



***Society of Cable
Telecommunications
Engineers***

**ENGINEERING COMMITTEE
Network Operations Subcommittee**

AMERICAN NATIONAL STANDARD

ANSI/SCTE 175 2017

Multimedia Management (MMM)

**Recommended Practice for Qualifying Network Devices
for High Availability Streaming Video**

NOTICE

The Society of Cable Telecommunications Engineers (SCTE) Standards and Operational Practices (hereafter called “documents”) are intended to serve the public interest by providing specifications, test methods and procedures that promote uniformity of product, interchangeability, best practices and ultimately the long term reliability of broadband communications facilities. These documents shall not in any way preclude any member or non-member of SCTE from manufacturing or selling products not conforming to such documents, nor shall the existence of such standards preclude their voluntary use by those other than SCTE members.

SCTE assumes no obligations or liability whatsoever to any party who may adopt the documents. Such adopting party assumes all risks associated with adoption of these documents, and accepts full responsibility for any damage and/or claims arising from the adoption of such documents.

Attention is called to the possibility that implementation of this document may require the use of subject matter covered by patent rights. By publication of this document, no position is taken with respect to the existence or validity of any patent rights in connection therewith. SCTE shall not be responsible for identifying patents for which a license may be required or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

Patent holders who believe that they hold patents which are essential to the implementation of this document have been requested to provide information about those patents and any related licensing terms and conditions. Any such declarations made before or after publication of this document are available on the SCTE web site at <http://www.scte.org>.

All Rights Reserved

© Society of Cable Telecommunications Engineers, Inc. 2017
140 Philips Road
Exton, PA 19341

Table of Contents

1. Introduction.....	4
2. Scope.....	5
3. Informative References.....	5
4. Glossary, Acronyms, and Abbreviations	7
5. Introduction.....	8
5.1 Availability Metrics and Measurements	8
5.2 Availability Reports.....	9
5.3 Pass/Fail Criteria.....	10
6. Test Configurations:.....	10
6.1 Baseline Tests:	12
6.1.1 Test 1: Headend Broadcast Baseline	13
6.1.2 Test 2: Core Network Switch/Router Baseline Operation	15
6.1.3 Test 3: Edge/Hub Network Switch/Router Baseline Operation.....	17
6.1.4 Test 4: Edge/Hub Network Switch/Router Baseline Multicast Operation.....	20
7. Summary.....	23

1. Introduction

This document is identical to SCTE 175 2011 except for informative components which may have been updated such as the title page, NOTICE text, headers and footers. No normative changes have been made to this document.

Given the rapidly developing competitive market environment and subscriber sensitivity to program stream impairments, network devices in these applications are required, ideally, to operate without errors of any kind – *ever*. Or, should an error be unavoidable, it ought to be short and, therefore, minimize its impact to subscribers. IP Network devices including data switches and routers are increasingly being deployed in the transport of large numbers of flows of streaming video over IP in MSO transport environments. Program streams have somewhat different characteristics and transport requirements than VoIP and data. MSO Operations Groups now use measurements of per-program *Availability* in evaluating the operation of deployed equipment. Thus, availability Metrics are being deployed by major video service providers to measure and monitor the end-to-end per-program systems performance in their plants. Achieving low numbers of per-program errored seconds to the subscriber through monitoring and management techniques is also the subject of recent SCTE standards such as SCTE 168-6 2010. This RP is intended to specify means to evaluate whether a network device is capable of delivering a desired per-program availability performance level; i.e., 3 nines (99.9%), 4 nines (99.99%), etc. The testing and evaluation specified in this document is intended to be applied during the product evaluation and selection phase before deployment. Comparing test results from different products will help guide the selection of products best suited to delivering high availability streaming media over IP networks typical of MSO deployments.

Because of the differences in video over IP streams' characteristics compared to other payloads such as voice and data, using traditional data application oriented tests exclusively may not reflect the network performance actually experienced when used in streaming media applications. With an emphasis on deriving longer term program availability results, this note proposes and describes the following plan for network device evaluation:

- Identify a set of tests that describe common traffic characteristics and flow paths to be used during a switch and/or network device evaluation intended for streaming media based on common specific configurations found in many MSO streaming media transport networks as described in the ANSI/SCTE 168-7 2010 Recommended Practice.
- Identify realistic, repeatable traffic load test conditions and typical network stress events that will characterize the availability performance of network devices intended for video and streaming media transport.

- Identify test results in terms of program availability to help evaluate expected subscriber results and satisfaction when the DUT is deployed in streaming media networks. These results are also intended to provide a common means of comparing devices that perform similar functions in terms of how subscribers may be affected in the presence of common network operations and stress conditions.

These tests are intended to complement existing data oriented device performance evaluations such as described in RFC 2544, RFC 2889, etc. They include traffic loads with test generated video formatted streams or live streams originating at a service provider.

2. Scope

The ANSI/SCTE 168 series of Recommended Practices [17][18][19] describe IP video networks at MSO Headend, Core, and Hub networks. While all these configurations carry media over IP, the video flow types and distribution, link speeds, and possibly QoS policies are different at various locations. Different mixes of traffic types such as VoIP and data may be present in some locations and not in others. The recommended baseline tests in this document are intended to represent the operation of network devices in these three applications. These baseline tests provide a common set of reference results that can be used to preliminarily screen equipment and configurations during the selection, configuration, and deployment process. They will provide results to determine if the equipment under consideration provides a level of performance that is “in the ballpark” of that intended for an application. The test results indicate a per-program availability for extended “soak” tests of long term delivery of a number of realistic simultaneous video flows. They will also indicate how program availability will be affected by certain types of common realistic network events often experienced in operating complex live networks. Devices that take longer to recover from faults, that cannot accommodate simultaneously arriving persistent packet stream packets, experience crashes, or occasionally lose data for whatever reasons will fare less well in these tests compared to designs and configurations that are optimized for high availability.

