International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.873.1

(03/2006)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

Digital networks – Optical transport networks

Optical Transport Network (OTN): Linear protection

ITU-T Recommendation G.873.1

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TRANSMISSION MEDIA CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
General aspects	G.800–G.809
Design objectives for digital networks	G.810–G.819
Quality and availability targets	G.820–G.829
Network capabilities and functions	G.830–G.839
SDH network characteristics	G.840–G.849
Management of transport network	G.850–G.859
SDH radio and satellite systems integration	G.860–G.869
Optical transport networks	**G.870–G.879**
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DATA OVER TRANSPORT – GENERIC ASPECTS	G.7000–G.7999
ETHERNET OVER TRANSPORT ASPECTS	G.8000–G.8999
ACCESS NETWORKS	G.9000–G.9999

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

ITU-T Recommendation G.873.1

Optical Transport Network (OTN): Linear protection

Summary

This Recommendation defines the APS protocol and protection switching operation for the linear protection schemes for the Optical Transport Network at the Optical Channel Data Unit (ODUk) level. Protection schemes considered in this Recommendation are:

– ODUk subnetwork connection protection with inherent monitoring (1+1, 1:n);

– ODUk subnetwork connection protection with non-intrusive monitoring (1+1);

– ODUk subnetwork connection protection with sublayer monitoring (1+1, 1:n).

Source

ITU-T Recommendation G.873.1 was approved on 29 March 2006 by ITU-T Study Group 15 (2005-2008) under the ITU-T Recommendation A.8 procedure.

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 not 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 G.873.1

Optical Transport Network (OTN): Linear protection

1 Scope

This Recommendation defines the APS protocol and protection switching operation for the linear protection schemes for the Optical Transport Network at the Optical Channel Data Unit (ODUk) level. Linear protection schemes considered in this Recommendation are:

– ODUk subnetwork connection protection with inherent monitoring (1+1, 1:n);

– ODUk subnetwork connection protection with non-intrusive monitoring (1+1);

– ODUk subnetwork connection protection with sublayer monitoring (1+1, 1:n).

The APS protocol and protection switching operation for OTN ring protection is currently under study.

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.

– ITU-T Recommendation G.709/Y.1331 (2003), *Interfaces for the Optical Transport Network (OTN)*.

– ITU-T Recommendation G.798 (2004), *Characteristics of optical transport network hierarchy equipment functional blocks*.

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

– ITU-T Recommendation G.806 (2006), *Characteristics of transport equipment – Description methodology and generic functionality*.

– ITU-T Recommendation G.808.1 (2006), *Generic protection switching – Linear trail and subnetwork protection*.

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

– ITU-T Recommendation G.872 (2001), *Architecture of optical transport networks*.

3 Definitions

This Recommendation defines the following terms:

3.1 APS channel: See ITU-T Rec. G.870/Y.1352.

3.2 entity: See ITU-T Rec. G.870/Y.1352.

3.3 extra traffic signal: See ITU-T Rec. G.870/Y.1352.

3.4 head end: See ITU-T Rec. G.870/Y.1352.

3.5 normal traffic signal: See ITU-T Rec. G.870/Y.1352.

3.6 **null signal**: See ITU-T Rec. G.870/Y.1352.

3.7 **protection communication channel**: See ITU-T Rec. G.870/Y.1352.

3.8 **protection group**: See ITU-T Rec. G.870/Y.1352.

3.9 **signal**: See ITU-T Rec. G.870/Y.1352.

3.10 **tail end**: See ITU-T Rec. G.870/Y.1352.

4 Abbreviations

This Recommendation uses the following abbreviations:

APS	Automatic Protection Switching
DNR	Do Not Revert
EXER	Exercise
FS	Forced Switch
LO	Lockout for protection
MS	Manual Switch
NR	No Request
ODUk	Optical Channel Data Unit k
OTN	Optical Transport Network
OTUk	Optical Channel Transport Unit k
PCC	Protection Communication Channel
RR	Reverse Request
SD	Signal Degrade
SF	Signal Fail
WTR	Wait-to-Restore

5 Protection characteristics

5.1 Monitoring methods and conditions

Protection switching will occur based on the detection of certain defects on the transport entities (working and protection) within the protected domain. How these defects are detected is the subject of the equipment Recommendations (e.g., ITU-T Recs G.806 and G.798). For the purpose of the protection switching controller, an entity within the protected domain has a condition of no defect = OK, degraded (signal degrade = SD), or failed (signal fail = SF).

