F5 BIG-IQ Centralized Management: Monitoring and Reports

Version 5.2

Table of Contents

Health monitoring and alerts using SMTP and SNMP alerts

Health and event monitoring using SMTP and SNMP alerts

You can use F5® BIG-IQ® Centralized Management to easily monitor the health of your managed devices, as well as BIG-IQ itself, using the following tools:

- Simple Mail Transfer Protocol - SMTP is a standard for email transmission used for monitoring and alerting you to the health of devices in your network.
- Simple Network Management Protocol - SNMP is an industry standard protocol for monitoring devices on IP networks. Once configured, the SNMP agent sends data collected from BIG-IQ Device to your third-party SNMP manager. BIG-IQ is compatible with SNMPv1, SNMPv2c, and SNMPv3.

After you configure SMTP and/or SNMP (which you typically do when you initially set up BIG-IQ), you can specify email recipients to receive alerts when certain events occur. These alerts are configurable; you can enable and disable them, and, for some alerts, you can set specific thresholds to prompt an alert.

Specify an SMTP server to send email alerts

You specify an SMTP server so F5® BIG-IQ® Centralized Management can send email to alert specified people when a certain condition happens, such as when an SSL certificate is about to expire.

1. At the top of the screen, click **System**.
2. On the left, click **SMTP configuration.**
3. Click the **Add** button at the upper right of the screen.
4. In the **Name** field, type a name for this SMTP configuration.
5. In the **SMTP Server Host** and **SMTP Server Port** fields, type the SMTP server and TCP port.
   By default, SMTP uses TCP 25.
6. In the **From Email Address** field, type the email address from which to send the alert email.
7. From the **Encryption** list, select the type of encryption to use for the email.
8. To require a user name and password, from the **Use Auth** list, select **Yes**, and type the required user name and password.
9. To verify that you can reach the server you configured, click the **Test Connection** button.
10. Click the **Save & Close** button at the bottom of the screen.

You can now specify email recipients and set up the alert conditions that prompt BIG-IQ to send an email when a certain event happens on a managed device.

How do I set up BIG-IQ to work with SNMP?

Simple Network Management Protocol (*SNMP*) is an Internet-standard protocol for collecting and organizing information about managed devices on IP networks. You can set up BIG-IQ® to work with SNMP so you can receive alerts when certain things happen on a managed device.

To set up BIG-IQ to work with SNMP, you must:

1. Set up the SNMP Agent

2. Configure SNMP Access
3. Specify settings for the SNMP Trap

Before you configure SNMP

Gather the following information before you start your SNMP configuration.

CONFIGURATION COMPONENT	CONSIDERATIONS	FOR MY CONFIGURATION
SNMP administrator contact information	Find out or decide who is responsible for SNMP administration. The contact information is a MIB-II simple string variable.	
Machine location	Find out the location of the BIG-IQ system. The location is a MIB-II simple string variable.	
BIG-IQ client allow list	Gather the IP or network addresses (with netmasks) of the SNMP managers from which the SNMP agent will accept requests.	
Access	Find the OID for the top-most node of the SNMP tree to provide access to.	
Community	Get the v1 and v2c communities and the IP addresses of the SNMP managers you want to grant access to.	
Users	Get the v3 users you want to grant access to SNMP data, along with the privacy protocols and passwords, Community, Destination, and Port.	

Configuring SNMP agent for sending alerts

This screen displays specified user addresses allowed to access your 3rd-party SNMP Manager BIG-IQ through the SNMP Agent. An agent can communicate with multiple managers, so you can configure BIG-IQ to support communications with one management station using the SNMP version1 protocol, one using the SNMP version 2C protocol, and another using SMNP version 3.

1. At the top of the screen, click **System**.
2. On the left, click .
3. At the top of the screen, click the **Download MIB** button to download the F5-required MIBs.
4. At the top of the screen, click **Edit**.
5. Edit the **Contact Information** and **Machine Location** fields to reflect your SNMP agent settings and click the **Save & Close** button at the bottom of the screen.
6. Click the **Save & Close** button at the bottom of the screen to save your changes.
7. For the **SNMP Access - Client Allowed List** setting, click the **Add** button.
8. In the **Addresses/Networks** and **Mask** fields, type the IP address and networks and the netmask (if applicable) that the SNMP manager is allowed to access.
9. To add another address, click the plus (+) sign.
10. At the bottom of the screen, click the **Save & Close** button.

You can now configure SNMP access and SNMP traps.

Configure Access and Traps for SNMP version 3 to send alerts

After you configure the SNMP agent, you can configure SNMP access and SNMP traps.

You configure SNMP access to allow the SNMP agent to accept requests from specific SNMP managers.

1. At the top of the screen, click **System**.
2. On the left, click **LOCAL HOST SETTINGS** > **SNMP Configuration** > **SNMP Access (v3)**.
3. Click the **Add** button at the upper right of the screen.
4. In the **Name** and **User Name** fields, type a name for this SNMP access and the user name.
5. If you want to specify the authentication protocol for SNMP traps, from the **Type** list, select an option.
   - **MD5** specifies digest algorithm.
   - **SHA** specifies secure hash algorithm.
6. If you selected an authentication protocol, in the **Password** and **Confirm Password** fields, type and confirm the password for access.

   The password must be between 8 and 32 characters, include alphabetic, numeric, and special characters, but no control characters.
7. If you want to encrypt the SNMP traps, from the **Protocol** list, select an option.
   - **AES** specifies Advanced Encryption Standard
   - **DES** specifies Data Encryption Standard
8. If you selected a privacy protocol, in the **Password** and **Confirm Password** fields, type the password to use for authentication.

   Alternatively, you can select the **Use Authentication Password** check box to use the authentication password.
9. In the **OID** field, type the object identifier (OID) you want to associate with this user.
10. From the **Access** list, select an option:
    - **Read Only** - This user can only view the MIB.
    - **Read/Write** - This user can view and modify the MIB.

    The most secure access level or type takes precedence when there is a conflict. When you set the access level to read/write, and an individual data object has a read-only access type, access to the object remains read-only.
11. Click the **Save & Close** button at the bottom of the screen to save your changes.
12. On the left, click **SNMP Traps**.
13. In the **Name** field, type a name for this SNMP trap.
14. From the **Version** list, select **V3**.
15. In the **Destination** and **Port** fields, type the IP address and the port for this trap destination.
16. For the **Security Level** setting, select an option. **Auth, No Privacy** processes SNMP messages using authentication, but no encryption. **Auth and Privacy** processes SNMP messages using authentication and encryption.
17. For the **Security Name** setting, specify the user name you want to use to handle SNMP version 3 traps.
18. For the **Engine ID** setting, specify the unique identifier (snmpEngineID) of the remote SNMP protocol engine.
19. In the **Password** and **Confirm Password** fields, type and confirm the password for the protocol.
20. Click the **Save & Close** button at the bottom of the screen to save your changes.

You can now specify email recipients for alerts.

Configuring Access and Traps for SNMP version 1 and 2C to send alerts

After you configure the SNMP agent, you can configure SNMP access and SNMP traps.

You configure SNMP access to allow the SNMP agent to accept requests from specific SNMP managers.

1. At the top of the screen, click **System**.
2. On the left, **LOCAL HOST SETTINGS** > **SNMP Configuration** > **SNMP Access (V1, V2C)**
3. At the top left of the screen, click the **Add** button.
4. In the **Name** field, type the SNMP manager's user name.
5. From the **Type** list, select the format for the IP address.
6. In the **Community** field, type the community string (password) for access to the MIB.
7. From the **Source** list, select a source or select **Specify** and type the source address for access to the MIB.
8. In the **OID** field, type the object identifier (OID) you want to associate with this user.
9. From the **Access** list, select an option:
   - **Read Only** - This user can only view the MIB.
   - **Read/Write** - This user can view and modify the MIB.

   The most secure access level or type takes precedence when there is a conflict. When you set the access level to read/write, and an individual data object has a read-only access type, access to the object remains read-only.
10. Click the **Save & Close** button at the bottom of the screen to save your changes.
11. On the left, click **SNMP Traps**.
12. At the top left of the screen, click the **Add** button.
13. In the **Name** field, type a name for this SNMP trap.
14. In the **Community**, **Destination**, and **Port** fields, type, respectively, the community name, IP address, and port for the trap destination.
15. At the bottom of the screen, click the **Save & Close** button.

You can now specify email recipients for alerts.

Add email recipients for SMTP and SNMP alerts

After you configure SMTP and/or SNMP, you can add email recipients.

Email recipients you add will get alert notifications when specified events happen on BIG-IQ or your managed devices

1. At the top of the screen, click **System**.
2. On the left, click **LOCAL HOST SETTINGS** > **Email Notification Recipients**.
3. At the top left of the screen, click the **Add** button.
4. In the **Name** the **Email** address fields, type the name and email address of the person you want to receive an alert.
5. In the **Description** field, you can type an optional description to help identify this user.
6. Select the check box next to each type of notification you want this user to receive an email about.
7. To add another email recipient, click +.
8. Click the **Save & Close** button at the bottom of the screen to save your changes.

You can now configure the alert settings that trigger BIG-IQ to send an email to the specified recipients.

How do I monitor SSL certificate expiration dates for my managed devices?

When you manage BIG-IP® devices that load balance SSL traffic, you must monitor their SSL traffic.

BIG-IQ® imports the certificates for every managed BIG-IP device you discover. This makes it easy to monitor the expiration dates all of your devices' SSL certificates from one location.

You can also:

- Set up alerts to let you know when a certain certificate is about to expire within a specified number of days.
- Download the data to a CSV file for reporting purposes.

Set up alert conditions that triggers BIG-IQ to send a notification

After you set up the SNMP and/or SMTP on F5 ®BIG-IQ® Centralized Management, you can select the alerts that prompt BIG-IQ to send an email to the people you specified.

1. At the top of the screen, click **Monitoring**.
2. On the left, click **ALERTS & NOTIFICATIONS.**
3. At the top of the screen, click the **Settings** button.
4. Select the **Enabled** check box next to each alert you want to receive and, if applicable, specify the **Threshold**.

   Only SNMP events specified as **Yes** are available for SNMP alerts. BIG-IQ uses SMTP for all other event types.
5. Click the **Save & Close** button at the bottom of the screen.

Statistics Monitoring Overview

What analysis can I perform using collected statistics?

You can use the statistics collected by F5® BIG-IQ® Centralized Management to visually analyze the performance of Local Traffic objects, device traffic, DNS traffic, and overall system statistics. The statistics are displayed in graphical charts and tables that you can drill down into for more specific details. You might want to track network performance on a device, or memory and CPU utilization; or you might want to compare the performance on two devices or a group of devices. You can focus the statistics in the charts on different categories such as virtual servers, pools, pool members, or DNS traffic.

How do I get started managing statistics data?

You can monitor statistics data generated by the devices you manage. You can monitor the performance of your BIG-IQ® devices as well, but that works differently and is discussed in the online help. There are a few things you need to do before you can start monitoring the statistics data generated by your managed devices.

- You need to install, configure, and discover a data collection device (DCD). The DCD stores the data from your devices, and routes the date to your BIG-IQ device. Refer to *Planning and Implementing a Centralized Management Deployment* for details.
- You need to enable statistics collection for the devices you want to monitor. There a couple of ways to do that. Refer to *Enabling Statistics Collection* for details.

Once you have your system set up and you are receiving statistics, you should take a minute or two to understand how the user interface works. The interface is set up so that statistics from Device, DNS, and Local Traffic use a common set of tools to manage how you access and manage your data. Once you understand how this common interface works, you should be ready to go.

Important: *Statistics for the Access component use a slightly different user interface. For details on monitoring these statistics, refer to Access Reporting and Statistics.*

Important: *Statistics for this BIG-IQ device also use a different user interface. For details on monitoring these statistics, go to* ***System*** > ***THIS DEVICE*** > ***Statistics****, or* ***System*** > ***THIS DEVICE*** > ***BIG-IQ Metrics*** *and refer to the online help.*

What makes up a statistics overview screen?

This figure shows a typical statistics overview screen. To view a screen similar to this, click **Monitoring** > **DASHBOARDS** > **Device** > **Health**. Until you configure statistics collection, there won't be any data. The three parts of the statistics overview screens work together so you can fine-tune the statistics display.

