

CCAMP Working Group
Internet Draft
Expiration Date: August 2002

CCAMP GMPLS P&R Design Team

Eric Mannie (Ebone)- Editor
Dimitri Papadimitriou (Alcatel) - Editor

Deborah Brungard (AT&T)
Sudheer Dharanikota (Nayna)
Jonathan Lang (Calient)
Guangzhi Li (AT&T)
Bala Rajagopalan (Tellium)
Yakov Rekhter (Juniper)

February 2002

Recovery (Protection and Restoration) Terminology for GMPLS

[draft-mannie-gmpls-recovery-terminology-00.txt](#)

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

For potential updates to the above required-text see: <http://www.ietf.org/ietf/1id-guidelines.txt>

1. Abstract

This document defines a common terminology for GMPLS based recovery mechanisms (i.e. protection and restoration) that are under consideration by the CCAMP Working Group.

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 [RFC-2119](#) [1].

E. Mannie et. al.

1

Draft-design-team-gmpls-recovery-terminology-00.txt

February 2002

3. Introduction

This document defines a common terminology for GMPLS based recovery mechanisms (i.e. protection and restoration) that are under consideration by the CCAMP Working Group.

The terminology proposed in this document is intended to be independent of the underlying transport technologies and borrows from an ITU-T ongoing effort (G.gps ù Generic Protection Switching [[G.gps](#)]) and from the G.841 ITU-T Recommendation. The restoration terminology and concepts have been gathered from numerous sources including IETF drafts.

In the context of this document we will use the term "recovery" to denote both protection and restoration. The specific terms "protection" and "restoration" will only be used when differentiation is required.

Note that this document focuses on the terminology for the recovery of LSPs controlled by a GMPLS control plane. We focus on end-to-end, segment and span (i.e. link) LSP recovery. Terminology for control plane recovery is not in the scope of this document.

Protection and restoration of switched LSPs under tight time constraints is a challenging problem. This is particularly relevant to optical networks that consist of TDM and/or all-optical (photonic) cross-connects referred to as GMPLS nodes (or simply nodes, or even sometimes "LSRs") connected in a general topology [[GMPLS-ARCH](#)].

Recovery typically involves the activation of a recovery (or alternate) LSP when a failure is encountered in the working (or primary) LSP.

A working or recovery LSP is characterized by an ingress interface, an egress interface, and a set of intermediate nodes and spans through which the LSP is routed. The working and recovery LSPs are typically resource disjoint (e.g. node and/or span disjoint). This ensures that a single failure will not affect both the working and recovery LSPs.

A bi-directional span between neighboring nodes is usually realized as a pair of unidirectional spans. The end-to-end path for a bi-

directional LSP therefore consists of a series of bi-directional segments (i.e. Sub-Network Connections, or SNCs, in the ITU-T terminology) between the source and destination nodes, traversing intermediate nodes.

4. Recovery Terminology Common to Protection and Restoration

This section defines the following general terms common to both protection and restoration (i.e. recovery). In addition, most of these terms apply to end-to-end, segment and span LSP recovery. Note that span recovery assumes that the nodes at each end of the span did not fail, otherwise end-to-end or segment LSP recovery is used.

The terminology and the definitions have been originally taken from G.gps. However, for generalization, the following language that is not directly related to recovery has been adapted to GMPLS and the common IETF terminology:

An LSP is used as a generic term to designate either an SNC (Sub-Network Connection) or an NC (Network Connection) in ITU-T terminology. The ITU-T uses the term transport entity to designate either a link, an SNC or an NC. The term "Traffic" is used instead of "Traffic Signal". The term protection or restoration "scheme" is used instead of protection or restoration "architecture".

The reader is invited to read G.841 and G.gps for references to SDH protection and ITU-T generic protection terminology. Note that restoration is not in the scope of G.gps.

4.1 Working and Recovery LSP/Span

A working LSP/span is an LSP/span transporting "normal" user traffic. A recovery LSP/span is an LSP/span used to transport "normal" user traffic when the working LSP/span fails. Additionally, the recovery LSP/span may transport "extra" user traffic (i.e. pre-emptable traffic) when normal traffic is carried over the working LSP/span.

4.2 Traffic Types

The different types of traffic that can be transported over an LSP/span in the context of this document are defined hereafter:

A. Normal traffic:

User traffic that may be protected by two alternative LSPs/spans (the working and recovery LSPs/spans).