The customary monitoring methods are as follows:

Inherent – Protection switching is triggered by defects detected at the ODUk link connection (e.g., server layer tail and server/ODUk adaptation function). No defect detection is performed at the ODUk layer itself.

NOTE – In contrast to SDH SNC/I, ODUk SNC/I can stretch only a single link connection, as the FDI defect resulting from further upstream server layer defects is not detected in the server/ODUk adaptation function.

Non-intrusive – Protection switching is triggered by a non-intrusive monitor of the ODUkP layer or ODUkT sub-layers at the tail end of the protection group.

Sublayer – Protection switching is triggered by defects detected at the ODUkT sublayer trail (TCM). An ODUkT sublayer trail is established for each working and protection entity. Protection switching is therefore triggered only on defects of the protected domain.

The protection switching controller does not care which monitoring method is used, as long as it can be given (OK, SD, SF) information for the transport entities within the protected domain. Some monitors or network layers may not have an SD detection method. Where this is the case, there is no need to use a different APS protocol – it would simply happen that an SD would not be issued from an equipment that cannot detect it. Where an APS protocol is used, the implementation should not preclude that the far end declares an SD over the APS channel, even if the monitor at the near end cannot detect SD.

6 Protection group commands

6.1 End-to-end commands and states

This clause describes commands that apply to the protection group as a whole. When an APS is present, these commands are signalled to the far end of the connection. In bidirectional switching, these commands affect the bridge and selector at both ends.

Lockout of protection – This command prevents a working signal from being selected from the protection entity. This effectively disables the protection group. An Extra traffic signal, if present on the protection entity, is dropped.

Force switch normal traffic signal #n to protection – Forces Normal traffic Signal #n to be selected from the protection entity after the required bridge is present.

Force switch null signal – For 1:n architectures, it switches the null signal to the protection entity, unless an equal or higher priority switch command is in effect. A normal traffic signal present on the protection entity is transferred to and selected from its working entity. For 1+1 architectures, it selects the normal traffic signal from the working entity.

Force switch extra traffic signal – It switches the extra traffic signal to the protection entity, unless an equal or higher priority switch command is in effect. A normal traffic signal present on the protection entity is transferred to and selected from its working entity.

Manual switch normal traffic signal #n to protection – In the absence of a failure of a working or protection entity, forces Normal traffic Signal #n to be selected from the protection entity after the required bridge is present.

Manual switch null signal – For 1:n architectures, it switches the null signal to the protection entity, unless a fault condition exists on other entities or an equal or higher priority switch command is in effect. A normal traffic signal present on the protection entity is transferred to and selected from its working entity. For 1+1 architectures, it selects the normal traffic signal from the working entity.

Manual switch extra traffic signal – It switches extra traffic signal to the protection entity, unless a fault condition exists on other entities or an equal or higher priority switch command is in effect. A normal traffic signal present on the protection entity is transferred to and selected from its working entity.

Wait-to-restore normal traffic signal #n – In revertive operation, after the clearing of an SF or SD on working entity #n, maintains Normal traffic Signal #n as selected from the protection entity until a Wait-to-Restore timer expires. If the timer expires prior to any other event or command, the state will be changed to NR. This is used to prevent frequent operation of the selector in the case of intermittent failures.

Exercise signal #n – Exercise of the APS protocol. The signal is chosen so as not to modify the selector.

Do not revert normal traffic signal #n – In non-revertive operation, this is used to maintain a normal traffic signal selected from the protection entity.

No request – All Normal Traffic Signals are selected from their corresponding Working transport entities. The protection entity carries either the null signal, extra traffic, or a bridge of the single normal traffic signal in a 1+1 protection group.

Clear – Clears the active near end Lockout of Protection, Forced Switch, Manual Switch, WTR state, or Exercise command.

6.2 Local commands

These commands apply only to the near end of the protection group. When an APS is present, they have not been signalled to the far end via the APS channel.

