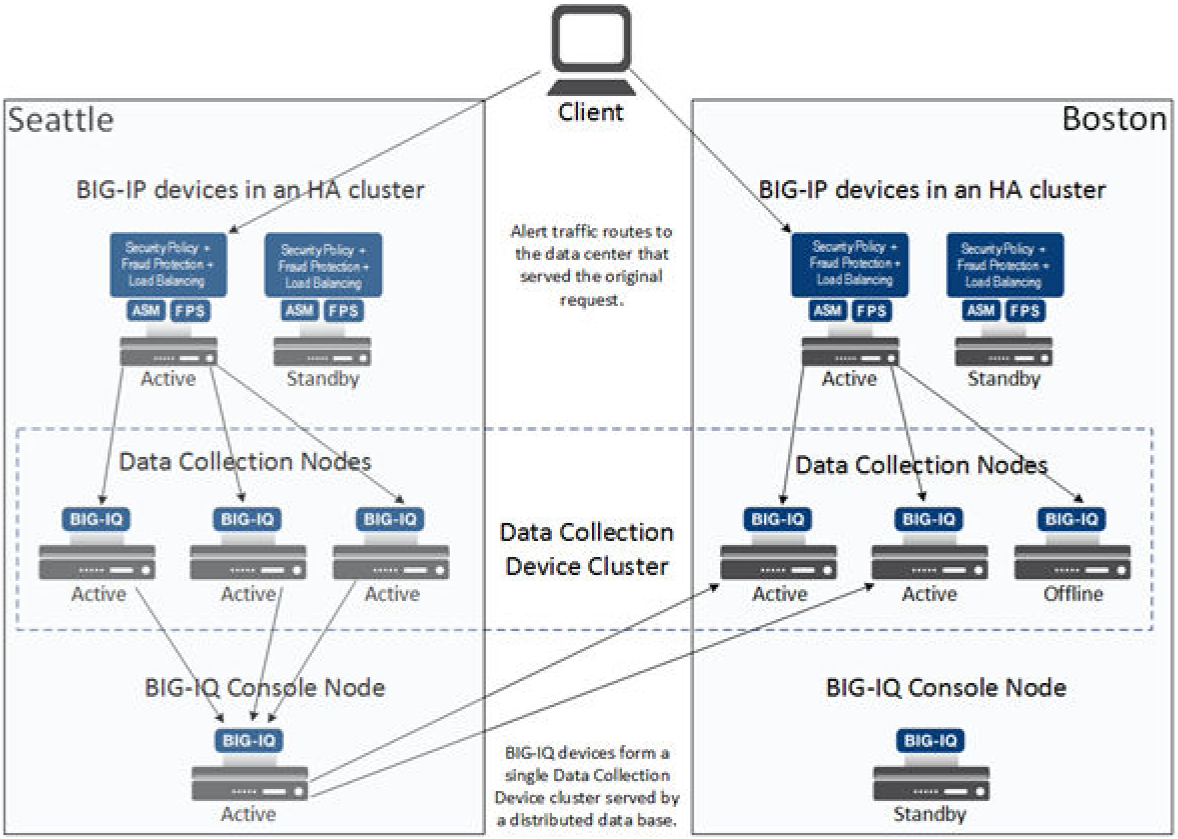
Client
Seattle
BIG-IP devices in an HA cluster
Security Policy + Fraud Protection + Load Balancing
ASM
FPS
Active
Standby
Data Collection Nodes
BIG-IQ
Active
BIG-IQ Console Node
Alert traffic routes to the data center that served the original request.
Data Collection Device Cluster
BIG-IQ devices form a single Data Collection Device cluster served by a distributed data base.
Boston
BIG-IP devices in an HA cluster
Offline
Standby

Figure 5: Two data centers, one DCD cluster example

The DCD cluster logic that governs the distribution of data between your DCDs identifies one node in the cluster as the master node. The master node monitors the cluster health and manages the cluster operation. It is elected by the cluster, and can reside on any node, including a console node. (Console nodes however, do not store any data). When the master node goes down, a new master is elected from among all the nodes in the cluster.

How is data handled when DCDs fail?

Here are some of the most common failure scenarios that can occur, and how the cluster responds:

- One of the DCDs fails. Alert data is automatically routed to another DCD in the zone.
- The master node fails. The cluster logic chooses a new master node. This process is commonly referred to as *electing* a new master node.
- All of the DCDs in a zone fail. BIG-IP devices route their data to DCDs in the other zone.

How is data handled when communication between the two data centers fails?

