PM5312

NETWORK SURVIVABILITY USING AUTOMATIC PROTECTION SWITCHING (APS) OVER SONET/SDH POINT-TO-POINT & RING NETWORKS

APPLICATION NOTE

PRELIMINARY

ISSUE 3: FEBRUARY 1998

PMC-Sierra, Inc. *105 - 8555 Baxter Place Burnaby, BC Canada V5A 4V7 604 .415.6000*

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 NETWORK SURVIVABILITY

1.1 Overview

Fibre optic technology and the ever increasing speeds of electronic processing are today fueling the creation of networks carrying more and more data along single strands of fibre. The loss of a high capacity path could knock out a large area causing disruption to crucial financial, medical and infrastructure services. To avoid such a disruption, today's networks must be designed to be fault tolerant or self healing. Figure 1 below shows how the network has evolved, and it is easy to see how a single fault could bring down a large area of the network in the absence of having built in network survivability.

The SONET/SDH standard has become widely accepted in the telecommunications industry throughout the world. One main reason for the success of this standard is the fact that it has provided key functionality to accommodate network survivability. This functionality is inherent in its transport overhead allowing intelligent circuits to be developed that can automatically select the better of a redundant incoming signal.

Figure 1 - Snapshot Of Today's Evolving Network

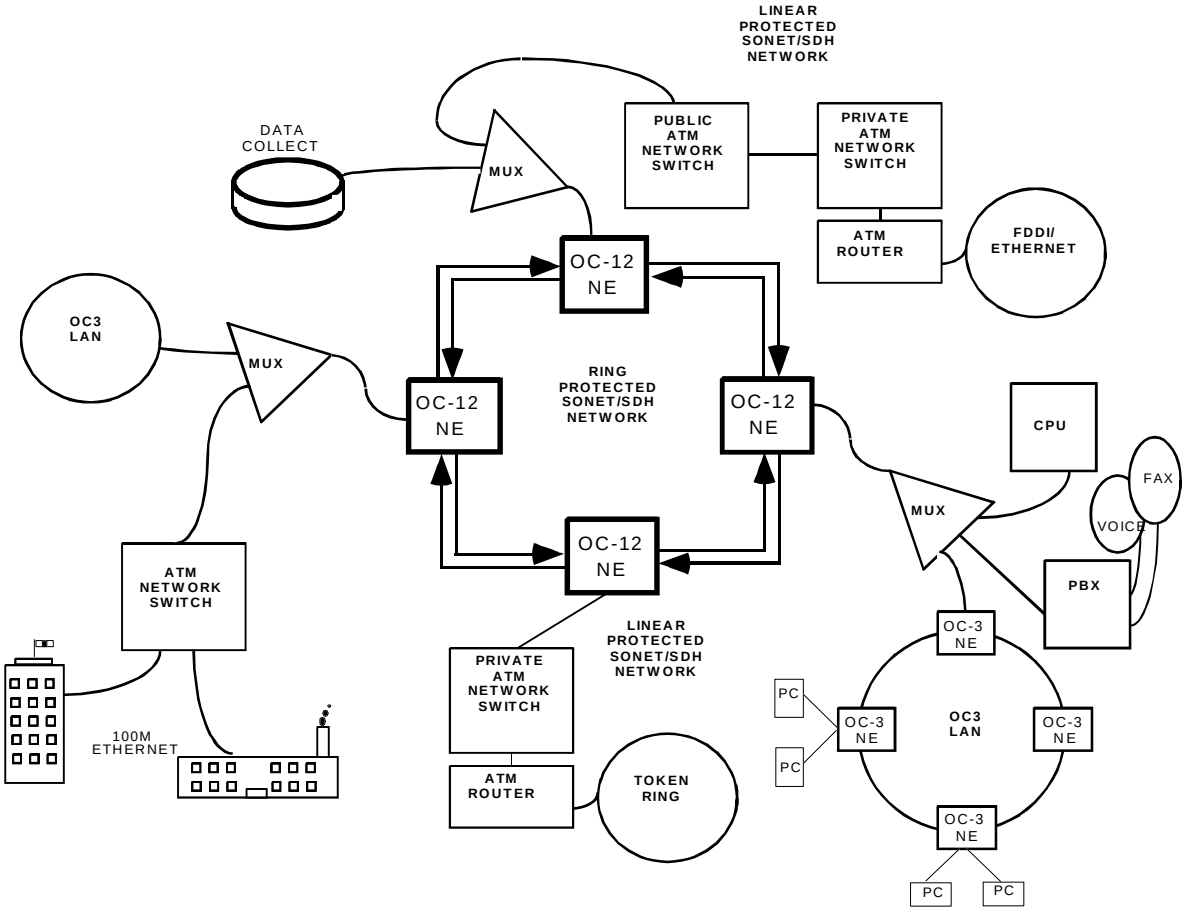

The network today has become a hybrid of both synchronous and plesiochronous transport systems. This has evolved from a point where service providers used SONET combined with a whole host of old technology overlay networks cobbled together with dedicated fibre optic cable runs to individual customers. Supporting network survivability in such a configuration was both costly and cumbersome. Also, other disadvantages such as having to support an internetworking infrastructure, the limitations of bandwidth choices (supporting only DS0, DS1, DS3, E1 and E3), the inflexibility of upgrading lines from one rate to another and the inefficient use of bandwidth, in cases such as where a 10 Mbit Ethernet connection requires a 45 Mbit DS3 link, soon became apparent. The solution to this has been addressed by today's hybrid synchronous and plesiochronous network architecture shown in figure 1. This network transforms the conventional Token Ring Protocol and LAN Ethernet to an ATM cell structure that can be transported easily over a SONET/SDH backbone. The same

SONET/SDH backbone is used to transport the voice traffic by mapping the asynchronous transport system into the SONET/SDH payload. Although ATM services are not yet widely available, they are being steadily implemented as required by demand. The ease with which this can be done using the hybrid network is accomplished by interfacing add on port cards into the SONET/SDH network rather than by re-engineering a completely new network together with its redundant survivability costs.

Although an OC-12 ring is the highest rate interconnect shown in figure 1, it is by no means the highest existing transport rate available today.

