Signalling Performance Case Studies using the Multiport UNI Signalling Performance Test Solution

Agilent Technologies Broadband Series Test System Application Note

Introduction

This application note looks at the task of testing switches or networks in order to determine their signalling performance.

The Agilent Technologies E1600A Multiport UNI Signalling Performance Test Module is specifically designed to suit this need.

Four case studies, collected from live networks, will be examined in detail in order to illustrate different aspects of signalling performance in switches and networks.

How well does your switch or network perform?

Understanding the performance of your switch or network is crucial to ensure the successful deployment of any ATM network.

Bottlenecks and reliability problems need to be identified, and their impact on the network investigated to ensure that the network will perform adequately.

The following questions need to be answered when evaluating signalling performance:

- Why is signalling performance important?
- Who needs to perform signalling performance tests?
- How does the switch architecture affect signalling performance?
- What are the key performance factors to consider?
- What is the call life-cycle? What is meant by *active calls*, *call cycles*, and *call setups*?
- What are the key signalling performance parameters?
- How do you measure them?

Why signalling performance is important

Characterization of signalling performance is important to understanding how your switch or network will perform under realistic traffic situations.

This helps you deliver the level of responsiveness and reliability that your customers demand. To do this, you need to observe and analyze a range of signalling performance parameters.

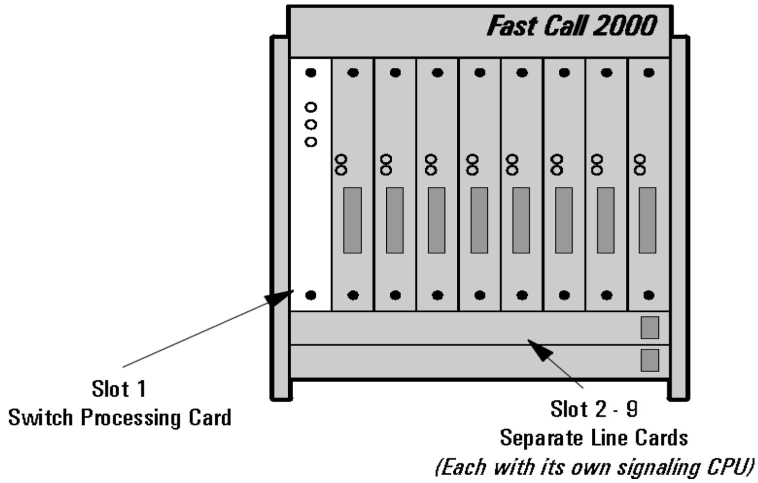

Switch architecture with a central processor and a signalling processor per line card.

Who needs to measure signalling performance

System test groups need to ensure that the performance of the System Under Test (SUT) meets the relevant UNI signalling specifications.

QA groups need to ensure that the SUT performs reliably under load for an extended period of time.

Equipment designers need to observe the signalling performance of the SUT to identify hot spots before they can attempt to fine-tune the performance.

Equipment manufacturers need to differentiate switch capabilities against the competition for competitive advantage.

Service Providers and other end-users have a range of applications for performance testing. They need to:

- independently verify performance claims made by manufacturers when determining which switch best suits their needs
- test a switch to determine which one is best suited for their specific application. Some switches are better suited to core networks, while others are best suited to edge-based applications.
- characterize and fine-tune the performance of the network after it has been installed and become operational
- independently verify the long-term reliability and performance of a switch or network

How ATM switch architecture affects signalling performance

Switch architectures are many and varied. Often the number of ports in a switch will vary depending on the number of line cards installed in the switch itself. Other switches have a fixed number of ports and do not allow the inclusion of extra cards.

The distribution of signalling processing within the equipment has a large impact on the signalling performance. For example, in the simplest architectures, the switch has one central processor managing the whole switch. This processor has to perform all of the signalling processing for every port under its domain. Therefore, the signalling performance will vary depending on how many ports are in use at one time.

More sophisticated switch architectures use a central processor to manage the backplane, as well as a processor for every line card in the system. The signalling processing is performed by the line-card processors. The central processor is only consulted for tasks such as addressing and routing (PNNI signalling). This architecture can dramatically increase the signalling performance of the entire switch.

The signalling performance will depend on how many ports are in use on a particular line card. As a secondary effect, the number of line cards in use will also influence performance. Other issues, such as the amount of memory and the processor bandwidth, can have an effect as well.

Key performance factors to consider

Call establishment time

At what level of signalling activity does my switch start taking longer to establish calls?

Your customers will get used to a certain level of network performance. If the network suddenly starts performing badly, they usually get upset and want to know what is wrong.

It is also important to know at what level you can load your network up before it needs upgrading in order to meet customer expectations.

Excessive signalling load

Does my switch recover from excessive signalling load?

Does my switch crash under excessive signalling load?

For many applications, it is critical that the network does not break down for any reason. It is important to evaluate how the switch or network behaves under excessive signalling load.

Call rejection

At what number of active calls does my switch start to reject calls?

This is a basic performance metric that is used for comparison and assessment of different switch equipment.

Hot spots

Are there any hot spots in the signalling performance of my switch?

Hot spots are critical areas of signalling performance in the SUT that could adversely affect the overall performance of a network. For example, a low figure for the maximum number of active calls could severely limit the capacity of the network to provide SVC connections.

Single-port versus multiport testing

Does the switch or network perform better with single testing compared to multiport testing?

What is the switch architecture and how is signalling processing distributed across the switch?

How well does the performance of my switch scale across multiple ports?

Does my switch handle signalling across multiple ports fairly?

What is the real signalling performance when all ports are loaded with signalling traffic?

