



MSF Performance Interoperability

Performance Management
Reporting IOT Event
(P-IOT 2010)

Contents

Executive summary	p. 3
The MultiService Forum (MSF)	p. 3
The P-IOT 2010 Interoperability Event	p. 3
Key Objectives of the P-IOT Interoperability Event	p. 3
Key Results	p. 4
Introduction	p. 6
Part I: Participants and Planning	p. 8
Part II: P-IOT 2010 Execution	p. 9
Part III: Results and Issues	p. 13
Future Work	p. 16
Appendix A: Test Scenarios	p. 17
Scenario 1 – Single Domain Performance Management Reporting	p. 17
Scenario 2 – Multi-Domain Performance Management Reporting	p. 27
Appendix B: The Benefits of MSF Membership	p. 33
Appendix C: Participants in the P-IOT 2010 Event	p. 35

Executive Summary

The MultiService Forum (MSF)

The MSF is a global association with a membership that includes the world's leading Internet Protocol (IP) communications companies. The MSF promotes the testing of interoperability based on open standards within a defined end-to-end architecture. The established Global MSF Interoperability (GMI) events are now being complemented by more targeted test events focused on critical and timely issues associated with deployment of IP communications. The Performance Interoperability Test (P-IOT) 2010 Event is the second test event of this nature (following the Long Term Evolution/Evolved Packet Core (LTE/EPC) IOT Event in March 2010). The P-IOT is the first of a planned series of IOT events covering key aspects of network robustness including:

- network monitoring,
- congestion and overload control,
- disruption testing,
- network feedback and instabilities,
- prioritization,
- End-to-end quality of service (QoS) and QoS over the Network-to-Network interface (NNI) in particular.

The P-IOT 2010 Interoperability Event

The P-IOT 2010 Interoperability Event focused on the ability of Session Border Gateways (SBGs) to monitor normal and priority voice and video sessions under various network loads, and report QoS statistics to a Network Operations Center (NOC). The P-IOT 2010 Event also tested the ability of a Data Border Gateway (DBG). A DBG is a class of gateway at the edge of the Service Provider's Network to provide ingress/egress points for non-IMS traffic and associated bearers. The DBG functionality includes topology hiding, security and traffic shaping.

The P-IOT 2010 event was conducted in the U.S. Federal Government's National Communications System's (NCS) eXperimental Testbed Environment (XTE) laboratory in Chantilly, Virginia.

Two physical scenarios were specified for the P-IOT 2010 event:

- Scenario 1 – Single Domain Performance Management Measurement,
- Scenario 2 – Multi Domain Performance Management Measurement.

Within these two physical scenarios, a total of 18 separate test cases were defined.

Key Objectives of the P-IOT 2010 Interoperability Event

Service Providers deploying Next Generation Networks (NGNs) aim to make them as resistant to failures, attacks and congestion as the legacy Public Switched Telephone Network (PSTN). The ability to define, verify and test the features, functions and attributes of a communications network architecture is key to ensuring such robustness is realized in practice. Two key attributes of such an architecture are the ability to accurately monitor and report network performance under varying load conditions, and the ability to prioritize various traffic classes under all load conditions.

The various traffic classes could include normal voice calls as well as priority Emergency Telecommunications Services (ETS) voice calls.

The P-IOT 2010 event was used to:

- Demonstrate network performance monitoring and reporting functions of SBGs for Voice over IP (VoIP), video and data connections,
- Investigate the ability to collect performance monitoring statistics on single and interconnected networks under normal, overload and congestion conditions,
- Examine the ability of SBGs to manage voice, video and data traffic of differing priority classes,
- Determine if sufficient information is collected to calculate the QoS Key Performance Indicators (KPIs) that have been proposed by the International IP Interconnection Forum (i3Forum).
- Investigate DBG behaviour and performance reporting for non-SIP (Session Initiated Protocol) sessions under varying network loads.

Specifically, the test cases include the following:

- Access side SBG performance and reporting for normal and priority VoIP sessions under varying network loads,
- Access side SBG performance and reporting for normal and priority video sessions under varying network loads,
- SBG NNI provisioning of IPSec tunnels,
- Network side SBG performance and reporting under varying network loads,
- SBG support of i3Forum QoS KPIs,
- DBG performance and reporting for data “sessions” under varying network loads.

Key Results

- **All three tested SBGs included packet loss, delay and jitter in the Call Detail Records (CDRs).** This provides the minimum base functionality necessary for the collection of any performance statistics.
- **Two of the tested SBGs included only locally measured values for packet loss, delay and jitter while the third SBG included two sets of values, one set being based on local measurement and the other on received Real-Time Control Protocol (RTCP) packets.** Without RTCP information, a single SBG may not be able to identify a problem that only impacts one direction of traffic. Correlation of CDRs from all SBGs can provide more insight into call quality; however, inclusion of RTCP statistics in the CDRs can provide additional means of identifying quality issues associated with a call.
- **Some of the tested SBGs inhibited collection of end-to-end statistics by not allowing RTCP packets to transit the SBG.** This is believed to be due to the SBG being based upon earlier IP-to-Time Division Multiplexer (TDM) Gateway (GW) products, where this would be an appropriate action. Additionally, these SBGs were also observed to renumber Real-Time Protocol

(RTP) packets before sending them into an IP domain. These actions were carried out even when the SBG was acting as an IP-IP Gateway between two domains.

- **Some endpoints do not generate RTCP, precluding the collection of true end-to-end metrics.** The MSF also notes that user equipment (UEs)/endpoints resident in the customer premises may be considered "untrusted" by the carriers' networks. Furthermore, the expected plethora of such devices will make it difficult to certify all UEs in terms of their RTCP support. Recognizing these difficulties, carriers may choose to use RTCP received from the end user devices, or may identify a "trusted" entity at the edge of its network to generate RTCP to reflect the portion of the media path in the operator's network. For this latter case, "end to end" would only apply to the carrier-access-edge to carrier-access-edge portion of the session.
- **There was no support of RTCP-XR (Extended Reports).** RTCP-XR packets were not generated by the tested endpoints and thus the ability of SBGs to transit RTCP-XR packets could not be tested. RTCP-XR information was not defined in the CDR formats of the three SBGs tested. The lack of XR support is a reflection of the current status of this work within the Internet Engineering Task Force (IETF) Audio/Video Transport (AVT) Working Group (WG).
- **Only one tested SBG reported video statistics in its CDRs for video teleconferencing calls.** This is believed to be a consequence of early implementations of SBGs only supporting a single audio stream. Additional functionality is required to recognize that multiple media streams can exist on a single session and that separate measurements are required for each stream.
- **In the event of processor overload, SBGs continue generating CDRs but some SBGs stop reporting delay, loss and jitter metrics.** Under processor overload, the completion and maintenance of calls is seen as more important than reporting on QoS, since call completion is higher priority. It should be noted that measurements of emergency communication data may be required in some markets during periods of congestion and overload.
- **The CDRs generated by all tested SBGs were sufficient to enable the calculation of the i3Forum statistics for call completion.** However, the CDR performance statistics may only cover a portion of the end-to-end connection (due to TDM-based SBGs in the call path) and the information in the CDRs may be incomplete (due to non-reported RTCP statistics or no RTCP packets being received). As a result, the performance statistics in the CDRs may not provide an accurate measure of the Service Level Agreements (SLAs).
- **Some SBGs provide direct support for priority communications, with the ability to add Differentiated Services Code Point (DSCP) and/or Class of Service (CoS) markings to packets associated with priority sessions.**
- **SBGs which do not explicitly recognise priority communications via control plane signalling can still be provisioned to provide an additional**

likelihood of session completion for priority session requests versus normal session requests by using SBG QoS capabilities. All of the tested SBGs can be provisioned with header manipulation rules that identify priority sessions and QoS mechanisms that provide preferential treatment to those sessions.

- **Differences were observed between the SBGs regarding support of IPsec on the media interface.** Some of the SBGs allow IPsec tunnels on the (combined) signaling and media interfaces while other SBGs permit an IPsec tunnel on the signaling interface only. The different configurations of the IPsec tunnels did not affect the performance reporting capabilities of the SBGs.
- **The “default” parameters that vendors use to assist in setting up IPsec tunnels are not consistent, making provisioning for interoperability between products cumbersome.** The difficulties observed in establishing IPsec tunnels illustrate the importance of profiles for specifying all parameters going across an NNI. One example of an activity to generate such a profile is the ongoing work in the Alliance for Telecommunications Industry Solutions (ATIS) Next Generation Carrier Interconnect (NG-CI) task force.
- **DBGs are needed to support priority communications for non-SIP-based flows.** The SBGs involved in the test event only supported SIP-based flows such as voice and video communications and did not support monitoring of non-SIP-based data exchanges, such as Web browsing via hypertext transfer protocol (http).
- **The DBG tested supports static QoS policies and provides device level statistics. To support priority communications, “real-time” policy updates and traffic flow statistics are required.** While the monitoring and “deep packet inspection” capabilities are sufficient to support priority communications, current policy mechanisms tend to be static and the policy interfaces tend to be vendor specific. To support priority communications, the DBG needs to be able to accept and incorporate “real-time” policy changes from a network entity like a Policy and Charging Rules Function (PCRF) using a standard interface.

The MSF P-IOT 2010 was the first ever example of interoperability testing focussed on performance management reporting in an end-to-end multi-vendor environment with network congestion. The overall results provide a snapshot of performance management reporting capabilities in current implementations of SBGs. The insight from P-IOT 2010 is being liaised to other fora (e.g. the i3Forum, the ATIS Packet Technologies and Systems Committee (PTSC)) to provide input into their ongoing work.

Introduction

The MultiService Forum

The MSF is a global association with a membership that includes the world's leading IP communications companies. The MSF promotes the testing of interoperability based on standards-compliant architectures and protocols to facilitate the deployment of next-generation multi-service networks.

The MSF has traditionally held bi-annual GMI events to assess interoperability. However, given the rapid pace of technology changes, the established GMI events are now being complemented by more targeted test events associated with specific focus areas. Four areas of focus are currently being pursued by the MSF:

- Long Term Evolution / Evolved Packet Core interoperability,
- Internet Protocol Television (IPTV) interoperability,
- Network Robustness,
- Content Distribution Networks (CDN).

An LTE/EPC IOT Event was conducted by the MSF in March 2010. The results of this test can be found at

<http://www.msforum.org/interoperability/MSF%20LTE%20Interoperability%202010%20White%20Paper.pdf>.

The Performance Management Reporting IOT Event, the second test event of this nature, was conducted from 8 November to 19 November 2010. P-IOT 2010 is the first in a planned series of Network Robustness IOT events.

This white paper presents the results of P-IOT 2010.

Network Robustness

The Network Robustness area of focus includes the following topics:

- Network monitoring,
- Congestion and overload control,
- Disruption testing,
- Network feedback and instabilities,
- Prioritization,
- End -to-end QoS and QoS over the NNI in particular.

P-IOT 2010

P-IOT 2010 focused on the performance management reporting interfaces of Session Border Gateways for voice, video and data sessions under varying conditions of load. The objective of the event was to:

- Verify the completeness and accuracy of the performance-related statistics collected for voice/video/data sessions under various load conditions,
- Determine if the SBG CDRs are sufficient to enable identification of the portion of the end-to-end connection where poor QoS performance was occurring (fault isolation),

- Understand how existing SBG QoS mechanisms can be used to support priority communications.
- Determine if the information provided by the SBGs is sufficient to support the QoS KPIs as defined by the i3Forum (www.i3forum.org),
- Investigate DBG behaviour and performance reporting for non-SIP sessions under varying network loads.

Specifically, the P-IOT 2010 event focussed on:

- Access side SBG behaviour and performance reporting for normal and priority VoIP sessions under varying network loads,
- Access side SBG behaviour and performance reporting for normal and priority video sessions under varying network loads,
- SBG NNI provisioning of IPsec tunnels,
- Network side SBG behaviour and performance reporting under varying network loads,
- SBG support of i3Forum QoS KPIs,
- Data Border Gateway behaviour and performance reporting for non-SIP sessions under varying network loads.

This white paper is organized into three parts and three appendices. Part I discusses the planning of the P-IOT 2010 Interoperability Event, Part II describes the execution of the two-week event, and Part III presents the results obtained from the P-IOT 2010 event. Appendix A provides more details on the physical scenarios; Appendix B discusses the benefits of MSF membership and Appendix C provides brief resumes of the participating companies.

