




Functional Testing of ATM Signalling Protocols

Agilent Technologies Broadband Series Test System
Application Note



Introduction

There are generally two parts to making a call within an ATM network. First, the signalling process indicates that a call is to be made, and that network resources should be allocated for this call. Second, the data transfer occurs, which allows the end users or applications to exchange information. The data transfer phase uses the network resources that were allocated during the signalling process to convey voice, data or video information.

Signalling, in turn, can be broken down into two parts – the signalling that occurs between customer premises equipment (such as a workstation) and the network access device (such as a switch), and signalling that occurs between network switches.



Agilent Technologies

Innovating the HP Way

Support for ATM's high-bandwidth applications such as multimedia impacts both the signalling and data transfer phases of the communication process in a significant way-higher complexity. For example, the signalling protocols have become increasingly complicated. In addition, the interaction between the protocols at the user-network interface (UNI) and the network-network interface (NNI), which are both illustrated in Figure 1, are quite detailed. And with that complexity comes a more stringent requirement to test the signalling capabilities for correct operation and function.

This type of functional, inter and intra network testing is essential in order to validate the new functions offered by these evolutionary signalling protocols. Not only do these new signalling protocols provide services that are directly useful by the end user, they also provide the basis for new types of services that will rely upon a solid ATM network foundation. Ensuring that these services have been properly implemented now will save testing headaches in the future.

This solution note will examine the many aspects of signalling through a broadband network, and will focus on the different areas that must be tested in order to create a functional and high performance network. Specific testing concerns will be addressed. Testing methodologies such as conformance and performance testing will also be introduced.




Figure 1: Interaction between the protocols at the user-network interface (UNI) and the network-network interface (NNI).

The focus of this paper is on signalling functional testing. Conformance testing is covered in the paper entitled *Protocol Verification and Automated Testing*. Performance testing is dealt with in the paper entitled *Testing ATM Signalling Performance*.




Figure 2: Protocols at the User-Network Interface.

Principles of Signalling and Signalling Testing

To fully appreciate the complexities of signalling, the operation of the signalling protocols must be examined.

The following sections will explore these details.

Signalling Protocols within an ATM Network

When ATM networks were first implemented, connections between endpoints were permanently established, such that the path between two devices at each end of an ATM network was always active. These permanent connections, or permanent virtual circuits (PVCs), were provisioned by manually configuring the connections at each switch.

PVCs were useful for transferring data between two locations continuously, such as between a remote office and a data center. However, users may want to route information to multiple locations. In addition, the message is not always data—voice and video communication must also be considered, and these signals may go to a large number of destinations.

As ATM technologies evolved, the signalling standards to support switched virtual circuits (SVCs) were developed. These signalling standards provided users with the ability to specify a particular destination for their call, and also to specify the network resources, such as a particular Quality of Service (QoS) that they wanted dedicated to this call. Signalling at the UNI gave users the ability to indicate to the network the parameters for their call. Signalling within the network would subsequently transfer this information between switches.

Protocols at the User-Network Interface

The User-Network Interface (UNI) requires support from a number of protocols, as shown in Figure 2. The lowest layer is, of course, the ATM cell. The next layer, the ATM Adaptation Layer (AAL), maps higher layer protocols into ATM cells. The Signalling ATM Adaptation Layer (SAAL) refers collectively to the AAL-5, Service Specific Connection Oriented Protocol (SSCOP) and the Service Specific Coordination Function (SSCF). The SAAL acts as a reliable transport mechanism for the UNI signalling protocol.

The UNI signalling protocol relays actual signalling information between the user and network. The UNI signalling messages carry the destination ATM address for the call, and call parameters, such as the QoS and the level of AAL support required, requested from the network.



Broadband Series Test System

The Integrated Local Management Interface (previously known as the Interim Local Management Interface), or ILMI, is a Management protocol used at the UNI. ILMI provides basic connectivity and configuration between the user and the network, by providing ATM address registration functions, link keep alive, parameter exchange, and other services. ILMI uses the services of the Simple Network Management Protocol (SNMP) to manipulate databases (known as Management Information Bases (MIBs)) at the user and network sides.

Protocols Within the Network

Similarly, a complete protocol stack is used within, and between networks, as illustrated in Figure 3. The configuration of the protocol stack depends upon the types of networks being connected: public or private.

Public networks are based on a network of switches, signalling transfer points (STPs) which transfer signalling information between switches, and service control points (SCPs), which provide address resolution and routing information, as shown in Figure 4. The STPs transfer signalling information between switches, and queries between switches and the SCPs. Signalling information transferred between switches via STPs uses the Broadband ISDN User Part (B-ISUP) signalling protocol. This protocol relays information to the switches, such as instructions to reserve bandwidth, and to configure virtual circuits between switches. Switches resolve addresses and routing by querying the SCPs via the STPs with the Transaction Capabilities (TCAP) protocol, which is carried within the Signalling Connection Control Part (SCCP) protocol. The Message




Figure 3: A complete protocol stack is used within, and between networks.

Transfer Part 3b (MTP-3b) protocol provides signalling link management, message routing and distribution, plus network management functions. The facilities between switches carry data only.

The private network is a different story. Because private networks are smaller than public networks, specialized STPs do not exist, but the functionality of STPs is maintained within the switches themselves. The Private Network-Network Interface (PNNI) protocol is used within private networks. The protocol consists of two parts – a signalling protocol and a routing protocol, also shown in Figure 3. The routing protocol is used to exchange network information such as topology, reachability of ATM addresses, and congestion, between switches.

By exchanging topological information, the switches organize themselves into a hierarchy composed of peer groups. Using the PNNI routing protocol, each peer group elects a leader. Depending on the size of the network, the peer group leaders can also organize themselves into a peer group, and elect a leader of this group, and so on. Peers in a group, and peer group leaders between groups, then share reachability and network capacity information with one another, so that each switch in the network has an image of where a particular ATM address is located, and how it can be reached through the network. The PNNI signalling protocol carries information pertaining to a single call through the network.

Finally, in order to place calls between networks, networks must be able to exchange signalling information. The Broadband Inter-carrier Interface (B-ICI) protocol does just this – it allows the exchange of signalling information between networks without revealing network topologies. B-ICI is a form of the B-ISUP protocol, and it is transported over the MTP-3b protocol, also shown in Figure 3.

Three Views of Signalling Testing

A thorough test of a signalling system involves three discrete, yet interrelated, aspects: functional testing, conformance testing and performance testing. These three types of tests are necessary to ensure that a protocol implementation works, that a protocol will interoperate with other implementations, and that the protocol implementation delivers the desired level of performance.