Often, switches perform well on a single port, but do not perform as well when signalling occurs on multiple ports. It is important to evaluate the performance of a switch when all ports are under load. As discussed earlier, the switch architecture can affect multiport signalling performance significantly.

Long-term considerations

Will my switch perform reliably in the long-term?

It is important that the switch can continue providing SVC services for extended periods of time under high-load conditions.

Other considerations

What other factors impact on the performance of my switch?

Other factors such as traffic parameters and ILMI MIB settings can affect signalling performance. These factors need to be taken into consideration when fine-tuning performance.

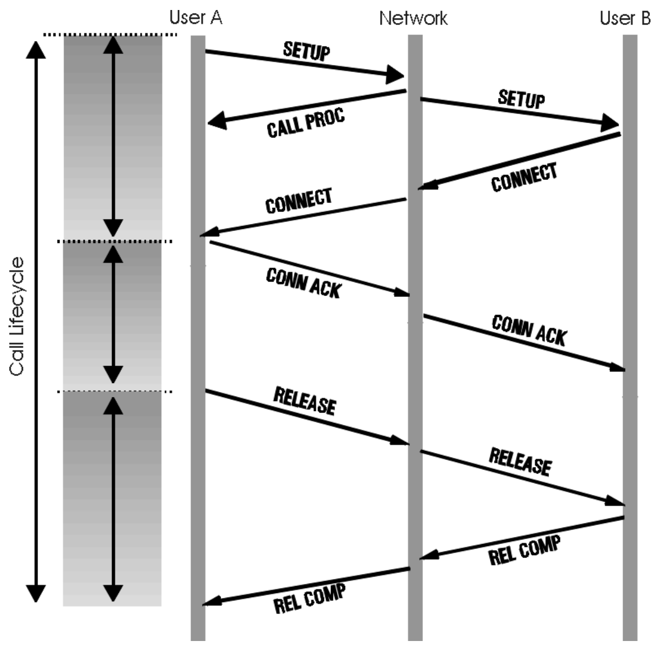

Call life-cycle diagram showing establishment, active, and inactive phases.

The call life-cycle

The call life-cycle consists of three phases.

1. An *establishment* phase, initiated by a UNI signalling *SETUP* message.
2. An *active* phase, where the call remains connected and data can be transferred across the SVC. The duration of the call is known as the *hold time*.
3. An *inactive* phase, initiated by a UNI signalling *RELEASE* message. The process of forcing the network to release an SVC is also called *tear-down*.

To thoroughly test a network, all three phases of the call life-cycle must be exercised.

Establishment phase

The establishment phase consists of three signalling message primitives. A *SETUP* message, which must be responded to with a *CONNECT* message, which in turn must receive a *CONNECT_ACKNOWLEDGE* message. Only after all three messages have travelled through the network is the establishment phase finished. The entire establishment phase is known as a *call setup*.

Active phase

Once the establishment phase has completed successfully, a call enters the *active* phase. More specifically, a call enters the active phase after the transmission of either a *CONNECT* message or a *CONNECT_ACKNOWLEDGE* message, depending on whether the end-point is the originator or the receiver of the call.

Inactive phase

The active phase ends whenever a *RELEASE* or *RELEASE_COMPLETE* message is sent or received by either party. This process is known as a *call release* or *call tear-down*.

All three phases in sequence are known as a *call cycle*.

Pending and rejected calls

After the first *SETUP* message has been sent and before the establishment phase completes successfully, the call is referred to as *pending*.

If the establishment phase does not complete successfully, the call is referred to as *rejected*.

Key signalling performance parameters

Some of the key signalling performance parameters in use today are:

- *Maximum active calls* is a fundamental performance parameter. If this figure is low, then it will severely affect most other aspects of signalling performance.
- *Call setup performance* is one of the is the most widely-used performance parameters today. It is often quoted by equipment manufacturers in sales brochures and technical specifications.
- *Burst response* is similarly important, but uses a different test methodology.
- *Call cycling performance* is the most thorough way of testing signalling performance. This method is able to test all aspects of signalling performance in the one test because it looks at both the call setup and the call release phases of the call life-cycle.

How to measure the key performance parameters

Introducing the Multiport UNI Signalling Performance Test Solution

The Multiport UNI Signalling Performance Test Solution is provided by the E1600A Multiport Test module (MTM) and the UNI signalling performance software application for the Broadband Series Test System. The MTM provides a total of 50 real-time statistics related to UNI signalling performance of your switch or network.

