




Testing Operations & Maintenance (OAM) Implementations for ATM

Agilent Technologies Broadband Series Test System
Application Note



Introduction

ATM is being rapidly deployed into the backbones of both corporate and public networks. The backbone is a mission-critical area which demands low downtimes and performance guarantees.

To assist with these critical functions, the International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) defined the Operations and Maintenance (OAM) protocol. This protocol provides for in-service fault detection, fault localization and performance monitoring for ATM networks. To optimize your network's performance, OAM should become an integral component of an overall ATM network monitoring and network management strategy.



Agilent Technologies

Innovating the HP Way

Networks constructed with equipment from multiple vendors require non-proprietary network monitoring strategies to provide a stable base upon which to construct a reliable network. OAM provides a key building block for this base, by defining standardized data elements which are sent within the user data stream to provide particular information about a network and its elements. For every virtual-path or virtual circuit connection OAM can provide:

- Fault detection
- Fault localization
- Performance monitoring

OAM provides a standardized method for in-service network monitoring and fault management. It offers automatic fault detection, and when faults are found, other features of OAM may be used to further isolate that fault. Thus, network managers using OAM can reduce their network's downtime.

OAM also offers performance monitoring which is crucial to the reliable transport of voice and video traffic signals which require a guaranteed quality-of-service (QoS).

This paper begins with a discussion of the features of OAM and the operation of the protocol. Then test strategies for validating OAM implementations are presented.

OAM Standards

OAM is specified in two standards:

- ITU-T Recommendation I.610, B-ISDN Operation and Maintenance Principles and Functions, December 1995.
- Bellcore GR-1248-CORE, Generic Requirements for Operations of ATM Network Elements, Issue 2, September 1995.

The ITU-T document presents the framework for implementing OAM. The Bellcore document expands on the ITU-T document by providing test methodologies for OAM.




Figure 1: ITU-T considers an ATM network to consist of five flow level

The OAM Protocol

The OAM Protocol, defined by the ITU-T, consists of a number of key elements, including logical flows of management information, and the formats for the cells that carry that management information. These two areas will be examined in the next two sections.

OAM Flows

The ITU-T considers an ATM network to consist of five flow levels. These levels are illustrated in Figure 1.

The lower three flow levels are specific to the nature of the physical connection. The ITU-T recommendation briefly describes the relationship between the physical layer OAM capabilities and the ATM layer OAM.

If the physical layer is SONET or Synchronous Digital Hierarchy (SDH), the F1 and F2 flows are carried by the section overhead. The F3 flow is carried by the path overhead.




Figure 2: Simplified network diagram showing two switches and four end stations

If the physical layer is from the Plesiochronous Digital Hierarchy (PDH), such as DS1, DS3 etc.), F1 is the frame alignment bytes and the remainder of the overhead is F3. There is no F2 flow in a PDH connection.

From an ATM viewpoint, the most important flows, and the ones that the remainder of the paper will concentrate on, are known as the F4 and F5 flows. The F4 flow is at the virtual path (VP) level. The F5 flow is at the virtual channel (VC) level. When OAM is enabled on an F4 or F5 flow, special OAM cells are inserted into the user traffic.

To understand why there are two flows at the ATM level and not just one, consider the simplified network diagram shown in Figure 2. In this network, there are two switches and four end stations.

For example, suppose the following connections have been set up between A1 and A2 and between B1 and B2. Additionally, these two channels are transported within the same virtual path between the switches.

If a service provider wanted to monitor the health of the connection between the switches a single F4 flow could be enabled. If each connection is to be monitored from end-to-end, use two F5 flows.

Why would a network manager want to only monitor a virtual path and not the entire end-to-end connection for all circuits? Two reasons: cost and complexity.

When OAM is enabled on a circuit special cells are inserted into the user data flow. The bandwidth of these cells is quite low, but every OAM cell transmitted is a lost opportunity for a user-data cell to be sent.

The complexity issue is one of trying to keep track of thousands of virtual channel connections (VCC) rather than hundreds of virtual path connections (VPC). Often monitoring of only F4 flows will provide sufficient information to a network manager.

As a further refinement, OAM may be enabled on only a portion (Segment) of a VCC or VPC. Segment OAM flows facilitate the management of a particular piece of the connection. This allows for problem localization such as failure detection or segment performance surveillance.

If different service providers are involved in an end-to-end connection, each provider may use segment OAM flows to support operation and maintenance of their particular segment in the link.

The F4 and F5 flows will now be described in greater detail.



