

Availability and Scalability

Ericsson Service-Aware Policy Controller

FACILITY DESCRIPTION

Copyright

© Ericsson España, S.A. 2017. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.

Abstract

This document provides a description of Availability and Scalability function provided by Ericsson Service-Aware Policy Controller (SAPC).



Contents

1	Availability and Scalability Introduction	1
2	Availability and Scalability Function	1
2.1	Availability	1
2.2	Scalability	7
3	Availability and Scalability Operational Conditions	9
3.1	Availability and Scalability External Conditions	9
3.2	Availability and Scalability Function Administration	9
3.3	Availability and Scalability Security	10
	Reference List	11





1 Availability and Scalability Introduction

This document provides a description of Availability and Scalability function provided by Ericsson Service-Aware Policy Controller (SAPC).

2 Availability and Scalability Function

2.1 Availability

The SAPC is built on a high available architecture where a single failure does not stop the operation of the cluster. It is built over a cluster of Virtual Machines of three types:

- System Controllers (SC): There are always two SCs in the cluster. They provide the OAM and provisioning services through the OAM virtual IP address.
- Traffic Processors (TP): The number of traffic processor VMs is determined depending on the operator needs (the minimum value is two TPs). TPs provide the SAPC traffic services through virtual IP addresses: they provide the traffic interface, balance the traffic processing among all TPs in the cluster and also store provisioning and dynamic data in a replicated way:
 - Primary copy and one replica for dynamic data (sessions, accumulators) and subscribers.
 - Totally replicated for remaining provisioning data.
- Virtual Routers (VRs): Only provided for Cloud Data Center deployments. Two VRs providing access to the SAPC OAM Virtual IP Address and two VRs providing access to the SAPC Traffic Virtual IP Address.

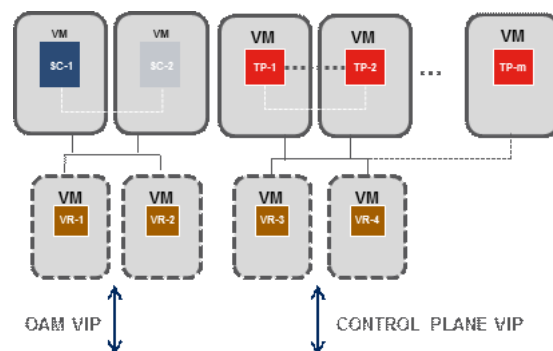


Figure 1 SAPC Cluster architecture.

The two SCs provide the OAM and provisioning services in an active-standby mode, which means that if an SC goes down, all services considering it as the active one, are managed by the other SC. The rest of Virtual Machine types work in an active-active mode. The incoming traffic is distributed by a maximum of six TPs (usually, the first six TPs) among all the available traffic processors in the cluster. These TPs are also the ones publishing the traffic virtual IP address to the external network. If one of these TPs goes down, the publishing of the virtual IP address and the traffic distribution functions are moved to another available TP. Also, this TP is not considered to receive traffic in the distribution until it is up again.

The following situations, in which multiple failures were produced simultaneously, would affect on the SAPC service availability:

- In case the two SC processors were restarted simultaneously, if, after a period of 15 minutes, none of the SC processors are recovered, then a cluster restart is produced until, at least, one of the SCs is recovered. During the SCs absence, Diameter traffic is processed, but the OAM traffic cannot be handled because of the OAM VIP is not available. Performance counters are stored in memory during SCs absence and the result is dumped in the performance files once SCs are up again. Alarms and logs are not available during SCs absence and recovered once are up again. For Geographical Redundancy scenario, the cluster reboot is produced immediately after the SC processors restart.
- In case two TP processors containing both replicas of dynamic data or subscribers were restarted, the SAPC database would be restarted to assure consistency. In this case, Gx session data (also including, the time-related session events and the internal IP session usage accumulator information) is not recovered and Diameter connections with PCEFs are closed.
- In case the two VRs providing access to the OAM virtual IP address are restarted simultaneously, OAM incoming traffic would be affected.
- In case the two VRs providing access to the Traffic virtual IP address are restarted simultaneously, incoming traffic would be affected.

