INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1720

(09/2003)

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT GENERATION NETWORKS

Internet protocol aspects – Operation, administration and maintenance

Protection switching for MPLS networks

ITU-T Recommendation Y.1720

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT GENERATION NETWORKS

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	**Y.1700–Y.1799**
Charging	Y.1800–Y.1899
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation Y.1720

Protection switching for MPLS networks

Summary

This Recommendation provides requirements and mechanisms for 1+1, 1:1, shared mesh, and packet 1+1 protection switching functionality for the user-plane in MPLS networks. The mechanism defined herein is designed to support end-to-end point-to-point LSPs. Protection switching functionality for multipoint-to-point and point-to-multipoint LSP are for further study. m:n protection switching is for further study. Hitless protection switching is outside the scope of this version of the Recommendation.

Source

ITU-T Recommendation Y.1720 was approved by ITU-T Study Group 13 (2001-2004) under the ITU-T Recommendation A.8 procedure on 13 September 2003.

Keywords

defect, failure, LSP, MPLS, PML, protection switching, PSL, rerouting

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

ITU-T Recommendation Y.1720

Protection switching for MPLS networks

1 Scope

This Recommendation provides requirements and mechanisms for 1+1, 1:1, shared mesh, and packet 1+1 protection switching functionality for the user-plane in MPLS networks. The mechanism defined herein is designed to support end-to-end point-to-point LSPs. Protection switching functionality for multipoint-to-point and point-to-multipoint LSP are for further study. m:n protection switching is for further study. Hitless protection switching is outside the scope of this version of the Recommendation.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[1] ITU-T Recommendation Y.1710 (2002), *Requirements for Operation and Maintenance functionality for MPLS networks.*

[2] ITU-T Recommendation Y.1711 (2002), *Operation and Maintenance mechanism for MPLS networks.*

[3] ITU-T Recommendation G.805 (2000), *Generic functional architecture of transport networks.*

NOTE – There is a limitation of the applicability of the architecture specified by Recommendation G.805. It is not applicable to LDP based multipoint-to-point LSP and the case where PHP is in effect with the egress not supporting MPLS data plane.

[4] ITU-T Recommendation G.841 (1998), *Types and characteristics of SDH network protection architectures.*

[5] ITU-T Recommendation I.630 (1999), *ATM protection switching.*

[6] ITU-T Recommendation M.495 (1988), *Transmission restoration and transmission route diversity: Terminology and general principles.*

[7] ITU-T Recommendation M.20 (1992), *Maintenance philosophy for telecommunication networks.*

[8] IETF RFC 3031 (2001), *Multiprotocol Label Switching Architecture*, *Category: Standards Track.*

[9] IETF RFC 3032 (2001), *MPLS Label Stack Encoding, Category: Standards Track.*

3 Definitions

3.1 1+1 protection: A protection mechanism in which the traffic is duplicated on the protection path (constantly bridged). The path merging LSR performs the switching of the traffic between the working and protection path.

3.2 1:1 protection: A protection mechanism in which the traffic is sent only on the working path or the protection path. The path switching LSR performs the switching of the traffic between the working and protection path.

3.3 shared mesh protection: Shared mesh protection switching can be viewed as an extension of 1:1 protection. It provides sharing of bandwidth between protection LSPs corresponding to working LSPs that belong to disjoint links, nodes, or SRGs.

3.4 Shared Risk Group (SRG): SRG is a group of links or nodes that can fail simultaneously due to a single failure incident. For example, fibres within a conduit belong to an SRG, because a single conduit breakage can cut all fibres passing through it.

3.5 packet 1+1 protection: Like 1+1 protection the traffic is fed onto both LSPs. Packet level 1+1 allows selection of incoming packet from any of the two LSPs irrespective of the LSP from which the last packet was selected. That is, packet 1+1 protection treats both LSPs as working as opposed to designating one LSP as working while the other as the protection LSP.

3.6 bidirectional protection switching: A protection switching architecture in which, for a unidirectional failure, both directions of the LSP, including the affected direction and the unaffected direction, are switched to protection.

3.7 bridge: The action or function of transmitting identical traffic on both the working and protection LSP.

3.8 defect: Interruption of the capability of an LSP to transfer user or OAM information (see Note 1).

3.9 extra traffic: Traffic that is purposely placed on the same network layer resource as a protection LSP (but in a separate LSP which is parallel to protection LSP) in the knowledge that, on failure, this (extra) traffic will be disconnected to make way for the protected traffic from the failed working connection.

3.10 failure: Termination of the capability of an LSP to transfer user or OAM information. A failure can be caused by a persisting defect (see Note 1).

3.11 forced switch for working LSP: A switch action initiated by an operator command. Switch action is conducted unless a higher priority switch request (i.e., LoP) is in effect.

3.12 hold-off time: The time between declaration of signal degrade or signal fail, and the initialization of the protection switching algorithm.

3.13 manual switch: A switch action initiated by an operator command. Switch action is conducted unless an equal or a higher priority switch request (i.e., LoP, FS, SF or MS) is in effect.

3.14 MPLS protection domain: The set of LSRs over which a working path and its corresponding protection path are routed.

3.15 non-revertive protection switching: A protection switching method where revertive action (switch back to the working LSP) is not taken after the working LSP is repaired.

3.16 no request: A state where no protection switching request exists.

3.17 path switch LSR: An LSR that is responsible for switching or replicating the traffic between the working LSP and the protection LSP.

3.18 path merge LSR: An LSR that is responsible for receiving the protection path traffic, and either merges the traffic back onto the working path, or, if it is itself the destination, passes the traffic on to the higher layer protocols.

3.19 protection LSP: The LSP within the protection domain from which working traffic is received at the sink of the protection domain where a working LSP has failed.

3.20 protection switching: A recovery mechanism in which the protection LSP or path segments are created prior to the detection of a fault on the working path. In other words, a protection mechanism in which the protection LSP is pre-calculated, its capacity is pre-assigned and the protection LSP is pre-established.

3.21 rerouting: A recovery mechanism in which the recovery path or path segments are created dynamically after the detection of a fault on the working path. In other words, a recovery mechanism in which the recovery path is not pre-established.

3.22 revertive protection switching: A protection switching method where revertive action (switch back to the working LSP) is taken after the working LSP is repaired.

3.23 selector: A switch which selects to receive the traffic from the working LSP or the protection LSP at the sink of the protection domain, or a switch which selects to send the traffic to the working LSP or the protection LSP at the source of the protection domain.

3.24 source of the protection domain: A transmitting endpoint (ingress) in a path switch LSR of the protection domain.

3.25 sink of the protection domain: A receiving endpoint (egress) in a path merge LSR of the protection domain.

3.26 transport entity: An architectural component which transfers information between its inputs and outputs within a layer network (see Note 2). An LSP is used as a transport entity in an MPLS network.

3.27 unidirectional protection switching: A protection switching architecture in which, for a unidirectional failure (i.e. a failure affecting only one direction of transmission), only the affected direction of the LSP is switched to protection.

3.28 wait to restore: An automatically initiated command that is issued when the working LSP exits SF condition. It is used to maintain the state until the wait to restore timer expires unless it is pre-empted by a higher priority bridge request.

3.29 wait to restore timer: A configurable timer which is used to delay before reversion.

3.30 working LSP: The LSP within the protection domain from which working traffic is received at the sink of the protection domain under fault-free condition in revertive mode.

NOTE 1 – ITU-T Rec. M.20 gives a more general and detailed definition.

NOTE 2 – ITU-T Rec. G.805 gives a more general and detailed definition.

4 Symbols and abbreviations

This Recommendation uses the following abbreviations:

APS	Automatic Protection Switching
BDI	Backward Defect Indication
CV Packet	Connectivity Verification Packet
FDI	Forward Defect Indication
FFD Packet	Fast Failure Detection Packet
FS	Forced Switch
LDP	Label Distribution Protocol
LOCV	Loss of Connectivity Verification
LoP	Lockout of Protection
LSP	Label Switched Path

LSR	Label Switch Router
MPLS	Multiprotocol Label Switching
MS	Manual Switch
OAM	Operation, Administration and Maintenance
PHP	Penultimate Hop Popping
PML	Path Merge LSR
PS	Protection Switching
PSL	Path Switch LSR
SDH	Synchronous Digital Hierarchy
SF	Signal Fail
SLA	Service Level Agreement
TTSI	Trail Termination Source Identifier

5 Requirements

Techniques to enhance reliability performance of a network by providing a capability to recover from service interruption (e.g., due to defects) are referred to as survivability techniques. Survivability techniques include protection switching and rerouting. This Recommendation is developed to specify protection switching techniques. In this Recommendation the difference between protection switching and rerouting is intended to mean the following:

- Protection switching: This implies that both routing and resources are pre-calculated and allocated to a dedicated protection LSP prior to failure. Protection-switching therefore offers a strong assurance of being able to re-obtain the required network resources post-failure.
- Rerouting: This implies that a dedicated protection LSP is not defined, and so neither routing nor resources are pre-calculated/allocated prior to failure. Rerouting is commonly used to refer to cases where there are routing and signalling functions in operation, and that when a "re-connection request" has to be instigated on failure (either by the network, or by the customer), that this "reconnect request" has to contend with other similar traffic types for obtaining the required resource. Rerouting, therefore, offers no assurance of being able to re-obtain the required network resources post-failure and is generally slower than protection switching.

Protection switching is necessary for fast recovery from failure, and thereby enhances the reliability and availability performance of MPLS networks. For protection switching, the following features are required:

1) Protection switching should be applied to an entire LSP.
2) Prioritized protection between Signal Fail (SF) and operator switch requests (see Table 1).
3) The possibility to achieve protection at the MPLS layer as fast as possible (subject to the temporal resolution of the defect detection mechanism) should be provided.
4) Protection ratio of 100%, i.e., 100% of impaired working traffic is protected for a failure on a single working LSP.
5) An extra traffic capability should be supported when possible.

