Deploying the F5 Analytics iApp Template

Welcome to the F5 deployment guide for deploying F5 BIG-IP for easy export of statistical data. This document contains guidance on configuring the BIG-IP system with the Analytics iApp template. You can use this iApp template to marshal statistical and logging data from the BIG-IP system. The iApp takes this data and formats it as a JSON object which is then exported for consumption by data consumers, such as F5 BIG-IQ or applications such as Splunk®.

The Analytics iApp allows you to configure several categories of data to be exported. For data consumers like Splunk, the iApp lets you configure the network endpoint to which the data is sent.

Products and applicable versions

Product	Versions
BIG-IP LTM, AAM, APM, ASM, AFM	11.4 - 14.1
iApp version	f5.analytics.v3.7.2rc8
Deployment Guide version	2.8 (see *Document Revision History on page 32*)
Last updated	04-01-2020

Important: *Make sure you are using the most recent version of this deployment guide, available at http://f5.com/pdf/deployment-guides/f5-analytics-dg.pdf*

To provide feedback on this deployment guide or other F5 solution documents, contact us at *solutionsfeedback@f5.com*.

Contents

What is F5 iApp?

F5 iApp is a powerful set of features in the BIG-IP system that provides a new way to architect application delivery in the data center, and it includes a holistic, application-centric view of how applications are managed and delivered inside, outside, and beyond the data center. For more information on iApp, see the White Paper *F5 iApp: Moving Application Delivery Beyond the Network:* *http://www.f5.com/pdf/white-papers/f5-iapp-wp.pdf*.

Prerequisites and configuration notes

The following are general prerequisites and configuration notes for this guide:

- The configuration described in this deployment guide is supported by F5 Networks. F5 Technical support can help validate the configuration described in this guide if necessary, but your environment may have other factors which may complicate the configuration.
  If you need additional guidance or help with configuration that is not included in this guide, we recommend you consult your F5 FSE, check DevCentral (*https://devcentral.f5.com/*) and AskF5 (*https://support.f5.com/*), or contact F5 Professional Services (*https://f5.com/support/professional-services*) to discuss a consulting engagement. If you believe you have found an error in this guide, contact us at *solutionsfeedback@f5.com*.
- For this implementation, you must be using BIG-IP version 11.4-13.1.
- This deployment guide is meant to accompany iApp template versions f5.analytics.v3.7.2rc5. Previous F5-contributed versions of the template were available on DevCentral.
- If you are using BIG-IP APM and want to use the iApp to gather statistics, you must be on BIG-IP version 12.0 or later. This is a new requirement as of iApp version 3.7.0.
- This version of the iApp removed the ability to gather APM statistics from BIG-IP versions prior to 12.0. If you need to gather APM statistics, you must upgrade your BIG-IP system to v12.0 or later.
- All fields in the iApp template that have a small blue bar on the left column are required.
- One of the data formats you can choose in the iApp template is **Splunk**. However, you can use the resulting JSON file with other data consuming applications.
- iApp version v3.7.1rc1 added support for sending data to Sumo Logic.
- If you have used a previous version of the Analytics iApp template, we recommend you upgrade to the current version. The upgrade process retains all of the values you entered in the previous version of the template. For instructions, see *Upgrading an Application Service from previous version of the iApp template on page 5*.
- If you choose to use the iApp to send iHealth snapshot information, your iHealth username and password may be stored in cleartext in the BIG-IP configuration and/or the scriptd.out log
- BIG-IP supports only one instance of the Analytics iApp.

Configuration example

The iApp template allows to you select the data format for the data consumer you are using, and then granularly select the type of statistics and information you want to collect. The iApp produces a JSON file which can be used by your data consumer. If using Splunk specifically, see *Appendix: Using Splunk with the data on page 30.*

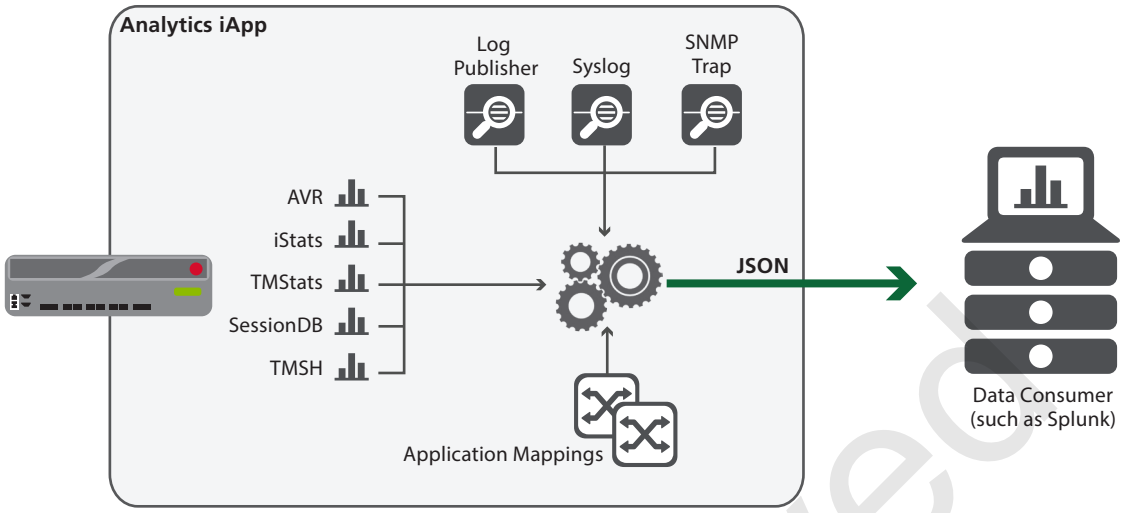

Configuring the Analytics iApp template

Use the following guidance to help configure the BIG-IP system with the iApp template.

Downloading and importing the new iApp

The first task is to download and import the Analytics iApp template.

To download and import the iApp

1. Open a web browser and go to *downloads.f5.com*.
2. Click **Find a Download**, and then in the **F5 Product Family > BIG-IP** section, click **iApp Templates**.
3. From the **Select a Product Version and Container** page, click **iApp-Templates**.
4. Accept the EULA, and then download the iapps zip file to a location accessible from your BIG-IP system. We strongly recommend using **f5.analytics.v3.7.1rc4** (or later) available in the **Release_Candidates** folder.
5. Extract (unzip) the **f5.analytics.v<latest version>.tmpl** file.
6. From the BIG-IP system web-based Configuration utility.
7. On the Main tab, expand **iApp**, and then click **Templates**.
8. Click the **Import** button on the right side of the screen.
9. Click a check in the **Overwrite Existing Templates** box.
10. Click the **Browse** button, and then browse to the location you saved the iApp file.
11. Click the **Upload** button. The iApp is now available for use.

Upgrading an Application Service from previous version of the iApp template

If you configured your BIG-IP system using a previous version of the f5.analytics iApp template, we strongly recommend you upgrade the iApp template to the most recent version.

When you upgrade to the current template version, the iApp retains all of your settings for use in the new template. In some new versions, you may notice additional questions, or existing questions asked in different ways, but your initial settings are always saved.

To upgrade an Application Service to the current version of the template

1. From the Main tab of the BIG-IP Configuration utility, expand **iApp** and then click **Application Services**.
2. Click the name of your existing f5.analytics application service from the list.
3. On the Menu bar, click **Reconfigure**.
4. At the top of the page, in the **Template** row, click the **Change** button to the right of the list.
5. From the **Template** list, select **f5.analytics.<latest version>**.
6. Review the questions in the new template, making any necessary modifications. Use the iApp walkthrough section of this guide for information on specific questions.
7. Click **Finished**.

Getting Started with the iApp

To begin the iApp Template, use the following procedure.

1. Log on to the BIG-IP system.
2. On the Main tab, expand **iApp**, and then click **Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **analytics-iapp_.**
5. From the **Template** list, select **f5.analytics.<latest version>**. The iApp template opens.

Advanced options

If you select **Advanced** from the **Template Selection** list at the top of the page, you see Device and Traffic Group options for the application. This feature is a part of the Device Management configuration. This functionality extends the existing High Availability infrastructure and allows for clustering, granular control of configuration synchronization and granular control of failover. To use the Device and Traffic Group features, you must have already configured Device and Traffic Groups before running the iApp. For more information on Device Management, see the product documentation.

We recommend you leave these settings at the defaults.

1. ***Device Group***
   To select a specific Device Group, clear the **Device Group** check box and then select the appropriate Device Group from the list.

2. ***Traffic Group***
   To select a specific Traffic Group, clear the **Traffic Group** check box and then select the appropriate Traffic Group from the list.

Welcome to the Analytics Template

This section contains information and general questions about the way you configure the iApp template.

1. ***Do you want to see inline help?***
   Choose whether you want to see informational and help messages inline throughout the template, or if you would rather hide this inline help. If you are unsure, we recommend having the iApp display the inline help.
   Important and critical notes are always shown, no matter which selection you make.

   - **Yes, show inline help text**
     Select this option to see all available inline help text.
   - **No, do not show inline help text**
     If you are familiar with this iApp template, or with the BIG-IP system in general, select this option to hide the inline help text.

2. ***Do you want to display advanced options?***
   Select whether you want to have more granular, advanced options presented in the iApp.

   - **No, do not show advanced options**
     Select this option if you do not need to configure the advanced options in the template. Continue with the next section.
   - **Yes, show advanced options**
     In advanced configuration mode, you see additional options that are not required for a working configuration.