F4 Flow

The F4 flow is used for operation and maintenance of VP connections.

There are actually two F4 flows: end-to-end flows and segment flows. An end-to-end flow, as the name implies; provides monitoring for a connection over its entire length. A segment flow is for monitoring a connection over a portion of the connection, e.g. between switch nodes.

OAM cells inserted into an F4 flow have the same VPI of the user traffic and a standardized VCI. The value of the VCI determines the type of OAM cells:

VCI Use

- 3 Segment OAM F4 flow cell
- 4 End-to-end OAM F4 flow cell

Intermediate points along the VPC may monitor OAM cells passing them through them. A node must be a termination point for it to be able to insert new OAM cells.

F5 Flow

The F5 flow is used for the operation and maintenance of virtual channel connections.

OAM cells for an F5 flow have the same VPI and VCI as the user cells of the VCC. The payload type indicator (PTI) is used to distinguish between user cells (PTI = 000-011 binary) and OAM cells (PTI = 100,101 binary).

As with the F4 flow there are both end-to-end and segment flows defined:

PTI Use

- 100 Segment OAM F5 flow cell
- 101 end-to-end OAM F5 flow cell




Figure 3: General format of an OAM cell

OAM Type	OAM Function	Main Application
Fault Management	AIS (alarm Indication signal)	For reporting defect indications in the forward direction
	RDI (remote defect indicator)	For reporting remote defect indications in the backward direction
	Continuity check	For continuously monitoring the availability of a link
	Loopback	For on-demand connectivity monitoring For fault localization For pre-service connectivity verification
Performance Management	Forward Monitoring	For estimating performance over a segment of end-to-end on a link
	Backward reporting	For reporting performance estimations in the backward direction
Activation /Deactivation	Performance monitoring A/D Continuity check A/D	For activation/deactivation of PM and/or CC functionality in a standard way
System Management	Not specified in I.610	For use by end-systems only

Figure 4: OAM-type and Function-type fields are used to distinguish the type of OAM cells. *Note: Not to be standardized by ITU-T I.610.*




Figure 5




Figure 6: Reporting defects using AIS and RDI cells

AM Cell Types

Four types of OAM cells are defined to support the management of VP/VC connections:

- **Fault Management OAM cells.** These OAM cells are used to indicate failure conditions. They can be used to indicate a discontinuity in VP/VC connection or may be used to perform checks on connections to isolate problems.
- **Performance Management OAM cells** used to monitor performance (QoS) parameters such as cell block ratio, cell loss ratio and mis-inserted cells on VP/VC connections.
- **Activation-deactivation OAM cells.** These OAM cells are used to activate and deactivate the generation and processing of OAM cells, specifically continuity check (CC) and performance management (PM) cells.

- **System management OAM cells.** These OAM cells can be used to maintain and control various functions between end-user equipment. Their content is not specified by I.610, and they are limited to end-to-end flows.

The general format of an OAM cell is shown in Figure 3. The header indicates which VCC or VPC an OAM cell belongs to. The cell payload is divided into five fields. The OAM-type and Function-type fields are used to distinguish the type of OAM cell, such as a Fault Management or Performance Management cell, as shown in Figure 4. The Function Specific field contains information pertinent to that cell type. A 10 bit Cyclic Redundancy Check (CRC) is at the end of all OAM cells. This error detection code is used to ensure that management systems do not make erroneous decisions based on corrupted OAM cell information.

Fault Management

Fault management cells consist of alarm indication signal (AIS), remote defect indication (RDI), continuity check (CC) and loopback cells.

AIS cells are used to report defects in the forward direction. RDI cells report defects in the backward direction. This is shown in Figure 6.

Loss of continuity (LOC) triggers AIS cells to be sent in the forward direction. The end node of a connection will send RDI cells in the backward direction.



Broadband Series Test System

Continuity check cells allow for continuous monitoring of a connection's availability. They are periodically inserted into a user data stream. Intermediate nodes can be set to look for the presence (or lack) of these CC cells.

Loopback cells are used to determine connectivity at specific points in a network, shown in Figure 7. This may be done for monitoring, fault localization, or pre-service connectivity verification.

Fault management has only been defined for end-to-end defects. Segment fault management remains to be defined.

Activation and Deactivation

The Activation/Deactivation cells, shown in Figure 8, enable and disable PM and CC OAM cell flows. OAM cell flows may also be enabled through the network management system of a switch.

In addition to enabling OAM flows, these cells are used to establish agreement on the block size (number of user cells between OAM cells), the nature of the OAM monitoring, and the direction of transmission. The nature of activation and deactivation requests is a negotiation between endpoints.