6 Principles

Protection switching is a fully allocated protection mechanism that can be used on any topology. It is fully allocated in the sense that the route and bandwidth of the protection LSP is reserved for a selected working LSP. To be effective under all possible failures of the working LSP however, the protection LSP must be known to have complete physical diversity over all common-failure modes. This may not always be possible. Also, this might require the working LSP not to follow its shortest path.

The MPLS PS architecture can be a 1+1 type, a 1:1 type, a shared mesh type, or a packet 1+1 type. Other types are for further study.

In the 1+1 architecture type, a protection LSP is dedicated to each working LSP with the working LSP bridged onto the protection LSP at the source of the protection domain. The traffic on working and protection LSPs is transmitted simultaneously to the sink of the protection domain, where a selection between the working and protection LSP is made, based on some predetermined criteria, such as defect indication.

In the 1:1 architecture type, a protection LSP is dedicated to each working LSP. The working traffic is transmitted either by working or protection LSP. The method for a selection between the working and protection LSPs depends on the mechanism. The protection LSP can be used to carry "extra traffic" when it is not used to transmit the working traffic.

In the shared mesh architecture type, possible sharing of protection capacity between disjoint link, node, or SRG failures in the network is achieved while guaranteeing recovery from a single failure. For each link in the network, it keeps track of all the working paths whose traffic will be switched onto it after a given failure. By keeping track of this, it only needs to reserve maximum of the protection capacity required to protect a single failure in the network.

In the packet 1+1 architecture type, the traffic is transmitted simultaneously onto two possibly disjoint routed LSPs to the sink of the protection domain. Each pair of duplicate transmitted packets is assigned the same identifier (sequence number) but distinct from the other pairs of duplicate packets. At the sink of the protection domain packet level selection mechanism is employed to select one of the two possibly received copies of each packet. The following list provides principles for MPLS protection architectures and mechanisms development.

1) Defects in layers above MPLS should not cause server layer protection switching. For example, in case of ATM over MPLS, defects in ATM layer should not cause MPLS protection switching.
2) In general, if lower layer (e.g., SDH or optical) protection mechanisms are being utilized in conjunction with MPLS layer protection mechanisms, then the lower layers should have a chance to restore working traffic before the MPLS layer initiates protection actions (e.g., using a hold-off timer). The objective here is to avoid duplicated protection switching in different layer networks.
3) Protection switching actions in one protection domain should not adversely affect network operations, performance and protection switching in other domains.
4) The protection switching mechanism should facilitate fast recovery of working traffic to minimize the network outage, and ideally recovery should be before the unavailability entry threshold is reached.

7 Mechanisms

This clause describes mechanisms of unidirectional and bidirectional protection switching.

7.1 Unidirectional protection switching

7.1.1 Application architectures

7.1.1.1 Application architecture of unidirectional 1+1 protection switching

The 1+1 linear protection switching architecture is as shown in Figure 1. In the case of unidirectional protection switching operation as described here, protection switching is performed by the selector at the sink of the protection domain based on purely local (i.e., at protection sink) information. The working traffic is permanently bridged to working and protection LSPs at the source of the protection domain. If CV packets, FFD packets, or other continuity probe packets are used to detect defects of working or protection LSP, they are inserted at the source of the protection domain of both working and protection side and detected and extracted at the sink of the protection domain. It is noted that they should be sent regardless of whether the LSP is selected by the selector or not.

For example, if a unidirectional defect (in the direction of transmission from PSL to PML) occurs for the working LSP as in Figure 2, this defect will be detected at the sink of the protection domain at PML and the selector at PML will switch to the protection LSP.

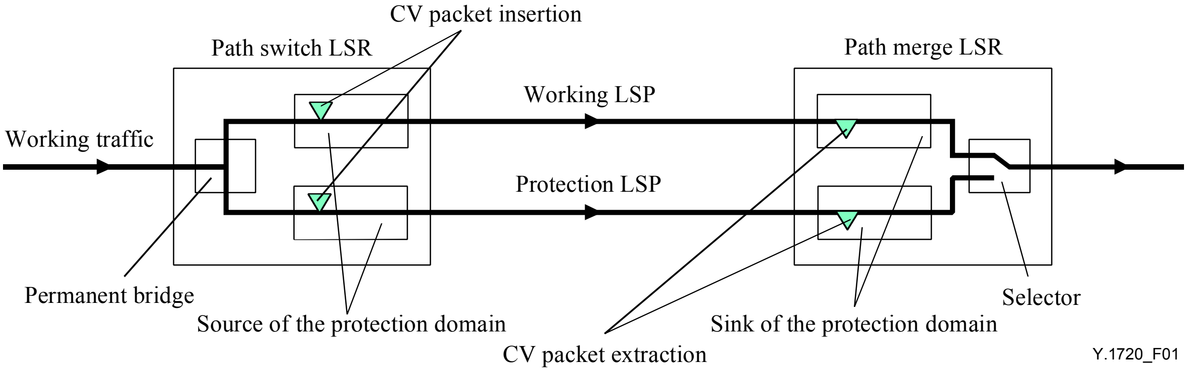

Figure 1/Y.1720 – Unidirectional 1+1 protection switching architecture

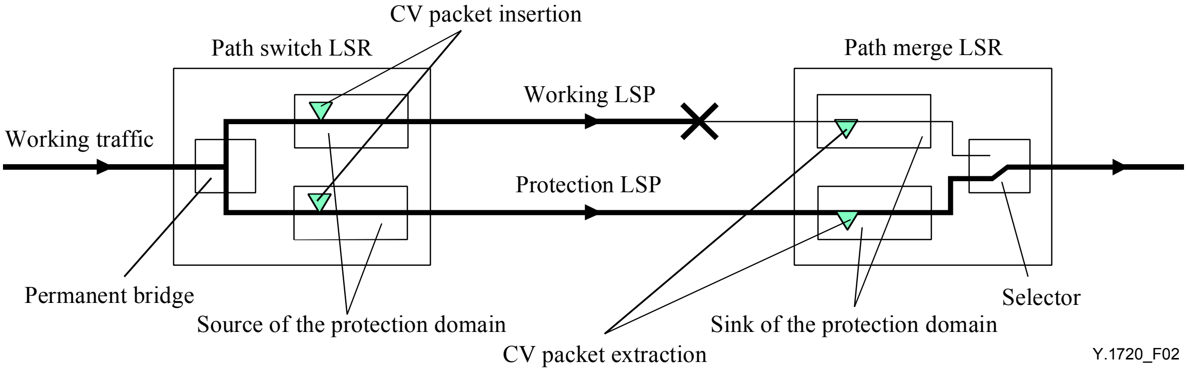

Figure 2/Y.1720 – Unidirectional 1+1 protection switching architecture – working LSP fails

7.1.1.2 Application architecture of unidirectional 1:1 protection switching

The 1:1 linear protection switching architecture is as shown in Figure 3. In the case of unidirectional protection switching operation as described here, protection switching is performed by the selector at the source of the protection domain based on purely local (i.e., at protection source) information. The working and protection traffic is permanently merged at the sink of the protection domain.

If CV packets, FFD packets, or other continuity probe packets are used to detect defects of working or protection LSP, they are inserted at the source of the protection domain of both working and protection side and detected and extracted at the sink of the protection domain. It is noted that they should be sent regardless of whether the LSP is selected by the selector or not.

For example, if a unidirectional defect (in the direction of transmission from PSL to PML) occurs for the working LSP as in Figure 4, this defect is detected at the sink of the protection domain at PML and then reported by BDI to the source of the protection domain at PSL. The selector at PSL switches to the protection LSP on reception of this report.

NOTE – dTTSI_Mismerge cannot be protected by 1:1 protection switching.

When SF for working LSP is declared and user traffic is transmitted by protection LSP, FDI packet and user traffic may be merged at the sink of the protection domain. Nodes in downstream may receive FDI packets, CV or FFD packets and user traffic at the same time. The same applies in the case where SF for protection LSP is declared. One way to solve this problem is to use a merging selector. The operation of the merging selector when a defect occurs on the working LSP, is the following:

1) Receive FDI packets or detect a lower layer defect at the egress of the working LSP.
2) Switch the merging selector at the egress (i.e., open the switch on working LSP and close the switch on protection LSP).
3) Send BDI packets on working LSP.
4) Switch the selector at the ingress (i.e., working LSP to protection LSP and cut off the extra traffic).