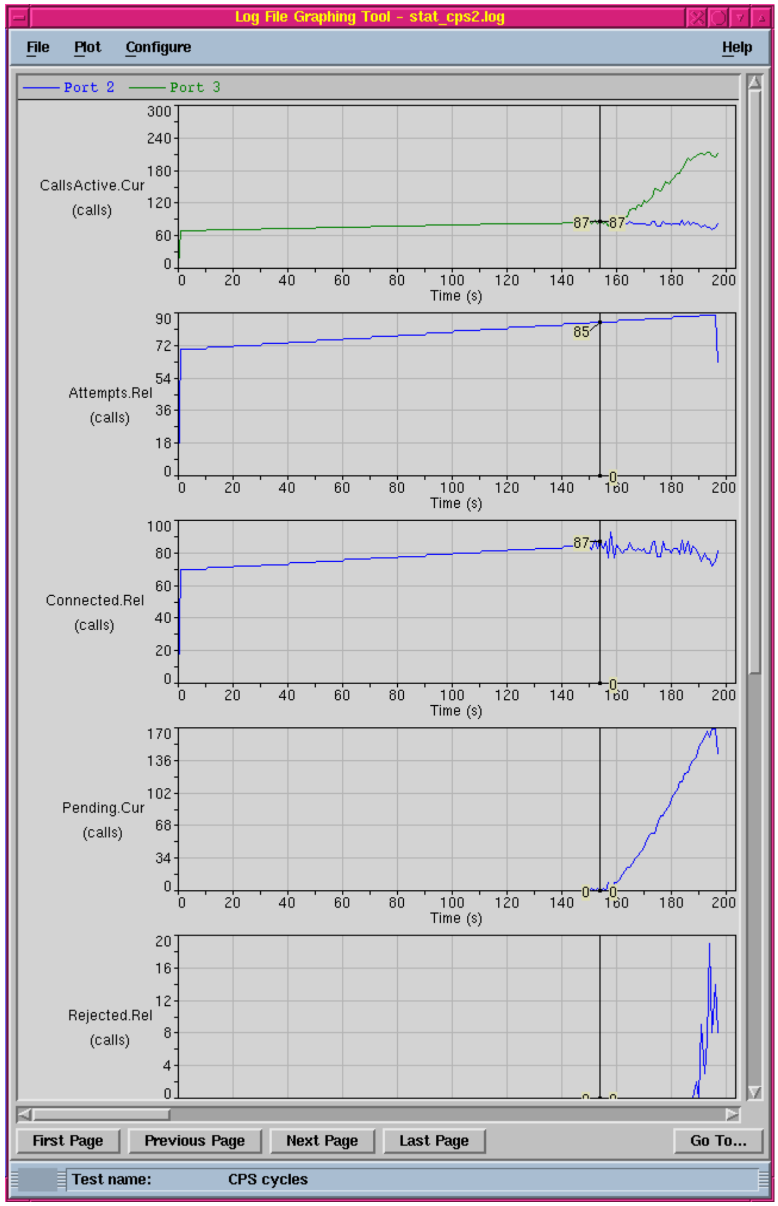

Log file graphing tool, supplied with the Multiport UNI Signalling Performance Test Solution.

MTM Log File Graphing Tool

Correlation across a number of statistics is required in order to derive the key performance parameters. It is important that you get a complete view of the signalling traffic so you are able to identify any hot spots in performance.

The Multiport UNI Signalling Performance Test Solution provides a log file graphing tool that allows you to select a group of statistics and display them as time-correlated graphs.

Key statistics of the E1600A MTM

Some of the most important statistics provided by the MTM are related to the call life-cycle discussed earlier.

- Call attempt statistics such as *Attempts.Cur* shows the number of calls currently being initiated by the MTM module. *Attempts.Rel* shows the call attempt rate in calls/s. This statistic, correlated with the number of connected, pending, and rejected calls is used to determine the call setup performance of the SUT.
- Active call statistics such as *CallsActive.Cur* measure the maximum number of active calls supported by the switch or hardware. This statistic is also correlated with setup and tear-down measurements. Typically, the setup and tear-down time increase as the number of active calls grows.
- Statistics for released calls and clear-pending calls give an accurate representation of the call release phase of the life-cycle. Often this is the ‘Achilles’ Heel’ of the switch’s overall performance and is much worse than the setup performance.

The following case studies show how to measure the four key performance parameters on a switch or network.

- Maximum active calls
- Call setup performance
- Burst response
- Call cycling performance

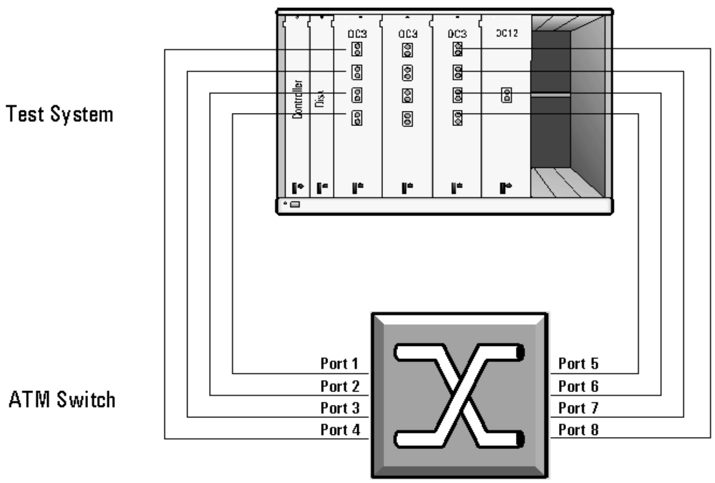

Performance Test Setup for a Switch.

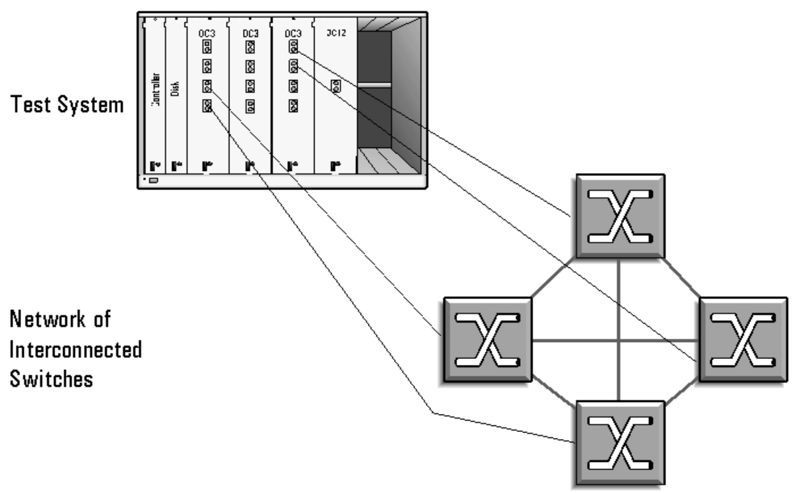

Performance Test Setup for a Network.