The activation request endpoint sends an OAM activation-deactivation cell to the other far end point. The far end then must reply with a confirmation or denial.

Performance Management

Most ATM traffic classes have associated QoS parameters. In production networks, end users and end-user applications will assume that the ATM network is meeting its pre-negotiated QoS.




Figure 7: Loopback cells are used to determine connectivity at specific points in a network. *Note: Values (except for default all 1's) are not subject to standardization and encoding of non-default values is optional.*




Figure 8: Activation/Deactivation cells

OAM performance management allows for in-service monitoring of virtual paths and virtual channels. Performance monitoring may be used as an ongoing monitor on selected channels for service monitoring or it may be individually activated in response to complaint from a customer or end-user.

OAM performance management allows a network provider to measure the QoS of a link without having to take the link out of service, or to use out-of-service test procedures which may not accurately predict the actual performance of a connection under actual traffic conditions. These in-service measurements may also be used for trend analysis and to pinpoint potential problems before they lead to major service outages.




Figure 9: Two kinds of OAM performance management cells are defined

Two kinds of OAM performance management cells are defined:

- the Forward Monitoring OAM PM cell,
- and the Backward reporting OAM PM cell,

as shown in Figure 9.

Forward monitoring PM cells transmit redundant data to the far side of a network. This allows network elements along the link to calculate quality metrics over the particular user-cell stream. OAM PM cells are inserted after a block of user cells.

The PM cell contains information about the cell block:

- Number of CLP0+1 user cells
- Number of CLP0 user cells
- A time stamp
- An even parity BIP-16 for bit error detection over the information fields of the user-cell block

Different block sizes are supported. The actual size of the user-cell block between OAM cells may be only

approximately the size of the defined block size. Depending on the user traffic, an OAM source can insert PM cells between $N/2$ and $3N/2$ (where N is the nominal block size). The standard allows a monitoring cell to be inserted up to $N/2$ cells after an insertion request has been initiated.

The end device counts the cells between forward monitoring PM cells and compares it to the forward PM cell content. It then checks the sequence number and recalculates the BIP-16 over the user cells. The result of this analysis together with some of the original forward monitoring PM data is mapped into a backward reporting OAM PM cell and sent back to the source. This data can be extracted by any entity on the link and analyzed to provide QoS information at the network management level.

Results that are available from OAM PM cells are:

- Severely Errored Cell Blocks (SECB)
- Errored cell count
- Lost CLP0+1 user information cells
- Lost CLP0 user information cells
- Mis-inserted user information cells
- Number of transmitted user cells (CPL0+1 and CLP0)
- Impaired blocks
- Transmission delay



Broadband Series Test System

OAM and Network Management

ATM networks offer significant challenges to network management. These include large network topologies with long distances and high node counts, high speeds, multiple vendors of network equipment and multiple service providers for a single connection.

OAM offers a standardized toolbox for fault detection, fault localization, performance monitoring and connectivity verification. With an OAM-enabled ATM network, a superior network management system may be built, capable of detecting and isolating network faults, and monitoring the QoS of any connection in real-time.

Using Continuity Check Cells

For example, consider the continuity check cell. In a non-OAM environment, a network node could try to detect disruption to its received user traffic by monitoring the regular arrival of user cells. This fault detection method assumes a relatively constant flow of user traffic.

But what if the traffic is bursty, as shown in Figure 10? This monitoring method would fail. Logical ATM connections may have a variety of traffic profiles: from constant to bursty. Simplistic monitoring methods can fail to detect the loss of traffic.

To provide constant monitoring of a connection, a network manager can configure an upstream network node to periodically insert continuity check (CC) OAM cells into the cellstream. Using the example shown in Figure 10, the traffic profile is very bursty. The insertion (and subsequent verification) of continuity is solely provided by the presence of CC cells.




Figure 10: A network node to periodically insert continuity check (CC) OAM cells into the cellstream

The reliability of the system from a customer perspective occurs at the cost of additional management bandwidth from the service provider of about 1% of cells per connection.

Using Loopback Cells

Loopback cells are a valuable diagnostic tool, and may be sent between any node in the network. This functionality is especially useful across boundaries between networks, such as between different service providers. In order to verify that the connection between customer premise equipment and the network is available, the CPE needs to support loopback functionality.

The utilization of loopback cells also allows network providers to localize the point of discontinuity, or play an active role in verifying a new connection or network topology. To take full advantage of loopback cells for end-to-end verification, the loopback function has to be deployed into both network and end user equipment.

