Network OAM Requirements for the New York City Transit Network

Zhi-Wei Lin, New York City Transit

ABSTRACT

New York City Transit's subway system is one of the largest and most complex mass transportation systems in the world. Because of the extensive subway system, an extensive communications network infrastructure is needed to allow for communications of the various subway services with the control centers to support mission-critical (and safety-critical) applications. The current network utilizes multilayer network technologies, including SONET, ATM, and IP (layer 2/3 device) layer networks. As a consequence of using these various technologies, extensive challenges are faced in trying to consolidate these separate networks into a single management view, simplifying (and automating) not only the provisioning aspects but also troubleshooting/fault management aspects of the network. Such automation will help to simplify network operations and allow for better handling and management of the reliability and availability of the network. To achieve this goal, various technologies were evaluated. Consideration is being given to multiprotocol label switching technology. MPLS offers the potential to help converge to a simpler networking model. In order to achieve this, certain capabilities must be available (e.g., the ability to quickly and automatically identify defects/failures and subsequently reroute around these failures). This article discusses the operational requirements for the MPLS network from the point of view of backbone networks that support a mass transportation system operator.

INTRODUCTION

New York City Transit's subway system is one of the largest and most complex mass transportation systems in the world, operating bus and subway services throughout the five boroughs of New York City 24 hours a day, 365 days a year. The subway system consists of approximately 722 track miles (240 route miles), and provides service to 468 stations in four counties (boroughs). Its 25 subway lines are interconnected, with free transfers between lines permitted at more than 50 locations. Each day, more than seven million people use New York City Transit; close to three billion customers annually.

Because of the extensiveness — large number of stations and continuous operation — of the subway system, an extensive communications network is required that allows for communications of the various subway services with our central control centers. Examples of types of mission- and safety-critical applications carried over this network include the automatic train supervision (ATS) system, station information management system (including public address, information signage, and CCTV/video systems), communications-based train control (CBTC) system, as well as voice/telephony systems. This network must provide the absolute highest reliability and support minimal manual/craftsperson intervention in maintaining and operating this network. In addition, because network equipment is deployed on both elevated (exposed) stations as well as underground (subway) stations, it must withstand harsh environments not expected in many central offices (e.g., ambient temperatures as high as 130°F, severe vibrations especially on elevated stations, airborne contaminants such as steel dust particles coating the equipment). Environmental requirements are not covered in this article.

At the time of network design and deployment, no single technology addressed the various communications needs. To support these requirements, a multilayered network was designed made up of synchronous optical network (SONET), asynchronous transfer mode (ATM), and IP networking technologies. Each of these technologies is used to support a specific need. For example, the SONET network was deployed to support very high reliability and fast recovery speed. This is accomplished via the SONET operations, administration, and maintenance (OAM) mechanisms (e.g., AIS and RDI signals carried within the SONET overhead) as well as fast protection mechanisms (e.g., BLSR and UPSR protection switching mechanisms). The ATM network was deployed as the ubiquitous access for most applications, providing the means for service aggregation as well as the ability to partition resources for different service level requirements. ATM OAM mechanisms are employed to support fast detection and notification of faults at the ATM layer (e.g., using continuity check cells). In addition fast restoration mechanisms (e.g., SPVCs, SVCs, and redundant PVCs) are used to support fast

