Technical White Paper written for Panduit Corporation.

Managing the Physical Layer

1. Overview

Service Providers and Enterprises are facing business challenges on many fronts. Amid intense competition and a tough economy they are being forced to accept reduced revenue while at the same time are being called upon to provide higher network up-time and service levels. Emerging technologies that improve the management of the network are therefore gaining interest, because they enable a more cost-effective use of the network infrastructure.

Network management tools that monitor the health and performance of network elements have been available for many years. Yet all the sophisticated fault and performance management tools are useless when a human mistakenly disconnects users from network service. The industry is truly plagued with stories of simple network maintenance activities bringing down critical connections – for extended periods of time – due to an unmanaged and poorly documented physical layer. We now know that a significant percentage of network downtime (up to 80% according to a study by Sage Research) can be attributed to physical layer connections.

An emerging trend in the industry is physical layer management systems that provide real-time monitoring of physical layer connections. Physical layer elements monitor and map all connections in the cross-connect field. The monitoring hardware can be integrated with the Element Management and Network Management layers in the Telecommunications Managed Network (TMN), and any inadvertent disconnection is immediately detected. Physical layer monitoring solutions also have the advantage of being able to localize any fault caused by the physical layer, including those at remote sites. When a circuit is disconnected causing a disruption of service, technicians know where the fault is located, providing for rapid fault resolution.

This paper will examine emerging physical layer management technologies, their implementation in the service provider and enterprise markets, and will describe solutions currently available that have incorporated these new technologies. The paper will also show that by effectively implementing these solutions, there will be a reduction in network downtime, security of the network will be increased, mean time to repair (MTTR) will be enhanced, and Service Level Agreements (SLAs) can be developed with a higher uptime commitment.

2. The Physical Layer Management Value Proposition

In the past few years, there has been a realization that network management is not a complete picture without a window into the physical layer. Access to important data and linkage to the overall network is fundamentally based on the physical layer, and security is therefore at its most vulnerable at this layer. All things considered, the physical layer is essentially the true “edge” of the network, but unfortunately has the least amount of visibility of all of the critical components making up the network infrastructure. Although this is the case, there has been a very slow adoption of the technology needed for the visibility and management of the physical layer. Although on the increase, to date, it is estimated that less than 1% of the entire copper and fiber infrastructure is being managed in real time. There are three key factors that come into play regarding this slow adoption:

Today’s economic conditions drive the lion share of decision-making regarding the purchase of these products and technologies. What is sometimes a very clear economic decision (positive ROI in a year or less) or a way to avoid penalties associated with downtime clauses in SLA’s is often overshadowed by the need to have expenditures directly related to revenue generation. A physical layer management system that does not deliver revenue generation is put in the “luxury” category.

1. Physical layer management is stuck in the middle. It is often viewed as inconsequential by the MIS and IT organizations, deserving little attention. Their attention is focused on provisioning and maintaining the network at the higher layers of the OSI model, and leaves little time to worry about Layer 1. Conversely, the cable installers and the frame craft either view it as a threat or as a complicated solution to what should be (in their minds) a simple one. It therefore falls in between, both in the decision-making and adoption once the investment is made.

2. Network elements and deployment architectures offer many features that obviate the need to keep the physical layer connectivity up 100% of the time. Protection switching in SONET transmission gear, dual-homed ring architectures, diverse routing, and other fault tolerant measures provide answers for a great deal of scenarios for problems with the physical layer. There is some measure of closure here, but it is often unrecognized that these capabilities are focused on the outside plant/horizontal cabling side of the CO’s, data centers, and headends where the majority of the problems occur, and are reactive in nature. Service providers delivering end user services will be impacted at the patch panels and other connectivity points regardless of their failover strategies.

We will consider these factors in the next sections in order to properly understand the barriers to the adoption of physical layer management, and will provide some reasoning that the value propositions outweigh the barriers.

2.1 Justifying The Expenditure In This Economy: Understanding What Is Needed

In order to address the initial financial barriers, we first need to define where expenditures are going, what the magnitude of the expenditures is, and what practices and methodologies can be created or changed to address the problem areas. The expenditures center around themes that recur throughout the service provider and enterprise world:

Provisioning Service

- Handling Service Affecting Events
- Detecting and Dealing With Intrusion
- Auditing and Maintaining Cross Connects

In the next sections, we will examine each one of these areas in detail, and begin to put together the attributes of the “ideal” physical layer management platform to address the needs.