     Advanced options in the template are marked with the Advanced icon: Advanced. If you are not using the Advanced settings, you can skip the questions with the Advanced icon.

Information Sources

In this section, you supply the information sources for implementation. This includes the format you want to use to send the data, and which type of statistics you want to send.

1. ***Data Format***
   Choose which data format you want the system to use to send data. This option allows you to specify the system type to which the BIG-IP sends data. If you do not see your logging format in the following list, we recommend using Splunk.

   - **Splunk**
     Select this option if you are using a data consumer like Splunk. If using Splunk specifically, see *Appendix: Using Splunk with the data on page 30* for information on setting up Splunk to display the analytics data.
   - **Sumo Logic**
     Select this option if you are using Sumo Logic. This option was introduced in iApp version v3.7.1rc1.
   - **F5 Analytics**
     This option supports certain F5 functionality which requires additional software from F5. Consult your F5 Sales or Professional Services representative for specific details.
   - **F5 Risk Engine**
     Select this option if you want to use F5 Risk Engine.
   - **BIG-IQ**
     Select this option if your data consumer is F5 BIG-IQ. The BIG-IP system you are configuring pulls the data and makes it available to the BIG-IQ system.

2. ***Unique ID*** **Advanced**
   This field does not appear if you selected BIG-IQ as the data format

   Type your unique ID for this implementation. The Unique ID advanced option accommodates the ability to deploy the f5.analytics iApp more than once. Each deployment of the iApp deploys virtual servers to process log-based traffic. These virtual servers use the same IP address for all deployments, but on different ports. This Unique ID is the seed number for port assignment.

   The first iApp deployment should use port 1000. You should increment this number by 1000 each time for additional deployments, for example 2000, 3000, 4000, and so on.

3. ***Log Stats Responses*** **Advanced**
   Choose whether you want to log stats responses. If you select Yes, a BIG-IP log message is written after every time a statistics update message is sent to a data consumer.

   - **No**
     Select this option if you do not want the system to log stats responses.
   - **Yes**
     Select this option if you want the system to log stats responses. This BIG-IP system writes a log message every time a statistics update message is sent to a data consumer.

4. ***System Statistics***
   Choose whether you want to enable system statistics. The System Statistics option enables analytics collection. This is the main component of the f5.analytics iApp. Full collection of statistics can be disabled for the entire iApp you only want the logging component, high speed logging streams, and so on. If you select No, system statistics are not collected. To ensure the proper functionality of most reports, we typically recommend you select **Yes**.

   - **No**
     Select this option if you do not want to enable system statistics. You can still enable other logging options.
   - **Yes**
     Select this option if you want to enable system statistics. When you select Yes, the *Analytics System Configuration on page 13* appears further down the template to configure the options.

5. ***Module High Speed Logging Streams***
   This field does not appear if you selected BIG-IQ as the data format
   Choose whether you want to capture module high speed logging (HSL) streams. Capturing Module HSL Streams creates the configuration objects necessary to allow for BIG-IP module logs collection. You should only enable this option if you want the system to create a Log Publisher that you can reference from any BIG-IP module. If you only want logs related to items such as system management plus messages automatically logged during TMOS traffic processing or conventionally logged by iRules, you can select No, and then configure syslog (in the next question).

   - **No**
     Select this option if you do not want to capture high speed logging streams using this iApp. You can still create a log publisher manually outside this iApp, but you have to manage it separately.

   - **Yes**
     Select this option to enable Module HSL Streams. When you select Yes, the *Module Log Stream Capture on page 16* appears further down the template to configure the options.
     Selecting Yes causes the iApp to create the following objects:

     » *A Logging Publisher* (logging-publisher-f5_analytics)
     This log publisher is available for assignment to BIG-IP APM, AFM and other modules that send to log publishers. To get module logs you must configure the BIG-IP to use this Log Publisher. For each BIG-IP module, such as APM or AFM, you select this log publisher in the module configuration after you complete the iApp. See the appropriate module documentation for details.

     » *A Pool*
     This pool can be used for High Speed Log (HSL) iRule commands to send Key Value Pair data to: <iApp Name>-hec-forwarder-tcp-log-stage0.

     » *A virtual server*
     This virtual server can be used to send ASM events to: 255.255.255.254:1001
     If you are using a custom Unique ID, the port number is the Unique ID + 1.

     To get BIG-IP ASM logs, you must configure ASM to send logs via TCP Key Value Pairs to this IP and Port.

6. ***Local System Logging (syslog)***
   This field does not appear if you selected BIG-IQ as the data format

   Choose whether you want to enable local system logging (syslog). Enabling Local System Logging (syslog) forwards syslog messages from the BIG-IP.

   - **No**
     Select this option if you do not want to enable syslog. Selecting No for syslog is an easy way to reduce the volume of data being sent. As an alternative, enabling System SNMP Alerts is a way to get a reduced amount of critical device data.

   - **Yes**
     Select this option to enable local system logging. When you select Yes, the *Local Logging Capture and System on page 19* appears further down the template to configure the options.

7. ***System SNMP Alerts***
   This field does not appear if you selected BIG-IQ as the data format

   Choose whether you want to enable system SNMP alerts. Enabling System SNMP Alerts forwards SNMP traps from the BIG-IP.

   - **No**
     Select this option if you do not want to enable system SNMP alerts.

   - **Yes**
     Select this option if you want the system to forward SNMP traps from the BIG-IP system. When you select Yes, the *SNMP Alert Capture on page 21* appears further down the template to configure the options.

8. ***iHealth Snapshot Information***
   This field does not appear if you selected BIG-IQ as the data format

   Choose whether you want to enable iHealth snapshot information. Enabling iHealth Snapshot Information allows the iApp to send a QKView to F5's iHealth infrastructure on a scheduled basis and for the analytics system to collect the iHealth Diagnostics data.

For more information on iHealth, see https://f5.com/support/tools/ihealth.

Note: *The BIG-IP must have access to ihealth.f5.com directly or via a proxy.*

- **No**
  Select this option if you do not want to enable iHealth snapshot information.
- **Yes**
  Select this option if you want the system to send a QKView to F5's iHealth system. When you select Yes, the *iHealth Snapshot Information Capture on page 23* appears further down the template to configure the options.

9. ***Facility Name***
   This field does not appear if you selected BIG-IQ as the data format

   Type the facility name to identify where the BIG-IP or BIG-IP cluster is located. This is just a string provided for your own use and does not affect operation.

   Spaces are permitted in this field, for example: SFO, New York, DC1, ORD-DC2, and so on.

10. ***Analytics System Tenant***
    This field only appears if you choose F5 Analytics or Sumo Logic from the Data Format list

    When using the F5 Analytics Data Format, you must specify an additional required Analytics System Tenant. The value in this field identifies 'this' BIG-IP to the analytics system, which refers to various BIG-IP devices as tenants.

    This Analytics System Tenant is different from the Default Tenant used for mappings. The Analytics System Tenant is designed for data isolation within the F5 Analytics Platform.

    Spaces and special characters are **not** permitted.

11. ***Default Tenant***
    This field does not appear if you selected BIG-IQ as the data format

    You can optionally specify the default Tenant. A Tenant is used to group statistics for Applications together based on logical sets. This optional field allows simple configuration of a tenant if all applications, virtual servers, wide IPs, and other objects are part of the same Tenant.

    If the BIG-IP or BIG-IP cluster is used by multiple tenants, then more advanced Tenant mapping can be done later in the template and this should remain blank. Note: Spaces and special characters are permitted.

    Things to note about tenants:

    - An managed service provider (MSP) may divide tenants by customer name, such as ACME or CompanyCo.
    - Others may divide tenants by environment, such as Production, Staging, and so on.
    - Further Business units could divide tenants, such as Web Hosting, Internal Service, or PCI.
    - A combination of all or some of those, such as “ACME - Web Hosting - Production”.

12. ***Alternative Device Group***
    You can ***optionally*** specify an alternate Device Group. This allows for the device group to be overridden from the device group name configured by the system outside the iApp. By default, the device group name is set to **Sync-Failover** which is not a unique or memorable representation of the device group/cluster. You can use an alternative Device Group to correct this without any changes needed to the cluster. No changes are made to the system Device Group configuration.
    Note: Spaces and special characters are permitted.

    For example, if a cluster is made up of DC-BIGIP21.dc.company.com and DC-BIGIP22.dc.company.com, the Alternative Device Group could be *DC-BIGIP20.company.com*. Or if the cluster is made up of PWZ1.nyc.company.com and PWZ2.nyc.company.com the Alternative Device Group could be *NYC Production Web Zone Cluster*.

13. ***Alternate Hostname*** Advanced
    This field does not appear if you selected BIG-IQ as the data format

You can ***optionally*** specify an alternate hostname. This option only functions in a standalone deployment. It allows the host name of the BIG-IP to be overridden when the iApp sends statistics and events, without any changes needed to the BIG-IP.

14. ***Analytics Callback Integration*** **Advanced**
This field does not appear if you selected BIG-IQ as the data format

Choose how you want links back to the BIG-IP system to be generated. Some items displayed in generated reports contain links that direct you back to the BIG-IP for additional information. By default, the management IP address is used as the hostname in each link. However, if you select Static, you can specify a Callback URL. If configuration changes should be made via a floating Self IP address or a specific URL, you should select Callback URL.

