

Independent Submission
Request for Comments: 7347
Category: Informational
ISSN: 2070-1721

H. van Helvoort, Ed.
Huawei Technologies
J. Ryoo, Ed.
ETRI
H. Zhang
Huawei Technologies
F. Huang
Philips
H. Li
China Mobile
A. D'Alessandro
Telecom Italia
September 2014

Pre-standard Linear Protection Switching in
MPLS Transport Profile (MPLS-TP)

Abstract

The IETF Standards Track solution for MPLS Transport Profile (MPLS-TP) Linear Protection is provided in RFCs 6378, 7271, and 7324.

This document describes the pre-standard implementation of MPLS-TP Linear Protection that has been deployed by several network operators using equipment from multiple vendors. At the time of publication, these pre-standard implementations were still in operation carrying live traffic.

The specified mechanism supports 1+1 unidirectional/bidirectional protection switching and 1:1 bidirectional protection switching. It is purely supported by the MPLS-TP data plane and can work without any control plane.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7347>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	4
2. Conventions Used in This Document	5
3. Acronyms	5
4. Linear Protection-Switching Overview	6
4.1. Protection Architecture Types	6
4.1.1. 1+1 Architecture	6
4.1.2. 1:1 Architecture	6
4.1.3. 1:n Architecture	7
4.2. Protection Switching Type	7
4.3. Protection Operation Type	7
5. Protection-Switching Trigger Conditions	8
5.1. Fault Conditions	8
5.2. External Commands	8
5.2.1. End-to-End Commands	8
5.2.2. Local Commands	9
6. Protection-Switching Schemes	10
6.1. 1+1 Unidirectional Protection Switching	10
6.2. 1+1 Bidirectional Protection Switching	11
6.3. 1:1 Bidirectional Protection Switching	12
7. APS Protocol	13
7.1. APS PDU Format	13
7.2. APS Transmission	16
7.3. Hold-Off Timer	17
7.4. WTR Timer	17
7.5. Command Acceptance and Retention	18
7.6. Exercise Operation	18
8. Protection-Switching Logic	19
8.1. Principle of Operation	19
8.2. Equal Priority Requests	21
8.3. Signal Degrade of the Protection Transport Entity	22
9. Protection-Switching State Transition Tables	22
10. Security Considerations	24
11. Acknowledgements	24
12. References	24
12.1. Normative References	24
12.2. Informative References	25
Appendix A. Operation Examples of the APS Protocol	26

1. Introduction

The IETF Standards Track solution for MPLS Transport Profile (MPLS-TP) Linear Protection is provided in [RFC6378], [RFC7271], and [RFC7324].

This document describes the pre-standard implementation of MPLS-TP Linear Protection that has been deployed by several network operators using equipment from multiple vendors. At the time of publication, these pre-standard implementations were still in operation carrying live traffic.

This implementation was considered in the MPLS WG; however, a different path was chosen.

This document may be useful in the future if a vendor or operator is trying to interwork with a different vendor or operator who has deployed the pre-standard implementation, and it provides a permanent record of the pre-standard implementation. It is also worth noting that the experience gained during deployment of the implementations of this document was used to refine [RFC7271].

MPLS-TP is defined as the transport profile of MPLS technology to allow its deployment in transport networks. A typical feature of a transport network is that it can provide fast protection switching for end-to-end transport paths and transport path segments. The protection-switching time is generally required to be less than 50 ms to meet the strict requirements of services such as voice, private line, etc.

The goal of a linear protection-switching mechanism is to satisfy the requirement of fast protection switching for an MPLS-TP network. Linear protection switching means that, for one or more working transport entities (working paths), there is one protection transport entity (protection path), which is disjoint from any of the working transport entities, ready to take over the service transmission when a working transport entity has failed.

This document specifies a 1+1 unidirectional protection-switching mechanism for a unidirectional transport entity (either point to point or point to multipoint) as well as a bidirectional point-to-point transport entity and a 1+1/1:1 bidirectional protection-switching mechanism for a point-to-point bidirectional transport entity. Since bidirectional protection switching needs the coordination of the two endpoints of the transport entity, this document also specifies the Automatic Protection Switching (APS) protocol, which is used for this purpose.

The linear protection mechanism described in this document is applicable to both Label Switched Paths (LSPs) and Pseudowires (PWs).

The APS protocol specified in this document is based on the same principles and behavior of the APS protocol designed for Synchronous Optical Network (SONET) [T1.105.01] / Synchronous Digital Hierarchy (SDH) [G.841], Optical Transport Network (OTN) [G.873.1], and Ethernet [G.8031] and provides commonality with the established operation models utilized in transport network technologies (e.g., SDH/SONET, OTN, and Ethernet).

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Acronyms

This document uses the following acronyms:

APS	Automatic Protection Switching
DNR	Do not Revert
EXER	Exercise
G-ACh	Generic Associated Channel
FS	Forced Switch
LO	Lockout of Protection
LSP	Label Switched Path
MPLS-TP	MPLS Transport Profile
MS	Manual Switch
MS-P	Manual Switch to Protection transport entity
MS-W	Manual Switch to Working transport entity
NR	No Request
OAM	Operations, Administration, and Maintenance
OTN	Optical Transport Network
PDU	Protocol Data Unit
PW	Pseudowire
RR	Reverse Request
SD	Signal Degrade
SD-P	Signal Degrade on Protection transport entity
SD-W	Signal Degrade on Working transport entity
SDH	Synchronous Digital Hierarchy
SF	Signal Fail
SF-P	Signal Fail on Protection transport entity
SF-W	Signal Fail on Working transport entity
SONET	Synchronous Optical Network
WTR	Wait to Restore

4. Linear Protection-Switching Overview

To guarantee the protection-switching time for a working transport entity, its protection transport entity is always preconfigured before the failure occurs. Normally, traffic will be transmitted and received on the working transport entity. Switching to the protection transport entity is usually triggered by link or node failure, external commands, etc. Note that external commands are often used in transport networks by operators, and they are very useful in cases of service adjustment, path maintenance, etc.

4.1. Protection Architecture Types

4.1.1. 1+1 Architecture

In the 1+1 architecture, the protection transport entity is associated with a working transport entity. The normal traffic is permanently bridged onto both the working transport entity and the protection transport entity at the source endpoint of the protected domain. The normal traffic on working and protection transport entities is transmitted simultaneously to the destination sink endpoint of the protected domain, where a selection between the working and protection transport entity is made based on predetermined criteria, such as signal fail and signal degrade indications.

4.1.2. 1:1 Architecture

In the 1:1 architecture, the protection transport entity is associated with a working transport entity. When the working transport entity is determined to be impaired, the normal traffic MUST be transferred from the working to the protection transport entity at both the source and sink endpoints of the protected domain. The selection between the working and protection transport entities is made based on predetermined criteria, such as signal fail and signal degrade indications from the working or protection transport entity.

