Signaling Delivery Controller

Troubleshooting Guide

5.1

Catalog Number: RG-016-51-42 Ver. 13

Publication Date: November 2020

Legal Information

Copyright

F5, 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

AskF5, F5, F5 [DESIGN], F5, OpenBloX, OpenBloX (design), Rosetta Diameter Gateway, Signaling Delivery Controller, SDC, Traffix, and Traffix [DESIGN] are trademarks or service marks of F5, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

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

Patents

This product may be protected by one or more patents indicated at: *http://www.f5.com/about/guidelines-policies/patents*

Confidential and Proprietary

The information contained in this document is confidential and proprietary to F5. The information in this document may be changed at any time without notice.

About F5

F5 (NASDAQ: FFIV) makes the connected world run better. F5 helps organizations meet the demands and embrace the opportunities that come with the relentless growth of voice, data, and video traffic, mobile workers, and applications—in the data center, the network, and the cloud. The world's largest businesses, service providers, government entities, and consumer brands rely on F5's intelligent services framework to deliver and protect their applications and services while ensuring people stay connected. For more information, visit *www.F5.com* or contact us at *Tfx_info@f5.com*.

About this Document

Document Name: F5 Signaling Delivery Controller Troubleshooting Guide

Catalog Number: RG-016-51-42 Ver. 13

Publication Date: November 2020

Document Objectives

This document describes the troubleshooting issues and their solutions for F5 SDC Release 5.1.

Document History

Revision Number	Change Description	Change Location
Ver. 2 - January 2017	Added Log Configuration section	*Configuring Log Files*
Ver. 3 - February 2017	Added procedure how to import Cassandra configuration files	*Importing Configuration Files for Lab Troubleshooting*
Ver. 4 - March 2017	Added procedures how to import Linux Routing, IPTables and SALT configuration files	*Importing Configuration Files for Lab Troubleshooting*
Ver. 5 - April 2017	Added troubleshooting workaround for adding routes	*Adding Routes*
Ver. 6 - August 2017	Added section how to export TDR reports	*Exporting TDR Reports*
Ver. 7 – September 2017	Added a procedure for changing a FEP IP address and a VLAN ID.	*Changing a FEP's IP Address; Changing a Site's VLAN ID*
Ver. 8 – November 2017	Added procedure how to configure number of displayed alarms	*Configuring the Number of Displayed Alarms*
Ver. 9 – December 2017	Changed wording in procedure step. Addedprocedure for adding	*Changing a FEP's IP Address; Changing a Site's VLAN ID;*

Revision Number	Change Description	Change Location
	additional supported application IDs	*Adding Additional Supported Application IDs*
Ver. 10 - January 2018	Added procedure to reduce disk usage for Cassandra	*Reducing Cassandra Disk Usage*
Ver. 11 - August 2019	Added procedure to delete server peers if new client peers were given the same name	*Configuring Peers*
Ver. 12 – July 2020	Added note about configuring TTA logs	*Collecting System Data*
Ver. 13 – November 2020	Splunk replaced with ELK	*Throughout document*

Conventions

The style conventions used in this document are detailed in Table 1.

Table 1: Conventions

Convention	Use
Normal Text Bold	Names of menus, commands, buttons, user-initiated CLI commands and other elements of the user interface
Normal Text Italic	Links to figures, tables, and sections in the document, as well as references to other documents
`Script`	Language scripts
`Courier`	File names
Note:	Notes which offer an additional explanation or a hint on how to overcome a common problem
Warning:	Warnings which indicate potentially damaging user operations and explain how to avoid them

Table of Contents

List of Figures

List of Tables

1. Troubleshooting Basics

This section describes recommended best practices to avoid errors and to help with troubleshooting when they do occur. To resolve specific issues, refer to the relevant chapter in this guide.

1.1 Referencing the SDC Documentation

The F5® Traffix® Signaling Delivery Controller™ (SDC) product documentation provides a comprehensive overview of system functionality. Some issues may be solved by consulting the relevant product documentation:

- *F5 SDC VMWare Virtual System Installation, Upgrade, and Maintenance Guide*
- *F5 SDC Openstack Virtual System Installation, Upgrade, and Maintenance Guide*
- *F5 SDC Bare Metal System Installation Guide*
- *F5 SDC Bare Metal System Upgrade Guide*
- *F5 SDC Bare Metal System Maintenance Guide*
- *F5 SDC CLI Application Guide*
- *F5 SDC Web Services API Guide*
- *F5 SDC User Guide*
- *F5 SDC SNMP Guide*
- *F5 SDC Release Notes*

Note: From 5.1 CF 30, EMS deployments will use ELK components, instead of Splunk components, to manage all SDC reporting functionalities. This change is reflected in version 13 and higher of the *Troubleshooting Guide*.

1.2 Verifying System Setup

Each SDC build supports specific SDC and third-party software (browsers, operating systems, etc.). Refer to the Release Notes to verify that your installation includes the recommended versions.

1.3 General Prevention

Make sure that all relevant machines are up and running, that all nodes are online and all relevant resources are started.

1.4 Collecting System Data

The SDC logs contain valuable data about your system activity. In order to collect all raw data, configuration, and logs, from the SDC, a specific TTA support script has been developed.

To collect system data:

1. Run the following script on one of the VInstaller servers:

   /opt/traffix/scripts/getTta.sh

   The script will run the TTA tool on all managed machines. Once the script has finished running, the following message will appear:

   “Create /var/log/rsyslog/tta-<yyyy-mm-dd-hh-mm-ss>.tar.gz on oam instances.”

2. Go to the `/var/log/rsyslog/` directory on one of the OAM (NMS) machines and download the .tar.gz archive.

3. Once the .tar.gz archive is validated, delete it from all the OAM (NMS) machines.

Note: From CF 29 and higher, during a rolling upgrade, you can configure the TTA script, with a flag (-x, -<exclude-old logs>, for example: ./tta-ng.sh -x -t /data -c 1) to ignore the logs of older SDC versions, so only logs from the most recent versions are retrieved.

1.5 Importing Configuration Files for Lab Troubleshooting

The following procedures describe how to import configuration files from a production to an F5 lab environment for troubleshooting. The following configuration files can be imported:

- Cassandra
- Linux routing
- IP tables rules
- Salt configuration

1.5.1 Prerequisite: Creating an SDC Lab Site Environment

Prior to importing any configuration files, you need to set up a lab environment.

1. Set up a new SDC lab site environment.
2. Start the system and allow the data folders to be written.

   Note: This may take a few minutes until folders are synchronized.

3. Stop all components (with the following command: **monit stop all**).

Note: Upon completing any import, restart all the SDC lab site components.

1.5.2 Importing Cassandra Configuration Files from an SDC Site

The following Cassandra configuration files can be imported:

- flowmanager.xml
- peer.xml
- peer-profile xml
- DEFAULT_LB_CONFIGURATION.xml

To import the Cassandra configuration files:

Note: When replacing the folder and file content, always locate and replace the latest version.

1. Set up a new SDC lab environment.
2. Start the system and allow the data folders to be written.

   Note: This may take a few minutes until folders are synchronized.

3. Stop all components (with the following command: **monit stop all**).
4. Delete the entire data folder in one of the configuration manager servers
5. On the other configuration manager server (in which the data folder was not deleted), locate the flowManager > ALL and DEFAULT folders.
   a. Within each of these folders, open the latesetVersion.xml to identify the most recent folder, with the following command:

      cat latestVersion.xml

      The latestVersion.xml folder points to the latest folder.

   b. Go to the folder that is identified from the cat latestVersion.xml and open it to locate the flowManager.xml.

Figure 1: Example of flowManager Folder Location

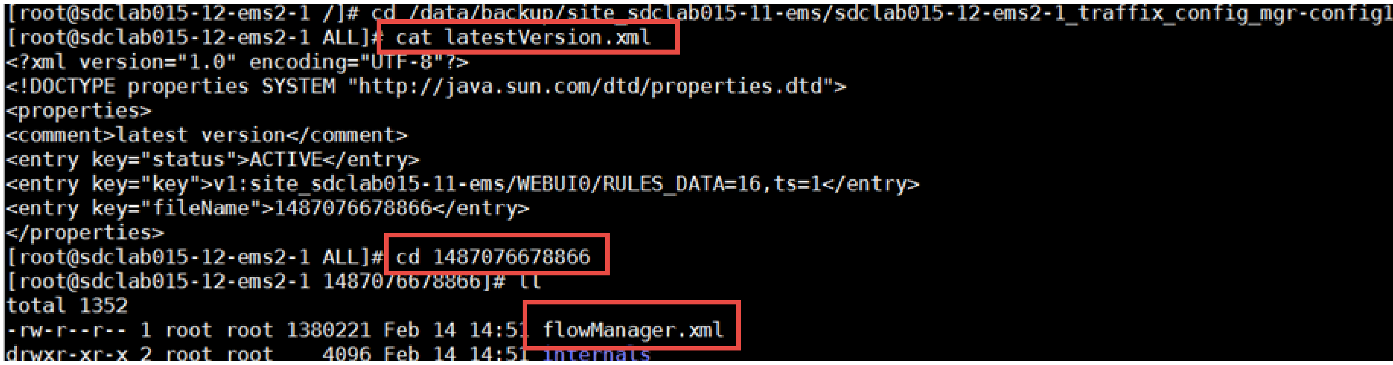

```
[root@sdclab015-12-ems2-1 /]# cd /data/backup/site_sdclab015-11-ems/sdclab015-12-ems2-1_traffix_config_mgr-config1
[root@sdclab015-12-ems2-1 ALL]# cat latestVersion.xml
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE properties SYSTEM "http://java.sun.com/dtd/properties.dtd">
<properties>
<comment>latest version</comment>
<entry key="status">ACTIVE</entry>
<entry key="key">v1:site_sdclab015-11-ems/WEBUI0/RULES_DATA=16,ts=1</entry>
<entry key="fileName">1487076678866</entry>
</properties>
[root@sdclab015-12-ems2-1 ALL]# cd 1487076678866
[root@sdclab015-12-ems2-1 1487076678866]# ll
total 1352
-rw-r--r-- 1 root root 1380221 Feb 14 14:51 flowManager.xml
drwxr-xr-x 2 root root    4096 Feb 14 14:51 internals
```

c. Copy the flowManager.xml from the production (customer) environment to the identified folder in the lab site.

d. Repeat these steps for the flowManager > DEFAULT folder.

6. On the configuration manager server (in which the data folder was not deleted), locate the latest version of the peerConfigurations folder and replace all of the content of the existing folder with the content of the peerConfigurations folder from the production environment.

7. On the configuration manager server (in which the data folder was not deleted), locate the of the DEFAULT_LB_CONFIGURATION folder and replace the content of the existing folder with the content of the latest version DEFAULT_LB_CONFIGURATION.xml from the production environment.

   a. Open the latestVersion.xml to identify the most recent folder, with the following command:

   cat latestVersion.xml

   The latestVersion.xml folder points to the latest folder.

   b. Go to the folder that is identified from the cat latestVersion.xml and open it to locate the DEFAULT_LB r.xml.

8. Copy the DEFAULT_LB_CONFIGURATION.xml from the production (customer) environment to the identified folder in the lab site.

9. On the configuration manager server (in which the data folder was not deleted), locate the latest version of the global folder and replace the content of the existing global.xml with the content of the global.xml from the production environment

    a. Open the latestVersion.xml to identify the most recent folder, with the following command:

    cat latestVersion.xml

    The latestVersion.xml folder points to the latest folder.

    b. Go to the folder that is identified from the cat latestVersion.xml and open it to locate the global.xml.

    c. Copy the global.xml from the production (customer) environment to the identified folder in the lab site

10. Start the configuration manager server (with the recently imported new files).

11. Start the other configuration manager server (that you deleted the data folder from).

12. Start all other components.

Note: To import configuration files from an EMS Site, delete the data folder in one of the EMS configuration manager servers and delete all data folders from each lab SDC site configuration manager server. In the EMS data folder, the global.xml and the flowmanager.xml files are by EMS site, and the DEFAULT_LB_CONFIGURATION.xml and the peerConfiguration folders are by SDC site.

Note: As not all of an environment's configuration files are imported, the **Administration** > **Backup And Restore** feature cannot be applied in the lab environment. and any component (i.e. FEP) specific routing rule will fail.