B. Extra traffic:

User traffic carried over recovery resources (e.g. a recovery LSP/span) when these resources are not being used for the recovery of normal traffic; i.e. when the recovery resources are in standby mode. When the recovery resources are required to recover normal traffic from the failed working LSP/span, the extra traffic is pre-empted. Extra traffic is not protected by definition, but may be restored.

C. Null traffic:

Traffic carried over the recovery LSP/span if it is not used to carry normal or extra traffic. Null traffic can be any kind of traffic that conforms to the signal structure of the specific layer, and it is ignored (not selected) at the egress of the recovery LSP/span.

4.3 LSP/Span Protection and Restoration

The following subtle distinction is generally made between the terms "protection" and "restoration", even though these terms are often used interchangeably [TEWG].

The distinction between protection and restoration is made based on the resource allocation done during the recovery LSP/span establishment. The distinction between different types of restoration is made based on the level of route computation, signaling and resource allocation done during the restoration LSP/span establishment.

A. LSP/Span Protection

LSP/span protection denotes the paradigm whereby one or more dedicated protection LSP(s)/span(s) is/are fully established to protect one or more working LSP(s)/span(s).

For a protection LSP, this implies that route computation took place, that the LSP was fully signaled all the way and that its resources were fully selected (i.e. allocated) and cross-connected

between the ingress and egress nodes.

For a protection span, this implies that the span has been selected and reserved for protection.

Indeed, it means that no signaling takes place to establish the protecting LSP/span when a failure occurs. However, various other kinds of signaling may take place between the ingress and egress nodes for fault notification, to synchronize their use of the protecting LSP/span, for reversion, etc.

B. LSP/Span Restoration:

LSP/span restoration denotes the paradigm whereby some restoration resources may be pre-computed, signaled and selected a priori, but not cross-connected to restore a working LSP/span. The complete establishment of the restoration LSP/span occurs only after a failure of the working LSP/span, and requires some additional signaling.

Both protection and restoration require signaling. Signaling to establish the recovery resources and signaling associated with the use of the recovery LSP(s)/span(s) are needed.

E. Mannie et. al. Internet Draft - August 2002 4

Draft-design-team-gmpls-recovery-terminology-00.txt February 2002

4.4 Recovery Scope

Recovery can be applied at various levels throughout the network. Local (span) recovery refers to the recovery of an LSP over a link between two nodes. Segment recovery refers to the recovery of an LSP segment (i.e. an SNC in the ITU-T terminology) between two nodes, i.e. the boundary nodes of the segment. End-to-end recovery refers to the recovery of an entire LSP from its source to its destination. An LSP may be subject to local (span), segment, and/or end-to-end recovery.

4.5 Recovery Domain

A recovery domain is defined as a set of nodes and spans over which one or more recovery schemes are provided. A recovery domain served by one single recovery scheme is referred to as a "single recovery domain", while a recovery domain served by multiple recovery schemes is referred to as a "multi recovery domain".

4.6 Recovery Types

The different recovery types can be classified depending on the

number of recovery LSPs/spans that are protecting a given number of working LSPs/spans. The definitions given hereafter are from the point of view of a working LSP/span that needs to be protected by a recovery scheme.

A. 1+1 type: dedicated protection

One dedicated protection LSP/span protects exactly one working LSP/span and the normal traffic is permanently duplicated at the ingress node on both the working and protection LSPs/spans. No extra traffic can be carried over the protection LSP/span.

This type is applicable to LSP/span protection, but not to LSP/span restoration.

B. 0:1 type: unprotected

No specific recovery LSP/span protects the working LSP/span. However, the working LSP/span can potentially be restored through any alternate available route/span, with or without any pre-computed restoration route. Note that there are no resources pre-established for this recovery type.

This type is applicable to LSP/span restoration, but not to LSP/span protection. Span restoration can be for instance achieved by moving all the LSPs transported over of a failed span to a dynamically selected span.

C. 1:1 type: dedicated recovery with extra traffic

One specific recovery LSP/span protects exactly one specific working LSP/span but the normal traffic is transmitted only over one LSP (working or recovery) at a time. Extra traffic can be transported using the recovery LSP/span resources.

This type is applicable to LSP/span protection and LSP restoration, but not to span restoration.

D. 1:N (N>1) type: shared recovery with extra traffic

A specific recovery LSP/span is dedicated to the protection of up to N working LSPs/spans. The set of working LSPs/spans is explicitly identified. Extra traffic can be transported over the recovery LSP/span. All these LSPs/spans must start and end at the same nodes.