Functional testing is used to describe those tests which determine if a protocol or group of protocols function correctly. Functional testing can be as simple as testing call connect procedures at the UNI, or can be as complex as testing interworking between networks. For example, a functional test would ensure that a

call can be made from one port on a switch, through to another port. Functional testing is crucial in order to ensure the correct functionality of a service.

The goal of conformance testing is to give confidence that every aspect of a protocol implementation behaves in accordance to the specifications. To name a few of those aspects: format and syntax of messages exchanged, sequence of messages, abnormal situations, timer durations, and so on. Based on an internationally-recognized methodology (ISO9646), conformance testing provides a series of pre-designed tests to cover all features of a given protocol, such as UNI 4.0 Signalling on a switch. By ensuring that a protocol has been correctly implemented in a device, the interoperability between devices can be assured.

Performance testing tests a device in order to determine what kind of load it can handle, in terms of volume and speed. For example, a switch needs to be tested to determine the number of calls per second that it can handle, as well as the number of simultaneous calls that can be handled. These types of measurements are important in order to characterize the performance of a switch, and to act as important comparison information for evaluating switches. Also, performance information is one metric that is vital for ensuring that subsequent changes to the switch software or hardware are not detrimental to the performance of the switch.



Broadband Series Test System

ILMI Connection, Parameter Exchange and Address Registration

ILMI uses the Simple Network Management Protocol (SNMP) which, as its name suggests, performs network management duties. The basic operation of SNMP can be described as a means through which objects in a database can be manipulated. The user and network both contain databases, known as Management Information Bases, or MIBs, which store a variety of information related to the link between the two endpoints, as shown in Figure 5. The MIB holds information such as the time the link has been up, the last time the link was checked for connectivity, parameters governing the link, and statistics on the performance of the link.

ILMI uses these capabilities to provide a number of services, such as link keep alive, registration and deregistration of end system addresses, methods for exchanging operating parameters and methods for handling the physical movement of end systems between ports on an ATM switch (also known as the change of attachment point procedure).

ILMI offers a number of powerful services to UNI signalling entities, at the usual cost of increased complexity. In this section, some of the nuances of address registration and the change of attachment point procedure will be examined.




Figure 4: An SS7 network with a complex signalling network involving many nodes and multiple networks.




Figure 5: The user and network both contain databases, known as Management Information Bases, which store a variety of information related to the link between the two endpoints.

Address Registration and Deregistration

ATM addresses come in three different types, as shown in Figure 6. Each address consists of three basic parts – the network prefix, the end system identifier, and a selector. The network prefix is the prefix assigned to that switch, or to that region, similar to an area code used within the telephone network. The end system identifier is unique to the terminal equipment, and can either be hard-wired into the equipment, like a Media Access Control (MAC) address used with a local area network, or can be set by the user. The selector may identify a specific process within the terminal equipment (for example, a unique software application).




Figure 6: Each ATM address consists of three basic parts – the network prefix, the end system identifier, and a selector.

There needs to be some way of putting the three parts of the address together, and to make this address known to the network. The address registration procedure is a means through which the network provides the network prefix to the terminal equipment, and the terminal equipment provides the complete address back to the network. Similarly, address deregistration provides the means through which addresses are deregistered with the network, such as when terminal equipment is disconnected from the network.

Implementing Address Registration/Deregistration

The procedure for address registration is as follows:

- Upon start up, the network checks to see if it's prefix is registered with the user by sending an ILMI request to read the status of the network prefix status within the end system's MIB.
- If the network prefix status indicates that the network prefix is not set in the end system, the network resets the end system, and then registers the network prefix with the end system by sending the network prefix to the end system.

- Upon receiving the network prefix, the end system updates its MIB to indicate that the network prefix has been registered. It then appends its end system identifier to it, and sets the ATM address in the network MIB.

Figure 7 shows a captured ILMI address registration exchange from a Broadband Series Test System (BSTS). This data file shows data transmitted (Tx) from the analyzer (the user) and received from the network (Rx).

In this captured data file, the initial parameter exchange can be seen between the two peers (the first object identifier in the list, the `atmfAtmLayerMaxVpiBits` is shown - the other parameters are not). Then, the address prefix status (`atmfNetPrefixStatus`) is read from the user, and the address status is read from the network (`atmfAddressStatus`). Each side returns an indication indicating that the prefix and address respectively have not been set. The network then sets the network prefix (`atmfNetPrefixStatus`). The user responds by prefixing the network prefix to the end system identifier, and by returning the complete address and by setting this address's status to valid. Keep alive functionality continues with each side reading the system up time object (`sysUpTime`) from the peer.

Just like addresses can be registered, and they can also be deregistered, such as when equipment is disconnected from the network. To deregister an address, the end system or network sets the address status in the peer to invalid.



Broadband Series Test System

The ILMI address registration procedures themselves need to be tested. A conformance test suite is useful for fully exercising the address registration protocol. However, from a network standpoint, methods to truly check whether the address has been registered, or deregistered are also necessary.

Testing Address Registration/ Deregistration

The logical connection between the address registration services and the network must be tested. For example, has the deregistration of the address been communicated to the call control mechanisms on the switch? This type of test is necessary in order to ensure that changes in address information are indeed processed correctly by the switch. The ramification of failing to test address registration and deregistration is that the switch (or network) may not be aware of the addition or deletion of an address to its address space. The end result being that calls are not properly routed to destination addresses.

One simple solution is to try to make a call to the deregistered address from another port on the same switch. First, an analyzer with two ports can be connected to a switch. ATM addresses are registered on both ports. In order to ensure that the addresses were registered properly with the switch, a call is made between both ports. Successful completion of the call indicates proper address registration. Next, an ATM address is deregistered on one of the ports. From the other port, a call is made. If the call comes through on the other port, there's something wrong with the address deregistration procedure within the switch.




Figure 7: A captured ILMI address registration exchange from a BSTS shows data transmitted (Tx) from the analyzer (the user) and received from the network (Rx).

Link Keep Alive and Change of Attachment Point Procedures

ILMI is also capable of detecting a change in attachment point. Such a change would occur when two ATM devices are swapped between different ports on a switch. For example, routine network administration or rewiring in a switching closet could swap the connections between two different devices.




Figure 8: Using a protocol analyzer connected to two ports on the switch.