1.5.3 Importing the Linux Routing Files

To import the Linux routing files, you need to perform a backup of the Linux routing files on the SDC production site server and then copy the relevant archived configuration files that were created during the back-up process to the SDC lab site server:

To import the Linux routing files:

1. Run the following commands to create archived files on the production site:

   ```
   tar cvf keepalived.tar /etc/keepalived/*
   tar cvf routing.tar /etc/sysconfing/network-scripts/*
   ```

   The archived files are created.

2. Copy the archived files to the SDC lab site environment.

1.5.4 Importing the IPTables

The firewall rules are stored in the */etc/sysconfig/iptables* file and are applied whenever the service is started or a server is rebooted.

To import the IPTables:

1. Save the firewall rules with the following command on the SDC production server:

   ```
   service iptables save
   ```

2. Backup the */etc/sysconfig/iptables* and/or */etc/sysconfig/ip6tables* files.
3. Copy the saved iptables/ip6tables file to the SDC lab site server.
4. Start the service on the SDC lab site server with the following command:

service iptables start

The firewall rules from the copied iptables files are now saved on the SDC lab site server.

1.5.5 Importing the SALT Files

You can import the master server SALT files for troubleshooting,

To import the SALT files:

1. On one of the production master servers, locate the following file:

   *tar cvf SALT.tar /srv/salt/<version>/**

2. Copy the selected tar cvf *SALT.tar /srv/salt/<version>/** file to both SDC lab site master servers.
3. Extract the archived files to the lab site, with the following command:

   tar xvf /tmp/SALT.tar

4. On the SDC lab site server, replace the current (default) */srv/salt/<version>* content with the content from the imported tar file.

1.6 F5 Support Contact Information

Contact technical support at: HTTP://www.f5.com/support/support-services/contact//.

2. General Commands for Troubleshooting

2.1 General SALT Commands

The following table summarizes the most used salt commands that you will use when troubleshooting the deployment and site infrastructure. These commands must be run on one of the VInstaller machines:

Command	Description
salt '*' monit.summary	Shows the status of all processes on the managed machines, including the VInstaller machine.
salt '*' cmd.run "ls"	Runs the specified command on all managed machines.
salt 'hostname*' monit.status	Shows extensive information for all processes running on the specific machine.
salt '*' saltutil.sync_all	Synchronizes the VInstaller information with all managed machines.
salt-key	Shows which machines the VInstaller is managing, including the VInstaller machine.
salt '*' state.apply system.hosts	Applies a specific salt state.

2.2 General Monit Commands

The following table summarizes the most used Monit commands that you will use when troubleshooting the SDC site status and performance. These commands can be run on any site machine:

Command	Description
Monit reload	Reinitializes Monit.
Monit summary	Shows the status of all processes.
Monit status “processname”	Shows the status of the specific process.
Monit stop “processname”	Stops the specific process.

Monit start "processname"	Starts the specific process.

3. IP Connectivity

This section describes troubleshooting issues and solutions relating to IP connectivity.

3.1 SCTP Client Shows in GUI Too Many Addresses

3.1.1 Error Description

In general, SCTP's multi-homing enables any number of IP connections in one SCTP channel. The F5 SDC can only support at most two IP addresses for one SCTP channel. Each virtual server and each static client/server peer needs to be configured with at most two IP addresses.

3.1.2 Causes

Client tried to connect to more than two addresses.

3.1.3 Resolution

System's IP tables need to be configured such that only two IP addresses are connected to each SCTP connection.

4. SDC-Server Routes

Using API commands or the add/deleteRoute scripts, you can add or remove routes when servers are added/removed with networks that were not previously recognized by the SDC.

Note: Refer to the relevant F5 SDC Bare Metal or OpenStack Maintenance Guides for more information about adding routes.

4.1 Adding Routes

When running the addroute script or the addRoute API, the pillar (the dynamic component managed by Salt) may be missing vipInterfaces, causing the action to fail.

4.1.1 Error Description

If in the case the vipInterface parameter is missing perform the following workaround so that the route can be added.

1. On both master Installer servers in your setup, run the following command:

```
cd /srv/traffix/pillar/
rm .lastRev .oamdb_*
/etc/init.d/salt-master restart
```

2. On all FEP servers, run the following command:

```
cd /etc/sysconfig/network-scripts; rm -rf ifcfg-eth* ifcfg-bond*
```

3. Login to one of master servers and run the following command:

```
salt '*' state.highstate
```

5. SDC Pipeline

5.1 Licensing and Access Control

Remote client peer fails to open a link to the SDC.

5.1.1 Error Description

Client peer sends a proper CER to the SDC, but link is not established.

5.1.2 Causes

- The SDC is configured not to allow connection to unknown peers.
- The remote peer is sending the CER messages to an IP address not licensed by F5.
- No common application-IDs between client peer to SDC.

5.1.3 Resolution

- Check the CPF log for error messages (verify in all CPF nodes), the FEP logs for messages relating to licensing issues, and the NMS platform for SNMP traps
- If “peer rejected” messages appear:
- In the Web UI, select Topology and then in the Access Control List section, enable the Accept Unknown Peers checkbox.

Or

- Configure the remote peer manually and validate the remote peer connection to the destination IP. For example, run **netstat** on the FEP with the port of the Virtual Server.
- In the Web UI, select Administration tab and verify that a valid license exists for the relevant peer

5.2 Topology

5.2.1 Configuring Peers

5.2.1.1 Error Description

If a client peer is accidentally created with the same name that was used previously for a deleted server peer, there will be synchronization problems upon restarting the SDC components.

5.2.1.2 Causes

The synchronization problem occurs since a server peer is persistent, meaning that its data is persisted by the Configuration Manager (CM) in the file system, while a client peer is not persisted in the file system and is stored in memory only by the CM, FEP, CPF, and NMS.

5.2.1.3 Resolution

1. Delete the server peer in the Web UI.
2. Stop all the SDC components on the EMS and site (make sure that all traffic is moved to the other site).
3. Delete the server peer configuration from all the file systems in the EMS and site.
   The peer configuration is located in */data/backup/<site name>/<cm name>/peerConfigurations/<peer name>*
4. Start all components.

Note: After executing this procedure, it will not be possible to restore the configuration for any point before deleting the server peer files.

5.3 CPF Routing

This section describes commonly found errors, their causes, symptoms, and resolutions related to CPF routing functionality.

5.3.1 Request is Not Routed Using the Routing Rows as Expected

5.3.1.1 Error Description

When traffic is sent to SDC, and then a change is made to the routing table (i.e. added new rows, or edited existing rows), and then another request is sent to one of the servers, without it being routed using one of the added routing rows.

5.3.1.2 Causes

This is due to session stickiness for pools. The already existing session is being reused for the incoming request, and the sessions bypasses the routing table.

5.3.1.3 Symptoms

- The pool that should be the destination for the traffic does not accept it (if it is a different pool than before the routing change).
- Routing scripts, (Check Error in Answer, Handle Server Error, etc.) that belong to the new routing row that should be selected, do not run.
- The logging is in TRACE mode for Transaction Management, and a message of the format “Pool {0} was selected for {1}” did not appear in CPF logs since the routing change.

5.3.1.4 Resolution

Change the session ID for the next request or wait for session expiration (default is 30 seconds).

5.3.2 Session Management Rule is Missing

5.3.2.1 Error Description

Each event that passes through the SDC, should have a session management rule that applies to it, whether it is a stateless or stateful session. In the event that there is no such correlating session management rule, a WARN message is generated in the CPF logs.

5.3.2.2 Symptoms

- The following WARN message is generated in the CPF logs: "not found in storage or mgmt table"

5.3.2.3 Resolution

Configure a session management rule for each stateless or stateful session.

5.3.3 No Pools are Selected for Routing

5.3.3.1 Error Description

No messages are routed to the pools configured in the routing rule, even though the rule was correctly configured with the ROUTE action. Error messages display in the log as described in the Symptoms section below.

5.3.3.2 Causes

- A pool was not selected in a routing row.
- All pools are in “Out of Service” state (since all peers of each of these pools are in a “Close” state (disabled or not yet connected) or “Out of Service”).
- At least one pool is in an “Open” state – but all of its peers are overloaded (reached maximum rate limit), and all other pools, if exist, are in “Out of Service” state.

Note: A pool will be in “Open” state when at least its “Minimum Number of Peers” (configurable, default is 1) is reached. That means that there are no “Minimum Number of Peers” peers in this pool such that their state is “Open”.

5.3.3.3 Symptoms

The following INFO message may appear in the CPF logs:

- “Unable to choose pool: {0}, reason: {1}” for each of the pools that belong to the selected routing row.

The following WARN message may appear in the CPF logs:

- "Failed to select a Pool to handle a request received from {0}. The selected routing row index is {1} with policy {2}, the incoming message is {3}".

5.3.3.4 Resolution

- Check the condition of the peers of each pool of the selected routing row to make sure they are not overloaded/closed/disabled, etc.
- Change the logging of TransactionManagement to DEBUG and follow the CPF logs to trace the pool selection.

5.3.4 Endless Pending Request Timeouts toward Client

5.3.4.1 Error Description

When CPF reaches a state where it cannot route a request to a server, it sends an error answer to the client. This error answer can be edited for each row of the routing table using the "Handle Server Error" scripts. Faulty scripts can cause SDC to behave very strangely. For example, if a script returns answerFromServer and the RemoteNodeEvent is CANNOT_ROUTE, then the message that is sent to client is the request that was sent by it.

Since it is a request that goes downstream, now CPF will set it as a pending request with a timeout, and since clients usually do not respond to requests, the pending timeout will be invoked and cause another endless cycle of requests that will be sent to the client.

5.3.4.2 Causes

A bad "Handle Server Error" script is configured for the routing row that was selected for the request.

5.3.4.3 Resolution

In the “Handle Server Error” script, reconfigure the “answerFromServer” parameter for any RemoteNodeEvent. For example, for CANNOT_ROUTE and TIMEOUT there are no answers from server, so answerFromServer cannot be returned.

5.3.5 Routing of Server Side Request (CLR) Fails

5.3.5.1 Error Description

A CLR request arrives and is then forwarded (based on a forward routing rule) to a client peer whose name appears as the Destination-Host AVP of the request. When this routing rule is configured as a Roaming Proxy, and then post transformation is done on the Destination-Host AVP before the request searches for the destination peer with this name. If after post transformation, the peer with this name is not found, then routing fails.

5.3.5.2 Causes

A server side request, such as CLR is generated in response to a client side request, such as ULR (when a CLR is sent by the server then the Destination-Host AVP is taken from the Origin-Host of the ULR). The routing rule of the client side requests must also be marked as roaming proxy otherwise the CLR routing will fail.

Each client peer at the SDC from where the ULR messages come from must have a peer profile (not default), because a peer profile name is used at translation of the destination peer.

5.3.5.3 Symptoms

The following section describes the error conditions and their relevant error messages.

- Client side request, like ULR, routed through SDC, with Roaming Proxy is enabled. The routing of ULR was successful, but peer profile at the client peer is not configured.

- Error Message: “Diameter client peer {some peer name} must have a peer profile for Roaming Proxy full functionality. Routing of future requests from server will fail!”

Routing of “Forwarded” message of server side request, such as CLR, when roaming is enabled is failed.

- Error Message: “Routing of roaming request from server failed, no suitable client peer found, at event {some event description}”
- Error Message: After the above message, there will be a regular SDC routing failed message.

5.3.5.4 Resolution

The following are ways resolve the issue:

- Verify that the client peer of the SDC through which the ULR messages have been sent has a non-default peer profile.
- Verify that the “Forward” routing rule of server side request, such as CLR, is configured as a roaming proxy and that its related client side, such as ULR, is also configured with a roaming proxy.
- Verify that there are no changes were made to SDC configuration in the time gap between the ULR request and its matching CLR request routing. For example, that the peer profile name of a client peer was not changed.

5.3.6 Request does not Arrive to the Correct Server Inside its Pool

5.3.6.1 Error Description

Traffic is sent to CPF and is routed to the correct pool. However, after changing the pool to contain different servers, the traffic that arrives to CPF is still routed to the previously selected servers that were contained by the pool.

5.3.6.2 Causes

Session stickiness of the destination peer may cause the previously selected session to be reused for the incoming request. Sessions bypass the more recently selected peer selection from within pool.