Part I: Participants and Planning

The P-IOT 2010 event involved a diverse group of IP communications professionals whose common goal was to test the current performance management reporting and prioritization capabilities of SBG and DBG products operating in real-world Service Provider environments. The results of this testing will allow vendors to improve their products, Service Providers to accelerate their service deployment strategies, and the MSF to identify standards shortfalls and communicate them to the appropriate Standards Development Organizations (SDOs) and industry fora.

Planning for the P-IOT 2010 Interoperability Event began in April 2010. A formal Work Item was produced which defined a range of items to be tested in support of network robustness. These items were prioritised and the performance management reporting capabilities of SBGs were selected as the initial focal point. Subsequently, the required scenarios and test cases for the event were identified and specified. The scope of the event is documented in the Physical Scenarios document MSF-P-IOT-SCN-001-FINAL which is publicly available at <http://www.msforum.org/techinfo/approved.shtml>. Two distinct scenarios were defined covering performance management reporting in a single domain (involving access-side SBGs) and performance management reporting in a multi-domain environment (involving both access- and network-side SBGs). Each of the two scenarios included a number of sub-scenarios. Test plans were developed for the identified sub-scenarios.

Engineers arrived at the host site prior to the event to ensure that the components were installed, configured and functioning properly. In addition, the P-IOT 2010 event made use of equipment that was already present at the NCS laboratory as part of its ongoing testing activities for priority communications associated with Emergency Telecommunications Services.

Host Site

NCS provided the host site in Chantilly, Virginia, USA. Extensive use was made of test tools and emulation capabilities present in the lab to actively monitor equipment performance and track, identify and correct issues.

Testing was structured to perform Scenario 1 (single domain) tests prior to expanding the environment to the Scenario 2 (multi-domain) tests. In all cases, different instances of each defined test case were run in order to reflect different vendor equipment combinations.

The NCS has worked within the telecommunications industry to define the functionality needed to support ETS in U.S. networks. This includes recognition of an ETS session request and providing preferential treatment to that request and session. Preferential treatment includes exemption from network management controls, exemption from machine congestion controls except at the highest congestion levels, queuing for resources not currently available, and use of resources dedicated to ETS sessions.

The existing NCS lab configuration was used for testing ETS traffic. GENBAND provided its S3 SBG product for this event. Generic header manipulation features in the S3 were successfully configured to identify and prioritize ETS traffic. Future testing will assess the remaining ETS capabilities.

Part II: P-IOT 2010 Execution

P-IOT 2010 involved a single lab. However, the test scenarios identified the need for testing across multiple domains. This was achieved through the use of Virtual Local Area Networks (VLANs) and Virtual Machines (VMs) to create multiple domains. Figure 1 presents a high-level diagram of the test environment used at the NCS XTE host-site.

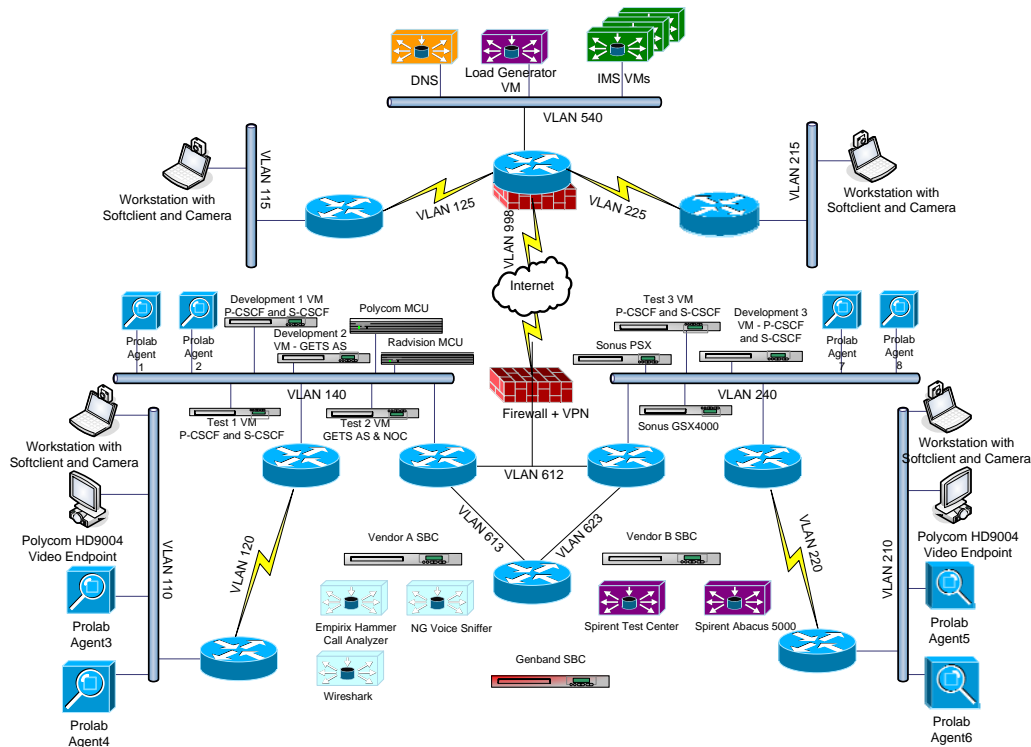


Figure 1 is a high-level view of the NCS XTE test environment.

The testing approach taken was to capture end-to-end performance statistics from endpoints and compare those to the results reported by the SBGs. Load was generated using commercially available load generation tools. Probes at the endpoints were used to measure call completion rates and to capture performance metrics needed to generate Mean Opinion Scores (MOS_{CQE}). The call performance statistics measured by the SBGs were extracted from the CDRs reported by all SBGs involved in each call. The results from these two types of measurement were compared for both normal and priority voice/video calls.

In the general case, when a session traverses a SBG, there are four distinct information flows per media stream. Thus, for a voice-only session traversing a SBG, there will be bi-directional RTP and RTCP flows between the originating and terminating domains. The SBGs can analyze the received RTP packets in both domains to calculate packet loss, delay and jitter for each direction. In addition, packet loss, delay and jitter (as measured by the remote endpoints) can also be obtained from received RTCP packets. This information provides a basis for the P-IOT 2010 performance management reporting.

Network Test Scenarios

Two test scenarios represented different physical configurations of the network components, enabling focused testing of specific interfaces and functionality that arise in real-world situations. The following paragraphs and figures provide an overview of the scenarios, which are described in more detail in Appendix A.

Figure 2 shows the test configuration for Scenario 1 in the P-IOT 2010 event. This baseline scenario is comprised of a single network domain. Within this domain, QoS was provided in both the core and access regions:

- At the edge of the core IP network the session border controller marked traffic with the appropriate differentiated services (DiffServ) code point based on traffic priority level. The core network was configured to provide appropriate treatment based on these markings.¹
- In the access network, DiffServ² was used to mark voice and signalling traffic and to provide proper relative prioritization, e.g., with respect to data traffic.

¹ The NCS does not dictate the DSCP or CoS markings to be used for ETS packets, but requires carriers to specify such markings and the router/switch treatments (e.g., per hop behavior, maximum and minimum flow rates) that these markings will invoke.

² In access networks, an alternative prioritization marking such as 802.1p could be used to provide priority towards the endpoint. Due to security considerations, priority markings for the flow from the endpoint into the core would typically be provided by the first “policy-enabled” device in the access tile. In the XTE, the routers connecting the access VLANs provided this capability.

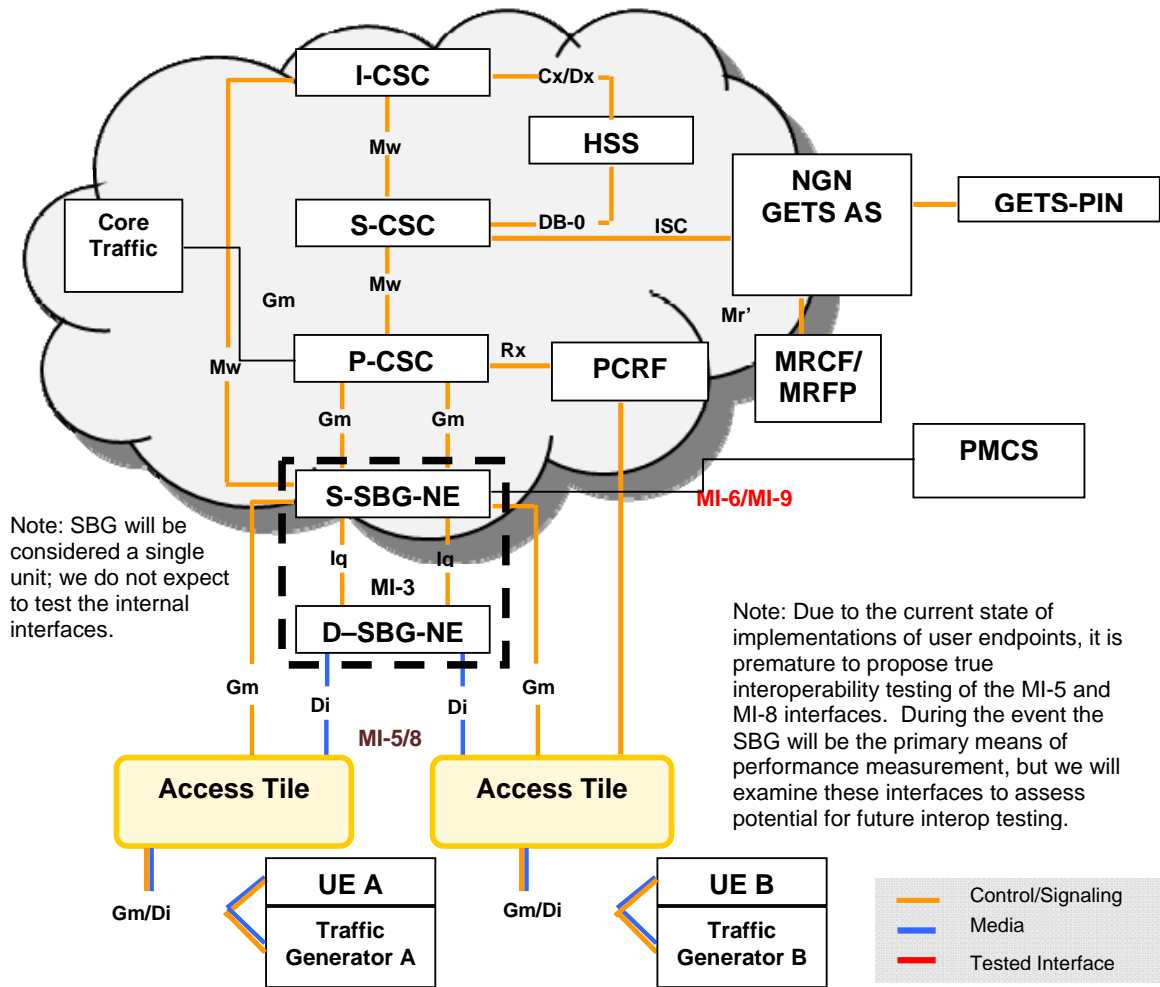


Figure 2 illustrates the network test configuration used for the Scenario 1 test cases.

Scenario 1 included the following test scenarios and related sub-scenarios:

- Access SBG Testing
 - Voice Performance Management Reporting under varying load conditions,
 - Voice Session Establishment Congestion,
 - Voice Bearer Throughput Congestion,
- Priority Video Testing
 - Video Performance Management Reporting under varying load conditions,
 - Video Session Establishment Congestion,
 - Video Bearer Throughput Congestion.
- i3Forum QoS KPIs testing

- SBG's ability to support i3Forum QoS KPIs.

Figure 3 shows the test configurations for Scenario 2 in the P-IOT 2010 event. This scenario, comprised of multiple network domains, enabled specific testing of the performance management, prioritization and reporting capabilities of SBGs at the NNI.