Using Performance Management Cells

ATM networks carry time sensitive voice and video. For these types of services mere presence of a connection does not guarantee that user information will arrive in a useable form. This type of traffic requires QoS. OAM may be used to passively monitor the QoS of a network connection and detect potential problems.

PM OAM cells are used to monitor performance parameters such as cell block error ratio, cell loss ratio and mis-inserted cells on an active connection by inserting special cells after a block of user cells. These PM OAM cells contain information about the user-cell block in-between, such as number of cells and a parity check sum over these user-cells. It may also include a time stamp to make delay and delay variation measurements. At the endpoint of an OAM PM flow, the data is recalculated and a comparison to the OAM PM cell determines the performance parameters.

ATM OAM functionality provides a network manager with a set of important tools to manage and troubleshoot a network from a customer perspective (per connection) and without the need to take a link out of service and therefore without disturbing well functioning connections of thousands of other users. The standardized nature of OAM allows these tools help to locate issues and problems even over the boundaries of networks and right into the customer premise.

Testing OAM Implementations

Testing an OAM implementation is a complex undertaking. In the first place, the testing must be performed in a real-time environment which demands dedicated test hardware. In addition, OAM is a state-machine type of protocol, in which certain stimuli require specific responses. It is a complex task to ensure that a piece of equipment transitions through every possible stimuli-response combination. In this section, testing methodologies for validating fault management and performance management will be presented.

Fault management OAM function testing requires the generation of a large number of test cases. This can be tedious if an ATM analyzer does not support the automatic generation of OAM cells.

OAM performance management testing requires OAM-specific test equipment to create the time-sensitive fields and do real-time turn-around of OAM cells.

Fault Indication and Alarm OAM Testing

The network manager must ensure that fault indications and alarms are functional before commissioning a network. The fault detection and alarm tests may be grouped into three areas:

- Physical layer test. OAM alarm propagation into the ATM layer (F1 to F3 alarms triggering F4/F5 OAM generation).
- ATM layer test. OAM alarm generation and distribution (F4 alarm distribution, performance and management accuracy.)



Broadband Series Test System

- OAM information reporting into the network management system.

The objective of these tests is to validate that the system-under-test (SUT) enters, sustains and exits fault states correctly.

To do this, the SUT must be subjected to faults which cause state changes. These include:

- Propagating lower layer errors, such as detaching the cable from the SUT.
- Creating error messages from physical levels.
- Creating an F4 AIS on a VP.
- Removing continuity check cells from a channel.

Figure 11 illustrates how a physical layer alarm eventually causes both path (F4) and channel (F5) OAM RDI's. Note the sequence of operations:

1. An intermediate node detects loss of continuity.
2. The intermediate node issues a F4 or F5 AIS OAM cell.
3. The end node detects this AIS cell and enters into the AIS state.
4. The end node, now in the AIS state, issues an F4 or F5 RDI for as long as it receives AIS cells from the intermediate node.

Testing the correct propagation of lower level alarms requires physical OAM testing capabilities. A sophisticated ATM analyzer can generate and analyze F1 to F3 OAM AIS on different SDH and PDH interfaces, in order to stimulate the generation of F4 and F5 OAM AIS or RDI cells.

For example, the Agilent Broadband Series Test System (BSTS) is well-equipped with these F1 to F3 level testing capabilities.




Figure 11: A physical layer alarm eventually causes both path (F4) and channel (F5) OAM RDI's




Figure 12: AIS OAM cells should be transmitted on the downstream connection to alert the connection end point

Loss of continuity is the only ATM-specific alarm that is currently defined, but more are expected in the future.

In the following sections a number of test procedures will be described to clarify how to test OAM Loss of Continuity (LOC).




Figure 13: Stops generating AIS or RDI cells after receiving the first continuity check OAM cell

Test scenario: Verify LOC Detection

If neither user traffic, nor continuity check OAM cells are received on a virtual connection for a long period of time (3.5 seconds +/- 0.5 seconds, according to I.610), a network connection point should assume loss of continuity (LOC). The network's node response to the detected fault will depend on its location.

If the fault is detected by an intermediate connection point, AIS OAM cells should be transmitted on the downstream connection to alert the connection end point. This is shown in the upper portion of Figure 12.

If the fault is detected by a connection end point; RDI OAM cells should be sent on the matching upstream connection to alert all connection nodes of the defect. This is shown on the lower portion of Figure 12. The transmission of RDI and or AIS OAM cells should end when user traffic resumes.