- Active calls ramp up steadily until T = 20s, when the switch reaches its call connection limit (1580 calls)
- Test stimulus: The tester generates call setup attempts at a steady rate (80 calls/s).
- The switch reaches its active call limit at T = 20 s and stops connecting calls - connection rate falls to zero over a 2 s interval.
- After T = 20 s, the call attempt rate is steady, but the call connection rate is zero; therefore pending calls ramp up steadily.
- After T = 27 s, UNI signalling protocol timers start to expire and the switch starts rejecting calls.

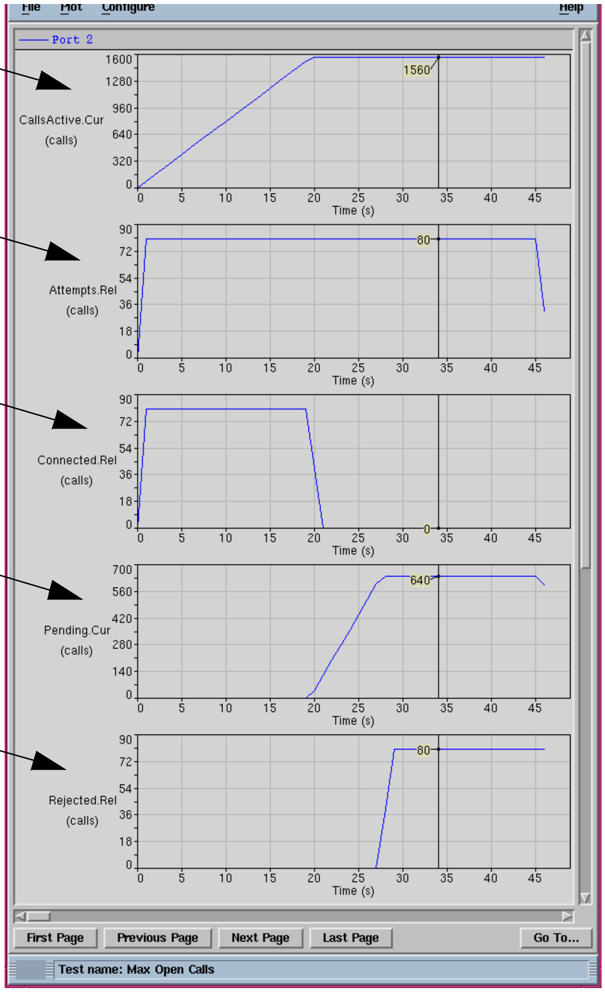

Graphs showing a switch reaching its active call connection limit (1580 calls).
Note: .Cur = current value (calls), .Rel = rate (calls/s).

Case Study #1: Maximum active calls

Test purpose

During a normal working day, businesses use your network to access vital information. Expected network load is a vital point to consider when evaluating networks or switches.

How many simultaneous users can access my network?

To answer this question, you must determine the *maximum number of active calls* that you switch or network can support.

For example, on your network, typical users have three channels open at any one time; a video link, file transfer, and a terminal connection. Therefore, you need to build in enough active connection capability to support this requirement.

Brief test procedure

Configure the test equipment to do the following:

- Initiate calls at a constant rate.
- Ensure there is no call clearing by selecting an infinite call duration.
- Enable automatic Layer 3 *RESTART* at the start of your test to clear any existing calls.
- Use the log file graphing tool of the MTM to determine the maximum number of active calls that can be handled by the switch.

Test results

The graph shown here illustrates typical switch behavior under constant load.

In this example, the *Connected.Rel* and *Pending.Cur* statistics indicate the point in time at which the SUT has reached its active call limit. The value of active call limit is given by the *CallsActive.Cur* statistic.

Affect of switch architecture

The number of active calls that a switch or network can handle will vary depending on the architecture.

If the signalling processing is distributed across line cards, then it is likely that the measured performance is determined by the maximum number of active ports per line card.

However, if there is a single dedicated signalling processor, then the measured performance is determined by the total number of active ports per for the entire switch.

To fully understand how the architecture of your switch affects the signalling performance, you should measure the call setup performance for:

- a single port
- a fully loaded line card
- a fully loaded switch
- an entire network.

Affect of switch configuration

The number of open calls supported by a switch can vary depending on a number of external influences.

Some switches allow configuration of the maximum number of active calls supported per port or switch. Typically, signalling performance will degrade as the number of open calls increases.

Affect of ILMI MIB settings

The settings in the ILMI MIB can affect the number of open calls supported per port.

For example, the number of VPI/VCI bits in the cell header needs to be configured the same for the switch and the end-station device. If the configurations do not match, then the lowest common denominator will be chosen between the settings of the switch and that of the end station device.

This affects the maximum number of active calls. For example, if the number of VPI bits resolves itself to 1, and the number of VCI bits resolves itself to 8, then the total number of calls supported by the port will be 256.

Other ILMI MIB parameters can have an effect on the total number of calls that can be set up for a port.

Affect of service category

The service category information in the *SETUP* message can also affect the number of active calls. There are two important factors:

- The service category selected (CBR, VBR, ABR, UBR). Sometimes the number of connections for a particular traffic type is limited due to resource constraints within the switch fabric itself.
- The specified traffic parameters for the selected service category. For example, if the physical link between the switch and the end station device is an OC-3 link, supporting approximately 353000 cells per second, then specifying a CBR traffic type with a PCR of 1000 will only give you 353 active connections in total.

To fully understand how the service category affects the performance of your switch or network, you should measure the call setup performance for a variety of service classes and traffic parameters.