Other tests should also be executed to complement these baseline tests before final equipment selection and deployment. Such tests would typically include, but not be limited to, the intended load levels, number-type-and speed of active ports, level and type of non-video converged traffic expected, forwarding protocols, management protocols, etc. that are expected in the operational environment.

3. Informative References

[1] “A Proposed Media Delivery Index”, RFC 4445

[2] “High Availability Computer Systems”, J. Gray and D. P. Siewiorek, IEEE Computer 24, pp. 39-48, Sept. 1991.

[3] “Building a Highly-Available Enterprise Network with Juniper Networks EX-series Switches”

[4] “High Availability Network Architectures for Triple Play Services, An Overview of RESIP Core, Edge, and Access Design Guidelines”, A white paper issued by: Siemens Communications and Juniper Networks, Inc.

[5] “Cisco – Baseline Process: Best Practices White Paper”, Document ID: 15112, Updated: Oct 04, 2005.

[6] “Cisco IOS Management for High Availability Networking: Best Practices White Paper”, Document ID: 25562, Updated: Jan 10, 2006.

[7] “Cisco IOS High Availability Curbs Downtime with Faster Reloads and Upgrades”, White Paper, 2004

[8] “Triple-play Services Quality of Experience (QoE) Requirements”, DSL Forum Technical Report TR-126

[9] “Recommended Practices for Monitoring Multimedia Distribution Quality”, ANSI/SCTE 168-6 2010

[10] ITU-T recommendation G.826 (2002), End-to-end error performance parameters and objectives for international, constant bit rate digital paths and connections

[11] ITU-T recommendation I.357 (2000), B-ISDN semi-permanent connection availability

[12] ITU-T recommendation Y.1540 (2007), Internet protocol data communication service – IP packet transfer and availability performance parameters

[13] ITU-T recommendation Y.1561 (March 2004), Performance and availability parameters for MPLS networks

[14] ITU-T recommendation Y.1731 (May 2006), OAM functions and mechanisms for Ethernet based networks

[15] ITU-T recommendation G.7710/Y.1701 (June 2007), Common equipment management function requirements

[16] ITU-T SG12 contribution C.116 (October 2007), Comparison of availability definitions for Y.ETHperf

[17] ANSI/SCTE 168-4 2010, Recommended Practice for Transport Stream Verification Metrics

[18] ANSI/SCTE 168-6 2010, Recommended Practice for Monitoring Multimedia Distribution Quality

[19] ANSI/SCTE 168-7 2010, Recommended Practice for Transport Stream Verification in an IP Transport Network

[20] ISO/IEC 13818-1 (2000-12-01), Information technology – Generic coding of moving pictures and associated audio information: Systems

4. Glossary, Acronyms, and Abbreviations

Availability – The number of unimpaired seconds of a multimedia stream divided by the number of tested seconds expressed as a percentage.

CBR – Constant Bit Rate

DUT – Device Under Test

Errored Second – A second during which a measured stream has one or more missing packets.

HD – High Definition; An MPEG 2 encoded HD stream is usually transmitted at greater than 10 Mb/s

IP – Internet Protocol

Link Flap – A network interface that repeatedly becomes enabled and disabled due to a layer 1 fault or other reasons

Multicast -- A method for transmitting a single source stream to many receivers. Multicast uses network devices to replicate packets that are sent to multiple receivers

Packet – an IP Datagram as defined in RFC 791

Program – A program is a collection of program elements. Program elements may be elementary streams. Program elements need not have any defined time base; those that do have a common time base are intended for synchronized presentation. See ISO 13818-1.

Program Element – A generic term for one of the elementary streams or other data streams that may be included in a program. See ISO 13818-1.

Program Stream -- a succession of packets (IP Datagrams) carrying multimedia content and related data for that content.

SD – Standard Definition; An MPEG 2 encoded SD stream is usually transmitted at less than 4 Mb/s

Stream – A steady flow or succession of packets (IP Datagrams) addressed to the same destination.

Unicast -- Transmission of a stream from one source to a single receiving destination

VLAN – Virtual LAN

5. Introduction

5.1 Availability Metrics and Measurements

Availability/unavailability status is defined in various standards in different ways depending on the type of network. Traditionally, entering the *unavailable* state occurs when the performance of a service is *highly degraded*. Using this definition, if the service is *slightly degraded*, the service is identified as *available* but with *degraded performance*. Sometimes the degradation must be completely removed before reentering the *available* state. Further, the criteria to enter a *degraded* or *unavailable* state may require the persistence of a degraded condition for a specified number of seconds. Likewise, the *degraded* or *unavailable* state may require the persistence of no degradation for a specified number of seconds and the durations to enter and leave the impaired states may not be equal. For availability definitions for SDH networks see ITU-T G.826 which describes a period of *unavailable* time at the onset of ten consecutive severely errored seconds. For ATM networks see ITU-T I.357 where *availability* is defined as the fraction of time a transaction can be supported. For MPLS networks see ITU-T Y.1561 where *unavailability* begins with the occurrence of ten consecutive *severe loss* blocks. For general IP networks see ITU-T Y.1540 which bases *availability* on a threshold of IPLR (IP Loss Ratio) performance, and for general Ethernet networks see MEF 10.1 and ITU-T G.7710 where *unavailability* is reported by Spanning Tree Protocol convergence status or OAM flows to detect losses of continuity.

However, video service networks are different. As described in other references [8], [18], even a single lost packet or loss second can cause a user perceptible video and/or audio impairment. An errored second – a second during which one or more program errors occur – may well be considered an unacceptable highly degraded condition by a subscriber. With this criteria, a program is considered as entering an unavailable state for *any* errored seconds. This approach is recommended for video service networks.

An errored second is any second that includes one or more lost packets. This approach is recommended for video service networks. A user *may* consider an errored second as one in which other stream characteristics exceed a preset threshold or duration such as out of order packets, duplicate packets, unacceptable packet jitter, or other stream impairment

but should identify the test conditions along with the test results if errored seconds are counted based on metrics other than packet loss.