The bridge at the source endpoint can be realized in two ways: it is either a selector bridge or a broadcast bridge. With a selector bridge, the normal traffic is connected either to the working transport entity or the protection transport entity. With a broadcast bridge, the normal traffic is permanently connected to the working transport entity, and in case a protection switch is active, it is also connected to the protection transport entity. The broadcast bridge is recommended to be used in revertive mode only.

4.1.3. 1:n Architecture

Details for the 1:n protection-switching architecture are out of scope of this document and will be provided in a different document in the future.

It is worth noting that the APS protocol defined here is capable of supporting 1:n operations.

4.2. Protection Switching Type

The linear protection-switching types can be a unidirectional switching type or a bidirectional switching type.

- o Unidirectional switching type: Only the affected direction of the working transport entity is switched to the protection transport entity; the selectors at each endpoint operate independently. This switching type is recommended to be used for 1+1 protection in this document.
- o Bidirectional switching type: Both directions of the working transport entity, including the affected direction and the unaffected direction, are switched to the protection transport entity. For bidirectional switching, the APS protocol is required to coordinate the two endpoints so that both have the same bridge and selector settings, even for a unidirectional failure. This type is applicable for 1+1 and 1:1 protection.

4.3. Protection Operation Type

The linear protection operation types can be a non-revertive operation type or a revertive operation type.

- o Non-revertive operation: The normal traffic will not be switched back to the working transport entity even after a protection switching cause has cleared. This is generally accomplished by replacing the previous switch request with a "Do not Revert (DNR)" request, which has a low priority.
- o Revertive operation: The normal traffic is restored to the working transport entity after the condition(s) causing the protection switching has cleared. In the case of clearing a command (e.g., Forced Switch), this happens immediately. In the case of clearing a defect, this generally happens after the expiry of a "Wait to Restore (WTR)" timer, which is used to avoid chattering of selectors in the case of intermittent defects.

5. Protection-Switching Trigger Conditions

5.1. Fault Conditions

Fault conditions mean the requests generated by the local Operations, Administration, and Maintenance (OAM) function.

- o Signal Fail (SF): If an endpoint detects a failure by an OAM function or other mechanism, it will submit a local signal failure (local SF) to the APS module to request a protection switch. The local SF could be on the working transport entity (Signal Fail on Working transport entity (SF-W)) or the protection transport entity (Signal Fail on Protection transport entity (SF-P)).
- o Signal Degrade (SD): If an endpoint detects signal degradation by an OAM function or other mechanism, it will submit a local signal degrade (local SD) to the APS module to request a protection switching. The local SD could be on the working transport entity (Signal Degrade on Working transport entity (SD-W)) or the protection transport entity (Signal Degrade on Protection transport entity (SD-P)).

5.2. External Commands

The external command issues an appropriate external request to the protection process.

5.2.1. End-to-End Commands

These commands are applied to both local and remote nodes. When the APS protocol is present, these commands, except the Clear command, are signaled to the far end of the connection. In bidirectional switching, these commands affect the bridge and selector at both ends.

- o Lockout of Protection (LO): This command is used to provide the operator a tool for temporarily disabling access to the protection transport entity.
- o Manual Switch (MS): This command is used to provide the operator a tool for temporarily switching normal traffic to the working transport entity (Manual Switch to Working transport entity (MS-W)) or to the protection transport entity (Manual Switch to Protection transport entity (MS-P)), unless a higher priority switch request (i.e., LO, FS, or SF) is in effect.

- o Forced Switch (FS): This command is used to provide the operator a tool for temporarily switching normal traffic from the working transport entity to the protection transport entity, unless a higher priority switch request (i.e., LO or SF-P) is in effect.
- o Exercise (EXER): Exercise is a command to test if the APS communication is operating correctly. The EXER command SHALL NOT affect the state of the protection selector and bridge.
- o Clear: This command between management and the local protection process is not a request sent by APS to other endpoints. It is used to clear the active near-end external command or WTR state.

5.2.2. Local Commands

These commands apply only to the near end (local node) of the protection group. Even when an APS protocol is supported, they are not signaled to the far end.

- o Freeze: This command freezes the state of the protection group. Until the freeze is cleared, additional near-end commands are rejected, and condition changes and received APS information are ignored. When the Freeze command is cleared, the state of the protection group is recomputed based on the condition and received APS information.

Because the freeze is local, if the freeze is issued at one end only, a failure of protocol can occur as the other end is open to accept any operator command or fault condition.

- o Clear Freeze: This command clears the local freeze.

6. Protection-Switching Schemes

6.1. 1+1 Unidirectional Protection Switching

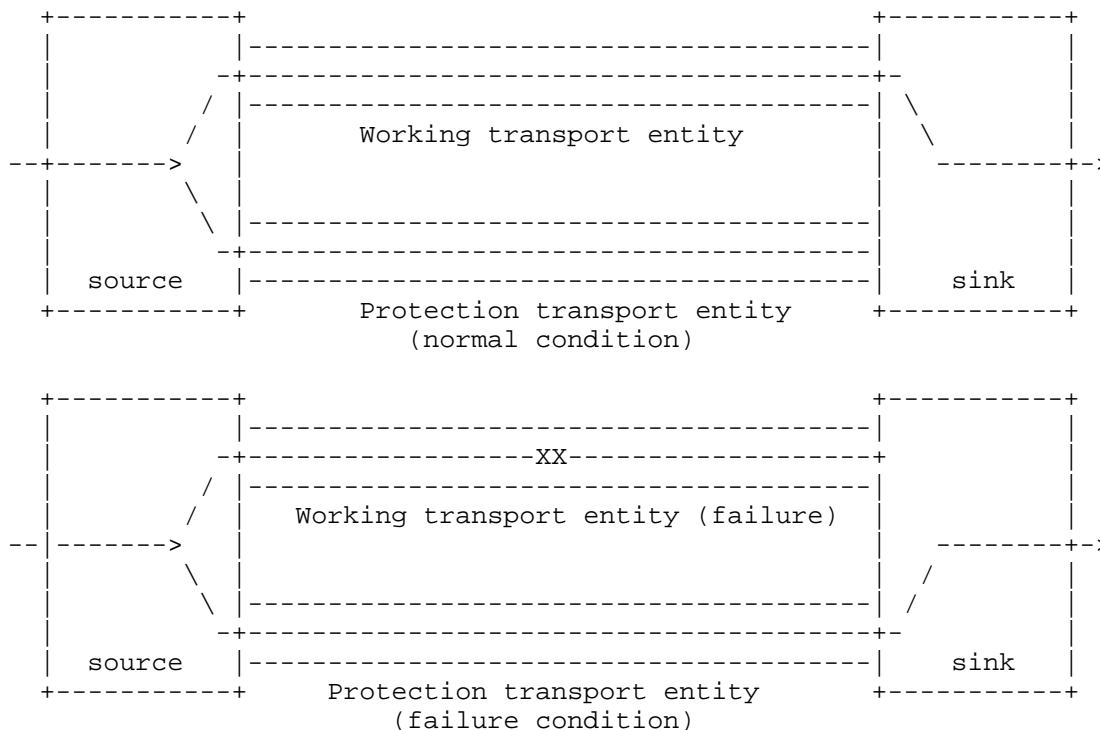


Figure 1: 1+1 Unidirectional Linear Protection Switching

1+1 unidirectional protection switching is the simplest protection switching mechanism. The normal traffic is permanently bridged on both the working and protection transport entities at the source endpoint of the protected domain. In the normal condition, the sink endpoint receives traffic from the working transport entity. If the sink endpoint detects a failure on the working transport entity, it will switch to receive traffic from the protection transport entity. 1+1 unidirectional protection switching is recommended to be used for unidirectional transport.