The Change of Attachment Point Procedures

If some sort of regular polling is not present between a switch and an end system, a switch might not detect changes in the status of an end system. In this case, the switch does not detect, for example, the movement of the equipment between ports. Consequently, calls that are open would be left open and the data sent over the virtual circuits would be lost, and calls to a particular ATM address would be routed to the wrong port. Upper layer applications would ultimately find out about the loss of connectivity. Unfortunately, that information would be of little value, as the Application Layer cannot signal the ILMI layer that it should re-register its address.

The physical event would, of course, be detected by the switch along with the associated loss of SAAL connectivity. However, unless such an event exceeds the timeout for reestablishing the SAAL layer (up to a 15 second duration, after which the polling timer (Timer_POLL) will expire, and polling will begin, followed by 7 seconds (Timer_NO-RESPONSE), during which the link is attempted to be reestablished), the event will not be detected. The layer 3 signalling layer will not be affected, nor will the associated virtual circuits be torn down.

A related function of ILMI is to perform the link keep alive. It does this by regularly querying three objects within the peer's MIB: the port identifier, the system identifier, and the system up time. The port identifier object simply holds the number of the port on the switch that the end system is attached to. The system identifier is an identifier unique to the switch or end system. The system up time object keeps track of the total amount of time that the link has been up between the network and the end system.

The ILMI management entity compares these values with the values kept in the local MIB, and if the port or system identifiers are different, or if the value of the system up time is less than the value kept in the local MIB, then the attachment point has changed. The ILMI management entity will then declare that ILMI connectivity has been lost, clear all virtual circuits set up over the link, reinitialize ILMI, and perform address registration again.

As can be seen, there are a number of different aspects that need to be tested. First, the change of attachment point detection procedure must be tested. Next, confirmation is required that the virtual circuits set up between the peers are cleared. Finally, link reinitialization and address re-registration must be tested. The ILMI 4.0 specification, in section 8.3.2, describes these procedures in further detail.



Testing the Change of Attachment Point Procedure

As was mentioned in the previous section, it is necessary to test the change of attachment point procedure in order to ensure that a change at the physical port results in the clearing of all calls setup over that physical port. Failure to ensure this correct operation could result in serious protocol failures at higher layers with the end result that a call is not cleared properly and that data disconnect happens for a longer period than is necessary.

There are two ways to test the change of attachment point detection process. A simple method is for two terminals to be attached to a switch, and then the terminals swapped between the ports on the switch. However, this method requires operator intervention, which makes accurate time-correlation difficult. A better method is to use a protocol analyzer connected to two ports on the switch, as shown in Figure 8. Values in the MIB on the analyzer can be changed, and the effects of these changes on the switch can be noted. For example, the port identifier value in the MIB of the analyzer can be altered, and it can be noted whether this revision causes the switch to undergo the change of attachment point procedures. Also, the system up time (the sysUpTime object) on the protocol analyzer can be changed, and again it can be noted whether this revision causes the change of attachment point procedures to be invoked.

As mentioned above, when the change of attachment point procedures are invoked, all virtual circuits that have been established between the terminal and the network (and, the far end terminal) must be cleared. To test this aspect, a protocol analyzer can be connected as shown in Figure 8. Multiple calls can be established between the two ports. Then, the port identifier on one port of the protocol

analyzer can be changed, and this will invoke the change of attachment point procedure on the switch. The switch should release all calls between the two ports.

After link reinitialization and address re-registration has occurred on the port that changed, a test to confirm that calls may still be made to the terminal should occur. This can be easily tested using the same configuration as illustrated in Figure 8. Make the call from one port to the other, and ensure that the call is accepted, and that data transfer can take place.

In summary, we can see how a protocol analyzer can assist in the testing of the ILMI change of attachment point procedure by giving precise control over the conditions that could cause this procedure to be invoked, and also by providing the appropriate stimuli to ensure that the procedure worked properly.

Other ILMI Tests

There are many other ILMI tests that can be performed. One of these tests is to test what happens when two identical ATM addresses are registered with the network. Upon receiving a set request with an already registered address,

the network should respond indicating a badValue error. A protocol analyzer will provide the control necessary to register two identical addresses, and to evaluate the result.

Another test is to test the cold start procedure. After a cold start, both sides should reset their tables and clear all address information. Both sides should then exchange configuration information and perform address registration.




Figure 9: Timeline for ATM forum deployment of ATM products and services.

Testing Layer 3 Signalling

Another important aspect of functional testing is the verification of proper UNI signalling procedures. The ATM Forum's new UNI 4.0 signalling specification has introduced a number of new services. Each new set of signalling services will ultimately be necessary in order to support a wide range of services to be delivered to the customer. For example, UNI 4.0 introduces the ability of an end station to join a point to multipoint call already in progress. The ultimate customer service could be broadcast video – the end user joins and leaves point to multipoint calls as easily as changing television channels. Ensuring the proper operation of the basic signalling processes will provide a solid foundation for new customer services.

In this section, we will review the evolution of the ATM signalling standards; examine some of the newest UNI signalling features, focusing on leaf initiated joins to point-to-multipoint calls; and examine how to test this service.

The Evolution of ATM Signalling Standards

UNI signalling provides the means through which a user can tell the network that it wishes to make a call, define the intended destination, and specify key call parameters. In Figure 2, we saw how UNI signalling is transported over the SAAL layer. What this diagram does not illustrate are the different variations of UNI signalling.

Two groups were involved in the development of ATM standards: the ITU-T and the ATM Forum. The ATM Forum is motivated to facilitate rapid deployment of ATM products and services, which is why there is a different timeline for their version of signalling specifications. Figure 9 shows a timeline of this development.

The ATM Forum finished their specification first, called the UNI 3.0 specification, published in September 1993. The ATM Forum based their SAAL layer on Q.SAAL1 and Q.SAAL2, the draft precursors to Q.2110 and Q.2130, respectively.

The ITU-T proceeded to evolve the SAAL layer, which became the Q.2110 and Q.2130 recommendations. Unfortunately, the final versions of Q.2110 and Q.2130 diverged from the draft Q.SAAL1 and Q.SAAL2 versions, which meant that the SAAL layers defined by the ATM Forum and the ITU-T were incompatible.



The ATM Forum updated their specification to version 3.1, in September 1994. This version used Q.2110 and Q.2130 SAAL specifications. In February 1995, the ITU-T completed their signalling recommendation, which was published as Q.2931. However, there were a number of differences in the signalling message set and structures between the ATM Forum UNI 3.1 specification and the ITU-T Q.2931 recommendation. For example, UNI 3.1 does not support some message types (ALERTING, NOTIFY etc.); the applicability of some information elements to some messages is different; some fields have been added to information elements (the ATM traffic descriptor in UNI 3.1 contains subfields describing the traffic, such as the forward sustainable cell rate, etc.); and some procedures are different. These differences are significant enough that UNI 3.1 will not interwork with Q.2931.