5.3.6.3 Symptoms

- The new servers of the pool do not accept the traffic.
- The logging is in DEBUG mode for TransactionManagement, and the following message is displayed: "Destination {0} had been selected for {1}" in the CPF logs of the previously selected peers.

5.3.6.4 Resolution

- Change the session ID for the next requests or wait for session expirations (default is 30 seconds).

5.3.7 Adding Additional Supported Application IDs

Diameter application IDs define the Diameter messages that the SDC Component may handle. Refer to the *F5 SDC 5.1 User Guide* Appendix B: Supported Application Identifiers, for a full list of the supported applications. From 5.1 CF-12, you can add additional application IDs in the standard_dynamic_example.txt file (located in the /opt/traffix/sdc/config folder) or using the Web Service API Method: SetDiameterPropertiesforNode, in the supportedApplicationIds parameter. For more information on using the Web Services API, see the *F5 SDC Web Services API Guide*.

5.3.7.1 Resolution

1. Open the standard_dynamic_example.txt file (located in the /opt/traffix/sdc/config folder).

Note: The file is an example file.

For EMS deployments, you need to update the standard_dynamic.txt file in the EMS site and not the SDC site. Changes made to the file in the EMS site will be implemented in the SDC site and reflected in the SDC site Web UI.

2. Update the standard_dynamic.txt file for any additional Application Identifiers. The following fields need to be defined:

   Name

   applicationId

   protocol

   vendorId

   applicationType

   fsmType

Figure 2: Example of Added Application IDs to a standard_dynamic.txt File

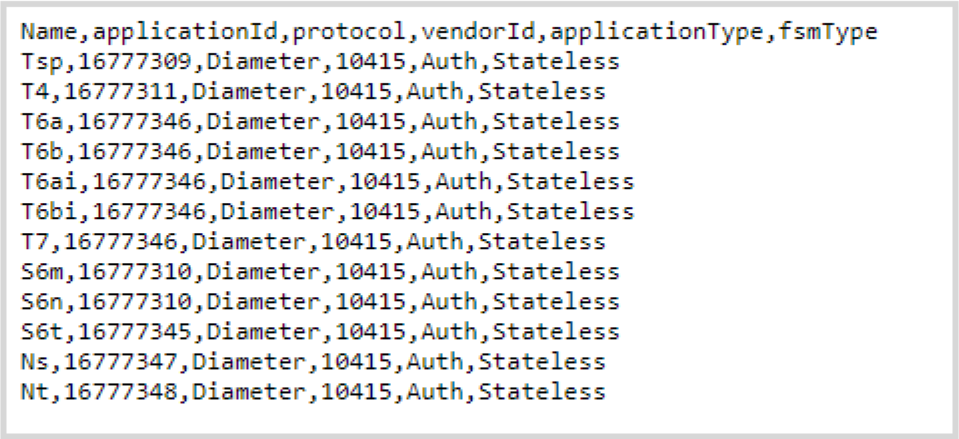

```
Name,applicationId,protocol,vendorId,applicationType,fsmType
Tsp,16777309,Diameter,10415,Auth,Stateless
T4,16777311,Diameter,10415,Auth,Stateless
T6a,16777346,Diameter,10415,Auth,Stateless
T6b,16777346,Diameter,10415,Auth,Stateless
T6ai,16777346,Diameter,10415,Auth,Stateless
T6bi,16777346,Diameter,10415,Auth,Stateless
T7,16777346,Diameter,10415,Auth,Stateless
S6m,16777310,Diameter,10415,Auth,Stateless
S6n,16777310,Diameter,10415,Auth,Stateless
S6t,16777345,Diameter,10415,Auth,Stateless
Ns,16777347,Diameter,10415,Auth,Stateless
Nt,16777348,Diameter,10415,Auth,Stateless
```

3. Rename the file to standard_dynamic.txt and save it to */opt/traffix*

4. Check that any newly added application IDs are defined in the Dictionary. If not, you need to update the Dictionary.

5. Restart the Web UI, tomcat services, CFP and FEP components.

Note: Upon restarting the Web UI, remove any cookies from the web browser.

6. Verify that any added Application ID appears in the SDC Web UI (**Topology** > **Specific Site Settings** > **Site** > **SDC Components** > **Diameter**).

Note: When removing an application ID from the standard_dynamic.txt file, it will not be removed from the SDC Web UI (**Topology** > **Specific Site Settings** > **Site** > **SDC Components** > **Diameter**), until you restart the Web UI.

5.4 CPF Transformation

5.4.1 CPF Dictionary

Each Diameter network element holds its own dictionary. A successful Diameter connection between two network elements requires compatible dictionaries that maintain the same AVP data message format. All AVPs included in a Diameter dictionary must have a unique AVP name.

5.4.1.1 Mismatch between Multiple AVPs

5.4.1.2 Error Description

When CPF loads a Diameter dictionary, that has two AVPs defined with the same name, but with different commands or vendor IDs, only the first AVP is saved in the application.

5.4.1.3 Causes

Using a Diameter dictionary in which the AVP name is not unique.

5.4.1.4 Symptoms

The following section describes the error conditions and their relevant error messages.

The Diameter dictionary used by CPF contains several AVPs that do not have a unique AVP name.

For example, the following two AVPs contain the same name and command, but different vendor IDs.

- <avp name="Service-Selection" code="493" format="utf8String" mRule="must" vendorId="10415"/>
- <avp name="Service-Selection" code="493" format="utf8String" mRule="must" vendorId="0"/>

This results in a mismatch between the two AVPS, as they are not unique, though they have the same name and the following message appears in the log file:

- ERROR [10155] Diameter Dictionary (SDC dictionary v9): Mismatch between AVP definition <Content Definition 3GPP-Charging-Characteristics 13 10415 UTF8String> and <Content Definition 3GPP-Charging-Characteristics 13 UTF8String>: Trying to load two AVPs with identical names and different codes. [jmsContainer-1_11] [DiameterDictionary.mismatchDetected()]
- 2013-10-15 15:23:23,006 ERROR [10155] Diameter Dictionary (SDC dictionary v9): Mismatch between AVP definition <Content Definition Service-Selection 493 10415 UTF8String> and <Content Definition Service-Selection 493 UTF8String>: Trying to load two AVPs with identical names and different codes. [jmsContainer-1_11] [DiameterDictionary.mismatchDetected()]

5.4.1.5 Resolution

For Diameter application messages, use the super dictionary. The super dictionary contains only AVPs with unique names.

The base part of the super dictionary contains messages with application ID 0. These messages are used by more than one Diameter Interface (application ID).

The following are options on how to resolve different AVP dictionary issues:

To add a message with an application ID according to spec (not 0):

- Create a message so that the name of this message will be built as concatenation of name and needed interface. For example, for a S6b application:
  - <! -- applicationId="16777272" -->
  - <message name="RAR-S6b" code="258" applicationId="16777272" isProxiable="true" isRequest="true" sentByClient="false"/>
  - <message name="RAA-S6b" code="258" applicationId="16777272" isProxiable="true" isRequest="false" sentByClient="true"/>

To add AVPs for a specific vendor whose codes are used by other interfaces:

- Add a prefix (vendor name) to the AVP name. For example:
  - <! -- vendorId="12645" -->
  - <avp name="Vodafone-Radio-Access-Technology" code="260" format="enumerated" mRule="may" vendorId="12645"/>

To distinguish between two AVPs that have the same name, but a different code in the interface:

- Add suffix that defines the data type of AVP. For example, in 3GPP Vendor-Id 10415 Service-Type AVP has three commands code, so can add its data type at the end:
  - <avp name="Service-Type-Grouped" code="1483" format="grouped" mRule="must" vendorId="10415"/>
  - <avp name="Service-Type-Unsigned32" code="2031" format="unsigned32" mRule="must" vendorId="10415"/>
  - <avp name="Service-Type-Enumerated" code="6" format="enumerated" mRule="must" vendorId="10415"/>

If there is more than one AVP whose name ends with an “Address” string and their data type are addresses too, then the first character before the “Address” should be changed. For example:

- <avp name="Served-Party-IP-Address" code="848" format="address" mRule="must" vendorId="10415"/>
- <avp name="Served-Party-Ip-Address" code="248" format="address" mRule="must" vendorId="10415"/>

5.4.2 Message Parsing Failures

Message parsing can fail due to multiple reasons. This section presents some parsing troubleshooting issues including providing instructions on how to investigate them in order to find out the root cause of the issue.

5.4.2.1 Error Description

CPF fails to parse some Diameter AVPs with an internal error due to a thrown InvalidAvpLengthValidationException. The message handling continues but a Wireshark capture shows that the incoming message to SDC and outgoing message from SDC are different (no transformation has occurred). The outgoing message is marked as an “Unreassembled Packet” and Wireshark is not able to parse all the AVPs in the message, which were parsed successfully in the incoming message, (as shown in the screenshots below).

Figure 3: Unreassembled Packet

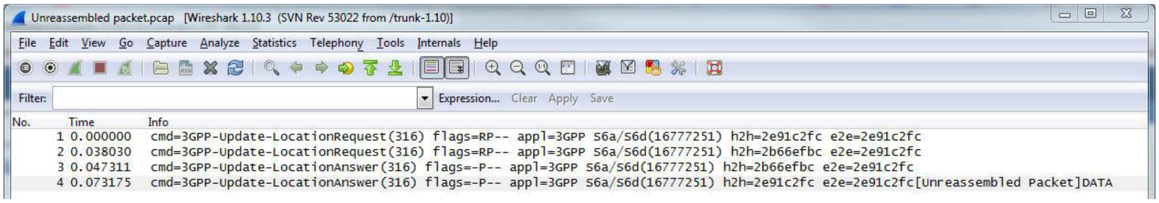
Unreassembled packet.pcap [Wireshark 1.10.3 (SVN Rev 53022 from /trunk-1.10)]
File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help
Filter:
Expression... Clear Apply Save

```
No.     Time      Info
      1 0.000000  cmd=3GPP-Update-LocationRequest(316) flags=RP-- appl=3GPP S6a/S6d(16777251) h2h=2e91c2fc e2e=2e91c2fc
      2 0.038030  cmd=3GPP-Update-LocationRequest(316) flags=RP-- appl=3GPP S6a/S6d(16777251) h2h=2b66efbc e2e=2e91c2fc
      3 0.047311  cmd=3GPP-Update-LocationAnswer(316) flags=-P-- appl=3GPP S6a/S6d(16777251) h2h=2b66efbc e2e=2e91c2fc
      4 0.073175  cmd=3GPP-Update-LocationAnswer(316) flags=-P-- appl=3GPP S6a/S6d(16777251) h2h=2e91c2fc e2e=2e91c2fc[Unreassembled Packet]DATA
```