- **Use Management**
Select this option if you want to use the Management IP address for links back to the BIG-IP system.

- **Static**
Select this option if you want to specify a static IP address for building these links that would otherwise go back to the BIG-IP system.

    a. *Callback URL*
    Specify the callback URL you want to use. This allows an alternative URL to be used when generating links within the visualization dashboards. Note: This URL is shared between all members of a device cluster. You can use an IP address (for example, https://10.1.1.1/ or https://10.1.1.1:9000/ or an FQDN (such as https://bigipcluster1.example.com/).

15. ***Role Based Access Controls***
*This option only appears if you selected **Splunk** as the Data Format*

Choose whether you can to use Role Based Access Controls (RBAC). This option allows for isolation of statistics data by tenant. When sending data to Splunk, RBAC causes the data to be associated to a specific index name. This allows Splunk to control which indexes can be accessed by which users. By default no Splunk index is sent, which means the system uses the default HTTP Event Collector index for the authentication token passed.

- **No**
Select this option if you do not require RBAC. Continue with *Analytics System Configuration on page 13*.

- **Yes**
Select this option if you want to use RBAC. You configure specific options in *Role Based Access Control (RBAC) on page 11*.

Role Based Access Control (RBAC)

This entire section only appears if you selected Splunk as the Data Format, and to enable RBAC.

If you are using Role Based Access Controls (RBAC), indexes are used as containers to which users have selective access. There are many options with RBAC design based on your architecture and desired permission isolation. At a high level, there are two main paths.

- If this BIG-IP or BIG-IP Cluster is used for a single tenant, then setting a Default Tenant and setting all values to Use Default Tenant is suggested
- If this BIG-IP or BIG-IP Cluster is used for multiple tenants than a more complex index model is needed. Selecting Specify is suggested. An overview of each index option is provided below for this multi-tenant use case.

Note: Dots (.) and spaces ( ) are replaced with underscores (_) when being set in the index name.

1. ***Use an Index Prefix?***
   Select whether you want to set an Index prefix. The Index Prefix allows an optional prefix to be configured for all index values sent from the iApp. This allows simple index search configuration within the f5_index macro, such as **index=f5-*** which lets you select BIG-IP data more efficiently in Splunk or other data consumers.

   - **No**
     Select this option if you do not want to specify an index prefix.
   - **Yes**
     Select this option if you want to set an index prefix. You must specify the prefix in the following field.

     a. *Index prefix*
        Type the prefix for the index you want to use.

2. ***Systems Statistics Index***
   Choose whether you want to use the default tenant or a static tenant for the system statistics index. System statistics, which are those not capable of being mapped based on virtual name, wide IP name, partition, and so on (such as CPU Interface Statistics), can be configured to use a specific index. Regardless of the deployment scenario you can always specify a complex name allowing for more index isolation, such as *production_system_stats*.

   The recommended index depends on whether you have a single tenant or multi tenant BIG-IP:

   Single tenant
   Select **Use Default Tenant** to set the index used for System Statistics. This sets the index to *<IndexPrefix><Default Tenant>*.

   Multi Tenant BIG-IP
   Select **Specify** and set an index name to be used for System Statistics, such as **system_stats**. This example would set the index to *<IndexPrefix>system_stats*.

   - **Use Default Tenant**
     Select this option to use the default tenant. The index is set to *<IndexPrefix><Default Tenant>*.
   - **Specify**
     Select this option if you want to specify a static tenant for system statistics. You must specify index name in the following field.

     a. *System Statistics index*
        Type the index name you want to use.

3. ***Default Statistics Index***
   Choose whether you want to use the default statistics index or a static tenant for the default statistics index. The default statistics are statistics for objects that could not be mapped to a Tenant for whatever reason. This ensures there is an index to use for these unmapped statistics. Regardless of the deployment scenario you can always specify a complex name allowing for more index isolation, such as *companyx_unknown_stats*.

   The recommended index depends on whether you have a single tenant or multi tenant BIG-IP:

   Single tenant
   Select **Use Default Tenant** to set the index used for default statistics. This sets the index to *<IndexPrefix><Default Tenant>.*

Multi Tenant BIG-IP
Select **Specify** and set an index name to be used for System Statistics, like **unknown_stats**. This example would set the index to *<IndexPrefix>unknown_stats*.

- **Use Default Tenant**
  Select this option to use the default tenant. This sets the index to *<IndexPrefix><Default Tenant>*.
- **Specify**
  Select this option if you want to specify a static tenant for default statistics. You must specify index name in the following field.

  a. *System Default index*
     Type the index name you want to use.

4. ***Shared Statistics Index***
   Choose whether you want to use the default shared statistics index or a static tenant for the default statistics index. Shared statistics, which are those for objects that are used across tenants, can be configured to use a specific index. Note regardless of the deployment scenario you can always specify a complex name allowing for more index isolation, such as *development_shared_stats*.

   The recommended index depends on whether you have a single tenant or multi tenant BIG-IP:

   Single tenant
   Select **Use Default Tenant** to set the index used for default statistics. This sets the index to *<IndexPrefix><Default Tenant>.*

   Multi Tenant BIG-IP
   Select **Specify** and set an index name to be used for Shared Statistics, such as **shared_stats**. This example sets the index to *<IndexPrefix>shared_stats*.

   - **Use Default Tenant**
     Select this option to use the default tenant. This sets the index to *<IndexPrefix><Default Tenant>*.
   - **Specify**
     Select this option if you want to specify a static tenant for shared system statistics. You must specify index name in the following field.

     a. *Shared Statistics index*
        Type the index name you want to use.

5. ***Event Index***
   Choose whether you want to use the default event index or a static tenant for the event index. Events such as syslog, SNMP Traps, or module logs are not currently capable of being isolated to a specific tenant on a multi tenant BIG-IP. However, an index can be configured to allow for RBAC permission configuration to event data. On a single tenant BIG-IP this is not an issue. Regardless of the deployment scenario you can always specify a complex name allowing for more index isolation such as *development_events*.

   The recommended index depends on whether you have a single tenant or multi tenant BIG-IP:

   Single tenant
   Select **Use Default Tenant** to set the index used for events. This sets the index to *<IndexPrefix><Default Tenant>.*

   Multi Tenant BIG-IP
   Select **Specify** and set an index name to be used for events, such as events. This example sets the index to <IndexPrefix>events.

   - **Use Default Tenant**
     Select this option to use the default tenant. This sets the index to <IndexPrefix><Default Tenant>
   - **Specify**
     Select this option if you want to specify a static tenant for the event index. You must specify index name in the following field.

     a. *Shared Statistics index*
        Type the index name you want to use.

Analytics System Configuration

In this section, you configure the analytics system. This is the "Main Configuration" referred to in some of the options in the iApp.

1. ***IP Address or Hostname***
   This field does not appear if you selected BIG-IQ as the data format

   Specify the the IP Address or hostname of the Splunk Indexing server to which you want to send data. This can be a virtual server or a DNS name that distributes requests to multiple index servers, but the Authentication (API) Key must be the same for all indexers. Note BIG-IP version 11.6.0 and higher supports FQDNs. Previous versions require an IP address.

   If using **Splunk Cloud** this is **input-<yourinstancename>.cloud.splunk.com** for BIG-IP version 11.6.0 and higher. For previous versions, use the IP address to which this resolves.

2. ***Hostname when destination is an IP address*** **Advanced**
   This field does not appear if you selected BIG-IQ as the data format

   You can optionally provide a HTTP host header to be used inside of the communication. This is useful for BIG-IP versions before 11.6 where FQDNs are not supported. Providing a host header ensures that proxies and other security devices intercepting this traffic, including the HTTP Event Collector, understand the communication.

3. ***Port***
   This field does not appear if you selected BIG-IQ as the data format

   Specify the port number to send data to HTTP Event Collector (HEC). This is configured in the global settings for HEC. By default this is **8088**, including for Splunk Cloud.

4. ***Protocol***
   This field does not appear if you selected BIG-IQ as the data format

   Select the Protocol to use for transmission. If using Splunk, this is configured in the global settings for HEC; by default this is **HTTPS**, including for Splunk Cloud. Using the HTTPS-based protocol encrypts all transmission of data.

5. ***API Key***
   This field does not appear if you selected Sumo Logic or BIG-IQ as the data format

   Specify the API key to use when authenticating.

   If using Splunk, you also need to have enabled Splunk API access and obtained your Splunk instance's API Key. To do this, log into your Splunk instance as administrator and enable API Access in global settings. Your API key is found in **Settings -> Data Input -> HTTP Event Collector**. See *Splunk HTTP Event Collector configuration on page 31* for more information.

6. ***Push Interval (in seconds)***
   Specify a number of seconds for the push interval. This controls how often statistics are sent. In most cases it also controls the granularity of statistics, for example a setting of 60 will allow for 60 second granularity. AVR generated statistics are the only statistics that this setting does not control, as AVR granularity is set to a default of 300 seconds.

   Performance and data-volume tip:
   Increasing this interval from 60 to 300 seconds will dramatically reduce the amount of data that is sent. Increasing this to 10 minutes (600 seconds) reduces it further. A value lower than 30 seconds is not advised; the average collection and sending process takes roughly 5-15 seconds depending on size of the configuration.

7. ***Randomize the start time within the push interval?***
   Choose whether you want to randomize the start time within the push interval you specified. For example, if you select Yes, and chose 60 seconds, instead of starting the push at 0 seconds, then 60 seconds, then 120 seconds, the system may start the push at 15 seconds, then 75 seconds, then 135 seconds, and so on.