More recently, the ITU-T has created a number of additional standards for UNI signalling, collectively known as Capability Set 2, phase one (CS2.1). This set of specifications covers direct dialing in, subaddressing, user to user signalling, and other functions. Another ITU-T specification, Q.2971, defined point-to-multipoint (P-MP) signalling.

The ATM Forum has currently advanced their specification to version 4.0 (UNI 4.0). This version adds support for Available Bit Rate (ABR) service, expanded point-to-multipoint capabilities, including leaf initiated joins, parameterized quality of service (QoS) support, and support of the ITU-T's CS2.1. Rather than superseding UNI 3.1, UNI 4.0 is designed to interoperate with UNI 3.1 implementations.




Figure 10: A P-MP call is initiated by one user (called the root) making a normal call to another user, (called a leaf node).

Point-to-Multipoint Calls

One of the key features for ATM service is voice/video conferencing, with applications such as video conference calls, and distance learning. The next sections will discuss the signalling requirements that accompany these features.

UNI 3.1/Q.2931 Point-to-Multipoint Calls

A P-MP call is a means through which one may communicate with, or exchange data with, many destinations (one-to-many communication). UNI 3.1 and Q.2971 both describe the operation of P-MP connections.

A P-MP call is initiated by one user (called the root) making a normal call to another user, (called a leaf node), as shown in Figure 10. The calling party specifies that a P-MP call is desired by including key parameters within the Endpoint Reference and a Broadband Bearer capability information elements that are sent as part of that initial SETUP message.




Figure 11: At the end of the P-MP, the root can remove leaves with the DROP PARTY message, or releases the entire call by sending a RELEASE message to the network.




Figure 12: In a root LIJ, leaves join the call by sending a LEAF SETUP REQUEST to the network.

After the call has been connected to the first party, additional parties can be added to the call with the ADD PARTY message. The ADD PARTY message is transferred through the network to the target user. If the target user can accept the call, an ADD PARTY ACKNOWLEDGE is transferred back to the root, otherwise the call is rejected with an ADD PARTY REJECT.

At the end of the P-MP, the root can remove leaves with the DROP PARTY message, or release the entire call by sending a RELEASE message to the network, which in turn, causes the network to release all parties in the call (see Figure 11).

UNI 4.0 Extensions Supporting P-MP Calls

The P-MP process described above is limited in that it requires the root to add each party (or leaf node) as a separate function.

Such capability is acceptable for a conference call, wherein the addresses of all of the attendees are known by the root node beforehand. However, many conference calls are arranged by sending out a bridge number to all possible attendees, so that they can join the call at their leisure. Also, for video broadcast, the root, or broadcast, node could not possibly know who would be interested in receiving the video feed.

Consequently, the ATM Forum extended the UNI specification by adding the leaf initiated join (LIJ) capability. This allows the users (leaf nodes) to join P-MP calls without an intervention from the root.

There are two ways that LIJ calls can be handled. In the first method, which is called root LIJ, the root node is informed of each leaf that wishes to join the call, and can accept or reject each leaf. This method could be used for conference calls, wherein the root would want to know who is listening in on the call. In the second method, which is called network LIJ, the network can add leaves without intervention from the root. This method could be used for unrestricted video broadcasts, like the 6 o'clock news.

In a root LIJ, leaves join the call by sending a LEAF SETUP REQUEST to the network, as shown in Figure 12. The network transfers the request to the root, and the root decides whether to accept the call or not. If the root accepts the call, it sends a SETUP through the network to the leaf, and then normal call setup procedures are followed.



Broadband Series Test System

In a network LIJ, leaves join the call by sending a LEAF SETUP REQUEST to the network, as shown in Figure 13. Instead of routing the request to the root, the network sets up the call to the leaf using normal call setup procedures. Another unique feature of LIJ calls allows a leaf to join a call before a call is even in place. This feature would be used for joining a broadcast before it starts, or connecting to a conference call in anticipation of it actually starting. To join a call (that has not yet begun), the leaf sends a LEAF SETUP REQUEST as described previously. The request is forwarded to the root node, which responds by sending a SETUP message to the network. Normal call setup procedures are then used between the root and the leaf.

Testing Leaf Initiated Joins

Since LIJ procedures are a relatively new service, it will be necessary to test these functions from the ground up. Testing procedures are complex, and involve separate tests for the network LIJ and root LIJ operations. These two options will be examined in the following sections.

Testing Root Leaf Initiated Joins

Convenience and simplicity are the features of network LIJ. Network LIJ provides the facilities through which a user can easily join, and leave, point to multipoint calls, such as video or audio broadcasts, without the intervention of the root. However, customer satisfaction should be the end objective in ensuring the correct operation of network LIJ - network LIJ services could be used to support broadcasts to hundreds or thousands or easily more end stations. A solid implementation of network LIJ is required to alleviate any future problems when delivering these services.




Figure 13: In a network LIJ, leaves join the call by sending a LEAF SETUP REQUEST to the network. Instead of routing the request to the root, the network sets up the call to the leaf using normal call setup procedures.

When testing network LIJ, the SETUP message sent from the root through the network to the first party must be checked to ensure that it contains a LIJ Call Identifier IE and a LIJ Parameters IE. The LIJ Call Identifier IE contains a unique identifier that is combined with the calling party number and calling party subaddress (if present) to make a value that distinguishes this P-MP call from other calls. This identifier is used by the leaves to indicate the call that it wishes to join. The LIJ Parameters IE is included to indicate to the network that the call is a network LIJ -- that the root is not to be notified when other parties join the call. The LIJ Parameters IE contains a single parameter _ called the screening indication— that indicates this. In addition, the SETUP message must be examined to ensure that it contains the Endpoint Reference and the Broadband Bearer Capability IEs, which indicate a P-MP call.

The next step is to ensure that the SETUP message delivered to the first party is delivered properly. If the Endpoint Reference value is set to zero, the first party has the option of negotiating the values within the Broadband Low Layer Information Element and/or the ATM Adaptation Layer Information Element. The endpoint identifier in all subsequent SETUP messages to other leaves must be unique and must be non-zero (only the first party has the option of negotiating the afore mentioned parameters).

Then, the next step is to add a leaf to the call. The leaf sends out a LEAF SETUP REQUEST message that must contain the ATM address and subaddress of the root, and the LIJ call identifier. It must also include the Leaf Sequence Number IE, which contains a unique number which is relevant only between the leaf and the network, and valid for this one call only.