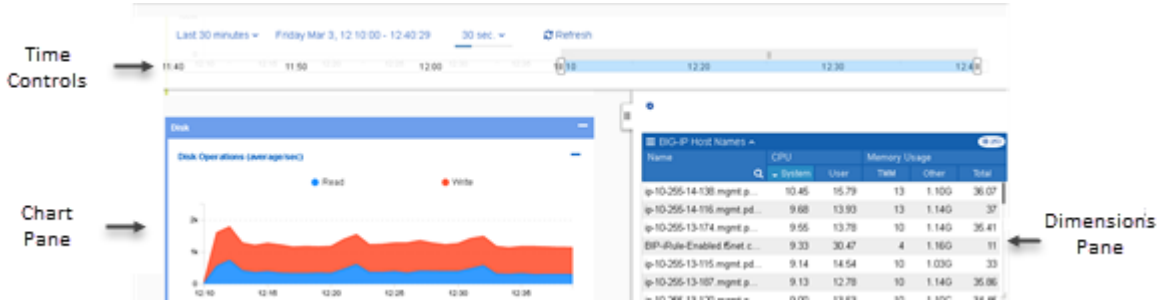

Figure 1: Typical statistics overview screen

Key elements of this screen are defined in the table.

User interface control	What does this part of the screen do?
Time Controls	Adjusts the time window for which statistics are displayed. For details on how these controls work, see *How do the time controls work?*
Chart Pane	Displays a series of charts that plot the collected statistics. For details on how to manipulate these charts, see *How does the chart pane work?*
Dimensions Pane	Determines the objects that you display statistics for. For details on how the controls on this pane work, see *How does the dimensions pane work?*

How do the time controls work?

This figure shows a close up of the time controls on a typical overview screen. To view a screen similar to this, click **Monitoring** > **DASHBOARDS** > **Device** > **Health**.The four time controls work together to give you control of the specific time period for which you wan to see statistics.

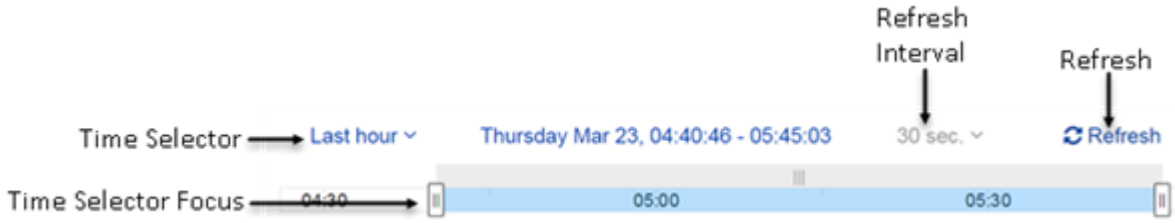

Figure 2: Overview screen time control detail

Key elements of this screen are defined in the table.

User interface control	What does this control do?
Time Selector	Use this control to specify the length of time for which you want to view statistics data. When you first start looking at statistics, only the **All** option is available. Then as you gather additional data, additional time period options become available. Data is displayed from the instant the last refresh occurred, back to time interval you specify. For example, if the last refresh occurred at 11:00, and the Time Selector is set to 30 minutes, the charts display data from 10:30 to 11:00.
Time Selector Focus	Use this control to focus on a specific window of time within the currently selected time period. Use the sliders at either end of this control to define the time segment you want to examine. The time segment you select here is indicated along the lower horizontal axis of the chart pane to provide a

User interface control	What does this control do?
	reference. For example, if the last refresh occurred at 11:00, and the Time Selector is set to 30 minutes, you could use the sliders to look at the period from 10:45 to 10:50. When you adjust the sliders, the time markers along the bottom of each chart axis update to indicate your selection. Also, note that you can adjust both ends of this control. If you adjust the right side of the control, the auto refresh stops, effectively freezing the display so you can focus on a particular data point. ***Tip:*** *Alternatively, you can click on the chart axis to specify the focus. Click on a point in the axis and drag in the direction that you want to view. As you drag, a highlighted window shows what you have selected for the time selector focus, and a small magnifying glass icon appears. When you have the time selector focus you want, click the magnifying glass.*
Refresh Interval	Use this control to specify how frequently the data on this page is refreshed.
Refresh	Use this control to trigger an immediate refresh of the data on this screen.

How does the chart pane work?

This figure shows a closer look at the elements that make up the chart pane on a typical Overview screen. To view a screen similar to this, click **Monitoring** > **DASHBOARDS** > **Device** > **Health**. You can re-order the charts by dragging and dropping them into place.

Figure 3: Overview screen, chart pane detail

Key elements of this screen are defined in the table.

User interface control	What does this control do?
Chart Title	Each chart displays a title that identifies the statistic that plots on that chart. Each title includes the units of measure that apply to these plots.
Statistics Legend	These colored dots identify the specific plots displayed on the chart. When you move your cursor over a chart, the value of each plot displays adjacent to these dots. If there is a multiplier applied to a value, it displays as well. For example, if you hover over one of the New Connections plots and the value `31.9k` displays, it means that there were an average of 31,900 connections per second at that point in time.
Statistics Values	These plots display the value of the statistics collected for the selected time period. Data is aggregated for the objects or devices that are currently selected. Initially, the selection is all of the managed objects or devices, but you can

User interface control	What does this control do?
	use the dimensions pane to change the selection. If you select one device, the charts shows statistics for just that device. If you select two devices, the charts plot aggregated statistics for those devices. If you then select just one device, and five virtual servers, the charts plot aggregated statistics for the five virtual servers and the single device. For more information on using the dimensions pane refer to *How does the dimensions pane work?*
Hide/Display Chart	Use this control to hide or display a chart. When you hide a chart, the chart title remains. If you create a comparison chart, an additional control appears that you can use to delete that chart.

How does the dimensions pane work?

In the BIG-IQ® user interface, a dimension is a statistical category (for example BIG-IP® host name or iRule event types). Each dimension is broken up into sub-categories that you can view when you expand the dimensions pane to display your statistical data as a table. The other primary use for the controls in the dimensions pane is to filter the objects for which statistics are displayed in the chart window. To view a screen similar to the following illustration, click **Monitoring** > **DASHBOARDS** > **Device** > **Health** > , and then click the down arrow on the BIG-IP Host Names dimension. After you expand a dimension, you can select individual objects, or multiple objects, or create a comparison graph that displays statistics for selected objects. The figure shows the key elements that make up the dimensions pane on a typical Overview screen. You can re-order the dimensions by dragging and dropping them into place.

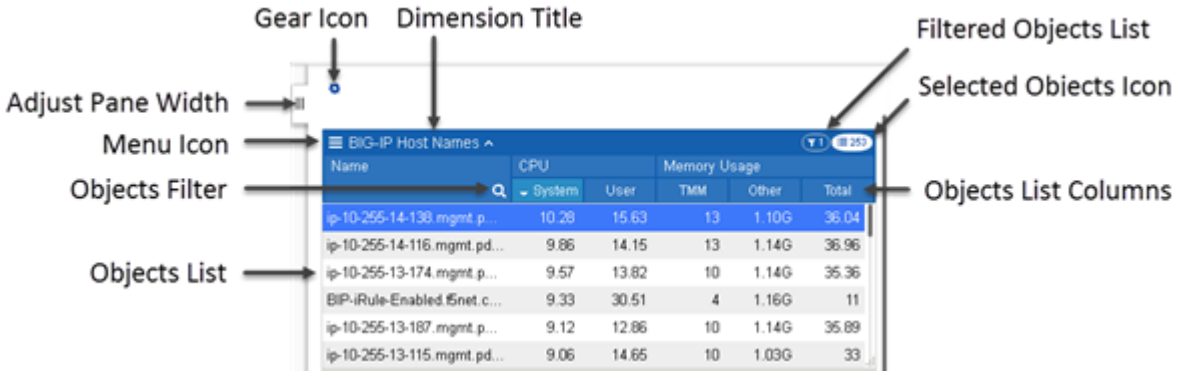

Figure 4: Overview screen dimensions pane detail

Key elements of this screen are defined in the table. Except for the name of each dimension and the pane width adjustment, these controls display only when you expand a dimension to display all of its members.

User interface control	What does this control do?
Adjust Pane Width	With this tab, you can adjust the width of the chart and dimensions panes. To adjust the dimensions pane width, click this tab and drag the pane to the width you want. To extend the dimensions pane to full screen, single click this tab.
Menu Icon	When you click the menu icon, you can choose between several options that you can use to change what is displayed in the Dimensions pane. • **Sort by** lists the columns defined for the selected object. You can choose which column you want to sort the listed objects by.

User interface control	What does this control do?
	• **Columns** lists the available columns that can be displayed for the selected object. You can choose which columns you want to display. For a list of what each column in the dimension pane specifies, see Dimension Pane Columns. ***Note:*** *You can also sort a dimension or specify which columns display by right-clicking the column header.* If you select one or more objects for a dimension, you can choose two additional options: • If you choose **Add Comparison Chart**, you can create a graph that plots values for two or more selected metrics. You can create multiple comparison charts and for each comparison chart, you change the metric that is being compared. • If you choose **Clear Selection**, you choose which objects you want to de-select.
Objects Filter	Click the magnifying glass to open a filter control. To filter the list of objects that display in the Objects List, type the name you want to find and click the magnifying glass again. Note that the filter is a prefix match (it starts with the first character of the object name), and the match is case-sensitive.
Objects List	The first 100 objects that meet the filter criteria display here. To display an object not in the top 100, you can change your filter criteria or sort order. Objects selected in the objects list, control the data that plots in the charts. That is: • When you select just one object, the data plotted on the charts is only for that object. • If you select additional objects, the data that plots is the aggregate for those selected objects. • When no objects are selected, the data that plots is the aggregate for all objects in the dimension. ***Note:*** *If you select two or more objects in this list, you can create a comparison chart that plots values for a selected parameter.*
Objects List Columns	The default number of columns that display depends on the type of dimension. You can also change which columns display using the menu icon. You can sort the entries in a dimension by clicking an individual column title. .
Dimension Title	The title of the dimension displays here adjacent to an up arrow/down arrow toggle. This toggle collapses and expands the list of objects of this dimension type for the devices you are currently managing.
Gear Icon	The gear icon provides several options that you can use to change how the objects you have selected in the Dimensions panel display. • You can use **Clear Filters**to de-select all of the objects for the selected dimension. • You can use **Reset Layout** to undo any custom sorting you have applied and reset the selected dimension to the default layout.

User interface control	What does this control do?
	• You can use **Sort Selected** to resort the list of objects for the selected dimension. With this filter applied, the objects you have selected move to the top of the column. Any type of sort you use to reorder this dimension change the sort order for the remaining objects, but the selected objects stay at the top of the list.
Selected Objects Icon	This icon displays the number of objects in this dimension that match the current filter settings. Note that the objects you select in one dimension impact the number of objects in the other dimensions. As an example, consider a BIG-IQ managing 20 BIG-IP devices that each have 100 virtual servers. Initially, on the Virtual Servers overview screen, the selected objects icon in the BIG-IP Host Names dimension reads 20, and the selected objects icon on the Virtual Servers dimension reads 2000. If you select one BIG-IP device, the icon for virtual servers changes to 100. On the other hand, if you select one of the virtual servers, the icon for BIG-IP devices would change to 1 (unless that virtual server happens to reside on more than one device).
Filtered Objects Icon	This icon displays the number of objects you have selected in this dimension. You can also click this icon to de-select all objects in this dimension.

Comparison Charts

Comparison charts allow you to plot data values for selected items in a new chart. When you initially create a comparison chart, you select the statistical metric that you want to compare. When the comparison chart displays, the title for the new chart displays which items are selected for comparison, along with the metric being compared, followed by a down arrow icon. You can click that down arrow if you want to change the comparison metric for the selected objects. To compare additional items, you can create additional comparison charts by again selecting multiple items, right clicking and choosing Create Comparison Chart, and selecting the statistical metric you want to compare.

Configuring Statistics Collection

How do I start viewing BIG-IP device statistics from BIG-IQ?

To start viewing statistics for a BIG-IP® device, you must enable statistics collection for that device. You can do that either during or after adding the device to the BIG-IP Devices inventory list on the BIG-IQ® system. You also need to install, configure, and add a data collection device before you can view device statistics.

Enabling statistics collection during device discovery

Before you can enable statistics for BIG-IP® devices:

- There must be a BIG-IQ® data collection device configured for the BIG-IQ device.
- The BIG-IP device must be located in your network and running a compatible software version. Refer to *https://support.f5.com/kb/en-us/solutions/public/14000/500/sol14592.html* for more information.
- Port 22 and 443 must be open to the BIG-IQ management address, or any alternative IP address used to add the BIG-IP device to the BIG-IQ inventory. These ports and the management IP address are open by default on BIG-IQ.

If you are running BIG-IP version 11.5.1 up to version 11.6.0, you might need `root` user credentials to discover and add the device to the BIG-IP devices inventory. You don't need `root` user credentials for BIG-IP devices running versions 11.6.1 - 12.x.