Other less demanding policies that would subject subscribers to potentially more severe and/or more frequent impairments before unavailability is defined would be to set a specified duration of error persistence in seconds before considering the program to be in the unavailable state and, perhaps a different duration of no errors before considering the program to again move to the available state. An errored second might also be defined such that a minimum count of packet loss events (loss period length) must occur before being tallied as an errored second.

In any case, whatever definitions and policies are adopted by a provider should be simple for these baseline tests in order to facilitate determining whether equipment under consideration meets the criteria as well as for actual monitoring and evaluation.

5.2 Availability Reports

Availability reporting is intended to reflect the user acceptability of delivered programs and indicates the availability of “good” program time. Availability reports indicate the number of unimpaired per-program seconds delivered by a device or system under test as a percent of evaluated seconds:

$$\% \text{ Availability} = ([\text{measured time interval} - \text{unavailable time}] / \text{measured time interval}) \times 100$$

For example, for 2 impaired seconds (unavailable time) in a 1 day period (measured time interval):

$$\% \text{ Availability} = ([86400 - 2] / 86400) \times 100 = 99.998\%$$

An example report would typically include information shown below (Figure 1) including program name, measurement location if relevant, and percent availability. The tested configuration along with specifications about how the availability is calculated should also be included.

	Data Collection Point					
	Core_A	Core_A Egress Hub_1	Core_A Egress Hub_2	Core_B	Core_B Egress Hub_1	Core_B Egress Hub_2
	Measures	Measures	Measures	Measures	Measures	Measures
Program	Availability	Availability	Availability	Availability	Availability	Availability
(A&E (37))	99.442%	99.443%	99.443%	99.443%	99.269%	99.443%
(ABCFMLY (26))	99.453%	99.454%	99.454%	99.454%	99.280%	99.454%
(ACTNMAX (343))	99.443%	99.444%	99.444%	99.444%	99.444%	99.270%
(AMC (59))	99.443%	99.443%	99.444%	99.444%	99.444%	99.444%
(ANIMLPL (63))	98.967%	99.181%	99.181%	99.179%	98.925%	99.179%
(BIOGRPHY (243))	99.442%	99.443%	99.443%	99.443%	99.443%	99.269%

Figure 1.

5.3 Pass/Fail Criteria

The ANSI/SCTE 168-6 2010 Recommended Practice suggests an acceptable per-program performance criteria of 1 errored second per 4 Hours for HD and 1 errored second per Hour for SD with an equivalent minimum availability of 99.993% and 99.972% per day respectively. Different systems may have different targets based on a variety of considerations such as available CapEx and OpEx resources, age of plant and environment, etc. Refer to your corporate policy for your specific system requirements. Evaluate if the measured results of the tests described in this note will meet deployed system goals. Some service providers are targeting Five 9s or 99.999% availability for their systems currently. Note that each deployed component must have higher availability than the target availability for the component ensemble. For example, if ten devices are connected in series and each has 99.99% availability, the ensemble would have 0.9999^{*10} or 99.90% availability.

If the tests are completed in a lab environment, the results may be better than a field deployed system subject to more unpredictable sources, physical interconnect stresses, environmental condition stresses such as temperature and humidity, power line transients, human errors, etc. This may cause the tester to set the desired availability results as measured in these tests somewhat higher than might otherwise be considered.

If the baseline tests' functionality cannot be met by a device or system being tested (for example, the specified nominal loads cause immediate and continuous packet loss or other regular, repeatable delivery failures), the device design may not be suitable for service provider networks intended to carry program streams. Alternatively, if the baseline tests are more demanding than the intended application for the device, then the baseline tests can be de-rated. For example, the number of simultaneous streams during the test may be reduced or the stream type mix changed. If this approach is taken, one should not attempt to use the device in an application with higher numbers of streams or types and expect availability results produced with the lower loads. The availability results reports must also show the de-rated test conditions so that accurate performance comparisons can be made between devices.

If the baseline tests complete properly but do not achieve the desired availability performance levels, the device or system might still be deployed though with a more appropriate expectation for performance level or another system may be chosen that better meets the desired performance target.

6. Test Configurations:

As shown in the figures below from the ANSI/SCTE 168 series, IP switch/routers are central to typical Headend system unit-to-unit communications as well as in forwarding

all program streams to downstream Hub and subscriber locations. At Hub locations, IP switch/routers are again central to the system's unit communications and in forwarding video downstream to the HFC network.

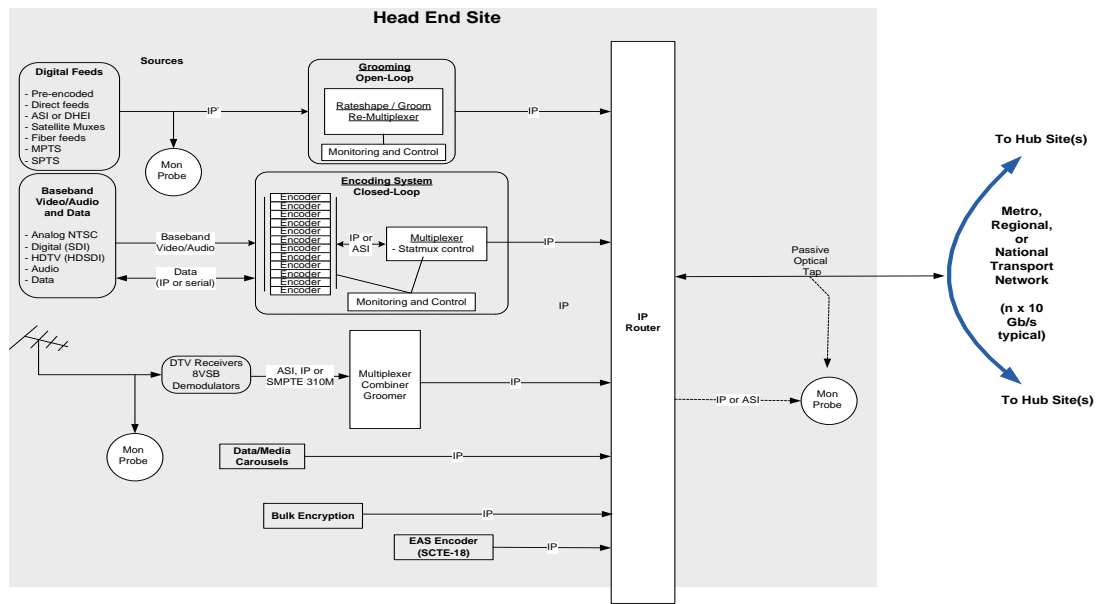


Figure 2.