Freeze – Freezes the state of the protection group. Until the freeze is cleared, additional near end commands are rejected. Condition changes and received APS bytes are ignored. When the Freeze command is cleared, the state of the protection group is recomputed based on the condition and received APS bytes.

Clear freeze

Lockout normal traffic signal #n from protection – Prevents Normal Traffic Signal #n from being selected from the protection entity. Commands for Normal Traffic Signal #n will be rejected. SF or SD will be ignored for Normal Traffic Signal #n. In bidirectional 1:n switching, remote bridge requests for Normal Traffic Signal #n will still be honoured to prevent protocol failures. As a result, a Normal Traffic Signal must be locked out from protection at both ends to prevent it being selected from the protection entity as a result of a command or failure at either end. Multiple of these commands may coexist for different Normal Traffic Signals.

Clear lockout normal traffic signal #n from protection

7 Protection architectures

In a linear protection architecture, protection switching occurs at the two distinct endpoints of a protected trail or protected subnetwork connection. Between these endpoints, there will be both "working" and "protection" entities.

For a given direction of transmission, the "head end" of the protected signal is capable of performing a bridge function, which will place a copy of a normal traffic signal onto a protection entity when required. The "tail end" will perform a selector function, where it is capable of selecting a normal traffic signal either from its usual working entity, or from a protection entity. In the case of bidirectional transmission, where both directions of transmission are protected, both ends of the protected signal will normally provide both bridge and selector functions.

The following architectures are possible:

1+1 – In a 1+1 architecture, a single normal traffic signal is protected by a single protection entity. The bridge at the head end is permanent. Switching occurs entirely at the tail end.

1:n – In a 1:n architecture, 1 or more normal traffic signal(s) are protected by a single protection entity. The bridge at the head end is not established until a protection switch is required. In the case where n > 1, it cannot be known which of the normal traffic signals should be bridged onto the protection entity, until a defect is detected on one of the protected signals. 1:n architectures are capable of carrying an extra (low priority, preemptable) traffic signal on the protection entity when it is not being used to protect any normal traffic signal. A 1:n architecture can be used even for n = 1 (1:1). This might be chosen over the simpler 1+1 architecture (which requires no head end actions of the protection algorithm) since 1:1 is capable of carrying extra traffic, where 1+1 is not.

m:n – In this architecture, m protection entities are used to protect n working entities. This is for further study.

With the assumption of a larger APS channel, the coding for the entity number "n" will use a full byte rather than the few bits in SDH. Two of the 256 values are reserved: 0 is used to indicate a null signal or the protection entity, and 0xFF (255) is used to indicate extra traffic.

The architecture at each end of the connection must match.

7.1 Unidirectional and bidirectional switching

In the case of bidirectional transmission, it is possible to choose either unidirectional or bidirectional switching. With unidirectional switching, the selectors at each end are fully independent. With bidirectional switching, an attempt is made to coordinate the two ends so that both have the same bridge and selector settings, even for a unidirectional failure. Bidirectional switching always requires an APS and/or PCC channel to coordinate the two endpoints. Unidirectional switching can protect two unidirectional failures in opposite directions on different entities.

7.2 Need for an APS/PCC channel

The only switching type that does NOT require an APS and/or PCC channel is 1+1 unidirectional switching. With a permanent bridge at the head end and no need to coordinate selector positions at the two ends, the tail end selector can be operated entirely according to defects and commands received at the tail end.

Bidirectional switching always requires an APS channel. 1:n unidirectional switching requires an APS channel to coordinate the head end bridge with the tail end selector.

7.3 Revertive and non-revertive switching

In revertive operation, traffic is restored to the working entities after a switch reason has cleared. In the case of clearing a command (e.g., Forced Switch), this happens immediately. In the case of clearing of a defect, this generally happens after the expiry of a "Wait-to-Restore" timer, which is used to avoid chattering of selectors in the case of intermittent defects.

In non-revertive operation, normal traffic is allowed to remain on the protection entity even after a switch reason has cleared. This is generally accomplished by replacing the previous switch request with a "Do not Revert (DNR)" request, which is low priority.