Note: *A BIG-IP device running versions 10.2.0 - 11.5.0 is considered a legacy device and cannot be discovered from BIG-IQ version 5.2. If you were managing a legacy device in previous version of BIG-IQ and upgraded to version 5.2, the legacy device displays as impaired with a yellow triangle next to it in the BIG-IP Devices inventory. To manage statistics for it, you must upgrade it to version 11.5.1 or later. For instructions, refer to the section titled, Upgrading a Legacy Device.*

One way to enable statistics collection for BIG-IP devices is to do it when you add those devices to the BIG-IQ system inventory. Adding devices to the inventory is referred to as *device discovery*. If the devices you want to enable have already been discovered, refer to *Enabling collection after device discovery*.

Note: *The ADC component is automatically included (first) any time you discover or import services for a device.*

Note: *You do not need to discover and import a device's configuration to collect and view statistics for it. You just need to establish trust between your BIG-IQ and the device. If you do not discover and import the device configuration, the virtual servers, pool, pool members, and iRules will be visible in the statistics dimension panes, but these objects will not appear in the configuration page for those objects. Also, you will not be able to manage these objects in BIG-IQ. If you decide you want to manage these objects, you can discover and import the BIG-IP device's configuration later without interrupting statistics collection.*

---

1. At the top of the screen, click **Devices**.
2. Click the **Add Device** button.
3. In the **IP Address** field, type the IPv4 or IPv6 address of the device.
4. In the **User Name** and **Password** fields, type the user name and password for the device.
5. If this device is part of a DSC pair, from the **Cluster Display Name** list, select one of the following:

- For an existing DSC pair, select **Use Existing** from the list and select the name DSC group from the list.
- To create a new DSC pair, select **Create New** from the list, and type a name in the field.

For BIG-IQ to properly associate the two devices in the same DSC group, the **Cluster Display Name** must be the same for both members in a group.

There can be only two members in a DSC group.

6. If this device is configured in a DSC pair, select an option:
   - **Initiate BIG-IP DSC sync when deploying configuration changes (Recommended)** Select this option if this device is part of a DSC pair and you want this device to automatically synchronize configuration changes with the other member in the DSC group.
   - **Ignore BIG-IP DSC sync when deploying configuration changes** Select this option if you want to manually synchronize configurations changes between the two members in the DSC group.
7. Click the **Add** button at the bottom of the screen.
   The BIG-IQ system opens communication to the BIG-IP device, and checks its framework.

   Note: *The BIG-IQ system can properly manage a BIG-IP device only if the BIG-IP device is running a compatible version of the REST framework.*
8. If a framework upgrade is required, in the popup window, in the **Root User Name** and **Root Password** fields, type the root user name and password for the BIG-IP device, and click **Continue**.
9. If in addition to basic management tasks (like software upgrades, license management, and UCS backups) you also want to centrally manage this device's configurations for licensed services, select the check box next to each service you want to discover.

   You can also select these service configuration after you add the BIG-IP device to the inventory.
10. To enable statistics collection for this BIG-IP device, under Statistics monitoring, select the check box next to each service you want to collect statistics for, and then click **Continue**.

    Note: *If you want to enable statistics collection without managing any services, just clear the check boxes for all services.*
11. Click the **Add** button at the bottom of the screen.

Enable statistics collection for devices

Before you can enable statistics collection for a BIG-IP® device using this method:

- The device must already be in the BIG-IQ system inventory.
- There must be a BIG-IQ data collection device configured for the BIG-IQ device.

Generally, if you want to collect statistics for a BIG-IP device, you enable statistics collection when you discover it. But you can enable or disable statistics collection for a device any time it is convenient for you.

1. At the top of the screen, click **Devices**.
2. Click the name of the device you want to enable statistics collection for.
3. On the left, click **Statistics Collection**.
4. To begin statistics collection, for **Collect Statistics Data**, select **Enabled**.
5. For **Modules/Services**, click the check box for the types of statistics you want to collect.
6. For **Frequency**, next to **Collect every**, select the interval at which you want to collect statistics from this device.

After you enable statistics collection for a device, data for that device begins aggregating along with any other devices for which you are collecting data. Two buttons (**View Health Statistics**, and **View Traffic**

Statistics) are added to the properties page for enabled devices. Clicking either of these takes you directly to the overview page for the statistics type you clicked.

Manage the retention policy for your statistics data

Before you can set the statistics retention policy, you must have added a data collection device.

You can manage the settings that determine how your statistics data is retained. The highest quality data is the raw data, (data that has not been averaged), but that consumes a lot of disk space, so you need to consider your needs in choosing your data retention settings. When you choose how much raw data to retain, you need to consider how much disk space you have available. The controls on this screen are simple to set up, but understanding how they work takes a bit of explanation.

The fields on the Statistics Retention Policy screen all work in similar fashion. One way to understand how these fields work is to think of your data storage space as a set of containers. The values you specify on this screen determine how much storage space each container consumes. Because data is saved for the time periods you specify, the longer the time period that you specify, the more space you consume. The disk storage that is consumed depends on several factors.

- The number of BIG-IP® devices you manage
- The number of objects on the BIG-IP devices you manage (for example, virtual servers, pools, pool members, and iRules®)
- The frequency of statistics collection
- The data retention policy
- The data replication policy

There are three key concepts to understand about how the retention policy works.

How long is data in each container retained?	Data is retained in each container for the time period you specify. When the specified level is reached, the oldest chunk of data is deleted. For example, if you specify a raw data value of 48 hours, then when 48 hours of raw data accumulate, the next hour of incoming raw data causes the oldest hour to be deleted.
When does data from one container pass on to the next?	Data passes from one container to the next in increments that are the size of the next (larger) container. That is, every 60 minutes, the last 60 minutes of raw data is aggregated into a data set and passed to the **Hour(s)** container. Every 24 hours, the last 24 hours of hourly data is aggregated into a data set and passed to the **Day(s)** container, and so on for the **Month(s)** container.
What about limits?	**Limit Max Storage to** specifies the percentage of total disk space that you want data to consume on the data collection devices in your cluster. If more disk space is consumed than the percentage you specified, BIG-IQ takes two actions: **1.** New statistical data is not accepted until the available disk space complies with the **Limit max storage to** setting.

2. Statistical data not required to calculate the next higher time layer is removed (for example, you need 60 minutes of raw data to aggregate to the Hours level). Data is removed starting with the raw data container, then the hourly data container, then the daily time container. This process stops when storage consumption is below the **Limit max storage to** setting.

The BIG-IQ takes this action to prevent data corruption when storage is completely exhausted.

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 **Statistics Collection**.
   The Statistics Collection Status screen displays the percentage of available disk space currently consumed by statistics data for each container.
5. To change the retention settings for your statistics data, click **Configure**.
   The Statistics Retention Policy screen opens.
6. In the **Keep real-time (raw) data up to** field, type the number of hours of raw data to retain.

   You must specify a minimum of 1 hour, so that there is sufficient data to average and create a data point for the **Keep hourly data up to** container.
7. In the **Keep hourly data up to** field, type the number of hourly data points to retain.

   You must specify a minimum of 24 hours, so that there is sufficient data to average and create a data point for the **Keep daily data up to** container.
8. In the **Keep daily data up to** field, type the number of daily data points to retain.

   You must specify a minimum of 31 days, so that there is sufficient data to average and create a data point for the **Keep monthly data up to** container.
9. In the **Keep monthly data up to** field, type the number of monthly data points to retain.

   Once the specified number of months passes, the oldest monthly data set is deleted.
10. In the **Limit max storage to** field, type the percentage of disk space that you want collected data to consume before the oldest monthly data set is deleted.
11. 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. By default, data replication for statistics is not enabled. Disabling replication reduces the amount of disk space required for data retention. However, this provides no protection from data corruption that can occur when you remove a data collection device. You should enable replicas to provide this protection.
12. When you are satisfied with the values specified for data retention, click **Save & Close**.

Analyzing Statistics Data

Browse through your managed devices looking for high resource usage

Before you can analyze BIG-IP® device performance data:

- There must be a BIG-IQ data collection device configured for the BIG-IQ® system that manages your BIG-IP devices.
- You must have discovered the BIG-IP devices that you want to analyze, and statistics collection must be enabled for those devices.
- It is also a good idea, though not a requirement, to define the retention policy for the statistics data you are collecting.

You can use the Device Health overview screen to review the resource usage for your BIG-IP devices.

1. At the top of the screen, click **Monitoring**.
2. On the left, click **DASHBOARDS** > **Device** > **Health**.
   The Device Health overview screen opens.
3. On the Device Health overview screen, adjust the dimensions pane so that it fills up at least half of the screen.
4. In the dimensions pane, click the down arrow on the BIG-IP Host Names dimension to expand the list of BIG-IP devices.
5. In the CPU column click the **System** heading to sort the list of BIG-IP devices by CPU usage.
6. If you find a device that is consuming a high level of CPU cycles, select it.
   The chart pane displays statistics for only the selected device.
7. Scan the CPU chart for the distressed device, and look for the point where the cycles jumped up.

   Tip: *If you find a spike, you will want to focus in to narrow the time focus. There are two ways you can zoom in on an area of interest:*

   - Click either end of the focus element on the time selector to define the interval you want to examine.
   - Click and drag over the interval. A small magnifying glass pops up over the area you highlight. When you click the magnifying glass, the time selector focus is set to the area you highlighted.

   Note: *When you zoom in to focus on the area of interest, the focus changes for all charts on the Device Health overview screen.*

8. Make a note of both the host name of the BIG-IP device, and the time that the CPU cycles climbed.
   You have found a device that is in distress and identified the time that the issue started.
9. On the left, click **DASHBOARDS** > **Device** > **Traffic**.
   The Device Traffic overview screen opens.
10. On the Device Traffic overview screen, adjust the dimensions pane so that it fills up at least half of the screen.
11. In the dimensions pane, click the down arrow on the BIG-IP Host Names dimension to expand the list of BIG-IP devices.
12. In the Name column, click the magnifying glass icon (), type the host name of the BIG-IP device noted in step 8, then click the magnifying glass icon again.
    The chart pane displays statistics for only the selected device.

13. Use the time selector controls to focus in on the time noted in step 8, and look through the traffic charts to see if there was a spike in traffic that corresponds to the spike in CPU cycles.
    If you find that spike, you have found the reason for the spike in CPU cycles. Traffic on that device is very high. You probably want to find out more about that traffic.
14. On the left, click **DASHBOARDS** > **Local Traffic** > **Virtual Servers**.
    The Virtual Servers overview screen opens.
15. On the Local Traffic Virtual Servers overview screen, adjust the dimensions pane so that it fills up at least half of the screen.
16. In the Name column, click the magnifying glass icon ( , type the host name of the BIG-IP device noted in step 8, then click the magnifying glass icon again.
    The chart pane displays statistics for only the selected device. Additionally, only the virtual servers that reside on that device are available on the Virtual Servers dimension.
17. In the dimensions pane, click the down arrow on the Virtual Servers dimension to expand the list.
18. In the Bytes Avg/s column, click one of the column headings (for example **C-->**) to sort the list of virtual servers by average traffic level in bytes per second on this virtual server.
19. Find the virtual server that has the highest number of new connections, and select it.
    The chart pane now displays statistics for the virtual server on the distressed device.

With the ID of the virtual server, you can figure out which application is triggering the traffic spike and figure out what your next step is from there.

Analyze application performance issues

Before you can analyze application performance data:

- There must be a BIG-IQ Data Collection Device configured for the BIG-IQ® device that manages your BIG-IP® devices.
- It is also a good idea, though not a requirement, to define the retention policy for the statistics data you are collecting.
- You will need to know the name of the virtual server that hosts the application that is having performance issues.

When you learn that an application is having performance issues, you can analyze the objects that serve that application to find the likely cause.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Click **Virtual Servers**.
   The Virtual Servers screen opens showing a list of virtual servers managed by this BIG-IQ.
4. On the Virtual Servers screen, in the Filter in the upper right corner, type the name of the virtual server that is hosting the troubled application.
5. When you find the virtual server, select the virtual server name.
   The properties screen for the selected virtual server opens.
6. Click **View Statistics**.
   The Virtual Servers overview screen opens, but the only items selected are the virtual server you selected, and the BIG-IP device on which it runs.
7. Scan the statistics charts plotting data for the distressed server, and look for data (packet throughput) that indicates a problem.

Tip: *There are two ways you can zoom in on an area of interest:*

- Click either end of the focus element on the time selector to define the interval you want to examine.
- Click and drag over the interval. A small magnifying glass pops up over the area you highlight. When you click the magnifying glass, the time selector focus is set to the area you highlighted.

If scanning the charts does not reveal an obvious cause, your next step might be to look for a troubled pool member or node.