In a network LIJ call, when the network receives a LEAF SETUP REQUEST message, it alone decides whether to accept or reject the call. This decision is based on whether the P-MP call specified in the request is valid, and whether network resources are available. If the call is accepted, normal call SETUP procedures are invoked towards the leaf, otherwise the leaf is sent a LEAF SETUP FAILURE message.

In this sequence, it is important to test usage and presence of the LIJ call identifier and the leaf sequence number. How does the network know whether the call identifier specified in the LEAF SETUP REQUEST message is valid? How a leaf learns about this identifier is not enumerated in the UNI 4.0 specification.

The identifier could be obtained through a yellow pages service, a well-known resource, or some directory service. Without that identifier, potential customers would

not be able to access the service. This will be an important network addressing and service-related problem.

The leaf sequence number, contained in the LEAF SETUP REQUEST message, will not likely require such rigorous examination. The leaf sequence number is only used to identify a SETUP or ADD PARTY message received in response to a LEAF SETUP REQUEST message. Unless a leaf is joining multiple calls at once (for example, to display multiple video images on a screen), does the sequence number becomes important. The UNI 4.0 specification states that the sequence number should increment each time a separate request is made. If the sequence number does not increment properly, the possibility exists that multiple LEAF SETUP REQUEST messages may not be properly distinguished one another – resulting in possibly incorrect video images on the screen.

Aspects of Testing Root Leaf Initiated Joins

Root LIJs, wherein each LEAF SETUP REQUEST is forwarded to the root node for acceptance/rejection, is another area of concern for testing. The importance of ensuring the correct operation of root LIJ is clear-security. If there is a failure or incorrect operation in the root LIJ procedure, a potential security hole could be left open through which unauthorized callers could join point to multipoint calls, such as a conference call. A secure audio/video conferencing service will rely upon the proper operation of the root LIJ point to multipoint capabilities.



Broadband Series Test System

As described previously, each request is sent from the leaf, through the network, to the root. Then, the root must examine the calling party number, calling party subaddress, and LIJ call identifier to determine whether the leaf should be added to the call. The root then sends out a SETUP message containing the leaf sequence number received in the LEAF SETUP REQUEST, and uses normal call procedures to setup the call to the leaf.

Testing root LIJ is similar to testing network LIJ. The initial SETUP message sent from the root will not contain the LIJ Parameters IE, indicating that the call is a root LIJ.

Subsequent LEAF SETUP REQUEST messages are passed through the network to the root, rather than being handled at the ingress to the network from the leaves: the root will handle call acceptance and rejection itself.

Facilitating Leaf Initiated Join Testing

As can be seen in this rather short description of testing some of the features of LIJ procedures from UNI 4.0, a lot of tests are necessary—from testing the root setup through to testing the addition of leaves to the call, and ultimately stress testing the ability of the switch and root to handle many leaf joins and drops. There are also a number of tests that haven't been discussed here, but they include testing the quality of service from root to leaf, and ensuring that, if enabled, data from leaves is transmitted to the root and other leaves in the network.

This kind of testing will require automation in order to save time. A protocol analyzer with a custom programming environment is one of the best ways of implementing these kind of tests. Each test can be written to test one particular aspect, then these tests can be combined into a large suite that can be used repeatedly to check subsequent releases of network software. If bugs

are found in the switch, a program can be written to test specifically for this bug, and the program can be added to the test suite.

Testing Inside the Network: UNI/PNNI/NNI Challenges

In addition to an analysis of the end user signalling operation, the various interfaces and functions within an ATM network must be tested as well. In this section, signalling interworking issues will be examined, including tracing a call through a PNNI network, and some of the intricacies about NNI testing.

Testing Crankback in PNNI Networks

The architecture of the PNNI was introduced in section 1.1.2. In this section, we will review the PNNI concepts, and then discuss the concepts of testing crankback in PNNI networks.

A PNNI network is simply a network of switches that share address reachability and resource availability information with each other. The PNNI protocol is composed of two portions—a signalling protocol (which is based on UNI 4.0), and a routing protocol. Switches use the routing protocol to: determine who they are connected to (the 'Hello' protocol); to organize themselves into a hierarchical structure (peer group leaders and the election process); and to exchange address reachability information and information on resources available throughout the network. Figure 14 illustrates a logical PNNI network.




Figure 14: Logical PNNI network.

It is important to understand that although each switch has an understanding of address reachability and resource availability throughout the network, this information is 'summary' information. For example, in Figure 14, node A.1 understands that end stations prefixed with 'B' are reachable through the network, but does not know that 'B.1' is in the network. Furthermore, A.1 knows that 'B' addresses are reachable, but it only has knowledge of paths through 'A' to get to 'B'. It does not know the route through B, nor does it know whether there are sufficient resources to allow the call through to B.1.

How Crankback Works

When a call is placed from A.1 to B.1, a SETUP message is sent. This SETUP message contains a special information element unique to PNNI signalling—the Designated Transit List (DTL) IE. This IE contains the routing information that is needed to route the call through the network. Initially, the DTL is composed of: A.1 (the originating node), A.2, A, B

Note that the DTL only contains 'B', rather than the route through B. As described earlier, 'A' knows that it has connectivity to 'B', but does not know the internal routing used within B.

As the SETUP message is transferred through the network, the DTL is

modified at each point. At A.2, the DTL is: A.2, A, B

The SETUP message is transferred to the 'B' group at node B.2.

Furthermore, B.2 knows how to route to B.1 — either through B.4 or B.3. If B.3 is chosen, B.2 replaces the DTL with: B.2, B.3, B.1.

At B.3, the DTL is replaced with: B.3, B.1

And then at B.1 the call is placed through to the end station.

Crankback occurs when the network is congested, and the call must be routed in some other direction. For example, assume the path B.3/B.1 is congested. In this case, B.3, upon receipt of the SETUP message will respond with a RELEASE COMPLETE message indicating that the route is congested. B.2 must then choose an alternate path. For example, it chooses B.4 and routes the call. If B.4/B.1 is also congested, then the call cranks back to B.2, then to A.2. At A.2, the call is routed through C, since it is known that C also has connectivity to B.

Testing Crankback

Why test crankback? It is necessary to test this function in order to ensure that the network can still attempt to route calls even if the network (or links in the network) are loaded. Thus, the network is more robust, and able to function properly under heavy loads.

Testing crankback is not easy. A protocol analyzer is needed to observe the flow of SETUP messages within the network in order to gauge the correct operation of the network.