Figure 2 - Hybrid Network Showing APS Protected Links

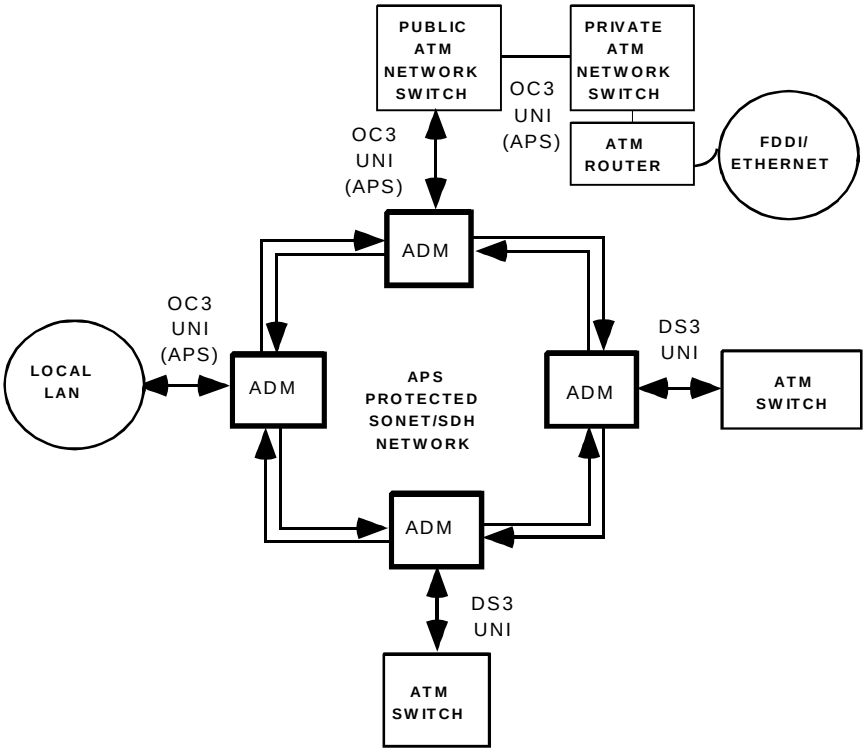

There are telecommunication manufacturers deploying OC-48 rings today with OC-192 in the planning stages for deployment in 1997-1998. The OC-3 rate is presently considered the standard for local access. The higher OC-N network paths will act as the backbone network for carrying the multimedia services that will result with the advent of the World Wide Web, the rest of the Internet, people working and running their businesses from home, and other services yet to be seen.

This document deals with the Automatic Protection Switching (APS) technique provided by the SONET/SDH standards. The APS functionality was originally developed to accommodate point-to-point connections and the path switching functionality was developed to accommodate the ring topology. However, today the standards have evolved to a point where the ring topology is also accommodated using APS line switching protection. This document will therefore address both point-to-point and ring APS protection in detail. The alternative technique of Path Switching will be introduced for future consideration but will not be the main focus of this document. Furthermore, the protection techniques discussed in this paper are applicable only where the network utilizes a SONET/SDH transport system. They are not applicable at separate ports that are non SONET/SDH based. Figure 2 shows a hybrid network and identifies the links that are applicable to the APS technique described in this document.

1.2 Basic APS Network Objectives

The point-to-point and ring APS signaling provides protection switching by providing a redundant protection path. All of the objectives applicable to a point-to-point linear APS protection architecture are also applicable to the ring network architecture. There are however, extra objectives that are applicable to the ring APS protection. This section outlines a basic idea of the objectives and requirements required to implement APS over a SONET/SDH network. A full description of the objectives and requirements can be found in GR-253-CORE, ITU-G.782 and ANSI T1.105.01.

1.2.1 Linear Network

The network will restore all traffic during a single point failure on a span between two nodes within a maximum time period of 50 msec after switch initiation. The time to accomplish switch initiation is dependent on the rate of the optical link and the provisioned bit error rate associated with the detection of the defect (such as Signal Fail, SF, and Signal Degrade, SD). Figure 6.5 in GR-253-CORE describes this in more detail.

A signal Fail (SF) condition will be triggered when the line detects a loss of signal (LOS), loss of frame (LOF), line AIS (AIS-L) or a line bit error rate (on byte B2) in excess of a user provisionable value of 10^{-3} to 10^{-5} . A signal Degrade (SD) criteria will also trigger a protection switch when the bit error rate exceeds a user provisionable value of 10^{-5} to 10^{-9}.

Link revertion back to the working channel from the protection channel will be accomplished by detecting a clearing BER threshold of one tenth the value required for detection. Once the clearing threshold has been achieved the line

must revert to a normal state after a revertion time and a Wait To Restore (WTR) time period. The clearing time is determined from the same curves that determine switch initiation time, i.e. figure 6.5 in GR-253-CORE (Issue 1, Dec. 94). As an example, this turns out to be 10 seconds for an OC-3 link provisioned to a clearing BER threshold of 10^{-7}. After a switch revertion period has elapsed a WTR period of 5-12 minutes (programmable in one minute increments) must elapse before switching back to the working channel.

Note: Programmable detection and clearing BER threshold monitoring is supported by integral circuitry included in PMC-Sierra chipsets.

1.2.2 Ring Network

All of the above objectives which apply to the linear network topology also apply to the ring topology. In addition, the following ring related objectives must be met. The ring will attempt to service multiple failures in a predictable fashion, in particular for switch requests of equal priority that cause multiple node isolation the network shall recover by segmenting into multiple sub-rings. Also, as an additional degree of protection for ring networks operating over four fibre interconnection, a mechanism to perform APS span switching is supported. All spans must have equal priority, therefore any protection switching on one span may impact the ability to execute protection switching (due to a failure of the same priority) on another span.

The maximum number of nodes on a ring must not exceed 16 due to the 4 bit node identification field in the K1 and K2 bytes of the APS protocol. The location of each node must be known by every other node on the ring network in order to avoid misconnected traffic. This is accomplished by way of a network map. Also, the state of every node must be known by all other nodes since protection channels are shared amongst multiple spans and multiple spans may be required to accomplish a protection switch. In order to accommodate this ring state knowledge, signaling over the long path and short span must be conveyed to each node. For example, although short span bridges may be established with only short span signaling, a bridge indication is sent on the long path in order to inform other nodes of the state of the ring. In addition, non time critical operations and administration messages may be transported over the DCC (Data Communications Channels of the SONET/SDH) to determine details regarding the condition of the ring.

2 APS PROTECTION SCHEMES

The three main protection schemes are line (APS) protection over a linear network, line (APS) protection over a ring network and path protection over a ring network. The two APS line protected schemes are the focus of this document and the path protection scheme will not be mentioned in the rest of this document.

2.1 Line APS In a Linear Network

This scheme utilizes the line overhead (K1 and K2) bytes of the SONET/SDH signal to protect a working channel with a protection channel. The K1 and K2 bytes are located on the line overhead of the SONET/SDH frame as shown in figure 3 below.