Detailed test procedure using the E1600A MTM

To test for the number of active calls on a single port:

- Select two ports on an MTM module.
- Configure one port to call the other port.
- Select a Constant call profile with a rate that will not cause the switch or network to be overloaded.
- Select the Call Duration to be Infinite.
- Enable an Automatic Layer 3 Restart at the start of a test through the Module Wide Parameters screen.

Detailed analysis using the E1600A MTM

The test can be stopped as soon as calls start to be rejected. Observing the listed statistics and their graph with the log file graphing tool:

- The number of active calls will flatten out when it has reached its maximum. This indicates the number of active calls supported by a particular port.
- The point at which calls first stop being connected is found by identifying the point at which pending calls first pending start to increase.

Case Study #2: Call setup performance

Test purpose

During normal network usage, the number of call establishments is generally relatively stable. However, at the commencement of a working day, workers will log on to a network, and establish connections with computers of interest, and download vital information such as e-mail from a central server. All this activity can place a large burden on a network.

At what peak level of activity does my switch or network start slowing down, stop altogether, or crash?

To answer this question, you need to quantify the *call setup performance* of your network or switch. If the performance is inadequate, then some users will be unable to log onto a network or obtain vital business information when they need it.

Brief test procedure

Configure the test equipment to do the following:

- Initiate calls at a gradually increasing rate.
- Ensure there is no call clearing by selecting an infinite call duration.
- Enable automatic Layer 3 *RESTART* at the start of your test to clear any existing calls.
- Use the log file graphing tool of the MTM to determine the call setup performance.

Test results

The graph shown here illustrates typical switch behavior under increasing load. In this example, the *Pending.Cur* statistic indicates the point in time at which the SUT has reached its call setup performance limit. The value of call setup performance is given by the *Connected.Rel* statistic.

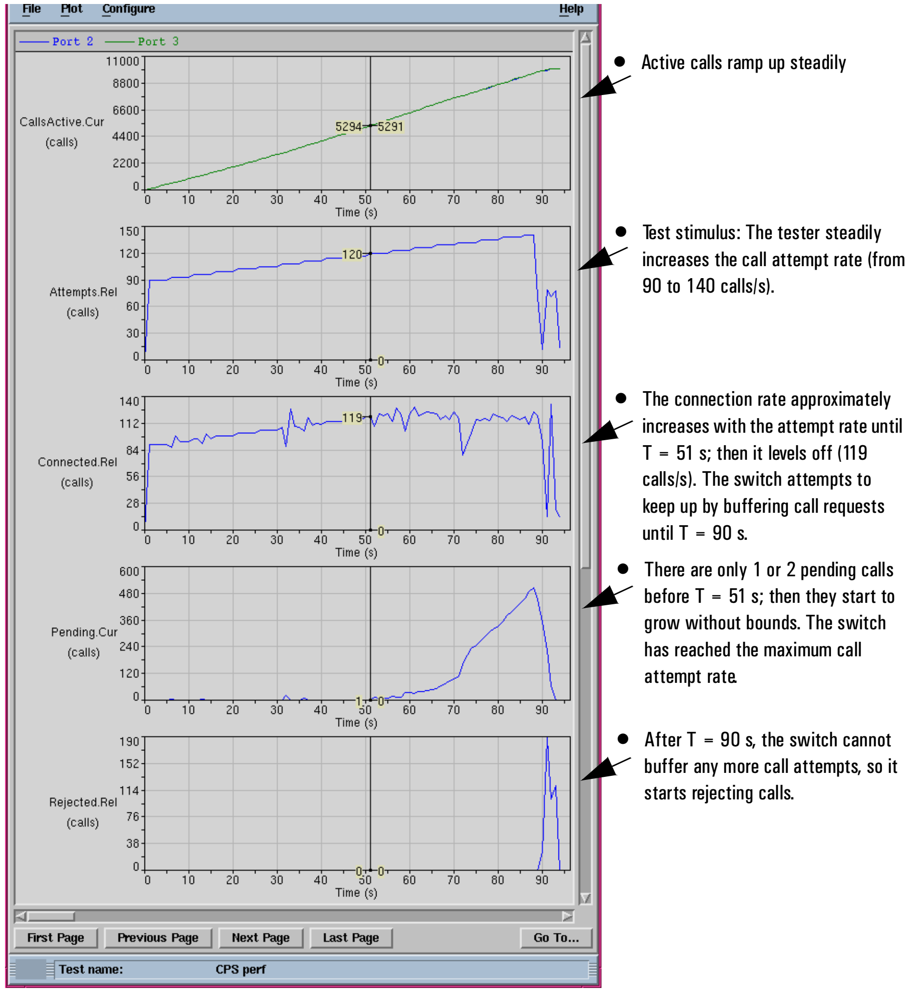

Graphs showing a switch reaching its maximum call connection rate (119 calls/s).
Note: .Cur = current value (calls), .Rel = rate (calls/s).

This same test can be performed when a number of ports are calling each other to determine the call setup performance supported by:

- an entire line card
- an entire switch
- an entire network

using multiple ports on multiple MTM modules.

Affect of the *maximum active calls* performance parameter

One of the factors that can affect the test results is the maximum number of active calls the switch is capable of handling (measured in case study #1). If this number is very low, then it will be difficult to test the call setup performance.

In this situation, each call setup will become an active call that takes the switch close to the active call limit. If the active call limit is reached before the maximum call setup performance is determined, then the test will not provide accurate information.

To measure the call setup performance accurately when the maximum active calls parameter is very low, you can use a different approach, such as the burst response test (case study #3).