Figure 4: Incoming Message AVPs

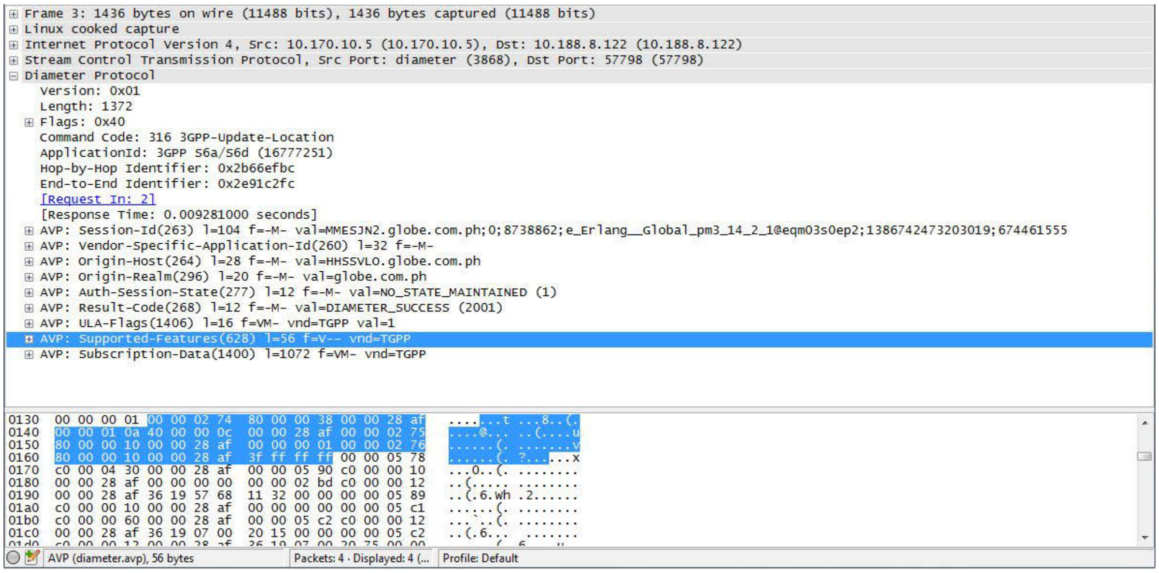

```
⊞ Frame 3: 1436 bytes on wire (11488 bits), 1436 bytes captured (11488 bits)
⊞ Linux cooked capture
⊞ Internet Protocol Version 4, Src: 10.170.10.5 (10.170.10.5), Dst: 10.188.8.122 (10.188.8.122)
⊞ Stream Control Transmission Protocol, Src Port: diameter (3868), Dst Port: 57798 (57798)
⊟ Diameter Protocol
    Version: 0x01
    Length: 1372
  ⊞ Flags: 0x40
    Command Code: 316 3GPP-Update-Location
    ApplicationId: 3GPP S6a/S6d (16777251)
    Hop-by-Hop Identifier: 0x2b66efbc
    End-to-End Identifier: 0x2e91c2fc
    [Request In: 2]
    [Response Time: 0.009281000 seconds]
  ⊞ AVP: Session-Id(263) l=104 f=-M- val=MMESJN2.globe.com.ph;0;8738862;e_Erlang__Global_pm3_14_2_1@eqm03s0ep2;1386742473203019;674461555
  ⊞ AVP: Vendor-Specific-Application-Id(260) l=32 f=-M-
  ⊞ AVP: Origin-Host(264) l=28 f=-M- val=HHSSVLO.globe.com.ph
  ⊞ AVP: Origin-Realm(296) l=20 f=-M- val=globe.com.ph
  ⊞ AVP: Auth-Session-State(277) l=12 f=-M- val=NO_STATE_MAINTAINED (1)
  ⊞ AVP: Result-Code(268) l=12 f=-M- val=DIAMETER_SUCCESS (2001)
  ⊞ AVP: ULA-Flags(1406) l=16 f=VM- vnd=TGPP val=1
  ⊞ AVP: Supported-Features(628) l=56 f=V-- vnd=TGPP
  ⊞ AVP: Subscription-Data(1400) l=1072 f=VM- vnd=TGPP

0130  00 00 00 01 00 00 02 74  80 00 00 38 00 00 28 af   .......t ...8..(.
0140  00 00 01 0a 40 00 00 0c  00 00 28 af 00 00 02 75   ....@... ..(....u
0150  80 00 00 10 00 00 28 af  00 00 00 01 00 00 02 76   ......(. .......v
0160  80 00 00 10 00 00 28 af  3f ff ff ff 00 00 05 78   ......(. ?......x
0170  c0 00 04 30 00 00 28 af  00 00 05 90 c0 00 00 10   ...0..(. ........
0180  00 00 28 af 00 00 00 00  00 00 02 bd c0 00 00 12   ..(..... ........
0190  00 00 28 af 36 19 57 68  11 32 00 00 00 00 05 89   ..(.6.Wh .2......
01a0  c0 00 00 10 00 00 28 af  00 00 00 00 00 00 05 c1   ......(. ........
01b0  c0 00 00 60 00 00 28 af  00 00 05 c2 c0 00 00 12   ...`..(. ........
01c0  00 00 28 af 36 19 07 00  20 15 00 00 00 00 05 c2   ..(.6...  .......

○ AVP (diameter.avp), 56 bytes | Packets: 4 · Displayed: 4 (... | Profile: Default
```

Figure 5: Outgoing Message AVPs

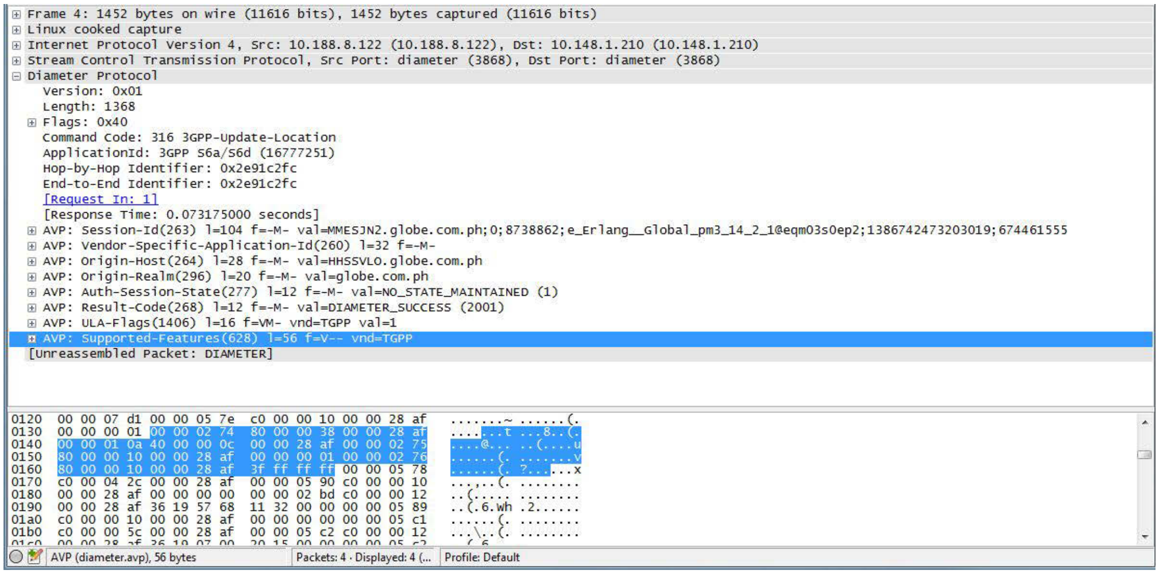

```
⊞ Frame 4: 1452 bytes on wire (11616 bits), 1452 bytes captured (11616 bits)
⊞ Linux cooked capture
⊞ Internet Protocol Version 4, Src: 10.188.8.122 (10.188.8.122), Dst: 10.148.1.210 (10.148.1.210)
⊞ Stream Control Transmission Protocol, Src Port: diameter (3868), Dst Port: diameter (3868)
⊟ Diameter Protocol
    Version: 0x01
    Length: 1368
  ⊞ Flags: 0x40
    Command Code: 316 3GPP-Update-Location
    ApplicationId: 3GPP S6a/S6d (16777251)
    Hop-by-Hop Identifier: 0x2e91c2fc
    End-to-End Identifier: 0x2e91c2fc
    [Request In: 1]
    [Response Time: 0.073175000 seconds]
  ⊞ AVP: Session-Id(263) l=104 f=-M- val=MMESJN2.globe.com.ph;0;8738862;e_Erlang__Global_pm3_14_2_1@eqm03s0ep2;1386742473203019;674461555
  ⊞ AVP: Vendor-Specific-Application-Id(260) l=32 f=-M-
  ⊞ AVP: Origin-Host(264) l=28 f=-M- val=HHSSVLO.globe.com.ph
  ⊞ AVP: Origin-Realm(296) l=20 f=-M- val=globe.com.ph
  ⊞ AVP: Auth-Session-State(277) l=12 f=-M- val=NO_STATE_MAINTAINED (1)
  ⊞ AVP: Result-Code(268) l=12 f=-M- val=DIAMETER_SUCCESS (2001)
  ⊞ AVP: ULA-Flags(1406) l=16 f=VM- vnd=TGPP val=1
  ⊞ AVP: Supported-Features(628) l=56 f=V-- vnd=TGPP
  [Unreassembled Packet: DIAMETER]

0120  00 00 07 d1 00 00 05 7e  c0 00 00 10 00 00 28 af   .......~ ......(.
0130  00 00 00 01 00 00 02 74  80 00 00 38 00 00 28 af   .......t ...8..(.
0140  00 00 01 0a 40 00 00 0c  00 00 28 af 00 00 02 75   ....@... ..(....u
0150  80 00 00 10 00 00 28 af  00 00 00 01 00 00 02 76   ......(. .......v
0160  80 00 00 10 00 00 28 af  3f ff ff ff 00 00 05 78   ......(. ?......x
0170  c0 00 04 2c 00 00 28 af  00 00 05 90 c0 00 00 10   ...,..(. ........
0180  00 00 28 af 00 00 00 00  00 00 02 bd c0 00 00 12   ..(..... ........
0190  00 00 28 af 36 19 57 68  11 32 00 00 00 00 05 89   ..(.6.Wh .2......
01a0  c0 00 00 10 00 00 28 af  00 00 00 00 00 00 05 c1   ......(. ........
01b0  c0 00 00 5c 00 00 28 af  00 00 05 c2 c0 00 00 12   ...\..(. ........

○ AVP (diameter.avp), 56 bytes | Packets: 4 · Displayed: 4 (... | Profile: Default
```

5.4.2.2 Causes

The cause for the illegal parsed value which then failed the parsing can be due to a buffer offset that was incorrectly incremented while parsing the message. A wrong buffer offset can cause AVPs to be parsed from the wrong index which can cause illegal values to be read. Such errors in a buffer offset can be caused by a wrong AVP description in the used dictionary. Finding the location from which the buffer started incrementing incorrectly can lead to the root cause of the issue

Example:

The CPF Error message displays the failed AVP values. For example, the value in the error message in Section 5.3.2.3 is code=83886080 (0x5000000) and length=1435 (0x059b).

Searching the incoming message for the length hex value found that the failed AVP length is an AVP code of "AMBR" AVP. The failed AVP code was found to be the data of the AVP which came before the "AMBR" AVP.

Figure 6: Parsing Offset

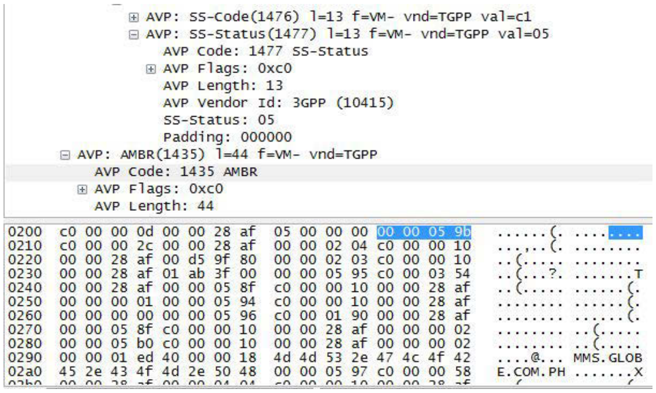

```
        ⊞ AVP: SS-Code(1476) l=13 f=VM- vnd=TGPP val=c1
        ⊟ AVP: SS-Status(1477) l=13 f=VM- vnd=TGPP val=05
            AVP Code: 1477 SS-Status
          ⊞ AVP Flags: 0xc0
            AVP Length: 13
            AVP Vendor Id: 3GPP (10415)
            SS-Status: 05
            Padding: 000000
    ⊟ AVP: AMBR(1435) l=44 f=VM- vnd=TGPP
        AVP Code: 1435 AMBR
      ⊞ AVP Flags: 0xc0
        AVP Length: 44

0200  c0 00 00 0d 00 00 28 af  05 00 00 00 00 00 05 9b   ......(. ........
0210  c0 00 00 2c 00 00 28 af  00 00 02 04 c0 00 00 10   ...,..(. ........
0220  00 00 28 af 00 d5 9f 80  00 00 02 03 c0 00 00 10   ..(..... ........
0230  00 00 28 af 01 ab 3f 00  00 00 05 95 c0 00 03 54   ..(...?. .......T
0240  00 00 28 af 00 00 05 8f  c0 00 00 10 00 00 28 af   ..(..... ......(.
0250  00 00 00 01 00 00 05 94  c0 00 00 10 00 00 28 af   ........ ......(.
0260  00 00 00 00 00 00 05 96  c0 00 01 90 00 00 28 af   ........ ......(.
0270  00 00 05 8f c0 00 00 10  00 00 28 af 00 00 00 02   ........ ..(.....
0280  00 00 05 b0 c0 00 00 10  00 00 28 af 00 00 00 02   ........ ..(.....
0290  00 00 01 ed 40 00 00 18  4d 4d 53 2e 47 4c 4f 42   ....@... MMS.GLOB
02a0  45 2e 43 4f 4d 2e 50 48  00 00 05 97 c0 00 00 58   E.COM.PH .......X
```

This analysis shows that the error in the offset started when parsing the AVP which came before the "AMBR" AVP which is the "SS-Status" AVP.