Provisioning Service

Rapid, high quality, and cost effective service provisioning is essential to the financial health and viability of any organization.

For Service Providers, it is essential that their customers are satisfied with each of these areas. For example, traditional POTS customers will not tolerate delays in the activation of a DS0 or for DSL service. There is plenty of competition these days in the form of competitive access providers for POTS, and the cable companies provide a viable alternative to DSL through DOCSIS based services for high-speed Internet access. Customers have choices, and will vote with their dollars (and they often do). This competitive environment puts the provider in a position where pricing comes under pressure, and accordingly the cost components of the service must be minimized. In addition, the customer will not tolerate services that are not at a high quality level, from either
an uptime perspective or a quality of service perspective.

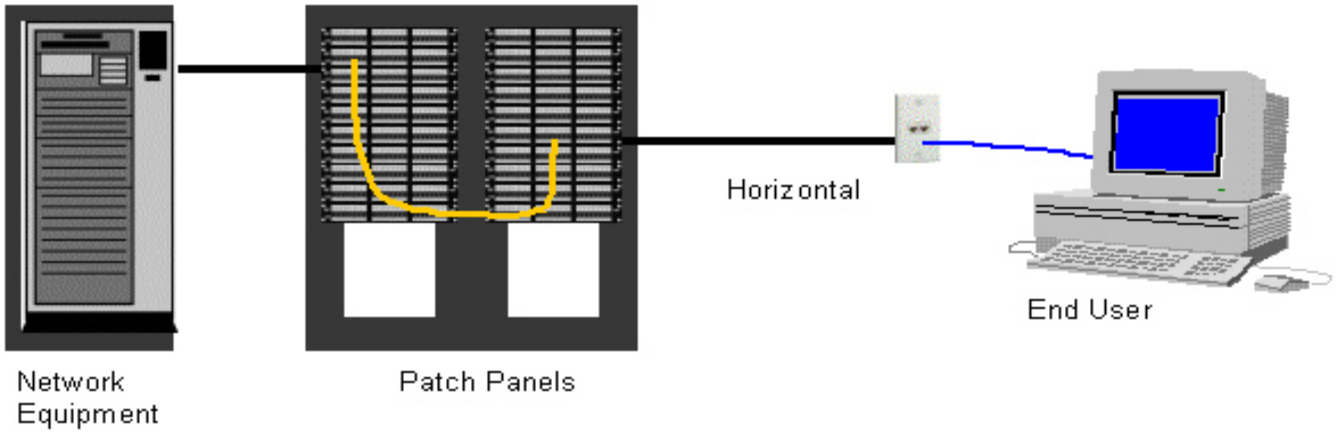

Figure 1. Service Path Via Patch Panels

For the Enterprise, the data center must be able to provide and maintain a myriad of services to the different organizations within the Enterprise in much the same way as the service provider. The tie-in to revenue generation is not as clear as the Service Providers, but the measurement here is the ability for the organization to perform moves, adds, and changes (MACs) in a rapid, high quality manner in order to keep the availability of the product or service they offer on track. In both of these entities, the provisioning of service includes the need to cost-effectively move connections between the equipment that will provide the service and the media delivering the service to customers. This is shown in Figure 1.

The provisioning process will often include the following steps:

Receive service activation, a service removal request, or a service movement request. This can come in electronic, paper, or verbal form.

1. Identify the location of the customer (OSP cable or horizontal cable in a building). This can be provided by either paper records or from a database, but in some cases may be verbal or can be done only after a visual inspection process to determine access points.
2. Identify an available connection to the equipment that will provide the service (or will disconnect the service in the case of a removal request). Again, provided by either paper records or from a database, but in some cases may be verbal or can be done only after a visual inspection process to determine access points.
3. Add (or remove) a copper patch cord or fiber jumper to activate (or remove) the service.

Each one of the steps described above abounds in the potential for errors based on the human intervention required in each step.

