Signaling Delivery Controller

Overload Control Overview

4.4

Catalog Number: FD-015-44-27 Ver. 2

Publication Date: June 2015

Legal Information

Copyright

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

AskF5, F5, F5 [DESIGN], F5 Networks, OpenBloX, OpenBloX (design), Rosetta Diameter Gateway, Signaling Delivery Controller, SDC, Traffix, and Traffix [DESIGN] are trademarks or service marks of F5 Networks, 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 Networks. The information in this document may be changed at any time without notice.

About F5 Networks

F5 Networks (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 Overload Control Overview

Catalog Number: FD-015-44-27 Ver. 2

Publication Date: June 2015

Document Objectives

This document provides an overview of the overload control capabilities provided by the SDC.

Document History

Revision Number	Change Description	Change Location

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. SDC Traffic Flows

The F5® Traffix® Signaling Delivery Controller™ is installed in and between networks, and processes traffic between different network elements.

Figure 1: Basic Traffic Flow between the SDC and Networks

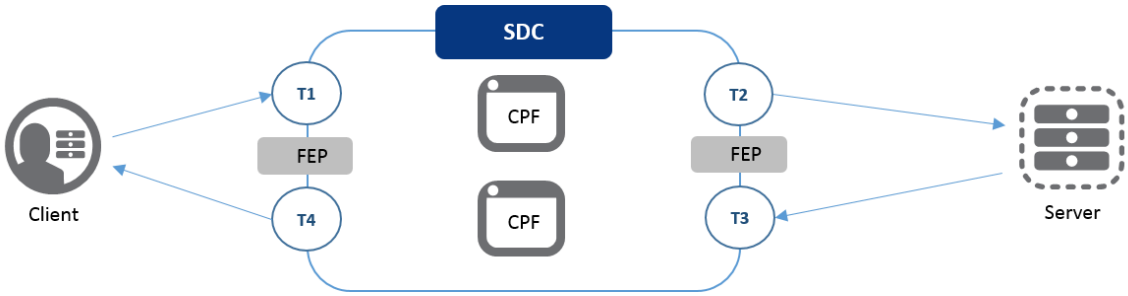

The basic traffic flow between the SDC and the networks is illustrated in *Figure 1*. In this flow, message requests are sent from clients, received by the SDC, and then sent by the SDC to a server. Message answers are then sent from the server back to the SDC, and then sent by the SDC to the client.

This flow includes two types of traffic– incoming (from the client/server to the SDC) and outgoing (from the SDC to the client/server). The volume of traffic received by the SDC at an entry point (T1, T3) or exit point (T2, T4) is monitored and can be limited. These limits ensure that the overall traffic flow performance is constantly under control and no service degradation will occur in overload conditions.

There are multiple possible reasons for overload, such as signaling storms caused by faulty peers, unexpected memory demands, or CPU or other resource utilizations that exceed the engineered capacity of the SDC. The implemented overload control mechanisms assure that in the event of an overload, traffic processing continues with minimal disruption.

These mechanisms control and limit the resource usage and allocation, by controlling the number of incoming/outgoing message requests and traffic rates per destination peer.

The overload protection provided for incoming and outgoing traffic (at the SDC entry and exit point, respectively) is described in the following chapters.

2. Incoming Traffic Overload Control

2.1 How are SDC Nodes Protected?

Incoming traffic flows are flows in which messages are sent from network elements (client or server peers) to the SDC. The two SDC entry points that receive this incoming traffic (T1 and T3) are illustrated in *Figure 2*. The SDC protects itself from incoming traffic overload using a sliding window mechanism that enforces the rate limit configured during system setup. The SDC further protects itself from excessive traffic by constantly adapting the transport layer configuration to only receive the amount of traffic that can be buffered in the SDC queues.