To increase the reliability and availability of the system, the SAPC includes several control mechanisms, as restoration procedures, overload control, session clean up procedure, and mechanisms to overcome connectivity loss.

2.1.1 Restart and Restore Procedures

The SAPC provides mechanisms to handle restart situations both for the SAPC itself and for the peer traffic plane nodes, and also procedures to restore.

2.1.1.1 SAPC Restart

Even when the SAPC provides a high level of availability, in case both SCs fail simultaneously during more than 15 minutes, the SAPC is restarted. Once the SAPC recovers from a restart, the last database information is recovered from stored backups. The information recovered may not be fully up to date and for this reason some actions are performed by the SAPC to consolidate this information.



Next sections describe the actions performed by the SAPC , in PCC deployment scenario, after a cluster restart.

The SAPC increments its own Origin-State-Id and includes the new value in every response message alerting the peer nodes about the loss of previous session state.

The IP-CAN sessions stored before the restart, are not recovered. And therefore all dynamic data related to sessions is neither recovered:

Note: The Origin-State-Id is a monotonically increasing value that is advanced whenever a Diameter entity restarts with loss of previous state.

- Time Trigger reauthorization events are lost.
- For IP-CAN sessions with end-user notifications, the same notification message can be sent again after the restart.
- For Fair Usage, the usage accumulators may be not accurate.

In case fair usage feature is active and usage accumulator data was stored before the SAPC restart, absolute usage accumulator is recovered from the stored backups. After the restart, the recovered data only contains usage activity until the time the last backup was performed. This usage information is used for the quota calculation for the new and ongoing sessions after the restart.

Gx, Sy and Rx sessions are identified by the Diameter `session_id`. A session is considered unknown if the SAPC does not find in its internal database a session with the same `session_id`. After a SAPC restart, requests sent from PCEFs, AFs or Online Charging Systems for an unknown session, will be answered by the SAPC with the `DIAMETER_UNKNOWN_SESSION_ID` error code.

All subscriber-related data is recovered from the stored backups.

2.1.1.2

Peer Restart

The SAPC is able to detect diameter peer nodes restarts based on the standard mechanism described for Diameter nodes in RFC 6733 (refer to [Diameter Base Protocol, IETF RFC 6733](#)).

The SAPC provides the following mechanisms to handle restart situations for the peer traffic plane nodes.

- PCEF Restart:

The SAPC detects that a Gx Diameter client has restarted after comparing the Origin-State-Id information received in every Gx client's activation request with its locally stored state information for that client. When a received Origin-State-Id is different from the stored one, the SAPC considers that this peer has restarted and, after a period of time of 2 hours, identifies all the IP sessions established by restarted client and removes them (see Section 2.1.2.2 for a more detailed procedure). The Gx massive clean up mechanism starts



2 hours after PCEF restart is detected. During this period of time the SAPC removes the obsolete sessions by means of Basic Clean up mechanism (see Section 2.1.2.1 for further information). This delay in starting massive clean up avoids or at least minimizes collisions between both clean up mechanisms being executed at the same time.

A PCEF can be a cluster of several master nodes and a backup node, and each master and backup can contain several Diameter client applications, each one identified by its origin-host-id. In this scenario, both master and backup nodes should be configured with the same diameter node id that is used in the SAPC to identify the session. In case there is a failover from one of the master nodes to the backup node the massive clean up mechanism is not used since in that case the origin-host-id changes. The SAPC performs the basic clean up mechanism when the backup node sends again a session establishment message for each session.

— AF Restart:

The SAPC detects that an AF has restarted after comparing the Origin-State-Id information received in every Rx client's request (initial or update) with its locally stored state information for that client. When a received Origin-State-Id is different from the stored one, the SAPC considers that this peer has restarted, identifies immediately all the AF sessions established by restarted client and removes them (see Section 2.1.2.2 for a more detailed procedure).