Sometimes, the working LSPs/spans are assumed to be resource

disjoint in the network so that they do not share any failure probability, but this is not mandatory. Obviously, if more than one working LSP/span in the set of N are affected by some failure(s) at the same time, the traffic on only one of these failed LSPs/spans may be recovered over the recovery LSP/span. Note that N can be arbitrarily large (i.e. infinite). The choice of N is a policy decision.

This type is applicable to LSP/span protection and LSP restoration, but not to span restoration.

Note: a shared recovery where each recovery resource can be shared by a maximum of X LSPs/spans is not defined as a recovery type but as a recovery scheme. The choice of X is a network resource management policy decision.

E. M:N ($M, N > 1; M \leq N$) type:

A set of M specific recovery LSPs/spans protects a set of up to N specific working LSPs/spans. The two sets are explicitly identified. Extra traffic can be transported over the M recovery LSPs/spans when available. All the LSPs/spans must start and end at the same nodes.

Sometimes, the working LSPs/spans are assumed to be resource disjoint in the network so that they do not share any failure probability, but this is not mandatory. Obviously, if several working LSPs/spans in the set of N are concurrently affected by some failure(s), the traffic on only M of these failed LSPs/spans may be recovered. Note that N can be arbitrarily large (i.e. infinite). The choice of N and M is a policy decision.

This type is applicable to LSP/span protection and LSP restoration, but not to span restoration.

4.7 Bridge Types

A bridge is the function that connects the normal traffic and extra traffic to the working and recovery LSP/span.

A. Permanent bridge

Under a 1+1 type, the bridge connects the normal traffic to both the working and protection LSPs/spans. This type of bridge is not

applicable to restoration types. There is of course no extra traffic connected to the recovery LSP/span.

B. Broadcast bridge

For 1:N and M:N types, the bridge permanently connects the normal traffic to the working LSP/span. In the event of recovery switching, the normal traffic is additionally connected to the recovery LSP/span. Extra traffic is either not connected or connected to the recovery LSP/span.

C. Selector bridge

For 1:N and M:N types, the bridge connects the normal traffic to either the working or the recovery LSP/span. Extra traffic is either not connected or connected to the recovery LSP/span.

4.8 Selector Types

A selector is the function that extracts the normal traffic either from the working or the recovery LSP/span. Extra traffic is either extracted from the recovery LSP/span, or is not extracted.

A. Selective selector

Is a selector that extracts the normal traffic from either the working LSP/span output or the recovery LSP/span output.

B. Merging selector

For 1:N and M:N protection types, the selector permanently extracts the normal traffic from both the working and recovery LSP/span outputs. This alternative works only in combination with a selector bridge.

4.9 Recovery GMPLS Nodes

This section defines the GMPLS nodes involved during recovery.

A. Ingress GMPLS node of an end-to-end LSP/segment LSP/span

The ingress node of an end-to-end LSP/segment LSP/span is where the working traffic may be bridged to the recovery end-to-end LSP/segment LSP/span. Also known as source node in the ITU-T terminology.

B. Egress GMPLS node of an end-to-end LSP/segment LSP/span

The egress node of an end-to-end LSP/segment LSP/span is where the working traffic may be selected from either the working or the recovery end-to-end LSP/segment LSP/span. Also known as sink node in the ITU-T terminology.

C. Intermediate GMPLS node of an end-to-end LSP/segment LSP

A node along either the working or recovery end-to-end LSP/segment LSP route between the corresponding ingress and egress nodes. Also known as intermediate node in the ITU-T terminology.

4.10 Switching Mechanism

A switch is an action that can be performed at both the bridge and the selector. This action is as follows:

A. For the selector:

The action of selecting normal traffic from the recovery LSP/span rather than from the working LSP/span.

B. For the bridge:

In case of permanent connection to the working LSP/span, the action of connecting or disconnecting the normal traffic to the recovery LSP/span. In case of non-permanent connection to the working LSP/span, the action of connecting the normal traffic to the recovery LSP/span.

4.11 Reversion operations

A revertive recovery operation refers to a recovery switching operation, where the traffic returns to (or remains on) the working LSP/span if the switch requests are terminated; i.e. when the working LSP/span has recovered from the failure.

Therefore a non-revertive recovery switching operation is when the traffic does not return to the working LSP/span if the switch requests are terminated.

4.12 Failure Reporting

This section gives (for information) several signal types commonly used in transport planes to report a failure condition. Note that fault reporting may require additional signaling mechanisms.