1+1 protection is often provisioned as non-revertive, as the protection is fully dedicated in any case, and this avoids a second "glitch" to the traffic. There may, however, be reasons to provision this to be revertive (e.g., so that the traffic uses the "short" direction around a ring except during failure conditions. Certain operator policies also dictate revertive operation even for 1+1).

Usually, 1:n protection is revertive. Certainly in the case where an extra traffic signal is carried on the protection entity, the operation would always be revertive so that the pre-empted extra traffic signal can be restored. It is certainly possible to define the protocol in a way that would permit non-revertive operation for 1:n protection, but the expectation is that it is better to revert and glitch the traffic when the working entity is repaired than when some other entity in the group fails that requires use of the protection entity to carry a different normal traffic signal.

In general, the choice of revertive/non-revertive will be the same at both ends of the protection group. However, a mismatch of this parameter does not prevent interworking – it just would be peculiar for one side to go to WTR for clearing of switches initiated from that side, while the other goes to DNR for its switches. See also 8.4.

7.4 Provisioning mismatches

With all of the options for provisioning of protection groups, there are opportunities for mismatches between the provisioning at the two ends. These provisioning mismatches take one of several forms:

- Mismatches where proper operation is not possible.
- Mismatches where one or both sides can adapt their operation to provide a degree of interworking in spite of the mismatch.
- Mismatches that do not prevent interworking. An example is the revertive/non-revertive mismatch discussed in 8.4.

Not all provisioning mismatches can be conveyed and detected by information passed through the APS channel. With a potential for up to 254 working entities in a 1:n protection group, there are simply too many combinations of valid entity numbers to easily provide full visibility of all of the configuration options. What is desirable, however, is to provide visibility for the middle category, where the sides can adapt their operation to interwork in spite of the mismatch. For example, an equipment provisioned for bidirectional switching could fall back to unidirectional switching to allow interworking. An equipment provisioned for 1+1 switching with an APS channel could fall back to operate in 1+1 unidirectional switching without an APS channel. The user could still be informed of the provisioning mismatch, but a level of protection could still be provided by the equipment.

8 APS Protocol

8.1 APS channel format

An APS channel is carried over the first three bytes of the APS/PCC field of the ODUk overhead. The fourth byte of the APS/PCC field is reserved. Eight independent APS channels are available to support protection at the ODUkP, the six ODUkT (TCM) levels and one level of ODUk SNC/I protection as defined in 15.8.2.4/G.709/Y.1331.

The format of the four APS bytes themselves within each frame is defined in Figure 1. The field values for the APS channels are defined in Table 1.

<table>
<tr><td colspan="8">1</td><td colspan="8">2</td><td colspan="8">3</td><td colspan="8">4</td></tr>
<tr><td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td><td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td><td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td><td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td></tr>
<tr><td colspan="4">Request/ state</td><td colspan="4">Protection type</td><td colspan="8">Requested Signal</td><td colspan="8">Bridged Signal</td><td colspan="8">Reserved</td></tr>
<tr><td colspan="4"></td><td>A</td><td>B</td><td>D</td><td>R</td><td colspan="8"></td><td colspan="8"></td><td colspan="8"></td></tr>
</table>

Figure 1/G.873.1 – APS channel format

Table 1/G.873.1 – Field values for APS channel

Field		Value	Description
Request/State		1111	Lockout of Protection (LO)
		1110	Forced Switch (FS)
		1100	Signal Fail (SF)
		1010	Signal Degrade (SD)
		1000	Manual Switch (MS)
		0110	Wait-to-Restore (WTR)
		0100	Exercise (EXER)
		0010	Reverse Request (RR)
		0001	Do Not Revert (DNR)
		0000	No Request (NR)
		Others	Reserved for future international standardization
Protection Type	A	0	No APS Channel
		1	APS Channel
	B	0	1+1 (Permanent Bridge)
		1	1:n (no Permanent Bridge)
	D	0	Unidirectional switching
		1	Bidirectional switching
	R	0	Non-revertive operation
		1	Revertive operation
Requested Signal		0	Null Signal
		1-254	Normal Traffic Signal 1-254
		255	Extra Traffic Signal
Bridged Signal		0	Null Signal
		1-254	Normal Traffic Signal 1-254
		255	Extra Traffic Signal

8.2 Transmission and acceptance of APS protocol

The APS/PCC protocol is transmitted via the protection entity. Although they may also be transmitted identically on working entities, receivers should not assume so, and should have the capability to ignore this information on the working entities.