Figure 3 - A SONET/SDH STS-3c/STM-1 Frame Showing Transport Overhead

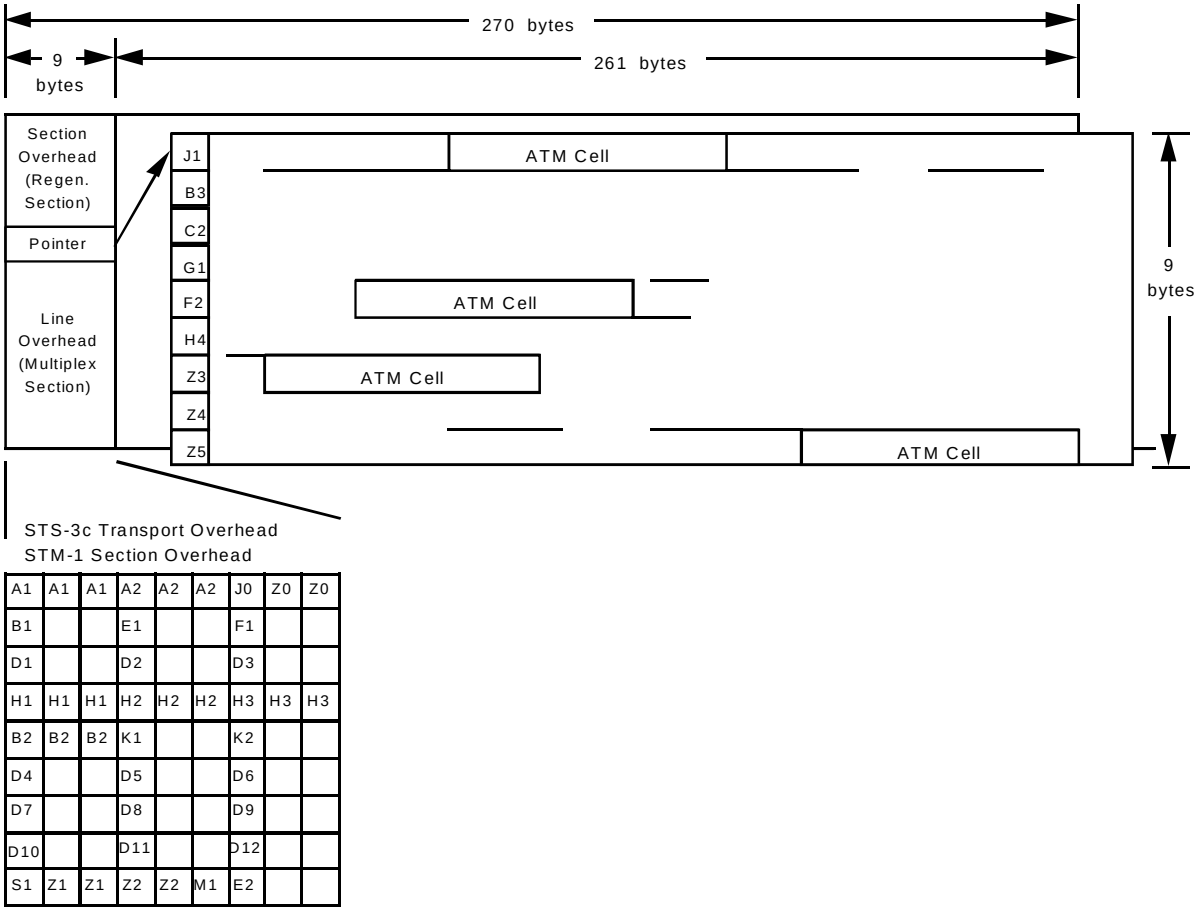

A1	A1	A1	A2	A2	A2	J0	Z0	Z0
B1			E1			F1		
D1			D2			D3		
H1	H1	H1	H2	H2	H2	H3	H3	H3
B2	B2	B2	K1			K2		
D4			D5			D6		
D7			D8			D9		
D10			D11			D12		
S1	Z1	Z1	Z2	Z2	M1	E2		

The APS protocol may either switch unidirectionally or bidirectionally and in either a revertive or non revertive mode depending upon the network management. In the case of bidirectional switching, the channel is bridged to the protection channel in both directions of transmission, and switching of only one direction is not allowed. The bridged data is then selected at the destination (in both directions) by selecting the protection channel. In the case of unidirectional switching, the switching is completed when the channel in the failed direction is switched to protection. Two configurations are defined: 1+1 (one plus one) and 1:n (one for n).

2.1.1 1:n Architecture

This is shown in figure 4, and shows a protection channel protecting up to 'n' working channels. The values for 'n' range from 1 to 15. The APS controller monitors the K1 and K2 bytes received on the protection channel and controls the 'bridging' and 'selection' of the appropriate SONET/SDH working/protection channels. To be more specific, the 'bridging' action takes place when the node transmits one of the "n" working channels over the protection channel and the 'selection' action takes place when the node selects the protection channel in place of the working channel. The reverse process takes place when the problem that caused the APS switch is fixed; the bridge is dropped and the working channel is selected instead of the protection channel.

Figure 4 - 1:n Point-to-Point APS Protection

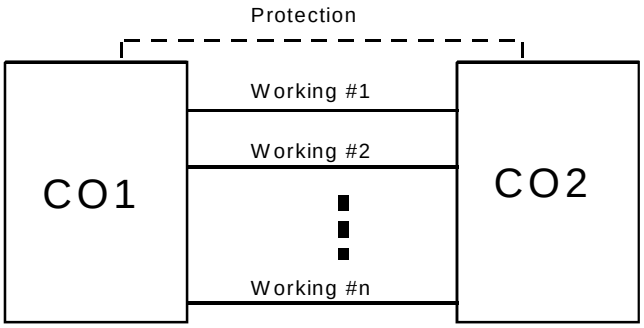

When more than one working channel is in a fault condition the channel with the higher priority is selected for protection. The priority of a working channel decreases with the channel identification number.

2.1.2 1+1 Architecture

This architecture is a simplification the 1:n architecture with the working channel permanently bridged over to the protection channel. When a fault is detected, the switch over to the protection channel takes place. There is no revertion process required for this network.

2.2 Line APS in a Ring Network

A line switched ring network can be configured as a unidirectional ring or a bidirectional ring. For the unidirectional case the traffic in both fibres travel in the same direction. In the bidirectional case both directions of a full duplex

connection travel through the same ring nodes but in opposite directions. The two directions of flow in the bidirectional ring network allows the ring some flexibility to share the load away from saturated spans. A bidirectional ring network can therefore support a higher maximum load than a unidirectional ring. Unidirectional rings have no extra advantages over bidirectional rings. The rest of this application note will therefore focus on bidirectional rings for line protected APS switching.

Bidirectional line switched rings (BLSR) can be categorized into two types; two fibre and four fibre. As in the linear line switched APS protocol, the ring working channels are protected by an unused protection channel.