   - **No**
     Select this option if you do not want the system to randomize the start time of the push interval
   - **Yes**
     Select this option if you want the system to randomize the start time.

8. ***Use an HTTP Proxy?***
   This field does not appear if you selected BIG-IQ as the data format

   Choose whether you want to use an HTTP proxy in this implementation. Use of an HTTP Proxy is supported for sending all data other than events to Splunk. To send data via a web proxy server to Splunk select Yes. In most cases and environments, a proxy server is not required. You may use the same or another proxy server to send iHealth data (see *iHealth Snapshot Information Capture on page 23)*.

   - **No**
     Select this option if you are not using an HTTP proxy. Continue with #8.
   - **Yes**
     Select this option if you want to include an HTTP proxy in this implementation. You must complete the following.

     a. *Proxy IP Address*
        Type the IP address or FQDN of the HTTP proxy server.

     b. *Proxy Port*
        Type the port used by the HTTP proxy server.

     c. *Proxy Username*
        Type a user name with access to the proxy. This field is optional.

     d. *Proxy Password*
        Type the password for the username you specified. This field is optional.

The rest of this section allows you to modify the statistical data is that is sent by the system. The iApp allows for control of what statistical data is collected and sent as well as some ability to customize some of the statistical options. This is yet another way to reduce/remove data collection and transfer of data that you do not want visualized.

9. ***Push Device State?*** **Advanced**
   Choose whether you want to push device state statistics. These are typically lightweight/small: up to approximately 1000 characters. Device state information includes statistics such as active/standby, serial number, module status, iApp version and so on. The source is *bigip.tmsh.system_status*.

10. ***Push Device Statistics?*** **Advanced**
    Choose whether you want to push device statistics. These are typically medium weight/medium: up to approximately 10,000 characters. Device statistics include tmstat data such as interface, CPU, virtual statistics, profile statistics, DoS L3 statistics, and so on. The source is *bigip.tmstats.<tablename>*.

11. ***Push iStats?*** **Advanced**
    Choose whether you want to push iStat statistics. These are typically lightweight/small. iStats are custom statistics created within iRules and other scripting (SSL Intercept uses these statistics as well). The source is *bigip.istats*. For more information on iStats, see *https://devcentral.f5.com/articles/introduction-to-istats-part-1-overview*.

12. ***Push SSLi Stats?*** **Advanced**
    Choose whether you want to push SSL Orchestrator (including SSL Intercept) statistics. These are typically medium weight/small. The source is *bigip.avrstats.AVR_STAT_SSLI_FQDN*.

13. ***Push AVR-sampled Data?*** **Advanced**
    Choose whether you want to push AVR-sampled data. This is typically heavy weight/large: typically more than 10,000 characters. AVR data includes items such as client/server latency, hit rates, transfer sizes, ASM violations, DNS request, SWG activity by URL, ClientIP, Method, FQDNs and so on. The configuration of what dimensions/tables are set is performed within the Analytics profile. The source is *bigip.avrstats.<tablename>.*

14. ***Push SessionDB Statistics (APM)?*** **Advanced**
    This question does not appear if using a BIG-IP version prior to 13.0

    Choose whether you want to push SessionDB statistics for BIG-IP APM. These are typically medium weight/medium. SessionDB statistics are a collection of APM session information and users. Only a set list of core session.x variables are captured as statistics. The source is *bigip.sessiondb*.

15. ***Include Custom SessionDB Variables in APM Statistics?*** Advanced
This question does not appear if using a BIG-IP version prior to 13.0

Choose whether you want to push SessionDB variables for BIG-IP APM. These are typically light weight/small. SessionDB variables allow for the configuration of an additional five user-specified session.x variables. The source is *bigip.sessiondb*.

- **No**
  Select No if you do not want to include custom SessionDB variables in the APM statistics.
- **Yes**
  Select Yes if you want to include custom SessionDB variables (like session.ldap.last.attr.myvar) in the APM statistics. Five fields for custom variables appear.

  Important *If you select Yes, you must put data into these variables in the BIG-IP APM policy or iRule(s) you create. See the APM documentation for specific details.*

  a. *Custom SessionDB 1-5* Advanced
  This field does not appear if you selected BIG-IQ as the data format

  Type your custom variables in the fields. You can specify up to five custom variables.

16. ***Push Configuration Map?*** Advanced
This field does not appear if you selected BIG-IQ as the data format

Choose whether you want to push the Configuration Map. This is a collection of the device configuration and mapping of the configuration into Tenants, Apps, AppComponents, and AppDependencies. This is similar to sending a schema map with each data transmission. By default the configuration map is pushed, however, you can choose to disable it in this question. This is typically medium weight/medium. The source is *bigip.objectmodel.<Object structure>*.

Important *This is required for the majority of the dashboards to function, so only choose No if your data consumer will persist the map from one update to the next or is separately configured with suitable mappings.*

- **No**
  Select No if you do not want the iApp to push the configuration map.
- **Yes**
  Select Yes if you want the iApp to push the configuration map (recommended).

17. ***Push Certificate Expiration Information?*** Advanced
This field does not appear if you selected BIG-IQ as the data format

Choose whether you want to push certificate expiration information. This is typically medium weight/small, This extracts certificate expiration dates for SSL certificates.

- **No**
  Select No if you do not want the iApp to push the certificate expiration information.
- **Yes**
  Select Yes if you want the iApp to push certificate expiry information.

18. ***Manually Manage Stat Collection*** Advanced
Choose whether you want to manually manage the collection of statistics. If you select Yes, a list of every known type of statistic appears in the next field. You can manually add or remove statistics from the collection.

- **No**
  Select this option if you do not want to manually manage statistics collection.
- **Yes**
  Select Yes if you want to manually manage statistics collection. You must add or remove statistics in the next field.

  a. *Custom Stat collection*
  All available statistics appear by default in the Selected list. You can select individual statistics to move to the Options list, which removes them from collection. You can move statistics from Options to Selected to include statistics you had previously removed.

Module Log Stream Capture

Module log streams allows for the collection of log events from the various BIG-IP modules such as ASM, APM, AFM and so on. In addition, you can use this section to enable Risk Engine log data collection as well as enable a conduit for sending custom iRule generated High Speed Logging (HSL) log messages to Splunk. Additional configuration is required to connect the created objects to the modules producing data.

Note: *This entire section only appears if you chose **Yes** from the **Module High Speed Logging Streams** list in the Information Services section. It does not appear at all if you selected BIG-IQ as the Data Format.*

The following objects are created as a part of this section:

- A new Log Publisher
  The Log Publisher named *logging-publisher-f5_analytics*, is available for assignment to APM, AFM and other modules that send to log publishers. Note: To get module logs you must configure the BIG-IP module to use this Log Publisher. See the appropriate module documentation for details.

- A new pool
  A new pool is created that can be used for High Speed Log (HSL) iRule commands to send Key Value Pair data to: *<iApp Name>-hec-forwarder-tcp-log-stage0*.

- A new virtual server
  This virtual server can be used to send ASM events to: 255.255.255.254:1001 Note: if using a custom Unique ID this port number will be the Unique ID + 1 Note: To get ASM logs you must configure ASM to send logs via TCP Key Value Pairs to this IP and Port.

1. ***Use the same configuration as in the Analytics System Configuration section?***
   Each data source section can customize the destination or the event data or can use the same destination as the Analytics System configuration from the main section. Specify Yes to use the same configuration from above.

   - **Yes**
     Select this option to use the same analytics system destination and port you used earlier in this template.

   - **No**
     Select this option to use a different destination. You must specify the following information.

     a. *IP address or Hostname*
     Specify the IP Address or Hostname of the data consumer's index server where you want to send data. This can refer to an LTM virtual server or (with TMOS version 11.6.0 or later) a DNS round-robin FQDN to allow data to be distributed to multiple index servers.

     If selecting an HTTP Event Collector (in Splunk or compatible applications) the Authentication (API) Key must be the same for all index servers. Note BIG-IP version 11.6.0 and higher supports FQDNs. Previous versions require an IP address.

     If using Splunk Cloud this will be input-<yourinstancename>.cloud.splunk.com for BIG-IP version 11.6.0 and higher. For previous versions use the IP address where this resolves.

     b. *Port*
     Specify the port to use for the indexing server. If using the HTTP Event Collector this is generally 8088.

     c. *Protocol*
     Choose the protocol of the indexing server. If sending data to Splunk via the HTTP Event Collector, select HTTP Event Collector - HTTPS. Other methods of sending data are supported by the iApp. Currently event correlation within the Splunk App relies on data being received by HEC.

     - **TCP**
       Select this option to use TCP. If you select TCP, all data is sent directly from the BIG-IP. Note that the BIG-IP may create more than one simultaneous connection to the configured destination via TCP.

     - **UDP**
       Select this option to use UDP. If you select UDP all data is sent directly from the TMM to the configured destination via UDP. Note: UDP transmission can cause some longer event messages to be truncated. The maximum UDP message length is 64 kilobytes.

- **HTTP Event Collector - HTTP** or
  HTTP Event Collector - HTTPS
  If you are using Splunk or a similar application, we recommend using one of these options. In this case, event messages are converted into JSON messages with additional metadata attached allowing for better correlation of event data to the environment.

  a. *Hostname when the destination is an IP address* Advanced
  You can optionally provide a HTTP host header to be used inside of the communication.