— OCS Restart

The SAPC does not detect when the OCS has restarted. In case OCS restarts losing charging and Sy session information, the network should be configured to terminate the affected IP sessions, or alternatively, to terminate the affected Gx sessions; so the SAPC removes Sy sessions at IP session or Gx session termination and create new Sy session at IP session or Gx session establishment.

2.1.1.3

SAPC Restore

The SAPC provides the System Data type of restore.

System Data backup is used to do a system data fallback to recover to a former version of the whole system with consistency.

After restoring a System Data backup, the SAPC reestablishes the following information:

- Software installed.
- Node configuration.
- Static Subscriber profile and provisioning information stored in the SAPC database.



- Dynamic Subscriber information (that is, usage accumulators) stored in the SAPC database.

And the SAPC loses the following data:

- IP-CAN sessions.
- Time Trigger events.

2.1.2 Session Cleanup Mechanisms

The following mechanisms are implemented in the SAPC to remove the obsolete information.

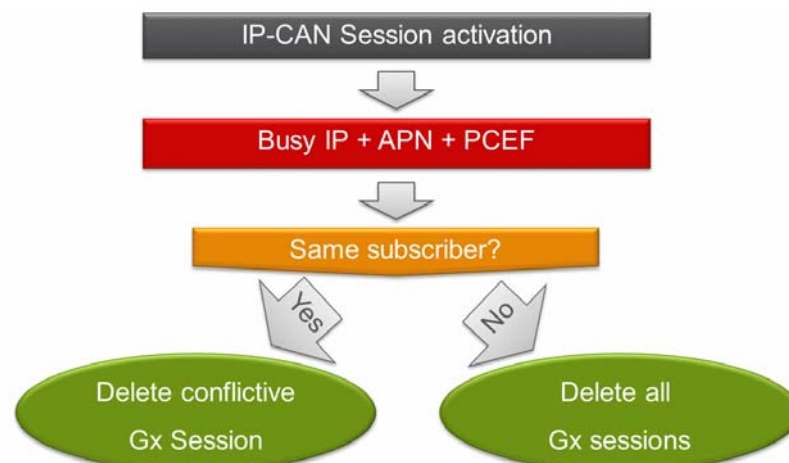
2.1.2.1 Basic Session Cleanup Mechanism

Following mechanism is related with the removal of specific obsolete sessions:

- Gx Mechanism

When the SAPC receives an IP-CAN session activation request from a given PCEF including an APN and an IP address, the SAPC checks if exists any IP-CAN session that was established from the same PCEF with the same IP address and the same APN. If any IP-CAN session is found:

- If the subscriber received in the activation request is different, the SAPC removes the IP-CAN session with all its Gx sessions and a new one is created with the new IP-CAN session request information.
- If the subscriber received in the activation request is the same, just the Gx session matching the received PCEF is removed and a new Gx session is created with the new IP-CAN session activation request information.



Finally in any case when the SAPC sends a RAR for a session towards the PCEF but the session no longer exists in that PCEF, the SAPC receives an RAA message with UNKNOWN_SESSION_ID error and removes the session.

IP session usage accumulator is reset.

As a consequence of the removal of each Gx session, the AF sessions, and Sy sessions are removed as it is done when an IP session is terminated. The SAPC requests the AF to revoke all the identified AF sessions, and to the Online Charging System to terminate the indicated Sy sessions.

2.1.2.2 Massive Cleanup Mechanism

Massive Gx Session Clean up at PCEF Restart

This clean up mechanism consists of deleting all the obsolete IP-CAN sessions existing in the SAPC for a restarted PCEF considering also:

- The AF sessions corresponding to the IP-CAN sessions are deleted and a request is sent to the corresponding AF to revoke the identified AF sessions.
- In a Multiple Gx scenario, the SAPC does not remove associated IP sessions of the other PCEF Diameter clients.
- Since the sessions are lost in the PCEF, IP session usage accumulator is removed.
- The Sy sessions corresponding to the removed IP sessions are also deleted if no more IP sessions are bound to the subscriber. When the Sy session is deleted, session termination message is sent to the Online Charging System.
- Gx massive clean up mechanism is load regulated as explained in [Overload Control](#).