The test goal is to verify that an intermediate connection point or a connection end point in the SUT correctly enters and exits a F4/F5 LOC fault state. This may be done by temporarily suspending the

transmission of continuity check OAM cells to the SUT.

Criteria for a successful test is the verification that the SUT:

- Enters a LOC state 3.5 +/- 0.5 seconds after the last received continuity check OAM cell,
- starts transmitting valid AIS or RDI cells on the appropriate VPC or VCC,
- transmits AIS or RDI cells at a nominal rate of one cell per second during the LOC or AIS alarm condition, and
- stops generating AIS or RDI cells after receiving the first continuity check OAM cell (shown as delta T2 in Figure 13). Typical values of T1 and T2 would be 3.5 seconds +/- 0.5 seconds, and 100 milliseconds, respectively.

This is only one example of testing AIS and RDI of an ATM system. Other tests include:

- Verification of AIS state transitions. Confirm that AIS and RDI cells are forwarded and handled properly and that each element correctly enters, sustains and exits alarm conditions.
- Verification that errors and error statistics are reported to the management system properly.
- Verification of protection switching functions.

Test scenario: Loopback

Loopback testing allows a faulty ATM network element to be isolated. Figure 14 show some of the possible loop-back scenarios. The simplest loopback is an end-to-end test which verifies the connectivity of an entire connection (shown in the upper portion of Figure 14).



Other loopback may be used to isolate the source of a connection fault. For example, the lower portion of Figure 14 illustrates an access line loopback. This test could be initiated by a customer to verify that their equipment maintains a connection to the outside world.


A network device can be configured as connection or segment source and sink for OAM loopback cells. A device under test does not have to be a termination point but can support the loopback ID feature by responding to loopback cells which contain a dedicated loopback ID.

To verify loopback sources, a SUT must send a loopback cell, and then receive and interpret the returned loopback cell. A segment or end-to-end loopback cell needs to be sent on an existing connection. Either the segment, or the end point, should take the cell and send it back to the source. In order to verify that the cell was looped back on the ATM layer and not on the physical layer, the loopback indication bit, which is in the first octet of the loopback cell, should be examined (review Figure 7).

Testing an OAM loopback source requires that a loopback request is activated from the management environment of the switch. The SUT should generate loopback cells and then wait for a response.

A test system needs to emulate the loopback response. It needs to have a behavior which can be configured; such as responding to every second or third loopback request cell, or responding with a delay. The test system also needs to be able to respond with custom-defined loopback response cell. This allows a stress test of the SUT's implementation. For example configuring a test instrument to not return OAM loopback cells would be used to verify that the time-out function of the SUT works correctly.

- **End-to-end loopback**



- **Access line loopback**

- b1 = Initiated by the Customer
- b2 = Initiated by the Network




Figure 14: Possible loopback scenarios

In many cases, a switch's management interface may have the capability of generating loopback report that may be used to determine if the SUT is behaving properly.

To test loopback sink, a test system is required to emulate a source which generates loopback requests and monitors the correct response. The test system validate that the loopback response cell was sent and measure the response time.

User defined loopback cells allow for testing of the response to a dedicated loop back ID.

Here is a summary of the loopback functions that need to be validated:

- Verifying correct loopback generation and analysis by using a test system to emulate the requesting or responding device on many virtual channels
- Configurable response behavior (such as responding to only every second or third loopback request cell, or responding with a user-configurable delay)

- Loading loopback requests on many VCCs and VPCs to stress loopback response emulation on an interface (including response time analysis)
- Editing loopback cell content to test optional loopback addressing

Testing Performance Management

Performance management functions need to be tested before the results they produce can be trusted by a network manager. To test a performance management (PM) implementation, the following need to be verified:

- Are forward monitoring OAM PM cells generated correctly?
- Are backward reporting OAM PM cells generated correctly?
- Is the block size maintained?
- Does forward monitoring PM generation add any delay?
- Are the right cells considered for PM cell content calculation (F4 includes F5)?
- Does a forward monitoring OAM PM cell receiver calculate the difference between the actual stream and the OAM correctly and accurately?
- Are these results reported to the network management system accurately?
- Is the accuracy of a receiver affected when there are multiple PM streams?

Test Scenario: Forward Monitoring PM Cell Generation

The first test in the verification of forward monitoring PM functions is to see if a network element can generate these types of cells accurately. This test involves more than simple syntax testing; it should verify that the content of a PM cell was calculated accurately, considering the user cell traffic.