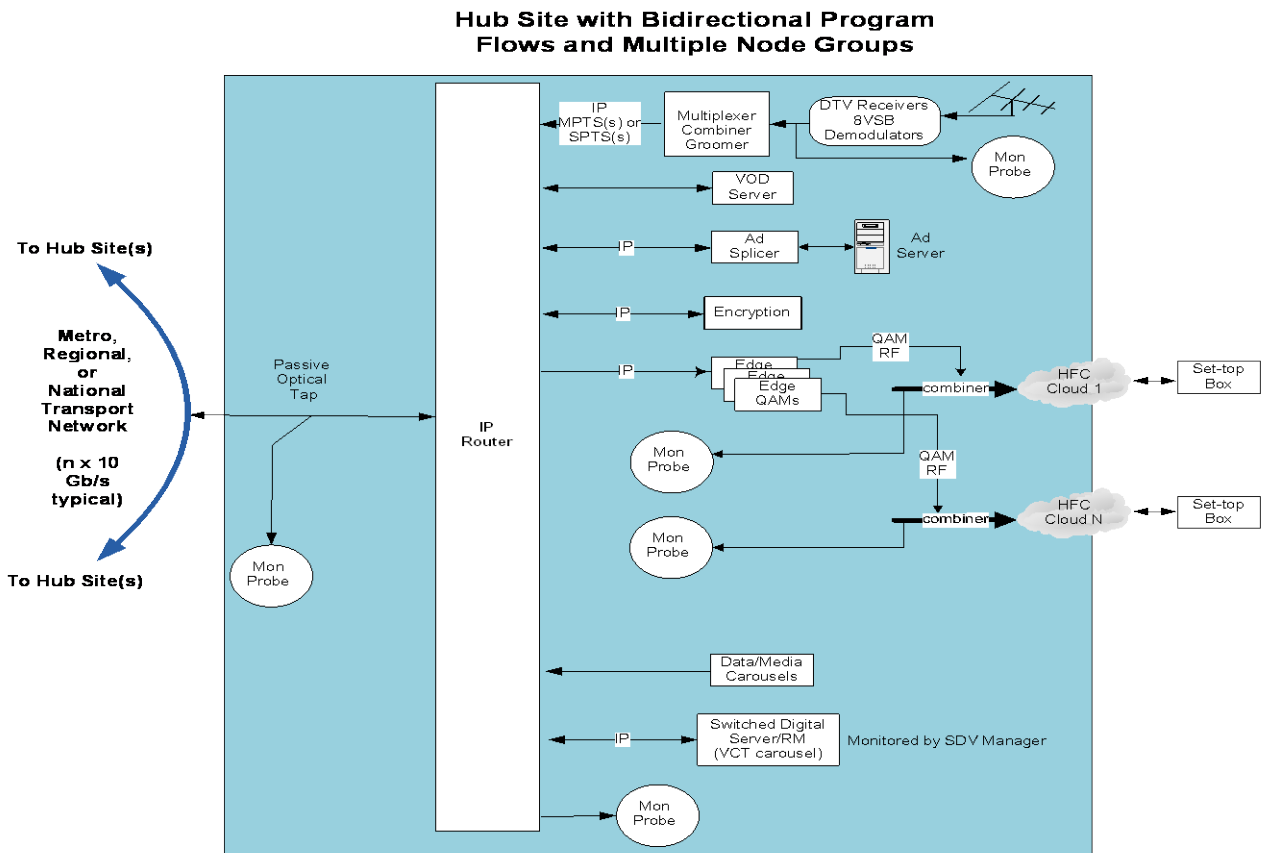


Figure 3.

To represent how a device will perform in a deployed configuration, the lab test configuration ideally should closely emulate the actual deployment configuration. For modern switch/routers with their flexible forwarding, security, and QoS policies, the configuration should be set as it would be used in the service provider's plant. Loads with similar numbers and types (speeds, physical layer types, etc.) of active physical ports will best represent how the device is likely to perform when deployed in the field. Such configurations will either be customized to that planned by the service provider or be supersets of the planned deployment configurations to stress the device beyond its planned initial operating configuration. The use of streaming video test programs should be used to reflect expected system loads. Streaming video analyzers, as used in service provider deployed networks, are used to continuously analyze every packet of the test streams both from the test sources at the input to the device or system under test as well as at the output of the device under test.

Analyzing the streams from the test sources as well as the outputs from the device under test allows for use of live streams as sources, if available, which may more closely represent the expected deployment environment. Use of a few live sources which are replicated by test instrumentation up to line rate or up to a volume that represents the worst case loads for a target system provides a realistic set of load conditions during the test while the source monitoring allows any transient impairments from the live source to be detected and properly discounted in the evaluation of the system under test.