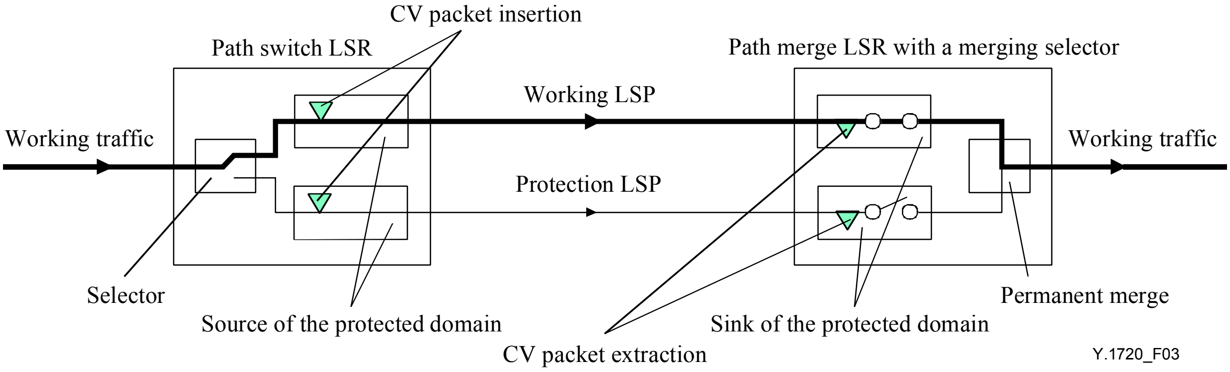

Figure 3/Y.1720 – Unidirectional 1:1 protection switching architecture

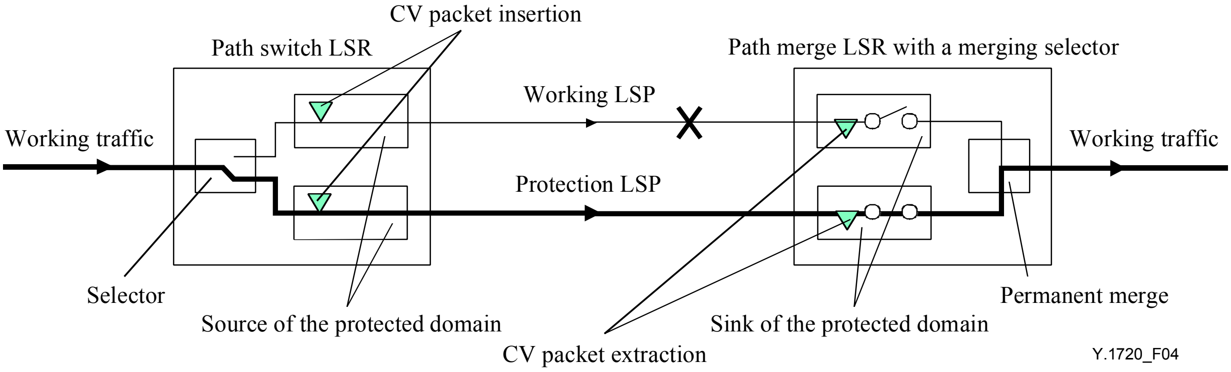

Figure 4/Y.1720 – Unidirectional 1:1 protection switching architecture – working LSP fails

Extra traffic

The 1:1 architecture can support extra traffic. As the traffic from the working and the protection LSPs is merged at the sink point of the protection domain, extra traffic must be transported via a separate LSP for which the physical route is the same as the protection LSP (see Figure 5) in order to avoid the extra traffic and the working traffic being merged, and to share the bandwidth between them. When the working traffic is switched over to the protection LSP, the extra traffic is disconnected to make way for the protected traffic from the failed working connection (see Figure 6). This generally requires a protection switching coordination protocol. In this Recommendation, BDI is used as the 1-phase protocol (see also ITU-T Rec. I.630). Connectivity verification of an extra traffic LSP is optional. In case notification of disconnection of extra traffic is required, connectivity verification should be used.

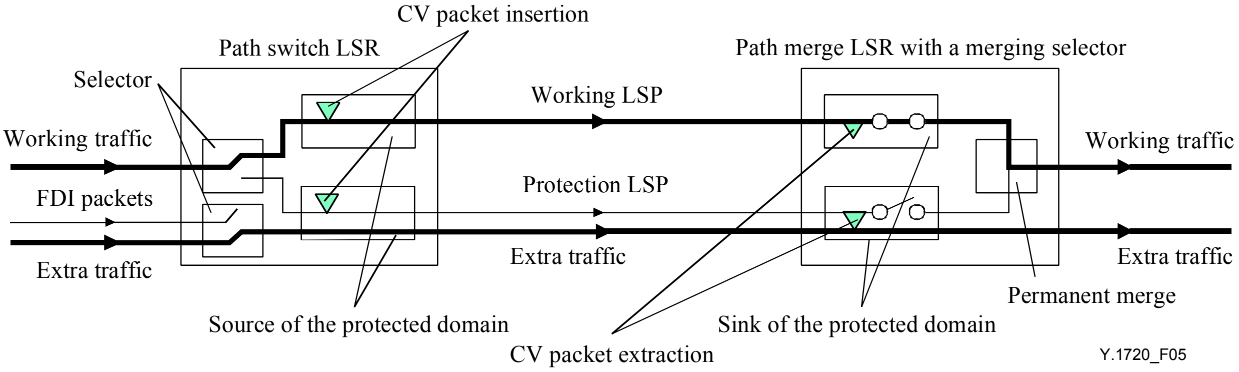

Figure 5/Y.1720 – 1:1 architecture with extra traffic

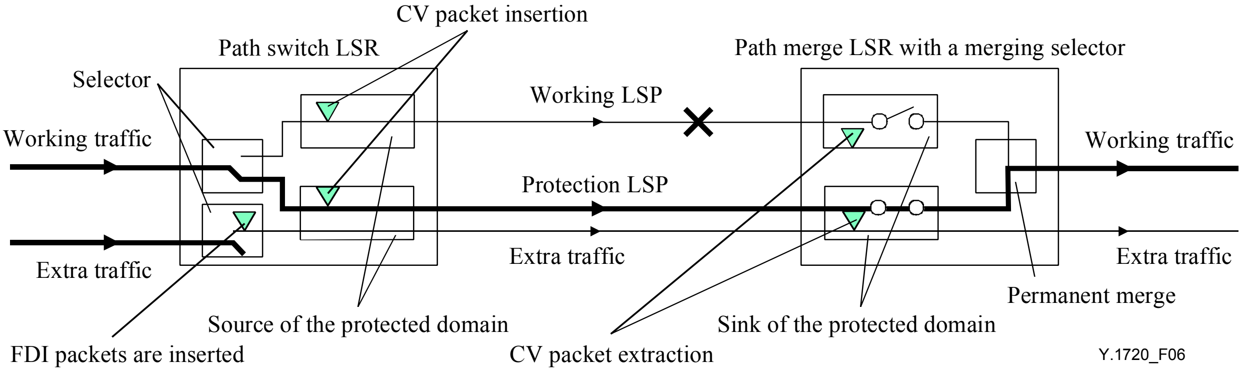

Figure 6/Y.1720 – 1:1 architecture with extra traffic – working LSP fails

7.1.1.3 Application architecture of unidirectional shared mesh protection switching

Shared mesh protection switching can be viewed as an extension of 1:1 protection. It requires all the functionality to realize 1:1 protection as well as additional functionality to achieve sharing among disjoint link, node, or Shared Risk Group (SRG) failures.

Functional requirements

Following are the necessary functional requirements to realize a real-time shared mesh protection scheme:

1) Shared mesh protection switching should have a capability to realize the possible sharing of protection capacity between disjoint link, node, or SRG failures in the network while guaranteeing recovery from a single failure.

2) Shared mesh protection switching should have a capability to reserve (set aside) the required capacity for protection on each link without allocating it to any LSPs.

3) Shared mesh protection switching should have a capability to detect (notify) the failure at (to) end nodes (ingress and egress).

4) Shared mesh protection switching should have a capability to allocate protection capacity to protection LSPs at the time of failure.

5) Shared mesh protection switching should have a capability to switch feeding of the traffic at the ingress and selection of traffic at the egress from working (protection) to protection (working) LSP.

6) Shared mesh protection switching should support recovery within bounded time constraints and may be compliant with generally used recovery times.

7) Shared mesh protection switching should allow efficient use of working LSP bandwidth using such measures as route optimization, taking into account route dependencies between a working path and its protection path.

Application architecture

A shared mesh protection scheme is targeted to provide guaranteed recovery while using a minimal amount of protection bandwidth in a general mesh topology. It sets aside a pool of dedicated protection capacity, sufficient to recover all the protected traffic from any single possible failure in the network. For each protected working connection, the protection capacity is allocated at the time of its activation. Arrival of a request to establish shared mesh protection service between two nodes prompts computation of a pair of disjoint paths between them with two necessary constraints. First, sufficient bandwidth should be available along the route of the working connection to accommodate the requesting traffic. Second, either already reserved protection bandwidth along the protection path must be sufficient to guarantee recovery from any single failure along the primary route, or the available bandwidth along the protection path must be enough to accommodate the additional bandwidth needed for protecting the new working connection. Note that sharing is achieved by always first trying to accommodate a new request with already allocated protection capacity. This can be achieved by keeping track, for every link in the network, of the required amount of capacity to recover from each node or link failure in the network. Note that, since it is well accepted and verified that the probability of multiple concurrent failures in most networks is small; the scheme has been described to ensure protection from any single failure in the network. Protection from multiple failures can be achieved through a straightforward extension.

The protection switching and the trigger mechanism for each affected working LSP is similar to 1:1 protection scheme. An example of shared mesh protection is illustrated in Figure 7. Figure 8 illustrates the protection of traffic after a specific failure.