8. On the Virtual Servers overview screen, click the back arrow ( ) to return to the properties screen for the virtual server. Click it again to return to the Virtual Servers screen.
9. Find the virtual server name again, and select the check box that corresponds to it.
   At the bottom of the Virtual Server screen, a two panel preview pane opens displaying information about the virtual server.
10. On the right panel of the preview pane, under Related Items, click **Show**.
11. Select one of the Pool Members to display its properties screen, and then click **View Statistics**.
    The Pools & Pool Members overview screen opens, and again, the chart data that is displayed is only for the BIG-IP, pool, and pool name that you selected.

Analyze load balancing issues

Before you can analyze load balancing data:

- There must be a BIG-IQ data collection device configured for the BIG-IQ® device that manages your BIG-IP® devices.
- It is also a good idea, though not a requirement, to define the retention policy for the statistics data you are collecting.

When you get a report that an application is having performance issues, you can quickly determine if the cause is related to a load balancing problem.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Click **Virtual Servers**.
   The Virtual Servers screen opens showing a list of virtual servers managed by this BIG-IQ.
4. On the Virtual Servers screen, in the filter in the upper right corner, type the name of the virtual server that is hosting the troubled application.
5. Find the virtual server name, and select the check box that corresponds to it.
   At the bottom of the Virtual Server screen, a two panel preview pane opens displaying information about the virtual server.
6. On the right panel of the preview pane, under Related Items, click **Show**.
7. Select one of the virtual servers pools to display its properties screen, and then click **View Statistics**.
   The Pools & Pool Members overview screen opens, and the chart data that displays is only for the pool you selected.
8. In the dimensions pane, click the down arrow on the **Pool Members** dimension to expand the list.
9. Click the name of each pool member until all of the members are selected.
10. Right click selected names, and select **Add Comparison Chart**.
    A new chart is added to the top of the chart pane. The chart plots the performance of the pool members, graphing the load balancing performance over time.
11. To change the comparison metric that plots for the selected pool members, click the down arrow in the title of the chart and select the new metric.

    You might want to try looking at the average bytes going to or from the server, or maybe the average new server connections.

12. If the virtual server has multiple pools, you might have to repeat the last four steps a couple of times to get the full picture.

Managing Audit Logs

About audit logs

You use audit logs to review changes in the BIG-IQ® system. All BIG-IQ system roles have read-only access to the audit log, and can view and filter entries. Any user with the appropriate privileges can initiate an action.

All API traffic on the BIG-IQ system, and every REST service command for all licensed modules, is logged in a separate, central audit log (`restjavad-audit.n.log`) which is located in `/var/log` on the BIG-IQ system.

Considerations when using the audit log

When using the audit log, consider the following:

- The audit log does not record an entry for every generation of a task. It only records an entry when the task status changes.
- When an object is deleted and then recreated with the same name, partition, and other information, the difference between those objects may show the deleted object as being the previous generation of the new object.
- By default, not all columns are displayed by the audit log to conserve space. To review what columns are displayed, click the gear icon in the upper right of the Audit Logging screen.

Actions and objects that generate audit log entries

BIG-IQ® records in the audit log all user-initiated changes that occur on the BIG-IQ system. A change is defined as when certain objects are modified, when certain tasks change state, or when certain user actions are performed. For example, when the admin account is used to log in to the BIG-IQ system, the audit log records the time, the user (admin), the action (New) and the object type (Login). The log does not include changes that occurred on BIG-IP® devices that were imported.

Changes to working-configuration objects generate audit log entries. In addition, these actions generate log entries:

- Creating or deleting a user account.
- Users logging in and logging out, including when the user is logged out due to inactivity.
- Creating or cancelling a device discovery or a device reimport.
- Creating a Network Security Change Verifications action to verify the changes to a specific BIG-IP device or group.
- Deleting a previously discovered device.
- Creating or deleting a deployment task.
- Creating a difference task.
- Creating, restoring, or deleting a snapshot.
- Editing some system information (such as editing a host name, a root password, a DNS entry, or an SNMP entry).

Audit log entry properties

The audit log displays the following properties for each log entry.

Property	Description
Source	IP address of the client machine that made the change. This property is blank for actions that were initiated by an internal process. For example, when a user invokes a deployment action, the deployment action then invokes a difference task to find the differences between the current configuration and the one to be deployed. The difference task has no Source IP address.
Service	Indicates whether the change was made by the internal object synchronization service. This service synchronizes shared objects, such as virtual servers, from the Local Traffic & Network service to the Network Security or Web Application Security services. • If a check mark is displayed, the change was made by the internal object synchronization service, and no IP address is shown in the Source column. The check mark is only displayed in the Network Firewall Audit Log or the Web Application Security Audit Log screens. • If a check mark is not displayed, the change was not made by the internal object synchronization service.
Time	Time that the event occurred. The time is the BIG-IQ system local time and is expressed in the format: mmm dd, yyyy hh:mm:ss (time zone); for example: `Apr 19, 2016 13:09:03(EDT)`.
Node	Fully qualified domain name for the BIG-IQ system that recorded the event. This appears as the **Hostname** at the top of the BIG-IQ user interface.
User	Name of the account that initiated the action, such as an account named `Admin` for an administrative account.
Action	Type of modification. For operation changes, the action types include New, Delete, and Modify. For task changes, the action types include Start, Finish, Failed, and Cancelled.
Object Name	Object identified by a user-friendly name; for example: `newRule1`, `deploy-test`, or `Common/global`. When the name `RootNode` is listed, that indicates that the object is associated with a BIG-IP device. `RootNode` is typically seen when creating, deleting or updating log profiles, service policies, or firewall policies.
Changes	Indicates whether there was a change in the object. If **View** occurs in this column, there is a change to the object. To view the detailed differences of the change, click **View**.
Object Type	Classification for this action. When the type `Root Node` is listed, that indicates that the object is associated with a BIG-IP device. `Root Node` is typically seen when creating, deleting or updating log profiles, service policies, or firewall policies.
Parent	The administrative partition and name of the parent object. This property is displayed for firewall rules, logging profiles, and DoS profiles. For firewall rules, the parent shows the rule list, firewall, or policy that contains the rule. A change in a firewall rule often also affects the rule's parent object.
Parent Type	Class or group of the parent object.

Property	Description
Version	Version of the configuration object. Typically, when a configuration object changes, the version is increased by 1. However, other audit entries, such as those for finishing snapshot creation or finishing deployment, may increase the version by more than 1.

Viewing audit entry differences

In the audit log, when potential changes to an object are logged, the **View** link is shown in the Changes column for that entry. You can click **View** to examine the differences between generations of that object.

1. At the top of the screen, click **Monitoring**.
2. On the left, expand **LOGS**, then expand **Audit Logs**, and then , click the component that you want to view audit entries for.
3. To display differences for an object, click **View** in the Changes column.

   A popup screen opens, showing two columns that compare the differences between the two generations of the object in JSON. In these columns, additions to an object generation are highlighted in green, and differences are highlighted in gold.

   If the system cannot retrieve a generation of an object, the column displays either `Generation Not Available` or `Generation No previous generation`. Object information may not be available if it has been automatically purged from the system to conserve disk space, or if it has been deleted.

   The JSON difference displayed for a delete entry in the audit log shows the JSON difference from the previous operation because the generation identifier is not incremented when an object is deleted.
4. When you are finished, click **Close** on the popup screen to return to the Audit Logging screen.

Filtering entries in the audit log

You can use the Filter field at the top right of the Audit Logging screen to rapidly narrow the scope displayed, and to more easily locate an entry in the audit log.

- Filtering is text-based.
- Filtering is not case-sensitive.
- You can use wild cards, or partial text.
- All BIG-IQ® roles can filter entries.
- To clear the filter, click the **X** to the right of the search string in the **Filtered by** field on the left.

1. At the top of the screen, click **Monitoring**.
2. On the left, expand **LOGS**, then expand **Audit Logs**, and then , click the component that you want to view audit entries for.
3. Use the Filter field in the upper right corner to narrow your search:
   a) Select the field that you want to specify filter options for.
   b) Type the information specific to the object you want to filter on.
   c) Select **Exact** if you want to view only logs that completely match the filtering content you typed. Or, if you want to view any logs that include the filtering content, select **Contains**.
   d) Press **Enter**.

Option	Description
All	Specifies that all objects should be filtered using the filter text. When this option is used, both the user-visible and the underlying data are searched for a match, so you may see matches to your filter text which do not appear to match it.
Client Address	For **Filter**, type the IP address of the device that generates the logs. Log entries from devices with a different IP address will not be displayed.
Time	Type both a date and a time. Displayed times are given in the local time of the BIG-IQ system. Supported time formats are highly Web browser-dependent. Time formats other than those listed might appear to filter successfully but are not supported. Entering a single date and time results in a filter displaying all entries from the specified date and time to the current date and time. For time formats that use letters and numbers, enter the date time in one of the following formats: • mmm dd yyyy hh:mm:ss. Example: `Jan 7 2014 8:30:00` • mmm dd, yyyy hh:mm:ss (time zone). Example: `Apr 28, 2016 13:09:03(EDT)` • mmm dd, yyyy. Example: `Apr 28, 2016` • mmm dd, yyyy hh:mm:ss. Example: `Apr 28, 2016 16:09:06` • ddd mmm dd yyyy hh:mm:ss. Example: `Thu Jan 16 2014 11:13:50` For time formats that use only numbers, enter the date time in one of the following formats: • mm/dd/yy hh:mm:ss. Example: `01/01/16 12:14:15` • m/d/yy hh:mm:ss. Example: `1/1/14 12:14:15` • mm/dd/yyyy hh:mm:ss. Example: `1/1/2014 12:14:15`
Node	Type the node name in the filter.
User	Type the user account name in the filter.
Action: Operation	Type the operation action name in the filter. Operation actions include: New, Delete, and Modify.
Action: Task Status	Type the task status action name in the filter. Task status actions include: Start, Finish, Cancelled, and Failed.
Object Name	Type the full or partial name of the object in the filter. If a partition name is displayed, do not include it in the filter. For example, Common/AddressList_4 would be entered as `AddressList_4`. Because the device-specific object name includes the BIG-IP® host name, you can enter a full or partial device name to get all objects for a specific BIG-IP device.
Object Type	Type the object type in the filter.
Parent	Type the parent name in the filter. Only appears for rules to show the rule list, firewall, or policy that contains the rule.
Parent Type	Type the Parent Type name in the filter. Only appears when the Parent field contains a value.
Contains	Specifies that the filter text is contained within the object specified. When you select **Contains**: • If the filter text is a string, the filter text matches an entire string or only a part of a string.

Option	Description
	• If the filter text is an IP address, the filter text matches an IPV4 or IPV6 address that is the same as the filter text, or matches an IPV4 address range or subnet that includes the filter text. IPV6 addresses can not be found within a range or subnet. • If the filter text is a port number, the filter text matches a port number that is the same as the filter text, or matches a port number range that includes the filter text.
Exact	Specifies that the filter text is exactly contained within the object specified. When **Exact** is selected: • If the filter text is a string, the filter text matches only the entire string. • If the filter text is an IP address, the filter text matches only an IPV4 or IPV6 address that is the same as the filter text. • If the filter text is a port number, the filter text matches only a port number that is the same as the filter text.

The result of a search filter operation is a set of entries that match the filter criteria, sorted by time.

Customizing the audit log display

You can customize the audit log display to assist you in locating information faster.

- To customize the order of columns displayed, click any column header and drag the column to the location you want.
- To sort by column, click the name of the column you want to sort. Not all columns can be sorted. When sorting items in the Object Name column, partition names are ignored. For example, the object name `Common/rule1` would be sorted without the common partition name, as if it were named `rule1`.
- To resize columns, click the column side and drag it to the preferred location.
- To select what columns are displayed, click the gear icon in the upper right of the Audit Logging screen. In the popup screen, select columns you want to display and clear columns you do not want to display. Move your cursor away from the screen to dismiss it.

Managing audit log archive settings

You can view or change the audit archive settings. The archived audit log files are stored in the `/var/config/rest/auditArchive/` directory on the BIG-IQ® Centralized-Management system.

Note: *Before changing audit archive settings, verify that audit logs are not currently being archived or deleted. Changing the audit archive settings while logs are being processed can lead to unexpected numbers of logs being archived or deleted.*

1. At the top of the screen, click **Monitoring**.
2. On the left, expand **LOGS**, then expand **Audit Logs**, and then , click the component that you want to view audit entries for.
3. Click the **Archive Settings** button in the upper left of the Audit Logging screen to display the audit log settings.
4. Complete or review the properties and status settings, and click **Save**.