After debugging the parsing of the "SS-Status" AVP, it was found that this AVP was parsed as a grouped AVP while it did not hold any grouped information. The reason for the parsing error was an error in the Diameter dictionary in which the "SS-Status" AVP was marked as a grouped AVP instead of as an octet sting AVP. Changing the AVP type in the dictionary resolved the issue.

5.4.2.3 Symptoms

The following section describes the error condition and its relevant error message.

- The buffer offset is out of sync with the message AVPs during parsing. CPF starts parsing AVPs from the wrong buffer index causing illegal field values to be read.
- Error Message:

  Internal warning: An attempt to create an AVP list from SlicedChannelBuffer(ridx=0, widx=1328, cap=1328) had failed. [Client Worker_4_11] [DiameterParsedGroupedAvp.getVendorSpecificAvpSet()] com.traffix.openblox.Diameter.exceptions.InvalidAvpLengthValidationException: Invalid avp length 1435 to avp with code 83886080 and vendorId 0.

5.4.2.4 Resolution

When encountering such parsing issues, you need to find the root cause (with Wireshark) to resolve the issue.

To find the root cause of the AVP parsing error:

1. Take a Wireshark capture of the failed transaction.
2. Compare the message coming into the SDC to the message coming out of the SDC to see at what step Wireshark could not parse the AVPs.
3. Search for the failed AVP code and length value in the message coming into the SDC at the segments where Wireshark was not able to parse the message coming out of the SDC.
4. Identify the last AVP before the failed AVP from the error log.
5. Check the AVP definition in the used dictionary and compare it to the AVP which was sent to SDC.
6. Correct any dictionary mistake according to the actual sent data and the latest 3GPP Diameter application codes and identifiers specification document (for example, TS 29.230).

Note: In case of conflict between the received data and the latest Diameter application specification document, more detailed specification documents (such as TS 29.272 in the described example) should be referenced to verify that there are no mistakes in the Diameter application specification document.

5.4.3 Configured Transformation Does Not Take Effect

5.4.3.1 Error Description

After configuring a Diameter Identity for some Routing Rules, the configured transformation does not take place. For example, the Destination Host/Realm of some request was not changed as expected when the request was routed through SDC.

5.4.3.2 Causes

After the configuration, the request on which the transformation rule was configured was wrong, as it was a request from an existing session.

5.4.3.3 Symptoms

There are no special log errors/warnings for this problem.

5.4.3.4 Resolution

The transformation rule change is only visible once a new session is initiated and the request from the client is sent, as only then the change takes place.

When a new session is created, the Diameter identity values that were configured for that session message are stored in Tripo.

For subsequent messages with the same session Id, the Diameter Identity values are taken from the stored data and not from the currently configured values in the Diameter identity section of the relevant routing row.

5.4.4 3GPP Destination Realm Normalization Does Not Work

5.4.4.1 Error Description

A routed request's destination realm is not normalized although it was configured.

5.4.4.2 Symptoms

The following section describes the error conditions and their relevant error messages.

The AVP containing the IMSI from which MNC and MCC is calculated, but was not found at the request.

- Error Message: "IMSI avp is not found, destination realm normalization will not work, at message {some request description}."

The parsing of the IMSI number to MNC/MCC failed.

- Error Message: "IMSI parsing failed, 3GPP realm normalization was canceled. Imsi: {the IMSI number}. Cause: {The cause for the fail}."

5.4.4.3 Resolution

The following are options on how to resolve the issue:

- The SDC takes the IMSI number from the following AVP's: For S6a or S6b application ID's the AVP is "User-Name". For Gx, Gy, Rx, Rf, Sy, S9 the AVP is inside the grouped AVP "Subscription-Id" and the one with AVP of " Subscription-Id-Type" equals to 1.

Note: Make sure that this AVP (User Name/Subscription-Id) is present at the request.

- Make sure the number is legal according to the RFC specification.
- When the "IMSI parsing failed…" error appears in the logDescription.txt file, look for the error, for example: IllegalArgumentException ("IMSI length must be 14 or 15 digits"), to correct the relevant input parameters.

- Make sure this error was not caused by a configured transformation error (as described in *Section 5.3.3*.)

6. Performance

6.1 HTTP Performance is Degraded

6.1.1 Error Description

The TPS of the HTTP routing is much slower than expected.

6.1.2 Causes

- Keep-alive: The SDC is not using keep-alive.
- Number of maximum connections (**Max Connections Count Limit (Per Server)** configured when adding a Remote Peer) between the SDC and the server is too small to support the traffic load.
- Number of maximum connections between the SDC and the client is too small to support the traffic load.
- The HTTP virtual server disconnects after each response.
- Hosting machines are not strong enough. VM is being used.
- Size of messages is too big.

6.1.3 Symptoms

- Timeouts appear for the client or in the CPF. Timeouts by our side appear in the CPF log with the format "Peer Timeout event occurred for {0}, message {1}".
- On the server side, there are many TCP connections that are in a state: TIME_WAIT.

6.1.4 Resolution

- Check that the HTTP virtual server's configuration for **Close Connection on Answer** is not enabled.
- Verify that **Keep Alive** is enabled when configuring an HTTP peer on the SDC (server/client).

- Increase the **Max Connections Count Limit (Per Server)** configuration for the HTTP server peer.

Note: The default value is 10. Generally, there should be 20% more connections between the SDC and the server than between the client peer

- Verify that the Max Connections Count Limit (Per Client) configuration for the HTTP client peer value is configured to support expected traffic load.

Note: The default connection size is 1024.

6.2 Logging Causes System Performance Degradation

6.2.1 Error Description

The SDC logging mechanism can impact the performance of SDC machines (especially machines running Tripo application instances).

6.2.1 Error Description

To minimize the system impact of logging, apply a logging rate limit:

1. On both Installer machines, locate the following file:

```
/srv/salt/<ISO_Version>/rsyslog/rsyslog.conf
```

2. In the `rsyslog.conf` file, modify the following lines with the desired rate limits:

Note: By default, these lines are configured with a value of “0”, which means that the rate limit mechanism is disabled.

$SystemLogRateLimitInterval <rate limit interval in seconds>

$SystemLogRateLimitBurst <maximum messages in interval>

$IMUXSockRateLimitBurst <maximum messages in interval>

$IMUXSockRateLimitInterval <rate limit interval in seconds>

For example, defining a rate limit of 200 messages per 5 seconds is configured as follows:

$SystemLogRateLimitInterval 5

$SystemLogRateLimitBurst 200

$IMUXSockRateLimitBurst 200

$IMUXSockRateLimitInterval 5

3. On one of the Installer machines, run the following command to apply the changes:

salt '*' state.apply rsyslog

7. Installation and Upgrade

7.1 Unexpected Reboots during Upgrade Cause Data Loss

7.1.1 Error Description

Data loss occurs when the SDC unexpectedly reboots while an SDC upgrade is in process and the logical drives are still split.

7.1.2 Causes

When a reboot occurs while an SDC upgrade is in progress, the split mirror procedure is still in effect. The backed up logical drive is used to install the system, instead of the upgraded logical drive.

7.1.3 Symptoms

Data collected and configured since the split mirror procedure was performed is lost.

7.1.4 Resolution

When encountering a reboot during an SDC upgrade, ensure that the correct logical drive is used to install the system:

1. Run the following command to identify the active logical drive:

   hpssacli controller slot=0 ld all show

2. Run the following command to mark the active logical drive, based on the output from the previous command:

   hpssacli controller slot=0 ld <active Id> show

3. Run the following command to define which disk should be used to boot from:

   hpssacli controller slot=1 ld <active Id> modify bootvolume=primary

4. Run the following command to verify that the correct disk is defined to boot from:

   hpssacli controller slot=0 show |grep Boot

8. Changing a FEP's IP Address

You can change an IP address of a FEP component that is defined in the Topology file.

Note: FEPs only use signaling networks. Make sure that you are changing a FEP's IP address and not an IP address that belongs to another SDC component. Verify that the IP address that you are changing to is allowed by your firewall system. This procedure is applicable to IPv4 and IPv6 protocols.

It is recommended to perform this procedure during a maintenance window.

8.1 Removing the Current Virtual Server

Prior to changing the IP address of a FEP component, you need to remove the current virtual server that listens to the FEP component that you are changing its IP address. After changing the IP address, you then need to create a new virtual server so that the existing port number will be available and is recognized with the new IP address.

To remove the current virtual server:

1. Go to **Topology** > **Specific Site Settings** > **Virtual Servers**.
2. Select the relevant virtual server that listens to the FEP component that you are changing its IP address, and click **Remove**.

8.2 Changing a FEP IP Address

To change a FEP's IP address:

1. Log in to Cassandra on one of master Installer servers with the following command:

   ```
   /opt/cassandra/bin/cqlsh <mgmt ip>
   ```

2. Run the following command:

   ```
   SELECT * from topology.interface;
   ```

   The following is an example of the displayed output:

Figure 7: Example of Topology Interface Output

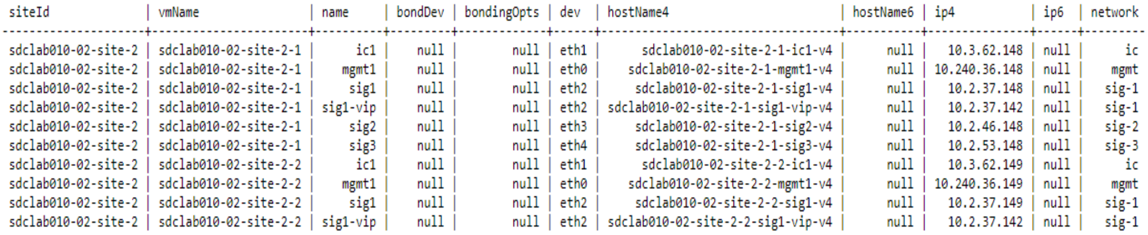

```
 siteId              | vmName                | name     | bondDev | bondingOpts | dev  | hostName4                         | hostName6 | ip4           | ip6  | network
---------------------+-----------------------+----------+---------+-------------+------+-----------------------------------+-----------+---------------+------+---------
 sdclab010-02-site-2 | sdclab010-02-site-2-1 |      ic1 |    null |        null | eth1 |      sdclab010-02-site-2-1-ic1-v4 |      null |   10.3.62.148 | null |      ic
 sdclab010-02-site-2 | sdclab010-02-site-2-1 |    mgmt1 |    null |        null | eth0 |    sdclab010-02-site-2-1-mgmt1-v4 |      null | 10.240.36.148 | null |    mgmt
 sdclab010-02-site-2 | sdclab010-02-site-2-1 |     sig1 |    null |        null | eth2 |     sdclab010-02-site-2-1-sig1-v4 |      null |   10.2.37.148 | null |   sig-1
 sdclab010-02-site-2 | sdclab010-02-site-2-1 | sig1-vip |    null |        null | eth2 | sdclab010-02-site-2-1-sig1-vip-v4 |      null |   10.2.37.142 | null |   sig-1
 sdclab010-02-site-2 | sdclab010-02-site-2-1 |     sig2 |    null |        null | eth3 |     sdclab010-02-site-2-1-sig2-v4 |      null |   10.2.46.148 | null |   sig-2
 sdclab010-02-site-2 | sdclab010-02-site-2-1 |     sig3 |    null |        null | eth4 |     sdclab010-02-site-2-1-sig3-v4 |      null |   10.2.53.148 | null |   sig-3
 sdclab010-02-site-2 | sdclab010-02-site-2-2 |      ic1 |    null |        null | eth1 |      sdclab010-02-site-2-2-ic1-v4 |      null |   10.3.62.149 | null |      ic
 sdclab010-02-site-2 | sdclab010-02-site-2-2 |    mgmt1 |    null |        null | eth0 |    sdclab010-02-site-2-2-mgmt1-v4 |      null | 10.240.36.149 | null |    mgmt
 sdclab010-02-site-2 | sdclab010-02-site-2-2 |     sig1 |    null |        null | eth2 |     sdclab010-02-site-2-2-sig1-v4 |      null |   10.2.37.149 | null |   sig-1
 sdclab010-02-site-2 | sdclab010-02-site-2-2 | sig1-vip |    null |        null | eth2 | sdclab010-02-site-2-2-sig1-vip-v4 |      null |   10.2.37.142 | null |   sig-1
```

3. From the displayed table, identify the line with the FEP IP address you want to change and run the following command with the values from the table, for example, as follows:

   UPDATE topology.interface SET ip4 = '<new IP address value>' WHERE "siteId" = 'sdclab010-02-site-2' AND "vmName" = 'sdclab010-02-site-2-1' AND name = 'sig1-vip'

Note: If the FEP IP address is a VIP and running on more than one server, then you need to run the UPDATE command for each relevant line.