For each of the eight levels, an independent acceptance process shall be performed. As the APS protocol is carried via the first three bytes of the four APS/PCC bytes, only these three bytes are taken into account for the acceptance process. A new APS protocol value shall be accepted if an identical value is received in these three bytes of a given level three times consecutively.

NOTE – Since the fourth byte of the APS message is 'reserved', it has not to be taken into account for the acceptance process of APS bytes.

8.3 Request type

The request types that may be reflected in the APS bytes are the "standard" types traditionally supported by protection switching for SONET and SDH. These requests reflect the highest priority condition, command, or state (see Tables 2 and 3). In the case of unidirectional switching, this is the highest priority value determined from the near end only. In bidirectional switching, the local request will be indicated only in the case where it is as high or higher than any request received from the far end over the APS channel. In bidirectional switching, when the far end request has the highest priority, the near end will signal Reverse Request.

Table 2/G.873.1 – Request/state priorities with APS protocol

Request/state	Priority
Lockout for Protection (LO)	1 (highest)
Signal Fail (SF) – protection	2 (see 8.9)
Forced Switch (FS)	3
Signal Fail (SF) – working	4
Signal Degrade (SD)	5
Manual Switch (MS)	6
Wait-to-Restore (WTR)	7
Exercise (EXER)	8
Reverse Request (RR)	9
Do Not Revert (DNR)	10
No Request (NR)	11 (lowest)

Table 3/G.873.1 – Request/state priorities without APS protocol

Request/state	Priority
Lockout for Protection (LO)	1 (highest)
Forced Switch (FS)	2
Signal Fail (SF)	3
Signal Degrade (SD)	4
Manual Switch (MS)	5
Wait-to-Restore (WTR)	6
Do Not Revert (DNR)	7
No Request (NR)	8 (lowest)

8.4 Protection types

The valid Protection Types are:

000x	1+1 Unidirectional, no APS
100x	1+1 Unidirectional w/APS
101x	1+1 Bidirectional w/APS
110x	1:n Unidirectional w/APS
111x	1:n Bidirectional w/APS

The values are chosen such that the default value (all zeros) matches the only type of protection that can operate without APS (1+1 Unidirectional).

Note that 010x, 001x, and 011x are invalid since 1:n and Bidirectional require APS.

If the "B" bit mismatches, the selector is released since 1:n and 1+1 are incompatible. This will result in an alarm.

Provided the "B" bit matches:

If the "A" bit mismatches, the side expecting APS will fall back to 1+1 unidirectional switching without APS.

NOTE 1 – In the case where a node does not support the APS channel, an all-0's pattern will be present in the APS/PCC field as specified in clause 15/G.709/Y.1331.

If the "D" bit mismatches, the bidirectional side will fall back to unidirectional switching.

If the "R" bit mismatches, one side will clear switches to "WTR" and the other will clear to "DNR". The two sides will interwork and the traffic is protected.

NOTE 2 – Each side signals always its maximum capabilities in the protection type field even if it falls back to operate with less capabilities (i.e., a side which supports bidirectional switching falls back to operate as unidirectional switch in case of interworking with a side that supports unidirectional switching only, but still signals "1" in the "D" bit).

NOTE 3 – Reporting of mismatch conditions is for further study.

8.5 Requested signal

This indicates the signal that the near end requests to be carried over the protection entity. For NR, this is either the Null Signal (0) or Extra Traffic (255). For LO, this can only be the Null Signal (0). For Exercise, this can be the Null Signal (0) or the Extra Traffic Signal (255) when Exercise replaces NR, or the number of a normal traffic signal in the case where Exercise replaces DNR. For SF or SD, this will be the number of a normal traffic signal, or the Null Signal (0) to indicate that protection is failed or degraded. For all other requests, this will be the number of the normal traffic signal requested to be carried over the protection entity.

8.6 Bridged signal

This indicates the signal that is bridged onto the protection entity. For 1+1 protection, this should always indicate normal traffic signal 1, accurately reflecting the permanent bridge. This allows a 2-phase rather than a 3-phase switch in the case of 1+1 architecture. For 1:n protection, this will indicate what is actually bridged to the protection entity (either the Null Signal (0), Extra Traffic (255), or the number of a normal traffic signal). This will generally be the bridge requested by the far end.