Note that 1+1 unidirectional protection switching does not use the APS coordination protocol since it only performs protection switching based on the local request.

6.2. 1+1 Bidirectional Protection Switching

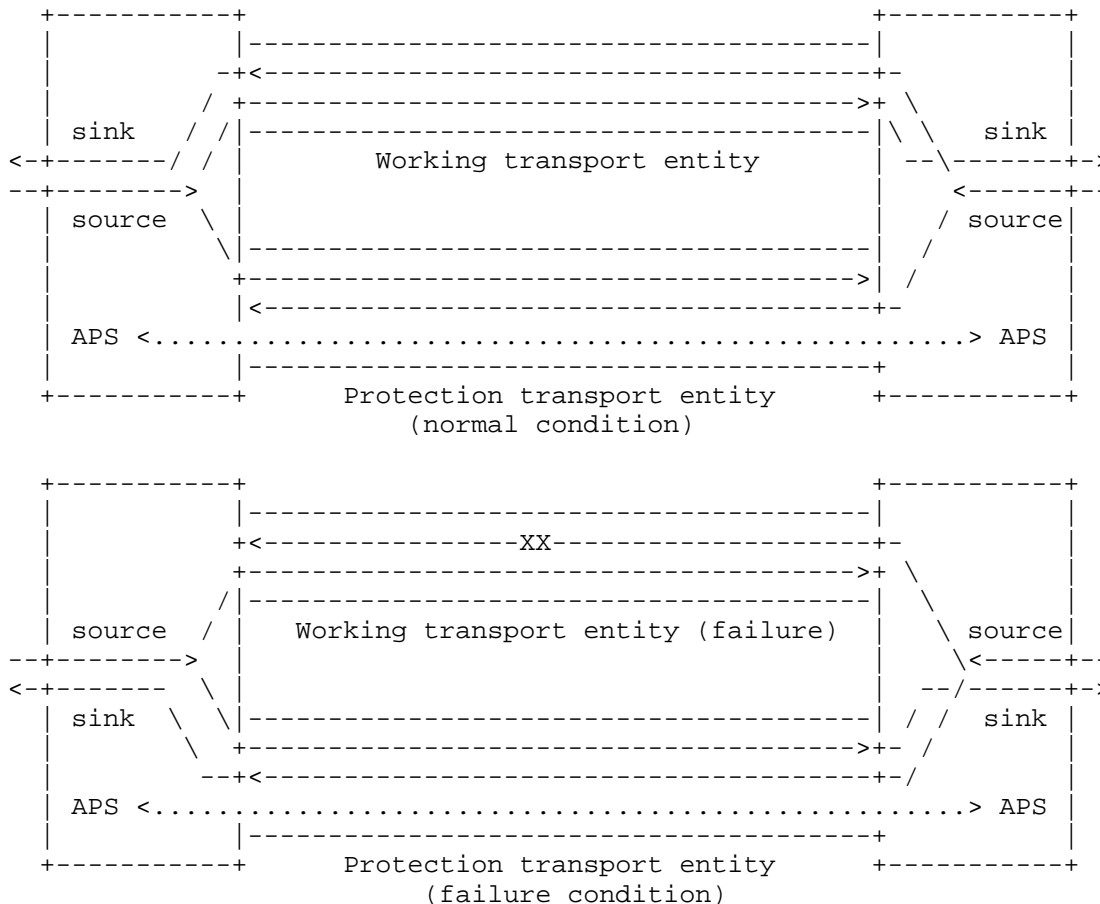


Figure 2: 1+1 Bidirectional Linear Protection Switching

In 1+1 bidirectional protection switching, for each direction, the normal traffic is permanently bridged on both the working and protection transport entities at the source endpoint of the protected domain. In the normal condition, for each direction, the sink endpoint receives traffic from the working transport entity.

If the sink endpoint detects a failure on the working transport entity, it will switch to receive traffic from the protection transport entity. It will also send an APS message to inform the sink endpoint on the other direction to switch to receive traffic from the protection transport entity.

The APS mechanism is necessary to coordinate the two endpoints of the transport entity and to implement 1+1 bidirectional protection switching even for a unidirectional failure.