2.2.1 Two Fibre Line APS Switched Ring

In this configuration both fibres carry both protection channels and working channels; 50% of the bandwidth is allocated for working traffic and the other 50% is allocated for protection. The working channels in one fibre are protected by the protection channels in the other fibre. These channels are identified on a time slot basis. To visualize the traffic traveling on an OC-N ring (for "N" not equal to 3 or 9) the channels can be thought of as a multiplex of N/3 STS-3c's (or STM-1's for SDH). These timeslots/channels are numbered from 1 to N/3. Channel numbers 1 to (N/3)/2 are assigned as working channels and channel numbers (N/3)/2 + 1 to N/3 are assigned as protection channels. This means that a working channel "m" will be protected by a protection channel (N/3)/2 + m. As an example, for an OC-12 SONET ring consisting of four STS-3c the first two STS-3c's would be allocated to working traffic and the last two STS-3cs would be protection channels. For "N" = 3 or 9 two fibre line switched application is not yet specified. Note that in these cases the line switched 4 fibre network is an alternative.

During a ring switch the working channels transmitted towards the failed span are switched to the protection channels traveling away from the failure. The bridged traffic (bridged from working to protection) travels the long way around the ring until it reaches the destination node. At this node they are switched back to the working channels traveling in the opposite direction.

Figure 5 - Ring Switch In A Two Fibre Ring Network

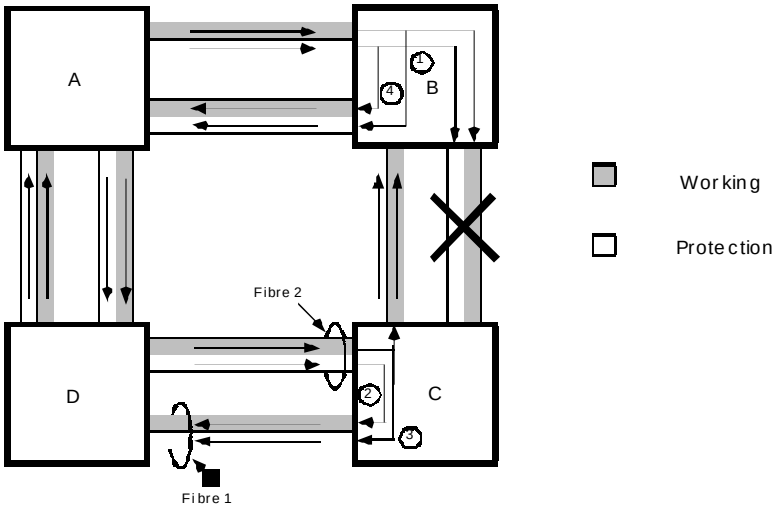

This switching action, controlled by the APS protocol, identically switches the other direction in this same manner regardless of whether that path itself is faulty or not. This is illustrated in figure 5 in a four step process.

Effectively, the failed span between node "B" and "C" has been replaced by the protection fibre between nodes "B","A", "D" and "C". Direct span switching (i.e. protection switching directly between nodes "B" and "C" is not possible in a two fibre network such as this.

In a two fibre ring network the APS codes are active only on the channels carrying protection traffic. The APS switching codes are transported over the first STS-1 (within an STS-N) line overhead K1/K2 bytes as shown in figure 3. This transport accommodates time critical switching functions to accommodate switching within 50 msec after initiation. The data communication channels in the line overhead may be utilized for non time critical functions at the users discretion.

2.2.2 Four Fibre Line APS Switched Ring

This configuration utilizes four fibre communication between nodes. The protection and working traffic is subdivided in a very different manner to the two fibre ring network. Working and protection channels are carried over different fibres. Two counter-rotating fibres are dedicated to working traffic and two counter-rotating fibres are dedicated to protection traffic. A working traffic fibre

traveling in one direction is supported by a protection fibre traveling in the opposite direction. Because of the separate working and protection fibres, this configuration can support span switching (directly between two nodes) as well as ring switching (along the long path between two nodes). Figure 6 shows this type of ring in more detail.

Figure 6 - A Four Fibre Line Switched Ring.

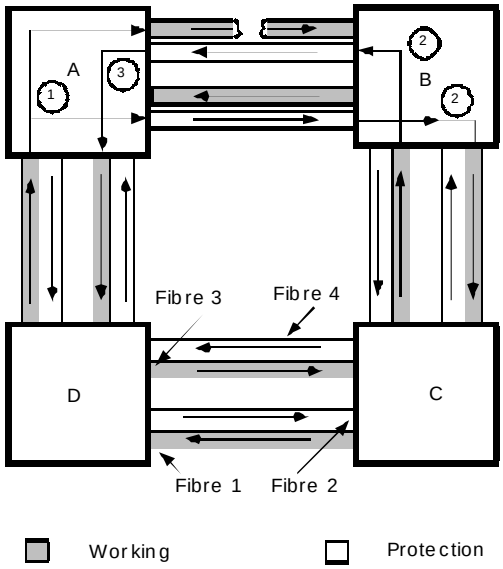

In a four fibre ring network the APS codes are active only on the fibres carrying protection traffic. The APS codes are transported over the first STS-1 (within an STS-N) line overhead K1/K2 bytes as shown in figure 3. This transport accommodates time critical switching functions to accommodate switching within 50 msec after initiation. The data communication channels in the line overhead may be utilized for non time critical functions at the users discretion. Figure 6 shows a possible fault that can occur in such a network. The switching actions required to protect against such a fault are shown in this diagram by steps labeled 1 to 3. This will be discussed in detail in the 'Ring Switch Operation' section.

3 APS FUNCTIONALITY

APS functionality over the K1 and K2 bytes is quite different in the linear line protection scheme compared to the ring line protection scheme. This section describes the format of the K1 and K2 bytes for both cases and also describes the switching actions by way of an example.

3.1 Linear APS K1 and K2 Byte Functionality

The K1 and k2 bytes signal three main types of functionality: the state of the channel, the type of fault detected by a channel or the request issued by the channel. The channel identity is carried in a 4 bit field in both the K1 and K2 bytes. The four least significant bits of the K1 byte indicates the channel identity of the channel issuing the request. The remaining 4 bits in this byte indicate the function requested by the issuing channel. Similarly, the K2 byte indicates (in the upper most significant 4 bits) the identity of the channel that is bridged, and a one bit field (the fourth most significant bit) indicates the architecture (1:n or 1+1 configuration) of the line switched APS network. The remaining three bits have a multiplicity of functionality. Line FERF is indicated when these bits are 110 binary. Line AIS is indicated when these bits are 111 binary. In SONET the value 101 binary indicates bidirectional switching and a value of 100 binary indicates unidirectional switching; in SDH applications there is no definition for these bits to indicate bidirectional or unidirectional modes. In a unidirectional system (such as video broadcast to the home) the lower three bits of the K2 bytes would indicate 100 binary. Figure 7 shows this more clearly.

Figure 7 - K1 and K2 Format.