The initial request itself can be in error due to records of the physical layer. They may be out of date or inaccurate due to a transcription error, making the high level request to provide service impossible or impractical due to the cost of purchasing new equipment or pulling new cables. The access point to the OSP cable or the horizontal cable fiber or copper data pairs may not be known at the patch panel, as well as the access point to the equipment on the panels associated with the service needed. Assuming both of these are known, there still exists the need for someone to make the actual connection (or removal) at the patch panels. This can be prone to error due the high probability of someone inserting a patch cord or jumper in the wrong port on the panel, or worse, removing the wrong patch cord or jumper. All of these considerations add to the time and cost associated with the provisioning process, and drive profitability of the entity in the wrong direction.

In order to perform these steps in the most expedient, cost effective, and highest quality manner, a platform or methodology would need to exist that provides the following attributes:

The ability to handle the service requests in an efficient manner, through an interface to a Network Management Layer system tasked with end-to-end provisioning.

• The ability to deliver up-to-date (ideally real-time) information about the entire link between the equipment providing service and the equipment and/or user receiving the service to the Network Management Layer (NML) system or the person responsible for the provisioning request. This should be able to be done locally or remotely.
• A way to clearly identify the access points on the patch panels where the operations need to be performed, on a macro level (e.g., either a single operation such as adding or removing service or a multiple operation such as moving a service). This should be as foolproof as possible, providing additional direction to the person implementing the operation in cases where incorrect operations are performed.
• The ability to provide these capabilities in such a way as to impact the cost of the equipment and software required to drive the operations in a rapid and high quality manner while reducing the skill level required to a minimal level.

Handling Service Affecting Events

One of the most beneficial aspects of a physical layer management system is to provide a window to Layer 1 operation and events. What is usually a "dark secret" at the Layer 1 level is now elevated and visible to Layer 2 and 3 based NML systems. An "event" at the physical layer would be an unintended disconnection or connection of service associated with a service request, or an unauthorized connection or disconnection of service.

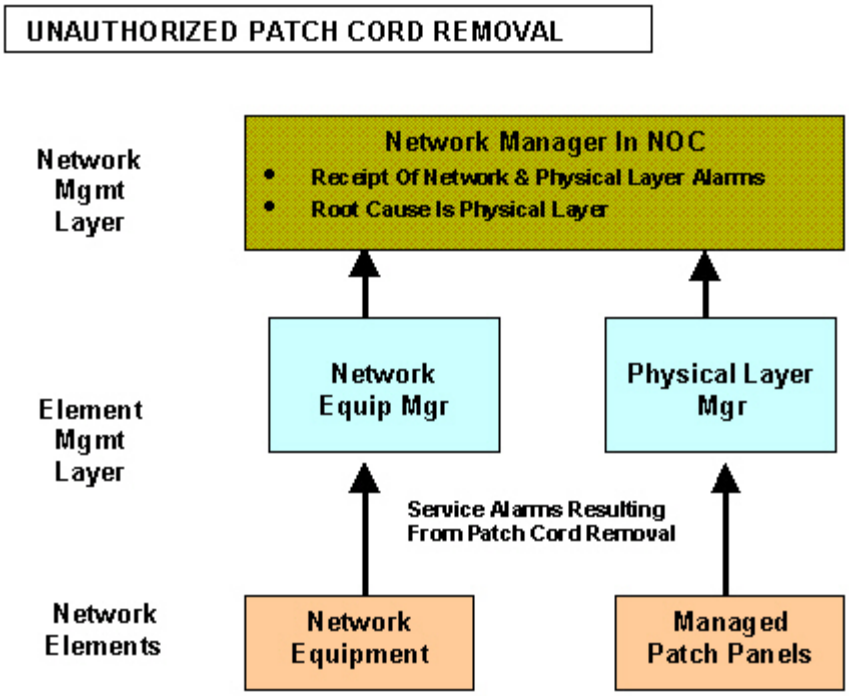

Figure 2 – Service Affecting Events

For example, a person involved in a requested disconnection of service mistakenly removes one end of a patch cord at an incorrect port. Assuming that this fiber was carrying live traffic, it is critical to send as much information about this event as quickly as possible to the entity responsible for the remediation of these types of events (typically the Network Operations Center, or NOC). In this way, the NOC would know where the event occurred, and the impact of the event can be minimized quickly. This, however, is the best case since a person is probably still at the point of failure (indeed, the person caused it in the first place), and with the provisioning aids that would be available (per the discussion in the previous **Provisioning Service** section) it should be resolved quickly.