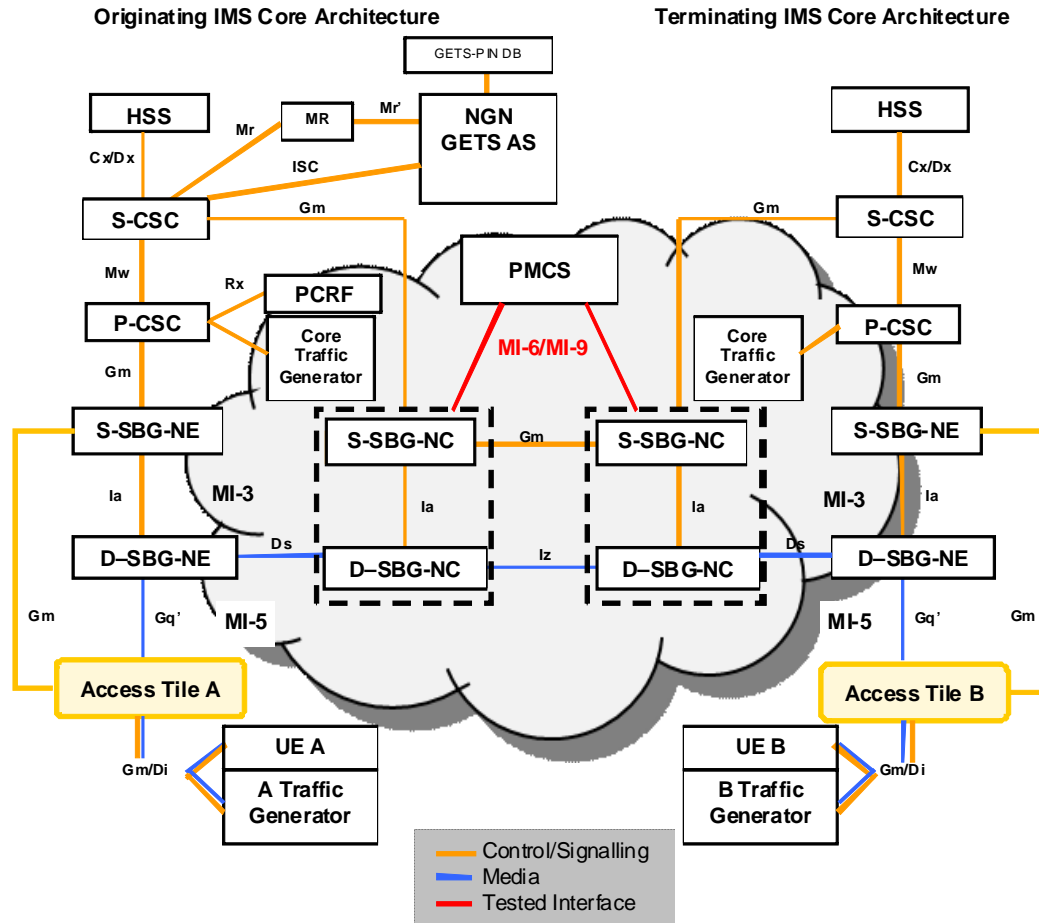


Figure 3 illustrates the network test configuration used for the Scenario 2 test cases.

Scenario 2 included the following test scenarios and related sub-scenarios:

- NNI Provisioning
 - Provisioning of encrypted and unencrypted IPsec tunnels between Network SBGs.
- Network SBG Testing
 - Voice/Video Performance Management Reporting under no load
 - Voice/Video Performance Management Reporting (IPsec tunnel congestion)

- Voice/Video Session Establishment Congestion
- Voice/Video Bearer Throughput Congestion
- Priority Data Testing
 - Dynamic Provisioning of the DBG to support priority data,
 - DBG Performance Management Reporting under no traffic load,
 - DBG Performance Management Reporting under internet congestion.
- i3Forum QoS KPIs testing
 - SBG's ability to support i3Forum QoS KPIs.

Part III: Results and Issues

The test event was designed to test several performance-related issues. The event objectives were:

- To test the performance management reporting interface of SBGs with regard to its accuracy under varying load conditions for voice/video/data sessions,
- To evaluate the behaviour of the SBGs under load conditions, and to specifically investigate their handling of priority and non-priority/normal sessions,
- To understand the RTCP capabilities of endpoints and SBGs, as those capabilities have a major influence on the ability to report the specific QoS KPIs that have been defined by the i3Forum³,
- To determine the extent to which SBGs based their reported statistics on local measurements or received RTCP or a combination of both.
- To investigate DBG behaviour and performance reporting for non-SIP sessions under varying network loads.

The following observations were made:

- **All three tested SBGs included packet loss, delay and jitter in the CDRs.** This provides the minimum base functionality necessary for the collection of any performance statistics
- **Two of the tested SBGs included only locally measured values for packet loss, delay and jitter while the third SBG included two sets of values, one set being based on local measurement and the other on received RTCP packets.** Without RTCP information, a single SBG may not be able to identify a problem that only impacts one direction of traffic. Correlation of CDRs from all SBGs can provide more insight into call quality; however, inclusion of RTCP statistics in the CDRs can provide additional means of identifying quality issues associated with a call. (For example, if the ingress calling-to-called statistics reported by SBG A are good, while the calling-to-called statistics reported by SBG B are bad, one can infer that there is a problem in the segment between SBG A and SBG B). A problem on the segment between the last SBG and the called party will not be identified if only ingress flows are captured in the CDRs.

Inclusion of RTCP statistics in the CDRs provides additional means of identifying quality issues associated with a call. Such CDRs can identify which flows through the SBG are having problems and whether the problems are unidirectional or bidirectional. Correlation of CDRs from all SBGs in the call path would still be required to identify the call path segment between SBGs where a problem is occurring.

As is noted below, not all end points generate RTCP packets, and certain

³ When measuring QoS on a media flow, it is important to understand the scope of a given RTCP flow in terms of the end-to-end IP based stream and whether RTCP is transited through SBGs or not. The i3Forum has recommended that RTCP should be supported by all media endpoints and transited through any intermediate IP-IP gateways in the media path.

SBGs do not transit RTCP packets. Because of this, one cannot be assured that the CDRs containing RTCP statistics provide a more complete and accurate picture of call quality than CDRs without RTCP statistics.

Some of the tested SBGs inhibited collection of end-to-end statistics by not allowing RTCP packets to transit the SBG. This is believed to be due to the SBG being based upon earlier IP-to-TDM Gateway products, where this would be an appropriate action. Additionally, these SBGs were also observed to renumber RTP packets before sending them into an IP domain. These actions were carried out even when the SBG was acting as an IP-IP Gateway between two domains

Ideally, the MSF believes that an IP-IP gateway should have the ability to transit an RTCP flow, act as a RTCP Back-to-Back User Agent (B2BUA) (terminate RTCP on one leg and generate RTCP on its other leg) or a combination of both. In the latter case, this would then enable two separate RTCP flows and would facilitate fault isolation. Such a concept has been raised in the IETF MONARCH Working Group ("Monitoring Architectures for RTP").

- **Some endpoints do not generate RTCP, precluding the collection of true end-to-end metrics.** The MSF also notes that UEs/endpoints resident in the customer premises may be considered to be "untrusted" by the carriers' networks. Furthermore, the expected plethora of such devices will make it difficult to certify all UEs in terms of their RTCP support. Recognizing these difficulties, carriers may choose to use RTCP received from the end user devices, or may identify a "trusted" entity at the edge of its network to generate RTCP to reflect the portion of the media path in the operator's network. For this latter case, "end to end" would only apply to the carrier-access-edge to carrier-access-edge portion of the session.
- **There was no support of RTCP-XR.** RTCP-XR packets were not generated by the tested endpoints and thus the ability of SBGs to transit RTCP-XR packets could not be tested. RTCP-XR information was not defined in the CDR formats of the three SBGs tested. The lack of XR support is a reflection of the current status of this work within the IETF AVT WG.
- **Only one tested SBG reported video statistics in its CDRs for video teleconferencing calls.** This is believed to be a consequence of early implementations of SBGs only supporting a single audio stream. Additional functionality is required to recognize that multiple media streams can exist on a single session and that separate measurements are required for each stream.

Given the wide variability in video packet sizes and burstiness when compared to the relatively "constant" nature of the audio stream, one can have acceptable audio quality on a teleconference call while experiencing "unwatchable" video. Thus one cannot accurately infer overall multi-media session quality on the basis of measuring only a single media stream.

- **In the event of processor overload, SBGs continue generating CDRs but some SBGs stop reporting delay, loss and jitter metrics.** Under processor overload, the completion and maintenance of calls is seen as more

important than reporting on QoS, since call completion is higher priority. It should be noted that measurements of emergency communication data may be required in some markets during periods of congestion and overload.

- **The CDRs generated by all tested SBGs were sufficient to enable the calculation of the i3Forum statistics for call completion.** However, the CDR performance statistics may only cover a portion of the end-to-end connection (due to TDM-based SBGs in the call path) and the information in the CDRs may be incomplete (due to non-reported RTCP statistics or no RTCP packets being received). As a result, the performance statistics in the CDRs may not provide an accurate measure of the SLAs.

The i3Forum has recognized this issue and notes that QoS measurements should be considered meaningful only in a complete Voice over IP environment (no TDM-IP translation) and in a tested and verified voice path where RTCP flows are not blocked. The results of this IOT suggest that these conditions may not be satisfied in many near term deployment scenarios.

- **Some SBGs provide direct support for priority communications, with the ability to add DSCP and/or CoS markings to packets associated with priority sessions.** This capability allows media packets associated with ETS calls to be protected from dropping when an SBG's line card experiences congestion.
- **SBGs which do not explicitly recognise priority communications via control plane signalling can still be provisioned to provide an additional likelihood of session completion for priority session requests versus normal session requests by using their QoS capabilities.** Specifically, the following functionality can be used to provide priority communications during times of network congestion:

- Recognition of priority markings in a session request and acting on these markings. All tested SBGs had Proxy Call Session Control Function (P-CSCF) functionality to recognize specific ETS destination addresses and to treat these requests differently from normal calls,
- Exemption from network management controls for recognized destination addresses,
- Exemption from machine congestion controls for recognized destination addresses, except at the highest congestion level,
- Queuing for resources not currently available and/or use of dedicated resources to complete priority calls,
- Recognition of priority markings in session packets and acting on these markings (e.g., by not dropping these packets).

The three SBGs tested all recognized dialled digits; this capability could be used to distinguish between priority calls and normal calls. Based on this information, the priority session request could be routed to a different domain, and/or different processing and SIP INVITE manipulations could occur. All SBGs support IP header manipulation rules which enable the identification of priority calls in the SIP INVITE. The IP header manipulation rules can be used to detect dialled digits associated with priority calls, or the presence of a SIP parameter, such as "RPH: ets.0" to identify such calls.

- **Differences were observed between the SBGs regarding support of IPsec on the media interface.** Some of the SBGs allow IPsec tunnels on the (combined) signaling and media interfaces while other SBGs permit an IPsec tunnel on the signaling interface only. The different configurations of the IPsec tunnels did not affect the performance reporting capabilities of the SBGs.
- **The “default” parameters that vendors use to assist in setting up IPsec tunnels are not consistent, making provisioning for interoperability between products cumbersome.** The difficulties observed in establishing IPsec tunnels illustrate the importance of profiles for specifying all parameters going across an NNI. One example of an activity to generate such a profile is the ongoing work in the ATIS NG-CI task force. The ATIS NG-CI template and test scenarios may be included in a future IOT.
- **DBGs are needed to support priority communications for non-SIP-based flows.** The SBGs involved in the test event only supported SIP-based flows such as voice and video communications and did not support monitoring of non-SIP based data exchanges, such as Web browsing via http.
 - **The DBG tested supports static QoS policies and provides device level statistics. To support priority communications, “real-time” policy updates and traffic flow statistics are required.** While the monitoring and “deep packet inspection” capabilities are sufficient to support priority communications, current policy mechanisms tend to be static and the policy interfaces tend to be vendor specific. To support priority communications, the DBG needs to be able to accept and incorporate “real-time” policy changes from a network entity like a Policy and Charging Rules Function (PCRF) using a standard interface. While the firewall available in the XTE could be provisioned to support most of the functions of a DBG, the statistics reported were in the form of operational measurements for the device as a whole and not “call detail records” for a given flow. Statistics need to be captured for specific priority “sessions.” The need to capture these statistics would likely be defined by the policy associated with the session. Additional discussions are needed with industry to determine what statistics should be collected.