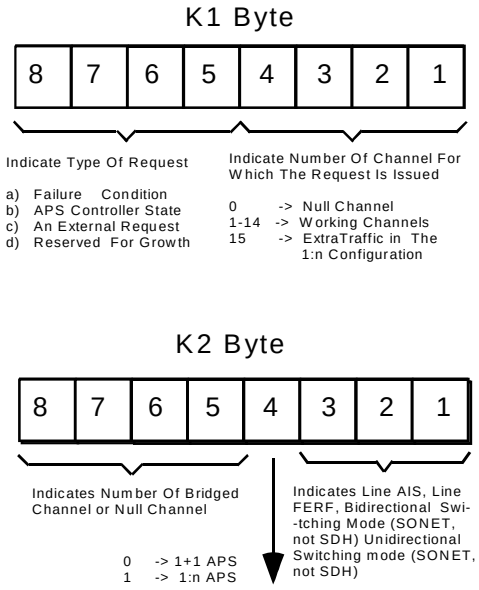

During operation, the requests and indications received on the K1 byte are evaluated in a descending priority basis as indicated in the following table.

Table 1 - Linear APS Request Types Ordered In Priority

Code Received On Upper Nibble	Condition, State or Request	Priority Order
1111	Lockout Protection	Highest
1110	Force Switch	|
1101	Signal Fail (SF_H) High Priority	|
1100	Signal Fail (SF_L) low Priority	|
1011	Signal Degrade (SD_H) High Priority	|
1010	Signal Degrade (SD_L) Low Priority	|
1001	Unused	|
1000	Manual Switch	|
0111	Unused	|

Code Received On Upper Nibble	Condition, State or Request	Priority Order
0110	Wait To Restore (WTR)	|
0101	Unused	|
0100	Exercise	|
0011	Unused	|
0010	Reverse Request	|
0001	Do Not Revert	|
0000	No Request	Lowest

The K1 and K2 bytes travel over the protection line to the APS controller. After receiving these bytes the APS controller must apply a three frame persistency check on all received K1 and K2 values before acting on the request. Invalid codes must be discarded without further action.

3.2 Linear APS Switch Operation

This operation describes the APS protocol applied to a 1:n APS network consisting of nodes A, B and C. In order to describe the APS protocol a series of fault conditions applied to the network shown in figure 8 will be analyzed. Not all of the faults and requests shown in Table 1 are analyzed since the number of all possible combinations are too many, however, the principle of the switching protocol will apply in a similar fashion to all these requests on a prioritized basis.

Figure 8 - 1:n Linear APS Network In No Fault Steady State Condition

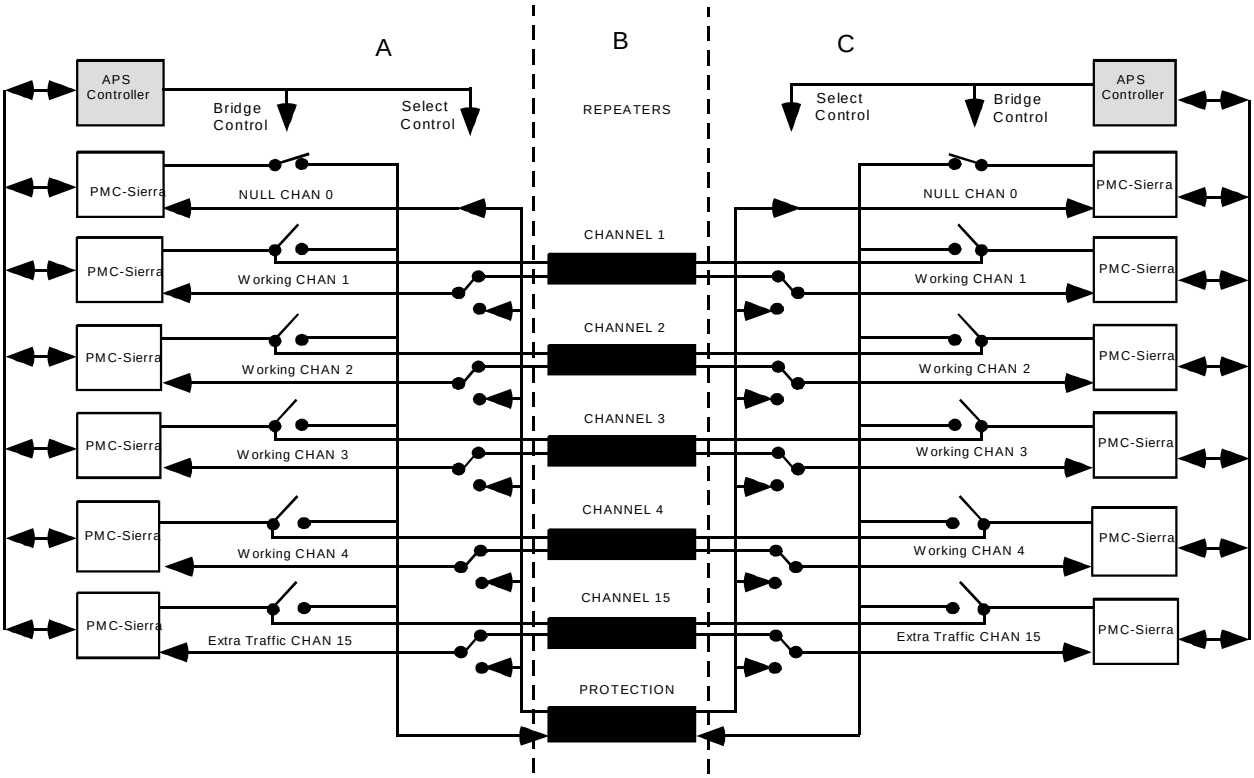

Figure 8 shows a linear network consisting of nodes A, B and C. Node B is just a repeater station and does not terminate line overhead and is therefore uninvolved in the APS protocol. Initially the network is in a idle state and no requests are active on the K1 and K2 bytes; node C to node A, K1 value is 00000000b' and the K2 byte is 00001101. The same values are transmitted by node A to node C. Note that in this initial state the APC controller on each end of the network is listening to the APS bytes on the protection channel from the other end. The source of the protection channel APS bytes can be selected arbitrarily to be any one of the working channels or the null channel. In the example above in figure 8, all working channels are communicating without error and the null channel is connected over the protection channel. The APS controller listens to all channels for the determination of SD or SF alarms (generated by monitoring for excessive bit error rates) but only needs to listen to the protection channel for the K1 and K1 APS bytes. For those PMC-Sierra devices that support BER monitoring, this is an easy task since they will not interrupt the APS controller unless there has been a change of state in the BER threshold monitoring. The standards call for the monitoring and transmission of the APS bytes over the protection channel, however, the K1 and K2 byte

functionality within the PMC-Sierra parts allows the APS controller to optionally listen to the K1 and K2 bytes received on all of the "n" channels instead of the single protection channel. The advantage would be to make the null channel circuitry redundant (assuming that no 'Extra Traffic' is carried on the null channel).