Property	Description
Retain Entries	Specifies the number of days to keep audit log entries. The field must contain an integer between 1 and 366. The default is 30.
Weekly Update	Specifies which days of the week to update the audit log. Select the check box to the left of each day that you want the audit log to be updated. The default is every day.
Start Time	Specifies when the audit archiving should begin. The default is 12:00 am.
Items Expired	Displays the read-only number of entries that have expired.
Last Error	If an error has occurred, displays the read-only error text for any errors found.
Last Error Time	If an error has occurred, displays a read-only value that contains the time the last error was found. The time in the field is the BIG-IQ system local time and is expressed in the format: ddd mmm dd yyyy hh:mm:ss, for example, `Fri Jan 17 2014 23:50:00`.

About archived audit logs

You can view or change how audit logs are archived by clicking the **Archive Settings** button on the Audit Logging screen.

Archived audit log files are stored in the `archive-audit.n.txt` file in the appropriate subdirectory of the `/var/config/rest/auditArchive` directory on the BIG-IQ® Centralized Management system:

- Network Security audit log: `/var/config/rest/auditArchive/networkSecurity/`
- Web Application Security audit log: `/var/config/rest/auditArchive/webAppSecurity/`
- Fraud Protection Service audit log: `/var/config/rest/auditArchive/websafe/`
- Local Traffic and Network audit log: `/var/config/rest/auditArchive/adc/`
- Device Management audit log: `/var/config/rest/auditArchive/device/`
- Access audit log:`/var/config/rest/auditArchive/access/`

Audit entries are appended to the `archive-audit.0.txt` file. When the `archive-audit.0.txt` file reaches approximately 800 MB, the contents are copied to `archive-audit.1.txt`, compressed into the `archive-audit.1.txt.gz` file, and a new empty `archive-audit.0.txt` file is created, which then has new audit entries appended to it.

Up to five compressed archived audit files can be created before those files begin to be overwritten to conserve space. The compressed audit log archive is named `archive-audit.n.txt.gz`, where n is a number from 1 to 5. As the audit log archives are created and updated, the content of the archives is rotated so that the newest archive is always `archive-audit.1.txt.gz` and the oldest is always the highest numbered archive, typically, `archive-audit.5.txt.gz`.

The file content rotation occurs whenever `archive-audit.0.txt` is full. At that time, the content of each `rchive-audit.n.txt.gz` file is copied into the file with the next higher number, and the content of `archive-audit.0.txt` is copied into `archive-audit.1.txt` and then compressed to create `archive-audit.1.txt.gz`. If all five `archive-audit.n.txt.gz`files exist, during the rotation the contents of `archive-audit.5.txt.gz` are overwritten, and are no longer available.

About audit logs in high-availability configurations

In high-availability (HA) configurations, there is a primary and secondary BIG-IQ® system. During failover, the audit log entries and the audit archive settings are copied from the primary to the secondary system before the secondary system becomes the new primary system.

However, archived audit logs are not copied from the primary to the secondary BIG-IQ system.

About the REST API audit log

The REST API audit log records all API traffic on the BIG-IQ® system. It logs every REST service command for all licensed modules in a central audit log (`restjavad-audit.n.log`) located on the system.

Note: *The current iteration of the log is named `restjavad-audit.0.log`. When the log reaches a certain user-configured size, a new log is created and the number is incremented. You can configure and edit settings in `/etc/restjavad.log.conf`.*

Any user who can access the BIG-IQ system console (shell) has access to this file.

Managing the REST API audit log

The REST API audit log contains an entry for every REST API command processed by the BIG-IQ® system, and is an essential source of information about the modules licensed under the BIG-IQ system. It can provide assistance in compliance, troubleshooting, and record-keeping. With it, you can review log contents periodically, and save contents locally for off-device processing and archiving.

1. Using SSH, log in to the BIG-IQ Network Security system with administrator credentials.
2. Navigate to the `restjavad` log location: `/var/log`.
3. Examine files with the naming convention: `restjavad-audit.n.log`.

   The letter *`n`* represents the log number.
4. Once you have located it, you can view or save the log locally through a method of your choice.

Access Reporting and Statistics

About Access and SWG reports

Access reports focus on session and logging data from Access devices (managed devices with APM licensed and provisioned). F5® Secure Web Gateway Services reports focus on user requests (for URLs or applications, for example) from Access devices with Secure Web Gateway Services provisioned. BIG-IQ® Centralized Management Access also supports high availability. Thus, users can view both Access and SWG reports on a secondary BIG-IQ system.

Access reports and SWG reports provide the following features.

- Reports on any combination of discovered devices, Access groups, and clusters
- Graphs for typical areas of concern and interest, such as cross-geographical comparisons or top 10 issues
- Tabular data to support the graphs
- Ability in some screens to drill down from summarized data to details
- Ability to save data to CSV files

Setup requirements for Access and SWG reports

Before you can produce Access reports and SWG reports, you must ensure that these tasks are already complete.

- Set up the BIG-IQ® Centralized Management data collection devices.
- Add the BIG-IP® devices to BIG-IQ inventory.
- Discover the devices. (Devices with the Access service configuration are what you need.)
- Run the data collection device configuration setup on the devices from the Access Reporting screen.

What data goes into Access reports for the All Devices option?

The **All Devices** option for Access reports includes data from the devices that are currently managed (discovered) in the BIG-IQ® system. This is in addition to data from devices that were managed at some point during the report timeframe, but that are not currently managed. With **All Devices** selected, if data from unmanaged devices exists, it displays in reports.

An unmanaged device might be unmanaged temporarily or permanently. Any time a configuration management change causes APM® to be undiscovered, the device and its data are moved to **All Devices** until APM is re-discovered on the device.

You cannot generate a report for an unmanaged device. However, you can generate a report for the timeframe when the device was managed, and then search the report for the unmanaged device name. In the Summary report, All Active Sessions includes the number of sessions that were active on the device when it became unmanaged. Those sessions stay in the Summary and in the Active sessions reports until the next session status update, which occurs every 15 minutes.

About upgrades affecting reports

When you upgrade a BIG-IQ® Centralized Management system without taking a snapshot, it deletes all reporting data, including both Access and SWG reports. After upgrading, users cannot obtain these reports from the BIG-IP® devices. To prevent the lost of reports, users should take an Elasticsearch snapshot before upgrading, and restore the snapshot after upgrading. For more information on elasticsnapshots, refer to *F5 BIG-IQ Centralized Management: Upgrading Logging Nodes to Version* x.x.

About the application dashboard

The Application Summary screen is your starting point to view and download general reports for BIG-IQ Access.

Viewing the application dashboard

The BIG-IQ® Centralized Management Access application dashboard displays information regarding the applications linked to the system.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access** > **Application Summary**.

The Application Summary screen displays, showing detailed information for specific applications.

About user visibility

You can monitor your user base by viewing the BIG-IQ® Centralized Management Access user dashboard for data on specific users. The system displays which users created the most sessions, were denied the most sessions, and had the longest total session duration. The administrator can enter a specific user name to get the following details for the user:

- The user login locations on a world map
- The total sessions, denied sessions, and session duration
- The Access denied sessions.
- The top authentication failures, including AD Auth and LDAP only
- The device type users used to log into the system
- The reason the system terminated the session
- The login history showing the success and failures over time
- The most accessed applications
- The most accessed URLs
- The login failure attempts over time, sorted by the reason
- The client session duration over time
- The Access denied reason over time

About application visibility

You can monitor your applications by viewing the BIG-IQ® Centralized Management Access user dashboard for data on which applications are linked to the BIG-IQ Access component. The system displays the top applications used and the application usage time. Administrators can expand the GUI for a specific application and view the following information:

- The application access history
- The users who use the application the most
- The access history
- The world map, showing where the user is access the application

About denied sessions

You can monitor the sessions that BIG-IQ® Centralized Management denies. By using the Access Monitoring option, you can view the following information:

- The history of denied sessions
- The reasons why sessions were denied
- The top denied users, sorted by session count
- The top authentication failures
- The top denied policies
- The top denied sessions by country of origin
- The top denied session by the virtual server
- The denied sessions, sorted by the client platform

Viewing denied sessions

You can use the BIG-IQ® Centralized Management Access reporting features to see which sessions were denied by the system, as well to create a report.

1. Log in to the BIG-IQ system with your user name and password.
2. Click **Monitoring** > **DASHBOARDS** > **Access** > **Sessions** > **Denied**.
3. From the **ACCESS GROUP/DEVICE** list at upper left, select **Managed Devices** or select one or more of these options:
   - ***<Access group name>*** - Select to include all devices in the Access group.
   - ***<Cluster display name>*** - Select to include the devices in the cluster.
   - ***<Device name>*** - Select to include the device. You can select any device from **Managed Devices**, ***<Access group name>***, or ***<Cluster display name>***.
4. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
5. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

From here, you can view details regarding denied sessions and create a report.

Managing a specific user in Access reporting

You can use the BIG-IQ® Centralized Management Access reporting tools to view the user dashboard for data on a specific user.

1. Log in to the BIG-IQ system with your user name and password.
2. Click **Monitoring** > **DASHBOARDS** > **Access** > **User Summary**.
   The User Summary screen displays, showing detailed information for specific users.

Running Access reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.

You can create Access reports for any device with the APM® service configuration on it that has been discovered on the BIG-IQ system, whether or not the device is a member of an Access group. To create a report, you can select any combination of Access groups, clusters, and devices.

1. At the top of the screen, click **Monitoring**.
2. On the left, select **DASHBOARDS** > **Access**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
3. From the left, select any report that you want to run.
4. At the top left of the screen, from the **ACCESS GROUP/DEVICES** list, either select one of the first two options (**All Devices** and **All Managed Devices**) or, select one or more of the other options (**<*Access group name*>**, **<*Cluster display name*>**, and **<*Device name*>**).
   - **All Devices** Includes Access devices that are currently managed, and Access devices that were managed at one time but are not managed now. (A managed device is one that has been discovered with the APM service configuration.)
   - **All Managed Devices** Includes all Access devices that are currently discovered.
   - **<*Access group name*>** - Select to include all devices in the Access group.
   - **<*Cluster display name*>** - Select to include the devices in the cluster.
   - **<*Device name*>** - Select to include the device. You can select any device from **Managed Devices**, **<*Access group name*>**, or **<*Cluster display name*>**.
5. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
6. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

Getting the details that underlie an Access report

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.

From the Summary report, and from most session reports, the initial display includes graphs that summarize the report data. You can get successively more detailed information by clicking a bar or a point on a graph or clicking a link if one is displayed on the screen.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access**.

   The Summary report is an example of the type of report that presents high-level data, and provides access to underlying data.

   A Summary report (for all devices and a default timeframe) starts to generate and display.
4. Click anywhere in a summary to get more information.

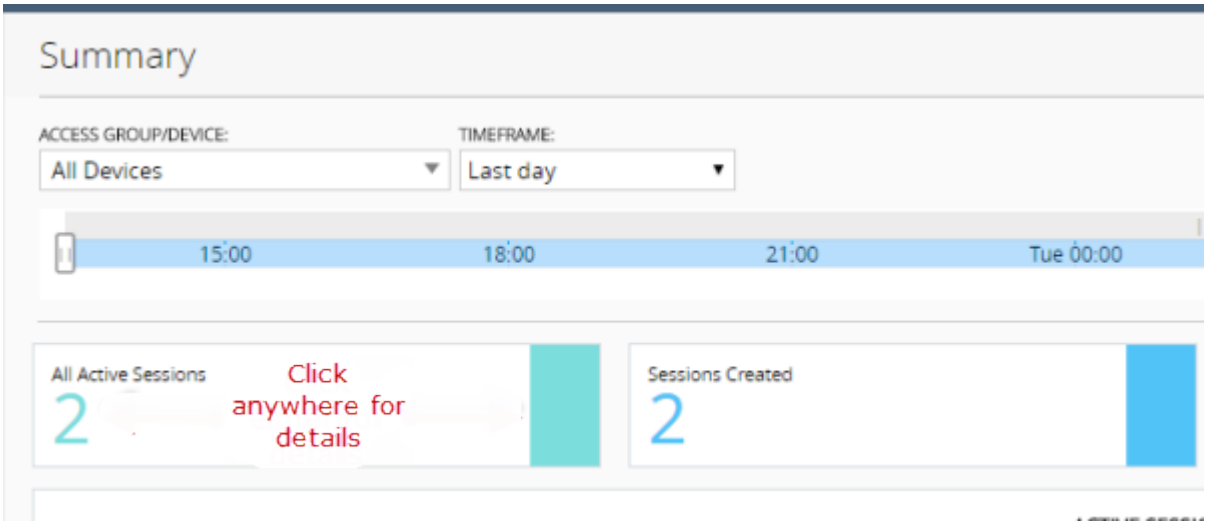

Figure 5: Top left portion of the Summary report display

Additional graphs display, and supporting data displays in a table at the bottom of the screen.