6.1 Baseline Tests:

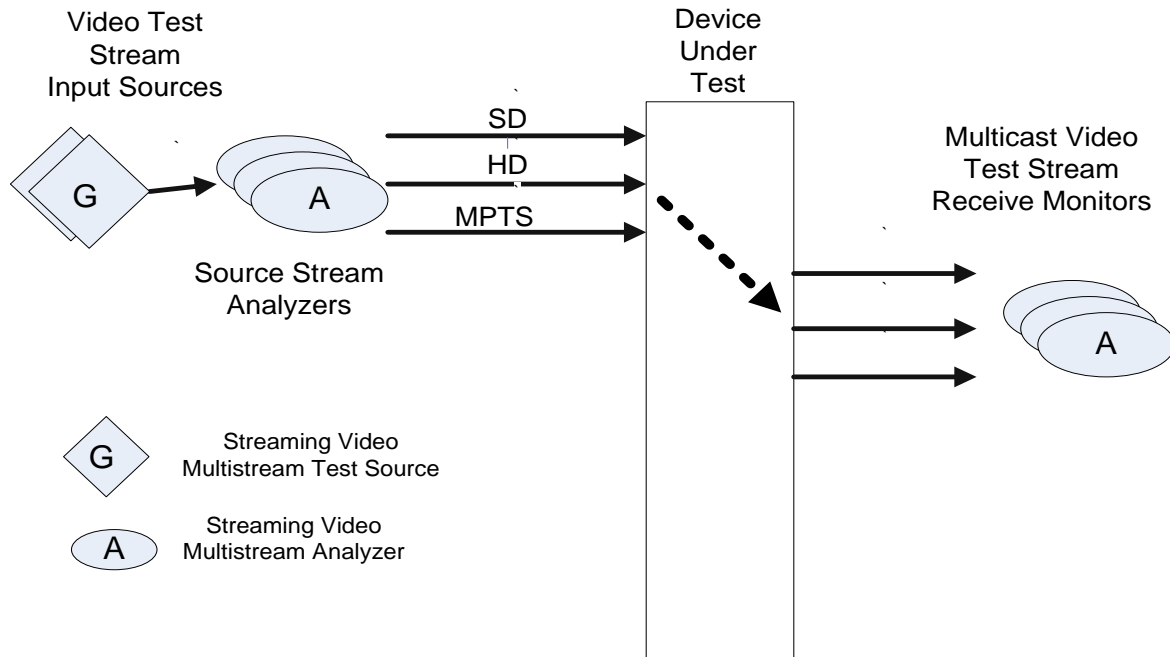
Since many devices of interest to service providers may never have been subject to longer duration, controlled video source availability tests, the following three test configurations are provided as a baseline. They should be used as a first step, before more complex and realistic configurations are constructed. The results of the baseline tests can also be used as a reference for comparative systems analysis in the absence of more detailed tests such as in preliminary evaluation during the selection of equipment and in evaluating different configurations of the same gear. The objectives are to obtain an availability performance baseline including:

- Performance stability – Verify that the device under test is capable of error-free forwarding of nominal video loads over continuous, multiday tests while providing acceptably high program availability.
- Recovery times for common fault types and operations variabilities – Verify the program availability impact for common, transient network

fault types including a Link Flap, power up (cold) reset, adding a new active output link, firmware update, updating a QoS policy, etc.

As these are baseline performance tests, any device configuration that enables the functionality described can be used for the baseline tests though it is always preferable to use an operational configuration that is expected to be deployed.

6.1.1 Test 1: Headend Broadcast Baseline



Test 1: Headend Baseline Configuration with SD, HD, and MPTS Multicast Aggregation Traffic

Test Description:

This baseline performance test uses 6 ports as shown in the figure above. The specific ports used on the device under test are not specified though it is preferable to use an expected operational port configuration subset. These tests assume that all ports operate at 1 Gb/s or faster. For slower or faster speed ports, traffic loads may be correspondingly scaled.

Port 1: 180 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (450 Mb/s) from test source [emulating ingest feeds]

Port 2: 23 HD SPTS MPEG 2 CBR multicast program streams of 19.0 Mb/s (450 Mb/s) from test source [emulating ingest feeds]

Port 3: 12 MPTS MPEG 2 CBR multicast program streams of 38.8 Mb/s (450 Mb/s) from test source [emulating MUX outputs]

Port 4: SPTS/HD Monitor Joining streams from source Ports 1 and 2 (900 Mb/s)

Port 5: SPTS/MPTS Monitor Joining streams from source Ports 1 and 3 sources (900 Mb/s)

Port 6: SPTS/HD/MPTS Monitor Joining selected streams from source Ports 1, 2, and 3 sources (520 Mb/s)

All test streams should be configured with nominal but realistic and common MDI Delay Factors [1] (burstiness) of 2.0 times ideal.

Test Execution:

1. Configure the device under test to enable the required forwarding and multicast functionality
2. Start up program stream sources on Ports 1, 2, and 3
3. Configure the program stream monitors on Ports 4, 5, and 6 to Join the test streams and log results to a test monitor
4. Run test for at least 5 continuous days
5. Generate Daily per-program Availability Report from program monitors

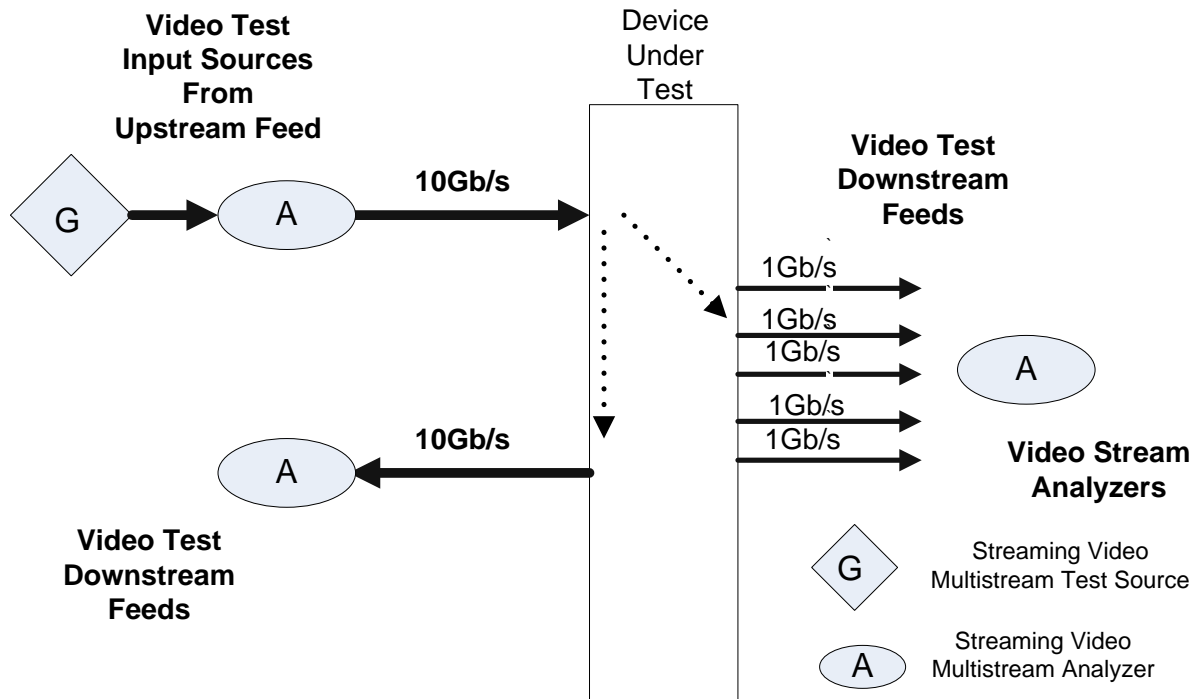
Evaluate if the results will meet deployed system goals. Many service providers are targeting Five 9s or 99.999% availability for their systems currently. Note that each deployed component must have higher availability than the target for the component ensemble in order to meet the required availability for the end-to-end system.