3.2.1 Response To Signal Degrade Detected

Next, let's assume that a signal degrade on working channel 2 has been detected by node C. Node C immediately sends a bridge request to node A by transmitting K1 = 10100010b' and K2 = 00001101b' (K2 stays the same).

When node A receives this request it bridges the working channel 2 over to the protection channel and sends back K1 = 00100010b'; meaning that it is requesting node C to do the same as node A (a reverse request) for channel 2. The K2 byte sent back to C indicates 00101101b; meaning that node A has bridged channel 2.

When node C receives the K1 and K2 bytes from A, it performs a switch (because node A has bridged working channel 2) and a bridge (because of the reverse request from node A) on working channel 2. After having performed these actions, node C sends back K1 = 10100010b' and K2 = 00101101b'. This indicates that node C is still detecting a SD on working channel 2 and that it has bridged working channel 2 to protection.

When node A receives the K1 and K2 values transmitted from node C it switches over to select protection on working channel 2. This completes the switching protocol for a signal degrade on working channel 2. In this state, the APS controller must signal further switching messages through channel two equipment since this channel now has control of the protection channel.

3.2.2 Response To Signal Fail Detected

The next event to occur is a signal fail detected by node A on channel 1. Node A transmits K1 = 11000001b' (signal fail on working channel 1) and an unchanged K2 byte. Node A releases the working channel 2 switch that was initiated during the previous SD failure.

On receiving the SF indication from node A, Node C bridges working channel 1 onto protection and releases the protection switch on channel 2 from the previous SD switching. Node C sends back a reverse request on its K1 byte and informs Node A that it has bridged working channel 1 onto protection by the value in its K2 byte. Therefore K1 = 00100001b' and K2 = 00011101b'.

When Node A receives the new K1 and K2 bytes from node C, node A switches to select the protection channel for working channel 1 (due to the bridge at node C indicated by the received K2 byte) and bridges working channel 1 to the protection channel (as requested by node C on its K1 byte). After taking these actions node A transmits K1 = 11000001b' (as before) and K2 = 00011101b' (signaling that node A has bridged working channel 1).

To complete the bidirectional APS switching protocol, node C switches to select the protection channel for it working channel 1 (due to the received K2 byte from node A).

3.2.3 Signal Fail Repaired

At this point a steady state is achieved where node A continually transmits K1 = 11000001b' and K2 = 00011101b, while node C transmits K1 = 00100001b' and K2 = 00011101b'. This state will now only change when the failure situation between the two nodes changes. If the SD fault is repaired there will be no APS action because the SF condition (of higher priority) will remain keeping the two nodes at the present state. However, if the SF condition is repaired before the SD condition, the APS protocol will terminate the switched condition due to SF on channel 1 and re-establish the SD switched condition on channel 2. To understand this process we consider the repair of the SF condition on working channel 1.

Node A detects that the SF condition has been repaired and enters a WTR (Wait to Restore) state by signaling K1 = 01100001b' and unchanged K2 byte to node C.

Because working channel 2 is still degraded, node C signals K1 = 10100010b' (signal degraded on working channel 2) and drops its previous switch to protection for working channel 1. It remains (for now) bridged on working channel 1 by transmitting K2 = 00011101b'.

Node A detects the new request and bridges working channel 2 to protection thereby transmitting K2 = 00101101b' and releases the selection of protection channel for channel 1 traffic. The K1 byte signals back a reverse request for channel 2 to node C by setting K1 = 00100010b'.

Node C responds to the reverse request by bridging working channel 2 to protection and switching to the received protection channel for traffic on channel 2. This is translated by sending back K1 = 10100010b' (still indicating SD on channel 2) and K2 = 00101101b' (indicating a bridge to channel 2) to node A.

To complete the bidirectional switching action for the SD failure, node A switches to select the protection channel for channel 2 (due to the received K2 byte from node C).

3.2.4 Signal Degrade Repaired

At this point the network enters another steady state where node A continually transmits K1 = 00100010b' and K2 = 00101101b', while node C transmits K1 = 10100010b' and K2 = 00101101b'. In order to analyze how the network returns to a no failure state, we consider the repair of the SD condition on channel 2.

On detecting a repair on channel 2, node C enters a WTR state and signals this to the other end by transmitting K1 = 01100010b'. The K2 byte remains unchanged from 00101101b'. After the expiration of the WTR period node C transmits no failure condition on its K1 byte and releases the switch that selects protection channel for channel 2. The K1 byte of 00000000b' indicating no request is transmitted to node A. The K2 byte still indicates a bridge of working channel 2 to protection because this has not yet been cleared, i.e. K2 still indicates 00101101b'.

In response to the newly received K1 byte, node A drops the bridge on working channel 2 releases the selection of protection on working channel 2. The K1 byte transmitted towards node C is changed to 00000000b' to reflect that there is no reverse request required.

Node C detects the all zeroes K1 byte from node A and clears its bridge to protection of working channel 2. A final no fault steady state results with both ends of the network transmitting K1 = 00000000b' and K2 = 0001101b'.

All other combinations of error states that may have arisen would be dealt with in a similar manner according to the priority of each request listed in table 1.

3.3 Ring APS K1 and K2 Byte Functionality

The functionality of the ring APS K1 byte is identical to the K1 byte in the linear APS switching application although there are new definitions of the request type carried in the upper nibble. The K2 byte is redefined. The upper nibble indicates the source node identification and the 4th most significant bit indicates whether the request is over the long path or the span (short path).

Figure 9 - Ring K1 and K2 Format.

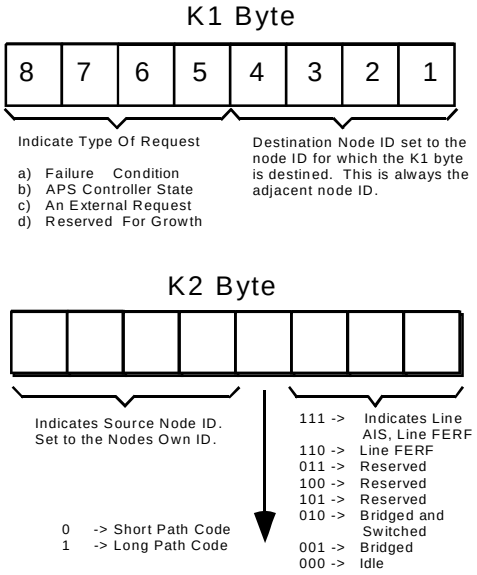

The three least significant bits in this byte indicate line AIS, line FERF, bridged & switched, bridged and idle conditions. Figure 9 shows this in more detail.

