

F5® BIG-IQ® Centralized Management: Data Collection Device Sizing Guide

Version 5.2.0



Table of Contents

Overview.....	3
What resources do I need, if I'm just collecting statistics?	4
What resources do I need to collect ASM event data?	5
What resources do I need to collect APM event data?	6
Using the APM data script.....	6
What resources do I need to collect FPS alert data?	8
Revision History.....	9
Legal notices	10
Trademarks.....	10
Usage disclaimer.....	10

Overview

This guide is intended to help you determine the resources that are required to handle the data generated by the BIG-IP® devices you manage. Requirements vary according to the type and amount of data you generate. Refer to the table(s) that correspond to the type of data that you want to collect from your BIG-IP devices.

Each table includes a Data Redundancy column to help you consider the additional data collection devices (DCDs) necessary to make sure that your data exists on more than one device.

- [What resources do I need, if I'm just collecting statistics?](#)
- [What resources do I need to collect ASM event data?](#)
- [What resources do I need to collect APM event data?](#)
- [What resources do I need to collect FPS alert data?](#)

What resources do I need, if I'm just collecting statistics?

What resources do I need, if I'm just collecting statistics?

Use this table to estimate the number of DCDs needed to collect and display statistical data from BIG-IP devices.

BIG-IP device statistics sizing detail

Number of BIG-IP devices	Data Redundancy	Recommended data collection device count
1-50	Yes	2
1-50	No	1
51-100	Yes	3
51-100	No	2
101-150	Yes	4
101-150	No	3
151+	Yes	5
151+	No	4 - 5

The following notes apply to this table:

- Provision all DCDs with 8 CPU and 32 GB virtual machines.
- The BIG-IQ[®] system can replicate statistical data across the DCD cluster. Data replication is useful for building in data redundancy in case one of your DCDs fail. Redundancy can also decrease the time it takes to render graphs, especially in environments with high scale or a large number of objects.
- If you have redundancy at the storage layer in the virtual environment infrastructure, redundancy at the BIG-IQ layer might not be necessary.
- Storage sizing is not estimated here, because storage requirements depend so much on configuration size, redundancy desires, retention policy, collection frequency, and BIG-IP count.

What resources do I need to collect ASM event data?

Use this table to estimate the number of DCDs that you need to collect and display ASM event data from BIG-IP devices.

If you are using an off-box logging solution to log ASM events, review a week's worth of alert logs to identify the logging load you need to support. Take note of both the average and maximum alerts per second.

BIG-IP device ASM sizing detail

Alerts/Sec	Data Redundancy	Recommended data collection device count
10000	No	1
14000	Yes	2
17000	Yes	3
20000	Yes	4
23000	Yes	5
26000	Yes	6

The following notes apply to this table:

- The sizing data in this table assumes that you provision your DCDs with 8 cores/CPU's and 32 GB of memory.
- When you choose an alert rate, consider both the average and peak loads that you anticipate. If your peak loads are frequent, consider provisioning additional data-collection nodes for the increased load. Also, when you determine your peak rates, consider the rate at which you expect your traffic load to increase, so you can plan for growth.
- In addition to the alert rate, consider the data retention policy you plan to use and the average size of the alerts you plan to store on the DCD. When you configure the retention and rotation settings for your alert data, you must consider the available storage on the DCD. Although DCD VE devices allow for expansion of the alert partition, we strongly recommend that you not exceed 50% of the available disk space, because without that disk space you will not be able to upgrade your software.

What resources do I need to collect APM event data?

What resources do I need to collect APM event data?

For APM®, we've created a spreadsheet and a script to help you determine the number of DCDs required to handle the data generated by the BIG-IP® devices you manage. You can download these tools from downloads.f5.com.

The script reads the last 10 days of logging data and provides the metrics you need for the spreadsheet. The spreadsheet uses those metrics to calculate your resource requirements. You enter information about the amount of data you expect to generate each day, and the spreadsheet calculates the number of DCDs required to handle that data.

Using the APM data script

F5 has provided a script that you can run to help you get a better idea of how much data you generate. The APM data script collects metrics for the data your APM devices generate and stores the data in a file on your BIG-IP device.

***Note:** You can only use the script if you have enabled local db logging. For information about enabling local-db logging, refer to [Overview: Configuring remote high-speed APM and SWG event logging on support.f5.com](#).*

If you use another logging mechanism, you can run reports (such as log reports or sessions reports) to get this data.

***Note:** Unless there has been a recent HA failure event (within the last 3 days), run this script on the active device. In this case, the active device may not have enough data and it is better to run the script on the standby device.*

Even with the script, it might be tedious to collect data from each of your BIG-IP devices, so we recommend that you take the values from one of the BIG-IP devices with the most users and access profiles to account for the worst case scenario.

To run the script, you must first download it.