Broadband Series Test System

First of all, assume that all of these switches are in a captive location such as a switch test lab. Hopefully, the connections between switches are made at a patch panel for easy access to the network connections, and also for easy reconfiguration of the network.

The call originating from the user can be emulated by the analyzer. Thus, the analyzer can provide precise control over the time of the call, and the parameters of the call (such as the requested QoS). Depending on the number of ports available on the analyzer, multiple points throughout the network can be monitored. For example, if monitor points are established at B.2/B.3 and B.2/B.4, the crankback can be observed, as shown in Figure 15.

As mentioned above, crankback occurs because nodes in the network are not aware of the congested links. Network congestion is communicated via PNNI routing protocol topology state packets, but the originating call may have started before receiving updated network topology information. In order to properly test crankback, this event must be forced to occur. One means to achieve this is to force the switch to busy-out the connections between B.3 and B.1. Or, if such control over the switch is not possible, then loads can be generated between B.3 and B.1 that consume available bandwidth. This could be performed by using terminal equipment connected to B.3 to request large amounts of bandwidth, if possible. Otherwise, another protocol analyzer (or, another port on the same analyzer) could be configured to force a load on B.3 to B.1.




Figure 15: If monitor points are established at B.2/B.3 and B.2/B.4, the crankback can be observed.




Figure 16: RELEASE COMPLETE message decoded on a Agilent Broadband Series Test System.

The analyzer should then be configured so that when it sends a SETUP message to switch A.1, it sets up triggers on SETUP messages appearing at B.2/B.3, and another message appearing at B.2/B.4. The analyzer will be used to capture and later display the content of these messages.

As soon as the link from B.3 to B.1 is loaded, the call is placed by the analyzer into A.1. The triggers that have been established at B.2/B.3 and B.2/B.4 should cause an event. A SETUP and RELEASE COMPLETE message pair would result, with the RELEASE COMPLETE message indicating that crankback is occurring. The RELEASE COMPLETE message should contain the Crankback IE with a blocked transit type #2, call or party has been blocked at the succeeding end of this interface. Figure 16 shows the RELEASE COMPLETE message decoded on a Agilent Broadband Series Test System. The resulting SETUP message between B.2 and B.4 should then contain the proper DTL indicating that the call is to be routed from B.2 to B.1 via B.4. The resulting CALL PROCEEDING and CONNECT messages should also be observed between B.2 and B.4.

In summary, the crankback procedure, although only a small portion of the entire PNNI functionality, requires careful testing. Control is needed to instigate the call, and to set up triggers to capture the crankback messages.

Testing Within and Beyond the Network - the NNI

A key part of making a call through an ATM network is what happens within the network itself. Previous sections have discussed PNNI—the protocols used within a private network, and how to test some aspects of PNNI. In this section, we'll take a look at public networks and examine some areas of testing concern within and between public networks.

Public networks are structured very differently from private networks. Public networks have evolved from supported plain old telephone service (POTS) to supporting ISDN, and now broadband services. Modern networks consist of a data network and a separate signalling network. Network signalling protocols within public networks are based on Signalling System number 7 (SS7), defined internationally by the ITU-T.

The SS7 network is a complex signalling network involving many nodes and multiple networks.

An SS7 network is as pictured in Figure 4. Access switches are called signalling points (SPs), as they are the source, or sink, for signalling information. Signalling information is transferred between Signalling Transfer Points (STPs), and relayed to SPs. Signalling Control Points (SCPs) provide address resolution and translation services to the access switches via the STPs. This type of structure makes the individual switches simpler—they just have to pass, and receive signalling information (rather than process, plot routes, and so on).



Broadband Series Test System

The protocols used within an ATM network are shown in Figure 3. Previous sections discussed the ATM and SAAL layers. The MTP-3b interface provides signalling message handling functions (which direct a signalling message to the proper signalling link or higher level function) and signalling network management functions (which control message routing and configuration of the signalling network facilities).

MTP-3b is defined in the ITU-T specifications Q.2210 and Q.704. B-ISUP, the broadband ISDN user part, provides ATM call and connection control management. B-ISUP is defined in the ITU-T recommendations Q.2761 through Q.2764. From an ATM call establishment perspective, B-ISUP is the protocol that we are keenly interested in.

The other two protocols in Figure 3, TCAP and SCCP, are the Transaction Capabilities and the Signalling Connection Control Part. TCAP provides database query and response functions between access switches and Signalling Control Points for the resolution of addresses. SCCP is a Transport layer for the TCAP protocol and provides services to locate the SCPs.

As shown in Figure 3, there is one more interface to consider, the interface between two networks. This interface is known as the broadband inter-carrier interface (B-ICI), and the protocol and procedures used across this interface are as defined in the ATM Forum's B-ISDN Inter-Carrier Interface Specification.

B-ICI utilizes the B-ISUP protocol (over MTP-3b) for communication between networks (note that the SCCP and TCAP protocols are not present across the ICI). The B-ICI carries information between networks and, as such, is the key to interworking between networks.




Figure 17: Configuring a protocol analyzer: at the UNI, the tester can either be used to emulate the calling party, or it can be used to monitor the signalling at the UNI. At the NNI, the analyzer can also be used to emulate a switch, or to monitor the link between two switches.

Testing within the Public Network

A typical problem within a public network may be ensuring that call parameters are accurately transferred within the network. For example, if a call is setup, but the requested quality of service has not been delivered, one possible cause could be the improper transfer of quality of service parameters through the network. The effect is clear—a call appears to be properly set up at the UNI, but the QoS parameters within the network have not been properly transferred. The result is a call quality that the customer may not have wanted.

This test can be done by configuring a protocol analyzer as shown in Figure 17. At the UNI, the tester can either be used to emulate the calling party, or it can be used to monitor the signalling at the UNI. At the NNI, the analyzer can also be used to emulate a switch, or it can be used to monitor the link between two switches.

At the UNI, a call can be made with the SETUP message specifying the appropriate parameters in the ATM traffic descriptor. If the analyzer has been configured to emulate a user, the traffic descriptor values can be changed as needed in order to create the exact circumstances surrounding a particular problem.




Figure 18: To generate potential non-conforming service requests into a network, an analyzer is connected.

At the NNI, the contents of the SETUP message are translated to an Initial Address Message (IAM) – the B-ISUP message passed between STPs and switches to indicate a call. At the NNI, a trigger can be set up to capture all IAMs. The IAM should first be inspected in order to ensure that the called number party parameter matches or is properly mapped from the ATM address specified in the SETUP message at the UNI. Then, the parameters in the ATM cell rate parameter in the IAM should be matched against those in the ATM traffic descriptor in the UNI SETUP message.