A more interesting situation is when there is no pending operation at the patch panel, and there is unintended and unknown damage or a break to the patch cord caused by a person working on another activity in the CO or data center. In this situation, if there were no visibility of the event from Layer 1 at the NOC, the technician would be left to deal solely with all of the Layer 3 events that all of the equipment involved in the end-to-end link generated. The NOC technician would need to go through a series of diagnostic operations to sectionalize the problem as much as possible, but would have no insight as to whether the events were generated by a problem with the equipment or the physical layer at the CO or the building where events are coming from. With the addition of the physical layer management system's ability to "see" the physical layer based problem, the NOC would get both Layer 1 and Layer 3 events. Since both events would be present, the problem must therefore be confined to the physical layer, and the root cause would be discovered quickly. This is represented on Figure 2.

Conversely, a Layer 3 generated event can be isolated to the equipment since no Layer 1 events are present. This is a huge benefit to the NOC manager, and dramatically reduces the downtime and lost revenues associated with these problems. Depending on the service being provided, this can be significant. This can be seen on Figure 3 below.

Aside from lost revenues, there may also be penalties associated with Service Level Agreements with customers that would need to be considered. This is especially true in the case where the system would be used for disaster recovery.

As in the previous **Provisioning Service** section, looking at the characteristics of a platform or methodology for handling service affecting events, the following would need to be supported to provide an effective solution:

The ability to detect patch cord insertions, removals, and damage and/or breakage.

- The ability to communicate these events to an Element Manager or a Network Manager, and provide information related to:
  - The physical location of the event
  - The entities affected by the event
  - The materials needed to correct the problem
- The ability to restore the patch field and associated end-to-end links to a state where they were prior to the event.

2.1.3. Detecting and Dealing With Intrusion

Intrusion detection is a current "top of mind" for many IT managers and CO personnel. The physical layer is by most measures the best way to access the equipment to be attacked. The security of telecommunications rooms can be enhanced through use of padlocks, badge readers, and other lockdown technologies, but these cannot prevent an "inside" job.

Assuming that the intrusion originated at a connector or port on a shelf or patch panel (the easiest way to access a piece of equipment), there is virtually no way to immediately detect the intrusion. The intrusion may eventually be detected by the Layer 3 security applications, but potentially not before some damage has been done. If a physical layer management system (such as one with the attributes described in the previous sections) is in place, an unauthorized access event would be generated and sent to the NOC. This would alert the NOC personnel of the event, and the appropriate measures could be taken. In addition, this event could trigger another activity such as the activation of a security camera at the location of the event, and the transmission of an image of the "intruder" to the NOC or local security.

2.1.4. Auditing and Maintaining Cross-Connects

The key to fast and effective provisioning of services lies in the immediate availability of accurate records of the link between the equipment and the entity where the service is provided. In order to keep this valuable information up to date and accurate, there must be a process of record keeping and auditing. It is a costly and error-prone activity, but a necessary one to maintain a high degree of service. To make matters worse, for the physical layer it is not the equipment side cabling or the OSP/horizontal cabling that is the item that is the most prone to changes and errors; it is the cross-connects on the patch panels.

In order to provide a solution to the above challenge, a physical layer management system must first and foremost be able to report the state of the patch panels in the link *in real time*. This relieves the need for any auditing since there is a real time indication of the cross-connects. In addition, the ability to move or share this information to a management system or database that is the source for any activities associated with these panels (e.g., MACs and other activities) would be necessary for optimal provisioning of service.