5. If more details are available, click the bars in the graphs to display more details.
6. Scroll down to the table to view the supporting data.
7. If the table includes a **Session ID** field, click the link in that field to open the session details.

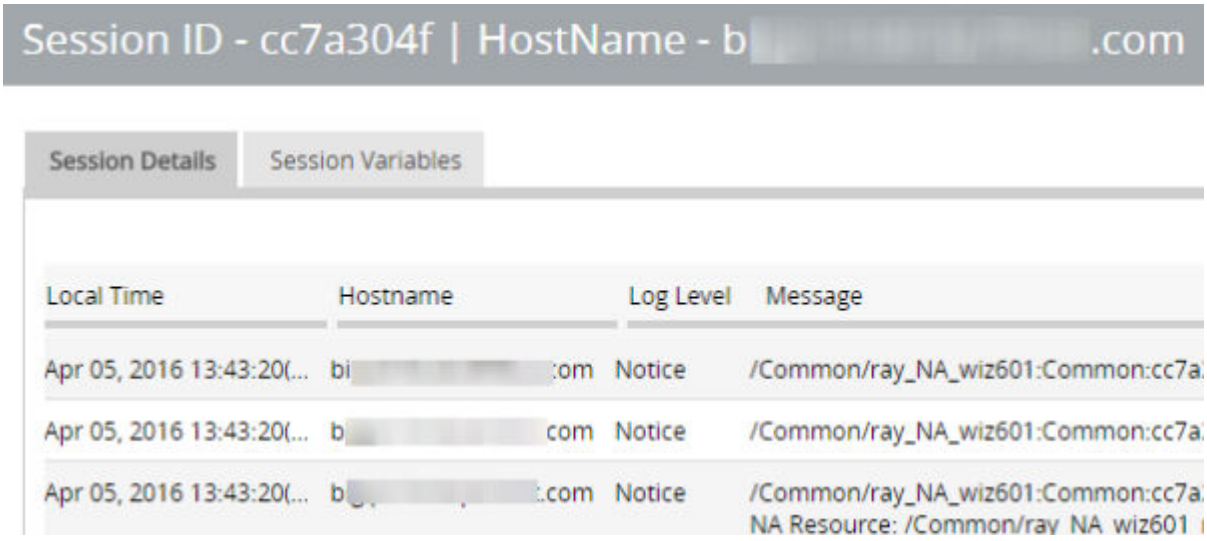

Figure 6: Session details popup screen (with addresses and host names blurred)

8. To change which records display on this screen, select a log level from the **LOG LEVEL** list at the top of the screen.

About the maximum number records for Access and SWG reports

When you run an Access report or an SWG report, Access can get up to 10,000 records to display to you. After you scroll to the end of those 10,000 records, Access displays a message. At that point, all you can do is select fewer devices or select a shorter timeframe.

Setting the timeframe for your Access or SWG report

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.

Use the **TIMEFRAME** list at the top of any Access or SWG report to change the report time period.

1. Log in to BIG-IQ Centralized Management with your admin user name and password.
2. At the top of the screen, click **Monitoring**.
3. To set a predefined timeframe, select one of these from the **TIMEFRAME** list: **Last hour**, **Last day**, **Last week**, **Last 30 days**, **Last 3 months**.
4. To set a custom timeframe, select one of these from the**TIMEFRAME** list:
   - **Between**: Click each of the additional fields that display to select dates and times. The report displays the records between those dates and times.
   - **Before**: Click the additional fields that display to select a date and a time. The report displays the records before that date and time.
   - **After**: Click the additional fields that display to select a date and a time. The report displays the records after that date and time.

Access report problems: causes and resolutions

Problem	Resolution
A session is over, but it continues to display in the Active sessions report.	If a session starts when logging nodes are up and working, but terminates during a period when logging modes are unavailable, the session remains in the Active sessions report for 15 minutes. After 15 minutes, the session status is updated and the session is dropped from the report.
Active sessions are included in the Summary and Active sessions reports for a device that is no longer managed.	Sessions were active on a device when it was removed from an Access group and became unmanaged. Sessions that were active when the device became unmanaged remain counted in All Active Sessions on the Summary screen and stay in the Active sessions report until the next session status update, which occurs every 15 minutes.
A session is over, but **Session Termination** and **Session Duration** are blank in a session report.	If a session starts when logging nodes are up and working but terminates during a period when logging nodes are unavailable, the session termination is not recorded and the session duration cannot be calculated.

What can cause logging nodes to become unavailable?

Logging nodes are highly available, but it is still possible for them to become unavailable. This could occur, for example, if all logging nodes are on devices in the same rack in a lab, and the power to the lab shuts down.

Sessions

Running Session reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.

You can create Session reports for any device with the APM® service configuration on it that has been discovered on the BIG-IQ system, whether or not the device is a member of an Access group. To create a report, you can select any combination of Access groups, clusters, and devices.

1. At the top of the screen, click **Monitoring**.
2. On the left, select **DASHBOARDS** > **Access** > **Sessions**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
3. From the left, select any report that you want to run.
4. At the top left of the screen, from the **ACCESS GROUP/DEVICES** list, either select one of the first two options (**All Devices** and **All Managed Devices**) or select one or more of the other options (**<*Access group name*>**, **<*Cluster display name*>**, and **<*Device name*>**).
   - **All Devices** Includes Access devices that are currently managed, and Access devices that were managed at one time but are not managed now. (A managed device is one that has been discovered with the APM service configuration.)
   - **All Managed Devices** Includes all Access devices that are currently discovered.
   - **<*Access group name*>** - Select to include all devices in the Access group.
   - **<*Cluster display name*>** - Select to include the devices in the cluster.
   - **<*Device name*>** - Select to include the device. You can select any device from **Managed Devices**, **<*Access group name*>**, or **<*Cluster display name*>**.
5. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
6. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

Stopping sessions on BIG-IP devices from Access

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.

You can stop currently active sessions on BIG-IP® devices, using the Active sessions report on the BIG-IQ system.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
4. On the left, from **Sessions**, select **Active**.
   The screen displays a list of active sessions for all devices.
5. To display sessions for particular devices, groups, or clusters only, select them from the **ACCESS GROUP/DEVICE** list at upper left.
   The screen displays the active sessions for the selected devices.
6. To stop specific sessions only, select the sessions that you want to end and click **Kill Selected Sessions**.
7. To stop all sessions, click **Kill All Sessions**.

Running Secure Web Gateway reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.Only a device with SWG provisioned on it can provide data for Secure Web Gateway reports.

You can create SWG reports for Access groups, clusters (in Access groups), or devices that you select from the Access groups and clusters (in Access groups) on the BIG-IQ system.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access** > **Secure Web Gateway**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
4. From the left, select any report that you want to run.
5. From the **ACCESS GROUP/DEVICE** list at upper left, select **Managed Devices** or select one or more of these options:
   - ***<Access group name>*** - Select to include all devices in the Access group.
   - ***<Cluster display name>*** - Select to include the devices in the cluster.
   - ***<Device name>*** - Select to include the device. You can select any device from **Managed Devices**, ***<Access group name>***, or ***<Cluster display name>***.
6. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
7. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

Getting the details that underlie an SWG report

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.Only a device with SWG provisioned on it can provide data for SWG reports.

From the Summary report, the initial display includes graphs that summarize the report data. You can get more detailed information by clicking a bar or a point on a graph to see additional graphs and tables with supporting entries.

1. Log in to BIG-IQ Centralized Management with your admin user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access** > **Secure Web Gateway**.
   The Summary starts to generate and display. A timeline and some summaries display across the top of the screen. Graphs display under the summaries. Each graph provide different views of the data.
4. Click any bar in a graph on the display to get more information.
   Additional graphs provide different views of the data, and supporting data displays in a table at the bottom of the screen.
5. If more details are available, click the bars in the graphs to display them.
6. Scroll down to the table to view the supporting data.

About VDI reports

You can monitor your virtual desktop infrastructure (VDI) by viewing the BIG-IQ® Centralized Management Access user dashboard for VDI applications, and creating a VDI report. The system

displays the top VDI applications used and the application usage time. Administrators can expand the UI for a specific application, and view the following information:

- The top 10 VDI applications
- The top users for the VDI applications
- The application usage history

Federation

Running OAuth reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.Only a device with OAuth provisioned on it can provide data for OAuth reports.

You can create OAuth reports for Access groups, clusters (in Access groups), or devices that you select from the Access groups and clusters (in Access groups) on the BIG-IQ® Centralized Management system.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access** > **Federation** > **OAuth**.
4. Select **Authorization Server**, **Client**, or **Resource**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
5. From the left, select any report that you want to run.
6. From the **ACCESS GROUP/DEVICE** list at upper left, select **Managed Devices** or select one or more of these options:
   - ***<Access group name>*** - Select to include all devices in the Access group.
   - ***<Cluster display name>*** - Select to include the devices in the cluster.
   - ***<Device name>*** - Select to include the device. You can select any device from **Managed Devices**, ***<Access group name>***, or ***<Cluster display name>***.
7. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
8. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

Running SAML reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® Centralized Management system, discovered, and had the Access remote logging configuration run for it.Only a device with SAML provisioned on it can provide data for SAML reports.

You can create SAML reports for Access groups, clusters (in Access groups), or devices that you select from the Access groups and clusters (in Access groups) on the BIG-IQ® Centralized Management system.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top of the screen, click **Monitoring**.
3. On the left, select **DASHBOARDS** > **Access** > **Federation** > **SAML**.

4. Select **SP Summary** or **IdP Summary**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
5. From the left, select any report that you want to run.
6. From the **ACCESS GROUP/DEVICE** list at upper left, select **Managed Devices** or select one or more of these options:
   - **<*Access group name*>** - Select to include all devices in the Access group.
   - **<*Cluster display name*>** - Select to include the devices in the cluster.
   - **<*Device name*>** - Select to include the device. You can select any device from **Managed Devices**, **<*Access group name*>**, or **<*Cluster display name*>**.
7. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
8. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

About Application Summary and traffic signatures

In BIG-IP Access, to view the Sharepoint and OWA application names in the Application dashboard correctly, update the classification signatures on the BIG-IP device. You can also customize your traffic signature by creating a custom signature.

Note: *Not updating the classification signatures on the BIG-IP device can prevent the Sharepoint and OWA application names from displaying.*

Overview: Updating classification signatures

Classification signatures define different types of traffic that Policy Enforcement Manager™ (PEM) can recognize, through Traffic Intelligence™. PEM™ recognizes a predefined set of signatures for common applications and application categories that are updated periodically. You can download signature updates from F5 Networks, and schedule the system to automatically update the signatures. You can also manually install the classification signatures and updates, for example, if the BIG-IP® system does not have Internet access.

Task Summary

Scheduling automatic signature updates

You can set up the BIG-IP® system to automatically update the classification signatures. This ensures that the system always has the latest classification signature files.

1. On the Main tab, click **Traffic Intelligence** > **Classification** > **Signature Update**.
   The Signatures screen opens.
2. Click **Check for Updates** to manually upload a signature file update if one is available.
   You see the current date and time in the **Latest Update Check** setting of the Signature Definitions area.
3. To upload a signature file update, in the Signature Definitions area, click **Import Signatures**.
   The Applications screen displays a **Signatures File** field where you can select the new signature file.
4. In the **Signatures File** field, click **Choose File** to navigate to the previously uploaded signatures file.
5. Click **Upload**.

A message displays indicating whether your upload was successful.

6. For the **Automatic Updates Settings**, in the **Signature Update** screen, select **Enabled**.
7. From the **Update Schedule** setting, select **Daily**, **Weekly**, or **Monthly** to specify how often you want the system to check for updates.
8. Click **Update** to save your settings.

The signature updates take effect immediately.

Overview: Creating custom classifications

Traffic Intelligence analyzes and identifies higher level protocols and applications. It has the ability to detect applications and protocols in Service Provider networks, for example, HTTP, popular P2P, and top categories (Audio/Video, File Transfer, Instant Messaging, Mail, P2P, Web). It provides an application update mechanism, which in turn, provides the ability to keep up with new, modified, or obsolete applications without going through software release upgrades. IP traffic classifications are based on the IP protocol field of the IP header (IANA protocol).

Note: *You can update the library (so) and signature definitions for web traffic (cpm) with hitless upgrade in Policy Enforcement Manager™ (PEM™).*

Task summary

Determining and adjusting traffic classifications

The BIG-IP® system classifies many categories of traffic and specific applications within those categories. You can determine which categories and applications of traffic the system can classify, and find out information about them such as their application or category ID.

1. On the Main tab, click **Traffic Intelligence** > **Applications** > **Application List**.
   The Applications screen displays a list of the supported classification categories.
2. To view the applications in each category, click the + icon next to the category.
3. To view or edit the properties of the application or category, click the name to open its properties screen.