6. On final day of test, introduce common transient impairments and generate per-program Availability Report for the day from program monitors. The common transient impairments should include:
 - a. A Link Flap on a port connected to test stream sources.
 - b. A Link Flap on a port connected to a test stream monitor.
 - c. Switch/Router power-on (cold start) Reset
 - d. Add one new forwarding link (port 7) and stream monitor that receives the same program lineup as port 6 emulating a simple system expansion.
 - e. For devices with redundant routing engines, induce a failover event
 - f. Software/firmware update event including any necessary system resets

Comparing the availability performance of competitive devices with various device configurations in the presence of these common impairments can help providers assemble and use components that will optimize the availability of their deployed end-to-end

system by minimizing the contribution of program unavailability by individual packet forwarding devices.

6.1.2 Test 2: Core Network Switch/Router Baseline Operation



Core Baseline test with 10 Gb/s and 1 Gb/s Transport

Test 2: Core Baseline Configuration with 10Gb/s Regional Pass Through and Hub Site Local Traffic

Test Description:

The core network baseline performance test uses 7 ports as shown in the figure. The specific ports used on the device under test are not specified though it is preferable to use an expected operational port configuration subset. These tests assume that two ports operate at 10 Gb/s and represent a metro, regional, or national interconnect. Five additional 1 Gb/s ports represent interconnects to local hubs.

10 Gb/s port 1: Sourcing 1000 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (2.5 Gb/s) and 131 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (2.5 Gb/s) from test source for 5 Gb/s total aggregate traffic ingest [emulating on-demand feeds from upstream server farm]

10 Gb/s Port 2: Monitoring 500 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (1.25 Gb/s) and 65 HD SPTS MPEG 2 CBR unicast streams of 19.0 Mb/s (1.235 Gb/s) from test source [emulating on-demand feeds to downstream hubs]

1 Gb/s Port 3: Monitoring 100 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (250 Mb/s) and 13 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (247 Mb/s) from test source.

1 Gb/s Port 4: Monitoring 100 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (250 Mb/s) and 13 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (247 Mb/s) from test source.

1 Gb/s Port 5: Monitoring 100 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (250 Mb/s) and 13 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (247 Mb/s) from test source.

1 Gb/s Port 6: Monitoring 100 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (250 Mb/s) and 13 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (247 Mb/s) from test source.

1 Gb/s Port 7: Monitoring 100 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (250 Mb/s) and 13 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (247 Mb/s) from test source.

All test streams should be configured with nominal but realistic and common MDI Delay Factors (burstiness) of 2.0 times ideal.

Test Execution:

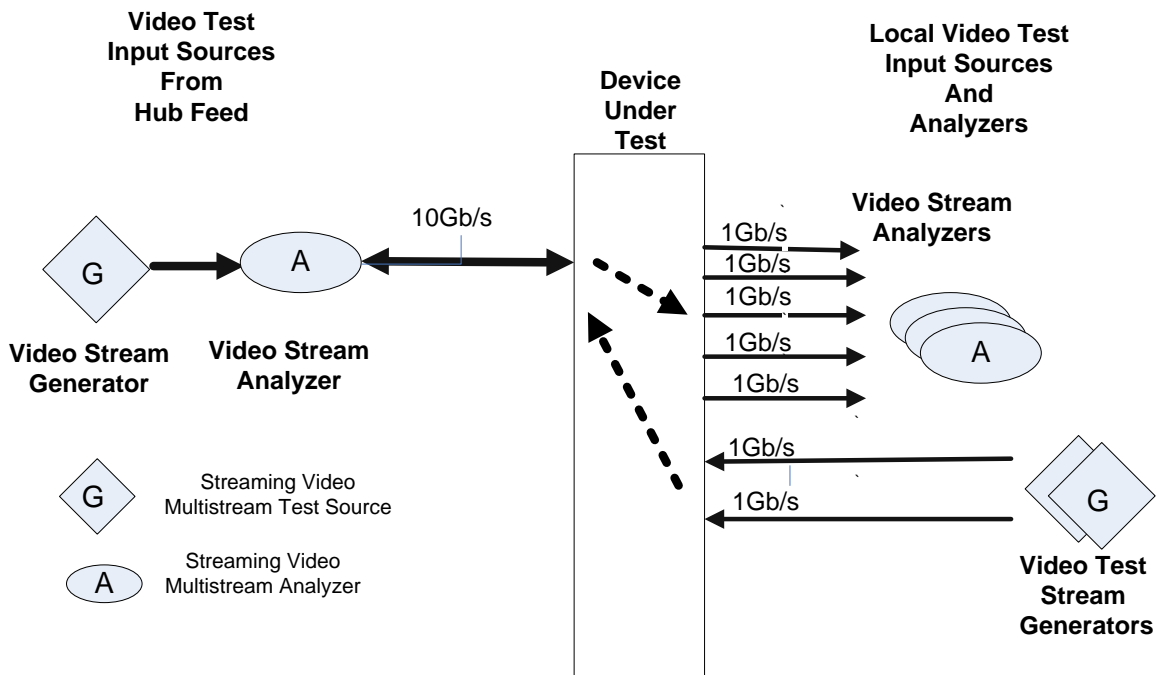
1. Configure the device under test to enable the required forwarding functionality based on the unicast addresses selected for the stream sources and monitors
2. Start up the program stream source on 10 Gb/s Port 1.
3. Configure the program stream monitors on Ports 1, 2, 3, 4, 5, 6 and 7 to monitor the test streams and log results to a test monitor
4. Run test for at least 7 continuous days
5. Generate Daily per-program Availability Report from program monitors