Massive Gx Session Clean up at PCEF Peer Removal

When a `diameterNode` peer is removed from the configuration data, the SAPC removes all the IP-CAN sessions established by that peer, using the same PCEF restart mechanism.

Massive Rx Session Clean up at AF Restart

This clean up mechanism consists of deleting all the obsolete AF sessions existing in the SAPC for a restarted AF considering also:

- The PCC rules from the IP sessions bound to the deleted AF sessions are removed, and corresponding PCEF is notified by means of a Gx RAR message.
- Rx massive clean up is load regulated as explained in Reference [2].



Note: The SAPC provides a robust mechanism that allows to clean the obsolete sessions even in case of geored switchover or scaling scenarios.

Both massive Gx and Rx clean up processes continue scanning and removing sessions until all the obsolete IP-CAN or AF sessions of the restarted peer have been removed.

2.2 Scalability

The SAPC is built on a scalable architecture providing the ability to, on runtime, increase the capacity for traffic processing adding additional processors (Scale-out) or reduce the capacity removing existing processors (Scale-in). SAPC is able to keep performance levels (few seconds impact on the ongoing traffic) when Scale-out or Scale-in functions are performed. VM types to be scaled are only the traffic processors (TPs).

Next figure shows a SAPC cluster initially installed with m TP VMs, that have been scaled-out up to z TP VMs.

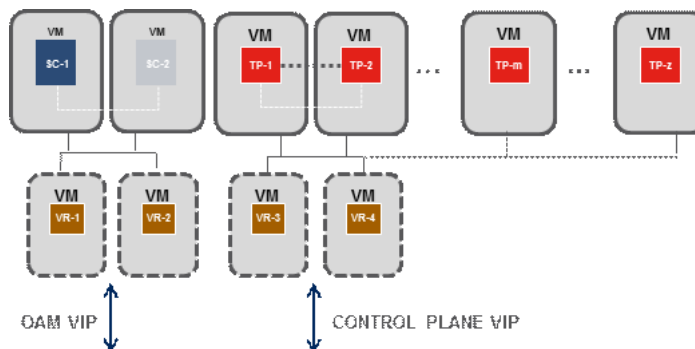


Figure 2 SAPC Cluster where new TP VMs have been added

When TPs are scaled, the traffic interface and traffic distribution functionalities are also included, up to six running instances in six different traffic processors. From this number onwards, the new scaled TPs provide these functions in a spare way (as standby to become ready if any other active instance gets down).

2.2.1 Multi-site Support

The SAPC supports geographical distribution (multi-site) configurations when a single SAPC does not have enough capacity to handle all the subscribers traffic in the following scenarios.

2.2.1.1 SAPC with Common Database

In this deployment, the operator has multiple SAPCs deployed in different sites and a common database to store the subscriber data. Hence, any SAPC can serve IP sessions from any subscriber. Fair Usage Accumulators must be stored in the

common database, so that any of the SAPCs can access and modify the data at any time