A. Signal Degrade (SD): a signal indicating that the associated data has degraded.

B. Signal Fail (SF): a signal indicating that the associated data has failed.

C. Signal Degrade Group (SDG): a signal indicating that the associated group data has degraded (under discussion at the ITU-T).

D. Signal Fail Group (SFG): a signal indicating that the associated group has failed (under discussion at the ITU-T).

4.13 External commands

This section defines several external commands, typically issued by an operator through the NMS/EMS, which can be used to influence or command the recovery schemes.

A. Lockout of recovery LSP/span:

A configuration action initiated externally that results in the recovery LSP/span being temporarily unavailable to transport traffic (either normal or extra traffic).

B. Lockout of normal traffic:

A configuration action initiated externally that results in the normal traffic being temporarily not allowed to be routed over its recovery LSP/span.

C. Freeze:

A configuration action initiated externally that prevents any switch action to be taken, and as such freezes the current state.

D. Forced switch for normal traffic:

A switch action initiated externally that switches normal traffic to the recovery LSP/span, unless an equal or higher priority switch command is in effect.

E. Manual switch for normal traffic:

A switch action initiated externally that switches normal traffic to the recovery LSP/span, unless a fault condition exists on other LSPs/spans (including the recovery LSP/span) or an equal or higher priority switch command is in effect.

4.14 Unidirectional versus Bi-Directional Recovery Switching

A. Unidirectional recovery switching:

A recovery switching mode in which, for a unidirectional fault (i.e. a fault affecting only one direction of transmission), only the normal traffic transported in the affected direction (of the LSP or span) is switched to the recovery LSP/span.

B. Bi-directional recovery switching:

A recovery switching mode in which, for a unidirectional fault, the normal traffic in both directions (of the LSP or span), including the affected direction and the unaffected direction, are switched to the recovery LSP/span.

4.15 Recovery Schemes Related Time and Durations

This section gives several typical timing definitions that are of importance for recovery schemes.

A. Detection time:

The time between the occurrence of the fault or degradation and its detection. Note that this is a rather theoretical time since in practice this is difficult to measure.

B. Correlation time:

The time between detection of the fault or degradation and the reporting of the signal fail or degrade. This time is typically used in correlating related failures or degradations.

C. Hold-off time:

The time between reporting of signal fail or degrade, and the initialization of the recovery switching algorithm. This is useful when multiple layers of recovery are being used.

D. Wait To Restore time:

A period of time that must elapse from a recovered fault before an LSP/span can be used again to transport the normal traffic and/or to select the normal traffic from.

E. Switching time:

The time between the initialization of the recovery switching algorithm and the moment the traffic is selected from the recovery LSP/span.

E. Mannie et. al. Internet Draft - August 2002 10

Draft-design-team-gmpls-recovery-terminology-00.txt February 2002

F. Recovery time:

The sum of the detection, correlation, hold-off and switching times.

4.16 Impairment

A defect or performance degradation, which may lead to SF or SD trigger.

4.17 Recovery Ratio

The quotient of the actually recovery bandwidth divided by the traffic bandwidth which is intended to be protected.

4.18 Hitless Protection Switch

Protection switch, which does not cause data loss, data duplication, data disorder, or bit errors upon recovery switching action.

4.19 Network Survivability

The set of capabilities that allow a network to restore affected traffic in the event of a failure. The degree of survivability is determined by the network's capability to survive single and multiple failures.

4.20 Survivable Network

A network that is capable of restoring traffic in the event of a failure.

4.21 Escalation

A network survivability action caused by the impossibility of the survivability function in lower layers.

5. Recovery Phases

It is commonly accepted that recovery implies that the following generic operations need to be performed when an LSP/span or a node failure occurs:

- Phase 1: Failure Detection

TBD.

- Phase 2: Failure Localization and Isolation

TBD.

- Phase 3: Failure Notification

TBD.

E. Mannie et. al. Internet Draft - August 2002 11

Draft-design-team-gmpls-recovery-terminology-00.txt February 2002

- Phase 4: Recovery (Protection or Restoration)

See above.

- Phase 5: Reversion (Normalization)

See above.

The combination of Failure Detection and Failure Localization and Notification is referred to as Fault Management.

5.1 Entities Involved During Recovery

The entities involved during the recovery operations can be defined as follows; these entities are parts of ingress, egress and intermediate nodes as defined previously:

A. Detecting Entity (Failure Detection):

An entity that detects a failure or group of failures; providing thus a non-correlated list of failures.