6. Evaluate if the results will meet deployed system goals. Many service providers are targeting Five 9s or 99.999% availability for their systems currently. Note that each deployed component must have higher availability than the target for the component ensemble in order to meet the required availability for the end-to-end system.

7. On final day of test, introduce common transient impairments and generate per-program Availability Report for the day from program monitors. The common transient impairments should include:
 - a. A Link Flap on a port connected to test stream sources.
 - b. A Link Flap on a port connected to a test stream monitor.
 - c. Switch/Router power-on (cold start) Reset
 - d. Add one new forwarding link (port 8) and stream monitor that receives the same program lineup as port 7 emulating a simple system expansion.
 - e. For devices with redundant routing engines, induce a failover event
 - f. Software/firmware update event including any necessary system resets

Comparing the availability performance of competitive devices with various device configurations in the presence of these common impairments can help providers assemble and use components that will optimize the availability of their deployed end-to-end system by minimizing the contribution of individual packet forwarding devices.

6.1.3 Test 3: Edge/Hub Network Switch/Router Baseline Operation



Hub Site Baseline test with 10 Gb/s Bidirectional Broadcast Traffic

Test 3: Hub Site Baseline Test Configuration with 10 Gb/s Bidirectional Broadcast Traffic

Test Description:

The Hub Site network baseline performance test uses 8 ports as shown in the figure above. The specific ports used on the device under test are not specified though it is preferable to use an expected operational port configuration subset. These tests assume that one port operates at 10 Gb/s and represent a metro, regional, or national interconnect. Five additional 1 Gb/s ports represent interconnects to local hub devices. This test uses unidirectional traffic fed from the upstream interconnect via the 10 Gb/s port representing the traffic that is intended to be delivered to Hub devices as well as unidirectional traffic originating from the Hub site bound for the upstream interconnect.

10 Gb/s port 1: Sourcing 1000 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (2.5 Gb/s) and 130 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (2.5 Gb/s) from test source for 5 Gb/s total aggregate traffic ingest [emulating on-demand feeds from upstream server farm]

1 Gb/s Port 2: Monitoring 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

1 Gb/s Port 3: Monitoring 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

1 Gb/s Port 4: Monitoring 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

1 Gb/s Port 5: Monitoring 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

1 Gb/s Port 6: Monitoring 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

1 Gb/s Port 7: Sourcing 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

1 Gb/s Port 8: Sourcing 200 SD SPTS MPEG 2 CBR unicast program streams of 2.5 Mb/s (500 Mb/s) and 26 HD SPTS MPEG 2 CBR unicast program streams of 19.0 Mb/s (494 Mb/s) from test source.

All test streams should be configured with nominal but realistic and common MDI Delay Factors (burstiness) of 2.0 times ideal.

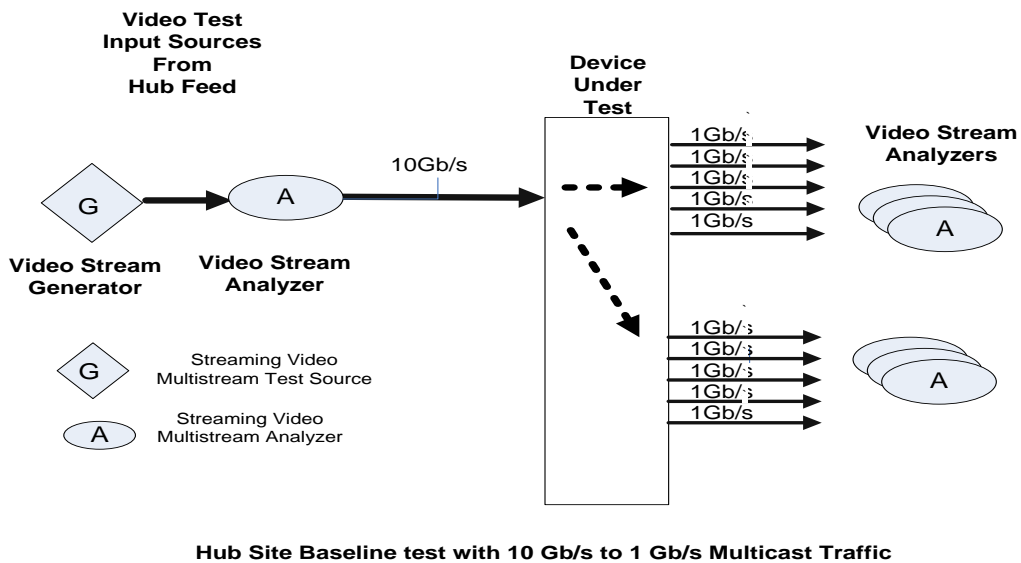
Test Execution:

1. Configure the device under test to enable the required forwarding functionality based on the unicast addresses selected for the stream sources and monitors. Static routes may be used to facilitate configuration.
2. Start up the program stream source on Ports 1, 7, and 8.
3. Configure the program stream monitors on Ports 1, 2, 3, 4, 5, and 6 to monitor the test streams and log results to a test monitor
4. Run test for at least 7 continuous days
5. Generate Daily per-program Availability Report from program monitors