0163-6804/04/$20.00 © 2004 IEEE

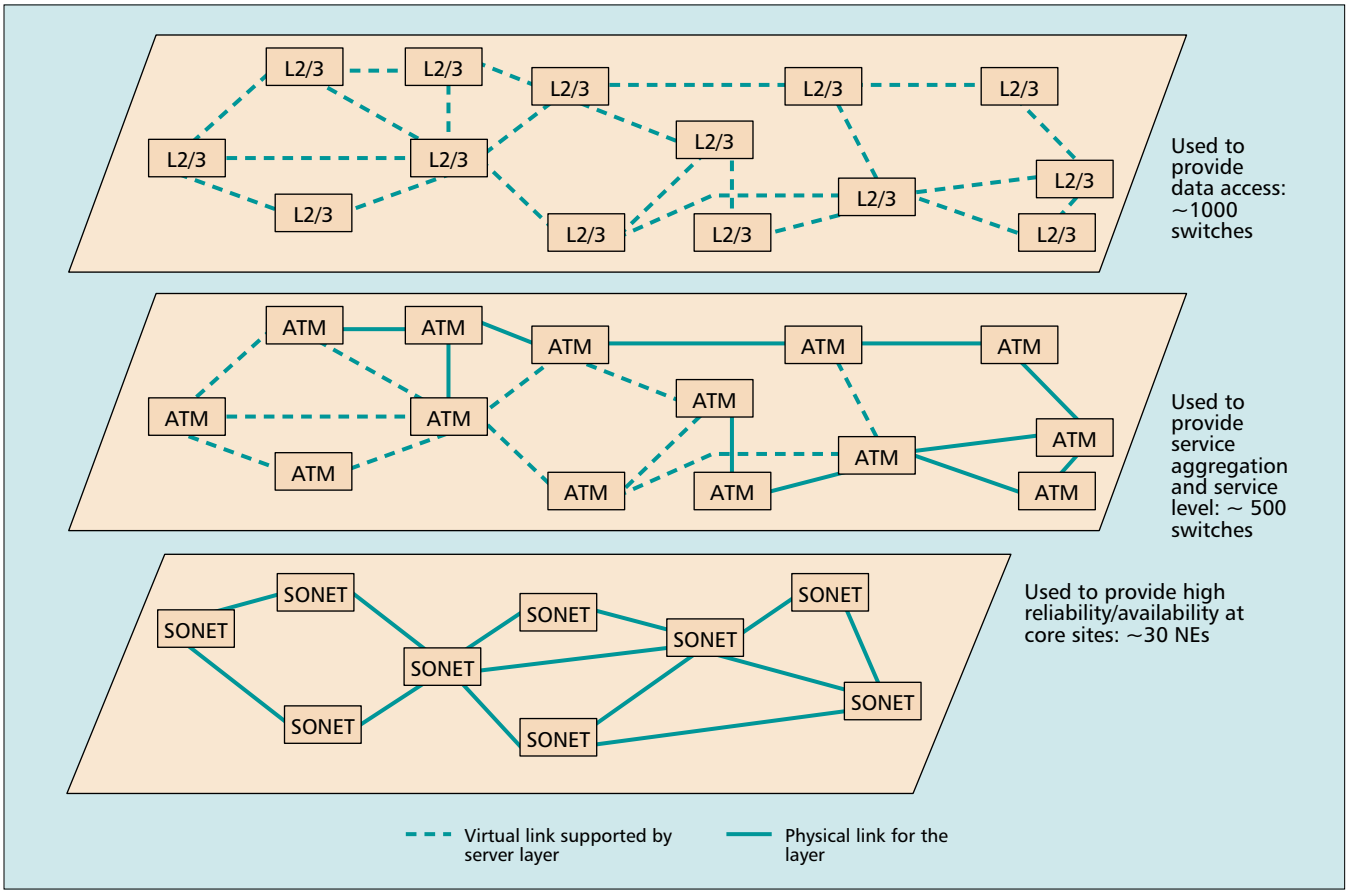

Figure 1. *A logical illustration of the MTA New York City Transit communications network.*

recovery. The IP network was deployed to handle the transport of data services. Within NYCT's IP layer network, basic IP mechanisms are used for detecting IP layer faults (e.g., Hello timer expiration). Figure 1 provides a logical illustration of the structure of the NYCT communications network system.

The network needs to support the communications needs of the 500-plus locations within the NYC Transit system. These include the 468 subway stations, plus various bus depots, train yards, office buildings, and so on that require access into the network. The design of the physical plant was constrained by the right of way of our subway systems, giving a fiber cabling infrastructure that mirrors the subway lines (see the MTA NYCT Web site [1] for a view of the transit system). The SONET NEs were deployed as the infrastructure for the backbone, while the ATM and IP switches/routers were deployed at all stations to support interfacing for various types of applications.