6.3. 1:1 Bidirectional Protection Switching

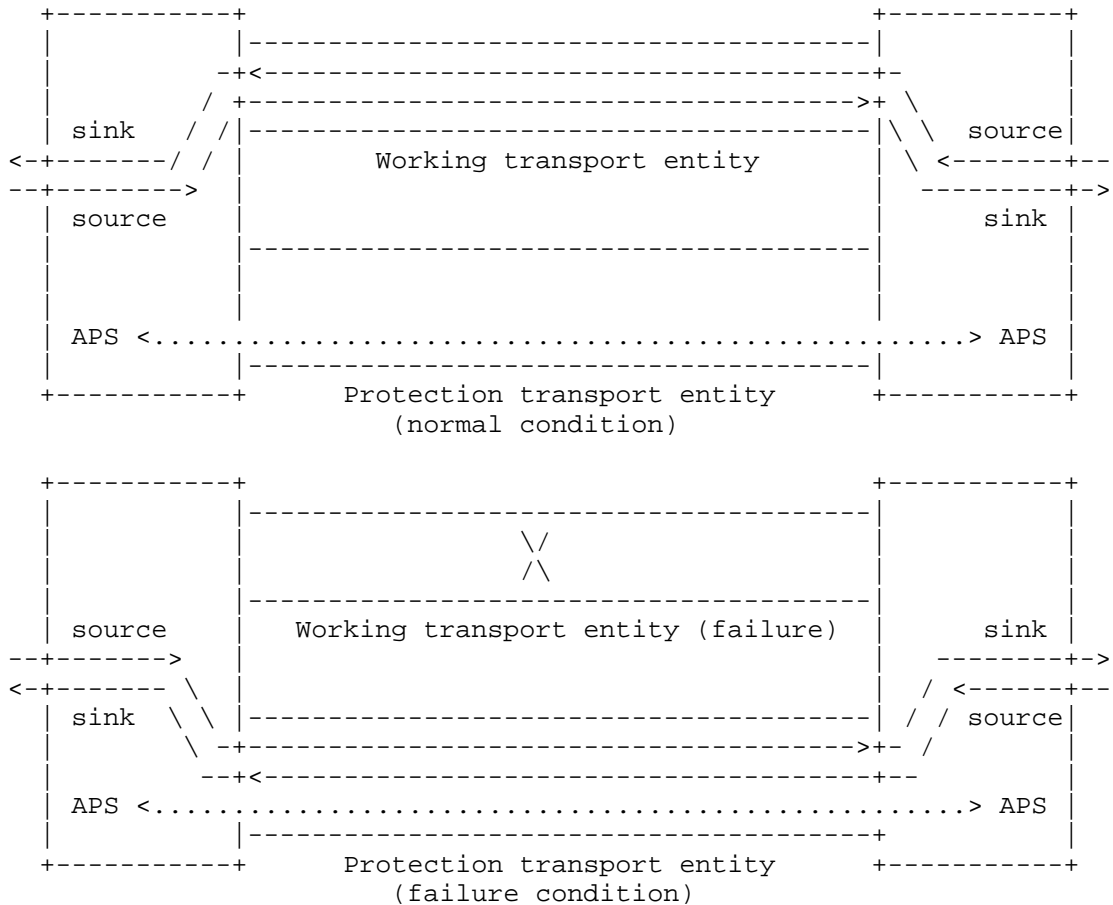


Figure 3: 1:1 Bidirectional Linear Protection Switching

In 1:1 bidirectional protection switching, for each direction, the source endpoint sends traffic on either the working transport entity or the protection transport entity. The sink endpoint receives the traffic from the same transport entity on which the source endpoint sends the traffic.

In the normal condition, for each direction, the source and sink endpoints send and receive traffic from the working transport entity.

If the sink endpoint detects a failure on the working transport entity, it will switch to send and receive traffic from the protection transport entity. It will also send an APS message to inform the sink endpoint on another direction to switch to send and receive traffic from the protection transport entity.

The APS mechanism is necessary to coordinate the two endpoints of the transport entity and implement 1:1 bidirectional protection switching even for a unidirectional failure.

7. APS Protocol

This APS protocol is based upon the APS protocol defined in Section 11 of [G.8031]. See that reference for further definition of the Protocol Data Unit (PDU) fields and protocol details beyond the description in this document.

7.1. APS PDU Format

APS packets MUST be sent over a Generic Associated Channel (G-ACh) as defined in [RFC5586].

The format of APS PDU is specified in Figure 4 below.

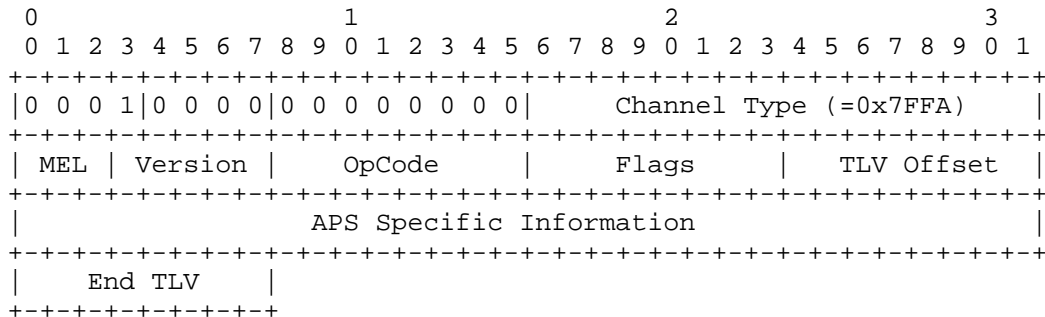


Figure 4: APS PDU Format

The following values MUST be used for APS PDU:

- o Channel Type: The Channel Type MUST be configurable by the implementation. During deployment, the local system administrator provisioned the value 0x7FFA. This is a code point value in the range of experimental Channel Types as described in RFC 5586, Section 10.

- o Maintenance Entity group Level (MEL): The MEL value to set and check MUST be configurable. The DEFAULT value MUST be "111". With co-routed bidirectional transport paths, the configured MEL MUST be the same in both directions.
- o Version: 0x00
- o OpCode: 0x27 (=0d39)
- o Flags: 0x00
- o TLV Offset: 4
- o End TLV: 0x00

The format of the APS-specific information is defined in Figure 5.




Figure 5: APS-Specific Information Format

All bits defined as "Reserved" MUST be transmitted as 0 and ignored on reception.

- o Request/State:
The four bits indicate the protection-switching request type. See Figure 6 for the code of each request/state type.

In case that there are multiple protection-switching requests, only the protection-switching request with the highest priority MUST be processed.

Request/State	Code/Priority
Lockout of Protection (LO)	1111 (highest)
Signal Fail on Protection (SF-P)	1110
Forced Switch (FS)	1101
Signal Fail on Working (SF-W)	1011
Signal Degrade (SD)	1001
Manual Switch (MS)	0111
Wait to Restore (WTR)	0101
Exercise (EXER)	0100
Reverse Request (RR)	0010
Do Not Revert (DNR)	0001
No Request (NR)	0000 (lowest)

Figure 6: Protection-Switching Request Code/Priority

o Protection Type (Pr.Type):

The four bits are used to specify the protection type.

A: reserved (set by default to 1)
 B: 0 - 1+1 (permanent bridge)
 1 - 1:1 (no permanent bridge)
 D: 0 - Unidirectional switching
 1 - Bidirectional switching
 R: 0 - Non-revertive operation
 1 - Revertive operation

- o Requested Signal:

This byte is used to indicate the traffic that the near-end requests to be carried over the protection entity.

value = 0: Null traffic
value = 1: Normal traffic 1
value = 2~255: Reserved

- o Bridged Signal:

This byte is used to indicate the traffic that is bridged onto the protection entity.

value = 0: Null traffic
value = 1: Normal traffic 1
value = 2~255: Reserved

- o Bridge Type (T):

This bit is used to further specify the type of non-permanent bridge for 1:1 protection switching.

value = 0: Selector bridge
value = 1: Broadcast bridge

- o Reserved:

This field MUST be set to zero.

7.2. APS Transmission

The APS message MUST be transported on the protection transport entity by encapsulation with the protection transport entity label (the label of the LSP used to transport protection traffic). If an endpoint receives APS-specific information from the working transport entity, it MUST ignore this information and MUST report the failure of protocol defect (see Section 8.1) to the operator.

A new APS packet MUST be transmitted immediately when a change in the transmitted status occurs. The first three APS packets MUST be transmitted as fast as possible only if the APS information to be transmitted has been changed so that fast protection switching is possible, even if one or two APS packets are lost or corrupted. The interval of the first three APS packets SHOULD be 3.3 ms. APS packets after the first three MUST be transmitted with the interval of 5 seconds.