  This is useful for BIG-IP versions before 11.6 where FQDNs are not supported. Providing a host header ensures that proxies and other security devices intercepting this traffic including the HTTP Event Collector understands the communication.

  b. *Use the same API key as the main configuration?*
  This field does not appear if you selected Sumo Logic or BIG-IQ as the data format

  Choose whether you want to use the same API key you specified in Step 5 of *Analytics System Configuration on page 13*, or if you want to specify a specific API key to be used for Module Log Streams. By specifying a new API key, you can set a different index for this data type within the data consumer.

  - **Yes**
    Select this option if you want to use the same API you specified earlier
  - **No**
    Select this option if you want to specify a different API key than the one in *Analytics System Configuration*.

    a. *API Key (Required)*
    Specify the API key to use when authenticating. For Splunk, this is the Token Key generated within the HEC settings.

2. ***Source Address Translation*** Advanced
   Choose how you want the system to handle Source Network Address Translation (SNAT). You can select SNAT auto map or a SNAT pool to perform source network address translation. When data is transmitted from the iApp, the source of the traffic is set to an IP address the BIG-IP owns. With SNAT auto map, the BIG-IP chooses one of its Self IP addresses using the scheme explained in F5 note K7336 (*https://support.f5.com/csp/article/K7336*). You can change that behavior using a SNAT pool you create outside the iApp. A SNAT pool allows you to specify otherwise unused IP addresses as translation addresses.

   Note that if you configure a SNAT Pool for this configuration, it must use non-floating addresses.

   - **Auto-Map**
     Select this option if you want the system to use SNAT auto map. In this case, the system automatically selects one of the systems self IP addresses (typically a floating self IP address of the egress VLAN), and maps it to the original IP address.
   - *Select the SNAT Pool you created*
     If you manually created a SNAT pool for this configuration, select it from the list.

3. ***Send Behavioral Anti-DDoS Stats***
   If using Behavioral Anti-DDoS advanced statistics on the machine, learning algorithms can be collected and viewed within the solution per application.

4. ***Send Risk Logs to Analytics System?***
   Choose whether you want to send risk logs to the analytics system. This option is used when the data format and destination is the F5 Risk Engine.

   - **No**
     Select this option if you do not want to send risk logs to the analytics system.
   - **Yes**
     Select this option if you want to send risk logs to the analytics system. You must specify the virtual servers to which you want to attach the risk logging profile.

     a. *Virtual Servers to Attach Risk Logging profile*
     You must select the virtual servers on the system that you want to attach the risk logging profile. From the Options list, select the virtual server(s) you want to include and then click Add (<<) to move the virtual servers to the Selected box.

5. ***Send Login Events to Analytics System?***
   Choose whether you want to send login events recognized by BIG-IP ASM to the data consumer.

   - **No**
     Select this option if you do not want to send login events to the analytics system.
   - **Yes**
     Select this option if you want to send login events to the analytics system.

     a. *Configured Logins*
        You must complete the following information about the logins:

        - **Virtual Server**
          From the list, select either ALL-FROM-ABOVE, or select a specific virtual server from the list.
        - **Login Path**
          Type the login path. When a web application user accesses this path, BIG-IP ASM watches for a login method to be processed.
        - **Login Method**
          Choose the appropriate HTML login method from the list.

          a. *HTML-Form*
             Select this option if you are using an HTML-Form login method.

          b. *HTML - Basic Auth*
             Select this option if you are using Basic Auth.

          c. *HTML - Digest*
             Select this option if you are using HTML-Digest.
        - **Username Field**
          Type what you want to appear in the username field.
        - **Validation Method**
          Select the validation method, which is the way ASM recognizes that a login method has been successful for a particular web application user.

          - **Cookie Name**
            Select this option if you want to use the cookie name as the validation method.
          - **Cookie Name - Value**
            Select this option if you want to use the cookie name and value as the validation method.
          - **Response Code**
            Select this option if you want to use the response code as the validation method.
          - **Redirect Path**
            Select this option if you want to use the redirect path as the validation method.

          d. *Validation Check String*
             Type the validation check string you want to use. For a detailed explanation, see the BIG-IP ASM documentation for your version of TMOS.

          e. *Validation Check Cookie Value*
             Type the validation check cookie value.

If you selected **BIG-IQ** as the Data Format, this completes the configuration. Continue with *Finished on page 28*.

Local Logging Capture and System

In this section, you configure the local system logging (syslog) information.

It is important to keep in mind that capturing local logs to send them to your Analytics data consumer prevents them from going to any other consumer, such as a syslog-based event recorder. Furthermore, the Analytics data consumer will only get local log data when the BIG-IP device is on-line for traffic processing, even though the BIG-IP management plane can potentially generate log messages while traffic processing (TMM) is offline.

Note: *This entire section only appears if you chose **Yes** from the **Local System Logging (syslog)** list in the Information Services section. It does not appear at all if you selected BIG-IQ as the Data Format.*

1. ***Use the same configuration as in the Analytics System Configuration section?***
   Each data source section can customize the destination or the event data or can use the same destination as the Analytics System configuration from the main section. Select Yes to use the same configuration from above.

   - **Yes**
     Select this option to use the same analytics system destination and port you used earlier in this template.
   - **No**
     Select this option to use a different destination. You must specify the following information.

     a. *IP address or Hostname*
     Specify the IP Address or Hostname of the data consumer's index server where you want to send data. This can refer to an LTM virtual server or (with TMOS version 11.6.0 or higher) a DNS round-robin FQDN to allow data to be distributed to multiple index servers.

     If selecting an HTTP Event Collector Protocol the Authentication (API) Key must be the same for all index servers. Note BIG-IP version 11.6.0 and higher supports FQDNs. Previous versions require an IP address to be specified.

     If using Splunk Cloud this will be input-<yourinstancename>.cloud.splunk.com for BIG-IP version 11.6.0 and higher. For previous versions use the IP address where this resolves.

     b. *Port*
     Specify the port to use for the indexing server. If using the HTTP Event Collector this is generally 8088.

     c. *Protocol*
     Choose the protocol of the indexing server. If sending data to Splunk via the HTTP Event Collector, select HTTP Event Collector - HTTPS. Other methods of sending data are supported by the iApp. Currently event correlation within the Splunk App relies on data being received by HEC.

     - **UDP**
       Select this option to use UDP. If you select UDP all data is sent directly from the TMM to the configured destination via UDP. Note: UDP transmission can cause some longer event messages to be truncated. The maximum UDP message length is 64 kilobytes.
     - **HTTP Event Collector - HTTP** or
       HTTP Event Collector - HTTPS
       If you are using Splunk or a similar application, we recommend using one of these options. When selected, syslog/SNMPTraps messages are converted into JSON messages with additional metadata attached allowing for better correlation of event data to the environment.

       Note: syslog messages are routed to TMM for this conversion requiring TMM to be online to process messages.

       a. *Hostname when the destination is an IP address* **Advanced**
       You can optionally provide a HTTP host header to be used inside of the communication.

       This is useful for BIG-IP versions before 11.6 where FQDNs are not supported. Providing a host header ensures that proxies and other security devices intercepting this traffic including the HTTP Event Collector understands the communication.

b. *Use the same API key as the main configuration?*
This field does not appear if you selected Sumo Logic or BIG-IQ as the data format

Choose whether you want to use the same API key you specified in Step 5 of *Analytics System Configuration on page 13*, or if you want to specify a specific API key to be used for Local Logs (syslog). By specifying a new API key, you can set a different index for this data type within the data consumer.

- **Yes**
  Select this option if you want to use the same API you specified earlier.
- **No**
  Select this option if you want to specify a different API key than the one in the Analytics System Configuration section.

  a. *API Key (Required)*
  Specify the API key to use when authenticating. For Splunk, this is the Token Key generated within the HEC settings.

2. ***Source Address Translation*** **Advanced**
Choose how you want the system to handle Source Network Address Translation (SNAT). You can select SNAT auto map or a SNAT pool to perform source network address translation. When data is transmitted from the iApp, the source of the traffic is set to an IP address the BIG-IP owns. With SNAT auto map, the BIG-IP chooses one of its Self IP addresses using the scheme explained in F5 note K7336 (*https://support.f5.com/csp/article/K7336*). You can change that behavior using a SNAT pool you create outside the iApp. A SNAT pool allows you to specify otherwise unused IP addresses as translation addresses.

Note that if you configure a SNAT Pool for this configuration, it must use non-floating addresses.

- **Auto-Map**
  Select this option if you want the system to use SNAT auto map. In this case, the system automatically selects one of the systems self IP addresses (typically a floating self IP address of the egress VLAN), and maps it to the original IP address.
- *Select the SNAT Pool you created*
  If you manually created a SNAT pool for this configuration, select it from the list.

SNMP Alert Capture

In this section, you configure the SNMP Alert capture information.

It is important to note that capturing SNMP Alerts to send them to your Analytics data consumer prevents them from going to any other consumer, such as an SNMP-based network-management server. Furthermore, the Analytics data consumer will only get SNMP Alert data when the BIG-IP device is on-line for traffic processing, even though the BIG-IP management plane can potentially generate alerts while traffic processing (TMM) is offline.

Note: *This entire section only appears if you chose **Yes** from the **System SNMP Alerts** list in the Information Services section. It does not appear at all if you selected BIG-IQ as the Data Format.*

1. *Use the same configuration as in the Analytics System Configuration section?*
   Each data source section can customize the destination or the event data or can use the same destination as the Analytics System configuration from the main section. Select Yes to use the same configuration from above.