Testing Between Public Networks

The interface between networks is a potential source of a number of interoperability problems. Testing at the interface between networks will be crucial for ensuring the proper operation of ATM services between service providers, nationally and globally. A few of the problems that may be encountered are clear—address mapping is one obvious source of problems, as are the translation of QOS parameters, similar to the problem examined in section 4.2.1. There are a few other causes for concern, such as the role of firewalls between networks in order to protect networks. We'll take a look at these internetwork concerns in this section.

The internetwork interface as a network firewall is an area of significant concern for testing. Understandably, every network will

have it's own criteria for accepting or rejecting calls, and unique contracts will exist between networks, so we can only discuss testing this interface in rather general terms.

One test to perform at the inter-carrier interface is to ensure that only calls within traffic contracts are accepted. To generate potential non-conforming service requests into a network, an analyzer is connected as shown in Figure 18. Calls can be generated from the analyzer to request varying amounts of bandwidth. The response of the network under test can be measured in order to ensure appropriate call acceptance or rejection.

A second test to perform at the inter-carrier interface is that of protocol verification – namely, the blocking of network specific messages between networks. Networks often implement carrier-specific features, which are supported within their signalling networks. These features are particular to networks and, as such, must be blocked from being transferred between networks. The procedures to test these features again are network- specific. Such a test would require an analyzer at the UNI to generate calls to destinations outside the network but with calls specifying the use of network specific services. The analyzer should also be connected at the ICI. When a non-conforming call is generated from the analyzer, the call should be rejected within the network.



If the call is detected at the ICI, the call should be checked to see whether the call is progressing with the network specific services request still inside the IAM, or whether the request for network specific services has been dropped. Similarly, if the call is being progressed, has the user been notified about the absence

of the network specific services pertaining to the call?

Although the tests in this section are rather general, this is an area where tests must be performed on the basis of very specific information known about the networks under test, and of the communication permitted between networks.

Is the Connection Really Established?

Setting up a call between two endpoints, as was discussed in the previous sections, is only sufficient to ensure that we have connectivity. It is then necessary to ensure that the call was setup as specified. In other words, we must test for the proper values of the desired quality of service.

This is important in order to ensure that requested parameter have been properly transferred to each of the switches along the path of the call. Performing these type of tests ensures that the customer does indeed get what he or she asks for. Failing to actually check the quality of service delivered by a connection is similar to building a car, checking each component, starting the engine, but not actually driving the car to ensure that it truly works!

This section will discuss the verification of these important parameters.

Traffic Policing

Traffic policing regulates the data flowing into the switch, both from an end user and from within the network. At call setup, the SETUP message contains a number of parameters that indicate the quality of service desired by the caller. Traffic policing continuously monitors the data stream and either changes the priority of non-conforming cells to be available for discard, or discards these cells. Those cells whose priorities are changed, or are discarded, are those that exceed the parameters defined by the Generic Cell Rate Algorithm (GRCA). The GCRA was defined by the ATM Forum, and adopted from the Peak Cell Rate Reference Algorithm, defined in ITU I.371.

B-ISDN services can be broadly categorized into four types: variable bit rate (VBR) services, unspecified bit rate (UBR), constant bit rate (CBR) services and available bit rate (ABR). VBR services are categorized by the variability of the bit rate delivered to the network. UBR traffic is non-realtime traffic that is not sensitive to delay or delay variations. File transfers are an example of UBR traffic. CBR traffic defines bits that are delivered to the network at a constant rate. For example, a 64 Kb/s PCM voice signal would be an example of a CBR service. ABR is a special service wherein traffic sources limit their transmission rate in response to the availability of bandwidth within the network. ABR services are specifically designed for the efficient handling of data traffic.




Figure 19: Stream distribution dialog from a protocol analyzer that controls the transmission of cells according to GCRA parameters.




Figure 20: Results of a GCRA test on the Agilent BSTS.

The parameters that dictate a requested quality of service are the peak cell rate (PCR), and cell delay variation tolerance (CDVT). Variable bit rate traffic has two additional parameters: the sustained cell rate (SCR) and the maximum burst size (MBS). The peak cell rate, in cells/second is the expected maximum rate of transmission of cells. The sustained cell rate (cells per second) is the average cell transmission rate. The maximum burst size sets the maximum number of cells that may be transmitted at the peak cell rate. The cell delay variation tolerance is the maximum amount of delay that can be expected between cells.

These parameters are specified within the SETUP message. The PCR, SCR and MBS are specified within the ATM Traffic Descriptor IE. The CDVT is specified either within the quality of service classes specified in the Quality of Service Parameter IE, or, in UNI 4.0, is specified explicitly within the Extended QoS Parameters IE.

In order to assist with traffic policing, ATM cells contain a cell loss priority (CLP) bit. This bit is set to zero for high priority cells. If set to one, it indicates a lower priority cell that can be dropped, if necessary.

As data enters the switch, the profile of the traffic is compared to the parameters requested. If the profile is within the tolerances specified by the traffic, the data is forwarded to the destination.

If the profile does not match, non-conforming cells can be treated in three different ways: cells whose cell loss priority (CLP) is set to zero can be set to one, indicating that they can be dropped if congestion occurs within the network; cells with their CLP already set to one can be dropped if congestion is occurring; or non-conforming cells can just be dropped.

In order to test that a signalling connection has been established that conforms to the expected QoS parameters, it is necessary to provide a traffic source that conforms to these parameters.

To test the boundaries of the QoS, a non-conforming traffic source has to be generated.



Testing Call QoS

In order to detect the loss of cells between endpoints, special test cells are transmitted. These test cells contain information on the number of cells with a CLP of zero, and the number with a CLP of one, in a test sequence, along with the current CLP of the test cell (in case the CLP is changed by the network). The test cells can be generated and transmitted according to GCRA parameters. Figure 19 shows the stream distribution dialog from a protocol analyzer that controls the transmission of cells according to GCRA parameters. Note that full control over all GCRA parameters is provided, as well as the option to optimize (stress) the cell rate or the burst size.

These test cells are captured and analyzed by the test application. The test application can display the cell loss count for cells with CLP of zero or one, tagged cells (cells whose CLP has been changed), cell misinsertion rate, and errored cells. The received traffic can also be analyzed to determine whether it exceeded the bounds of the GCRA, the cell delay for cells in the measurement period, and the cell delay variation. Figure 20 shows the results of a GCRA test on the Agilent BSTS.