The overall results provide a snapshot of performance management reporting capabilities in current implementations of SBGs. The insight from P-IOT 2010 is being liaised to other fora (e.g. the i3Forum, the ATIS PTSC etc.) to provide input into their ongoing work

Future Work

Additional IOT events are being planned to look at other aspects of network robustness. As stated in the existing MSF multi-part Work Item, these include:

- congestion and overload control,
- disruption testing,
- network feedback and instabilities,
- prioritization,
- QoS in general and QoS over the NNI in particular.

In particular, IOT events are being considered to test network disruption cases, as well as addressing congestion via cooperation and traffic-sharing among individual carriers. In addition, the test cases in this event may be re-run where modifications have been made to the SBGs and DBGs to improve their performance management reporting capabilities.

The timeframe of any subsequent event will be dependent on the availability of equipment supporting the required functionality as well as the interest of MSF members in seeing such functionality tested in a multi-vendor environment.

Appendix A: Test Scenarios

This appendix provides a detailed description of each of the two test scenarios and presents the test results on a per-scenario basis.

Scenario 1 – Single Domain Performance Management Reporting

Scenario 1 concentrated on the performance management reporting capabilities of an access side SBG in a single IP domain under varying load conditions for voice and video sessions. The ability of the SBG to support the i3Forum-defined QoS KPIs was also tested. The test configuration for Scenario 1 is shown in figure A-1 below.

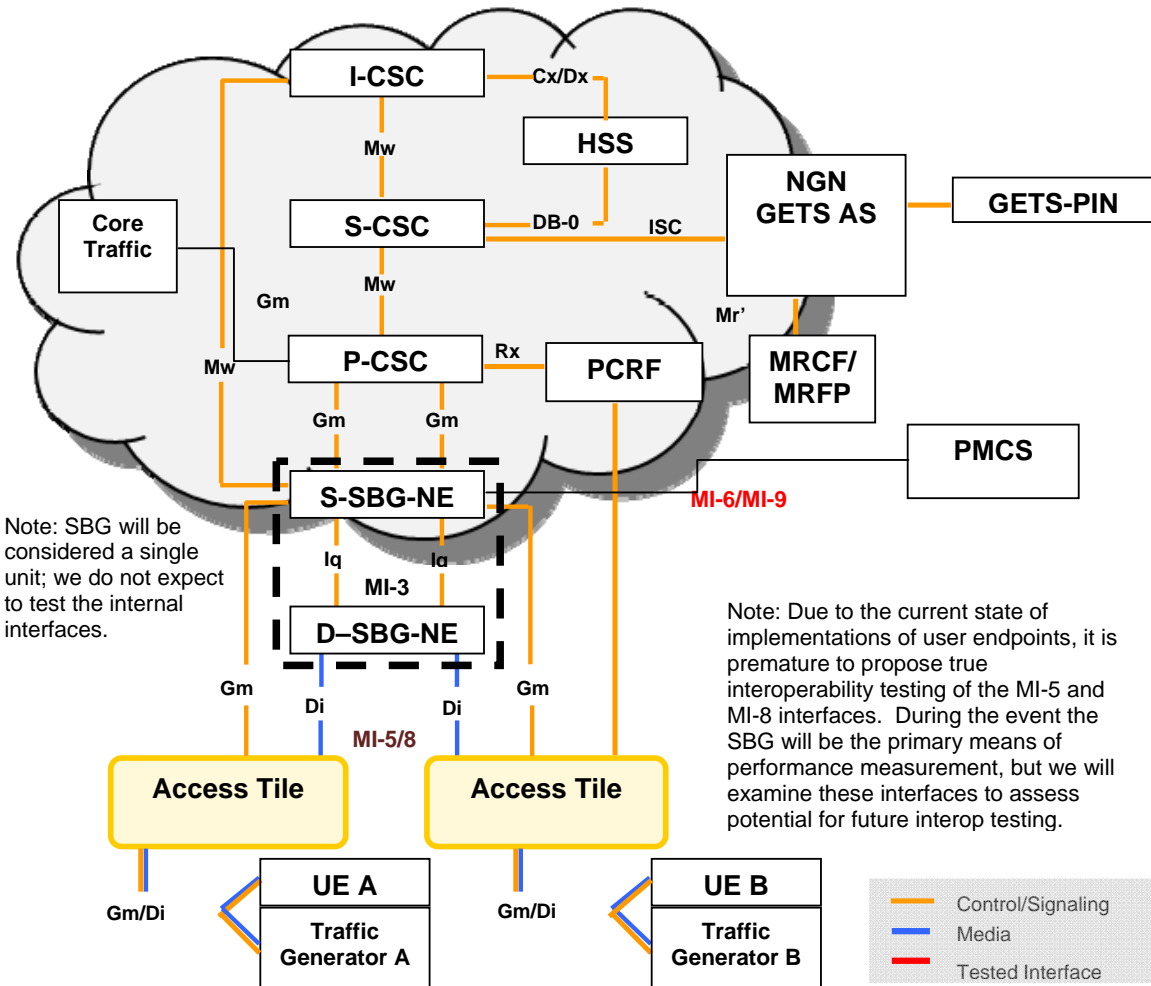


Figure A-1: Test Configuration for Scenario 1

This scenario was divided into a number of sub-scenarios that are detailed below.

Scenario 1.1 – Access SBG Testing

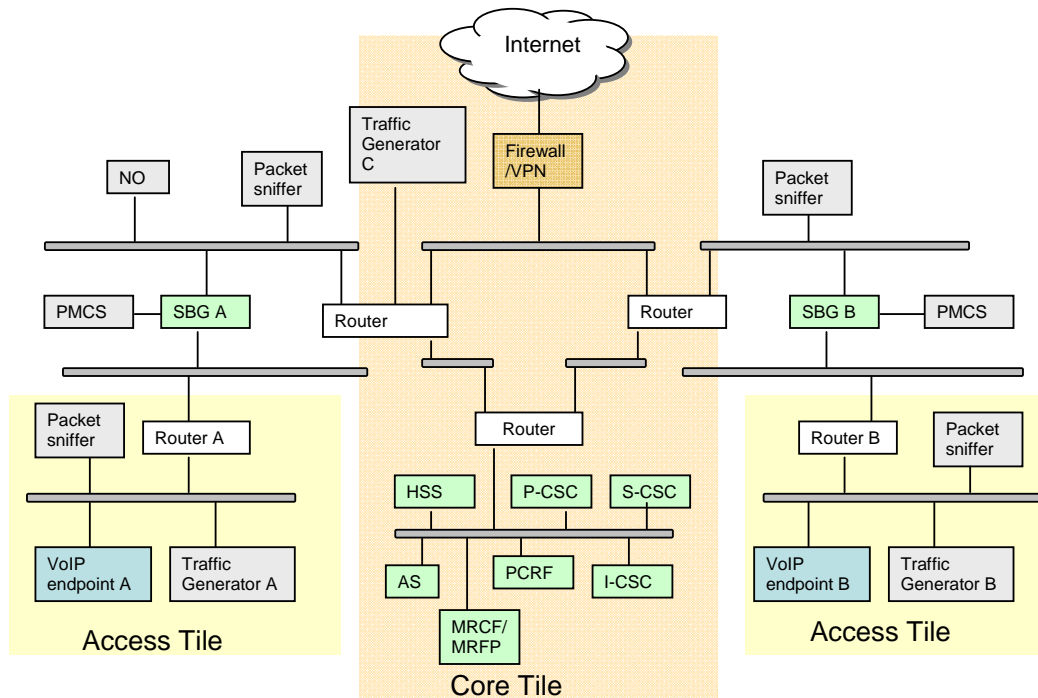


Figure A-2: Scenario 1.1 - Access SBG Testing

The network architecture for scenario 1.1 is shown in Figure A-2 above.

Scenario 1.1 Objectives

Four tests were defined in Scenario 1.1 with the objectives of verifying:

- The ability of the SBG to report performance of VoIP call traversing SBG under no traffic load (S1.1.1),
- The ability of the SBG to report performance of VoIP call traversing SBG under link traffic congestion (S1.1.2),
- The ability of the SBG to reject new VoIP sessions when maximum number of established sessions is exceeded. Different thresholds may be applied for normal and priority voice sessions (S1.1.3),
- The ability of the SBG to reject new VoIP sessions when threshold for maximum bearer layer throughput/bandwidth is exceeded. Different thresholds may be applied for normal and priority voice sessions (S1.1.4).

Scenario 1.1 Test Results and Observations

Figure A-3 identifies the terminology used in reporting the test results and observations. A session is assumed to traverse two domains within the SBG; these domains are referred to as Domain A and Domain B. Each domain has an ingress media stream from an upstream endpoint, and an egress media stream to a downstream endpoint. In the following discussion, Domain A is assumed to be the calling side for the session, while Domain B is assumed to be the called side for the session.

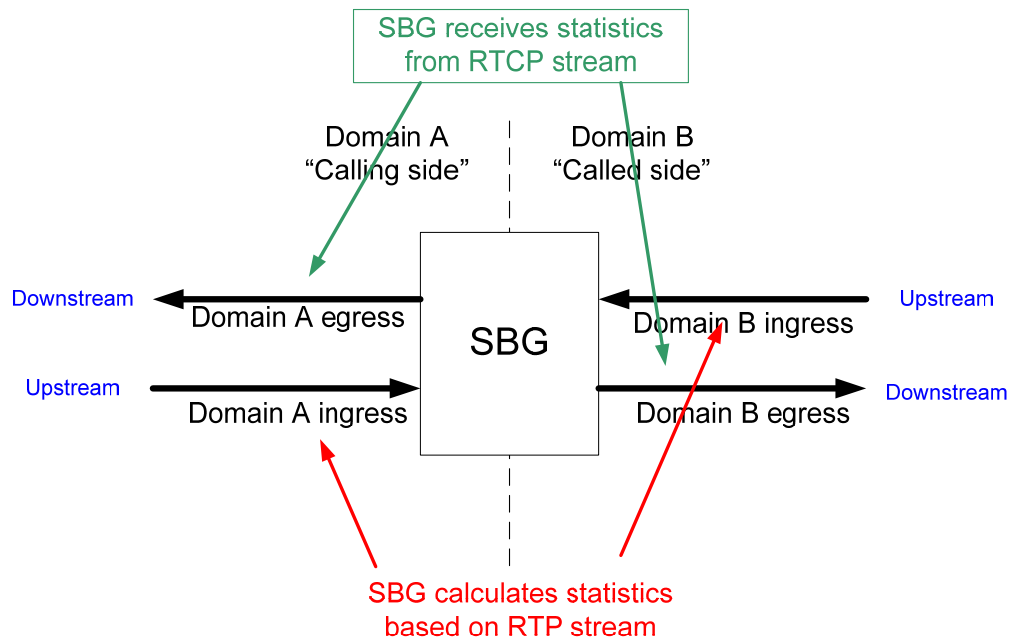


Figure A-3: SBG Terminology

Scenario 1.1.1 – Performance Management Reporting for a VoIP call (no load)

This scenario looks at the performance management reporting of the SBGs for voice sessions with no background load. Each SBC was tested as the only SBC within the call path, and in conjunction with another SBC in the call path, for a total of six test instances.

The results were as follows:

- All three SBCs tested reported the following QoS statistics in their CDRs:
 - “Calling” (i.e., Domain A ingress) packets received
 - Calling packets lost
 - Calling packet latency
 - Calling packet jitter
 - “Called” (i.e., Domain B ingress) packets received
 - Called packets lost
 - Called packet latency
 - Called packet jitter

- Only one SBC tested reported RTCP information captured in its CDRs:
 - “Calling” RTCP (i.e., Domain A egress) packets lost
 - Calling RTCP packet latency
 - Calling RTCP packet jitter
 - “Called” RTCP (i.e., Domain B egress) packets lost

- Called RTCP packet latency
- Called RTCP packet jitter

The reason for collecting RTCP information in a CDR is:

- Within a given domain, the upstream and downstream media need not traverse identical paths from the SBG to the “endpoint.” Thus CDRs with information obtained from RTCP packets allow one to identify that a problem occurred in a single direction within a session.
- RTCP reports provide quality measurements from a point remote from the SBG and may give a different indication of overall quality from local measurements on the SBG based on received RTP (see scenario 1.1.2).
- If there are quality issues on a call, the combination of local and RTCP measurements does help the diagnosis of the part of overall media path which is the source of the problems (see scenario 1.1.2).