B. Reporting Entity (Failure Correlation and Notification):

An entity that can make an intelligent decision on fault correlation and report the failure to the deciding entity. Fault reporting can be automatically performed by the deciding entity detecting the failure.

C. Deciding Entity (part of the failure recovery decision process):

An entity that makes the recovery decision or select the recovery resources. This entity communicates the decision to the impacted LSPs/spans with the recovery actions to be performed.

D. Recovering Entity (part of the failure recovery activation process):

An entity that participates in the recovery of the LSPs/spans.

6. Protection Schemes

This section clarifies the multiple possible protection schemes and the specific terminology for the protection.

To be completed with references to ITU-T protection schemes and a table summarizing the multiple ITU-T protection schemes.

E. Mannie et. al. Internet Draft - August 2002 12

Draft-design-team-gmpls-recovery-terminology-00.txt February 2002

7. Restoration Schemes

This section clarifies the multiple possible restoration schemes and the specific terminology for the restoration.

To be completed when an agreement on restoration scheme definitions and mechanisms has been achieved in other drafts.

8. Security Considerations

This document does not introduce or imply any specific security consideration.

9. References

- [G.707] ITU-T Recommendation G.707, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)", April 2000.
- [G.709] ITU-T Recommendation G.709, "Network Node Interface for the Optical Transport Network (OTN)", February 2001 (+ revision 1, Oct 2001).
- [G.783] ITU-T Recommendation G.783, "Characteristics of

Synchronous Digital Hierarchy (SDH) Equipment
Functional Blocks

- [G.798] ITU-T Recommendation G.798, "Characteristics of Optical Transport Network (OTN) Equipment Functional Blocks"
- [G.806] ITU-T Recommendation G.806, "Characteristics of Transport Equipment - Description Methodology and Generic Functionality"
- [G.841] ITU-T Recommendation G.841, "Types and Characteristics of SDH Network Protection Architectures"
- [G.842] ITU-T Recommendation G.842, "Interworking of SDH network protection architectures"
- [G.gps] ITU-T on-going work G.gps, "Generic Protection Switching", ITU-T Draft (October, 2001).
- [ANSI-T1.105] "Synchronous Optical Network (SONET): Basic Description Including Multiplex Structure, Rates, and Formats" ANSI T1.105, 2000.
- [BALA] Rajagopalan, Bala et al, "Signaling for Protection and Restoration in Optical Mesh Networks", Internet Draft, Draft-bala-protection-restoration-signaling-00.txt.
- [GMPLS-ARCH] E.Mannie Editor, "GMPLS Architecture", Internet Draft, [draft-ietf-ccamp-gmpls-architecture-01.txt](#), November 2001.

E. Mannie et. al. Internet Draft - August 2002 13

Draft-design-team-gmpls-recovery-terminology-00.txt February 2002

- [TEWG] W.S Lai, et al., "Network Hierarchy and Multilayer Survivability", Internet Draft, [draft-team-tewg-restore-hierarchy-00.txt](#), July 2001.
- [SUDHEER] Sudheer Dharanikota et al., "NNI Protection and restoration requirements", OIF 2001.507.

10. Acknowledgments

Valuable comments and input were received from many people.

11. Author's Addresses

Deborah Brungard
AT&T

Rm. D1-3C22
200 S. Laurel Ave.
Middletown, NJ 07748
USA
Email: dbrungard@att.com

Sudheer Dharanikota
Nayna Networks Inc
481 Sycamore Drive
Milpitas, CA 95035
USA
Email: sudheer@nayna.com

Jonathan P. Lang
Calient Networks
25 Castilian
Goleta, CA 93117
Email: jplang@calient.net

Guangzhi Li
AT&T
180 Park Avenue,
Florham Park, NJ 07932
gli@research.att.com
973-360-7376

Eric Mannie
Ebony
Terhulpesteenweg 6A
1560 Hoeilaart
Belgium
Phone: +32 2 658 56 52
Email: eric.mannie@ebony.com

E. Mannie et. al. Internet Draft - August 2002

14

Draft-design-team-gmpls-recovery-terminology-00.txt

February 2002

Dimitri Papadimitriou
Alcatel
Francis Wellesplein, 1
B-2018 Antwerpen
Belgium
Phone: +32 3 240-84-91
Email: dimitri.papadimitriou@alcatel.be

Bala Rajagopalan

Tellium, Inc.
2 Crescent Place
P.O. Box 901
Oceanport, NJ 07757-0901
USA
Phone: +1 732 923 4237
Email: braja@tellium.com

Yakov Rekhter
Juniper
Email: yakov@juniper.net