This scenario is a little more complex and needs a little more detail to understand. It's also much less likely to occur. In the two data center DCD cluster scenario referenced previously, there is one master node located in one of the two data centers. The admin doesn't control which node is the master, but the master node is identified on the Logging Configuration page. For this example, let's assume that the master node is elected in the Seattle data center. If communication goes down between the data centers, the Seattle data center continues to function as before. In the Boston data center, a new master node is elected because without communication between the two data centers, each center forms its own DCD cluster. So, initially, BIG-IP devices in both data centers are sending data to the DCDs in their clusters, and a master node in each zone is controlling that zone.

When communication is resumed, the original cluster reforms and the master node from one of the zones is re-elected. The master node then syncs its data with the (new) nodes in the cluster. This data sync overwrites the data on the new nodes. If the Boston node is elected, its data (which only includes data collected during the communications failure) overwrites the data on the Seattle node (this means that most of the data for the cluster is lost). The Seattle data set includes both the data collected from before, and from during the communications failure. To prevent overwriting the more comprehensive data set and losing logging data, you can perform two precautionary steps.

- When a communication failure occurs, change the target DCDs for BIG-IP devices in the zone that did not include the original master node (Boston, in our example) to one of the DCDs in the other zone.
- When communication is restored, in the zone that did not include the original master node (Boston in our example), use SSH to log in to the master node as `root`, then type `bigstart restart elasticsearch` , and press Enter. Restarting this service removes this node from the election process just long enough so that the original master node can be elected.

The result is that during the failure, all data is sent to nodes in the zone that contained the original master node. Then when communication is restored, the DCDs in the zone in which the master node was restarted rejoin the cluster, and the data is synced to all of the DCDs. All data is preserved.

How does the minimum master eligible nodes setting work?

One parameter of special significance in determining the behavior of the data collection device (DCD) cluster is the minimum master eligible nodes (MMEN) setting. All of the nodes in the cluster (including the primary and secondary console nodes) are eligible to be the master node.

When a node is added or removed from the DCD cluster, the system performs a calculation to determine the optimum default value. You can override the default value to suit your requirements.

This setting determines how many DCDs in the cluster must be online for the cluster to continue to process alert data. If your goal is to keep operating regardless of node failures, it would seem like the obvious choice would be to set this number to as low a value as possible. However, you should keep in mind a couple of factors:

- The BIG-IQ® console is counted as a node in the cluster, so a cluster size of 1 does not make sense.
- Similarly, a cluster size of 2 (a DCD and the BIG-IQ console) is not a good idea. Because the DCD cluster logic uses multiple DCDs to ensure the reliability of your data, you need at least two logging nodes to get the best data integrity.
- It might also seem like a good idea to set the MMEN value to a higher value (for example, one less than the number in the entire cluster), but actually, best practice is to not specify a value larger than the number of nodes in one zone. If there is a communications failure, the nodes in each zone compose the entire cluster, and if the MMEN is set to a lower value, both clusters would stop processing data.

How is alert data handled when data collection devices fail?

Here are some of the most common failure scenarios that can occur to a data collection device cluster, and how the cluster responds to that scenario.

What failed?	How does the cluster respond?
One of the data collection devices fails.	All alert data, including the data that was being sent to the failed node, is still available. When a node is added, removed, or fails, the cluster logic redistributes the data to the remaining nodes in the cluster.
The master node fails.	The cluster logic chooses a new master node. This process is commonly referred to as *electing* a new master node. Until the new master is elected, there

What failed?	How does the cluster respond?
	may be a brief period during which alert processing is stopped. Once the new master is elected, all of the alert data is available.
All of the data collection devices in a zone fail.	Just as when a single data collection device fails, the cluster logic redistributes the data to the remaining nodes in the cluster

How is data handled when communication between the two data centers fails?

This scenario is a little more complex than the case where data collection devices fail and needs a little more discussion to understand. However, it's also much less likely to occur. The cluster behavior in this scenario is controlled by the MMEN setting. In a two data center data collection device (DCD) cluster scenario, the MMEN setting is `3` and there is one master node located in one of the two data centers. The admin doesn't control which node is the master, but the master node is identified on the Logging Configuration screen. For this example, let's assume that the master node is in the Seattle data center. If communication goes down between the data centers, the Seattle data center continues to function as before because with four nodes (the console and three DCDs) it satisfies the MMEN setting of `3`. In the Boston data center, a new master node is elected because without communication between the two data centers, communication with the master node is lost. Since there are also four master eligible nodes in the Boston data center, it satisfies the MMEN setting too. The Boston data center elects a new master and forms its own cluster. So, initially, BIG-IP® devices in both data centers are sending alerts to the DCDs in their clusters, and a master node in each zone is controlling that zone.

When communication resumes, the original cluster does not reform on its own, because both data centers have formed their own independent clusters. To reform the original cluster, you restart the master node for one of the clusters.