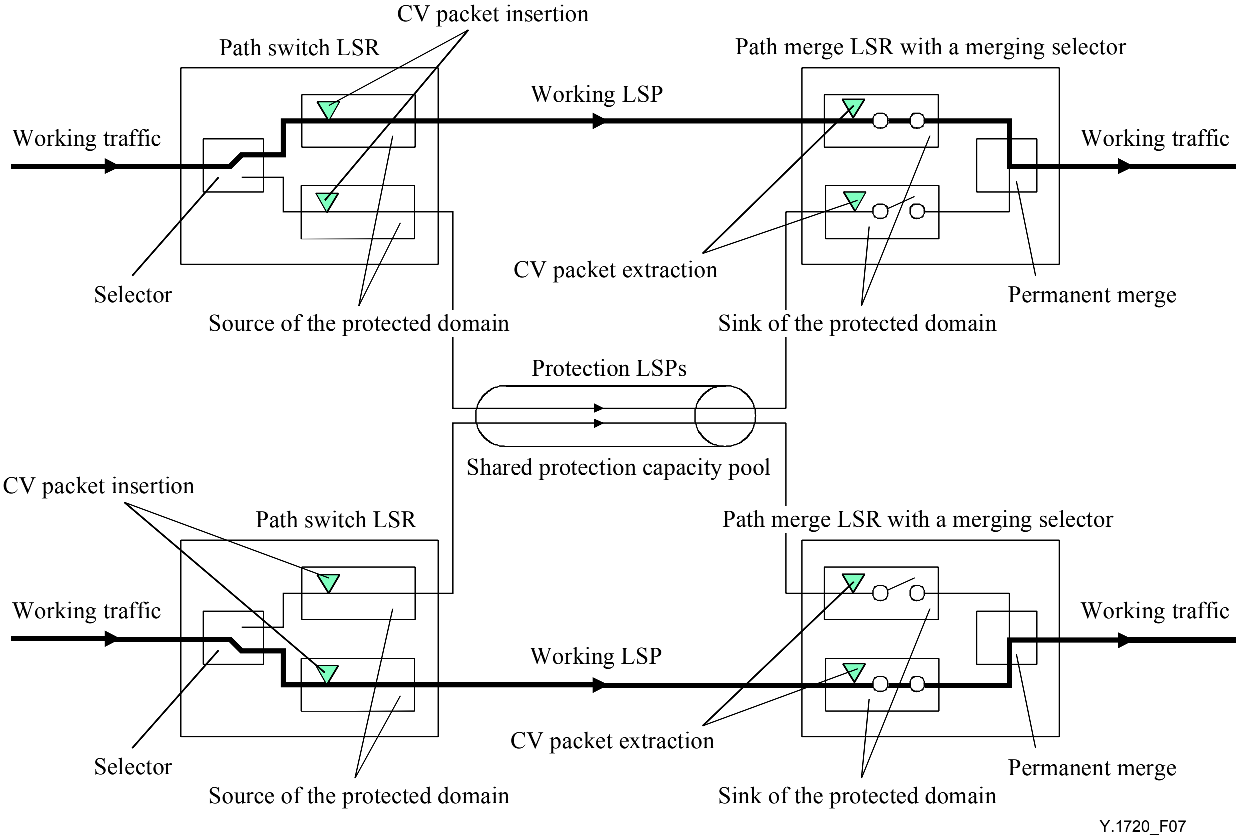

Figure 7/Y.1720 – Unidirectional shared mesh protection switching architecture

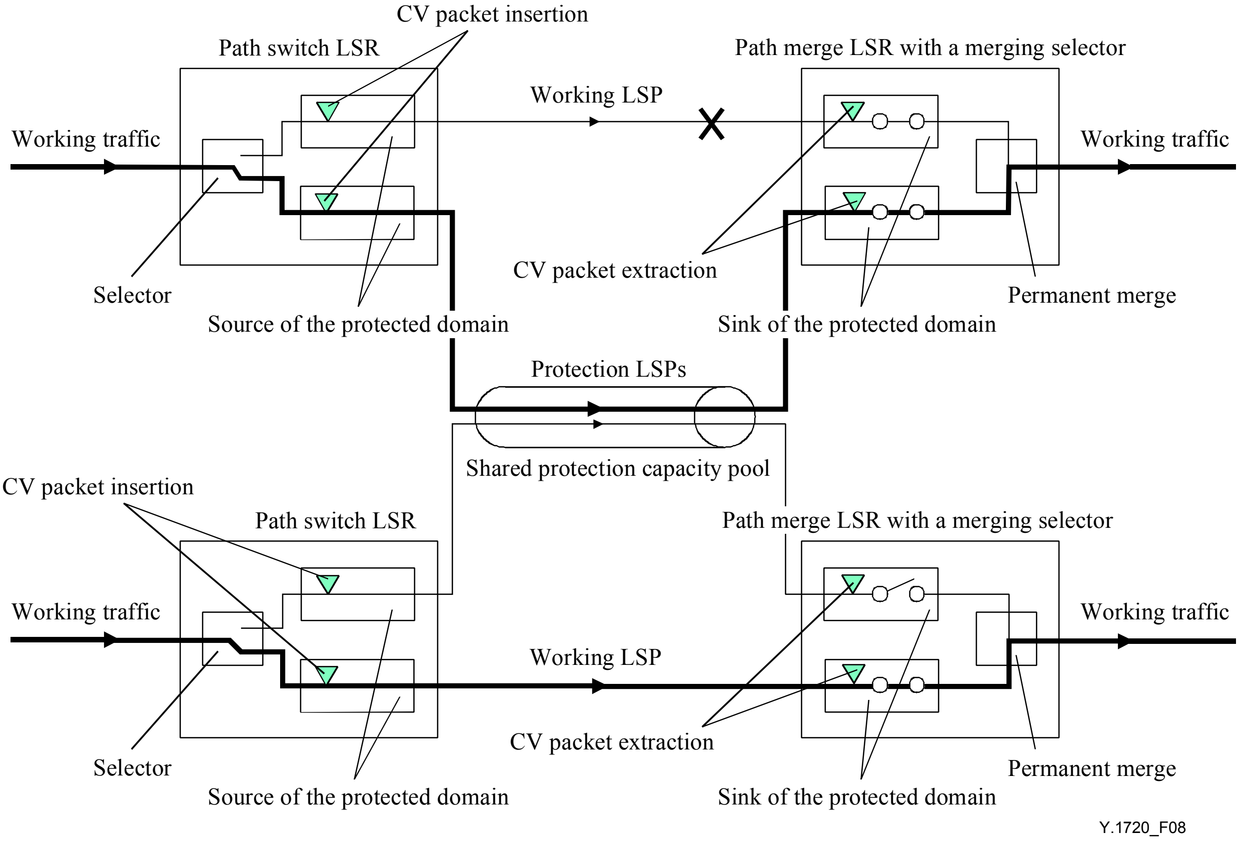

Figure 8/Y.1720 – Unidirectional shared mesh protection switching architecture – single failure

7.1.1.4 Application architecture of unidirectional packet 1+1 protection switching

Packet 1+1 path protection provides a packet level protection service similar in some respects to the conventional connection level 1+1 service with several important distinctions. Packet level 1+1 allows selection of incoming packet from any connection irrespective of the connection from which the last packet was selected. That is, packet 1+1 protection treats both connections as working connection as opposed to designating one connection as working while the other as the protection. In the latter, packets are selected from the working connection until a detection of failure on the working connection causes a switching to the protection connection. In contrast, packet 1+1 does not require explicit failure detection and protection switching. This allows the packet level 1+1 scheme to recover from any failure instantaneously and transparently. Similar to the connection level 1+1 protection, only edge nodes need to be service-aware.

To provide packet 1+1 protection service between two nodes in a MPLS network, a pair of LSPs is established between them along disjoint paths. Packets from a client flow subscribing to the service are dual-fed at the ingress node onto the two LSPs. Disjoint paths in the simplest case may be link or node disjoint but, in general, may involve a more complicated notion such as shared risk link groups. At the egress edge node, one of the two copies of the packets is selected and forwarded from the two possible received copies, each traversing a disjoint path. Given this, any single failure in the network, other than the ingress or egress node itself, can affect, at most, one copy of each packet. This allows the service to withstand a single failure transparently. In terms of restoration time, this can be characterized as an instantaneous recovery from a failure since there is no need to detect, notify and switch to a protection path explicitly.

Functional requirements

The minimum requirements to provide a packet 1+1 protection service are as follows:

a) There is no new requirement on the interior nodes of the network.

b) The network should support the establishment of diversely routed LSPs.

c) *Ingress Node*

   1) Must be able to associate the two LSPs that are used to provide packet level 1+1 protection between two end nodes.

   2) Must support the carrying of an identifier in the packet which will be used to identify duplicate copies of a packet at the egress node.

   3) Must be able to dual-feed each packet on the two mated LSPs.

d) *Egress Node*

   1) Must be able to associate the two LSPs that are used to provide packet level 1+1 protection between two end nodes.

   2) Must be able to identify the duplicate copies of a dual-fed packet using the identifier.

   3) Must be able to select and forward one and only one of the copies of a packet.

The above stated requirements describe the minimal functionality necessary to implement a packet level 1+1 protection scheme.

Reference model

Figure 9 illustrates a realization of the packet 1+1 protection scheme using sequence numbers as identifiers. After passing through the classifier, each packet that needs to be forwarded on the mated LSPs is assigned a distinct sequence number at the ingress node. This packet with the distinct identification is then duplicated and forwarded onto the two disjoint LSPs. On the egress node, a counter is used to keep track of the expected sequence number of the next packet. The details of an example implementation are described in the appendix.

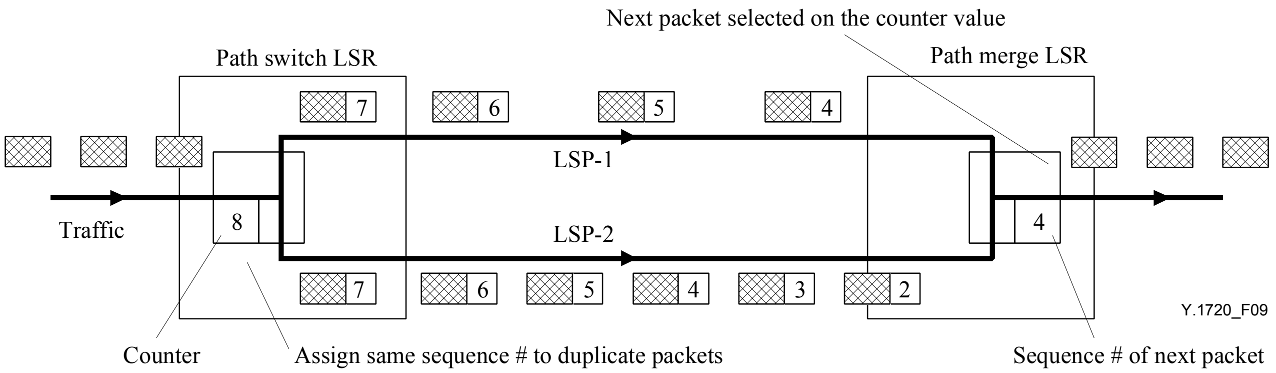
Next packet selected on the counter value
Path switch LSR
Path merge LSR
7
6
5
4
LSP-1
Traffic
8
4
LSP-2
7
6
5
4
3
2
Y.1720_F09
Counter
Assign same sequence # to duplicate packets
Sequence # of next packet

Figure 9/Y.1720 – Unidirectional packet 1+1 protection switching architecture

7.1.2 Protection switching trigger mechanism

Protection switching action should be conducted when:

1) initiated by operator control (e.g., manual switch, forced switch, and lockout of protection) without a higher priority switch request being in effect;

2) SF is declared on the connected LSP (i.e., working LSP or protection LSP) and is not declared on the other LSP and the hold-off timer has expired; or

3) the wait to restore timer expires (revertive mode) and SF is not declared on the working LSP.