   - **Yes**
     Select this option to use the same analytics system destination and port you used earlier in this template.
   - **No**
     Select this option to use a different destination. You must specify the following information.

     a. *IP address or Hostname*
     Specify the IP Address or Hostname of the data consumer's index server where you want to send data. This can refer to an LTM virtual server or (with TMOS version 11.6.0 or higher) a DNS round-robin FQDN to allow data to be distributed to multiple index servers.

     If selecting an HTTP Event Collector Protocol the Authentication (API) Key must be the same for all index servers. Note BIG-IP version 11.6.0 and higher supports FQDNs. Previous versions require an IP address to be specified.

     If using Splunk Cloud this is input-<yourinstancename>.cloud.splunk.com for BIG-IP version 11.6.0 and higher. For previous versions use the IP address where this resolves.

     b. *Port*
     Specify the port to use for the indexing server. If using the HTTP Event Collector this is generally 8088.

     c. *Protocol*
     Choose the protocol of the indexing server. If sending data to Splunk via the HTTP Event Collector, select HTTP Event Collector - HTTPS. Other methods of sending data are supported by the iApp. Currently event correlation within the Splunk App relies on data being received by HEC.

     - **UDP**
       Select this option to use UDP. If you select UDP all data is sent directly from the TMM to the configured destination via UDP. Note: UDP transmission can cause some longer event messages to be truncated. The maximum UDP message length is 64 kilobytes.

       a. *SNMP Community Name*
       Type your SNMP community name.

     - **HTTP Event Collector - HTTP** or
       HTTP Event Collector - HTTPS
       If you are using Splunk or a similar application, we recommend using one of these options. When selected, syslog/SNMPTraps messages are converted into JSON messages with additional metadata attached allowing for better correlation of event data to the holistic environment. Note: syslog messages are routed to TMM for this conversion requiring TMM to be online to process messages.

       a. *Hostname when the destination is an IP address* Advanced
       You can optionally provide a HTTP host header to be used inside of the communication.

       This is useful for BIG-IP versions before 11.6 where FQDNs are not supported. Providing a host header ensures that proxies and other security devices intercepting this traffic including the HTTP Event Collector understands the communication.

b. *Use the same API key as the main configuration?*
This field does not appear if you selected Sumo Logic or BIG-IQ as the data format

Choose whether you want to use the same API key you specified in Step 5 of *Analytics System Configuration on page 13*, or if you want to specify a specific API key to be used for SNMP Alerts. By specifying a new API key, you can set a different index for this data type within the data consumer.

- **Yes**
  Select this option if you want to use the same API you specified earlier.
- **No**
  Select this option if you want to specify a different API key than the one in *Analytics System Configuration on page 13*.

  a. *API Key (Required)*
  Specify the API key to use when authenticating. For Splunk, this is the Token Key generated within the HEC settings.

2. ***Source Address Translation*** **Advanced**
Choose how you want the system to handle Source Network Address Translation (SNAT). You can select SNAT auto map or a SNAT pool to perform source network address translation. When data is transmitted from the iApp, the source of the traffic is set to an IP address the BIG-IP owns. With SNAT auto map, the BIG-IP chooses one of its Self IP addresses using the scheme explained in F5 note K7336 (*https://support.f5.com/csp/article/K7336*). You can change that behavior using a SNAT pool you create outside the iApp. A SNAT pool allows you to specify otherwise unused IP addresses as translation addresses.

Note that if you configure a SNAT Pool for this configuration, it must use non-floating addresses.

- **Auto-Map**
  Select this option if you want the system to use SNAT auto map. In this case, the system automatically selects one of the systems self IP addresses (typically a floating self IP address of the egress VLAN), and maps it to the original IP address.
- *Select the SNAT Pool you created*
  If you manually created a SNAT pool for this configuration, select it from the list.

iHealth Snapshot Information Capture

In this section, you configure the iHealth snapshot capture details. By sending QKView data to iHealth, diagnostics data based on F5's Weekly Updated Heuristics can be provided. This provides insights into Security Vulnerabilities, know bugs, and other best practice guidance within the visibility solution. For more information on iHealth, see *https://f5.com/support/tools/ihealth*.

Note: *This entire section only appears if you chose **Yes** from the **iHealth Snapshot Information** list in the Information Services section.*

If you choose to use the iApp to send iHealth snapshot information, your iHealth username and password may be stored in cleartext in the BIG-IP configuration and/or the scriptd.out log

This section does not appear if you selected BIG-IQ as the Data Format.

1. ***iHealth Username***
   Type your iHealth username for the iHealth account data should be sent with. If you do not have an iHealth account, you can create one by visiting the f5 iHealth web site at *https://ihealth.f5.com/qkview-analyzer/*.

2. ***iHealth Password***
   Type the associated password.

3. ***Use an HTTP Proxy?***
   Choose whether you want to use an HTTP proxy for iHeath information. Use of an HTTP Proxy is supported for communicating to F5's iHealth servers. This proxy server can be different from the proxy server used to communicate to the data consumer. Use of an FQDN or IP address is permitted.

   - **No**
     Select this option if you are not using an HTTP proxy. Continue with #4.
   - **From Main Configuration**
     Select this option if you want to use the same proxy server you specified in the Analytics System Configuration section.
   - **Yes**
     Select this option if you want to include a different HTTP proxy for iHealth information. You must complete the following.

     a. *Proxy IP Address*
        Type the IP address or FQDN of the HTTP proxy server.

     b. *Proxy Port*
        Type the port used by the HTTP proxy server.

     c. *Proxy Username*
        Type a user name with access to the proxy. This field is optional.

     d. *Proxy Password*
        Type the password for the username you specified. This field is optional.

4. ***Schedule***
   Choose the frequency with which you want to send QKViews. After you choose the Start and End times, the iApp picks a moment between these two times to send data. This spreads out the load on F5's iHealth servers.

   - **Daily**
     Select this option if you want to send QKViews daily.

     a. *Start Time*
        Type the start time for the window for sending iHealth data.

     b. *End Time*
        Type the end time for the window. This should be at least two hours later than the start time.

- **Weekly**
  Select this option if you want to send QKViews weekly. We recommend this option.

  a. *Day of the week*
     Select the day of the week you want to send iHealth data.

  b. *Start Time*
     Type the start time for the window for sending iHealth data.

  c. *End Time*
     Type the end time for the window. This should be at least two hours from the start time.

- **Monthly**
  Select this option if you want to send QKViews monthly.

  a. *Day of the month*
     Type the day of the month you want to send iHealth data (1-31). If specifying the 31st, the system only sends data in months with a 31 days. You can use this option to send every other month.

  b. *Start Time*
     Type the start time for the window for sending iHealth data.

  c. *End Time*
     Type the end time for the window. This should be at least two hours from the start time.

Application Mapping

At the core of this solution is Application Mapping. Application Mapping allows for consolidation of all of the BIG-IP objects such as partitions, wide IPs, virtual servers, pool members, and so on into common constructs such as Tenants, App, App Components, and the Dependencies between them. The mapping configuration section allows for the configuration of these mappings by using regular expression (regex) pattern matching.

This mapping is exportable so it can be reused across a large number of BIG-IP devices without having to meticulously redefine the same mapping on each BIG-IP or BIG-IP cluster.

In most cases one or more naming conventions are used within an enterprise allowing this mapping to be defined and then not have to be touched as additional applications, virtual servers, or wide IPs are added, changed, or removed.

This section does not appear if you selected BIG-IQ as the Data Format.

1. ***Search iRules?***
   Choose whether you want the system to search iRules. iRules are commonly used for logically modifying the Data Plane. Within an iRule, there are generally statements that can affect how the application is defined, such as which pools it uses, if it sends 301/302 response codes, and so on. The iApp searches these iRules for those statements to allow a full Application Definition to be discovered.

   To get additional value out of this data collection, iRules can be modified to add additional metadata for collection. This additional metadata is not required. as the iApp will discover the Application Definition regardless but it does add context. This can be done for Pool, HTTP::respond, and DNS::answer statements within a rule by adding

   ```
   ;#context: "<Context/Definition Text>" istat: "<iStat Name>"
   ```

   The following code shows some examples:

   ```
   when HTTP_REQUEST {
     if {[HTTP::uri] ends_with ".js"} {
       ISTATS::incr "ltm.virtual [virtual name] counter jsrequests" 1
       pool varnishpool ;#context: "JS Request" istat: "jsrequests"
     }
   }

   when HTTP_REQUEST {
     if {[HTTP::uri] ends_with ".jpg"} {
       pool imagepool ;#context: "Image Request"
       }
   }

   when HTTP_REQUEST {
     if {[HTTP::uri] eq "/promo/tv123"} {
       ISTATS::incr "ltm.virtual [virtual name] counter promo-tv123" 1
       HTTP::Respond 301 location "https://[HTTP::host]/productoverview?id=1721" ;#context: "Promo Redirect" istat: "promo-tv123"
     }
   }
   ```

   - **No**
     Select this option if you do not want the system to search iRules.
   - **Yes**
     Select this option if you want the system to search iRules.

2. ***Configuration Mode***
   The iApp supports several configuration modes for the Application Mapping configuration. If this is the first BIG-IP on which this version of mapping is being defined, select **Define**.