- If you restart the master node in the Boston data center, the cluster logic sees that the Seattle data center already has an elected master node, so the Boston cluster joins the Seattle cluster instead of forming its own. The Seattle master node then syncs its data with the Boston nodes in the cluster. The data sync overwrites the Boston data with the Seattle data. The result is that Boston data received during the communication failure is lost.
- If instead, you restart the master node in the Seattle data center, the cluster logic would see that the Boston data center already has an elected master node, so the Seattle cluster would join the Boston cluster. The Boston master node would then sync its data with the nodes in the Seattle cluster. That data sync would overwrite the Seattle data with the Boston data. In this case, the result is that Seattle data received during the communication failure is lost.

To preserve as much data as possible, instead of just reforming the original cluster by restarting one of the clusters, we recommend you perform the following two precautionary steps.

1. When a communication failure occurs, change the target DCDs for the BIG-IP devices in the zone that did not include the original master node (Boston, in our example) to one of the DCDs in the zone that housed the original master node (Seattle, in our example).
2. When communication is restored, in the zone that did not include the original master node (Boston, in our example), use SSH to log in to the master node as `root`, and then type `bigstart restart elasticsearch` , and press Enter. Restarting this service removes this node from the election process just long enough so that the original (Seattle) master node can be elected.

After you perform these two steps, all alerts are sent to nodes in the zone that contained the original master node. Then when communication is restored, the DCDs in the zone in which the master node was restarted (Boston) rejoin the cluster. The resulting data sync overwrites the Boston data with the Seattle data. The Seattle data center has the data that was collected before, during, and after the communications failure. The result is that all of the data for the original cluster is saved and when the data is synced, all alert data is preserved.

How do I optimize my deployment for disaster recovery?

When you have data collection devices in multiple data centers, you can optimize your deployment to maintain data collection when some or all of the data collection devices in one data center fail. A few slight variations to the deployment process make it possible for you to take advantage of the BIG-IQ® system's zone awareness feature. The resulting data collection device (DCD) cluster deployment provides the optimum data collection performance in an outage scenario.

***Note:** In the user interface, a data center is identified by its zone name. The data center in Seattle and the zone named Seattle are the same thing.*

1. Install and configure a BIG-IQ console node in both data centers.
2. Configure the two console nodes so that one is the HA primary and the other is the HA secondary.
3. For the console node configured as the HA primary, specify the zone in which that device resides.
4. Deploy the DCD in each data center.
5. On the primary console node, add all of the DCDs to the cluster. As you add devices, specify the zone name appropriate for the data center in which each node is physically located.
6. Initiate an HA failover of the primary BIG-IQ console to the secondary BIG-IQ console.
   When the BIG-IQ HA primary fails over to the secondary, the logging configuration information propagates to both zones and a new primary is designated.
7. On the (newly designated) primary BIG-IQ console, specify the name of the zone in which that console resides.
8. Initiate an HA failover of the (new) primary BIG-IQ console to the secondary BIG-IQ console.
   The BIG-IQ console that was originally the HA primary is once again the primary.
9. Set the minimum master capable nodes on the primary console. To display the HA properties screen, click **System** > **BIG-IQ HA** and then select the device.

   The **Zone** field is located at the bottom of the screen.

How do I perform routine maintenance on a node?

Before you do routine maintenance on a data collection device (DCD), there are a couple of steps you should perform to make sure the DCD cluster operation is not impacted.

1. On the BIG-IQ Data Collection Devices screen, remove the node from the DCD cluster.
   The system recalculates and resets the **Minimum Master Eligible Nodes** setting.
2. If your DCD cluster configuration does not use the default value, override the **Minimum Master Eligible Nodes** setting to its previous value.
3. Perform the maintenance on the DCD.
4. Add the node back into the DCD cluster.
   The system recalculates and resets the **Minimum Master Eligible Nodes** setting.
5. If your DCD cluster configuration does not use the default value, override the **Minimum Master Eligible Nodes** setting to its previous value.

How do I change the zone for a data collection device?

If you decide to change the zone for a data collection device (DCD), you should perform a couple of extra steps to make sure that the cluster recognizes the change.

1. At the top of the screen, click **System**.

2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to display the currently defined data collection device cluster.
3. Under Device Name, select the DCD that you want to revise.
4. On the DCDs properties page, click **Edit** to change the zone for the DCD.
5. Use SSH to log in to DCD as `root`.
6. Type `bigstart restart elasticsearch`. and press Enter.
7. Repeat the last two steps for each DCD, and for each BIG-IQ ®system in the cluster.

   Note: *As you run this command on each DCD, it momentarily stops processing DCD data, so the data routes to another node in the cluster and no data is lost.*

Legal Notices

Legal notices

Publication Date

This document was published on April 14, 2017.

Publication Number

MAN-0666-00

Copyright

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.

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