8.7 Control of bridge

In 1+1 architectures, the normal traffic signal is permanently bridged to protection. The normal traffic signal number "1" will always be indicated in the bridged signal field of the APS channel.

In 1:n architectures, the bridge will be set to the one indicated by the "Requested Signal" field of the incoming APS channel. Once the bridge has been established, this will be indicated in the "Bridged Signal" field of the outgoing APS channel.

8.8 Control of selector

In 1+1 Unidirectional architectures (with or without APS), the selector is set entirely according to the highest priority local request. This is a single-phase switch.

In 1+1 Bidirectional architectures, the normal traffic signal will be selected from the protection entity when the outgoing "Requested Signal" and the incoming "Bridged Signal" both indicate normal traffic signal "1" (The incoming "Bridged Signal" should always indicate "1" in this architecture). This is a two-phase switch, as the far end does not switch until the APS bytes indicating that a bidirectional switch is initiated by the near end arrives.

In 1:n uni- or bidirectional architectures, a normal traffic signal "n" or extra traffic signal 255 will be selected from the protection entity when the same number "n" (or 255) appears in both the outgoing "Requested Signal" and the incoming "Bridged Signal" fields. This generally results in a three-phase switch.

8.9 Signal Fail of the protection entity

Signal Fail on the protection entity is higher priority than any defect that would cause a normal transport signal to be selected from the protection entity. For the case an APS signal is in use, a SF on the protection entity (over which the APS signal is routed) has priority over the Forced Switch. A Lockout command has higher priority than SF. During failure conditions, lockout status shall be kept active.

8.10 Equal priority requests

In general, once a switch has been completed due to a request, it will not be overridden by another request of the same priority (first come, first served behaviour). When equal priority requests occur simultaneously, the conflict is resolved in favour of the request with the lowest entity number. In bidirectional switching, a request received over the APS channel with a lower entity number will always override an identical priority local request with a higher entity number. Equal priority requests for the same entity number from both sides of a bidirectional protection group are both considered valid, and equivalent to a received "RR" from a near-end processing standpoint.

8.11 Command acceptance and retention

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

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

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

If a command is overridden by a condition or (in bidirectional switching) APS request, that command is forgotten.

8.12 Hold-off timer

In order to coordinate timing of protection switches at multiple layers or across cascaded protection domains, a hold-off timer may be required. The purpose is to allow either a server layer protection switch to have a chance to fix the problem before switching at a client layer, or to allow an upstream protection domain to switch before a downstream domain (e.g., to allow an upstream ring to switch before the downstream ring in a dual node interconnect configuration so that the switch occurs in the same ring as the failure).

Each protection group should have a provisionable hold-off timer. The suggested range and values are 0, 20 ms, and 100 ms to 10 seconds in steps of 100 ms (accuracy of ±5 ms as per ITU-T Rec. G.808.1).

The operation of the hold-off timer uses the "peek twice" method specified in SDH standards. Specifically, when a new defect or more severe defect occurs (new SD or SF, or SD becoming SF), this event will not be reported immediately to protection switching if the provisioned hold-off timer value is non-zero. Instead, the hold-off timer will be started. When the hold-off timer expires, it will be checked whether a defect still exists on the trail that started the timer. If it does, that defect will be reported to protection switching. The defect need not be the same one that started the timer.

8.13 Exercise operation

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

Exercise command shall issue the command with the same requested and bridged entity numbers of the NR or DNR request that it replaces. The valid response will be an RR with the corresponding requested and bridged entity numbers. To allow the RR to be detected, the standard response to DNR should be DNR rather than RR. When the exercise command is cleared, it will be replaced with NR if the requested entity number is 0 or 255, and DNR for any normal traffic signal number 1 to 254.

NOTE – Exercise operation for OTN has been defined differently from exercise operation defined for SDH.

8.14 APS channel alarming

"Failure of Protocol" situations for groups requiring APS are as follows:

– Fully incompatible provisioning (the "B" bit mismatch), described in 8.4.