Detailed test procedure using the E1600A MTM

To test for the call setup performance on a single switch port, the MTM should be configured in the following way:

- Two ports will be used. One port making calls to the other.
- Ensure there is no call clearing by selecting an infinite call hold time. If a call is cleared, then it becomes a call cycle. Call cycling performance is a different aspect and should be tested separately (case study #4).
- Enable an Automatic Layer 3 Restart through the Module Wide settings screen.
- Choose a stepped call initiation profile with a gradually increasing rate.

Detailed analysis using the E1600A MTM

The test can be stopped as soon as calls start to be rejected. Observe the statistics and the graph via the log file graphing tool:

- Observing the *Pending.Cur* statistic, there should be a point at which the number of pending calls starts growing without bounds. This is the point at which the call setup performance of the switch or network has reached its maximum.
- Correlating this point with the *Attempts.Rel* statistic will show the number of Call Setups that are being attempted by the MTM. At the same point in the graph, the *Connected.Rel* statistic should level out. This number is the Maximum Call Setup Performance of the SUT.
- If *Connected.Rel* drops off altogether, then it shows a flaw in the SUT. A switch or network should still connect calls at the maximum rate it is capable of, even if it cannot match the call attempt rate.
- Some time after the number of pending calls first starts to grow, calls will start to be rejected. This can be seen with the *Rejected.Rel* statistic. This will typically occur between 10 and 30 seconds after the number of pending calls first starts to grow. It represents protocol time-outs, and is a normal response when calls fail to be connected within a specified time interval.

Determining the cause of call rejection

It can be useful to observe the value of the *CAUSE* information element within the *REJECT* or *REJECT_ACKNOWLEDGE* message returned by the SUT. This can be seen via the *Clearing Cause Distribution* screen that is available through the MTM results dialog panel. This information helps you determine the cause of the performance limitation.

Variation of *call setup time* with number of active calls

Another interesting trend that can be observed with this test is the variation in the call setup times as the number of active calls grows. Typically, the call setup time, which can be observed through the *AvgCallSetupTime.Cur* statistic, will increase as the number of active calls increases.

A large variation in the call setup times may affect the measured value of call setup performance.

For example, if the call setup time is

- 40 ms when the number of active calls is very low, or
- 100 ms when the number of active calls is near the maximum active call limit,

then the call setup performance measurement will be dependent on the number of active calls.

This is an important trend to characterize because it indicates degradation of the network performance under realistic load conditions.

- Active calls ramp up steadily.
- Test stimulus: The tester generates a 10 s burst of call attempts at 200 calls/s.
- The switch connects calls at its maximum rate (125 calls/s). After T = 10 s, the switch is only handling buffered call requests so the connection rate increases.
- Pending calls ramp up steadily until T = 10 s, due to the difference between the attempt rate and the connection rate. By T = 14 s, the switch has connected all the pending calls, so the recovery time is 4 s.
- The call setup performance of the switch is 1250 calls / 10 s = 125 calls/s.

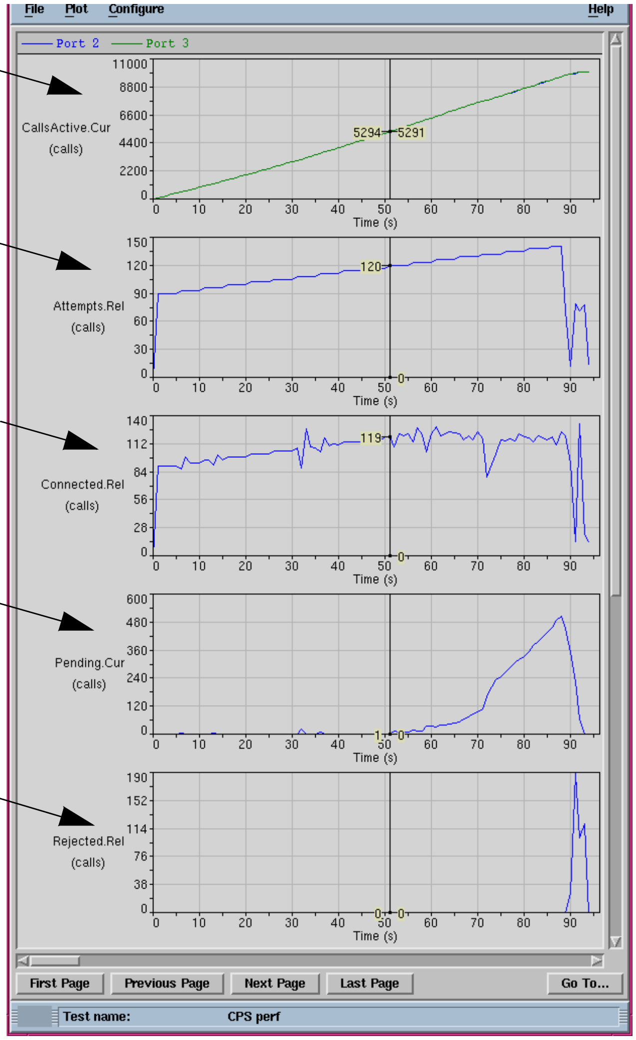

Graphs showing the call setup performance (125 calls/s) and recovery time (4 s) of a switch under overload conditions.
Note: .Cur = current value (calls), .Rel = rate (calls/s), .Cum = cumulative value (calls)

Case Study #3: Burst response

Test purpose

It is quite likely, and acceptable, that a network will experience periods where the rate of signalling traffic it is subjected to will temporarily exceed its maximum call setup capacity.

How does my switch or network handle a burst of signalling traffic?

The burst response parameter represents the ability of the switch or network to handle a load condition that temporarily exceeds its signalling processing capacity. It is important to understand the behavior of the switch under these conditions.

Does the switch keep attempting to connect calls at the maximum rate it is capable of?

Does the connection rate drop off slightly or fall to zero?