7.1.2.1 Manual control

Manual control of the protection switching function may be transferred from the operation system.

7.1.2.2 Signal fail declaration conditions

7.1.2.2.1 1+1 architecture

For 1+1 architecture, Signal Fail (SF) is declared when the sink point of the protection domain enters the LSP Trail sink Near-End Defect State by entering the dServer, dLOCV, dTTSI_Mismatch, dTTSI_Mismerge, dExcess, or dUnknown condition.

In order to achieve fast protection (the requirement for fast protection is under study) SF can be declared when an FDI packet is received by the sink of the protection domain before it enters other defect conditions (e.g., dLOCV). It allows fast protection against the defects sourced from layers below the MPLS layer (and this requires that the incoming FDI have the DT codepoint 0x0101).

In addition, the FDD function can be used to achieve a faster declaration of the signal fail condition.

NOTE – It is only to be used if the lower layer is not protected. If the lower layer is also protected, it may lead to unnecessary protection switching by declaring SF on reception of FDI packets.

In the case where the CV or FFD function is not activated, SF is declared when an FDI packet is received by the sink of the protection domain. It only applies to the defects sourced from layers below the MPLS layer (and this requires that the incoming FDI have the DT codepoint 0x0101).

7.1.2.2.2 1:1 architecture

For 1:1 architecture, Signal Fail (SF) is declared when:

- the source of the protection domain enters the Trail sink Far-End Defect State by receiving a BDI packet (from the return LSP or out of band).

7.1.2.2.3 Shared mesh architecture

Shared mesh architecture is an extension of 1:1 architecture. Signal Fail (SF) is declared as in 1:1 architecture.

NOTE – Protection against a bidirectional LSP defect is for further study.

7.1.3 Compliance with network objectives

The following network objectives apply:

1) *Operating modes*

   Revertive and non-revertive switching are provided.

2) *Manual control*

   Operator control via Lockout of Protection, Forced Switch and Manual Switch commands are supported.

3) *Other switch initiation criteria*

   Signal Fail, Wait to Restore, and No Request are supported in addition to the manual control commands listed above, as criteria for initiating (or preventing) a protection switch.

7.1.4 Switch initiation criteria

The following switch initiation criteria exist:

1) an externally initiated command (Clear, Lockout of Protection, Forced Switch, Manual Switch);
2) an automatically initiated command (Signal Fail) associated with a protection domain; or
3) a state (Wait to Restore, No Request) of the protection switching function.

All requests are local (i.e., protection sink for 1+1 architecture and protection source for 1:1 architecture). The priority of local requests is given in Table 1.

Table 1/Y.1720 – Priority of local requests

Local request (i.e. automatically initiated command, state, or externally initiated command)	Order of priority
Clear	Highest
Lockout of Protection	|
Forced Switch	|
Signal Fail	|
Manual Switch	|
Wait To Restore	|
No Request	Lowest

NOTE 1 – A forced switch for the working LSP should not be overridden by a Signal Fail on the protection LSP. Since unidirectional protection switching is being performed, and no APS protocol is supported over the protection LSP, Signal Fail on the protection LSP does not interfere with the ability to perform a forced switch for the working LSP.

NOTE 2 – A forced switch for the protection LSP is not defined because this function may be achieved via a lockout of the protection command.

7.1.4.1 Externally initiated commands

Externally initiated commands are listed below in descending order of priority. The functionality of each is described below.

clear: This command clears all of the externally initiated switch commands listed below.

Lockout of Protection (LoP): Fix the selector position on the working LSP. Prevents the selector from switching to the protection LSP when it is selecting the working LSP. Switches the selector from the protection to the working LSP when it is selecting the protection LSP.

Forced Switch (FS) for working LSP: Switches the selector from the working LSP to the protection LSP (unless a higher priority switch request (i.e., LoP) is in effect).

Manual Switch (MS) for working LSP: Switches the selector from the working LSP to the protection LSP (unless an equal or higher priority switch request (i.e., LoP, FS, SF or MS) is in effect).

Manual Switch (MS) for protection LSP: Switches the selector from the protection LSP to the working LSP (unless an equal or higher priority switch request (i.e., LoP, FS, SF or MS) is in effect).

7.1.4.2 FDI triggered protection switch

In the case of FDI triggered protection switching, if the LSP with SF never enters a near end defect state, there may be a need to prevent frequent transitions. If so, some time may be defined that must pass before taking another protection switching action. This is FFS.

7.1.4.3 States

Wait to Restore is only applicable for the revertive mode and applies to a working LSP. This state is entered by the local protection switching function in conditions where working traffic is being received via the protection LSP when the working LSP is restored, if local protection switching requests have been previously active and now become inactive. It prevents reversion back to select the working LSP until the Wait to Restore timer has expired. The Wait to Restore time may be configured by the operator in 1 minute steps between 1 and 30 minutes; the default value is 12 minutes.

No Request is the state entered by the local protection switching function under all conditions where no local protection switching requests (including Wait to Restore) are active.

7.1.5 Protection switching protocol

In the unidirectional 1+1, 1:1, and shared mesh protection switching architecture, there is no need for APS protocol.

7.1.6 Unidirectional protection switching algorithm operation

7.1.6.1 Control of the selector

In the 1+1 1:1, and shared mesh architecture in the unidirectional protection switching operation, the selector is controlled by the highest priority local (i.e., sink of the protection domain for 1+1 architecture; source of the protection domain for 1:1 architecture) request (automatically initiated command, state, or externally initiated command). Therefore, each end operates independently of the other. If a condition of equal priority (e.g., SF) exists on both LSPs, switching shall not be performed.

In packet 1+1, the packets are selected based on a packet level selector which uses identifiers (sequence numbers) carried within transmitted packets.

7.1.6.2 Revertive mode

In revertive mode of operation, under conditions where working traffic is being transmitted via the protection LSP and when the working LSP is restored, if local protection switching requests have been previously active and now become inactive, a local Wait to Restore state is entered.

This state normally times out and becomes a No Request state after the Wait to Restore timer has expired. Then, reversion back to select the working LSP occurs. The Wait to Restore timer deactivates earlier if any local request of higher priority pre-empts this state.

7.1.6.3 Non-revertive mode

When the failed LSP is no longer in an SF condition, and no other externally initiated commands are present, a No Request state is entered. During this state, switching does not occur.

7.2 Mechanisms of bidirectional protection switching

For further study.

8 Security aspects

This Recommendation does not raise any security issues that are not already present in either the MPLS architecture or in the architecture of its client layer protocols.

Protection switching could enhance the security of MPLS networks as it will automatically switch traffic from defected LSPs, that may have misbranched or misconfigured into other LSPs, onto properly working LSPs. This will prevent customers' traffic being exposed to other customers.

Appendix I

Example of protection capacity sharing for shared mesh protection switching

Guaranteed recovery, with a similar grade of service, requires that the proper amount of bandwidth be set-aside in the network for restoration purposes. A shared mesh scheme requires that this protection bandwidth be sufficient to accommodate all traffic affected by any single failure in the network. This can be achieved by computing and reserving protection capacity at the time of activation of the working LSP. A request to establish the shared mesh service between two nodes prompts computation of a pair of disjoint paths between them with two necessary requirements. First, sufficient bandwidth should be available along the primary route of the LSP to accommodate the requested bandwidth. Second, either an already reserved protection bandwidth along the protection path is sufficient to guarantee recovery of the LSP from any single failure along the primary route, or available bandwidth along the protection path must be enough to accommodate additional bandwidth needed to protect it.