The reasons for not collecting this RTCP information in a CDR include:

- A conscious decision was made to apply the processing that could be used to collect this information to handling sessions, especially under congestion scenarios.
- For many sessions, the upstream and downstream media follow the same path (e.g., via a Multiprotocol Label Switched (MPLS) path).
- Correlation of CDRs from multiple devices occurs for trouble shooting and accounting purposes. This correlation will allow one to identify a problem that occurred in a single direction within a session.
- Many end devices do not generate RTCP streams, so these fields will likely have a zero value in the CDR. Furthermore, end devices are typically untrusted (e.g. UEs in customer premises).

Scenario 1.1.2 – Performance Management Reporting for a VoIP call (link congestion)

In this scenario, link congestion was introduced at different portions of the end-end media path to investigate whether the resultant CDRs accurately reported the congestion condition and whether it was possible to diagnose the part(s) of the overall end-end path which were the cause of the quality issues. As in Scenario 1.1.1, a total of six instances were run.

It was noted that there is an important difference between SBGs that support TDM domains versus SBGs that support only IP domains in terms of their handling of RTP and RTCP. All the tested SBGs acted as B2BUAs. However, SBGs that support TDM domains were observed to drop received RTCP packets, due to RTCP packets having no meaning in a TDM domain. Similarly, these SBGs also renumbered RTP packets before sending them into an IP domain, once more due to the media being received from a TDM domain not being in the form of packets. These actions were carried out even when the SBG is acting as an IP-IP Gateway between domains. This seems to be an implementation issue where the SBG is based on an earlier IP-TDM GW and further functionality is required to enable RTP/RTCP packets to be transited when the equipment is acting as an IP-IP GW. Figure A-4 shows the two types of SBGs.

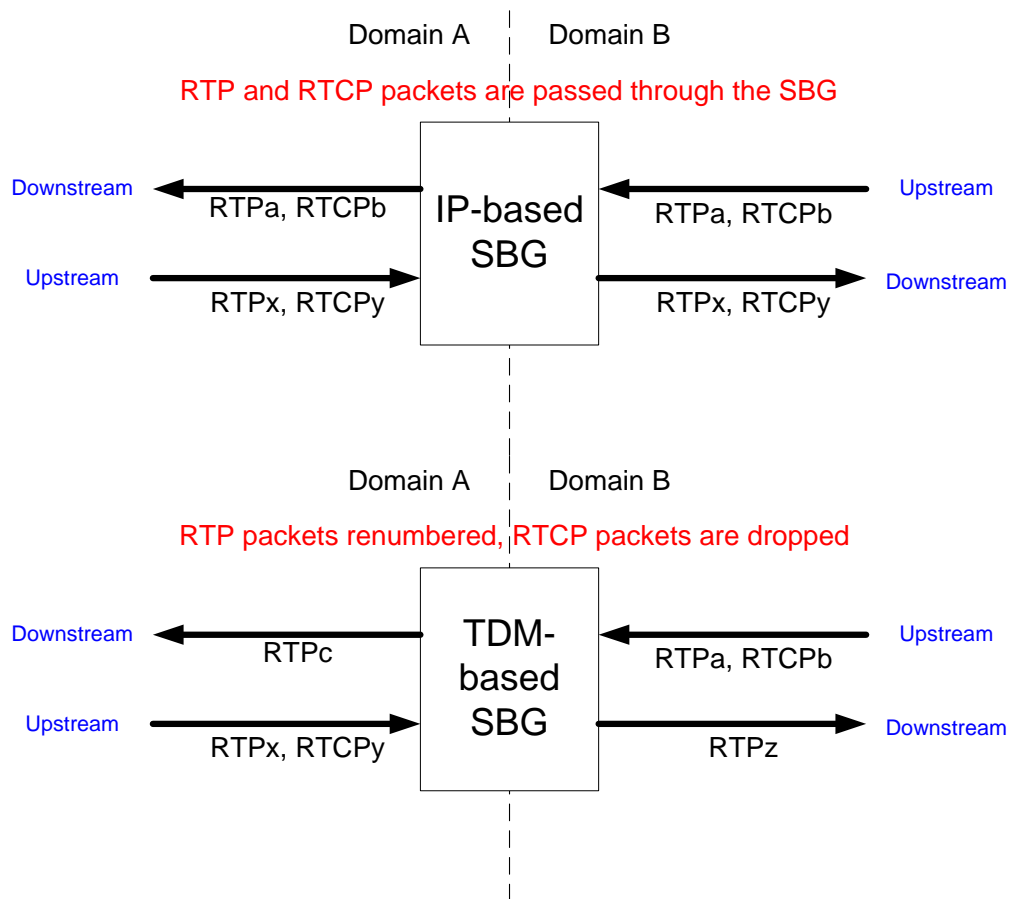


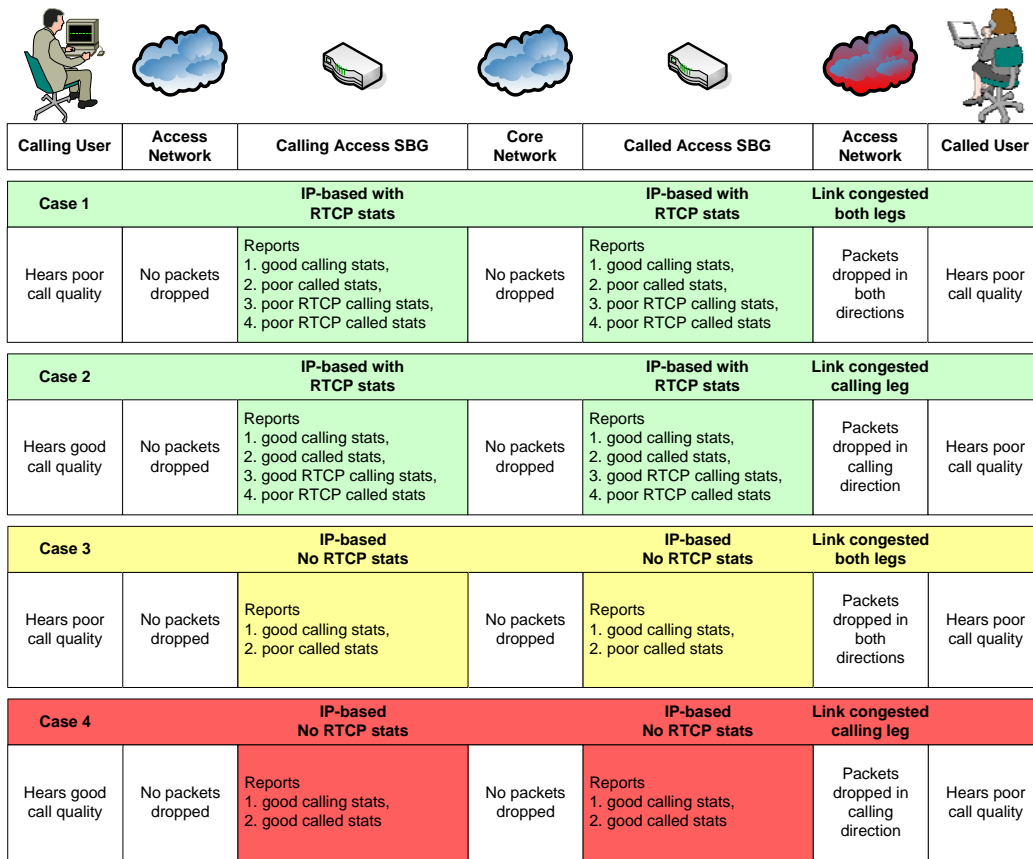
Figure A-4: SBG Types

The impact of not having RTCP statistics in a CDR can be seen in Figure A-5. In Case 1, RTCP statistics are provided. In this case, the QoS statistics from the CDR of the calling access SBC can be used to determine:

- There is no problem on the calling side ingress stream (from good calling statistics)
- There is a problem somewhere on the called side egress stream (from poor RTCP called statistics)
- There is a problem somewhere on the called side ingress stream (from poor called statistics)

In Case 2, only one direction of the session has dropped packets. In this case, the QoS statistics from the CDR of the calling access SBC can be used to determine there is a problem somewhere on the called side egress stream (from poor RTCP called stats).

In Case 3 and 4, RTCP statistics are not captured in the CDR. In Case 3, the CDR of the calling access SBC can identify there is a problem downstream on the called-to-calling media path, but cannot identify the problem on the calling-to-called media path. In Case 4, the CDR of the calling access SBC does not identify any problem, although the called party is listening to a poor quality call.



Legend

All results accurately reported
Accurate but incomplete results reported
Results reported are not accurate

Figure A-5: Accuracy of SBG CDR QoS Statistics

The impact of having TDM-based SBG in the call path can be seen in Figure A-6. Because of the packet renumbering performed by the TDM-based SBG, the CDR for the calling access SBG in Case 5 does not show any problems with the call, although both the calling and called parties are listening to a poor quality call. It should be noted that in Case 5, the TDM-based called access SBG does identify a problem with the called side ingress media stream.

In Case 6, the CDR for the calling access SBG shows no problems with the call, although the called party is listening to a poor quality call. Cases 7 and 8 are similar to Cases 3 and 4.

Scenario 1.1.2 demonstrated that there are many impediments to obtaining an accurate end-to-end view of call quality under congestion scenarios. Correlation of the CDRs from multiple devices can provide the best insight into where problems with call quality are occurring.

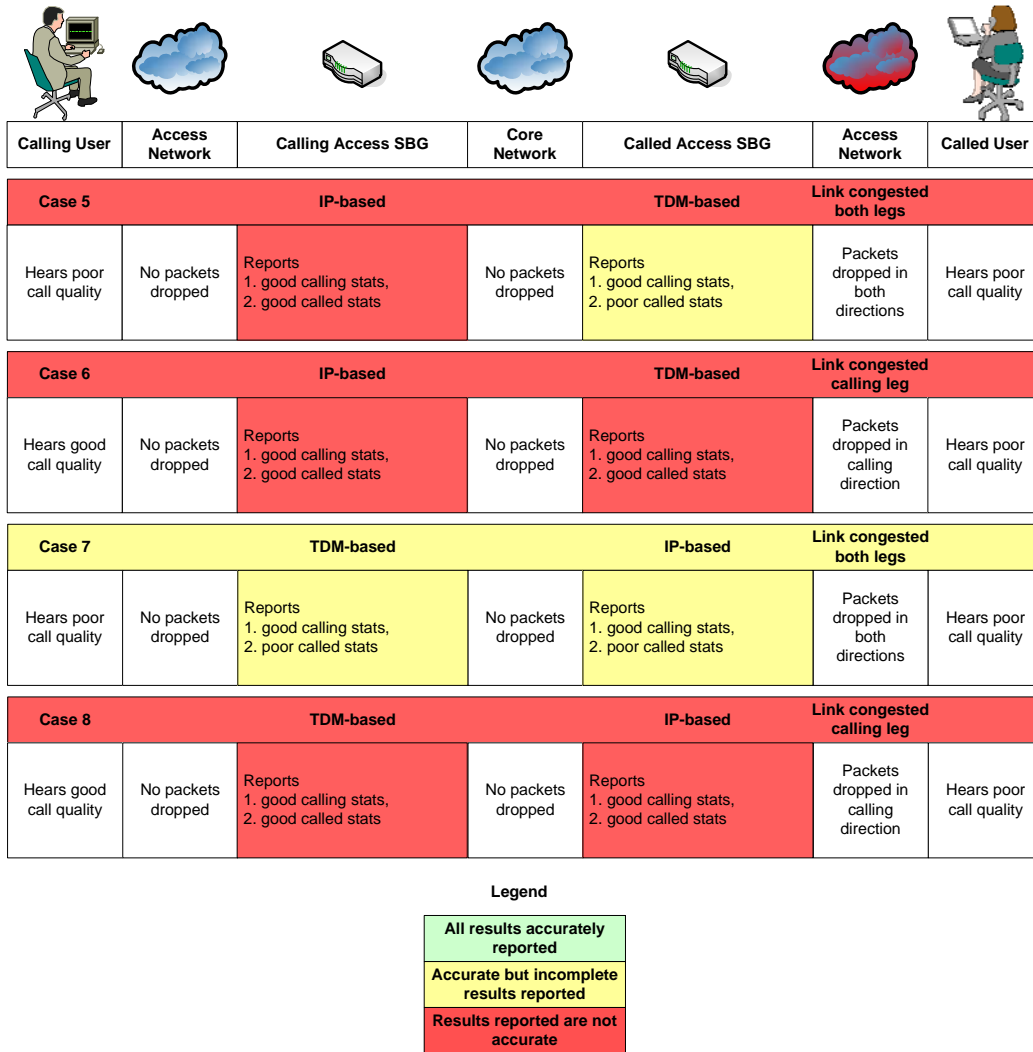


Figure A-6: Accuracy of SBG CDR QoS Statistics with TDM-Based SBG in Call Path

Scenario 1.1.3 – Rejection of new VoIP sessions when maximum session threshold hit

In this scenario, traffic generators are used to test the behaviour of SBGs when their session threshold limit is reached. Under such conditions, the SBG should ideally reject requests for normal (priority) sessions but permit priority sessions to proceed (typically within an additional threshold limit). The means by which the SBG differentiates between normal and priority sessions is also of interest in this scenario (e.g. dialled digit string, recognition of appropriate SIP headers). For Scenario 1.1.3, each SBG was tested individually (i.e., three test instances).