Does the switch crash?

The burst response test can also be used as an alternative method to measure call setup performance.

It is probably the most common method used by switch manufacturers to determine the published switch performance figure. However, the results from this test can be quoted in a number of different ways and can sometimes create a misleading impression about the true performance of the switch!

Brief test procedure

Configure the test equipment to do the following:

- Initiate a 10 second burst of call setups. The burst rate should be higher than the known *call setup performance* for the switch (case study #2)
- Ensure there is no call clearing by selecting an infinite call duration.
- Enable automatic Layer 3 *RESTART* at the start of your test to clear any existing calls.
- Use the log file graphing tool of the MTM to analyze the burst response of the switch.

Test results

The graph shown here illustrates a typical burst response.

The number of pending calls, measured by *Pending.Cur*,rises steadily during the burst period because the switch cannot keep up with the call activation rate initiated by the tester.

After the burst, the number of pending calls falls to zero, as the switch handles buffered call requests.

The cumulative number of connected calls, *Connected.Cum*, at the end of the burst, divided by the burst duration (10 seconds) gives the *call setup performance* for the switch.

Detailed test procedure using the E1600A MTM

To test for the burst response on a single switch port, the MTM should be configured in the following way:

- Two ports will be used. One port making calls to the other.
- Ensure there is no call clearing by selecting an infinite duration. If a call is cleared, then it becomes a call cycle. Call cycling performance is a different aspect and should be tested separately (case study #4).
- Select a stepped profile to generate a fixed burst of call setups. This can be achieved by specifying the final rate of the step to be zero, and the initial rate equal to the magnitude of the burst each second. The step size is equal to the initial rate, and the step duration is the length of the burst itself.
- Select the Call Duration (hold time) to be Infinite.
- Perform a manual Layer 3 Restart through the “Module Wide Controls” dialogue before the test is started.

Detailed analysis using the E1600A MTM

The test can be stopped as soon as all calls in the burst have been initiated, and there are no pending calls remaining. Observing the statistics and the graph via the log file graphing tool:

- The point where the *Attempts.Rel* statistic reaches the end of its burst is the main point of interest, and is where other statistics will also be analyzed.
- The *CallsActive.Cur* statistic, and the *Connected.Cum* statistic keep growing after this point. However the value of the *Connected.Cum* statistic at this point is of most interest and will be discussed shortly.
- The *Pending.Cur* statistic slowly drops off and returns to zero following the end of the burst. The time taken for the number of pending calls to fall to zero is the *recovery time* of the switch to this particular overload condition.
- It is also not unusual to observe rejected calls at some stage through this test. If the number of attempts in the burst grows too large, then the calls will not be responded to within the limits of the protocol timers, and the calls will be deemed rejected.

- The *Connected.Rel* statistic provides some additional information about the behavior of the switch. During the burst, the connection rate is fairly constant, and at the maximum level the switch is capable of. After the burst, the connection rate more than doubles! This is because the switch is clearing call request buffer. It performs better because it does not have to respond to signalling messages.
- On the graph shown, when the burst has finished at the ten second mark, the *Connected.Cum* statistic has reached a value of 1250 calls. If this is aggregated over the period of the burst itself (10 seconds), then we get a value of 125 calls/second.

$$CallSetupRate = \frac{ConnectedCum}{BurstDuration}$$

Advantages of the burst response test

The burst response test is the most reliable way to determine the call setup performance. It avoids most of the problems associated with a low active call limit. The burst size can be shortened as necessary to avoid the active call limit without distorting the call setup performance calculation.

Case Study #4: Call cycling performance

Test purpose

During normal network usage, calls are not just set up; they follow the complete call life-cycle described earlier (setup + hold + release). The hold time is controlled by the end-user application.

The release time is another performance aspect of the switch. Often, it is much longer than the setup time and can severely degrade the overall signalling performance.

At what peak level of call cycling activity does my switch or network start slowing down, stop altogether, or crash?

The call cycling performance test provides a much better indication of UNI signalling performance because it simulates more realistic loading conditions.

The *call cycling performance* figure will usually be significantly lower than the *call setup performance* parameter. Although call setup performance is the most widely-quoted performance parameter today, call cycling performance should become the standard benchmark figure in the future!

Brief test procedure

Configure the test equipment to do the following:

- Initiate calls at a gradually increasing rate.
- Ensure that complete call cycles are generated by selecting a finite call duration. A call duration of zero can be used to simulate worst-case conditions.
- Enable automatic Layer 3 *RESTART* at the start of your test to clear any existing calls.