To allow the PM cell insertion without delaying or buffering of user cells, cell block size is not fixed, but allowed to vary $\pm 50\%$ around the defined target size (such as 1024, 512, ... cells). A key measurement is the analysis of the actual block size used by the system.

As a network element, such as a switch, becomes loaded with more traffic or more active connections, it must calculate OAM PM cells over multiple flows. A load test may be used to ensure that the switch's implementation of PM cell generation is not affected by the number of channels. It is important that a test system calculates statistics on the average and maximum size of the user cell blocks to verify that the PM generation of a network element chooses a correct user cell block size.

PM cell generation should not be dependent on any internal buffers in the switching element. When PM cells are inserted they should not delay the user cell stream. To verify this the transfer time should be the same for a stream with and without PM cell generation enabled.

The tester needs to verify that the user cell counts reported by the forward monitoring OAM PM cell correspond to the actual user-cell stream content. In particular, the tester should verify that the correct cells are considered for PM calculation. Specifically, F5 flow is incorporated into the F4 stream when calculating the F4 OAM PM.



Broadband Series Test System

The criteria for a successful analysis are to verify that the SUT generates forward monitoring cells with:

- An average user cell block size that agrees with the requested block size
- The PM cell contains an accurate count of CPL0 and CLP0+1 user cells and that the BIP-16 is calculated in accordance with the standard.
- The F5 cells, when activated, are incorporated into the F4 stream calculation.

The OAM test system is connected to the output of a SUT where it analyzes the data stream and captures OAM PM cells.

If the test system's calculation does not correspond to the actual OAM PM cell contents the forward monitoring OAM PM generation is not working correctly. The test system should be directly connected to the output of the SUT to ensure that it detects coding errors, not transmission errors.

A key issue is the handling of the correct cells. Only user cells should be considered for the OAM calculation. The tables shown in Figures 15 and 16 illustrate the definition for user cells on the F5 and F4 levels.

To validate the correct selection of user cells, a data stream with different cell content must be generated. This requires a flexible transmitter in the analyzer. If there is a mismatch in the results provided by the OAM PM cell, and the actual user-cell stream, the transmission of each type may be used to evaluate the source of error.

An important variation of this test is the analysis of an F4 stream that contains an F5 stream. The test goal here is to verify that the F5 OAM cells are treated as F4 user cells.

PTI	Interpretation	Category
0,1	User Data cell, congestion not experienced	User
2,3	User Data cell, congestion experienced	User
4	Segment OAM F5 flow cell	Non-User
5	End-to-end OAM F5 flow cell	Non-User
6	Resource management cell	Non-User
7	Reserved for future use	Non-User

Figure 15: Definition for user cells on the F5 level

VCI	Interpretation	Category
0	Unassigned (VPI = 0)	Non-user
0	Unused (VPI > 0)	Non-user
1	Meta-signaling cell (UNI)	User
2	General broadcast signaling cell (UNI)	User
3	Segment OAM F4 Flow	Non-user
4	End-to-end OAM F4 Flow	Non-user
5	Point-to-point signaling cell	User
6	Resource management cell	Non-user
7-15	Reserved for future use	Non-user
16-31	Reserved for future use	User
> 31	Available for user data transmission	User

Figure 16: Definition for user cells on the F4 level

For this test, at least two OAM PM streams need to be generated. Discrepancies between the user-cell stream contents and the forward monitoring OAM PM cell data for the F4 and F5 flows need to be checked. If this error is suspected, it can be verified by toggling the F5 stream and looking at the F4 PM calculation produced by the management system.




Figure 17: Delta T1

PM Cells Processing

In-service performance measurements are accomplished by injecting forward monitoring OAM PM cells into transmitted user traffic. Transmission errors are detected at the segment or connection end points by comparing the received user cells with the error detection information contained within received forward monitoring OAM PM cells. Detected errors are usually reported directly to the external network management system by the PM end point. However, if the monitored segment or connection crosses a private network boundary, the management system cannot directly access measurement results. Instead, performance results must be returned to the source point within backward reporting OAM PM cells.

Testing should also cover erred links as an input where actual stream content does not comply with forward monitoring OAM PM cells. The generation source needs to be able to insert errors such as number of CLP0 cells lost, number of CLP0+1 cells lost, number of cells mis-inserted, BIP errors, and loss of OAM PM cells. The test system needs to have control on the number of blocks to be impaired.


Control over the number and type of errors allows comparison between what the analyzer generated, and what the management system reports.

For example the E6270A has a flexible and powerful OAM error generator and it can also test outside of the protocol to ensure that the SUT implementation is robust.