6. Evaluate if the results will meet deployed system goals. Many service providers are targeting Five 9s or 99.999% availability for their systems currently. Note that each deployed component must have higher availability than the target for the component ensemble in order to meet the required availability for the end-to-end system.
7. On final day of test, introduce common transient impairments and generate per-program Availability Report for the day from program monitors. The common transient impairments should include:
 - a. A Link Flap on a port connected to test stream sources.
 - b. A Link Flap on a port connected to a test stream monitor.
 - c. Switch/Router power-on (cold start) Reset
 - d. Add one new forwarding link (port 9) and stream monitor that receives the same program lineup as port 8 emulating a simple system expansion.
 - e. For devices with redundant routing engines, induce a failover event
 - f. Software/firmware update event including any necessary system resets

Comparing the availability performance of competitive devices with various device configurations in the presence of these common impairments can help providers assemble and use components that will optimize the availability of their deployed end-to-end system by minimizing the contribution of individual packet forwarding devices.

6.1.4 Test 4: Edge/Hub Network Switch/Router Baseline Multicast Operation



Test 4: Hub Site Baseline Multicast Test Configuration with 10 Gb/s Bidirectional Broadcast Traffic

Test Description:

The hub network baseline multicast performance test uses 11 ports as shown in the figure above. The specific ports used on the device under test are not specified though it is preferable to use an expected operational port configuration subset. These tests assume that one port operates at 10 Gb/s and represent a metro, regional, or national interconnect. Ten additional 1 Gb/s ports represent interconnects to local hub devices or end users. This test uses unidirectional traffic fed from the upstream interconnect via the 10 Gb/s port representing the traffic that is intended to be delivered to local Hub devices. As the purpose of this test is to provide baseline tests of multicast performance, the test should be executed using the expected IGMP version that is planned for deployment. If the

version is not yet determined, the DUT should be tested separately for both IGMP v2 and IGMP v3.

10 Gb/s port 1: Sourcing 2000 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (2.5 Gb/s) aggregate traffic ingest [emulating feeds from upstream headends]

1 Gb/s Port 2: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 3: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 4: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 5: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 6: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 7: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 8: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 9: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 10: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

1 Gb/s Port 11: Monitoring 200 SD SPTS MPEG 2 CBR multicast program streams of 2.5 Mb/s (500 Mb/s) from test source.

All test streams should be configured with nominal but realistic and common MDI Delay Factors (burstiness) of 2.0 times ideal.

Test Execution:

1. Configure the device under test to enable the required forwarding functionality based on the multicast addresses selected for the stream sources and monitors.
2. Start the program stream source on Port 1.

3. Configure the program stream monitors on Ports 2 through 11 to join and monitor the test streams and log results to a test monitor.
4. Run test for at least 1 day.
5. Configure the program stream monitors on Ports 2 through 11 to periodically join and leave the test streams and log results to a test monitor. First, configure Port 2 to join/leave its streams every 20 minutes while ports 3 through 11 remain joined to their streams. After 60 minutes of operation, configure Port 2 to join/leave its streams every 20 minutes while ports 3 through 11 remain joined to their streams. Continue incrementally adding dynamic join/leave stream monitors until all monitors are joining/leaving their streams every 20 minutes.
6. Generate Daily per-program Availability Report from program monitors.
7. Generate a report of per-group maximum join/leave times for each incremental test in step 5. Include the number of failures to join or leave for each incremental test.
8. Evaluate if the results will meet deployed system goals. Many service providers are targeting Five 9s or 99.999% availability for their systems currently. Note that each deployed component must have higher availability than the target for the component ensemble in order to meet the required availability for the end-to-end system.
9. Introduce common DUT transient impairments while performing step 12 and generate per-program Availability Report for the day from program monitors. The common transient impairments should include:
 - a. A Link Flap on a port connected to test stream sources.
 - b. A Link Flap on a port connected to a test stream monitor.
 - c. Switch/Router power-on (cold start) Reset
 - d. Add one new forwarding link (port 9) and stream monitor that receives the same program lineup as port 8 emulating a simple system expansion.
 - e. For devices with redundant routing engines, induce a failover event
 - f. Software/firmware update event including any necessary system resets

Comparing the availability performance of competitive devices with various device configurations in the presence of these common impairments can help providers assemble and use components that will optimize the availability of their deployed end-to-end system by minimizing the contribution of individual packet forwarding devices.

Test Discussion:

With the 2000 source streams configured as 2000 unique multicast groups, the analyzers in step 3 above should be configured in at least the following configurations:

- Each analyzer is configured to join/leave different groups such that the port 2 analyzer joins 200 different groups, the port 3 analyzer joins/leaves 200 different groups that are themselves different groups than those joined in the port 2 analyzer, etc. This configuration using analyzers with non-overlapping groups results in the maximum number of different streams being forwarded.

- Each analyzer is configured to join/leave the same groups such that the port 2 analyzer joins 200 different groups, the port 3 analyzer joins/leaves the same 200 groups as the port 2 analyzer, etc. This configuration using analyzers with overlapping groups results in increasing the replication load on the DUT.
- Optionally, the DUT may be configured with VLANs such that each group join/leave in an analyzer could emulate an end user. This configuration may be used to further increase the replication load on the DUT.

7. Summary

This note proposes and describes a set of baseline tests that can be used in the process of evaluating network devices for suitability for streaming video applications. The test conditions include both longer soak times than may have been used in evaluations to date, typical sustained video traffic loads to represent operational conditions, and some typical transient fault conditions. The results are reported in terms of program availability which reflects how subscribers will experience faults in order to highlight the performance of one device compared to another or to compare device configurations. In addition, multicast test configurations indicate DUT effects on changes to multicast join and leave times.

These tests are intended to complement existing industry tests used to evaluate conventional data oriented performance of switches and routers.