With a design of this type, management of the multilayer network becomes a challenge, especially considering that the various layer networks are all interrelated by virtue of their shared links; that is, a "physical" link of a server (e.g., SONET) layer provides the virtual link for its client (e.g., ATM) layer. A single cable defect (e.g., cable breakage) can affect multiple network layers, and in turn will cause different layers to respond in different ways to a failure. With our current network configuration, the SONET layer may take 50 ms to recover from the fiber break; the ATM layer may take anywhere from hundreds of milliseconds to several seconds (or longer), depending on the type and number of circuits (PVC or SPVC); the IP layer may take several seconds to much longer depending on the size, connectivity, and routing convergence of the network. It should be noted that current (and future) generation IP routers do exhibit much faster recovery from failures. This can be attributed to faster router processors as well as improved survivability mechanisms. Introduction of multiprotocol label switching (MPLS) (as discussed below) and its set of functionalities should also help to significantly decrease the recovery time via mechanisms such as fast reroute.

Because NYC Transit operates the full span of the layered network, from the physical infrastructure up to the applications, a consolidated means to provide automation of network management is needed, in terms of a single view of the entire network, end-to-end provisioning across the different networking technologies, fault detection, localization/isolation, notification, and correlation across the different network layers, service level description, and so on within a single network management system. The current network does not provide the desired level of integration; separate management systems handle the SONET, ATM, and IP networks. As such, a failure within the network

Consolidation of the existing network layers into a single layer network (in this case, the MPLS network) will be considered if the network provides all the required capabilities and meets the requirements necessary to support the various critical applications.

may generate multiple alarms that need to be correlated manually. The logical plan is to consolidate the functions currently handled by these multiple layer networks into a single-layer network that exhibits the capabilities critical to meeting the required availability levels. The current layer network that has the potential to suit this role is the MPLS layer network. MPLS's potential includes not only its flexibility (e.g., label switched path, LSP, hierarchy allows for arbitrary levels of network segmentation) but also its ability to integrate more closely with IP layer networks (where we expect much of the future traffic to be carried). However, to ensure that MPLS fulfills its potential, basic functionalities must be provided, chief among them OAM capabilities to support highly reliable operations.

Requirements for the New Network

Consolidation of the existing network layers into a single-layer network (in this case, the MPLS network) will be considered if the network provides all the required capabilities and meets the requirements necessary to support various critical applications. The discussion that follows attempts to lay out the NYCT requirements for OAM capabilities of the network.

(The term OAM as used in this article includes any data plane mechanism that supports the exchange of network status information. This exchange of data plane status information may be used for different purposes, one of which is initiation of an automated network recovery mechanism.)

The purpose of this status information exchange may range from notification of the management system for maintenance purposes to notification of a reroute mechanism for initiating connection reroute. Note the distinction made between status information exchange (by means of detection and notification) and traffic reroutes (as a consequence of received status information).

Fault Detection and Notification

The first and foremost requirement is the ability of the MPLS network to support automation for fast fault (defect) detection and notification. Automatic detection of faults should extend to include not only physical layer faults (e.g., node failures and cable cuts) but also LSP failures.[1] Notification of the fault condition should follow as soon as detection has occurred to allow fast network recovery.

The importance of this requirement stems from the particular application requirements (e.g., CBTC and ATS) that stipulate the need to transport this traffic quickly and reliably (e.g., any delays or data loss impact the monitoring and control of trains, which impacts the scheduling of services). Automatic fault detection refers to the ability of the network to automatically and quickly (in near real time) detect whether a defect or failure has occurred within the network (this does not necessarily require knowledge of where the failure has occurred; that task is delegated to fault localization, discussed below). Important defects that need to be handled include such items as:

- Cable cuts
- Interface card failures
- Misrouting or misconnections (e.g., as a result of transit node misconfiguration or switching matrix defect); can occur for either statically or dynamically set up connections
- Connection failures
- Node failures

Several detection methods may be employed; those based on reactive detection (e.g., noticing disruption of traffic and then running mechanisms to verify failure) and those based on proactive detection (e.g., constant monitoring of the network). In order to satisfy these criteria, a proactive detection mechanism is required. It is not appropriate to detect a network failure based on detection of application traffic disruption; the network must detect its own failure conditions. One of the OAM functions that must be maintained by the MPLS network (e.g., those that currently exist for SONET and ATM networks) will be the ability to support proactive fault detection and notification.