Figure 2: Incoming Traffic Received at SDC Entry Points

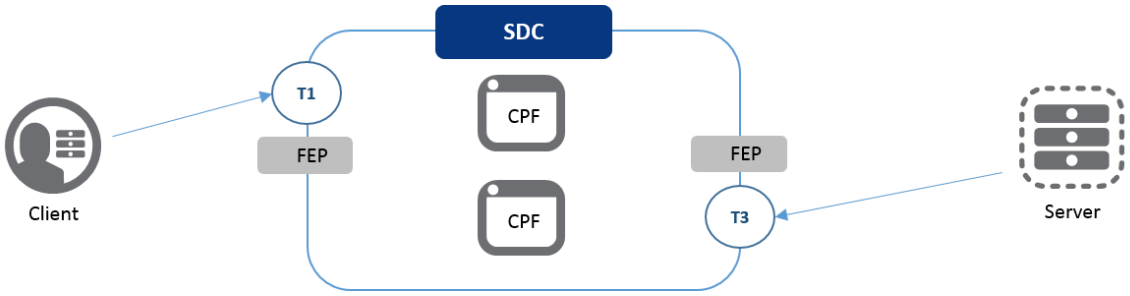

2.2 Message and Byte Rate Limits

Incoming rate limits are configured to control the amount of traffic that the SDC node receives from either a client or server peer. These limits are configured by the number of messages and/or bytes that the SDC can receive. This incoming traffic can either be limited per the client or server peer that the traffic is sent from, or per the SDC component (FEP/CPF) that receives the traffic.

2.2.1 Defining the FEP Rate Limit

In the incoming traffic flow, the FEP is one of the SDC components that receive traffic from the client or server peers. Configuring a “receiving rate limit” for a FEP limits the messages that the FEP is able to receive. This rate limit can be configured by number of messages and/or number of bytes, and is configured individually for each FEP in a site.

The FEP receiving rate limit is either configured globally or per peer. When configured globally, traffic received by the specific FEP – regardless of origin peer – is monitored and compared against the configured limit.

When configured per peer, traffic received by the specific FEP is monitored per origin peer, and compared against the configured limit.

Once the FEP receives more traffic than the configured limit, messages that are sent towards it are not processed. The SDC can either return a busy result code for these messages or silently discard them.

For more information about configuring the incoming rate limit, see the *F5 SDC User Guide*.

2.2.2 Defining the CPFs Rate Limit

Sometimes, in the incoming traffic flow, the client or server peers are connected directly to the CPFs, (instead of to a FEP) and send traffic towards them. Configuring a general "receiving rate limit" for all CPFs applies the defined value to each CPF, and limits the messages that each CPF is able to receive to the defined value. This rate limit can be configured by number of messages and/or number of bytes, and is configured per site.

Note: All CPFs are configured with the same rate limit. Unlike the FEP rate limit, unique rate limits cannot be defined for different CPFs in a site.

The receiving rate limit is either configured globally or per peer. When configured globally, traffic received by a CPF – regardless of origin peer – is monitored and compared against the configured limit.

When configured per peer, traffic received by a CPF is monitored per origin peer, and compared against the configured limit.

Once the CPF receives more traffic than the configured limit, messages that are sent towards it are not processed. The SDC can either return a busy result code for these messages or silently discard them. For more information about configuring the incoming rate limit, see the *F5 SDC User Guide*.

2.2.3 Defining the Origin Peer Rate Limit

In the incoming traffic flow, the client/server peer is the network element sending traffic to the SDC. Configuring the "Peer receive rate limits" for a peer limits the number of messages that the peer can send toward the SDC. The "peer received rate limits" is defined per peer, when configuring a peer profile.

Once the SDC component receives more traffic than the configured limit, messages that are sent towards it are not processed. The SDC can either return a busy result code for these messages or silently discard them. For more information about configuring the incoming rate limit, see the *F5 SDC User Guide*.

2.3 Incoming Traffic Congestion Control

Increasing the SDC ability to react efficiently to sudden occurrences of high traffic, the overload control mechanism also includes a high watermark threshold for the incoming traffic. Once it is exceeded, the SDC limits the allowed volume of incoming traffic by modifying the transport layer behavior (for example, changing the TCP window size).