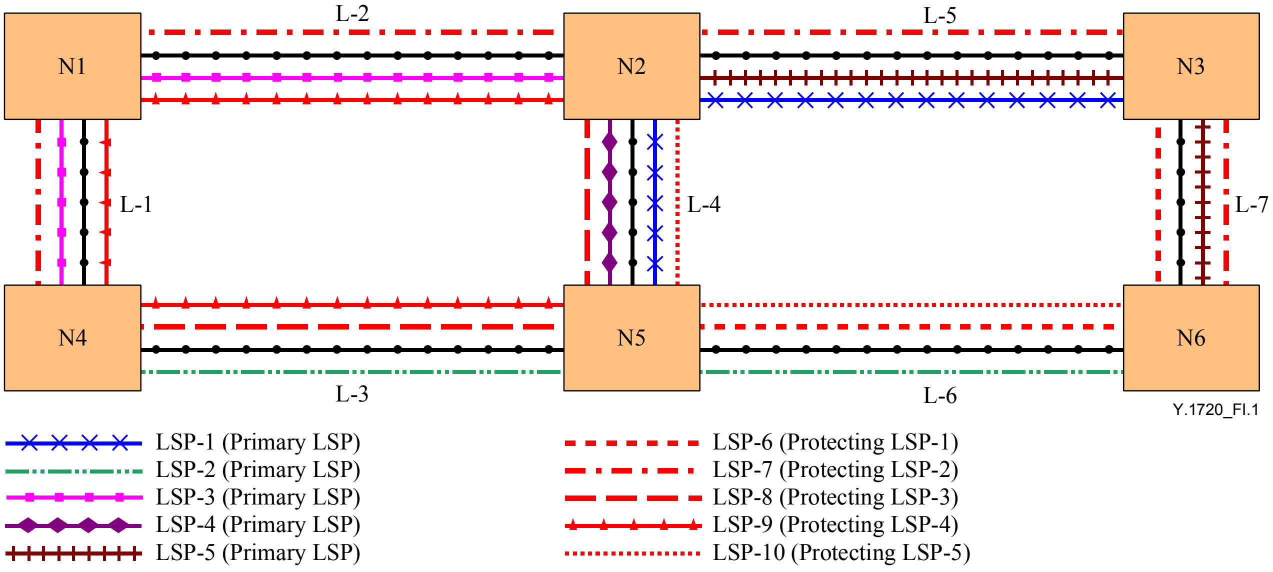

Figure I.1/Y.1720 – Sample working and protection connection paths

Realization of sharing of protection capacity among different failures can be achieved by keeping track for every link of the required amount of capacity to recover from each of them in the network. This can be illustrated using an example of an MPLS network. Figure I.1 shows an example network with five bidirectional working connections along with their five disjoint bidirectional protection connections. (Note that each connection consists of a pair of unidirectional LSPs.) For illustration purposes, assume that one unit of bandwidth is required by each working connection.

Table I.1 shows the amount of protection capacity required on each link for every possible single link or node failure in the network. To understand Table I.1, consider the first row associated with link L-1. The entry in column L-3 for that row indicates that there is one unit of traffic, due to LSP-2, on link L-3, which would employ link L-1 on its restoration route when link L-3 gets impacted by a failure. Similarly, the entry in column N5 addresses the case of a failure of node N5 and its impact on link L-1. The last column, titled Max, is the maximum value of all entries in that row. It is the amount of protection bandwidth that needs to be reserved on that link for the worst-case single failure in the network. This value for link L-6, for example, is 2 units to cover for failure of link L-5.

Table I.1/Y.1720 – Table tracking the failures and required protection bandwidth

Link	L-1	L-2	L-3	L-4	L-5	L-6	L-7	N1	N2	N3	N4	N5	N6	Max
L-1			1	1		1						1		1
L-2			1	1		1						1		1
L-3	1	1		1				1						1
L-4	1	1			1		1	1		1				1
L-5			1			1						1		1
L-6				1	2		1		1	1				2
L-7			1	1	1	1			1			1		1

The information in Table I.1 enables one to realize given the route and bandwidth of a working LSP, how much additional protection capacity needs to be reserved on each link along its route to guarantee its protection during a failure.

For a new protected connection, the contents of Table I.1 can be updated by updating the rows corresponding to the links along its protection path. This update involves incrementing, by the requested connection bandwidth, the value of every column corresponding to nodes and links along the working route. Then, the maximum of each updated row is computed as shown in the last column of Table I.1.

Table I.2/Y.1720 – Updated table accommodating additional gold service request

Link	L-1	L-2	L-3	L-4	L-5	L-6	L-7	N1	N2	N3	N4	N5	N6	Max
L-1			2	2		1						2		2
L-2			2	2		1						2		2
L-3	1	1		1				1						1
L-4	1	1			1		1	1		1				1
L-5			1			1						1		1
L-6				1	2		1		1	1				2
L-7			1	1	1	1			1			1		1

As an example, consider the arrival of a request for the shared mesh protection service between nodes N4 and N2 of the network in Figure I.1. Assume that, at the arrival of the request, the network was in a state illustrated in Figure I.1 and Table I.1. Further, assume that (N4-L3-N5-L4-N2) and (N4-L1-N1-L2-N2) are the computed working and protection routes, respectively, to service this request. Given these disjoint connection routes, links L-1 and L-2 of Table I.1 will be updated. The updated table is shown in Table I.2. Note that an additional one unit of bandwidth is now needed on both links L-1 and L-2 to guarantee the recovery of this new connection request from a failure along its working route.

Appendix II

Packet 1+1 example realization

The packet 1+1 scheme can be implemented by using a sequence as an identifier. The sequence number can be carried as the first four bytes inside the shim header of the LSP providing packet 1+1. Since the ingress and egress nodes must be aware of each LSP participating in the packet 1+1, the egress node will recognize that there is a sequence number inside the label. It will use the sequence number for selection purpose and then remove it before forwarding the accepted packet further. Note that packet 1+1 can be provided at any level of the hierarchy of a nested LSP. Figure II.1 illustrates the sequence number position behind the 4-bytes MPLS encapsulation header.

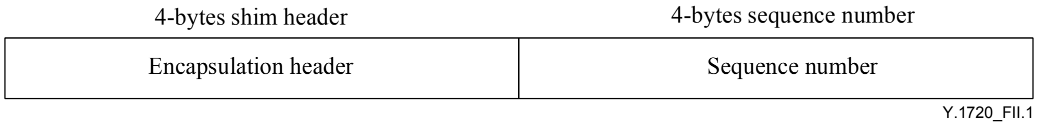

Figure II.1/Y.1720 – An illustration for sequence number transport

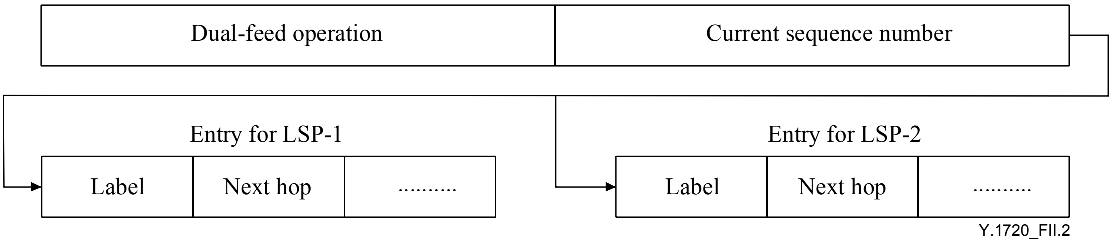

Figure II.2/Y.1720 – Enhanced NHLFE functionality to support dual-feed

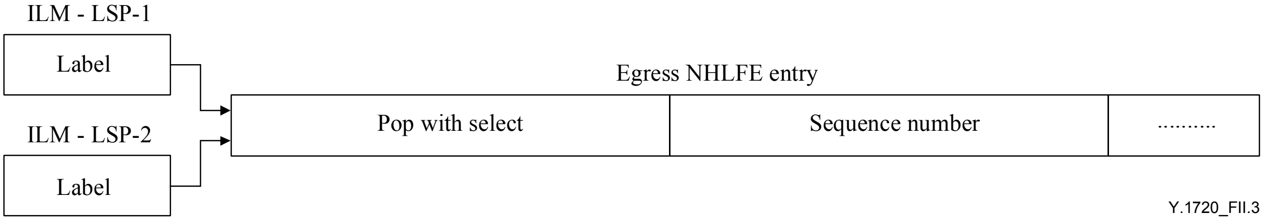

Figure II.3/Y.1720 – Enhanced NHLFE functionality to support selection

Dual-feed and select capabilities can be implemented at the MPLS shim layer by enhancing the Next-Hop-Label-Forward-Entry (NHLFE) entries. At the ingress node, to provide the dual-feed functionality, the NHLFE needs to support two instead of one outgoing LSP. This is easily achieved by using two next-hop/label entries instead of one, each corresponding to one of the mated diverse LSPs. Figure II.2 illustrates this case. Given this, when the client layer packet is forwarded to the NHLFE supporting dual-feed, it first duplicates the packet and then forwards it to the next hops with appropriate labels according to its two next-hop/label entries. In the middle of the network, each copy of the packet traverses the LSP in the standard way, as any other packet would traverse an LSP; thus transparent to the LSRs. On the egress node, the Incoming Label Map (ILM) needs to map the labels of the two diverse LSPs to a single NHLFE entry that enables the receive side to select one of possibly two received copies. Figure II.3 illustrates this case.

II.1 Dual-feed and select mechanism

Two components required for any dual-feed and select mechanism are:

1) the ability to dual-feed at one end; and

2) the ability to select appropriately from the dual-fed signal at the other end. Generally, realization of dual-feed is straightforward whereas, realization of select requires careful and often non-trivial treatment. At the source, packets can be dual-fed by copying on to two packet streams. At the destinations, each packet may be received twice at different times (or once only, or never), once from each of the two LSPs. In order to select each packet once, and once only, the destination must be able to identify the duplicate packets and to then select one, and to handle all possible variations. This selection process at the packet level is non-trivial as the duplicate packets may not arrive at the same time (due to propagation delay and buffering) and also these packets may get lost (due to transmission errors and buffer overflows).

The example algorithm below shows a method that addresses all these issues.

Algorithm

Variables:

```
N              /* number of bits to be used for sequence number */
```

```
rec_seq_no      /* the sequence number of the received packet */
select_counter  /* N bits counter at the receiver that keeps track of the sequence number of next
                   expected packet */
window_sz       /* size of the window; must be less than 2^N */
Initialization:
        Rec_seq_no = 0;
        select_counter = 0;
```

Algorithm:

```
Sender
        insert rec_seq_no to the inner "label" of the packet;
        transmit one copy of the packet on each mated LSPs;
        rec_seq_no ++;
Selector
        If(rec_seq_no is outside the sliding window defined by
          [select_counter, select_counter+window_sz])
           reject the packet;
        else    /* the rec_seq_no is in the window */
        {
                accept the packet;
                select_counter = rec_seq_no +1;
        }
```

II.2 Analysis of the packet 1+1 scheme

The ingress node inserts the sequence number. The packet is then duplicated and transported over diverse LSPs. Due to the diversity of the LSPs, there will be a leading LSP and a trailing LSP. The leading LSP will deliver the packets to the egress node faster than the trailing LSP. Therefore, under non-failure conditions, the egress node will select the packets from the leading LSP. The packets received on the trailing LSP will be duplicate packets and will, therefore, be discarded.