The following functionality is typically required for a communications device to provide support for priority communications during times of session threshold congestion:

- Recognition of a priority session via priority-related markings in the control plane signalling and applying preferential treatment to that session;
 - Exemption from network management controls,
 - Exemption from machine congestion controls except at the highest

congestion level,

- Queuing for resources not currently available and/or use of resources dedicated to priority communications.

SBGs which have QoS capabilities but no priority-communications-unique functionality can be provisioned to provide an additional likelihood of session completion to priority session requests versus normal session requests.

The results were as follows:

- The three SBGs tested all recognized dialled digits; this capability could be used to identify priority calls from normal calls. Based on this information, the priority session request could be routed to a different domain, and/or different processing and SIP INVITE manipulations could occur.
- Recognition of ETS calls via the “RPH: ets.0” parameter in a SIP INVITE was not tested. This feature was previously tested and verified in the two ETS-capable SBGs in the XTE.
- Two of the SBGs tested allowed a higher limit to be set for priority session requests than for normal session requests.
- One of the SBGs tested stopped generating delay, loss and jitter statistics in its CDRs when machine congestion control occurred.

Scenario 1.1.4 – Rejection of new VoIP sessions when maximum bearer throughput threshold hit

In this scenario, traffic generators are used to test the behaviour of SBGs when their bearer throughput threshold limit is reached. Similarly to scenario 1.1.3, the SBG ought to be able to identify and differentiate between normal and priority session requests. For Scenario 1.1.4, each SBG was tested individually (i.e., three test instances).

The results were as follows:

- None of the SBGs tested allowed a higher limit to be configured for priority sessions requests.
- One SBG did allow priority calls to exceed the maximum throughput up to a vendor-defined overage percentage.

Scenario 1.2 – Priority Video Testing

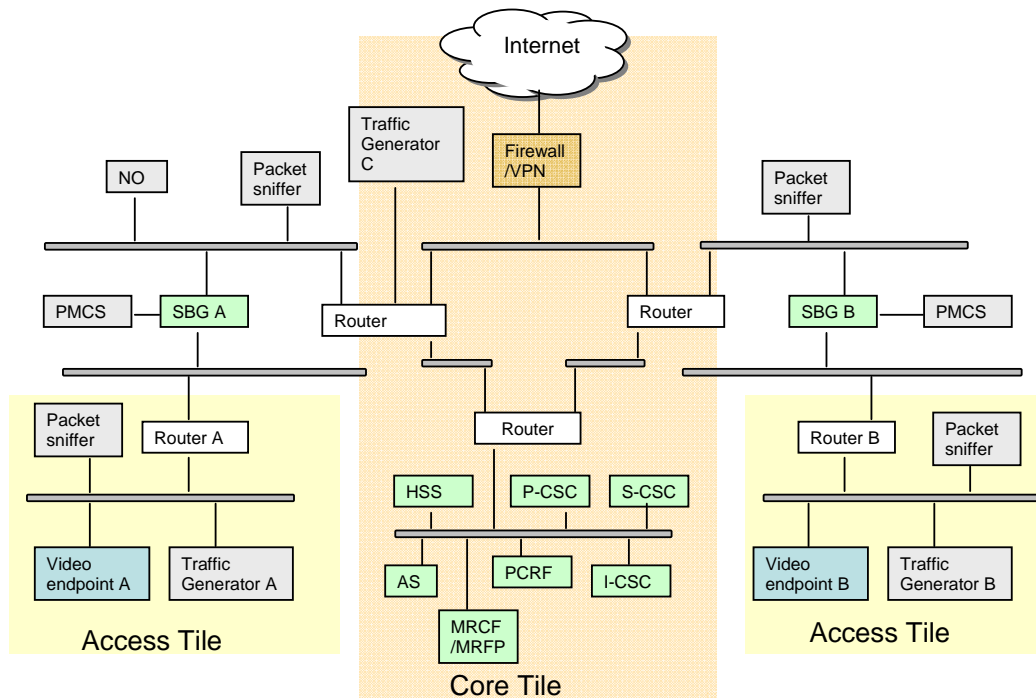


Figure A-7: Scenario 1.2 - Priority Video Testing

The network architecture for scenario 1.2 is shown in Figure A-7 above.

Scenario 1.2 Objectives

Four tests were defined in Scenario 1.2 with the objective of verifying:

- The ability of the SBG to report performance of a video teleconferencing call traversing SBG under no traffic load (S1.2.1),
- The ability of the SBG to report performance of a video teleconferencing call traversing SBG under link traffic congestion (S1.2.2),
- The ability of the SBG to reject new video teleconferencing sessions when maximum number of established sessions is exceeded. Different thresholds may be applied for normal and priority video teleconferencing sessions (S1.2.3),
- The ability of the SBG to reject new video teleconferencing sessions when threshold for maximum bearer layer throughput/bandwidth is exceeded. Different thresholds may be applied for normal and priority video sessions (S1.2.4).

Scenario 1.2 Test Results and Observations

The results for video session testing were similar to those for Scenario 1.1 with two notable exceptions:

- Of the three SBGs tested, only one captured video statistics in its CDRs. This is believed to be a consequence of many SBGs originally only supporting a single audio stream. Additional logic is required to recognize that (in the general case) multiple media streams can exist on a single session/connection and that separate measurements ought to be done per stream. The SBG that reported video statistics in its CDR also reported the RTCP video statistics it captured.
- Some SBGs did not initially support all video endpoints used in testing. In these cases, the video sessions were not established, as the SBG rejected the SIP INVITE. Upon identifying the problem to the vendors, patches were provided that allowed testing to proceed.

Scenario 1.3 – i3Forum QOS KPIs

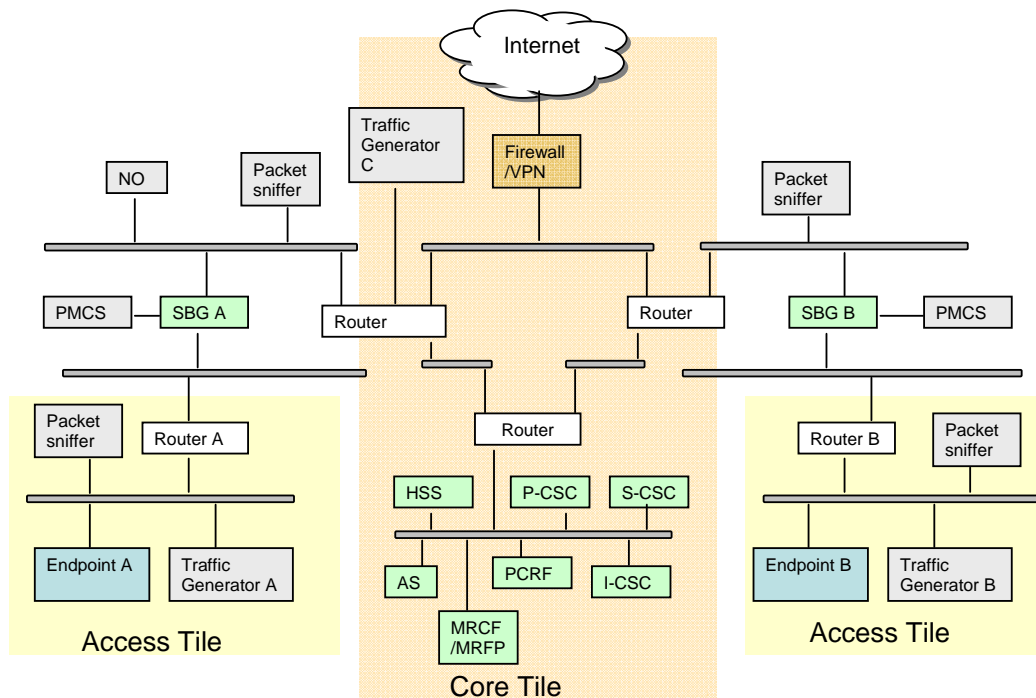


Figure A-8: Scenario 1.3 - i3Forum QOS KPIs

The network architecture for Scenario 1.3 is shown in Figure A-8 above.

Scenario 1.3 Objectives

One test was defined in Scenario 1.3 with the objective of verifying:

- The ability of the SBG to support the i3Forum QOS KPIs via information contained in its performance management reporting interface.

Scenario 1.3 Test Results and Observations

The i3Forum, in its document **Service value and process of measuring QoS KPIs (Release 1.0) May 2010** (see <http://www.i3forum.org/library>) defined the following measurements:

- Service Parameters
 - Network Efficiency Ratio (NER) – Call Establishment
 - Answer Seizures Ratio (ASR) – Called Party Answers
 - Average Length of Call (ALOC)
 - Post Gateway Ringing Delay (PGRD)
 - Mean Opinion Score (MOSCQE) / R-factor
- Call Attributes
 - Calling Line Identification (CLI) Transparency – Transmission of received CLI across network
- Transport Parameters
 - Round-Trip Delay
 - Jitter
 - Packet Loss

Scenario 1.3 reviewed the CDRs generated by the SBGs in the previous tests to determine how well the CDRs supported the i3Forum set of metrics. The MSF notes that the i3Forum is focussed on metrics on network side SBGs and this is covered in Scenario 2. However, for end-end QOS, then the same set of metrics are also applicable at access side SBGs and this is covered in Scenario 1.3.

According to the i3Forum, NER is measured by analyzing CDRs to determine the number of sent INVITEs and the SIP Release Causes (RC) for all delivered calls in a given period of time. All tested SBGs did identify the session release reason in their CDRs; however, a vendor-specified release code is likely to be used instead of the SIP RC. The vendor-specified release code provides more granularity on what occurred in the SBG, enabling troubleshooting. The vendor-specified release code is mapped to a SIP RC for transmission to other devices in the network.

The CDRs of all tested SBGs also supported derivation of the ASR, ALOC and PGRD metrics. The CDRs did not support the CLI metric. However, it is noted that the CLI metric is concerned with CLI transparency across a network and thus not dependent on any local measurement on the SBG.

The CDRs of all tested SBG contained delay, jitter and packet loss information. However, as was noted in Scenario 1.1.2, this information may not be complete (e.g., RTCP data may not be collected) nor accurate (e.g., packet loss may be filtered by an upstream TDM-based SBG).

The i3Forum notes that in calculating the MOSCQE / R-factor metric:

A carrier may obtain a MOS measurement for every call, using information generated by Session Border Controllers and/or Call Handling Function when the RTCP is activated; these measurements can be statistically represented with reference to a defined period of time, source-destination path, and can be specific for a given provider.

The measurement could be widely affected by the availability of RTCP along the RTP flow and by the location of the SIP end-point (e.g. a media gateway, an end user SIP terminal, a codec translation etc.) which can not be dynamically determined. This implies the measurement should be considered meaningful only in a complete Voice over IP environment (no TDM-IP translation) and in a tested and verified voice path. It is also to be

noted that the voice path is determined by the location and routing decision of the voice equipment: it has to be considered normal that the voice path can be different and with a longer distance than the IP best effort path. It is to be noted that it is not possible to detect on a per call basis if a carrier blocks the RTCP flow, or when the call is trans-coded or passed onto TDM to reach its destination. Therefore, it is important to use the MOS on routes that have been initially tested between partners in a trusted commercial environment.