3. Outgoing Traffic Overload Control

3.1 How are the Clients, Servers, and Server Pools Protected?

Outgoing traffic flows are flows in which messages are sent from the SDC to network elements (clients, servers, or pools of servers). The two SDC exit points that send this traffic (T2 and T4) are illustrated in *Figure 3*.

Figure 3: : Outgoing Traffic Sent from SDC Exit Points

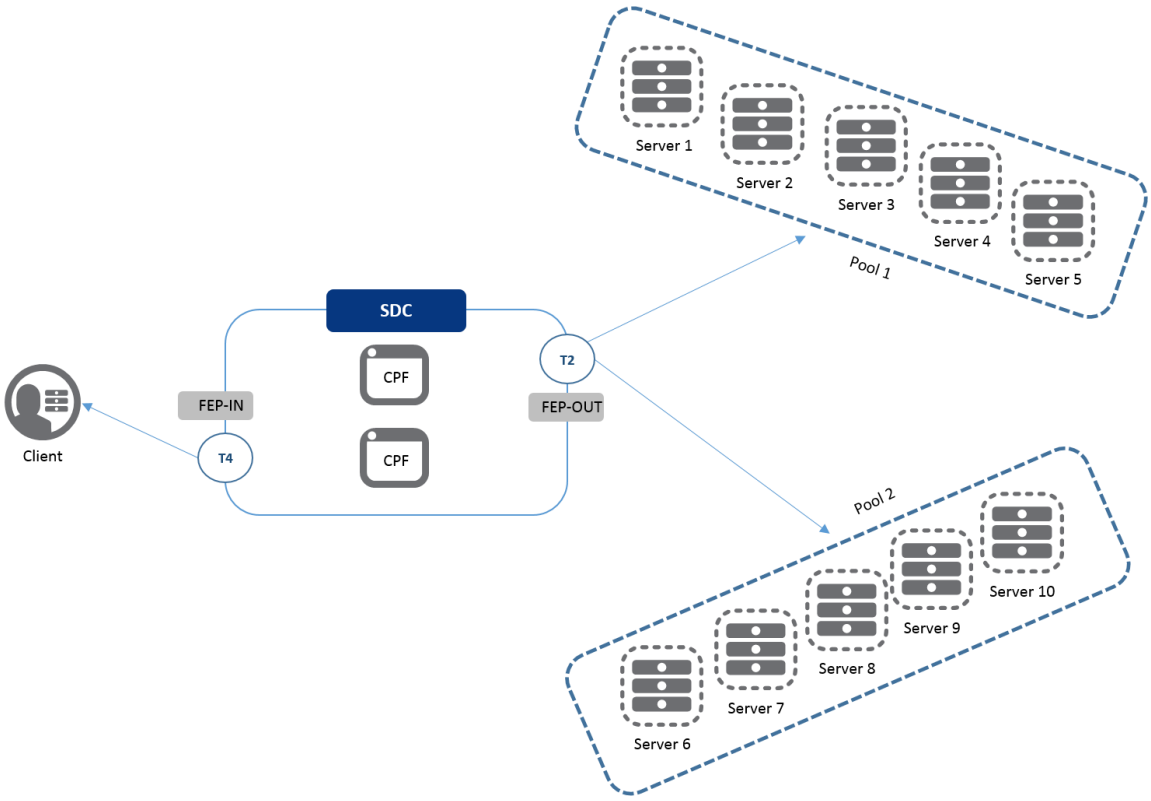

To ensure that the peers and server pools that the SDC sends messages to can efficiently receive the messages, rate limits can be configured. Load balancing policies can also be defined to efficiently distribute traffic sent from the SDC between servers in a pool. These policies are used to deal with and minimize server overload.

When a peer (or pool of server peers) nears or exceeds the configured rate limits, the traffic sent to it by the SDC is minimized, to ensure that minimal server degradation is experienced.

The SDC also recognizes that during pool initialization, the pool of peers is not yet capable of processing traffic to its full ability. Therefore, pool ramp up period can be configured, prioritizing the traffic sent and minimizing the volume of traffic sent.

The following sections describe these different mechanisms in detail.