4. Verify that the selected IP address was changed by running the following command:

   SELECT * from topology.interface;

   Per the example above, for the selected site, the IP address will show as 10.2.37.142 instead of the previous value.

5. Run the following commands on the two master Installer servers to rebuild the traffix pillar:

   cd /srv/traffix/

   echo 0 > .lastRev

   /etc/init.d/salt-master restart

   salt '*' saltutil.refresh_pillar

6. Clean-up the old configuration on any server that hosted the changed IP address:

   a. Back up to another location all files under */etc/sysconfig/network*-scripts that start with ifcfg-eth, ifcfg-bond and route.

   b. Remove the following files with the following command

   - **rm -rf /etc/sysconfig/network-scripts/ifcfg-eth***
   - **rm -rf /etc/sysconfig/network-scripts/ifcfg-bond***
   - **rm -rf /etc/sysconfig/network-scripts/route***

7. Go to one of master Installer servers and run:

   salt "*" state.highstate

   The networking files are now restarted.

8. Verify that all bonds and interfaces are located under: */etc/sysconfig/network-scripts/*

9. Reboot any server that hosted the changed IP address.

10. Generate a new license key for the new IP address:

    In the Web UI, go to **Administration** > **Specific Site Settings** > **License**

    For more information about generating a new license key, see the *F5 SDC 5.1 User Guide*.

11. Create a new virtual server (go to **Topology** > **Specific Site Settings** > **Virtual Servers >Add**) that will be the listener to the component with the changed IP address.

    For more information about creating a new virtual server, see the *F5 SDC 5.1 User Guide*.

9. Changing a Site's VLAN ID

You can change the defined VLAN ID for an SDC or EMS site. The Vlan ID is defined in the Topology File (networks) and it defines the VLAN-aware partitioning of the network.

To change a VLAN ID for a site:

1. Log in to Cassandra on one of master Installer servers with the following command:

   /opt/cassandra/bin/cqlsh <mgmt ip>

2. Run the following command:

   SELECT * from topology.network;

   The following is an example of the displayed output:

Figure 8: Example of Topology Network Output

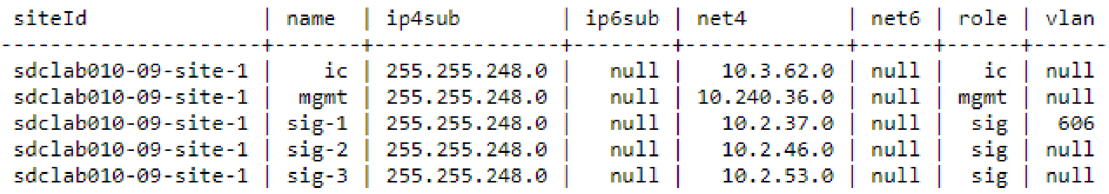

```
 siteId                | name  | ip4sub        | ip6sub | net4        | net6 | role | vlan
-----------------------+-------+---------------+--------+-------------+------+------+------
 sdclab010-09-site-1 |    ic | 255.255.248.0 |   null |   10.3.62.0 | null |   ic | null
 sdclab010-09-site-1 |  mgmt | 255.255.248.0 |   null | 10.240.36.0 | null | mgmt | null
 sdclab010-09-site-1 | sig-1 | 255.255.248.0 |   null |   10.2.37.0 | null |  sig |  606
 sdclab010-09-site-1 | sig-2 | 255.255.248.0 |   null |   10.2.46.0 | null |  sig | null
 sdclab010-09-site-1 | sig-3 | 255.255.248.0 |   null |   10.2.53.0 | null |  sig | null
```

3. From the displayed table, identify the line with the Vlan ID you want to change and run the following command with the values from the table, for example, as follows:

   UPDATE topology.network SET "vlan"=111 WHERE "siteId"='sdclab010-09-site-1' AND name ='sig-1' '

4. Verify that the selected Vlan ID was changed by running the following command:

   SELECT * from topology.network;

   Per the example above, for the selected site, the vlan value will show as 111 instead of 606.

5. Run the following commands on the two master Installer servers to rebuild the traffix pillar:

```
cd /srv/traffix/

echo 0 > .lastRev

/etc/init.d/salt-master restart

salt '*' saltutil.refresh_pillar
```

6. Clean-up the old configuration on any server that hosted the edited VLAN ID:

   a. Back up to another location all files under */etc/sysconfig/network*-scripts that start with ifcfg-eth, ifcfg-bond and route.

   b. Remove the following files with the following command

      - **rm -rf /etc/sysconfig/network-scripts/ifcfg-eth***
      - **rm -rf /etc/sysconfig/network-scripts/ifcfg-bond***
      - **rm -rf /etc/sysconfig/network-scripts/route***

7. Go to one of master Installer servers and run:

   ```
   salt "*" state.highstate
   ```

   The networking files are now restarted.

8. Verify that all bonds and interfaces are located under: */etc/sysconfig/network-scripts/*

9. Reboot any server that hosted the edited VLAN ID.

10. Cassandra Database

10.1 Resynching the Database

The Cassandra database is hosted on the master Installers and the OAMs. If one of the servers that host a Cassandra database is down for more than three hours, perform the following procedure to make sure the data is consistent with the other data saved on all other Cassandras.

Note: This procedure is relevant for both Bare Metal and Virtual environments.

The procedure includes removing the downed host server from the cluster of Cassandra servers, cleaning the existing data from the removed Cassandra server, uploading the cleaned server, followed by repairing the Cassandra server cluster so that all Cassandras can recognize the newly uploaded clean Cassandra server to synch all data.

To remove the downed Cassandra server:

1. Run the following command, with the name of the downed server, so that monit knows that the server is no longer active:

   # monit unmonitor <server host name>

2. Run the following command to identify the downed host server within the cluster and to get the matching Host ID:

   # /opt/cassandra/bin/nodetool status

For example:

```
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address         Load       Tokens       Owns    Host ID
Rack
UN  10.240.36.226  872.95 MB  256          ?       8fedbfab-9cf9-47d0-
83e7-67d2349fd1da  RAC1
```

```
UN  10.240.36.227  791.91 MB  256          ?       9fd5b181-201b-4d19-
b773-bf7bc75f9dc6  RAC1
DN  10.240.36.228  872.99 MB  256          ?       152e3202-1012-42f3-
bd0a-2e15644cf1c2  RAC1
UN  10.240.36.229  718.61 MB  256          ?       f6512e76-14e7-45b1-
ba07-c1e6b80c111e  RAC1
Datacenter: site-2
```

3. Run the following command with the identified host ID:

   # /opt/cassandra/bin/nodetool removenode <Host ID>

4. Run the status again to make sure that server is removed:

   # /opt/cassandra/bin/nodetool status

To clean any existing data from the removed host:

1. Run the following commands from the removed host:

   # cd /data/cassandra

   # pwd

   /data/cassandra

   # rm -rf data saved_caches commitlog

   Now, the Cassandra server is empty.

To reactivate the monitoring functionality of the downed server:

1. Run the following command:

   # monit monitor <server host name>

2. Check the status to make sure is up and running:

   # /opt/cassandra/bin/nodetool status

The downed Cassandra server, (in the example 10.240.36.228) is now up and running with a new Host ID.

For example:

```
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address         Load        Tokens       Owns    Host ID
Rack
UN  10.240.36.226  1020.39 MB  256           ?       8fedbfab-9cf9-47d0-
83e7-67d2349fd1da  RAC1
UN  10.240.36.227  881.95 MB  256          ?       9fd5b181-201b-4d19-
b773-bf7bc75f9dc6  RAC1
UN  10.240.36.228  97.14 KB   256          ?       aaffd25f-bc3a-4e13-
ac9a-7b423e746220  RAC1
UN  10.240.36.229  870.15 MB  256          ?       f6512e76-14e7-45b1-
ba07-c1e6b80c111e  RAC1
```

To reconnect the previously downed server with the Cassandra cluster:

1. Run the following command on the site where the server was down

   # /opt/cassandra/bin/nodetool repair -full

10.2 Reducing Cassandra Disk Usage

10.2.1 Error Description

As a result of running the Cassandra keyspace repair command, the Cassandra system_distributed keyspace tables fill-up with data, requiring high disk usage.

10.2.1 Resolution

To free up the disk usage, you can truncate the Cassandra system_distributed repair_history and parent_repair_history_ tables. In addition, you also need to delete the repair_history and parent_repair_history snapshots.

To truncate the system_distributed keyspace tables.

1. Verify that the repair procedure (the Cassandra keyspace repair command) is not currently running on any of the Cassandra hosted servers.

2. Run the following commands on one of the Cassandra hosted servers to delete the data from the repair_history and parent_repair_history tables:

```
cqlsh> TRUNCATE TABLE system_distributed.repair_history;
cqlsh> TRUNCATE TABLE system_distributed.parent_repair_history;
```

   After running these commands, snapshots of these tables are created.

3. Display all the snapshots from each Cassandra hosted server:

   a. Log on to each Cassandra hosted server

   b. Run the following command to display a list of all snapshots:

```
# /opt/cassandra/bin/nodetool listsnapshots
```

   All snapshot details are displayed.

   c. Repeat this command for each Cassandra hosted server.

   Example of Snapshot details:

Snapshot Name	Keyspace Name	Column Family Name	True Size (MB)	Size on Disk (MB)
1512998228632-repair_history	system_distributed	repair_history	107.54	107.54
1512998238390-parent_repair_history	system_distributed	parent_repair_history	1.1	1.1

   In this example, the total TrueDiskSpaceUsed: 108.64 MB

4. Remove the repair_history and parent_repair_history snapshots from the system_distributed keyspace table for each Cassandra hosted server.

a. Run the following command:

```
# /opt/cassandra/bin/nodetool clearsnapshot system_distributed
```

b. Repeat this command for each Cassandra hosted server.

5. Verify that the system_distributed repair_history and parent_repair_history snapshots have been removed:

```
# /opt/cassandra/bin/nodetool listsnapshots
```

11. Overload Control

11.1 Receive/Send Rate Limit is Half Than Expected

11.1.1 Error Description

Though you have configured a global message rate limit (Transaction receiving rate limit) or specific peer/profile message rate limits (Message sending rate limit), the TPS data graphs only show about half of the configured amount.

11.1.2 Causes

The discrepancy is because each counted message might be a request or a response, while each transaction of TPS is both a request and response.

11.1.3 Resolution

Refer to the Rate Limit table in the F5 SDC User Guide.

11.2 Web UI Statistics Memory Usage Increase

11.2.1 Symptoms

The Web UI statistics show increased memory usage.

11.2.2 Cause

Sessions are accumulating in the CPF memory due to session timeout being too long.

11.2.3 Resolutions

- Set the session timeout parameter to a lower value.
- Update the transformation scripts to release the sessions after receiving the last message of a session.

Note: This may not be possible if SDC is unable to identify the last session message.

12. Monitoring

12.1 Session Repository Logs

The SDC and Session Repository can be configured to generate logs for monitoring session life cycle events, session replication, and session errors. These logs can be used to help troubleshoot when stateful sessions fail to route or replicate. You can customize a session log by adding session attributes to a session log. For session life cycle events, logs are found in the CPF session logs and for session replication, logs are found in the Session Repository session logs.

The Session Repository (Tripo) maintains the following log files which are stored locally:

- */home/traffix/Tripo/log/session_output*
- */home/traffix/Tripo/log/session_error*
- */home/traffix/Tripo/log/session_mate_rep*
- */home/traffix/Tripo//log/session_srr*

12.1.1 Configuring the Log Level upon Session Repository (Tripo) Initialization

You can configure the Tripo log level for when the Session Repository (Tripo) is starting -up. If the log level is not configured, all generated log levels for Session Repository (Tripo) initialization are set to the default level, Alert.