The decision whether to accept or discard a received packet is based on the received packet's sequence number and a counter + sliding window at the egress node. The counter indicates the sequence number of the next packet it is expecting. The counter, plus sliding window, provides a window of acceptable sequence numbers. The sliding window is needed to properly accept and reject packets. If the received packet falls in the window, it is considered legitimate and can be accepted: otherwise, it is rejected. The size of the window should be larger than the maximum number of consecutive packets a working (an alive) LSP can lose.

The sliding window is used to solve the problem of losing packets on the leading LSP when the leading LSP's sequence number is very close to the wrap-around point. Figure II.4 illustrates a leading LSP (LSP-1) that delivers a packet with a sequence number 29. The packet is accepted and the counter is incremented to 30. If we assume that 2 consecutive packets are lost (i.e., packets with sequence numbers 30 and 31), the next received packet, on LSP-1 will be 0. Without a sliding window, the egress node will reject the packet since 0 < 30. By implementing a sliding window that is larger than the maximum number of consecutive packets, a working (an alive) LSP can lose: this problem can be solved. For example, let's say that the maximum number of consecutive packets that a working LSP can lose is 5, then a sliding window of 6 can be defined. Taking the same example as before, however, now using the sliding window, the egress node will accept packets in the range of {30, 31, 0, 1, 2, 3}. Therefore, even if 5 packets are lost (i.e., the maximum number of consecutive packets that can be lost on a working LSP) the next packet received will have a sequence number 3 and the packet will be accepted.

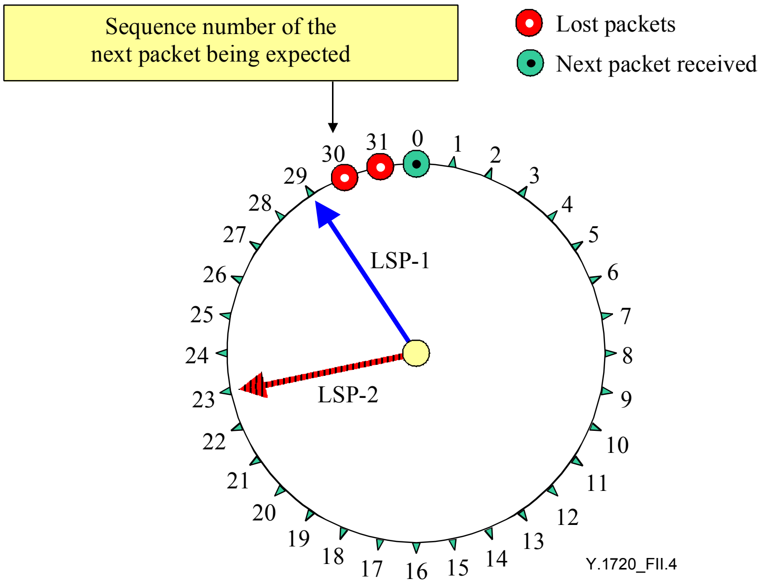

Figure II.4/Y.1720 – Packet loss in conjunction with wrap around

Note that this idea of a sliding window only works if the falling behind LSP cannot fall back in the sliding window range. If a packet with a sequence number in the range of the sliding window is received from the falling behind LSP, then it will be mistakenly accepted. A falling behind LSP can only receive a packet with a sequence number in the range of the sliding window if it falls back by more than (2^N – size of sliding window). Therefore, the number of bits N used for the sequence number must support the following equation:

$$2^N > \text{SlidingWindow} + \text{DelayWindow}$$

where:

SlidingWindow > maximum number of consecutive packets that can be lost on a LSP

and

DelayWindow = maximum number of packets the trailing LSP can fall behind the leading LSP

Note that the 4-byte field provides a sequence of more than 4 billion numbers which is large enough to accommodate worst-case consecutive packet losses and delay differentials.

One reasonable way of engineering the size of the sliding and delay windows is to make the size of the sliding window equal to the size of the delay window. (Note that it is assumed that the size of the delay window is generally larger than the size of the sliding window.) This guarantees selection of packets from the leading LSP in all scenarios after a failed LSP gets repaired. This point is further elaborated in the following clause which discusses various failure scenarios.

II.2.1 Operation of select mechanism under various failure scenarios

One way to view the operation of the select mechanism is to picture a clock with 2^N intervals. Figure II.5 illustrates an example where N = 4 (i.e., 4-bit sequence number) and, therefore, the sequence number ranges from 0 through 15. In this example, the sliding window is set equal to the delay window, which is 5.

Figure II.5 shows the leading LSP ahead of the trailing LSP by 3 sequence numbers. The leading LSP delivers a packet with a sequence number = 1 and the counter is now set to 2.

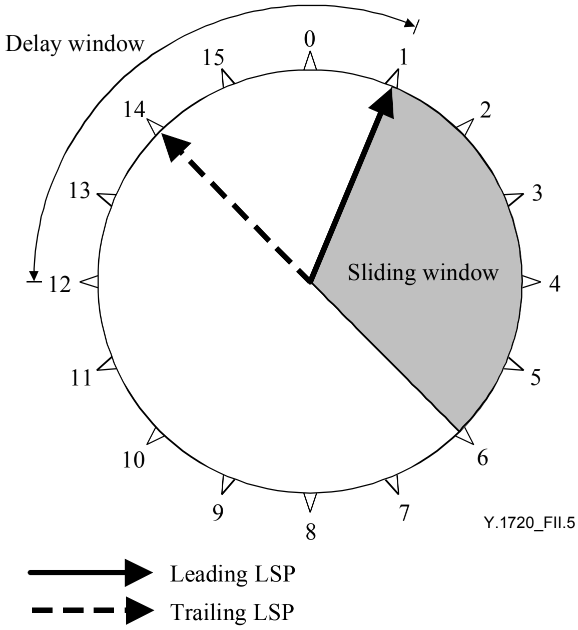

Figure II.5/Y.1720 – Sliding and delay windows concept

Figure II.6 shows that prior to receiving a packet with a sequence number equal to 2 on the leading LSP, the leading LSP fails. Until the packet with a sequence number equal to 2 is delivered from the trailing LSP, the egress node will not select any packets and the counter will remain equal to 2.

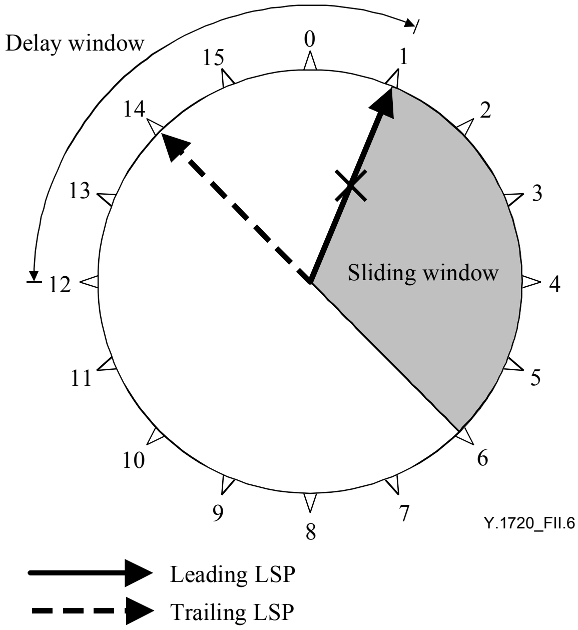

Figure II.6/Y.1720 – Leading LSP failure scenario

Figure II.7 illustrates that when the packet with a sequence number equal to 2 is received on the trailing LSP, the egress node increments the counter to 3 and the sliding window shifts so that a packet with a sequence number in the range of 3 through 7 can be accepted.

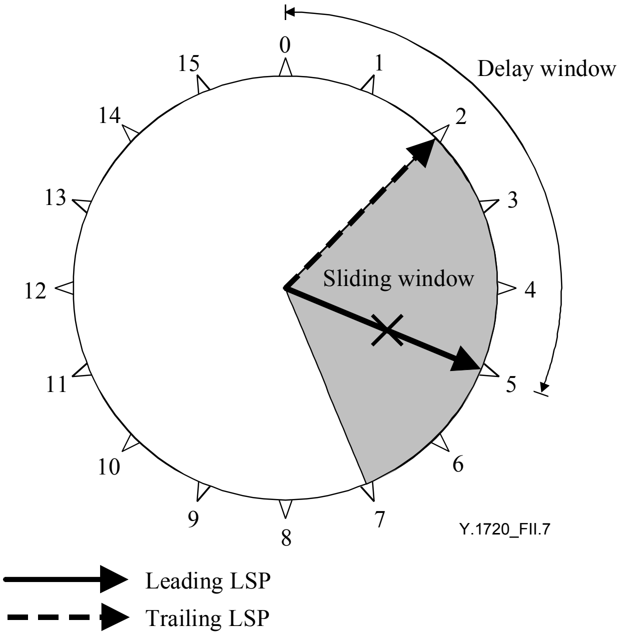

Figure II.7/Y.1720 – Traffic recovery after the leading LSP failure

Figure II.8 illustrates that prior to receiving a packet with a sequence number equal to 3 from the trailing LSP, the leading LSP is repaired and a packet with a sequence number equal to 6 is received from the leading LSP. Since 6 is within the sliding window range, the packet is accepted. Note that it is important that, so long as the leading LSP is working, packets are received from the leading LSP. Therefore, to ensure that when the leading LSP is repaired that it delivers a packet with a sequence number value that is within the sliding window range, the sliding window should be equal to or greater than the delay window which is the case for this example.