MOSCQE / R-factor can be calculated using the delay, jitter and packet loss information found in the CDRs. However, as noted above, this information may not be complete or accurate.

Scenario 2 – Multi- Domain Performance Management Reporting

Scenario 2 concentrated on the performance management reporting capabilities of a network side SBG in a multi-IP domain environment under varying load conditions for voice, video and data sessions. The provisioning interface on the SBG for setting up IPsec tunnels was also tested as was the ability of the SBG to support the i3Forum defined QOS KPIs. The test configuration for Scenario 2 is shown in figure A-9 below:

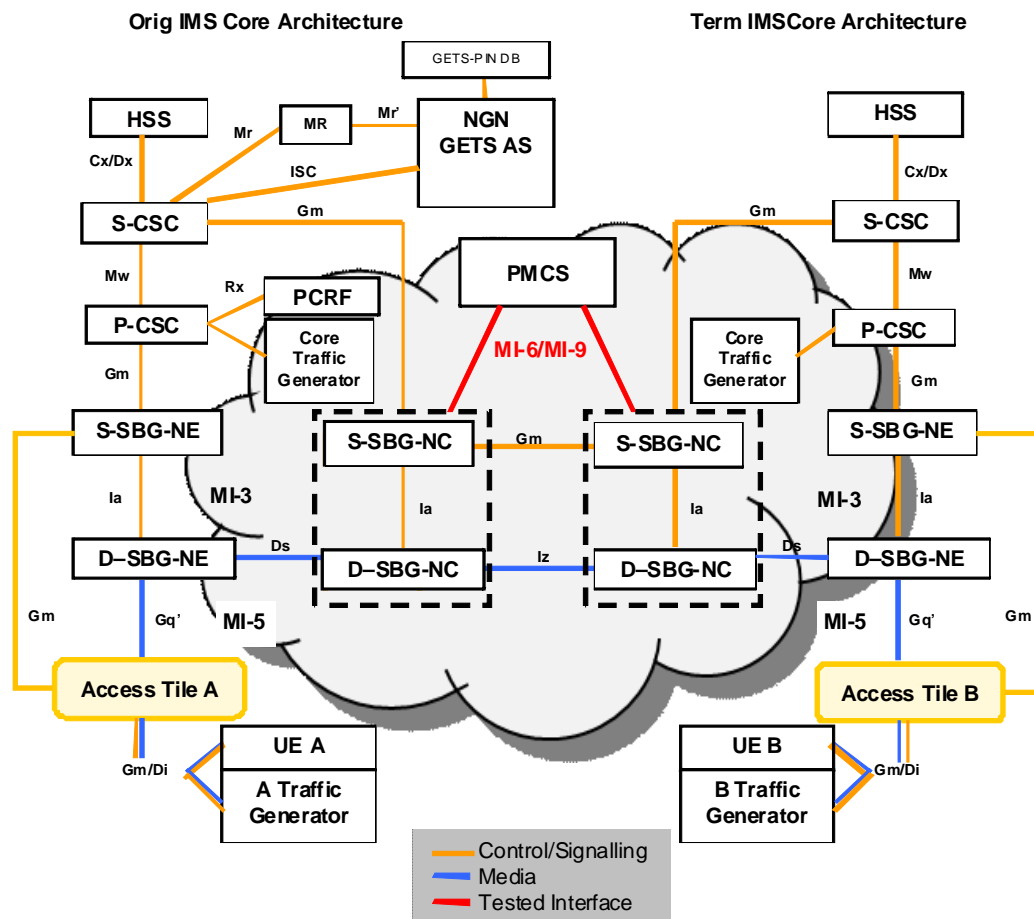


Figure A-9: Test Configuration for Scenario 2

This scenario was broken down into a number of sub-scenarios that are detailed below.

Scenario 2.1 – NNI Provisioning on the SBG

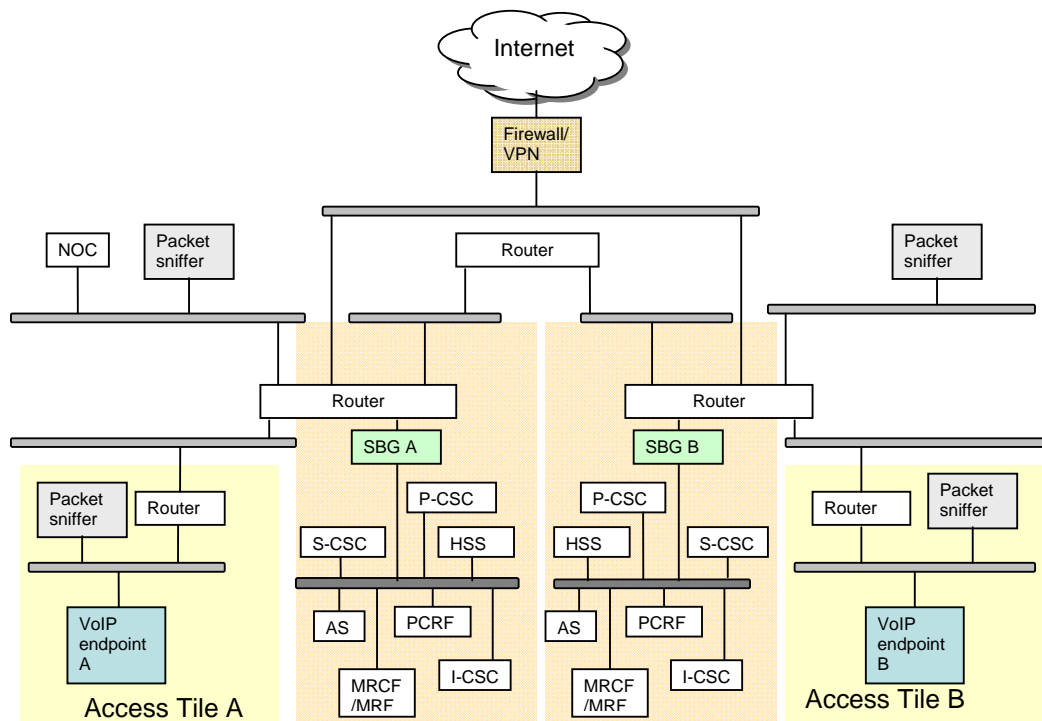


Figure A-10: Scenario 2.1 - NNI Provisioning

The network architecture for Scenario 2.1 is shown in Figure A-10 above.

Scenario 2.1 Objectives

One test was defined in scenario 2.1 with the objective of verifying:

- The ability of the SBG to set up an encrypted IPSEC tunnel to its peer node. Both normal and priority (ETS) traffic is sent over the same tunnel. Both IMS cores are provisioned to recognise the SIP RPH header. If a “REQUIRE: RPH” parameter is present in a SIP INVITE, then the SBG strips it off prior to sending the INVITE over the IPSEC tunnel. The “REQUIRE: RPH” parameter should be restored by the receiving SBG.

Scenario 2.1 Test Results and Observations

The XTE staff encountered errors while provisioning the “incumbent” SBGs at the XTE. For one SBG, the first error message indicated that additional licenses were required to enable IPsec. After obtaining the needed licenses, a second error indicated that additional hardware in the form of daughter-boards was needed to enable IPsec. The additional hardware was ordered but not delivered during the testing period. The second SBG required additional ports in order to support IPsec. This hardware could not be obtained in time for testing. These hardware issues prevented IPsec testing between different vendors’ SBGs during the IOT.

The Genband high availability SBG was divided into two separate SBGs to perform IPsec testing.

The following observations were made during the provisioning of SBGs for IPsec tunnels:

- One SBG only supported IPsec tunnels on their signalling interface, while the

others supported IPsec tunnels on both signalling and media.

- The provisioning of IPsec was very complex, with multiple IPsec-identified parameters which could interact with features provisioned elsewhere.
- The “default” parameters vendors use to assist in setting up IPsec tunnels are not consistent, making provisioning for interoperability between products cumbersome. The difficulties observed in establishing IPsec tunnels illustrate the importance of profiles for specifying all parameters going across an NNI. One example of an activity to generate such a profile is the ongoing work in the ATIS NG-CI task force
- All of the tested SBGs supported the stripping and restoral of the “REQUIRE: RPH” parameter via SIP manipulation rules associated with domains supporting the NNI. This is important since the stripping and restoral of this parameter is a relatively new requirement that was not in the original ETS requirements. Use of the SIP manipulation rules means that carriers do not need to wait for a vendor “patch” to implement this requirement.

Scenario 2.2 – SBG NNI Testing

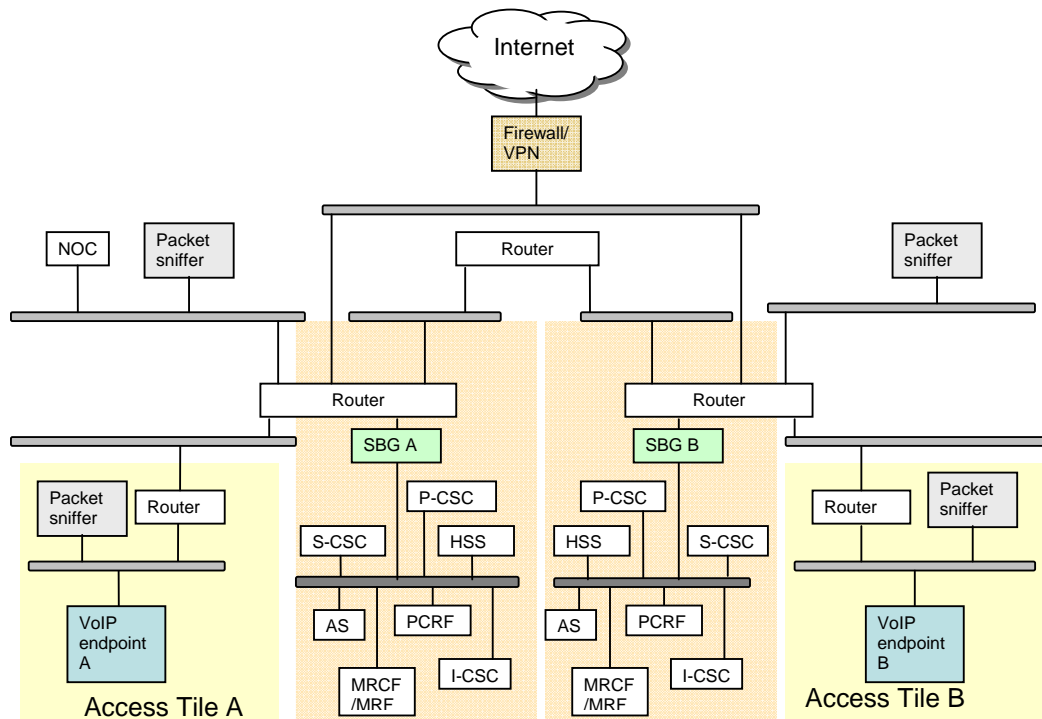


Figure A-11: Scenario 2.2 - SBG NNI Testing

The network architecture for scenario 2.2 is shown in Figure A-11 above.

Scenario 2.2 Objectives

Four tests were defined in scenario 2.2 with the objective of verifying:

- SBG NNI Voice and Video Performance Reporting Under “No Traffic Load” (S2.2.1)
- SBG NNI Voice and Video Performance Reporting For IPsec Tunnels Where the Link Supporting the IPsec Tunnel is Congested (S2.2.2)

- SBG NNI Voice and Video Session Establishment Congestion (S2.2.3)
- SBG NNI Voice and Video Bearer Throughput Congestion (S2.2.4)

Scenario 2.2 Test Results and Observations

The Scenario 2.2 tests repeated the experiments of Scenarios 1.1 and 1.2 across the NNI. The results were consistent with those seen when the SBG was functioning as an access SBG.