If no valid APS-specific information is received, the last valid received information remains applicable.

7.3. Hold-Off Timer

In order to coordinate timing of protection switches at multiple layers, a hold-off timer MAY be required. The purpose is to allow a server-layer protection switch to have a chance to fix the problem before switching at a client layer.

Each selector SHOULD have a provisioned hold-off timer. The suggested range of the hold-off timer is 0 to 10 seconds in steps of 100 ms (accuracy of +/-5 ms).

When a new defect or more severe defect occurs (new SF or SD) on the active transport entity (the transport entity that currently carries and selects traffic), this event will not be reported immediately to protection switching if the provisioned hold-off timer value is non-zero. Instead, the hold-off timer SHALL be started. When the hold-off timer expires, it SHALL be checked whether a defect still exists on the transport entity that started the timer. If it does, that defect SHALL be reported to protection switching. The defect need not be the same one that started the timer.

This hold-off timer mechanism SHALL be applied for both working and protection transport entities.

7.4. WTR Timer

In revertive mode of operation, to prevent frequent operation of the protection switch due to an intermittent defect, a failed working transport entity MUST become fault free. After the failed working transport entity meets this criterion, a fixed period of time SHALL elapse before a normal traffic signal uses it again. This period, called a WTR period, MAY be configured by the operator in 1 minute steps between 5 and 12 minutes; the default value is 5 minutes. An SF or SD condition will override the WTR. To activate the WTR timer appropriately, even when both ends concurrently detect clearance of SF-W and SD-W, when the local state transits from SF-W or SD-W to No Request (NR) with the requested signal number 1, the previous local state, SF-W or SD-W, MUST be memorized. If both the local state and far-end state are NR with the requested signal number 1, the local state transits to WTR only when the previous local state is SF-W or SD-W. Otherwise, the local state transits to NR with the requested signal number 0.

In revertive mode of operation, when the protection is no longer requested, i.e., the failed working transport entity is no longer in SF or SD condition (and assuming no other requesting transport entities), a local WTR state will be activated. Since this state becomes the highest in priority, it is indicated on the APS signal and maintains the normal traffic signal from the previously failed working transport entity on the protection transport entity. This state SHALL normally time out and become an NR state. The WTR timer deactivates earlier when any request of higher priority request preempts this state.

7.5. Command Acceptance and Retention

The commands Clear, LO, FS, MS, and EXER are accepted or rejected in the context of previous commands, the condition of the working and protection entities in the protection group, and (in bidirectional switching only) the APS information received.

The Clear command MUST be only valid if a near-end LO, FS, MS, or EXER command is in effect or if a WTR state is present at the near end and rejected otherwise. This command will remove the near-end command or WTR state, allowing the next lower-priority condition or (in bidirectional switching) APS request to be asserted.

Other commands MUST be rejected unless they are higher priority than the previously existing command, condition, or (in bidirectional switching) APS request. If a new command is accepted, any previous, lower-priority command that is overridden MUST be forgotten. If a higher priority command overrides a lower-priority condition or (in bidirectional switching) APS request, that other request will be reasserted if it still exists at the time the command is cleared. If a command is overridden by a condition or (in bidirectional switching) APS request, that command MUST be forgotten.

7.6. Exercise Operation

Exercise is a command to test if the APS communication is operating correctly. It is lower priority than any "real" switch request. It is only valid in bidirectional switching, since this is the only place where you can get a meaningful test by looking for a response.

The Exercise command SHALL issue the command with the same requested and bridged signal numbers of the NR, Reverse Request (RR), or DNR request that it replaces. The valid response will be an RR with the corresponding requested and bridged signal numbers. When Exercise commands are input at both ends, an EXER, instead of RR, MUST be transmitted from both ends. The standard response to DNR MUST be DNR rather than NR. When the exercise command is cleared, it MUST be

replaced with NR or RR if the requested signal number is 0 and DNR or RR if the requested signal number is 1.

8. Protection-Switching Logic

8.1. Principle of Operation

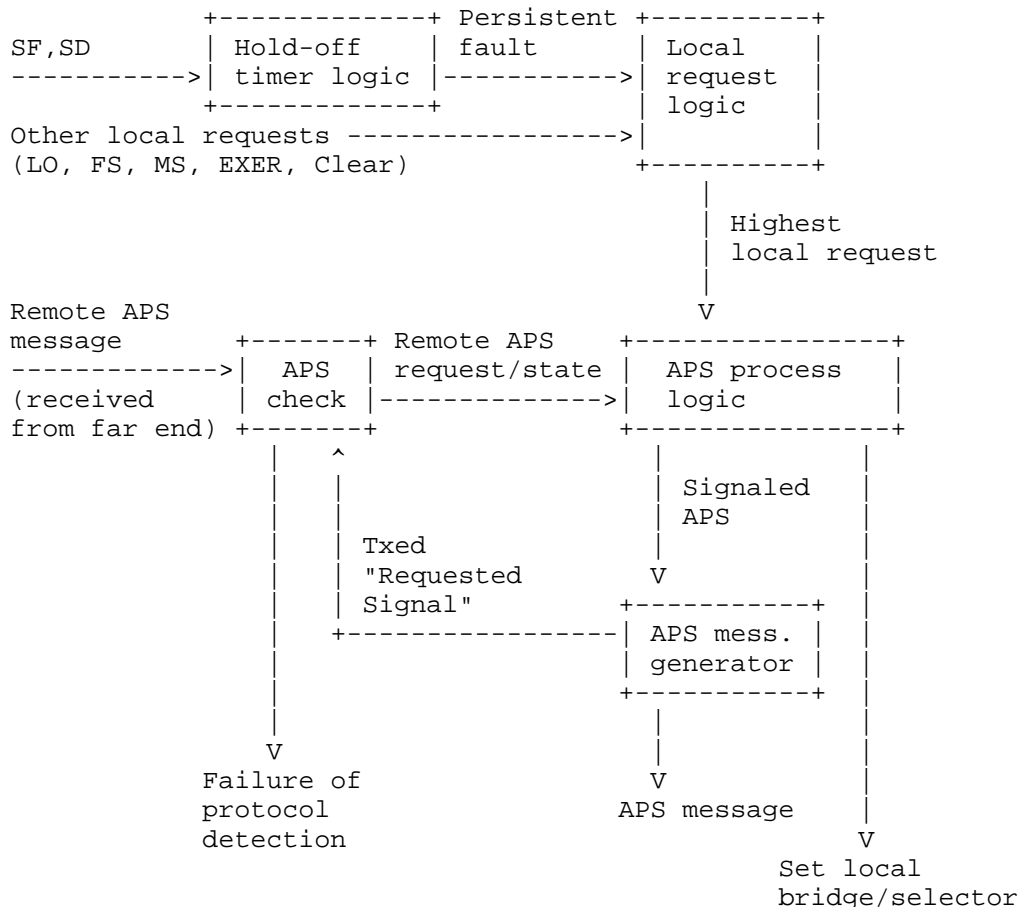


Figure 7: Protection-Switching Logic

Figure 7 describes the protection-switching logic.