Table 2 - Ring APS Request Types Ordered In Priority

Code Received On Upper Nibble Of K1 Byte	Condition, State or Request	Priority Order
1111	Lockout Protection (span) or Signal Fail On Protection (LP-S)	Highest
1110	Force Switch (Span)(FS-S)	|
1101	Force SwitchRing(FS-R)	|
1100	Signal Fail (Span)(SF-S)	|
1011	Signal Fail (Ring)(SF-R)	|
1010	Signal Degrade (Protection)(SD-P)	|
1001	Signal Degrade (Span)(SD-S)	|
1000	Signal Degrade (Ring)(SD-R)	|

Code Received On Upper Nibble Of K1 Byte	Condition, State or Request	Priority Order
0111	Manual Switch (Span)(MS-S)	|
0110	Manual Switch (Ring) (MS-R)	|
0101	Wait To Restore (WTR)	|
0100	Exercise (Span) (EXER-S)	|
0011	Exercise (Ring) (EXER-R)	|
0010	Reverse Request (Span) (RR-S)	|
0001	Reverse Request (Ring) (RR-R)	|
0000	No Request (NR)	Lowest

During operation, the requests and indications received on the K1 byte are evaluated in a descending priority basis as indicated in the following table.

The K1 and K2 bytes always travel over the protection line to the APS controller. The APS controller must apply a three frame persistency check on all received K1 and K2 values before acting on the request.

3.4 Ring Switch Operation

This operation describes the APS protocol applied to a multiple node ring network. The diagram of figure 10 shows a four fibre ring network with a signal failure detected on one of the working fibres between two nodes. The switching action required to protect against such a failure is discussed in the following section as well as the restoration process to an idle (no request) state.

Figure 10 - Ring SF-S Failure On Working Fibre From Node E To Node F

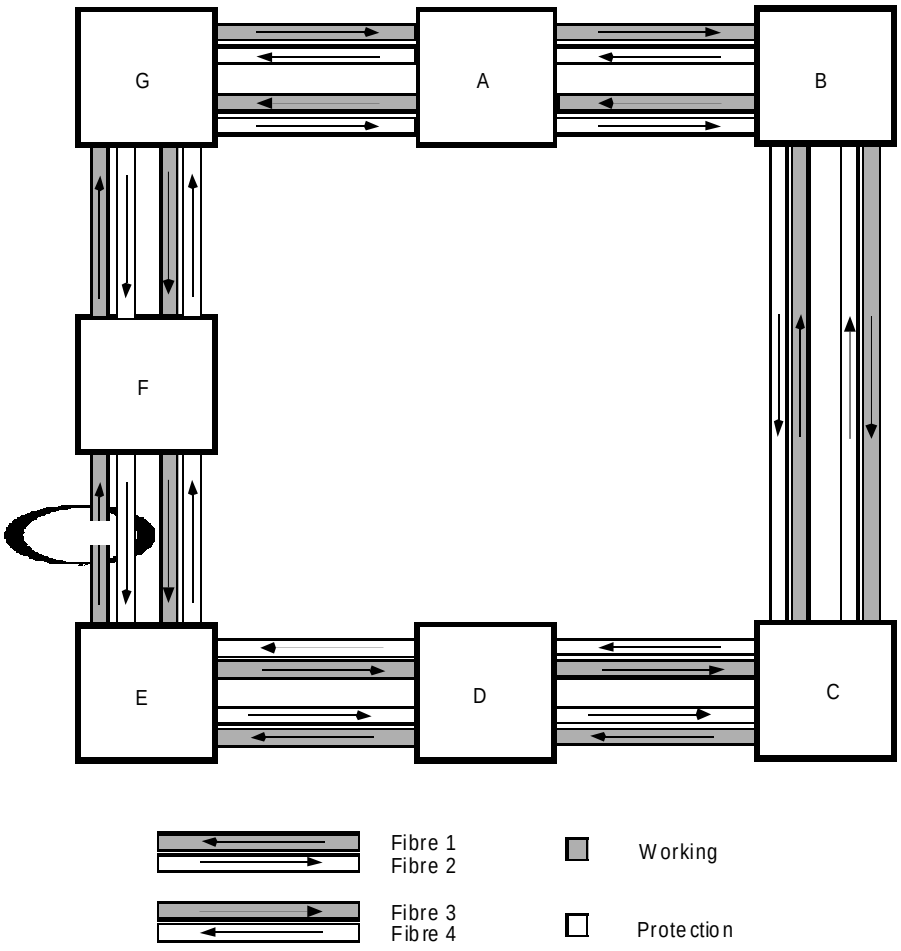

In the initial no error state, all nodes transmit the K1 and K2 bytes to the adjacent nodes with each node signaling NR (no request) and the destination node ID on the K1 byte and the source node ID and IDLE state on the K2 byte. For example node A transmits NR/B (the mapping of node ID's is established by the ring map data table that is provisioned by the network management software) on the K1 byte and A/S/IDLE on the K2 byte to node B. Node B in turn transmits NR/A on its K1 byte and B/S/IDLE on its K2 byte to node A. This is a steady state condition.

3.4.1 Signal Fail Detected on a Span

When node F detects a SF-S from node E, node F transmits SF-S/E on its K1 byte and F/S/IDLE on the K2 byte of the short path between nodes E and F. On the long path the same information is sent except the K2 byte indicates long path instead of short span. When node E detects the failure signal on the K1 and K2 byte over the short span it initiates a bridge of the working channel to the protection channel. In the meantime the same request is traveling round the long path but has not yet been received by node E. The long path transfer trickles through intermediate nodes transparently by placing them into K1/K2 byte bypass mode and will eventually be received by node E and will have no effect.

After node E bridges the working traffic to the protection channel it starts to send its own reverse request to F along the short path and the SF-S condition on the long path; i.e. K1:K2 = RR-S/F:E/S/Br and SF-S/F:E/L/Br. Again the span request will be received by Node F before the long path request. The long path node request will be received by node F eventually but will have no effect. NB: From now on, all long path communication will be ignored since in this example it will always be beaten by the quicker short span path. Long path communication is a factor only when the short path is also affected by the failure.

Node F receives the K1 and K2 bytes from node E and switches to the protection channel. It also obeys the RR-S request by bridging its working channel to the protection channel. Finally it send the new K1 and K2 bytes to E indicating SF-S/E:F/S/Br&Sw on the short path and SF-S/E:F/L/Br&Sw on the long path.

Node E receives the K1 and K2 bytes from node F and switches to select the bridged protection channel from F. Node E conveys its status to node F by sending the K1:K2 bytes RR-S/F:E/S/Br&Sw on the short path and SF-S/F:E/L/Br&Sw on the long path.

This completes the switching action required to protect an SF-S condition shown in figure 10. A steady state is reached where the K1 and K2 bytes are sourced by nodes E and F and all other nodes are transparently passing through these bytes.