To configure the initialization log level:

1. Add the following configuration to the SLF_SysConf.scf file, section id="LogSrv":

```
<SLF_INIT_LOG_LEVEL_HELPER>7</SLF_INIT_LOG_LEVEL_HELPER>
```

   Configuring this line with a "7" means that all the logs (Debug (7), Info (6), Notice (5), Warning (4), Error (3), Critical (2), and Alert (1)) can be generated. For other log levels, enter the relevant number.

2. Execute the following command to start the Session Repository:

```
>start.sh ;
```

The system is started.

3. Execute the following command to view the log level:

```
>UI_LogCfg –p
```

The logs are displayed with the different configured log levels.

12.2 Reporting

ELK is software that gathers, indexes, and arranges data from any application, server, or network device in your IT infrastructure. This data can then be generated into analytical reports with tables, charts, and graphs that are displayed in a Web UI.

12.3 Configuring Log Files

Note: It is recommended to use the pre-configured SDC log configurations. If you choose not to, then follow the procedures below and consult with *F5 Support*. The log levels are configured in the SDC Web UI as documented in the F5 SDC User Guide.

The SDC uses Log4j and rsyslogs compile logs for troubleshooting. Log4j collects the data and sends it to rsyslog. Once logs are sent to rsyslog, they are deleted from log4j.

In addition to customizing log parameters, such as included information, log name, log size, frequency of log generation, you can also configure the rotation mechanism for archiving logs.

12.3.1 Configuring the Log4J

Log4j is a flexible logging API framework written in Java, which is distributed under the Apache Software License. It views the logging process in terms of levels of priorities and offers mechanisms to direct logging information to a great variety of destinations, such as Linux RedHat Syslog (rsyslog).

For SDC installations, the Log4.xml is installed with the SDC*rpm package and is located in the following folder:

/opt/traffix/sdc5.0_<version>/config/log4j.xml

Log4j has three main components:

- Loggers: Responsible for capturing logging information.
- Appenders: Responsible for publishing logging information to various preferred destinations. The syslog appender is used to define the syslog properties.
- Layouts: Responsible for formatting logging information in different styles.

Table 2: Loggers Default Values

Logging Type	Log Level	Related Appender
com.traffix.mgmtconsole.server.logging	DEBUG	userActivityLogAppender
ElkTracedSessionsLogger	ALL	syslogAppender
ElkTracedSessionsLogger	ALL	syslogAppender
com.traffix.ElkTrafficTracingLogger	ALL	syslogElkAppender
com.traffix.ElkRoamingSessionsLogger	ALL	syslogRoamingAppender
userTraceLogger	ALL	syslogAppender
LoggerPerformanceListener	INFO	statisticsAppender
com.traffix.LoggerKpis	ERROR	kpisAppender
RawStatistics	ALL	rawStatisticsAppender
com.traffix.proxy	INFO	syslogAppender
com.traffix.openblox.nms	INFO	syslogAppender
com.traffix	ALL	syslogAppender
com.traffix.openblox.core.utils.license.LicensePerformanceListener	ERROR	licenseAppender

Logging Type	Log Level	Related Appender
com.traffix.openblox.core.utils.clustering.jmx	INFO	syslogAppender
org.apache.directory.shared.ldap.entry.String Value	ERROR	syslogAppender
com.traffix.openblox.core.storage.SessionLog gingUtil	INFO	sessionManagementLogAppender sessionErrorLogAppender

Table 3: Appender Default Values

Appender Type	File Type	Parameter	Default Value
mainLogAppender	<componentshortname>.log	Threshold	ALL
mainLogAppender	<componentshortname>.log	MaxBackupIndex	1
mainLogAppender	<componentshortname>.log	MaxFileSize	10MB
licenseAppender	license.dat	Threshold	ALL
licenseAppender	license.dat	MaxBackupIndex	1
licenseAppender	license.dat	MaxFileSize	10MB
sessionManagementLogAppender	session_output.log	Threshold	INFO
sessionManagementLogAppender	session_output.log	MaxBackupIndex	1
sessionManagementLogAppender	session_output.log	MaxFileSize	10MB
sessionErrorLogAppender	session_error.log	Threshold	ERROR
sessionErrorLogAppender	session_error.log	MaxBackupIndex	1
sessionErrorLogAppender	session_error.log	MaxFileSize	10MB
consoleCriticalAppender	System.err	Threshold	ERROR
kpisAppender	kpis.log	Threshold	ALL

Appender Type	File Type	Parameter	Default Value
kpisAppender	kpis.log	MaxBackupIndex	10
kpisAppender	kpis.log	MaxFileSize	50MB
rawStatisticsAppender	rawStatistics.csv	Threshold	TRACE
syslogElkAppender	Elk.log	Threshold	INFO
syslogRoamingAppender	Elk.log	Threshold	INFO
userActivityLogAppender	userActivityLogFile.log	Threshold	ALL

Table 4: Layout Default Values

Layout Category	Description
SimpleLayout	The level of the log statement, followed by " - " and then the log message itself. For Example, DEBOG – Check the configuration
HTMLLayout	Outputs events in an HTML table. Note: Appenders using this layout should have their encoding set to UTF-8 or UTF-16, otherwise events containing non-ASCII characters could result in corrupted log files.
PatternLayout	Generates logging information in a particular format based on a pattern. Note: As this layout option is known to have synchronization and other issues, it is recommended to use the SDC customized layout: CpfPatternLayout.

12.3.2 Configuring rsyslog

While the log name is initially assigned by Log4j, when the log is saved as rsyslog, the log name can be changed by being renamed or by adding a prefix or suffix.

In addition, you can configure the frequency in which rsyslogs are sent to the NMS server.

12.3.2.1 rsyslog Structure

The rsyslog is installed as a rsyslog.rpm which is delivered as part of the Red Hat (RHEL) Operating System.

The rsyslog files are managed by the Salt Master and for each component there is a corresponding template, as described in Table 5. The parent Salt Master server location for all rsyslogs is *salt/rsyslog.sls*. The rsyslog process is initiated by applying the *salt/rsyslog-init.conf* script on all the servers in the site.

Table 5: salt/rsyslog.sls Template Folders

Salt Server Location	Related SDC Component	Path	Related Managed Log Files
salt/rsyslog/rsyslog-server.conf	OAM Servers (1 and 2)	/etc/rsyslog.conf	Standard default OS local: logging /var/log/maillog /var/log/cron /var/log/spooler /var/log/boot /var/log/secure
salt/rsyslog/rsyslog-client-tmpl.conf	All Client Server (except Tripo)	/etc/rsyslog.conf	Standard default OS local: logging /var/log/maillog /var/log/cron /var/log/spooler /var/log/boot /var/log/secure
salt/rsyslog/rsyslog-client-tripo.conf	Tripo	/etc/rsyslog.conf	Standard default OS local: logging /var/log/maillog /var/log/cron

Salt Server Location	Related SDC Component	Path	Related Managed Log Files
			/var/log/spooler /var/log/boot /var/log/secure
salt/rsyslog/rsyslog.template.cm	Configuration Manager	/etc/rsyslog.d/	▪ /var/log/traffix_config_mgr-<nodeId>-stdout.log ▪ /opt/traffix/components/cm/logs//config_mgr/config_mgr.log ▪ /opt/traffix/components/cm/logs/gc/gc_<hostname>-<nodeid>.log.0.current ▪ /opt/traffix/components/cm/logs/gc/gc_<hostname>-<nodeid>.log.1.current
salt/rsyslog/rsyslog.template.cpf	CPF	/etc/rsyslog.d/	▪ /opt/traffix/components/cpf/logs/<nodeid>//session_output/session_output.log ▪ /opt/traffix/components/cpf/logs/<nodeid>//session_error/session_error.log ▪ /opt/traffix/components/cpf/logs/<nodeid>/cpf/statistics/CpfRawStatistic/rawStatistics.csv ▪ /opt/traffix/components/cpf/logs/<nodeid>/cpf/statistics/CpfRawStatistic/statistics.log ▪ /var/log/traffix_cpf-<nodeId>-stdout.log

Salt Server Location	Related SDC Component	Path	Related Managed Log Files
			▪ /opt/traffix/components/cpf/logs/<nodeid>/cpf/cpf.log ▪ /opt/traffix/components/cpf/logs/<nodeid>/cpf/license.dat ▪ /opt/traffix/components/cpf/logs/gc/gc_<hostname>-<nodeid>.log.0.current ▪ /opt/traffix/components/cpf/logs/gc/gc_<hostname>-<nodeid>.log.1.current
salt/rsyslog/rsyslog.template.fep	FEP	/etc/rsyslog.d/	▪ /opt/traffix/components/fep/logs/<nodeId>/fep/fep.log ▪ /opt/traffix/components/fep/logs/<nodeId>/session_output/session_output.log ▪ /opt/traffix/components/fep/logs/<nodeId>//session_error/session_error.log ▪ /var/log/traffix_fep-<nodeId>-stdout.log ▪ /opt/traffix/components/fep/logs/<nodeId>/fep/statistics/CpfRawStatistic/rawStatistics.csv ▪ /opt/traffix/components/fep/logs/<nodeId>/fep/statistics/CpfRawStatistic/statistics.log ▪ /opt/traffix/components/fep/logs/gc/gc_<hostname>-<nodeid>.log.0.current

Salt Server Location	Related SDC Component	Path	Related Managed Log Files
			▪ /opt/traffix/components/fep/logs/gc/gc_<hostname>-<nodeid>.log.1.current
salt/rsyslog/rsyslog.template.nmsagent	NMS Agent	/etc/rsyslog.d/	▪ /var/log/traffix_nmsagent-<nodeId>-stdout.log ▪ /opt/traffix/components/nmsagent/logs/nmsagent/nmsagent.log ▪ /opt/traffix/components/nmsagent /logs/gc/gc_<hostname>-<nodeId>.log.0.current ▪ /opt/traffix/components/nmsagent//logs/gc/gc_{{myHost}}-{{monitVars.nodeId}}.log.1.current
salt/rsyslog/rsyslog.template.oamDB	OAM	etc/rsyslog.d/	▪ /var/log/cassandra/logs/debug.log ▪ /var/log/cassandra/logs/system.log ▪ /var/log/traffix_cassandra-{{monitVars.nodeId}}-stdout.log /var/log/cassandra-repair.log
salt/rsyslog/rsyslog.template.tripo	Tripo	etc/rsyslog.d/	/home/traffix/Tripo/log/traffix_tripo-default-stdout.log
salt/rsyslog/rsyslog.template.webui	Web UI	etc/rsyslog.d/	▪ /var/log/traffix_webui-<nodeId>-stdout.log

Salt Server Location	Related SDC Component	Path	Related Managed Log Files
			▪ /opt/traffix/components/webui/logs/webui/webui.log ▪ /opt/traffix/components/webui//logs/webui/userActivityLogFile.log
salt/rsyslog/rsyslog.template.elasticsearch-master	ELK Master ELK Data	etc/rsyslog.d/	▪ /var/log/rsyslog/<vm>.elasticsearch_master.elasticsearch-plain.log ▪ /var/log/rsyslog/<vm>.elasticsearch_data.elasticsearch-plain.log

12.3.2.2 Prerequisites for Configuring the rsyslogs

The following are the prerequisites to configure the rsyslogs:

- rsyslog owner must have permission to the */var/log/rsyslog* folder
- rsyslog owner must have permission to the */var/spool/rsyslog* folder
- rsyslog supported version is: rsyslog 5
- Configured port is TCP 10514

The following is an example of an open 10514 port on an installed system:

```
netstat -anp | grep 10514
tcp        0      0 10.1.67.100:56035           10.1.67.102:10514
ESTABLISHED 6719/rsyslogds
tcp        0      0 10.1.67.100:57662           10.1.67.107:10514
ESTABLISHED 6719/rsyslogd
```

12.3.2.3 Configuring a Specific rsyslog File

You can configure a specific log file (from the list of Related Managed Log Files in *Table 5)*.

To configure parameters in an rsyslog file:

1. Connect to a Salt Master server, with the following command:

   cd /srv/salt<version>/rsyslog/

2. This command points to the parent syslog folder. From here, locate the folder that contains the relevant log file.

3. Edit the relevant parameters:

   - $InputFileName - The file that is being monitored. This must be an absolute name (no macros or templates).
   - $InputFileTag – The tag to be used for messages that originate from this file.
   - $InputFileStateFile

   rsyslog keeps track of which parts of the monitored file it already processed. This is done in the state file. This file is always created in the rsyslog working directory (configurable via $WorkDirectory).