Scenario 2.3 – Priority Data Testing

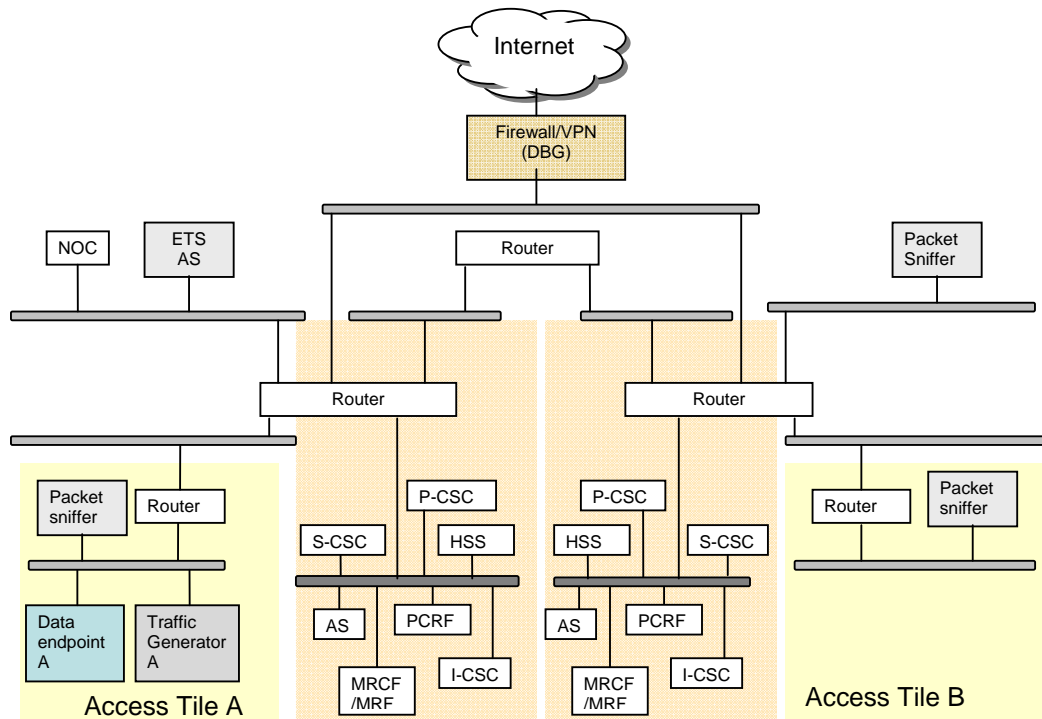


Figure A-12: Scenario 2.3 - Priority Data Testing

The network architecture for Scenario 2.3 is shown in Figure A-12 above.

Scenario 2.3 Objectives

Three tests were defined in Scenario 2.3 with the objective of verifying:

- Dynamic Provisioning of the Data Border Gateway (DBG) to Support Priority Data (S2.3.1)
- DBG Performance Reporting Under “No Traffic Load” (S2.3.2)
- DBG Performance Reporting Under Internet Traffic Congestion (S2.3.3)

The single “incumbent” XTE firewall was used in Scenario 2.3.

Scenario 2.3 Test Results and Observations

Of concern to the MSF is how non-SIP-based priority communications (e.g. HTTP) can be given priority treatment during network congestion. There is currently no secure protocol mechanism to indicate that specific non-SIP traffic should receive priority treatment. The current approach being taken to address this for the North American market is to establish a secure (HTTPS) session with an authentication server which then triggers the policy infrastructure to apply priority treatment and traffic shaping using normal policy rules and mechanisms. Although this approach is not explicitly standardised, it can be implemented today because it makes use of existing standardised policy interfaces to provide the required functionality.

The ATIS PTSC has defined Edge functions within its architecture. The Edge functions are used for media and traffic processing when aggregated traffic coming from different access networks is merged; they include functions related to support

for QoS and traffic control, and they support both session-based and non-session-based communications.

Edge functions can include detection and enforcement capabilities associated with Deep Packet Inspection, to achieve the following tasks:

- Detection
 - Service / flow identification
 - Statistics gathering
 - Gathering charging information
- Enforcement
 - Service tiering
 - Traffic shaping
 - Security
 - Filtering

Deep Packet Inspection in the Edge function can be integrated with other functions, or a standalone functional entity. The Edge functions are provided by an Access Border Gateway Function (A-BGF) in the PTSC architecture. The A-BGF is defined as a packet gateway between an access network and a core network that is used to mask a service provider's network from access networks, through which UE accesses packet-based services (e.g., IMS, Internet).⁴ The functions of the A-BGF have much in common with a SBG and may include:

- Opening and closing gate;
- Packet filtering-based firewall (including deep packet inspection capabilities);
- Traffic classification and marking;
- Traffic policing and shaping;
- Network address and port translation;
- Media Relay (i.e., media latching) for NAT traversal; and
- Collecting and reporting resource usage information (e.g., start-time, end-time, octets of sent data).

While the firewall available in the XTE could be provisioned to support most of the functions of an A-BGF, the statistics reported were in the form of operational measurements for the device as a whole and not on a per flow basis. The additional functionality required in the A-BGF is similar to that required for DBGs. Two areas of effort are seen as being needed to make DBGs support priority communications:

- While the monitoring and “deep packet inspection” capabilities are sufficient to support priority communications, current policy mechanisms tend to be static and the policy interfaces tend to be vendor specific. To support priority communications, the DBG needs to be able to accept and incorporate “real-time” policy changes from a network entity like a Policy and Charging Rules

⁴ The DBG was originally defined to protect the Service Provider network from the Internet. An MSF-member carrier representative at the P-IOT event observed that a DBG function could also be used to protect the Service provider network from the customer equipment. During crises like a pandemic, access networks are more likely to be congested than core network, due to telecommuting.

Function (PCRf) using a standard interface.

- Statistics need to be captured for specific priority media flows. The need to capture these statistics would likely be defined by the policy associated with the specific flow. Additional discussions are needed with industry to determine what statistics should be collected.

Scenario 2.4 – i3Forum QOS KPIs

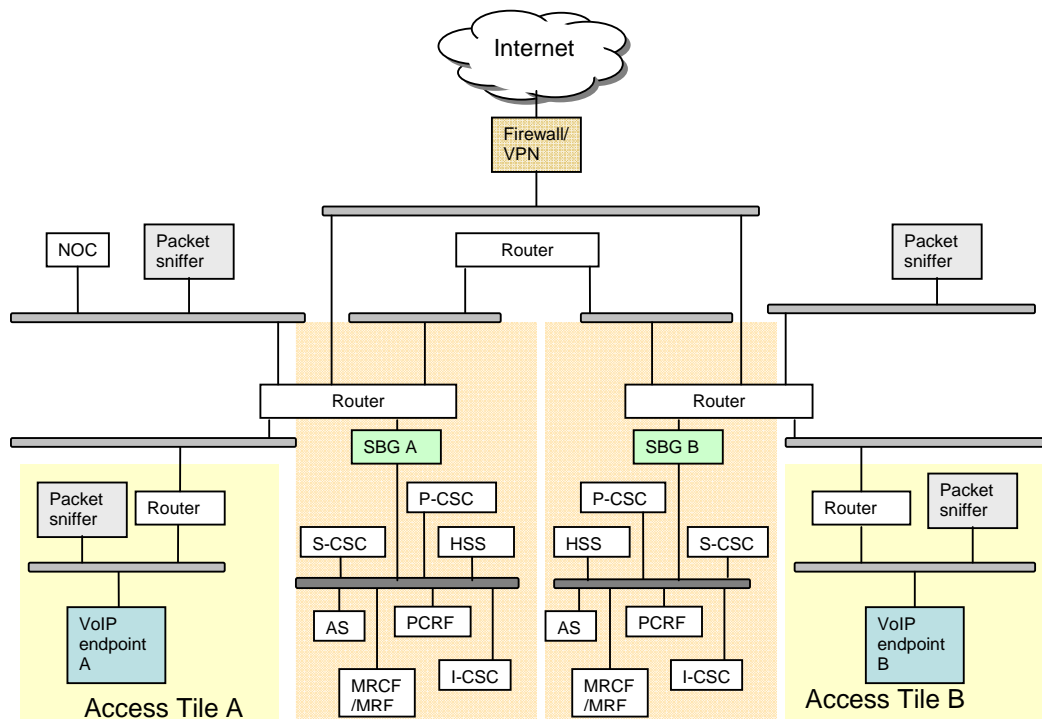


Figure A-13: Scenario 2.4 - i3Forum QOS KPIs

The network architecture for Scenario 2.4 is shown in Figure A-13 above.

Scenario 2.4 Objectives

One test was defined in Scenario 2.4 with the objective of verifying:

- The ability of the SBG to support the i3Forum QOS KPIs via information contained in its performance management reporting interface.

Scenario 2.4 Test Results and Observations

The i3Forum metrics supported by SBGs acting as network-side SBGs are identical to those observed in Scenario 1.3. .

Appendix B: The Benefits of MSF Membership

The MSF is a global association of service providers, system suppliers, test equipment vendors, and users committed to developing and promoting open-architecture, multiservice Next Generation Networks. Founded in 1998, the MSF is an open-membership organization whose members are drawn from the world's leading IP communications companies. Historically, the MSF developed Implementation Agreements (IAs), promoting worldwide compatibility and interoperability of network elements, and encouraging input to appropriate national and international standards bodies. However, this role has become less of a focus as a standard architecture has emerged and protocol profiles are being defined in other SDOs and industry fora (e.g. 3GPP). The MSF primarily focuses on running IOT events to facilitate the deployment of next-generation multi-service networks, driven by its member needs. In addition, in 2009, the MSF decided to complement its established biennial Global MSF Interoperability (GMI) events by more targeted test events focused on critical and timely issues associated with deployment of IP communications. The P-IOT 2010 Event is the second test event of this nature. For a history of MSF IOT events as well as information on planned upcoming IOT events, see <http://www.msforum.org/interoperability/GMI.shtml>.

The advantages of MSF membership include:

- Access to more than ten years of groundbreaking industry work with input from key service providers and vendors,
- The experience of some of the world's leading scientists and engineers,
- The opportunity to leverage the external talent pool active in the MSF to more efficiently implement a validated architecture built on industry-standard protocols,
- The ability to validate product implementations in industry-leading interoperability events.

In addition, service providers and equipment vendors that actively participate in MSF IOT events learn how multivendor next-generation products and networks will interoperate in the real world. That information translates into several financial benefits:

- Reduced time to market for deployment of interoperable solutions,
- Decreased costs and resources to resolve interoperability issues,
- Improved protocol documentation through facilitating clarifications in the tested standards via feedback to the appropriate SDOs ,
- Thoroughly evaluated architectural framework for cooperatively designing end-to-end networking solutions.

The MSF's IOT program is designed to facilitate implementation of NGNs and to deliver the Forum's mission statement that "We make Next Generation Networks work". MSF IOTs validate products in the latest standards-based architectural framework using network deployment scenarios that are meaningful to Service Providers.

Appendix C: The Participants

This appendix provides a brief resume of the participant organizations in the P-IOT 2010 event.

National Communications System (NCS)

The National Communications System (NCS) is an office within the [United States Department of Homeland Security](#) charged with enabling national security and emergency preparedness communications ([NS/EP telecommunications](#)) using the national telecommunications system. In fulfillment of its mission, the NCS has created a number of different services:

- [Government Emergency Telecommunications Service \(GETS\)](#) - provides emergency access and priority processing in the local and long distance segments of the public switched wireline network. GETS is used in an emergency or crisis situation during which the probability of completing a call over normal or other alternate telecommunication means has significantly decreased.
- **Telecommunications Service Priority (TSP)** - provides service vendors with a Federal Communications Commission (FCC) mandate for prioritizing service requests by identifying those services critical to NS/EP. A telecommunications service with a TSP assignment is assured of receiving full attention by the service vendor before a non-TSP service.
- [Wireless Priority Service \(WPS\)](#) - provides priority cellular network access. The WPS was approved by the FCC for NS/EP requirements on a call-by-call priority basis.

The focus of the NCS in the MSF is to demonstrate and test priority communications during IOT events. Previous GMIs have focused on component functionality; at these events, the NCS demonstrated the practicality of providing priority communications through an IMS architecture.

Since priority communications is critical during network congestion events, the NCS was the major proponent of the network robustness work item in the MSF.

The NCS operates an eXperimental Testbed Environment (XTE) to prototype and assess priority communications mechanisms that may be deployed in public NGNs to support priority voice, video and data communications.

About Genband

GENBAND is a global leader of IP infrastructure and application solutions, enabling fixed, mobile and cable service providers around the world to evolve communications networks through IP innovation. The Company offers market-leading Switching, Applications, Networking and Service solutions, with products deployed in over 600 customer networks spanning more than 80 countries. GENBAND is headquartered in Frisco, Texas, and has operations in 50 countries. To learn more, visit us on the web at www.genband.com.