3.4.2 Signal Fail Repaired on a Span

When node F detects a repair of the SF-S condition from node E, node F transmits WTR/E:F/S/Br&Sw on its K1 and K2 bytes along the short span and WTR/E:F/L/Br&Sw on the long path.

Node E detects the WTR request from F and reacts in a similar fashion. First it generates an WTR request of its own along the long path to node F (WTR/F:E/L/Br&Sw) and also it generates a reverse request (RR-S/F:E/S/Br&Sw) to node F in acknowledgment to the WTR received and to prepare node F to tear down its switching after the expiration of the WTR period.

After the WTR period expires, node F drops the span switch (in response to the RR-S request from node E) and generates a no request (NR) to node E by transmitting NR/E:F/S/Br on the short path and NR/E:F/L/Br on the long path.

Node E receives the latest K1 and K2 bytes from node F and drops the switch and bridge functions. It also generates new K1 and K2 bytes indicating NR/F:E/S/IDLE on the short path and NR/F:E/L/IDLE on the long path. Node E is now back to its no error (idle) state.

Node F inspects the new APS bytes from node E and clears down its bridge onto the protection fibre (since node E is now no longer listening to it).

Now that node F and E are not listening to the protection channel and they have both stopped bridging working channels to the protection channel, they both start transmitting NR and IDLE codes addressed to the adjacent neighbour. The adjacent neighours detect this condition and drop out of the APS bypass mode by returning to the condition where they too are sending NR and IDLE to their neighours. This causes a trickle down effect until all nodes are transmitting NR and IDLE to its adjacent neighbour. A normal "no fault" steady state now exists with all working channels operating without failure and the protection channels operating in the standby condition.

4 ON-CHIP APS SUPPORT FEATURES

The APS functionality is supported by all PMC-Sierra SONET and ATM chipsets to a varying degree of functionality. These features have evolved alongside the SONET/SDH standards and table 3 below shows the functionality covered for each PMC-Sierra device.

One of the on-chip features includes the extraction and insertion of the APS K1/K2 bytes. These bytes (K1, K2) are extracted into the Receive K1 Register and the Receive K2 Register after they have been filtered for three frames; at which point they are written to the internal holding registers. A protection switching byte failure (PSBF) alarm is declared when twelve successive frames have been received, where no three consecutive frames contain identical K1 bytes. The protection switching byte failure alarm is removed upon detection of three consecutive frames containing identical K1 bytes. The detection of invalid APS codes must be done in software by polling the Receive K1/K2 Registers. Interrupts are generated to signal when an alarm is triggered or when the K1/K2 value changes state.

The other main APS on chip functionality consists of BER threshold monitoring. The line level bit-interleaved parity (B2) is computed, and compared to the received B2 bytes. Line BIP-8 errors are accumulated in an internal counter and these can be accessed by software to determine excessive bit error rates and hence trigger an SF (signal Fail) or SD (signal degrade). In some cases, the SF and SD threshold crossing alarm functionality is integrated into the PMC-Sierra parts and automatically generates an interrupt when these thresholds are exceeded.

Table 3: PMC-Sierra Device Integrated APS Functionality

Device NAME	K1/K2 Insertion	K1/K2 Extraction	K1/K2 Persistency	'SF' DETECTION AND Clearing	'SD' DETECTION AND Clearing
PM5343 STXC	YES	YES	YES	YES	YES
PM5345 S/UNI-155	SERIAL PORT ACCESS	*SERIAL PORT ACCESS	NO	No	NO
PM5347 S/UNI-PLUS	YES	YES	YES	**BERM BLOCK Detection only	**BERM BLOCK Detection only
PM5346 S/UNI-LITE	No	NO	NO	NO	NO
PM5355 S/UNI-622	YES	YES	YES	**BERM BLOCK Detection only	**BERM BLOCK Detection only
PM5348 S/UNI-DUAL	YES	YES	YES	**YES	**YES
PM5312 STTX	YES	YES	YES	NO	NO
PM5342 SPECTRA-155	YES	YES	YES	YES	YES

Notes:

** One single BERM block is provided to detect either SD or SF thresholds, not both.

* There is no persistency applied to the K1/K2 bytes on the serial extract port.

The programmable registers allows the user to program the threshold value to declare and clear such alarms over a range of 10-3 to 10-9. The programmable bit error rate values associated with the declaration and clearing of the SF and SD alarms should be determined from the latest standards.

Table 3 above shows the degree of inbuilt APS functionality across the complete line of ATM and SONET/SDH chip-sets offered by PMC-Sierra.

The above chipsets also provide functionality to insert and extract the Z1 (S1) byte in the line overhead. This byte is not used in the APS protocol. However, it

provides a knowledge of the quality of clock being used by the far end transmitting node, and can be used the receiving node to select the better of two clocks in a looptimed application. For example, if the transmitting end has a better stratum clock than the receiving end, the receiving end could use the better clock by using looptimed mode of operation.

5 REFERENCES

- PMC-Sierra, Inc., PM5347 S/UNI-PLUS Data Sheet, Issue 4, Oct, 1995.
- PMC-Sierra, Inc., PM5343 SONET-STXC Data Sheet, Issue 3, February, 1995.
- PMC-Sierra, Inc., PM5345 S/UNI Data Sheet, Issue 3, February, 1995.
- PMC-Sierra, Inc., PM5346 S/UNI-Lite Data Sheet, Issue 6, Mar, 1996.
- PMC-Sierra, Inc., PM5355 S/UNI-622 Data Sheet, Issue 2, Apr, 1996.
- PMC-Sierra, Inc., PM5348 S/UNI-DUAL Data Sheet, Issue 3, Mar, 1996.
- PMC-Sierra, Inc., PM5348 SONET-STTX Data Sheet, Issue 4, Apr, 1995.
- ANSI, Synchronous Optical Network (SONET) Automatic Protection Switching, ANSI T1.105.01-1994.
- Fibre Network Survivability, Tsong-Ho Wu.
- ITU-T Recommendation G.782, Types and general Characteristics of Synchronous Digital Hierarchy (SDH) Equipment, January 1994.
- ITU-T Recommendation G.803, Architectures Of Transport Networks Based On The Synchronous Digital Hierarchy (SDH) Equipment Functional Blocks, March 1993.
- ITU-T Recommendation G.783, Characteristics Of Synchronous Digital Hierarchy (SDH) Equipment, March 1993.

NOTES

CONTACTING PMC-SIERRA, INC.

PMC-Sierra, Inc.
105-8555 Baxter Place Burnaby, BC
Canada V5A 4V7

Tel: (604) 415-6000

Fax: (604) 415-6200

Document Information: document@pmc-sierra.com
Corporate Information: info@pmc-sierra.com
Application Information: apps@pmc-sierra.com
Web Site: http://www.pmc-sierra.com

PMC-960505 (R3) Issue date October 1997