Note: Use unique names for each of the different files being monitored, as rsyslog does not check if the same name is specified multiple times.

   - $InputFileFacility

   The rsyslog facility is the way log file lines are read. This is specified in textual form (for example, “local0”, “local1”, ...) or as numbers (for example, 128 for “local0”). The textual form is recommended. The default is “local0”.

   - $InputFileSeverity

The rsyslog severity that is assigned to a specific facility (log file line format). This is specified in textual form (for example, “info”, “warning”, ...) or as numbers (for example, 4 for “info”). The textual form is recommended. The default is “notice”.

- $InputFilePersistStateInterval

Specifies how often the state file is written when processing the input file.

Note: The configured SDC default value is 20,000. This means that after 20,000 log lines, the state file is written and the monitored file is closed (end of rsyslogd execution). Frequent writing of the state file is very time consuming and, therefore, this setting can affect imfile performance, especially, when it is set to a low value. This setting can be used to prevent message duplication following a fatal error (such as, a shutdown or power failure), as the rsyslog does not restart reading the log lines from the beginning, but rather from the point when the fatal error occurred.

- $I nputRunFileMonitor

This enables the file monitoring. Without this setting, the text file will be ignored.

- $InputFileBindRuleset

This binds the listener to a specific rule set. The SDC default value is: RSYSLOG_DefaultRuleset

4. To save the changed configuration, run the following command from one of the Salt Master Servers:

```
salt "*" state.apply rsyslog
```

12.3.3 Configuring the logrotate Mechanism

The logrotate mechanism is installed as a logrotate.rpm. which is delivered as part of the Red Hat (RHEL) Operating System.

The log rotation mechanism, logrotate, is a way to handle a large number of logs created by Log4j and that are sent to rsyslog. logrotate allows automatic rotation, compression, archiving, removal, and mailing of log files. Each log file can be handled on a daily, weekly, or monthly basis. Once, a log file is larger than the defined log size, it is archived. You can configure the frequency of when the log size is checked, the allowable log size and the rotation of the archived logs.

The logrotate package is run by cron that automatically rotates log files as defined in the main configuration file: */etc/logrotate.conf*. The configuration files are written in the */etc/logrotate.d/directory*.

12.3.3.1 Configuring the Log Rotation Frequency

You can set when the script (see *Table 6*) will be run to check the log size. Any changes made, are applied to all servers.

To change the log rotation frequency:

1. Connect to a Salt Master server, with the following command:

   cd /srv/salt<version>/system

   This opens the salt/system/crontab.sls file that includes the defined log rotation frequency parameter: weekly, daily, or hourly:

Table 6: Log Rotation Frequency Scripts

Rotation Frequency	Script to be Run	Related Log File
Daily	▪ system/crontab/logrotate-traffix ▪ system/crontab/logrotate-salt	/etc/logrotate.d/logrotate-traffix.conf
Weekly	▪ system/crontab/logrotrafx ▪ system/crontab/logrotsalt	▪ /etc/logrotate.d/logrotate-traffix.conf
Hourly	▪ system/crontab/coredumps_cleanup.py	▪ /etc/logrotate.d/logrotate-traffix-hourly.conf

2. Select the relevant log files, search for the existing time frequency parameter that you want to change, and then delete it.
3. Open the other relevant folder that you want to change the frequency to and add in the new frequency parameter.

12.3.3.2 Configuring the Log Size and Rotation Schedule of Archived Logs

For each log folder, there is a default defined size and number of allowed archived rotations.

Note: Once the number of logs that are rotated exceed the defined number of rotations, no more logs are archived until that rotation is renewed.

To change the allowable log size and rotation index:

1. Connect to a Salt Master server, with the following command:

   cd /srv/salt<version>/system

   The *salt/system/logrotate.sls* file opens.

2. Open the *logrotate-traffix.conf*

   Note: To open all files under */var/log* with a “traffix” prefix and a “log” suffix, type /var/log/traffix*.log

3. Locate the log folder in which you want to change the default log size or rotation schedule and change the relevant parameters as in *Table 7.*
4. To save the changed configuration, run the following command from one of the Salt Master Servers:

   salt "*" state.apply system

Table 7: Default Log Rotation Values

SDC Log Type	Related Log File	Default Size (Megabytes)	Default Rotations Saved as an Archived log file
SDC Main logs	/var/log/rsyslog/*.main.log	10	10/Weekly
webui useractivity	/var/log/rsyslog/*.useractivity.log	10	10/Weekly
Tripo	/var/log/rsyslog/*.Storage.log	10	5/Weekly
Sessions	var/log/rsyslog/*.session*.log	100	20/Weekly
Stats	/var/log/rsyslog/*.statistics.log	50	10/Weekly
Stdout	/var/log/rsyslog/*.stdout.log	10	7/Daily
Messages	/var/log/rsyslog/*.messages	100	7/Daily
oamDB debug and system	/var/log/rsyslog/*.System.log	100	10/Daily
oamDB debug and system	/var/log/rsyslog/*.Debug.log	100	10/Daily
oamDB debug and system	/var/log/rsyslog/*.Stdout.log	100	10/Daily
Cron	/var/log/rsyslog/*.cron	100	7/Daily
Secure	/var/log/rsyslog/*.secure	100	7/Daily
CPF license	/var/log/rsyslog/*.license.log	10	10/Weekly
Tripo Comsrv	/var/log/rsyslog/*.ComSrv.log	10	5/Weekly
Tripo LogSrv	/var/log/rsyslog/*.LogSrv.log	10	5/Weekly
Tripo RepMgr	/var/log/rsyslog/*.RepMgr.log	10	5/Weekly

SDC Log Type	Related Log File	Default Size (Megabytes)	Default Rotations Saved as an Archived log file
Tripo SysWatchDog	/var/log/rsyslog/*.SysWatchDog.log	10	5/Weekly
Tripo UI_Config	/var/log/rsyslog/*.UI_Config.log	10	5/Weekly
Tripo UI_LogCfg	/var/log/rsyslog/*.UI_LogCfg.log	10	5/Weekly
Tripo CFGClient	/var/log/rsyslog/*.CFGClient.log	10	5/Weekly
Tripo WebStatMgr	/var/log/rsyslog/*.WebStatMgr.log	10	5/Weekly
Tripo default	var/log/rsyslog/*.default.log	10	5/Weekly
Keepalived default	/var/log/rsyslog/*.Keepalived_vrrp.log	10	5/Weekly
Elk	salt/rsyslog/rsyslog.template.elasticsearch-master.conf salt/rsyslog/rsyslog.template.elasticsearch-data.conf salt/rsyslog/rsyslog.template.fluent salt/rsyslog/rsyslog.template.kibana	10	10/Weekly
GC	/var/log/rsyslog/*.gc.log	16	10/Weekly
SNMP	/var/log/snmptrapd.log	100	10/Weekly
cassandra-repair logs	/var/log/cassandra-repair.log	10	7/Weekly
cassandra-repair logs	/var/log/rsyslog/*cassandra-repair*	10	7/Weekly

12.4 Exporting TDR Reports

In addition to TDR reports being displayed in the EMS Web UI (**Reports >Transaction Data Records**), TDR reports are also exported as tdr_export_<date>_<time>.csv.gz files to a defined location. For sites that are already updated with ELK, TDR reports are sent to the */opt/traffix/reports/elk* folder of the EMS site, and for those sites that have not yet been updated (going through the update Splunk-ELK process or running in mix-mode), TDR reports are sent to the */opt/traffix/reports/tdr* folder of the EMS site. The default setting is that the three most recently generated TDR .csv files are exported every 30 minutes.

Note: To generate TDR reports, the **Create Transaction Data Record** checkbox must be selected for the relevant routing rule (**Flow Management > Routing Rules > TDR Configuration**). For more information, refer to the *F5 SDC User Guide*.

12.4.1 Customizing the Default TDR Export Settings

You can set how often reports should be generated, how many reports should be saved, the file name format, and the export location of the reports.

To customize the default TDR export settings:

1. Go to *cd /srv/salt/<version>/elk/global/* and open the traffix_export_tdrs.sh. file.
2. Edit each variable with a custom value:

To change…:	Configure:

export file path	EXPORT_DIR="/opt/traffix/reports/elk"
	EXPORT_FILE_PREFIX="tdr_export"

name of the TDR report	EXPORT_FILE_NAME="${EXPORT_FILE_PREFIX}_${DATE}.csv"
How many files to keep	EXPORT_FILES_TO_KEEP
enable/disable compression	COMPRESS=true/false

how often TDR reports are generated	/srv/salt/<version>/elk/start.sls traffix-export-tdr-cron state

12.5 Configuring the Number of Displayed Alarms

While the maximum number of alarms that are displayed in the Web UI (Alarms > Alarm History Logs) is pre-defined, you have the option to change it. The default limit of displayed alarms is 2,000 for an EMS site and 200 for an SDC site.

To configure the maximum number of displayed alarms:

1. Go to the following file in the relevant EMS or SDC site that you want to change the default setting:

   /srv/salt/5.1<build number>/nms/traffix_nmsagent_init

2. Add the following parameter (where the -D indicates it is a parameter in the Java command line):

   ```
   -DAlarmsHistoryLimit=<user defined limit>\
   ```

   The following is an example of an added parameter in a command line:

   ```
   default_value   JAVA_INSTANCE_OPTS "\
       -DAlarmsHistoryLimit=1000
       -DcacheDnsForever=${CACHE_DNS_FOREVER} \
   ```

Note: Extending the display limit of alarms may impact upon the availability of Web UI resources and functionality.

Glossary

The following table lists the terms and abbreviations used in this document.

Table 8: Common Terms

Term	Definition
Answer	A message sent from one Client/Server Peer to the other following a request message
Client Peer	A physical or virtual addressable entity which consumes AAA services
Data Dictionary	Defines the format of a protocol's message and its validation parameters: structure, number of fields, data format, etc.
Destination Peer	The Client/Server peer to which the message is sent
Geo Redundancy	A mode of operation in which more than one geographical location is used in case one site fails
Master Session	The session for which the routing selection is performed based on the routing rules (Slave Sessions are applied with routing rules inherited from the Master Session)
Orchestrator	A workflow management solution to automate the creation, monitoring, and deployment of resources in your environment
Origin Peer	The peer from which the message is received
Pool	A group of Server Peers
QCOW2	A file format for disk image files
RADIUS	Remote Authentication Dial In User Service
REST	Representation of a resource between a client and server **(Re**presentational **S**tate **T**ransfer)
Request	A message sent from one Client/Server peer to the other, followed by an answer message

Table 9: Abbreviations

Term	Definition
AAA	Authentication, Authorization and Accounting
ACL	Access Control List
AF	Application Function
API	Application Programming Interface
AVP	Attribute Value Pair
CLI	Command Line Interface
CPF	Control Plane Function
DEA	Diameter Edge Agent
DRA	Diameter Routing Agent
EMS Site	Element Management System Site
FEP-In	In-Front End Proxy
FEP-Out	Out-Front End Proxy
HA	High Availability
HSS	Home Subscriber Server
HTTP	Hypertext Transfer Protocol
IaaS	Infrastructure as a Service
IMS	IP Multimedia Subsystem
JMS	Java Message Service
KPI	Key Performance Indicator
LDAP	Lightweight Directory Access Protocol
LTE	Long Term Evolution
MME	Mobility Management Entity

Term	Definition
NGN	Next Generation Networking
Node	Physical or virtual addressable entity
OAM	Operation, Administration and Maintenance
OCS	Online Charging System
OVF	Open Virtualization Format
PCEF	Policy and Charging Enforcement Function
PCRF	Policy and Charging Rules Function
PLMN	Public Land Mobile Network
SCCP	Signaling Connection Control Part
SCTP	Stream Control Transmission Protocol
SDC	Signaling Delivery Controller
SNMP	Simple Network Management Protocol
SS7	Signaling System No. 7
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
UE	User Equipment
URI	Universal Resource Identification.
VIP	Virtual IP
VM	Virtual Machine
VNFC	Virtualized Network Function Component
VPLMN	Visited Public Land Mobile Network
Web UI	Web User Interface

Term	Definition
WS	Web Service