This analysis determines whether the traffic received matched the traffic transmitted, and whether the connection established by the call does indeed match the requested QoS parameters.

Conclusion

In this paper, three different types of testing were introduced: functional testing, conformance testing and performance testing. Each of these types ask different questions: functional testing (does it work?), conformance testing (is the communications protocol properly implemented?), and performance testing (how well does it work?).

This solution note concentrated on the functional testing aspects, where the testing focuses on whether an aspect of an ATM signalling protocol works. Two other solution notes will discuss the Conformance and Performance testing aspects.

The emphasis here has been the breadth and extent of signalling testing, and signalling testing issues.

Testing the ILMI is particularly important, because of the functions that it provides to the UNI signalling layer (keep alive, address registration), and also of the effect that ILMI procedures can have on the network in terms of address presence.

UNI 4.0 has added many new capabilities to the user/network interface, and has added new test challenges. Leaf initiated joins are one of these new features, and we examined how leaf initiated joins work, and how they can be tested.

The Private Network Interface provides key network functionality for exchanging network topology information and address reachability information between switches, and provides call routing capabilities. Crankback is one of the facilities that PNNI uses to attempt to route a call even though routing information may be incorrect. Although testing crankback is difficult, they are necessary in order to determine the correct operation of this feature that allows a network to route calls even under a fast changing and dynamic network load.

In the public network, the translation of call quality parameters at the UNI to B-ISUP parameters is a potential source of error. Likewise, this same problem could occur at the inter-carrier interface. Also, at the inter-carrier interface, firewalls must be tested in order to ensure the correct acceptance and rejection of calls coming into the network, and the correct transmission of call parameters outside of the network.

Quite a few different areas have been covered in this paper – from ILMI testing at the UNI, through to testing between networks at the inter-carrier interface. For each of these areas, we've only taken a look at a few possible testing scenarios. Certainly there is a variety of other tests that can, and should, be performed in order to ensure proper functional operation of an ATM signalling service.




Broadband Series Test System

References

- Traffic Policing*,
BSTS Solution Note 5963-7510E
Q.2761,
- Testing Policing in ATM Networks*, Mohammad
Makarechian and Nicholas
Malcolm, to be published August
1997, Agilent Journal.
- ATM User-Network Interface (UNI)*
Specification, Version 3.1, The
ATM Forum, 1995.
Q.2762,
- ATM User-Network Interface (UNI) Signalling*
Specification, Version 4.0, The
ATM Forum, af-sig-0061.000, July,
1996.
- Private Network-Network Interface*
Specification Version 1.0 (PNNI
1.0), The ATM Forum,
af-pnni-0055.000, March, 1996.
Q.2764,
- B-ISDN Inter Carrier Interface (B-ICI)*
Specification Version 2.0
(Integrated), The ATM Forum,
af-bici-0013.003, December, 1995.
Q.2931,
- Integrated Local Management Interface (ILMI)*
Specification Version 4.0, ATM
Forum/95-0417R9, July 1996.
- Q.704,
Signalling System No. 7 – Signalling
Network Functions and Messages,
ITU-T, 03/93.
- Q.2110,
B-ISDN ATM Adaptation Layer –
Service Specific Connection
Oriented Protocol (SSCOP), ITU-T,
07/94.
Q.2971,
- Q.2130,
B-ISDN Signalling ATM Adaptation
Layer – Service Specific
Coordination Function For Support
of Signalling At the User Network
Interface (SSFC at UNI), ITU-T,
07/94.
- Broadband Integrated Services*
Digital Network (B-ISDN) –
Functional Description of the
B-ISDN User Part (B-ISUP) of
Signalling System No. 7, ITU-T,
02/95.
- Broadband Integrated Services*
Digital Network (B-ISDN) – General
Functions of Messages and Signals
of the B-ISDN User Part (B-ISUP) of
Signalling System No. 7, ITU-T,
02/95.
- Broadband Integrated Services*
Digital Network (B-ISDN) –
Signalling System No. 7 B-ISDN
User Part (B-ISUP) -Basic Call
Procedures, ITU-T, 02/95.
- Broadband Integrated Services*
Digital Network (B-ISDN) - Digital
Subscriber Signalling System No. 2
(DSS 2) – User-Network Interface
(UNI) Layer 3 Specification for
Basic Call/Connection Control,
ITU-T, 02/95.
- Broadband Integrated Services*
Digital Network (B-ISDN) – Digital
Subscriber Signalling System No. 2
(DSS 2) -User-Network Interface
(UNI) Layer 3 Specification for
Point-to-Multipoint Call/Connection
Control, ITU-T, 02/95.

Acronyms

AAL	ATM Adaptation Layer	PNNI	Private Network-Network Interface
ABR	Available Bit Rate	POTS	Plain Old Telephone Service
ATM	Asynchronous Transfer Mode	PVC	Permanent Virtual Circuit
B-ICI	BISDN Inter Carrier Interface	QoS	Quality of Service
B-ISUP	Broadband ISDN User Part	Rx	Receive
BSTS	Broadband Series Test System	SAAL	Signalling ATM Adaptation Layer
CBR	Constant Bit Rate	SCP	Service Control Point
CDVT	Cell Delay Variation Tolerance	SCCP	Signalling Connection Control Part
CLP	Cell Loss Priority	SCR	Sustained Cell Rate
CS2.1	Capability Set 2	SNMP	Simple Network Management Protocol
DTL	Designated Transit List	SP	Signalling Point
GRCA	Generic Cell Rate Algorithm	STP	Signalling Transfer Point
IAM	Initial Address Message	SS7	Signalling System Number 7
ICI	Inter Carrier Interface	SSCF	Service Specific Coordination Function
IE	Information Element	SSCOP	Service Specific Connection Oriented Protocol
ILMI	Integrated Local Management Interface	STP	Signalling Transfer Point
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector	SVC	Switched Virtual Circuit
LIJ	Leaf Initiated Join	TCAP	Transaction Capabilities
MAC	Media Access Control	Tx	Transmit
MBS	Maximum Burst Size	UBR	Unspecified Bit Rate
MIB	Management Information Base	UNI	User-Network Interface
MTP-3b	Message Transfer Part 3b	VBR	Variable Bit Rate
NNI	Network-Network Interface		
PCR	Peak Cell Rate		
P-MP	Point-to-Multipoint		



Broadband Series Test System



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 testing
- 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

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

Copyright ©2000 Agilent Technologies

Specifications subject to change.

5965-9331E 04/00 Rev C



Agilent Technologies

Innovating the HP Way