Network protection/restoration mechanisms should be initiated (immediately) upon fault notification in order to recover and reroute traffic away from the fault condition (requirements for network protection and restoration mechanisms are not discussed in this article). Independent of the survivability mechanism, network fault localization and isolation may also be initiated.

Fault Localization and Isolation

Fault localization (isolation) is another capability required. In the transit environment, because the cabling infrastructure runs along subway tunnels and communications rooms are located within station complexes, any mechanism that requires access to the infrastructure to localize faults will result in train reroutes that create scheduling delays and service degradations. Current networks provide manual methods for fault localization, typically by use of loopback mechanisms run iteratively to localize the fault to a particular section of the infrastructure. These mechanisms may be run remotely by operating personnel to locate the fault. As such, the MPLS network must support fault localization mechanisms.

A desirable capability beyond having a manual localization mechanism is to automate the localization process so that when a fault has been detected, the network automatically (initiated by either the automated OAM mechanism or an element management system) runs the localization process to determine the location of the fault, preferably immediately after the detection of the fault. Note that the localization process should be independent of the detection process, and must not impede the operation of fault detection. This would help free up operating personnel to perform other critical functions related to communications in support of mass transit service (train monitoring and control, coordinating incident responses, rerouting of trains due to incidents, etc.).

[1] In fact, in a multilayer network where the MPLS "physical" link may be supported by a server layer network, a physical layer failure at the server layer may not be detectable by the MPLS layer at all; the fault condition may only be translated to the MPLS layer as simultaneous failures of all the LSPs. This limitation is obviously dependent on the network configurations and implementation details for the MPLS equipment.

The NYC Transit communications network supports a variety of traffic profiles, as described above. These include applications that communicate via a point-to-point model (e.g., voice, ATS, CBTC) as well as via a point-to-multipoint model (e.g., video, public address announcements). As such, to support different traffic mixes, the fault detection, notification, and localization mechanisms must support the ability to detect, notify, and localize not only point-to-point faults but also point-to-multipoint faults.

The next criterion for the MPLS network relates to the traffic generated by the OAM mechanism (e.g., messages to support proactive monitoring, messages to support fault notification), especially as it relates to failure scenarios. As previously described, the current network deploys ATM switches to every station within the NYC Transit network. If MPLS were chosen as the technology for the future (assuming it meets all requirements), MPLS switches would also be deployed across the transit network (approximately 500 locations). Due to limited path diversity, many of the switches will be connected across common (aggregated) links, likely with multiple LSPs per location in support of the different traffic mixes. As such, certain failure scenarios may impact multiple MPLS nodal connections. As previously mentioned, a physical layer fault condition may not be detectable by the MPLS node (e.g., a cable cut not directly connected to the MPLS node, but to a server layer node such as a SONET network element, NE). In such a scenario the fault condition may trigger the failure of all LSPs traversing the MPLS node. This can result in detection and notification of a fault condition for each LSP, creating an alarm storm.

Alarm storm propagation can potentially impact the recovery speed of the network due to the need for MPLS switches to process these alarms, and may interfere with critical application traffic in terms of either disrupting application traffic flow or introducing delays into the traffic path. The OAM mechanism must therefore support the ability to prevent (or reduce) alarm storms from propagating throughout the network that subsequently interferes with critical applications traffic. Not only does the OAM mechanism need to support a reduction (or elimination) of alarm storms, but it must also support the ability to quickly notify downstream (and upstream) nodes of failure conditions (a failure may impact multiple paths, whose sources must all be notified of this failure). This is critical in allowing the network to quickly respond to a failure condition (e.g., enabling an alternate path or initiating rerouting of a path). Discussion of recovery requirements as applied to the NYC Transit system is not covered in depth in this article.