1. Log in to downloads.f5.com.
2. Navigate to BIG-IQ Centralized Management, v5.2.0, and click **I Accept** to accept the EULA.
3. Click **F5_Networks_BIG_IQ_APM_DCD_Sizing_Tools.zip**.
After a brief pause, the Download Locations screen opens.
4. Click a download location to start the download.
5. When the download completes, unzip the file. Two files will extract:
F5_Networks_BIG_IQ_APM_Dcd_Sizing_Tool.xlsx and
F5_Networks_BIG_IQ_Apm_Dcd_Sizing_Metrics.sh.
6. Copy the script file
(F5_Networks_BIG_IQ_Apm_Dcd_Sizing_Metrics.sh) to the
BIG-IP device on which you want it to run.

What resources do I need to collect APM event data?

7. In the folder where you copied the script file, type:
Run `F5_Networks_BIG_IQ_Apm_Dcd_Sizing_Metrics.sh`.
The script runs and collects the data. The output is stored in a file named `BIG_IQ_Apm_Dcd_Sizing_Metrics.txt`
8. After the script has run, navigate to the `/var/tmp` folder on the BIG-IP device and open the output file to see the collected metrics.
A sample output file is provided below:

```
NUMBER_OF_EVENTLOGS_PER_DAY
400,000
-----
NUMBER_OF_SESSIONS_PER_DAY
20,000
-----
LOGS_PER_SESSION
20
-----
NUMBER_OF_ACCESSSS_PROFILES
10
-----
NUMBER_OF_ACLs
50
```

9. Open the spreadsheet and type the values from the script output file into the appropriate fields on the spreadsheet.
For example, for the sample output above, you would type 20,000 in the Number of sessions per day field.
10. After you fill in all of the required fields, the recommended number of DCDs is calculated and displayed in the **Estimated Number of DCDs** field.

What resources do I need to collect FPS alert data?

What resources do I need to collect FPS alert data?

Use this table to estimate the number of DCDs required to collect and display FPS alert data from BIG-IP devices.

If you are using an off-box logging solution to log FPS events, review a week's worth of alert logs to identify the logging load you need to support. Take note of both the average and maximum alerts per second.

BIG-IP device FPS sizing detail

Alerts/Sec	Data Redundancy	Recommended data collection device count
100	No	1
200	Yes	2
300	Yes	3
400	Yes	4
500	Yes	5
600	Yes	6

The following notes apply to this table:

- The sizing data in this table assumes that you provision your DCDs with 8 cores/CPU and 32 GB of memory.
- When you choose an alert rate, consider both the average and peak loads that you anticipate. If your peak loads are frequent, consider provisioning additional data-collection nodes for the increased load. Additionally, when you determine your peak rates, consider the rate at which you expect your traffic load to increase, so you have the room you need to grow.
- In addition to the alert rate, consider the data retention policy you plan to use and the average size of the alerts you plan to store on the DCD. When you configure the retention and rotation settings for your alert data, you must consider the available storage on the DCD. Although DCD VE devices allow for expansion of the alert partition, we strongly recommend that you not exceed 50% of the available disk space, because that disk space is essential for the upgrade process.
- The outgoing data rate from your BIG-IP devices also depends on the FPS policy configuration, the number of simultaneous application users, and their expected transaction rate.
- The scale numbers provided assume the use of three forwarding rules (100% syslog forwarding, 10% SOC forwarding, 10% customer forwarding), and 17000 transform rules.

Note: Increasing the number of forwarding targets or the distance between your data collection devices lessens the performance level of your DCD cluster.

Revision History

Date	Change log
04/26/2017	Release Version

Legal notices

Legal notices

Copyright © 2017 F5 Networks, Inc. All rights reserved.

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

Trademarks

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

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

Usage disclaimer

THE CONFIGURATION GUIDE IS PROVIDED "AS IS" AND POSSIBLY WITH FAULTS. F5 DISCLAIMS ANY AND ALL WARRANTIES AND GUARANTEES, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING BUT NOT LIMITED TO THE WARRANTY OF MERCHANTABILITY, THE WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTY OF NON-INFRINGEMENT, WITH RESPECT TO THE CONFIGURATION GUIDE, YOUR USE THEREOF, OR ANY RESULTS. CUSTOMER WILL USE THE CONFIGURATION GUIDE AT ITS OWN RISK. F5 WILL NOT BE LIABLE TO CUSTOMER FOR ANY DIRECT OR INDIRECT DAMAGES INCURRED IN USING THE PRODUCT AND IN NO EVENT WILL F5 BE LIABLE FOR LOSS OF PROFITS, LOSS OF USE, LOSS OF DATA, BUSINESS INTERRUPTION, NOR FOR PUNITIVE, INCIDENTAL, CONSEQUENTIAL, OR SPECIAL DAMAGES OF ANY KIND, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.