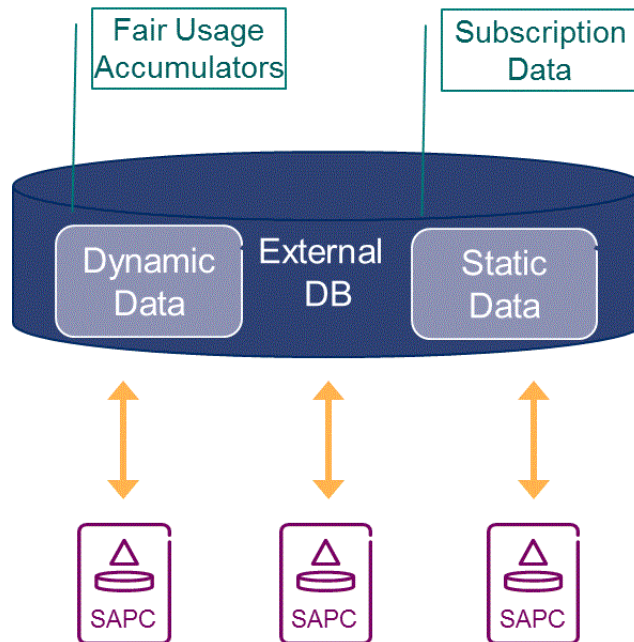


Figure 3 Multi-Site Deployment with Common Database

Subscribers static data and Fair Usage Accumulators are centralized in the external database. The rest of static data as Subscriber Groups, Services, and so on, together with operator defined policies, are provisioned in all the SAPCs.

2.2.1.2

Network Dependencies

The following considerations must be taken into account in deployments with multiple SAPCs.

- The SAPC keeps track of the state of the session, hence all the messages belonging to the same IP session must be routed to the same SAPC.
- In scenarios with dynamic policy control, successful session binding requires that all diameter sessions established over the Gx and Rx reference points for a certain IP session are routed to the same SAPC.
- Network deployments with multiple PCEF require that Gx sessions from two or more PCEFs that belong to the same user IP session, are routed to the same SAPC.
- Fair Usage with Multiple IP session functionality require that all the IP sessions for the same subscriber are routed to the same SAPC. In external database scenarios, concurrent updates of the subscriber usage accumulators as a result of usage reporting received from different IP-CAN sessions can lead to some concurrent usage reporting is not accumulated as SAPC does not perform concurrency control.



- Shared Devices Plans functionality requires that IP sessions that share a subscription are routed to the same SAPC.

- Shared Subscriber Plans functionality requires that IP sessions from all subscribers sharing usage quota are routed to the same SAPC.

Note: Shared Subscriber Plans are not supported with External Database.

- In scenarios where the SAPC is integrated with an Online Charging System for monetary spending limit reporting, it is recommended that all IP-CAN session establishments from the same subscriber are routed to the same SAPC. This allows the SAPC to establish a single session to the OCS per subscriber.

3 Availability and Scalability Operational Conditions

3.1 Availability and Scalability External Conditions

Notice that VIP Gateway routers are not part of a SAPC but are needed in whatever deployment of a SAPC.

3.2 Availability and Scalability Function Administration

The following sections list the relevant Operation and Maintenance related actions, alarms, logs, notifications, and statistics data related to the function.

3.2.1 Availability and Scalability Operation and Maintenance

3.2.2 Availability and Scalability Alarms

There are no specific SAPC alarms related to its availability, apart from the ones provided by the platform:

- “COM SA, AMF SI Unassigned” alarm is risen when a processor is restarted.
- “COM SA, CLM Cluster Node Unavailable” alarm is risen when the cluster cannot access a defined membership machine.
- “eVIP, Gateway Unavailable” alarm is risen when contact is lost with an external gateway. If contact is lost to all external gateways, all traffic is lost.



3.2.3 Availability and Scalability Logging

The following events are logged:

- Existing IP Session removed. (Basic clean up)
- Non-persistent data such as Time of Day, Gx session, and Subscriber notifications are empty
- Diameter peer restarted (Gx and Rx massive clean up)
- Start deleting old sessions (Gx and Rx massive clean up)
- End deleting old sessions (Gx and Rx massive clean up)

3.2.4 Availability and Scalability Notifications

There are no specific SAPC notifications related to service availability, apart from the ones provided by the platform.

- "Link down" and "Link up" notifications are risen when a processor is restarted.

3.3 Availability and Scalability Security

Not applicable.



Reference List

Ericsson Documents

- [1] Subscription and Policy Management
- [2] Overload Control

Standards

- [3] [Diameter Base Protocol, IETF RFC 6733](#)