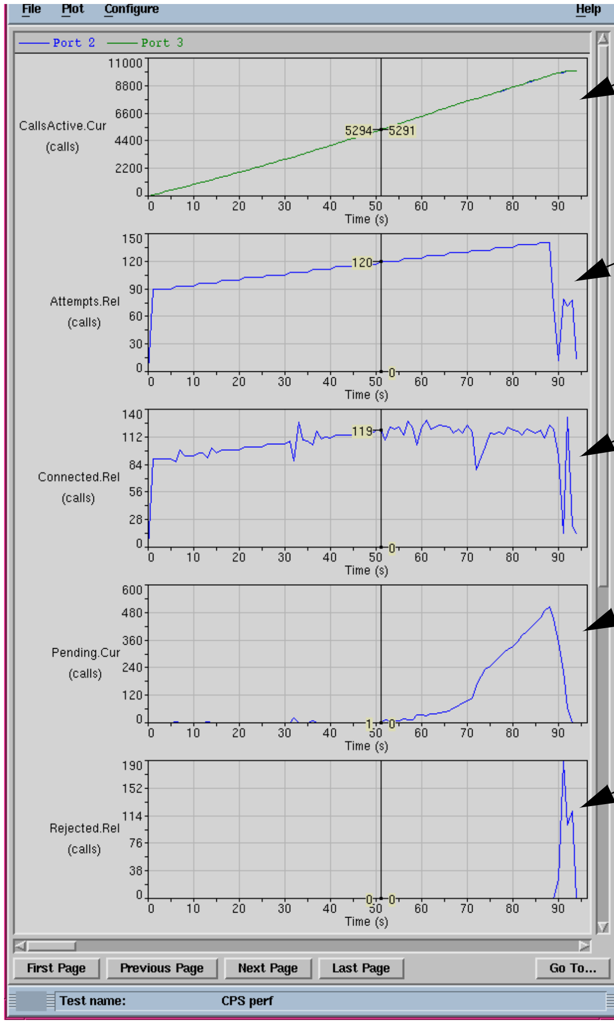

- Active calls approximately equals zero (each call setup is followed by a call tear down).
- Test stimulus: The tester generates setup + teardown cycles (with call hold time = 0). It steadily increases the call cycle rate (from 30 to 50 calls/s).
- The connection rate approximately increases with the attempt rate until it reaches its limit (46 calls/s) at T = 83 s.
- After T = 83 s, the number of pending calls starts to ramp up.
- After T = 88 s, UNI signalling protocol timers start to expire and the switch starts rejecting calls.

Graphs showing the call setup and teardown performance (46 calls/s).
Note: .Cur = current value (calls), .Rel = rate (calls/s).

- Use the log file graphing tool of the MTM to determine the call cycling performance.

Test results

The graph shown here illustrates typical switch behavior under increasing load. In this example, the *Pending.Cur* statistic indicates the point in time at which the SUT has reached its call cycling performance limit. The value of call cycling performance is given by the *Connected.Rel* statistic.

This same test can be performed when a number of ports are calling each other to determine the call cycling performance supported by:

- an entire line card
- an entire switch
- an entire network

using multiple ports on multiple MTM modules.

Choosing the call hold time value

The choice of call hold time value depends on the purpose of the test.

- For benchmarking applications, the hold time should be set to zero. This is the most convenient value to use for comparison purposes. It also measures the switch performance under worst-case conditions.
- For evaluation of network performance under realistic load conditions, the hold time should be set to the amount of time that is typically required by the end-user application. A further refinement is to configure the MTM to generate random call hold times, between specified upper and lower limits.

Detailed test procedure using the E1600A MTM

To test for the call cycling performance on a single switch port, the MTM should be configured in the following way:

- Two ports will be used. One port making calls to the other.
- Ensure the call hold time is a finite value. In this example, the call hold time has been set to zero.
- Enable an Automatic Layer 3 Restart through the Module Wide settings screen.
- Choose a stepped call initiation profile with a gradually increasing rate.

Detailed analysis using the E1600A MTM

The test can be stopped as soon as calls start to be rejected. Observe the statistics and the graph via the log file graphing tool:

- Observing the *Pending.Cur* statistic, there should be a point at which the number of pending calls starts growing without bounds. This is the point at which the call cycling performance of the switch or network has reached its maximum.
- Correlating this point with the *Attempts.Rel* statistic will show the number of call cycles that are being attempted by the MTM. At the same point in the graph, the *Connected.Rel* statistic should level out. This number is the Maximum Call Setup Performance of the SUT.
- If *Connected.Rel* drops off altogether (as it does in this example), then it shows a flaw in the SUT. A switch or network should continue to complete call cycles at the maximum rate it is capable of, even if it cannot match the call attempt rate.
- Some time after the number of pending calls first starts to grow, calls will start to be rejected. This can be seen with the *Rejected.Rel* statistic. This will typically occur between 10 and 30 seconds after the number of pending calls first starts to grow. It represents protocol time-outs, and is a normal response when calls fail to be connected within a specified time interval.

Acronyms

ABR	Available Bit Rate
ATM	Asynchronous Transfer Mode
BSTS	Broadband Series Test System
CBR	Constant Bit Rate
CPU	Central Processor Unit
ILMI	Intergrated Local Management Interface
MIB	Management Information Base
MTM	E1600A Multiport UNI Test Module with UNI signalling performance test software
PNNI	Public Network to Network Interface
QA	Quality Assurance
SUT	System Under Test
SVC	Switched Virtual Connection
UBR	Unspecified Bit Rate
UNI	User to Network Interface
VBR	Variable Bit Rate
VCI	Virtual Channel Indicator
VPC	Virtual Path Connection
VPI	Virtual Path Indicator

This page intentionally left blank.

This page intentionally left blank.

This page intentionally left blank.

Agilent Technologies Broadband Series Test System

The Agilent Technologies BSTS is the industry-standard ATM/BISDN test system for R&D engineering, product development, field trials and QA testing. The latest leading edge, innovative solutions help you lead the fast-packet revolution and reshape tomorrow's networks. It offers a wide range of applications:

- ATM traffic management and signalling
- Packet over SONET/SDH (POS)
- switch/router interworking and performance
- third generation wireless tesing
- complete, automated conformance testing

The BSTS is modular to grow with your testing needs. Because we build all BSTS products without shortcuts according to full specifications, you'll catch problems other test equipment may not detect.

www.Agilent.com/comms/BSTS

United States:

Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:

Agilent Technologies Canada Inc.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
1-877-894-4414

Europe:

Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-9999

Japan:

Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:

Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 267-4245
Fax: (305) 267-4286

Asia Pacific:

Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 2599-7889
Fax: (852) 2506-9233

Australia/New Zealand:

Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

Specifications subject to change.

5968-1235E 04/00 Rev B

Agilent Technologies
Innovating the HP Way