   Tip: *Here you can view the application or category ID number.*

   ---
4. Click **Update** to save any changes.

Creating a category

On the BIG-IP® system, you can create customized categories for classifying traffic if the predefined categories are not sufficient for your needs. For example, if you plan to create new application types unique to your organization, you can create a category to group them together.

1. On the Main tab, click **Traffic Intelligence** > **Applications** > **Application List**.
   The Applications screen displays a list of the supported classification categories.
2. Click **Create**.
   The New Application screen opens.
3. From the **Type** list, select **Category**.
4. In the **Name** field, type a name for the classification category.
5. In the **Description** field, type optional descriptive text for the classification presets.
6. In the **Category ID** field, type an identifier for this category, a unique number.
7. For the **Application List** setting, move applications that you want to associate with this category from the **Unknown** list to the **Selected** list.

If the applications are not listed yet, you can associate the applications with the category when you create them.

8. Click **Finished**.

You have created custom applications to handle traffic.

Creating classification presets

On the BIG-IP® system, you can create classification preset settings for a classification policy that you have previously created.

1. On the Main tab, click **Traffic Intelligence** > **Presets**.
   The Presets screen displays a list of the supported classification categories.
2. Click **Create**.
   The New Presets screen opens.
3. In the **Name** field, type a name for the application.
4. In the **Description** field, type optional descriptive text for the classification presets.
5. For the **Policy** setting, move the classification policies from **Available** list to the **Selected** list, to create a new preset.
6. In the **Allow Reclassification** list, **Enabled** is the default selection.
7. In the **Flow Bundling** list, **Enabled** is the default selection.
8. In the **Cache Results** list, **Enabled** is the default selection.
9. Click **Finished**.

Creating a custom URL database

You can create a customized URL database that can be used for adding custom URLs and categories.

1. On the Main tab, click **Traffic Intelligence** > **Categories** > **Feed Lists**.
   The URL DB feed list screen opens.
2. Click **Create**.
   The New Feed List screen opens.
3. In the **Name** field, type a unique name for the URL feed list.
4. In the **Description** field, type optional descriptive text for the URL feed list.
5. In the **Category ID** field, select a category name from the drop-down list.
6. In the URL DB Location area, select the appropriate option for URL DB location.

Option	Description
File	Click the **Browse** button, and select the `customdb` file. The `customdb` file should be present on your machine, and is not present on the BIG-IP system. The `customdb` file is a CSV file of the format. The format is: URL/IPv4 [,cat1] [,cat2]...

Note: *The non-IP URL should have a IANA registered top level domain. The URL category ID should be in the form of integer and the range is 24576 to 32767.*

For example, sample lines of customdb entry is:

```
weather.gov, 28678
pconline.com.cn, 28679
kannadaprabha.com, 28680
yandex.ru, 28677, 28676, 28681
pitt.edu,28682
```

FTP	Type the ftp location and the **User** and **Password**.

Option	Description
HTTP	Type the HTPP location and the **User** and **Password**.
HTTPS	Type the HTPPS location and the **User** and **Password**.

7. In the **Poll Interval** field, type the time interval in seconds at which the url needs to be polled.
8. On the Main tab, click **Traffic Intelligence** > **Policies**.
   The Policy list opens.
9. Click **Create**.
   The New Policy screen opens.
10. In the **Name** field, type a unique name for the URL category policy.
11. In the **Description** field, type optional descriptive text for the URL category policy.
12. In the Feed List area, select the feed list that you created to attach to the policy.
13. Click **Finished**.

The category lookup is done in the custom database, and the URL list is loaded into the custom database through file input. You can also perform URL categorization by looking up the server name indication (SNI) in SSL traffic.

Using iRules with classification categories and applications

If you are using custom classification categories or applications, you can use iRules® to identify the traffic for the custom classifications, or you can initiate an action based on how the traffic is classified.

1. On the Main tab, click **Local Traffic** > **iRules**.
2. Click **Create**.
3. In the **Name** field, type a 1- to 31-character name.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.
   For example, to classify traffic as `xxx_app`, a custom classification application that you created, you can use this iRule:

```
when HTTP_REQUEST {
                        if { [HTTP::header "Host"] contains "xxx" }  {
                        CLASSIFY::application set xxx_app
                        }
                        }
                        }
```

   For example, to perform an action (in this case, drop) on traffic classified as `xxx_app`, you can use this iRule:

```
when CLASSIFICATION_DETECTED {
                        if { [CLASSIFICATION::APP == "xxx_app"]}  {
                        drop
                        }
                        }
```

   For complete and detailed information about iRules syntax, see the F5 Networks DevCentral web site `http://devcentral.f5.com`.
5. Click **Finished**.

After creating the iRules, you must assign them as resources for each relevant virtual server on the BIG-IP® system.

Modifying iRule event for URL categories

On the BIG-IP® system, you can modify iRules® Event settings for URL categories.

1. On the Main tab, click **Traffic Intelligence** > **Categories** > **Category List**.
2. Select a URL category.
   The URL Properties screen opens.
3. In the **Name** field, type a unique name for the URL category policy.
4. In the **Description** field, type optional descriptive text for the classification presets.
5. In the **Category ID** field, type an identifier for this category, a unique number.
6. For the **Application List** setting, move applications that you want to associate with this category from the **Unknown** list to the **Selected** list.

   If the applications are not listed yet, you can associate the applications with the category when you create them.
7. Click **Finished**.
8. On the Main tab, click **Local Traffic** > **Profiles** > **Classification**.
   The Classification screen opens.
9. Select a classification profile or create one.
10. From the **URL Categorization** field, select **Enabled** from the drop-down list.
11. In the **iRule Event** field, select the appropriate setting.
    - To trigger an iRule event for this category of traffic, select **Enabled**. You can then create an iRule that performs an action on this type of traffic.
    - If you do not need to trigger an iRule event for this category of traffic, select **Disabled**.

    Note: *`CLASSIFICATION::DETECTED` is the only event that is supported.*

You have modified an iRule event setting for an existing URL category.

Classification iRule commands

When the BIG-IP® system identifies a specific type of traffic with iRules® enabled, it triggers a `CLASSIFICATION_DETECTED` event. You can use the commands within iRules for additional system flexibility to classify the flow as one or more of the application or category classifications. The `CLASSIFY` commands are available from the `HTTP_REQUEST` or `HTTP_RESPONSE` iRule events.

iRule Command	Description
`CLASSIFICATION::app`	Gets the name of the classified application (the most explicit classified application).
`CLASSIFICATION::category`	Gets the category of the application.
`CLASSIFICATION::disable`	Disables the classification for a flow.
`CLASSIFICATION::enable`	Enables the classification for a flow.
`CLASSIFICATION::protocol`	Gets the name of the classified protocol (the least explicit classified application).
`CLASSIFY::application set` *`appname`*	Classifies the flow as *`appname`* and associates the category that *`appname`* belongs to.
`CLASSIFY::category set` *`catname`*	Classifies the flow as *`catname`* and also associates the flow with the unknown category.

iRule Command	Description
`CLASSIFY::application add` *`appname`*	Adds the application *`appname`* to the classification statistics.
`CLASSIFY::category add` *`catname`*	Adds the category *`catname`* to the classification statistics.

Managing Security Reports

About security reporting

Reporting for BIG-IQ® Network Security

You can use BIG-IQ® Network Security Reporting to view reports for managed BIG-IP® devices that are provisioned for Application Visibility and Reporting (AVR). Reports can be for a single BIG-IP device or can contain aggregated data for multiple BIG-IP devices (that are of the same BIG-IP device version).

Network Firewall, DoS and IP Intelligence reports can be created. Analytic reports provide detailed metrics about application performance such as transactions per second, server and client latency, request and response throughput, and sessions. Metrics are provided for applications, virtual servers, pool members, URLs, specific countries, and additional detailed statistics about application traffic running through one or more managed devices. You can view the analytics reports for a single device, view aggregated reports for a group of devices, and create custom lists to view analytics for only specified devices.

Reporting for BIG-IQ® Web Application Security

You can use BIG-IQ® Web Application Security Reporting to view reports for managed BIG-IP® devices that are provisioned for Application Visibility and Reporting (AVR). Similar to the availability of the AVR reporting on a single device, you have the ability to get visibility into application traffic passing through a single managed BIG-IP device or an aggregated system (aggregated data for multiple BIG-IP devices.

You can generate reports and charts in the following areas:

- Application. You can view information about requests based on applications (iApps), virtual servers, security policies, attack types, violations, URLs, client IP addresses, IP address intelligence (reputation), client countries, severities, response codes, request types, methods, protocols, viruses detected, usernames, and session identification numbers.
- Anomalies. You can view charts of statistical information in graphs about anomaly attacks, such as brute force attacks and web scraping attacks. You can use these charts to evaluate traffic to the web application, and to evaluate the vulnerabilities in the security policy.
- DoS. If you have configured DoS protection on the BIG-IP system, you can view charts and reports that show information about DoS attacks and mitigations in place on the system.

Determine DNS Sync Group Health

How do I check my sync group health?

Using the tools available on the BIG-IP® user interface, it can be difficult to determine the health of your DNS sync groups. When you use F5® BIG-IQ® Centralized Management to manage your DNS sync groups, the task becomes quite straightforward. You can do a quick health check, diagnose health issues, and even set up an alert to notify you if a sync group health issue occurs.

Check DNS sync group health

Before you can monitor the sync group health, you must add a BIG-IP® device configured in a DNS sync group to the BIG-IP Devices inventory list, and import the LTM® and DNS services.

When you use F5® BIG-IQ® Centralized Management to manage your DNS sync group, you can monitor the health status of the group. Sync group health relies on complete alignment of a variety of device configuration elements. Using BIG-IQ simplifies the process of determining the health of your DNS sync groups.

1. At the top of the screen, click **Devices**.
2. On the left, click **BIG-IQ CLUSTERS** > **DNS Sync Groups**.
   The screen displays the list of DNS sync groups defined on this device. A health indicator icon and a message describes the status of each group.
3. To view the general properties for a sync group, click the sync group name.

   Note: *For a list of Health Status error messages, refer to DNS sync group messages.*

   The screen displays the properties for the selected group. This screen shows an overview of your DNS sync group health. Under Status, you can see the current state (for example, `Required Services Down`, or `Health Check(s) Passed`) for each device in the group.
4. To view the health for an individual sync group member, click **Health** .
   The Health screen displays detailed information for each factor that contributes to the health of a DNS sync group. Following a definition of each factor, a Status row provides additional detail. For each indicator, the most serious issues impacting that indicator are listed first. Finally, if the status for a health indicator is not `Health Check(s) Passed`, the **Recommended Action** setting describes what you can do to correct the issue.
5. Resolve any reported issues on the managed devices, and then return to the DNS Sync Groups screen and click **Refresh Status**.
   Once you resolve all reported issues, the status for the DNS sync group changes to `Health Check(s) Passed`.

DNS sync group status messages

When BIG-IQ® Centralized Management completes health checks for a DNS sync group, an icon and a message display to indicate the current status. There are four icons, each with its own associated meaning.

Table 1: Health indicator icons

Icon	Meaning
	Indicates that all health checks passed satisfactorily (green).

Icon	Meaning
	Indicates that the health status is unknown or uncertain (blue).
	Indicates a warning, or that the group health is sub-optimal (yellow).
	Indicates that a critical issue was found (red).