   - **Define**
     Select this option to define the configuration mode.

     a. *Mapping Table*
        Specify the following information for the mapping table. The Mapping Table allows for the definition of Object Mappings. It is important to ensure that the iApp has been configured such that at least a Tenant and App Name can be mapped for

all virtual servers and/or wide IPs. If this is a Single Tenant BIG-IP or BIG-IP cluster than a Default Tenant may have been configured earlier, it this is the case then only an App Name needs to be configured for mapping. In the case where a Default Tenant was not configured, ensure the Mapping Table accounts for a Tenant Mapping.

- **Order**
  The Order field is an optional field that forces the mapping engine to order which mapping occur first. The lower the number the earlier the mapping line is used; for example, 10 fires before 20. If no order is provided, then the mapping line executes after all the match lines with order specified have executed. This is done in the order listed in the Mapping Table.

  Tip: *A good practice is to increase these each by 10 to allow for an easy way to add additional mappings later. You can also split the orders by Type, for example all App Name mappings start with 100, Tenants 200, and so on.*

- **Type**
  The Type field specifies what type of mapping is to be done on the specific mapping line.

  a. *Tenant*
  The Tenant option is used to group (virtual servers and wide IPs) to Tenants within the dashboards. If RBAC is used, the Tenant Name is used to set the index name to allow for access controls to be set within Splunk. An object can only be mapped to one Tenant.

  b. *App Name*
  Most commonly used Type. The Application Name is used to group/align (virtual servers and wide IPs) to a displayed Application Name this allows multiple virtual servers and wide IPs across many BIG-IPs to be shown together as one Application Name. An object can only be mapped to one Tenant.

  c. *App Component Name*
  Allows a second tier of group to be done within an Application Name. This is a common way to sub divide an Application into multiple tiers of components such as Web, Middleware, Database or Web, Login Service, Image Rendering, Database, etc. An object can only be mapped to one App Component Name.

  d. *Application Dependencies*
  This choice is not currently useful but remains in the iApp template as a placeholder. Do not select this option.

- **From**
  The From field selects what value/string is used for the regex mapping

- **Regex**
  The regex field is used to specify how to match and capture data from the selected source. Use *http://regex101.com* as a tool to test regex matching. These can always be updated if mistakes are made. If a mistake is made display data in Splunk from the time of the corrected mapping onward to not show old data.

  The following table contains a few examples:

Source String	regex	Matched?	Captured String
/Common	`\/(.*)`	Yes	Common
vs_www.site.com-80	`vs_(.*)-[0-9]+`	Yes	www.site.com
/Common/exchange/exchange.app	`.*\/([^/]+)\.app`	Yes	exchange
www.site.com_VS	`(.*)[_-][vV][sS]`	Yes	www.site.com
anything.anything.com-80	`(.*)`	Yes	anything.anything.com-80
anything.anything.com-80	anything	Yes	

- **Action**
  The Action option specifies how the mapping line operates when a regex is matched.

  a. *Map*
  Choose this option and the system concatenates all of the captured regex data and sets the value of the Type to the captured data. If a match is found all further Mapping Lines of the same Type stops.

For example:

Type: *App Name* **Source**: *Virtual Name* **regex:** *vs_(.*)-.** Results:

- vs_site.com-80 -> App Name = site.com
- vs_site.com-443 -> App Name = site.com
- vs_app1-80 -> App Name = app1
- site.com-80 -> No Match

Type: *Tenant Name* **Source**: *Partition* **regex:** *\/(.*)* Results:

- /Common -> Tenant Name = Common
- /tenanta -> Tenant Name = tenanta

b. *Append*
Choose this option and the system concatenates all of the captured regex data and appends the captured data to the existing value. If a match is found other Mapping Lines of the same Type continue.

For example:

Type: *Tenant Name* **Source**: *Partition* **regex:** *\/(.*)* Results:

- /tenanta -> Tenant Name = tenanta

Type: *Tenant Name* **Source**: *Virtual Name* **regex:** *(.*)-.*-.** **Append:** -:

- pci-site.com-80 -> Tenant Name = tenanta-pci

c. *Ignore*
Choose this option and if any part of the regex is matched, the Type is set to **Ignore**. This is commonly used with the **App Name Type** to hide virtual servers that should not be displayed. If a match is found all further Mapping Lines of the same Type stops.

For example:

Type: *App Name* **Source**: *Virtual Name* **regex:** *testvirtual* Results:

- Vs_testvirtual-tom-80 -> App Name = Ignore

d. *Direct Mapping*
Choose this option and if any part of the regex is matched, the Type is set to the value of the Direct Mapping field. If a match is found all further Mapping Lines of the same Type stops.

For example:

Type: *App Name* **Source**: *Virtual Name* **regex:** *wwtypo.site.com* Results:

- vs_wwtypo.site.com -> App Name = www.site.com
- vs_ftp.site.com -> No Match

e. *Direct Mapping Append*
Chose this option, and if any part of the regex is matched, the Direct Mapping field is appended to the existing value. If a match is found other Mapping Lines of the same Type continue.

- **AppendPrefix**
The Append Prefix field appends the specified prefix before the captured regex data.

Note: *This field is only used if the Action is Map, Append or Direct Mapping Append.*

b. *Mapping Element for Export*
The Mapping Element for Export provides a way to export the mappings defined in the Mapping Table allowing it to be used for deployment on another BIG-IP or BIG-IP cluster.

Note: When you select **Use Existing with Additional** from the **Configuration Mode** list, this export is the combination of the imported string and the additional mappings. To ensure the latest version of the export string only, after you click

Finished on the iApp template, wait 60 seconds, click Reconfigure on the menu bar, and then copy the string that appears in this field.

- **Use Existing**
  Select this option from the **Configuration Mode** list and the system allows for an existing Mapping Table to be imported from another BIG-IP using the Mapping Import String. This is an easy way to maintain the same mapping for many BIG-IP devices without the risk of typos and misconfiguration. You must have run the iApp on the other BIG-IP device and copied the string.

  a. *Mapping Import String*
     Type or paste the mapping import string. The Mapping Import String is the string provided from another BIG-IP using the Mapping Element for Export (described in step b above). It is the Base64 encoded CSV representation of a defined Mapping Table.

- **Use Existing with Additional**
  Select this option from the **Configuration Mode** list and the system allows for an existing Mapping Table to be imported from another BIG-IP using the Mapping Import String. It then allows additional mappings to be defined in the Mapping Table to correct for virtual servers or other objects that don't align with the imported mapping.

  The order field can be used to place a new mapping in between mappings from the imported string.

  a. *Mapping Import String*
     Type or paste the mapping import string. The Mapping Import String is the string provided from another BIG-IP using the Mapping Element for Export (described in step b above). It is the Base64 encoded CSV representation of a defined Mapping Table.

     For guidance on the **Mapping Table** and **Mapping Element** for Export fields, return to *a. Mapping Table on page 25*.

Finished

Review the answers to your questions. When you are satisfied, click the **Finished** button. The BIG-IP system creates the relevant objects.

Troubleshooting

Use this section for common troubleshooting steps and workarounds for known issues.

- To view status of transfers via SSH, use the following command: **`tail /var/log/ltm -f | grep "Stats Response for"`**
- If you receive the following error message when you try to remove or reconfigure a log publisher:

  `01070635:3: The log destination (/Common/Splunk-hec-forwarder-tcp-log-splunkformat) is referenced by one or more publishers.`

  You can remove the iApp generated log destination by performing the following:

  a. From the BIG-IP Configuration utility, click **System > Logs > Configuration > Log Publishers**.

  b. Click the **logging-publisher-f5_analytics** publisher.

  c. In the Log Destinations section, click the Log Publisher created by the iApp, and move it from the Selected list to the Available list. For example, this publisher is named **Splunk-hec-forwarder-tc-log-splunkformat**.

- Reconfiguring the iApp application service to disable Module High Speed Logging Streams in the Information Sources section may cause an error like the following: **`01070265:3: The rule (/Common/ir-x-risk _ engine-log) cannot be deleted because it is in use by a virtual server (/Common/MyVS1)`**.

  If you experience this error, to disable Module High Speed Logging Streams after you have initially deployed the iApp, you must delete and then re-create the iApp application service.

Appendix: Using Splunk with the data

Use this appendix for information on configuring Splunk to use the data produced by this iApp.

Important *We provide this information for reference only; for complete information on how to deploy or configure the components of Splunk, consult the appropriate Splunk documentation. F5 cannot provide support for Splunk.*

Index Considerations

Indexes allow for data within Splunk to be isolated and managed separately. This enables performance gains for the F5 Analytics App. Several considerations should be taken when defining the indexes for this solution.

If you are not using Role Based Access Controls (RBAC) you can simply define a single index, **f5-default**, to be used for all data received by the HTTP Event Collector from the BIG-IP devices.

If you are using Role Based Access Controls (RBAC) indexes are used as containers that users have selective access to. There are many options with RBAC design based on your architecture and desired permission isolation. At a high level there are two main paths.

- If your environment is designed where each tenant is isolated to their own sets of BIG-IP devices then an index per tenant is suggested i.e. **f5-<tenant name>**
- If your environment is designed where multiple tenants share BIG-IP infrastructure than a more complex index model is needed.
  - Commonly this is made up of the same f5-<tenant name> indexes where data is collected by tenant.