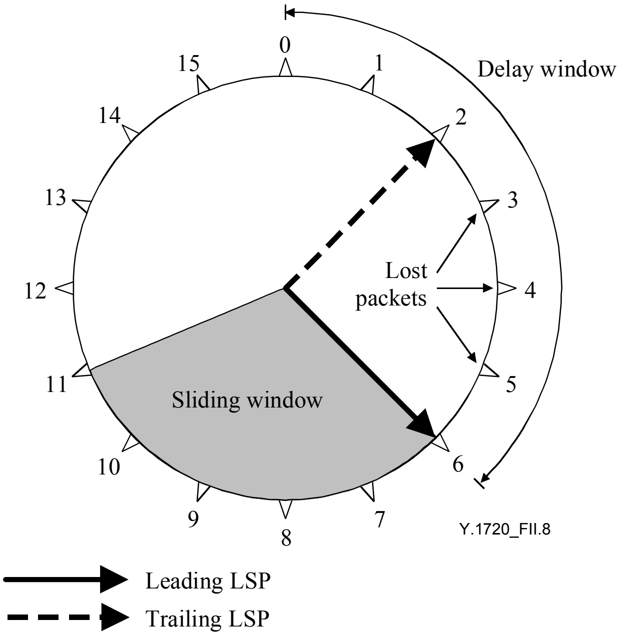

Figure II.8/Y.1720 – Leading LSP repair scenario

Figures II.9, II.10, and II.11 illustrate a problem if the sliding window is set smaller than the delay window. In this case, it is possible that when the leading LSP is repaired, it delivers packets with sequence numbers that fall outside the sliding window and, therefore, the egress node continues to accept packets from the trailing LSP. If, at a later time, the trailing LSP fails, there is a potential to lose many packets (worst case would be 2^N – *size_of_sliding_window*, where N is the number of bits used for the sequence number).

Figure II.9 shows an example where the sliding window is set to 3 while the delay window can be up to 6. In this example, the trailing LSP trails the leading LSP by 4 sequence numbers. Since the leading LSP has failed, the packets are selected from the trailing LSP.

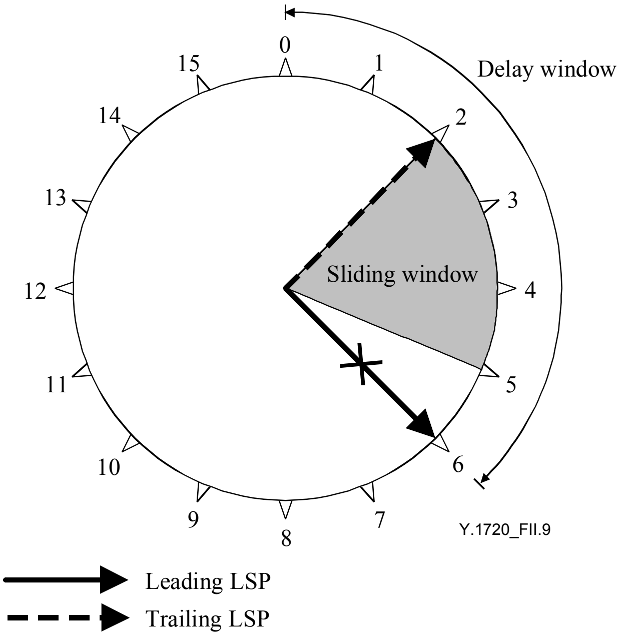

Figure II.9/Y.1720 – Scenario when sliding window < delay window

Figure II.10 illustrates that at the time when the leading LSP is repaired, it delivers a packet with a sequence number equal to 7 which is outside the sliding window and, therefore, rejected. The packets continue to be selected from the trailing LSP.

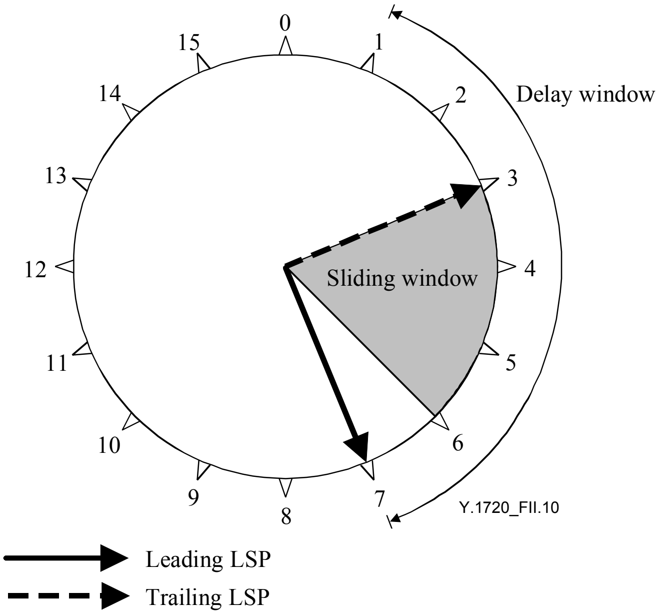

Figure II.10/Y.1720 – LSP repair: sliding window < delay window

Figure II.11 illustrates a failure to the trailing LSP. Since the leading LSP delivers packets outside the sliding window and, therefore, those packets are rejected, the egress node will not start accepting packets until the leading LSP comes all the way around and starts to deliver packets with a sequence number that falls within the sliding window. This can result in a significant loss of packets. Therefore to prevent such an occurrence, it is recommended that this type of selector algorithm set the sliding window equal to the delay window.

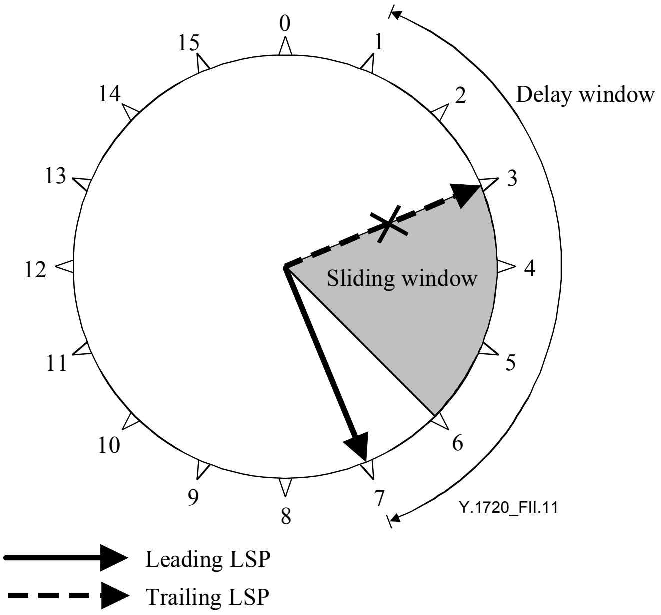

Figure II.11/Y.1720 – Possible problem: sliding window < delay window

II.2.2 Additional remarks

a) The scheme requires intelligence at the edge nodes only. Further, the scheme does not require any explicit fault detection or notification. This is implied by the packet selection scheme at the egress, which is carried out based on the sequence number and the locally maintained counters.

b) Dual feed requires duplication of packets at ingress. This introduces some additional minimum processing at the ingress. Selection requires comparisons of the sequence number carried in the packet with the counter value maintained at the receiver leading to a packet accept or reject condition. For hardware or software implementation, the processing cost is minimum. Another performance impact is the bandwidth cost due to the sequence number carried in the packets. This introduces some additional packet overhead depending on the length of the sequence number. With a 32-bit sequence number using the whole 4 bytes label, the bandwidth overhead is merely 4% for short 100 bytes packets.

c) The loss performance of the proposed service can be seen as follows. As the selection mechanism at the egress node takes packets from either LSP, the service in fact may compensate, although not required, the packet losses in the network. In the best case, this could result in zero loss, although each LSP may experience losses. On the other hand, in the worst case, the net packet loss would be the sum of the losses of both LSPs. In other words, the loss performance of the service is no worse and of the same order of magnitude as of the worst performing LSP, and sometimes could be much better.

d) The delay performance of the proposed service can be seen as follows. Since the algorithm always selects, without buffering, the first eligible arriving packet of the pair, the delay performance is always better than either of the LSPs.

e) The size of the window should be sized to be larger than the maximum number of consecutive packets a working LSP can lose. As a result, it is assured that the sequence number of the next packet from the same LSP will always fall within the window and will be accepted.

f) The size of the window should be sized such that the delay differential of the packet pairs traversing the mated LSPs, if not lost, is never more than (2^N – size of the window) packets. As a result, it is assured that an old packet will not be mistaken for a new one, thus causing mis-delivery.

g) In case of a single failure in the network, other than the ingress or egress nodes, only one of the mated diverse LSPs will be affected. The surviving LSP will continue delivering the packets. If the surviving LSP is the leading LSP, i.e., the last received and selected packet was from this LSP, then the select function at the egress node will continue to accept packets from it whereas, if the surviving LSP is the trailing LSP, then the select function rejects packets until it sees a packet whose sequence number falls within the sliding window. Upon successful repair of the failed LSP, if so desired, it may be brought back into service. In this "reverted restoration mode", the simplest approach would be to have the first dual-fed packet get the usual next sequence number, next to the one assigned to the last packet fed only on the surviving LSP alone. Various enhancements can be made to manage the service loss performance during this operation, if desired.

h) In case both LSPs have failed, additional mechanisms need to be defined to maintain the service and the LSP associated states to insure robust operations.

Appendix III

Bibliography

IETF, RFC 3469 (2003), *Framework for MPLS-based Recovery*, *Category: Informational*.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communications
Series Y	**Global information infrastructure, Internet protocol aspects and Next Generation Networks**
Series Z	Languages and general software aspects for telecommunication systems

Geneva, 2003