One or more local protection-switching requests may be active. The "local request logic" determines which of these requests is highest using the order of priority given in Figure 6. This highest local request information SHALL be passed on to the "APS process logic". Note that an accepted Clear command, clearance of SF or SD, or

expiration of the WTR timer SHALL NOT be processed by the local request logic but SHALL be considered as the highest local request and submitted to the APS process logic for processing.

The remote APS message is received from the far end and is subjected to the validity check and mismatch detection in "APS check". Failure of protocol situations are as follows:

- o The "B" field mismatch due to incompatible provisioning;
- o The reception of the APS message from the working entity due to working/protection configuration mismatch;
- o No match in sent "Requested Signal" and received "Requested Signal" for more than 50 ms;
- o No APS message is received on the protection transport entity during at least 3.5 times the long APS interval (e.g., at least 17.5 seconds), and there is no defect on the protection transport entity.

Provided the "B" field matches:

- o If the "D" bit mismatches, the bidirectional side will fall back to unidirectional switching.
- o If the "R" bit mismatches, one side will clear switches to WTR and the other will clear to DNR. The two sides will interwork and the traffic is protected.
- o If the "T" bit mismatches, the side using a broadcast bridge will fall back to using a selector bridge.

The APS message with invalid information MUST be ignored, and the last valid received information remains applicable.

The linear protection-switching algorithm SHALL commence immediately every time one of the input signals changes, i.e., when the status of any local request changes, or when different APS-specific information is received from the far end. The consequent actions of the algorithm are also initiated immediately, i.e., change the local bridge/selector position (if necessary), transmit new APS-specific information (if necessary), or detect the failure of protocol defect if the protection switching is not completed within 50 ms.

The state transition is calculated in the "APS process logic" based on the highest local request, the request of the last received "Request/State" information, and state transition tables defined in Section 9, as follows:

- o If the highest local request is Clear, clearance of SF or SD, or expiration of WTR, a state transition is calculated first based on the highest local request and state machine table for local requests to obtain an intermediate state. This intermediate state is the final state in case of clearance of SF-P; otherwise, starting at this intermediate state, the last received far-end request and the state machine table for far-end requests are used to calculate the final state.
- o If the highest local request is neither Clear nor clearance of SF or of SD nor expiration of WTR, the APS process logic compares the highest local request with the request of the last received "Request/State" information based on Figure 6.
 1. If the highest local request has higher or equal priority, it is used with the state transition table for local requests defined in Section 9 to determine the final state; otherwise,
 2. The request of the last received "Request/State" information is used with the state transition table for far-end requests defined in Section 9 to determine the final state.

The "APS message generator" generates APS-specific information with the signaled APS information for the final state from the state transition calculation (with coding as described in Figure 5).

8.2. Equal Priority Requests

In general, once a switch has been completed due to a request, it will not be overridden by another request of the same priority (first-come, first-served policy). Equal priority requests from both sides of a bidirectional protection group are both considered valid, as follows:

- o If the local state is NR, with the requested signal number 1, and the far-end state is NR, with the requested signal number 0, the local state transits to NR with the requested signal number 0. This applies to the case when the remote request for switching to the protection transport entity has been cleared.

- o If both the local and far-end states are NR, with the requested signal number 1, the local state transits to the appropriate new state (DNR state for non-revertive mode and WTR state for revertive mode). This applies to the case when the old request has been cleared at both ends.
- o If both the local and far-end states are RR, with the same requested signal number, both ends transit to the appropriate new state according to the requested signal number. This applies to the case of concurrent deactivation of EXER from both ends.
- o In other cases, no state transition occurs, even if equal priority requests are activated from both ends. Note that if MSs are issued simultaneously to both working and protection transport entities, either as local or far-end requests, the MS to the working transport entity is considered as having higher priority than the MS to the protection transport entity.

8.3. Signal Degrade of the Protection Transport Entity

Signal degrade on the protection transport entity has the same priority as signal degrade on the working transport entity. As a result, if an SD condition affects both transport entities, the first SD detected MUST NOT be overridden by the second SD detected. If the SD is detected simultaneously, either as local or far-end requests on both working and protection transport entities, then the SD on the standby transport entity MUST be considered as having higher priority than the SD on the active transport entity, and the normal traffic signal continues to be selected from the active transport entity (i.e., no unnecessary protection switching is performed).

In the preceding sentence, "simultaneously" relates to the occurrence of SD on both the active and standby transport entities at input to the protection-switching process at the same time, or as long as an SD request has not been acknowledged by the remote end in bidirectional protection switching.

9. Protection-Switching State Transition Tables

In this section, state transition tables for the following protection switching configurations are described.

- o 1:1 bidirectional (revertive mode, non-revertive mode);
- o 1+1 bidirectional (revertive mode, non-revertive mode);
- o 1+1 unidirectional (revertive mode, non-revertive mode).

Note that any other global or local request that is not described in state transition tables does not trigger any state transition.

The states specified in the state transition tables can be described as follows:

- o NR: NR is the state entered by the local priority under all conditions where no local protection-switching requests (including WTR and DNR) are active. NR can also indicate that the highest local request is overridden by the far-end request, whose priority is higher than the highest local request. Normal traffic signal is selected from the corresponding transport entity.
- o LO, SF-P, SD-P: The access by the normal traffic to the protection transport entity is NOT allowed in this state. The normal traffic is carried by the working transport entity, regardless of the fault/degrade condition possibly present (due to the highest priority of the switching triggers leading to this state).
- o FS, SF-W, SD-W, MS-W, MS-P: A switching trigger NOT resulting in the protection transport entity unavailability is present. The normal traffic is selected either from the corresponding working transport entity or from the protection transport entity, according to the behavior of the specific switching trigger.
- o WTR: In revertive operation, after the clearing of an SF-W or SD-W, this maintains normal traffic as selected from the protection transport entity until the WTR timer expires or another request with higher priority, including the Clear command, is received. This is used to prevent frequent operation of the selector in the case of intermittent failures.
- o DNR: In non-revertive operation, this is used to maintain a normal traffic to be selected from the protection transport entity.
- o EXER: Exercise of the APS protocol.
- o RR: The near end will enter and signal Reverse Request only in response to an EXER from the far end.

[State transition tables are shown at the end of the PDF form of this document.]

10. Security Considerations

MPLS-TP is a subset of MPLS and so builds upon many of the aspects of the security model of MPLS. MPLS networks make the assumption that it is very hard to inject traffic into a network and equally hard to cause traffic to be directed outside the network. The control-plane protocols utilize hop-by-hop security and assume a "chain-of-trust" model such that end-to-end control-plane security is not used. For more information on the generic aspects of MPLS security, see [RFC5920].