3.2 Load Balancing Policies

Load Balancing policies are used when messages are routed to a pool of server peers. The peer selection is based on the pool's defined load balancing policy. The load balancing policies provided by the SDC are described in *Table 2*.

Table 2: Load Balancing Policies

Load Balancing Policy	Description
By Precedence	Messages are sent to the first peer in the pool, as long as it is open for traffic. When the first peer goes out of service, messages are sent to the next peer in the pool, and so on. When the first peer gets back to service and reopens, messages are again sent to that first peer.
Round Robin	Messages are evenly distributed across the pool's available peers, in the order that the peers appear in the pool settings.
Weighted Round Robin	Messages are distributed across the pool's available server peers according to a predefined proportion. The weight of each server peer is set during peer configuration, and should be based upon its ability to handle incoming requests. Weighted Round Robin is a static algorithm. No

Load Balancing Policy	Description
	external parameters are taken under account upon request distribution. With Weighted Round Robin, new requests are distributed in a round robin pattern, but instead of sending the request to the next available server peer in line, requests are sent to the server peer that has not yet reached its quota.
Fastest Response Time	Messages are sent to the server peers according to the peer's response time. The response time is used as the weight of the server peer. Fastest Response Time is a dynamic algorithm since it takes external parameters (response time) into account upon request distribution.
Queue Size Ratio	Messages are sent to the servers peers according to the weight/queue length ratio. If Server A's weight is higher than Server B's weight, the policy assumes Server A has a higher traffic handling capacity and maintains a longer queue of pending requests, compared to other servers in the Pool. That is, the higher the server's weight, the greater the number of pending requests it will handle. After getting the performance figures from the active peers (RTT or the number of pending requests), they are normalized between the value 1 and the maximal ratio (the default value is 100): The highest value is 1 while the lowest value is the max ratio value. Queue Size Ratio policy is a dynamic algorithm and responds to external fluctuations upon request distribution.

Load Balancing Policy	Description
Load Based	Messages are distributed between servers based on the real-time performance and load experienced by the servers in the pool. Servers with the least load will be the first to receive requests.
Contextual	The Contextual load balancing policy maps the clients' session ID's to a list of available server peers. This way messages are sent to a specific server peer, according to their session ID.
Weighted Contextual	The Weighted Contextual load balancing policy maps the clients' session ID's to a list of available server peers. This way messages are sent to a specific server peer according to their session ID. In addition to the session ID parameter, traffic distribution is also controlled by a predefined proportion. The weight of each server peer is set during the peer configuration and should be based upon its ability to handle incoming requests.
External	The peer is selected according to an external script's rule.

3.3 Message Rate Limits for a Pool and/or Peer

Rate limits are configured to control the number of messages that are sent from an SDC exit point to either a client peer, server peer, or pool of server peers.

3.3.1 Defining the Peer Rate Limit

For outgoing traffic, the client/server peer is the network element receiving traffic from the SDC. Configuring a "send rate limit" for a peer limits the number of messages that the peer

can receive from the SDC. The "peer send rate limits" is defined per peer, when configuring a peer profile. For more information, see the *F5 SDC User Guide*.

3.3.2 Defining the Pool Rate Limit

For outgoing traffic, the pool of server peers is a collection of server peers receiving traffic from the SDC. Configuring "Rate Limit (TPS)" for a pool limits the number of messages that the pool can receive from the SDC. The "Rate Limit (TPS)" is defined per pool, during the pool configuration. For more information, see the *F5 SDC User Guide*.

3.3.3 Outgoing TPS Thresholds

The SDC offers configurable thresholds for outgoing traffic monitoring. These thresholds reflect the rate of the outgoing requests sent by SDC. These thresholds are defined per peer and per pool, and allow early detection of potential peer and pool overloads.

Upon reaching the defined thresholds, per pool and/or peer, the system sends an SNMP trap (Pool Rate Limit State Change and/or Peer Rate Limit State Change) indicating that the measured TPS has exceeded one of the configured peer and/or pool thresholds (Minor, Major, Critical).