Test Scenario: Verify Backward Reporting OAM PM Cell Generation

The goal here is to check that backward reporting OAM PM cells are generated accurately by a SUT whenever receiving a forward monitoring OAM PM cell. To accomplish this, the test system needs to be configured to emulate a segment or connection source. It needs to generate a stream of user cells and forward monitoring OAM PM cells with simulated connection errors. The test system then needs to:

- Analyze the backward reporting OAM PM cells generated by the SUT
- Check that the NMS accurately report connection errors
- Verify that backward reporting OAM PM cells are transmitted no more than one second after the received forward monitoring OAM PM cell (delta T1 in Figure 17)



Broadband Series Test System

Network Management Statistics

Two cases have to be considered when testing the accurate reporting of performance metrics through the network management system (NMS). In the first case test, the equipment needs to be set up to emulate a segment or connection. A known traffic pattern should be enabled and then the result from the NMS should be compared with the generated traffic pattern.

To analyze backward reporting OAM PM cells, the test equipment needs to generate a stream of backward reporting OAM PM cells and the result provided by the network management system compared with the analyzer's configuration to verify:

- Blocks with lost, mis-inserted and erred cells
- The total number of lost, mis-inserted and erred cells
- The number of backward reporting PM OAM cell sequence number violations

Performance Management Under Load

In a real world environment, it is expected that multiple OAM PM flows are active and correctly handled at the same time. Calculation of OAM PM cells is computation-intensive. Depending on the implementation architecture, only a few channels might be supported. From a user point of view, it is critical to know how many flows may be activated, without jeopardizing the accuracy of the performance measurements.

To verify OAM PM, a system under test needs to be loaded with a number of traffic streams with OAM flows. Then the performance reports which are generated by that system's management console need to be checked for accuracy. If this test is successful, the system itself can be used to validate that the generation process works for multiple active virtual connections.

Verification of Performance Management Under Load

Load testing is used to determine if the management system reports accurate results under a simulated live load.

To verify PM monitoring under load, the analyzer must be configured to generate multiple traffic streams with OAM PM on many VCCs and VPCs.

The number of connections is gradually increased until the management system reports a discrepancy. Since a known-good cell stream is being fed directly into the segment or connection end point, there are no transmission errors. Therefore, errors in the PM analysis can be attributed to performance limitations of the OAM-handling capability of the system under test. The result is a number of channels that can be accurately reported by the SUT's management system.

A second load test can be performed to ensure that a management system correctly detects and reports errors under load conditions. To do this, the analyzer is configured to inject a user-defined number of errors into a stream. The network management system should then report this value.

Activation-Deactivation of Functions

In order to use continuity check or OAM PM functionality, the far end has to be activated to support the requested functionality. A network management system might provide a proprietary activation process for these functions. In the real world, however, problems might cross network providers' boundaries. Therefore, the source for a problem might not lie within the realm of one network management system's (NMSs) capabilities. To make use of the discussed OAM functions, independent of proprietary network management systems, a standardized process for cell activation -deactivation has been defined.




Figure 18: A continuity check function

Both the continuity check and performance management OAM functions require coordination between segments or connection end points. For the continuity check function, a source process (labelled A1 in Figure 18) must insert continuity check cells into the transmitted user traffic. The sink process (labelled B1 in Figure 18) must monitor for gaps in its received traffic.

Similarly for the performance management function, a source process must insert forward monitoring PM cells into the user traffic it transmits, and the sink process must detect discrepancies between the received traffic profile and the source traffic profile described by the forward monitoring OAM PM cells.

The OAM activation-deactivation procedure enables and disables OAM source and sink processes. OAM functions are enabled by transmitting an activation request to the source or sink node.

If the target node has sufficient resources, an activation-confirmed cell is returned to the sender. However, if the target node is unable to satisfy the request, it responds with an activation request denied cell.

Similar negotiations are used during deactivation of an enabled OAM function. A deactivation request cell is sent to the target node which in turn responds with either a deactivation confirmed or deactivation request denied cell. Figure 18 shows one possible network configuration during an activation-deactivation procedure.

Activation and Deactivation Test

To test the cell activation-deactivation process, the analyzer can replace either side of the process. An activation-deactivation emulation is required in the test system to support this configuration.

In the case where the analyzer initialized the activation-deactivation process, the analyzer emulates a source (or sink) by:

- Generating OAM activate and deactivate request cells,
- monitoring received activation-deactivation response cells, and
- monitoring received continuity check or performance management cells.