This document describes a protocol carried in the G-ACh [RFC5586] and so is dependent on the security of the G-ACh, itself. The G-ACh is a generalization of the associated channel defined in [RFC4385]. Thus, this document relies heavily on the security mechanisms provided for the associated channel and described in those two documents.

11. Acknowledgements

The authors would like to thank Hao Long, Vincenzo Sestito, Italo Busi, Igor Umansky, and Andy Malis for their input to and review of the current document.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [RFC5586] Bocci, M., Vigoureux, M., and S. Bryant, "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [G.841] International Telecommunications Union, "Types and characteristics of SDH network protection architectures", ITU-T Recommendation G.841, October 1998.
- [G.873.1] International Telecommunications Union, "Optical Transport Network (OTN): Linear protection", ITU-T Recommendation G.873.1, May 2014.

[G.8031] International Telecommunications Union, "Ethernet linear protection switching", ITU-T Recommendation G.8031/Y.1342, June 2011.

[T1.105.01] American National Standards Institute, "Synchronous Optical Network (SONET) - Automatic Protection Switching", ANSI 0900105.01:2000 (R2010), March 2000.

12.2. Informative References

[RFC6378] Weingarten, Y., Bryant, S., Osborne, E., Sprecher, N., and A. Fulignoli, "MPLS Transport Profile (MPLS-TP) Linear Protection", RFC 6378, October 2011.

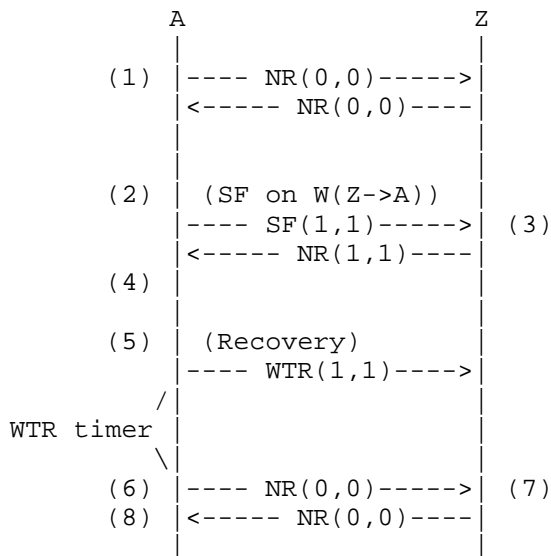
[RFC7271] Ryoo, J., Gray, E., van Helvoort, H., D'Alessandro, A., Cheung, T., and E. Osborne, "MPLS Transport Profile (MPLS-TP) Linear Protection to Match the Operational Expectations of Synchronous Digital Hierarchy, Optical Transport Network, and Ethernet Transport Network Operators", RFC 7271, June 2014.

[RFC7324] Osborne, E., "Updates to MPLS Transport Profile Linear Protection", RFC 7324, July 2014.

Appendix A. Operation Examples of the APS Protocol

The sequence diagrams shown in this section are only a few examples of the APS operations. The first APS message, which differs from the previous APS message, is shown. The operation of hold-off timer is omitted. The fields whose values are changed during APS packet exchange are shown in the APS packet exchange. They are Request/State, requested traffic, and bridged traffic. For an example, SF(0,1) represents an APS packet with the following field values: Request/State = SF, Requested Signal = 0, and Bridged Signal = 1. The values of the other fields remain unchanged from the initial configuration. The signal numbers 0 and 1 refer to null signal and normal traffic signal, respectively. W(A->Z) and P(A->Z) indicate the working and protection paths in the direction of A to Z, respectively.

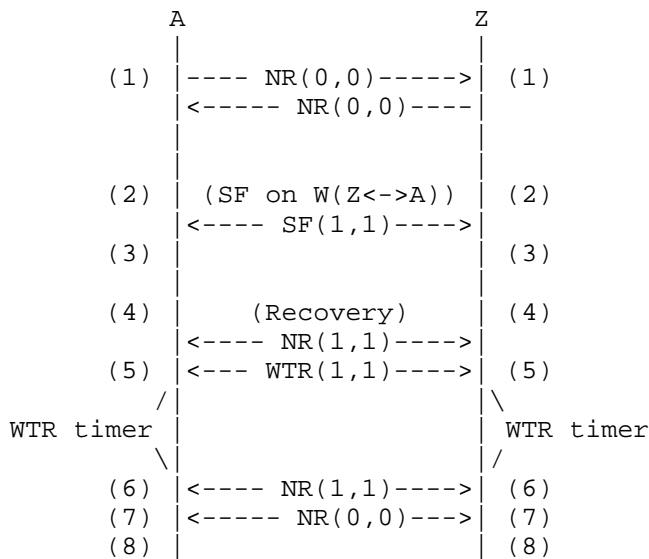
Example 1. 1:1 bidirectional protection switching (revertive mode) - Unidirectional SF case



- (1) The protected domain is operating without any defect, and the working entity is used for delivering the normal traffic.
- (2) Signal Fail occurs on the working entity in the Z to A direction. Selector and bridge of node A select protection entity. Node A generates an SF(1,1) message.

- (3) Upon receiving SF(1,1), node Z sets selector and bridge to protection entity. As there is no local request in node Z, node Z generates an NR(1,1) message.
- (4) Node A confirms that the far end is also selecting protection entity.
- (5) Node A detects clearing of the SF condition, starts the WTR timer, and sends a WTR(1,1) message.
- (6) At expiration of the WTR timer, node A sets selector and bridge to working entity and sends an NR(0,0) message.
- (7) Node Z is notified that the far-end request has been cleared and sets selector and bridge to working entity.
- (8) It is confirmed that the far end is also selecting working entity.

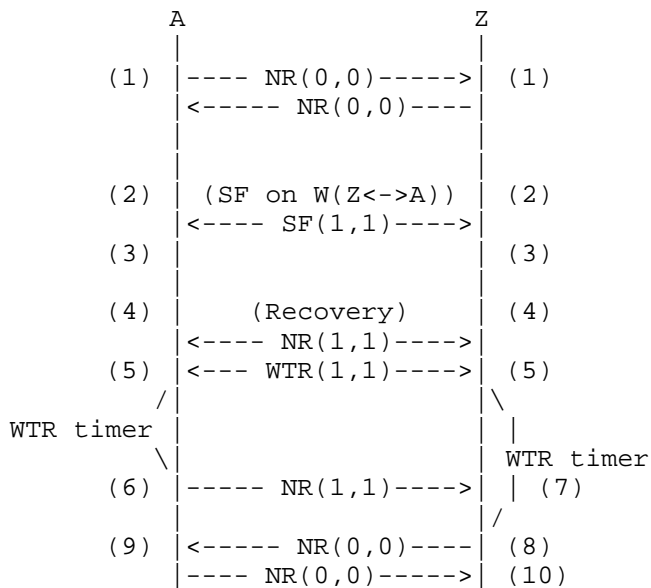
Example 2. 1:1 bidirectional protection switching (revertive mode) - Bidirectional SF case



- (1) The protected domain is operating without any defect, and the working entity is used for delivering the normal traffic.
- (2) Nodes A and Z detect local SF conditions on the working entity, set selector and bridge to protection entity, and generate SF(1,1) messages.