These thresholds are defined per SDC site. For more information, see the *F5 SDC User Guide*.

3.4 Pool Ramp Up

The Pool Ramp Up mechanism prevents a specific pool from overloading during startup after being out of service, busy, or partially out of service. During the ramp up period, which lasts a minimum of five seconds, traffic is sent to the pool on a gradual basis.

Pool Ramp Up is defined when configuring a pool. For more information, see the *F5 SDC User Guide*.

3.5 Monitoring Destination Peer Health

The SDC monitors the health of the remote peers it communicates with and detects when these peers are overloaded. This is based on the constant real time monitoring of Diameter traffic error events, such as:

- Timeouts
- Response time per peer
- Busy answers
- Other error responses

If the rate of the error events exceeds the user configurable threshold, the Diameter peer server status can be changed to “Out of Service” by the user for a defined time interval. When the Remote Peer state is “Out of Service,” no further requests are delivered to it. Users also have the option to apply an “Out of Service Partially” status. Peers in this state continue to process existing sessions while not accepting new sessions. The error event definition and the time interval duration for these peer states are configured by the user via Groovy script. For more information, see the *F5 SDC User Guide*.

4. Monitoring System Performance and Overload Control

The performance and status of the client/server peers, SDC components, and pools of server peers are constantly monitored by the SDC. The SDC also monitors the messages that are received and sent by the SDC. This information is communicated to the user through SNMP alarms that reflect the system behavior and potential traffic congestion, as well as SNMP statistics and Web UI reports that reflect system behavior,

Glossary

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

Table 3: Terms and Abbreviations

Term	Definition
AAA	Authentication, Authorization and Accounting
ACL	Access Control List
AF	Application Function
Answer	A message sent from one Client/Server Peer to the other following a request message
API	Application Programming Interface
AVP	Attribute Value Pair
CLI	Command Line Interface
Client Peer	A physical or virtual addressable entity which consumes AAA services
CPF	Control Plane Function
Data Dictionary	Defines the format of a protocol's message and its validation parameters: structure, number of fields, data format, etc.
DEA	Diameter Edge Agent
Destination Peer	The Client/Server peer to which the message is sent
DRA	Diameter Routing Agent
EMS Site	Element Management System Site
FEP-In	In-Front End Proxy
FEP-Out	Out-Front End Proxy
Geo Redundancy	A mode of operation in which more than one geographical location is used in case one site fails

Term	Definition
HA	High Availability
HSS	Home Subscriber Server
HTTP	Hypertext Transfer Protocol
IMS	IP Multimedia Subsystem
JMS	Java Message Service
KPI	Key Performance Indicator
LDAP	Lightweight Directory Access Protocol
LTE	Long Term Evolution
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)
MME	Mobility Management Entity
NGN	Next Generation Networking
Node	Physical or virtual addressable entity
OAM	Operation, Administration and Maintenance
OCS	Online Charging System
Origin Peer	The peer from which the message is received
PCEF	Policy and Charging Enforcement Function
PCRF	Policy and Charging Rules Function
PLMN	Public Land Mobile Network
Pool	A group of Server Peers
RADIUS	Remote Authentication Dial In User Service
Request	A message sent from one Client/Server peer to the other, followed by an answer message

Term	Definition
SCCP	Signaling Connection Control Part
SCTP	Stream Control Transmission Protocol
SDC	Signaling Delivery Controller
SDC Site	The entire list of entities working in a single site
Server Peer	A physical or virtual addressable entity which provides AAA services
Session	An interactive information interchange between entities
Slave (Bound) Session	A session which inherits properties from a master session
SNMP	Simple Network Management Protocol
SS7	Signaling System No. 7
TCP	Transmission Control Protocol
TLS	Transport Layer Security
Transaction	A request message followed by an answer message
Tripo	Session data repository
UDP	User Datagram Protocol
UE	User Equipment
URI	Universal Resource Identification.
Virtual Server	A binding point used by SDC to communicate with the Remote Peers (Clients and Servers)
VPLMN	Visited Public Land Mobile Network
Web UI	Web User Interface
WS	Web Service