    Dots (.) and spaces ( ) within the tenant names are replaced with underscores (_) when being set in the index name.
  - Because there are system level statistics for the shared infrastructure you must decide what tenant/index to use for these statistics. These system statistics can be sent to a separate index, for example f5-system_stats.
  - There are also shared statistics, which are statistics for shared objects that cannot be mapped to a specific tenant on the BIG-IP devices but need to be accessible to each tenant. These statistics should be send to a separate index and assigned to be accessible by all tenants, for example f5-shared_stats.
  - Events made up of Syslog, SNMPTraps, and Module Logs are not classified by tenant. These events need to get mapped to an index, like f5-events.
  - There is an unknown mapping which is an index to be used if the tenant mapping elicited no match: f5-unknown.

Splunk allows for control of access to data per user / user groups using Roles. More details on working with Users and roles can be found at: *http://docs.splunk.com/Documentation/Splunk/latest/admin/Aboutusersandroles*.

Further information on index and RBAC can be found in *Role Based Access Control (RBAC) on page 11*.

Splunk Index Configuration

Use this section for guidance on configuring a Splunk Index. For complete information, see the Splunk documentation.

1. On the Splunk main page, from the Menu bar click **Settings**, and then click **Indexes**.
2. Click **New Index**.
   a. In the **Index Name** field, type a name for the index.
   If you are *not* using Role Based Access Controls, we suggest using the name **f5-default**.
   If you are using Role Based Access Controls, we suggest using a name such as **f5-<your tenant name>** (an index per tenant), f5-shared_stats, f5-system_stats, and so on.
   b. Configure the rest of the settings as applicable for your configuration.
   c. Click **Save**.

Details on index configuration options surrounding performance and retention can be found at:
http://docs.splunk.com/Documentation/Splunk/latest/Indexer/HowSplunkstoresindexes.

Splunk HTTP Event Collector configuration

Use this section for guidance on configuring the Splunk HTTP Event Collector (HEC). For complete information, see the Splunk documentation.

HEC Overview

HEC is a way to send data to Splunk Enterprise and Splunk Cloud. Notably, HEC enables you to send data over HTTP (or HTTPS) directly to Splunk Enterprise or Splunk Cloud from your application.

HEC Configuration

1. From the Splunk indexing server receiving the data, click **Settings** and then click **Data Inputs** (in the Data area).
2. In the Local Inputs area, click HTTP Event Collector.
3. If this is the first time you are configuring HEC, perform the following steps:
    a. Click **Global Settings**.
    b. In the **All Tokens** area, click **Enabled**.
    c. Click the **Save** button.
4. Click the **New Token** button.
5. In the **New Token** area, in the **Name** field, type a unique name such as **F5-BIG-IP**, and then click **Next**.
6. In the Input Settings, **Index** area, from the **Default Index** list, select the Index you created in *Splunk Index Configuration on page 30*.
7. Review your settings and then click **Submit**.
8. Once the token has been created, copy the **Token Value** for use in the Analytics iApp. You enter this token in *5. API Key on page 13*.

F5 Splunk App Deployment and Settings

Use this section for guidance on configuring Splunk to use the Analytics iApp.

1. From the Splunk UI, click the **Apps** settings icon on the Search Head.
2. Click either **Browse more apps** or **Install app from file** as your method of installation.
    - If you chose **Browse more apps**, search for **F5 Networks** and then click the Install button on **F5 Networks – Analytics (New)**.
    - If you chose **Install app from file**, download the latest version of the app from *https://apps.splunk.com/apps/id/f5*.

Configuring the F5 Networks - Analytics Splunk App

To improve the performance of the application, it is best to modify the **f5_index** macro to point specifically to the indexes used to store the F5 data. This is not a required step, as the macro is preconfigured to search all indexes.

1. Click **Settings** and then click **Advanced Search** (in the Knowledge area).
2. Click Search macros.
3. From the **App context** list, select F5 Networks (f5).
4. Click the **f5_index** link from the list.
5. In the Definition area, modify the setting to the index you want and then click **Save**. This can be set to the exact index you configured in *Splunk Index Configuration on page 30* (**index=f5-default**, or if using RBAC **index=f5-***. If using a custom index, **name index=<your index name>**.
6. Click the **Save** button.

For complete information, see the Splunk documentation.

Document Revision History

Version	Description	Date
1.0	New Deployment Guide for the f5.analytics.v3.7.0 version of the iApp template.	03-29-2017
1.1	Added information about iHealth passwords being stored in cleartext to the prerequisites and the iHealth sections.	04-04-2017
1.2	Updated this guide for Release Candidate v3.7.1rc1. This iApp (and guide) contains the following changes: - Added support for sending data to Sumo Logic. Options for Sumo Logic are the same as for Splunk except "API Key" options do not appear, Sumo Logic does not use an API Key for authentication. - Corrected an issue where the iApp would send incorrect log messages about failures to connect to the remote analytics server. - Corrected an issue where the iApp would fail to apply Risk-engine logging iRules to virtual servers when 'Module High Speed Logging Streams' and 'Send Login Events to Analytics System?' were set to 'Yes', and 'Send Risk Logs to Analytics System?' was set to 'No.' - Corrected an issue where the iApp would leave old statistics files on the BIG-IP hard disk. - Corrected an issue where if an iHealth password included a dollar-sign '$' character (ASCII 0x24) the iApp would not upload data to iHealth. - Corrected an issue where any of the following characters in a TMOS configuration object name or attribute value could cause the iApp to generate improper JSON which would be rejected by data consumers such as Splunk: * ? [" ' \ SPACE (asterisk, question-mark, left-bracket, quote-mark (ASCII 0x22), apostrophe, backslash, or SPACE (ASCII 0x20)). Note that from version 3.7.1 forward quote-marks and backslashes in configuration object names or attributes will be replaced with underscore '_' characters in JSON data, and all control characters (ASCII 0x00-0x1f plus 0xff) will be replaced by exclamation-points '!'.	07-11-2017
1.3	Updated this guide for f5.analytics.v3.7.1rc2. This version of the template contains a single fix for an issue that was preventing statistical data from being sent because an array was not initialized.	08-24-2017
1.4	Updated this guide for f5.analytics.v3.7.1rc3. This version of the template contains no new features, but includes the following changes: - Corrected an issue where, on a VIPRION with multiple blades, secondary blades would go down after running the iApp. - Corrected an issue that caused parsing of qkviews for iHealth to fail. - Corrected an issue that caused the analytics solution API key to not be included in the authorization header. - Corrected an issue that caused the parsing of GTM configuration to fail on certain BIG-IP systems.	12-14-2017
1.5	Updated this guide for f5.analytics.v3.7.1rc4. This version of the template contains no visible changes to the iApp presentation. This release corrects an issue in the iApp code where the command for getting certain BIG-IP device information was different between BIG-IP v12.1 and earlier versions.	02-01-2018
1.6	Updated this guide for f5.analytics.v3.7.1rc5. This version of the template contains no visible changes to the iApp presentation. This release corrects an issue when using an FQDN pool member on systems older than 12.x. Previously, attempting to use an FQDN pool member was denied, but only logged with no error reported to the user.	03-22-2018
1.7	Updated this guide for f5.analytics.v3.7.1rc6. This version of the template contains no visible changes to the iApp presentation. This release contains the following fixes: - Corrected an issue in the iApp where some BIG-IP versions would give an error stating FQDNs are not supported. - Corrected an issue where not all syslog messages were delivered to the remote endpoint. - Updated the iApp so it checks to see if a management-route exists (management-routes are not a part of the Application Service Object (ASO)). If it does exist, the iApp updates it, if not it creates a new route.	06-07-2018
1.8	Updated this guide for f5.analytics.v3.7.1rc6. This version of the template contains no visible changes to the iApp presentation. This release adds support for BIG-IP versions 13.0-13.1.	06-08-2018
1.9	Updated this guide for the fully supported f5.analytics.v3.7.1 version of the iApp template available on downloads.f5.com. This iApp contains all of the features and fixes included in RC-1 through RC-6.	07-26-2018

Version	Description	Date
2.0	Updated this guide to add a note in the prerequisites that only one instance of the Analytics iApp is supported on any BIG-IP.	08-15-2018
2.1	Updated this guide for f5.analytics.v3.7.2rc3. This version of the template contains no visible changes to the iApp presentation.	09-13-2018
2.2	Updated this guide to correct the version note on page 3.	09-14-2018
2.3	Updated this guide for f5.analytics.v3.7.2rc4. This release allows the Splunk node address in the iApp to support route domains. Previously, the iApp included the route domain in the node name.	10-18-2018
2.4	Updated this guide for f5.analytics.v3.7.2rc5. This release fixes an issue that was introduced in f5.analytics.v3.7.2rc4 that caused the periodic iCall script to report an error message in /var/tmp/scriptd.out about missing variable “splunkdestinationip_withrd”. This error resulted in statistics not making it to Splunk or other consumers.	12-27-2018
2.5	Updated this guide for f5.analytics.v3.7.2rc6. This release fixes an issue that prevented the Analytics iApp from reporting certificate expiration dates to Splunk.	03-21-2019
2.6	Updated this guide for f5.analytics.v3.7.2rc7. This release fixes an issue that prevented using certain characters in the iHealth password.	05-23-2019
2.7	Updated this guide for f5.analytics.v3.7.2rc8. This release fixes an issue that prevented the iApp from pushing certificate data to Splunk.	07-11-2019
2.8	Updated this supported version in this guide to match the iApp template (11.4-14.1).	04-01-2020

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com