- (3) Upon receiving SF(1,1), each node confirms that the far end is also selecting protection entity.
- (4) Each node detects clearing of the SF condition and sends an NR(1,1) message as the last received APS message was SF.
- (5) Upon receiving NR(1,1), each node starts the WTR timer and sends WTR(1,1).
- (6) At expiration of the WTR timer, each node sends NR(1,1) as the last received APS message was WTR.
- (7) Upon receiving NR(1,1), each node sets selector and bridge to working entity and sends an NR(0,0) message.
- (8) It is confirmed that the far end is also selecting working entity.

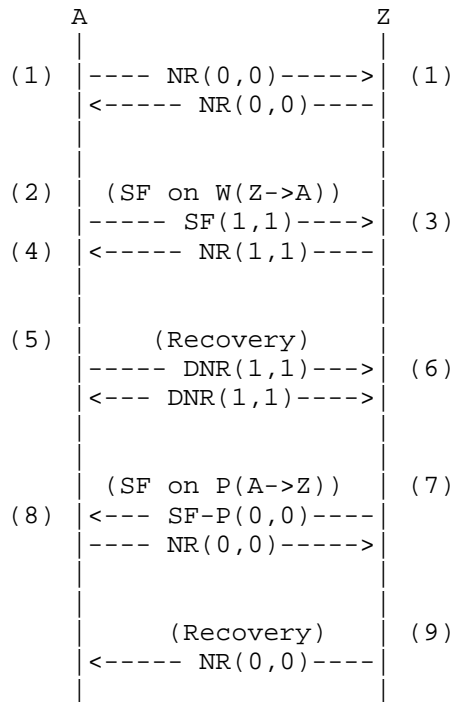
Example 3. 1:1 bidirectional protection switching (revertive mode) - Bidirectional SF case - Inconsistent WTR timers



- (1) The protected domain is operating without any defect, and the working entity is used for delivering the normal traffic.
- (2) Nodes A and Z detect local SF conditions on the working entity, set selector and bridge to protection entity, and generate SF(1,1) messages.

- (3) Upon receiving SF(1,1), each node confirms that the far end is also selecting protection entity.
- (4) Each node detects clearing of the SF condition and sends an NR(1,1) message as the last received APS message was SF.
- (5) Upon receiving NR(1,1), each node starts the WTR timer and sends WTR(1,1).
- (6) At expiration of the WTR timer in node A, node A sends an NR(1,1) message as the last received APS message was WTR.
- (7) At node Z, the received NR(1,1) is ignored as the local WTR has a higher priority.
- (8) At expiration of the WTR timer in node Z, node Z sets selector and bridge to working entity and sends an NR(0,0) message.
- (9) Upon receiving NR(0,0), node A sets selector and bridge to working entity and sends an NR(0,0) message.
- (10) It is confirmed that the far end is also selecting working entity.

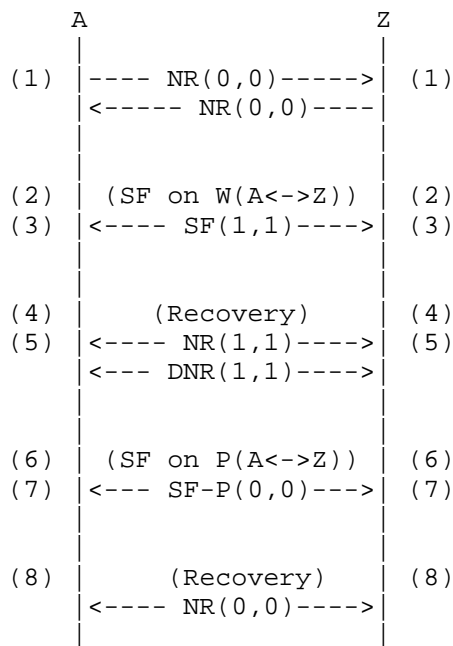
Example 4. 1:1 bidirectional protection switching (non-revertive mode) - Unidirectional SF on working followed by unidirectional SF on protection



- (1) The protected domain is operating without any defect, and the working entity is used for delivering the normal traffic.
- (2) Signal Fail occurs on the working entity in the Z to A direction. Selector and bridge of node A select the protection entity. Node A generates an SF(1,1) message.
- (3) Upon receiving SF(1,1), node Z sets selector and bridge to protection entity. As there is no local request in node Z, node Z generates an NR(1,1) message.
- (4) Node A confirms that the far end is also selecting protection entity.
- (5) Node A detects clearing of the SF condition and sends a DNR(1,1) message.
- (6) Upon receiving DNR(1,1), node Z also generates a DNR(1,1) message.
- (7) Signal Fail occurs on the protection entity in the A to Z direction. Selector and bridge of node Z select the working entity. Node Z generates an SF-P(0,0) message.

- (8) Upon receiving SF-P(0,0), node A sets selector and bridge to working entity and generates an NR(0,0) message.
- (9) Node Z detects clearing of the SF condition and sends an NR(0,0) message.

Exmample 5. 1:1 bidirectional protection switching (non-revertive mode) - Bidirectional SF on working followed by bidirectional SF on protection



- (1) The protected domain is operating without any defect, and the working entity is used for delivering the normal traffic.
- (2) Nodes A and Z detect local SF conditions on the working entity, set selector and bridge to protection entity, and generate SF(1,1) messages.
- (3) Upon receiving SF(1,1), each node confirms that the far end is also selecting protection entity.
- (4) Each node detects clearing of the SF condition and sends an NR(1,1) message as the last received APS message was SF.
- (5) Upon receiving NR(1,1), each node sends DNR(1,1).

- (6) Signal Fail occurs on the protection entity in both directions. Selector and bridge of each node selects the working entity. Each node generates an SF-P(0,0) message.
- (7) Upon receiving SF-P(0,0), each node confirms that the far end is also selecting working entity.
- (8) Each node detects clearing of the SF condition and sends an NR(0,0) message.

Authors' Addresses

Huub van Helvoort (editor)
Huawei Technologies

EEmail: huub@van-helvoort.eu

Jeong-dong Ryoo (editor)
ETRI

EEmail: ryoo@etri.re.kr

Haiyan Zhang
Huawei Technologies

EEmail: zhanghaiyan@huawei.com

Feng Huang
Philips

EEmail: feng.huang@philips.com

Han Li
China Mobile

EEmail: lihan@chinamobile.com

Alessandro D'Alessandro
Telecom Italia

EEmail: alessandro.dalessandro@telecomitalia.it