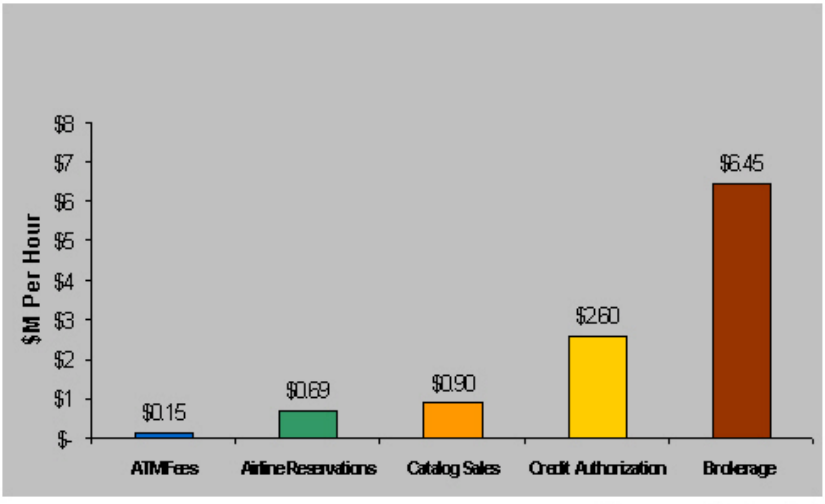

Figure 3. Financial Impact of Network Downtime

Return On Investment (ROI) For A Physical Layer Management System

Given the attributes described above, we can focus on the following areas of savings in order to calculate the ROI for the purchase of a physical layer management system:

1. "Guided" moves, adds, and changes (MACs), and the ability to do this remotely. With this capability, provisioning moves from hours to minutes. It has been shown that this area alone can provide a saving of up to 87% over present methods of operation (PMO).
2. Downtime avoidance, the compression of mean time to repair (MTTR), and avoidance of SLA penalties. Downtime is dramatically reduced, and the lost revenues associated with it. On average, IT and CO managers spend approximately 39% of their time in fault management.
3. Reuse of spare equipment ports based on the real time reporting of patch panel utilization. The average data center loses between 2% and 3% of equipment capacity due to incomplete MACs.
4. Elimination of patch panel audits for available inventory and to detect security breaches. The savings related to this item will vary greatly between locations, but could account for a significant savings depending on the size of the cross-connects and the frequency of the audits.

In most cases, the need will be to provide a ROI of less than 1 year. Given the above attributes, it is well within reason that complying physical layer management systems will meet this requirement. In come cases, the payback will come even sooner given the type of business that is being conducted (per the information contained in Figure 3 above). Additional capabilities and unique features will differentiate the systems in the market, but these are the key for ROI.

3. Current Approaches To Physical Layer Management

Physical layer management systems come in different "flavors".

There are systems that can automate the entire cross-connect operation from a remote location, including the pre-service testing of the link. These systems provide a DS3 (or higher) capable switching matrix that is suitable for remote unattended cross connections. There are also fiber-based cross-connect fabrics that perform the same functions at much higher data rates. This approach has many of the important attributes defined above, but is typically unsuitable for large installations since the cost would be prohibitive.

There are other approaches to physical layer management available today that can be categorized as a managed or an "intelligent" patch panel concept. These products provide a cross-connect or interconnect patch field where the patch panel ports are scanned on a continuous basis. All of these systems have the ability to provide real time information about the state of the patch panel connections, although different approaches are utilized to detect the presence of patch cords. Some systems are capable of detecting damaged or cut patch cords through the use of a continuity check for both copper and fiber patch cords and panels, while others depend on the presence of the ends of the cords to represent the state of the connection, assuming the cord integrity. Most of these systems have the features identified above, although guided patching is not provided with all.

All of the systems above can be monitored and controlled remotely, and have varying levels of integration with Network Management Systems for Fault Management. All provide a Configuration Management and Fault Management capable Element Manager to enable the value propositions discussed in the previous sections of this paper. The usability of these systems and the ROI are the measures by which they will be sold and implemented.

4. Conclusion

Although the first implementation was introduced over 9 years ago, physical layer management is still in its infancy.
Wise CO and IT managers are purchasing the management components and building them into mission critical links to provide the visibility needed in the critical areas of their networks, while proving in the ROI models provided by the vendors.

Clearly, we are at the beginning of this revolution, and we can expect many changes in the value proposition along the way as the different solutions are put into the working world. We can also expect that the scope of the management value proposition will grow over time, and will accordingly need to provide windows into each physical piece in the service path. In addition, the NOC managers will begin to embrace the concept of events from the physical layer as a part of their network visibility, along with the network equipment and the power infrastructure.

One can also expect that the cost of a management-capable physical layer component will decrease over time such that a non-managed component will become less attractive. These "management ready" components will be specified and installed in the physical layer from the beginning, providing a highly visible Layer 1 going forward.