The test is successful if the requested activation-deactivation is performed. In the case where system resources are not available an request-denied should be returned. This is a success since the SUT has responded correctly.



Broadband Series Test System

On the receiver side of an activation process the test system should respond to activation- deactivation requests. It should be able to:

- Generate request confirmed or denied response OAM cells,
- simulate response delays, and
- activate or deactivate the appropriate continuity check or performance management source process.

This allows verification of received activation and deactivation requests, and validation that the response OAM cells are processed correctly.

In-service Testing

OAM testing may also be used during network deployment.

An ATM test system which supports loopback emulation can be used to verify the connectivity of an ATM network. For example, it may be used to verify that a customer's CPE equipment has been connected to the network.

In an operating network environment, timing issues are critical. Measuring response delays for loopback and activation-deactivation requests can be used to detect inter-working and service problems.

By using the timestamp in an OAM PM cell (the TSTP field shown in Figure 9), a test system can determine network delay between any two points. This delay measurement allows network providers to characterize their networks and to estimate the performance of those networks.

Summary

OAM is an integral part of ATM. The procedures defined by the OAM protocol enable in-service tests of network health and operation which are of critical concern to both public and private network providers. Integrating the OAM functions, including fault and performance management, into network devices such as switches, can help reduce the downtime, and thus optimize the performance, of these mission critical networks. And when OAM is implemented, other higher layer network management applications can be more effective.

But to optimize overall network management system performance, correct operation of the OAM functions must be verified. By its nature, the OAM protocol is a complex, real-time, stimulus-response system. Thus, any testing of the OAM capabilities of a system, such as a switch, also becomes a complex, yet necessary, task. To verify these critical functions requires a comprehensive OAM analyzer with the following capabilities:

- Line interfaces for all defined transport specifications, including SDH and PDH. The analyzer requires line interfaces which allow access into the physical layer OAM F1 to F3 flows.
- For OAM fault indication, AIS/RDI cells need to be generated periodically, AIS/RDI conditions need to be analyzed and reported, and AIS states need to be emulated.
- For OAM continuity checking, the analyzer needs to be able to periodically generate CC cells with a customized interval, monitor CC cell streams, plus detect and report LOC states.

- For OAM loopback functions, the analyzer needs to be able to generate and monitor request and response cells, simulate network delays and report loopback delays.
- For OAM performance management function testing, the analyzer needs to be able to generate 120 user cell streams* with OAM PM cells inserted, simulate user-cell errors, losses and mis-insertions into the transmitted cells, emulate the backward reporting OAM PM response, and analyze user cell blocks and PM cells on at least two streams in real-time.
- For OAM activation- deactivation testing, the analyzer needs to emulate both sides of the activation-deactivation process, simulate network delay, and report request outcomes and response delays.
- An OAM analyzer should be easy and intuitive to use. It must contain both an intuitive user interface and an automated test environment that allows regression tests to be configured.

Conclusion

This paper has presented some insights into the ATM OAM capabilities and applications. Some selected examples were presented to illustrate how to test OAM and what measurements are required to characterize OAM functionality. However, there is much more to OAM implementation and analysis that was not covered by this text. Some of these areas are still topics of ongoing research, including:

- Verifying the correct functionality of OAM with traffic policing
- OAM information translation in an interworking environment (such as Frame Relay)
- OAM and ABR

Consult the proceedings of Study group 13 of the ITU-T for further information on the above subjects.

Agilent is a leader in providing test applications for emerging networking technologies. Agilent offers a wide range of ATM analyzers, including the Broadband Series Test System (BSTS). For further information on ATM OAM analysis products offered by Agilent Technologies, please contact your local sales representative.

Acronyms

AIS	Alarm Indication Signal
CC	Continuity Check
CLP	Cell Loss Priority
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
FM	Fault Management
LOC	Loss of Continuity
NMS	Network Management System
OAM	Operations and Maintenance
PDH	Plesiochronous Digital Hierarchy
PM	Performance Management
QoS	Quality of Service
RDI	Remote Defect Indication
SDH	Synchronous Digital Hierarchy
SECB	Severely Errored Cell Blocks
SUT	System Under Test
VC	Virtual Circuit
VCC	Virtual Circuit Connection
VP	Virtual Path
VPC	Virtual Path Connection



Broadband Series Test System



Agilent Technologies Broadband Series Test System

The Agilent Technologies BSTS is the industry-standard ATM/ISDN 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.

5965-9332E 05/00 Rev B



Agilent Technologies

Innovating the HP Way