Table 2: Health indicator messages

Message	Health indicator color	Description	Corrective Action
Awaiting Sync	Yellow	When considering the health of a DNS sync group, the single most important indicator of health is whether the devices in the sync-group have the same configuration in the master control program (MCP) daemon. *MCP* stores the configuration information for the BIG-IP® device. If the configuration is not the same (for devices in the sync group and MCP), then the devices could handle traffic differently, depending on what the configuration differences are.	Recommended Action: Wait a few minutes for synchronization to each member to occur. If synchronization does not complete, refer to troubleshooting solution. Related Solutions: SOL13690: Troubleshooting BIG-IP GTM synchronization and iQuery connections.
Certificate Expired	Red	BIG-IP DNS uses the device's Apache server certification to act as the server certification when establishing iQuery® connections. If this certificate expires, then all iQuery communication to and from this device is prevented. This indicator informs the DNS admin when one of the devices in a sync group has a device certificate that is near expiration, or is currently expired. This indicator only validates the expiration on the server certificate for each device. It does not examine the traffic certificates used in SSL profiles or DNSSEC certifications.	Renew the device certificate or import a new certificate. Related Solutions: SOL6353: Updating an SSL device certificate on a BIG-IP system.
Certificates Expiring	Yellow	The device certificate for this BIG-IP DNS device is near expiration. If the certificate expires, this BIG-IP DNS device will not be able to communicate with other BIG-IP devices using the iQuery protocol.	Either renew the device certificate or import a new certificate.
Changes Pending	Yellow	When considering the health of a DNS sync group, the single most important indicator of health is whether the devices	Recommended Action: Wait a few minutes for

Message	Health indicator color	Description	Corrective Action
		in the sync-group have the same configuration in the master control program (MCP) daemon. *MCP* stores the configuration information for the BIG-IP device. If the configuration is not the same (for devices in the sync group and MCP), then the devices could handle traffic differently, depending on what the configuration differences are.	synchronization to each member to occur. If synchronization does not complete, refer to troubleshooting solution. Related Solutions: SOL13690: Troubleshooting BIG-IP GTM synchronization and iQuery connections.
Collecting Data	Blue	Either the certificate has not yet been discovered by BIG-IQ or the device is unreachable.	If the certificate is the issue, the needed data should be collected automatically. If this condition persists, check the BIG-IQ logs for any error messages. If the device is unreachable, determine why BIG-IQ can not contact the BIG-IP device. There could be network issues, the device could be offline, or BIG-IQ Restjavad service could be is down.
Incompatible Device Versions	Red	A GTM sync group consists of one or more GTM devices. For sync to perform correctly, each device must have the same base version of TMOS installed. To determine the version of TMOS: view the version component of the output of `tmsh show sys version`.	Upgrade all BIG-IP devices in the sync group to the same version. Related Solutions: SOL8759: Displaying the BIG-IP Software Version. SOL13734: BIG-IP DNS synchronization group requirements.

Message	Health indicator color	Description	Corrective Action
Member Sync Disabled	Red	BIG-IP DNS devices have properties to control which sync group a device belongs to, and whether synchronization is enabled. A device can be a member of a sync group, but have synchronization disabled. Any changes made on a device on which synchronization is disabled cannot sync changes to the other devices. F5 recommends not having sync groups with synchronization disabled on some of the devices. We also recommend not making changes on devices if synchronization is disabled.	Enable synchronization on all devices in the group. Related Solutions: SOL13734: BIG-IP DNS synchronization group requirements.
Required Services Down	Red	For the BIG-IP DNS devices to be able to sync configuration changes, the following services (daemons) must be running on all the devices in the sync group: • `mcpd` • `gtmd` • `big3d` • `tmm` If any of these services is down, then configuration will not sync between the devices in the sync group. The sync group health is primarily concerned with reporting the health of only the sync group itself; not the health of the functionality provided by each device in the sync group.	Start stopped services Related Solutions: SOL13690: Troubleshooting BIG-IP DNS synchronization and iQuery connections Troubleshooting daemons.
Server Object Missing	Red	On the BIG-IP device, the DNS server objects define the IP address on which iQuery connections are made. There must be a server object for every DNS device in the sync group so that they can establish the necessary connections. This indicator validates that all devices have a server object, and that the necessary ports are open to allow the iQuery communication that happens over port 4353.	Verify that the DNS server objects have an associated self IP address. Related Solutions: SOL13734: BIG-IP DNS synchronization group requirements.
Syncing Changes	Yellow	When considering the health of a DNS sync group, the single most important indicator of health is whether the devices in the sync-group have the same configuration in the master control program (MCP) daemon. *MCP* stores the configuration information for the BIG-IP device. If the configuration is not the same (for devices in the sync group and MCP), then the devices could handle traffic	Recommended Action: Wait a few minutes for synchronization to each member to occur. If synchronization does not complete, refer to

Message	Health indicator color	Description	Corrective Action
		differently, depending on what the configuration differences are.	troubleshooting solution. Related Solutions: SOL13690: Troubleshooting BIG-IP GTM synchronization and iQuery connections.
Unknown Device Availability	Blue	The BIG-IQ device must collect data from each device in a sync group to be able to determine if the overall sync group is healthy. If BIG-IQ cannot reach one of the devices, then it cannot detect changes that make the overall group unhealthy. If a device cannot be reached, then the group is marked as unhealthy because there is no other way to know the health of the group.	Determine and fix loss of device availability. Related Solutions: SOL13690: Troubleshooting BIG-IP DNS synchronization and iQuery connections Troubleshooting daemons.
Unreachable Devices	Red	The BIG-IQ device must collect data from each device in a sync group to be able to determine if the overall sync group is healthy. If BIG-IQ cannot reach one of the devices, then it cannot detect changes that make the overall group unhealthy. If a device cannot be reached, then the group is marked as unhealthy because there is no other way to know the health of the group.	Determine and fix loss of device availability. Related Solutions: SOL13690: Troubleshooting BIG-IP DNS synchronization and iQuery connections Troubleshooting daemons.

How do I set up an alert for DNS sync group issues?

You can configure a BIG-IQ® SMTP alert to send email notifications when specific DNS sync group issues occur.

The following issues can trigger an alert:

- A new health status is generated for a DNS sync group. For instance, you might have just discovered a new sync group.
- The overall health status changes. For example, a device group that was healthy becomes unhealthy.
- The primary indicator (the most significant reason for the group's current health status) changed. (For example, the group is still unhealthy, but the reason is different than before.)

You enable or disable DNS alerts from the **System Management** > **Alerts** screen. For detailed instructions on creating an SMTP alert, refer to *How do I set up BIG-IQ to work with SMTP?* in the *F5 BIG-IQ Centralized Management: Licensing and Initial Setup* guide on `support.F5.com`.

Troubleshooting using iHealth

What is iHealth?

iHealth® is a tool that helps you troubleshoot potential issues. It does this by analyzing configuration, logs, command output, password security, license compliance, and so on.

From F5® BIG-IQ® Centralized Management, you can create a snapshot of a configuration in the form of a QKView file and then upload it to the F5® iHealth service. The file is compared to the iHealth database, which contains known issues, common configuration errors, and F5 published best practices. F5 returns an iHealth report you can use to identify any potential issues that you need to attend to.

Limit the number of simultaneous iHealth related file transfers to and from BIG-IQ

If you want to save system resources, you can easily limit how much traffic is dedicated to file activity related to iHealth® if you need to. You do this by specifying a limit of simultaneous file transfers to and from BIG-IQ® Centralized Management.

1. At the top of the screen, click **System**.
2. On the left, click **BIG-IQ iHealth** > **Configuration**.
3. Click the **Edit** button near the top of the screen.
4. In the **QKView Transfer Limit** field, type the greatest number of QKView files you want to occur on BIG-IQ at one time.
5. Click the **Save & Close** button at the bottom of the screen.

How do I get access to send QKView files for my managed devices to the F5 iHealth diagnostics server?

You'll need a single sign on (SSO) to the F5® Support site to access the F5 iHealth® diagnostics server. If you don't have one yet, register at *https://login.f5.com/resource/login.jsp*

With access to the F5 iHealth diagnostics server you can upload QKView files and download iHealth reports for your managed devices.

1. At the top of the screen, click **Monitoring**.
2. On the left, click **REPORTS** > **Device** > **iHealth** > **Configuration**.
3. Click the **Add** button.
4. In the **Name** field, type a name to identify this user.
5. In the **Username** and **Password** fields, type this user's F5 Support SSO user name and password.
6. In the **Description** field, type an optional description for this user.
7. Click the **Test** button to verify you can reach the iHealth diagnostics site.
8. Click the **Save & Close** button at the bottom of the screen.

You can now upload QKView files to the F5 iHealth server to get iHealth reports for your managed devices.

Troubleshoot issues for a managed device by uploading a QKView file to the F5 iHealth server

To upload a QKView file, you must have access to the F5® iHealth® server configured on BIG-IQ® Centralized Management.

You upload a QKView file to F5 Networks to create an iHealth diagnostics report. You can use that report to troubleshoot any potential issues with a managed device.

1. At the top of the screen, click **Monitoring**.
2. On the left, click **REPORTS** > **Device** > **iHealth** > **Uploads**.
3. Click the **Upload** button.
4. In the **Name** field, type a name to identify this report, and type an optional identifier in the **Description** field.
5. If you have (and want to associate) a support case number with this QKView file, type that into the **F5 Support Case Number** field.

   This step is not required.
6. From the **Credential** list, select the credentials to log in to the iHealth diagnostic site.
7. From the **Available** list, click the device you want to upload a QKView file for, and click **->** to move it to the **Selected** list.
8. Click the **Upload** button at the bottom of the screen.

When BIG-IQ finishes uploading the QKView file(s) to F5, it displays a status icon next to it on the Uploads screen.

If the upload fails, click the report's **Name** link and view the error message for more information. After F5 successfully receives the QKView file, it creates an iHealth report, which you can download from the Reports screen.

Download an iHealth diagnostics report for a managed device

F5® creates an iHealth® diagnostics report after you upload a managed device's QKView file to F5.

Downloading and reviewing an iHealth report helps you troubleshoot any potential issues with your managed devices.

Note: *The Reports screen displays a link only to the most recently created iHealth reports. The F5 iHealth server retains the report for approximately 30 days, after which it deletes the report, and the link from BIG-IQ® becomes invalid. This date is shown as the expiration date.*

---

1. At the top of the screen, click **Monitoring**.
2. On the left, click **REPORTS** > **Device** > **iHealth** > **Reports**.
3. In the **Report** column, click the **Download PDF** file link for the report you want.

BIG-IQ downloads the report you selected in the form of a PDF.

You can now open and review the iHealth diagnostics report for your managed device.

How do I get access to send QKView files for the BIG-IQ system to the F5 iHealth diagnostics server?

You'll need a single sign on (SSO) to the F5® Support site to access the F5 iHealth® diagnostics server. If you don't have one yet, register at *https://login.f5.com/resource/login.jsp*

With access to the F5 iHealth diagnostics server you can upload QKView files and download iHealth reports for your BIG-IQ® Centralized Management system.

1. At the top of the screen, click **System**.
2. On the left, click **BIG-IQ iHealth** > **Configuration**.
3. Click the **Add** button.
4. In the **Name** field, type a name to identify this user.
5. In the **Username** and **Password** fields, type this user's F5 Support SSO user name and password.

6. In the **Description** field, type an optional description for this user.
7. Click the **Test** button to verify you can reach the iHealth diagnostics site.
8. Click the **Save & Close** button at the bottom of the screen.

You can now upload QKView files to the F5 iHealth server to get iHealth reports for your BIG-IQ system.

Troubleshoot issues for the BIG-IQ system by uploading a QKView file to the F5 iHealth server

To upload a QKView file, you must have access to the F5® iHealth® server configured on BIG-IQ® Centralized Management.

You upload a QKView file to F5 Networks to create an iHealth diagnostics report. You can use that report to troubleshoot any potential issues with the BIG-IQ system.

1. At the top of the screen, click **System**.
2. On the left, click **BIG-IQ iHealth** > **Uploads**.
3. Click the **Upload** button.
4. In the **Name** field, type a name to identify this report, and type an optional identifier in the **Description** field.
5. If you have (and want to associate) a support case number with this QKView file, type that into the **F5 Support Case Number** field.

   This step is not required.
6. From the **Credential** list, select the credentials to log in to the iHealth diagnostic site.
7. From the **Available** list, click the device you want to upload a QKView file for, and click **->** to move it to the **Selected** list.
8. Click the **Upload** button at the bottom of the screen.

When BIG-IQ finishes uploading the QKView file(s) to F5, it displays a status icon next to it on the Uploads screen.

If the upload fails, click the report's **Name** link and view the error message for more information. After F5 successfully receives the QKView file, it creates an iHealth report, which you can download from the Reports screen.

Download an iHealth diagnostics report for the BIG-IQ system

F5® creates an iHealth® diagnostics report after you upload a BIG-IQ® Centralized Management QKView file to F5.

Downloading and reviewing an iHealth report for BIG-IQ helps you troubleshoot any potential issues for your BIG-IQ system.

Note: *The Reports screen displays a link only to the most recent BIG-IQ iHealth report created. The F5 iHealth server retains the report for approximately 5 days, after which it deletes the report, and the link from BIG-IQ becomes invalid. This date is shown as the expiration date.*

---

1. At the top of the screen, click **System**.
2. On the left, click **BIG-IQ iHealth** > **Reports**.
3. In the **Report** column, click the **Download PDF** file link for the report you want.

BIG-IQ downloads the report you selected in the form of a PDF.

You can now open and review the iHealth diagnostics report for the BIG-IQ system.

Legal Notices

Legal notices

Publication Date

This document was published on April 20, 2017.

Publication Number

MAN-0650-00

Copyright

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*.

Link Controller Availability

This product is not currently available in the U.S.

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

A

B

C

D

E

H

I

L

M

N

O

P

Q

R

S

T

U

V

W