Coexistence and Interoperability of OAM Mechanisms

MPLS is likely to be considered as a choice for the future network to support consolidation of the different layer networks; however, as with any network deployment, there will be a point where the existing and new networks will coexist and provide transport for critical application traffic. Migration of application traffic from the existing network to the new one will take some time to complete. As such, during this migration period there will be scenarios where particular application traffic may be transported across both networks (e.g., sourced in the existing network and sinked in the new MPLS network). Such transport is likely to easily be supported; however, the issue will be how path connectivity/continuity will be handled in terms of:

- Fault detection (e.g., detection mechanism spanning both networks, or partitioned into each network and the interconnection between the network)
- Fault localization (e.g., localization mechanism spanning the network, or partitioned into each network and the interconnection between the network)
- Fault notification (e.g., notification mechanism spanning the network, or partitioned into each network)

Solutions to support end-to-end detection, localization, and notification spanning both existing and new MPLS networks would be desirable. This would allow, from an operational perspective, a single mechanism running end-to-end in support of the common requirement. Partitioning the mechanism into segregated domains because of OAM incompatibility would require that an upper layer management system (e.g., network management system, NMS) handle the integration and correlation of various OAM traffic from different subnetworks to produce an end-to-end view of any failure condition. While this approach may be taken, it has been our experience that network management system features lag behind built-in OAM features, sometimes by significant time periods. As such, relying on NMS capability to support an end-to-end view would be undesirable from NYCT's perspective.[2]

Whether one method is favored over the other has not been decided; however, the goal of future deployment is to simplify the operational/management aspects as much as possible. Trade-offs in terms of handling end-to-end monitoring using an OAM mechanism vs. handling end-to-end monitoring using an EMS/NMS system need to be analyzed with respect to the impact on the amount of provisioning, simplification of operations, performance issues, as well as the scalability and flexibility of each method. Other nontechnical requirements will also come into play in deciding the methodology, such as the personnel training needed to support the various methods.

Trade-offs in terms of handling end-to-end monitoring using an OAM mechanism versus handling end-to-end monitoring using an EMS/NMS system needs to be analyzed with respect to the impact on the amount of provisioning, simplification of operations, performance issues, as well as the scalability and flexibility of each method.

MPLS as a Potential Networking Solution

MPLS has received much attention. Many experts view this as the layer network with the best potential to provide convergence of traditional time-division multiplexing (TDM)-based and packet-based systems. As a result of this view, much standardization activity is occurring to complete the MPLS capability set to support the necessary features that will make it a "carri-

[2] *Traditional service providers may have different views, since they may have better/tighter control of the development of NMS capabilities (largely developed in house).*

NYC Transit looks forward to the fast-paced and quality developments of the MPLS technology as a medium for network convergence, and will continue to monitor the standardization efforts within the industry.

er-grade" solution. One such effort is in the area of adding fault detection and localization mechanisms. Another effort is in the area of LSP reroutes. Please note that these areas are being treated as distinct in this discussion. Reroute mechanisms may be activated to allow for fast switchover from a disrupted LSP to a backup LSP; however, this switchover does not occur until a failure has been detected. As such, it is not adequate to specify one and not the other. Both fast detection and fast protection/restoration are needed to support fast *network* recovery.

There are two solutions currently positioned. One proposal is based on the tried and true IP traceroute mechanism (Internet Engineering Task Force, IETF, Request for Comments, RFC, 3609), which provides (as the name implies) requirements on tracing of a route or tunnel to verify proper operation (a companion document describes the protocol specifications for a ping and traceroute mechanisms). The second proposal is based on the tried and true ATM OAM mechanism (International Telecommunications Union — Telecommunications Standardization Sector, ITU-T, Y.1711), which provides connectivity verification (called *continuity check* under ATM) as well as fault notification.

