

F5 Integrations

Introduction

This document shows information related to F5 Integration.

The F5 BIG-IP integration allows users to monitor LTM, AFM, APM, ASM, and AVR activity. F5 BIG-IP covers software and hardware designed around application availability, access control, and security solutions.

The F5 BIG-IP integration can be used in three different modes to collect data:

HTTP Endpoint mode - F5 BIG-IP pushes logs directly to an HTTP endpoint hosted by users' Elastic Agent.

AWS S3 polling mode - F5 BIG-IP writes data to S3 and Elastic Agent polls the S3 bucket by listing its contents and reading new files.

AWS S3 SQS mode - F5 BIG-IP writes data to S3, S3 pushes a new object notification to SQS, Elastic Agent receives the notification from SQS, and then reads the S3 object. Multiple Agents can be used in this mode.

For example, users can use the data from this integration to analyze the traffic that passes through their F5 BIG-IP network.

Data streams

The F5 BIG-IP integration collects one type of data stream: log.

Log help users to keep a record of events happening on the network using telemetry streaming. The log data stream collected by the F5 BIG-IP integration includes events that are related to network traffic. See more details in the Logs.

This integration targets the five types of events as mentioned below:

LTM provides the platform for creating virtual servers, performance, service, protocol, authentication, and security profiles to define and shape users' application traffic. For more

information, refer to the link here.

AFM is designed to reduce the hardware and extra hops required when ADC's are paired with traditional firewalls and helps to protect traffic destined for the user's data center. For more information, refer to the link here.

APM provides federation, SSO, application access policies, and secure web tunneling and allows granular access to users' various applications, virtualized desktop environments, or just go full VPN tunnel. For more information, refer to the link here.

ASM is F5's web application firewall (WAF) solution. It allows users to tailor acceptable and expected application behavior on a per-application basis. For more information, refer to the link here.

AVR provides detailed charts and graphs to give users more insight into the performance of web applications, with detailed views on HTTP and TCP stats, as well as system performance (CPU, memory, etc.). For more information, refer to the link here.

Assumptions

The procedures described in Section 3 assumes that a Log Collector has already been setup.

Requirements

Elasticsearch is needed to store and search data, and Kibana is needed for visualizing and managing it. You can use our hosted Elasticsearch Service on Elastic Cloud, which is recommended, or self-manage the Elastic Stack on your hardware.

The reference link for requirements of telemetry streaming is here.

1. <https://clouddocs.f5.com/products/extensions/f5-telemetry-streaming/latest/prereqs.html>

The reference link for requirements of Application Services 3(AS3) Extension is here.

2. <https://clouddocs.f5.com/products/extensions/f5-appsvcs-extension/latest/userguide/prereqs.html>

This module has been tested against F5 BIG-IP version 16.1.0, Telemetry Streaming version 1.32.0 and AS3 version 3.40.0.

Setup

To collect LTM, AFM, APM, ASM, and AVR data from F5 BIG-IP, the user has to configure modules in F5 BIG-IP as per the requirements.

To set up the F5 BIG-IP environment, users can use the BIG-IP system browser-based Configuration Utility or the command line tools that are provided. For more information related to the configuration of F5 BIG-IP servers, refer to F5 support website [here](https://support.f5.com/csp/knowledge-center/software).

<https://support.f5.com/csp/knowledge-center/software>

Configuration of Telemetry Streaming in F5

For downloading and installing Telemetry Streaming, refer to the link [here](https://clouddocs.f5.com/products/extensions/f5-telemetry-streaming/latest/installation.html).

<https://clouddocs.f5.com/products/extensions/f5-telemetry-streaming/latest/installation.html>

Telemetry Streaming will send logs in the JSON format to the destination. Telemetry Streaming is compatible with BIG-IP versions 13.0 and later. Users have to prepare F5 servers for it and set up the Telemetry Streaming Consumer.

To use telemetry streaming, user have to send POST request on `https://<BIG-IP>/mgmt/shared/telemetry/declare` for declaration.

F5 BIG-IP modules named LTM, AFM, ASM, and APM are not configured by Telemetry Streaming, they must be configured with AS3 or another method. Reference link for setup AS3 extension in F5 BIG-IP is [here](https://clouddocs.f5.com/products/extensions/f5-telemetry-streaming/latest/event-listener.html?highlight=as3#configure-logging-using-as3).

To configure logging using AS3, refer to the [link here](https://clouddocs.f5.com/products/extensions/f5-telemetry-streaming/latest/event-listener.html?highlight=as3#configure-logging-using-as3).

<https://clouddocs.f5.com/products/extensions/f5-telemetry-streaming/latest/event-listener.html?highlight=as3#configure-logging-using-as3>

To collect data from AWS S3 Bucket, follow the below steps:

- Create an Amazon S3 bucket. Refer to the link [here](#).

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/create-bucket-overview.html>

- The default value of the "Bucket List Prefix" is listed below. However, the user can set the parameter "Bucket List Prefix" according to the requirement.

To collect data from AWS SQS, follow the below steps:

- If data forwarding to an AWS S3 Bucket hasn't been configured, then first set up an AWS S3 Bucket as mentioned in the above documentation.
- To set up an SQS queue, follow "Step 1: Create an Amazon SQS queue" mentioned in the [Documentation](https://docs.aws.amazon.com/AmazonS3/latest/userguide/ways-to-add-notification-config-to-bucket.html). <https://docs.aws.amazon.com/AmazonS3/latest/userguide/ways-to-add-notification-config-to-bucket.html>
 - While creating an SQS Queue, please provide the same bucket ARN that has been generated after creating an AWS S3 Bucket.
- Set up event notifications for an S3 bucket. Follow this [link](https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-event-notifications.html). <https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-event-notifications.html>
 - Users have to set the prefix parameter the same as the S3 Bucket List Prefix as created earlier. (for example, log/ for a log data stream.)
 - Select the event type as s3:ObjectCreated:*, select the destination type SQS Queue, and select the queue that has been created in Step 2.

Note:

- Credentials for the above AWS S3 and SQS input types should be configured using the [link](https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-input-aws-s3.html#aws-credentials-config). <https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-input-aws-s3.html#aws-credentials-config>
- Data collection via AWS S3 Bucket and AWS SQS are mutually exclusive in this case.

Enabling the integration in Elastic

1. In Kibana go to Management > Integrations.
2. In the "Search for integrations" search bar, type F5 BIG-IP.
3. Click on F5 BIG-IP integration from the search results.
4. Click on the Add F5 BIG-IP button to add F5 BIG-IP integration.
5. Enable the Integration to collect logs via AWS S3 or HTTP endpoint input.

F5 BIG-IP integration

The "Integration name" and either the "Description" and the following will need to be provided in the Configure integration when adding the F5 BIG-IP integration.

Procedures:

Please provide the following information to CyTech:

Collect F5 BIG-IP logs via HTTP Endpoint:

1. Listen Address - The bind address to listen for http endpoint connections. Set to 0.0.0.0 to bind to all available interfaces.

F5 BIG-IP logs via HTTP Endpoint:

1. Listen Port - The port number the listener binds to.

Revision #2

Created 23 April 2024 12:38:31

Updated 19 June 2024 06:54:01