For the purpose of New York City Transit network requirements, both methods may be used in complementary fashion within the network. A Y.1711-based mechanism provides automation for fast detection and notification of a fault condition, while the RFC 3609-based mechanism provides the requisite tool for subsequent diagnosis of the network to localize a fault (in fact, Y.1711 has provisions to complete a loopback mechanism, which should support a similar function). As such, the combination of these tools should provide a basis from which the critical fault monitoring function may be achieved. However, additional capability enhancements must be made to these tools in order to support their integration and automation to support automated fault detection, notification, and localization.

To qualify MPLS as the future networking technology within the New York City Transit communications network, not only does the above mechanism need to be enhanced and stabilized, but other non-OAM-based mechanisms must also be developed and enhanced. One such non-OAM requirement is the ability of the network to not only quickly detect a failure, but also quickly recover from one (others include but are not limited to support of quality of service, common support for both TDM-centric and data-centric traffic, operation in harsh environments, modular architecture). This will be critical in support of time-sensitive (and safety-critical) applications that require highly robust networks (e.g., the current SONET network) to support their needs. NYC Transit looks forward to fast-paced and quality developments in MPLS technology as a medium for network convergence, and will continue to monitor the standardization efforts within the industry.

Acknowledgments

The author would like to acknowledge helpful discussions and guidance from Kwame Asamoah, Anne O'Neil, and Morris Schwartz in preparing this article.

References

[1] MTA NYCT subway system: http://www.mta.info/nyct/maps/submap.htm

Additional Reading

[1] ATMF af-pnni-0055.000 (1996), "Private Network-Network Interface Specification Version 1.0 (PNNI 1.0)."
[2] IETF RFC 3031 (2001), "Multi-Protocol Label Switching Architecture."
[3] IETF RFC 3609 (2003), "Tracing Requirements for Generic Tunnels."
[4] ITU-T Rec. Y.1710 (2002), "Requirements for OAM Functionality for MPLS Networks."
[5] ITU-T Rec. Y.1711 (2002), "Operation and Maintenance Mechanism for MPLS Networks."
[6] ITU-T Rec. Y.1712 (2003), "OAM Functionality for ATM-MPLS Interworking."
[7] ITU-T Rec. I.326 (1995), "Functional Architecture of Transport Networks Based on ATM."
[8] Telcordia GR-253-CORE, Issue 3 (2000), "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria."
[9] IETF draft draft-ietf-mpls-lsp-ping-05.txt, "Detecting MPLS Data Plane Failures," Feb. 2004, work in progress.
[10] IETF draft draft-ietf-mpls-lsr-self-test-02.txt, "Label Switching Router Self-Test," Feb. 2004, work in progress.
[11] IETF draft draft-ietf-mpls-rsvp-lsp-fastreroute-05.txt, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels," May 2004, work in progress.
[12] IETF draft draft-ietf-ccamp-gmpls-recovery-terminology-04.txt, "Recovery (Protection and Restoration) Terminology for GMPLS," April 2004, work in progress.

Biography

ZHI-WEI LIN (zhiwlin@NYCT.com) is a manager with MTA New York City Transit, Department of Capital Programs Management, Communications Engineering. He has multidisciplinary expertise on network survivability strategies, network planning, architecture, and design. His networking knowledge encompasses a wide range of areas including ATM, SONET/SDH, OTN, GMPLS/ASON control plane specifications, and MPLS/Ethernet and the corresponding OAM mechanisms. From 1995 to 2002he worked at Telcordia as a senior consultant, where he was responsible for multitechnology network design and deployment, evolution planning, business plan development, and associated analysis and modeling. He joined Lucent Technologies in July 2000, providing support for standardization and systems engineering for Lucent's optical networking product line. Since joining NYCT in 2003, he has been involved with developing the requirements and network architecture for NYCT's communications infrastructure. He is also involved with various projects to modernize operation of the transit system, including overseeing deployment of the Automatic Train Supervision System. Over the past nine years he has provided over 50 technical contributions to various standards and industry fora, many of which have been used as input toward approved international Recommendations. He holds an M.S. degree in electrical engineering from Cornell University, Ithaca, New York.