– Lack of response to a bridge request (i.e., no match in sent "Requested Entity" and received "Bridged Entity") for > 50 ms.

If an unknown request or a request for an invalid entity number is received, it will be ignored. It will be up to the far end to alarm the non-response from the near end.

Appendix I

Examples of operation

I.1 1+1 Unidirectional switching

APS may or may not be present. Even if APS is not present, the bridge is assumed to be permanent, so switches are performed immediately according to the local request. The APS bytes, if present, are informational only and do not control the operation of the protection group. If they are present, an equipment may allow a query for the far end state.

This example shows overlapping SF and SD requests from opposite sides. For illustration, the example in Figure I.1 shows mismatched provisioning with side A being non-revertive and side B being revertive.

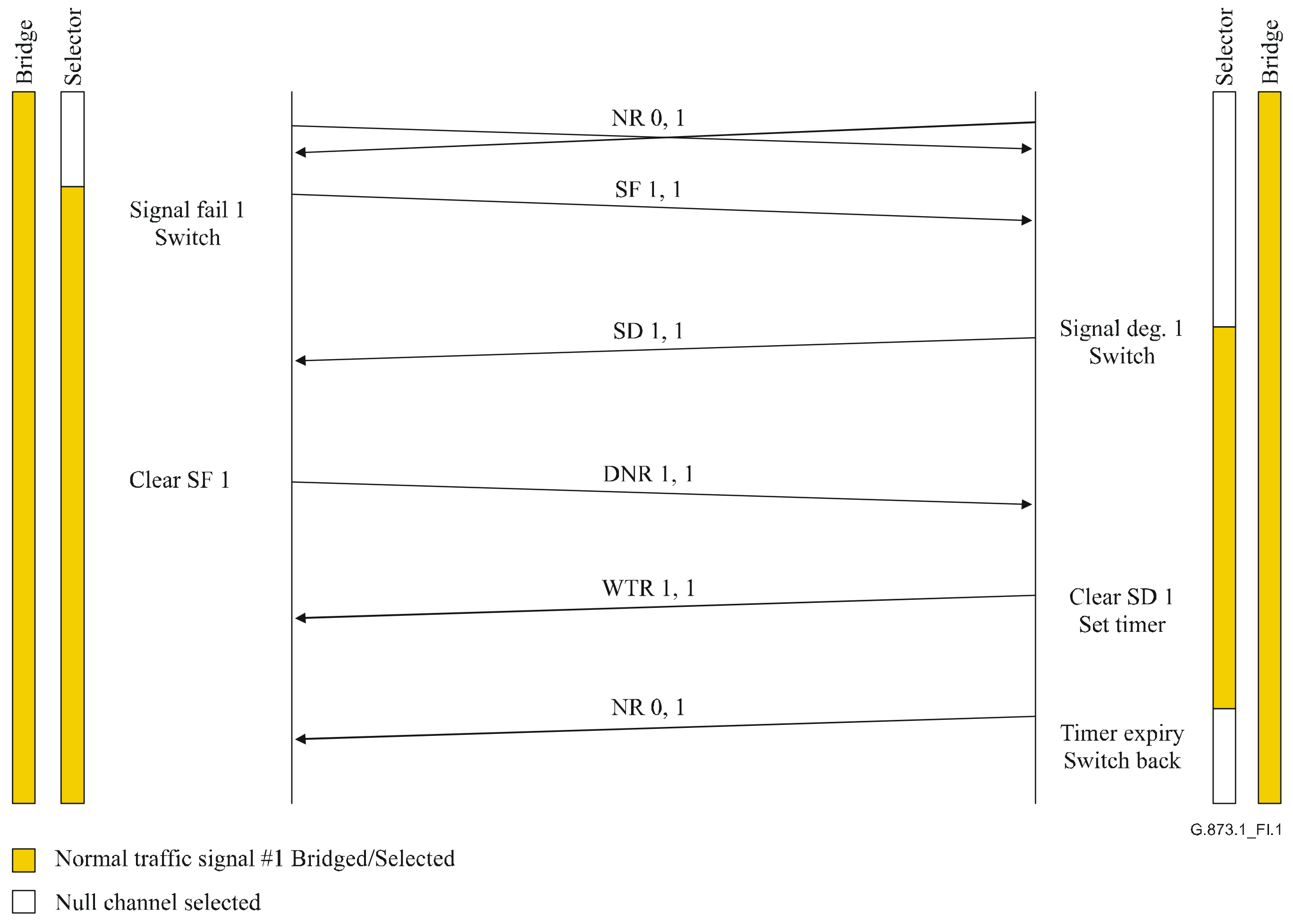

Figure I.1/G.873.1 – Example APS message flow for 1+1 unidirectional switching

I.2 1+1 Bidirectional switching

The example in Figure I.2 illustrates a 1+1, bidirectional, non-revertive switch. Because the permanent bridge is indicated in the APS bytes from the start, the switch can be two-phase instead of three-phase.

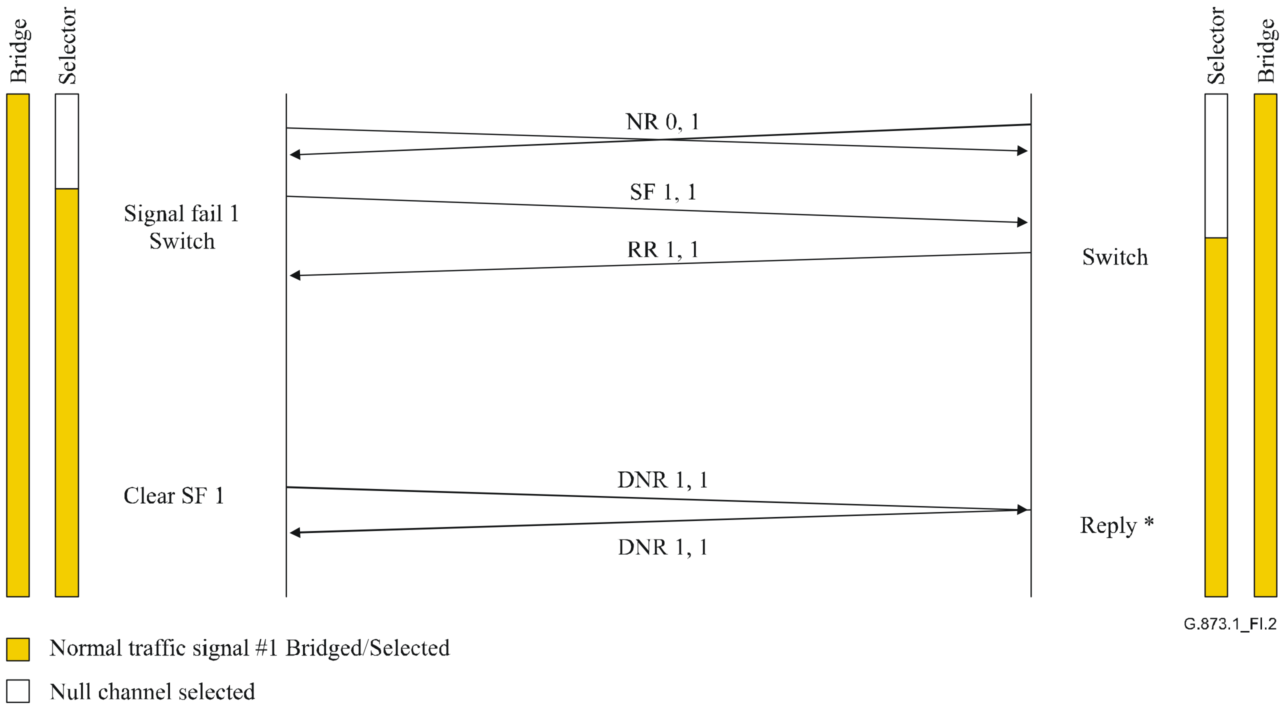

Figure I.2/G.873.1 – Example APS message flow for 1+1 bidirectional switching

NOTE – Historically, DNR was acknowledged with RR. Here, answering DNR with DNR makes no fundamental difference in the states of the two sides, and it allows for a meaningful exercise implementation.

I.3 1:n Bidirectional switching

Figure I.3 shows an example of 1:n Bidirectional switching with Extra Traffic. What is illustrated is the case where an SD on working #2 is pre-empted by an SF on working #3.

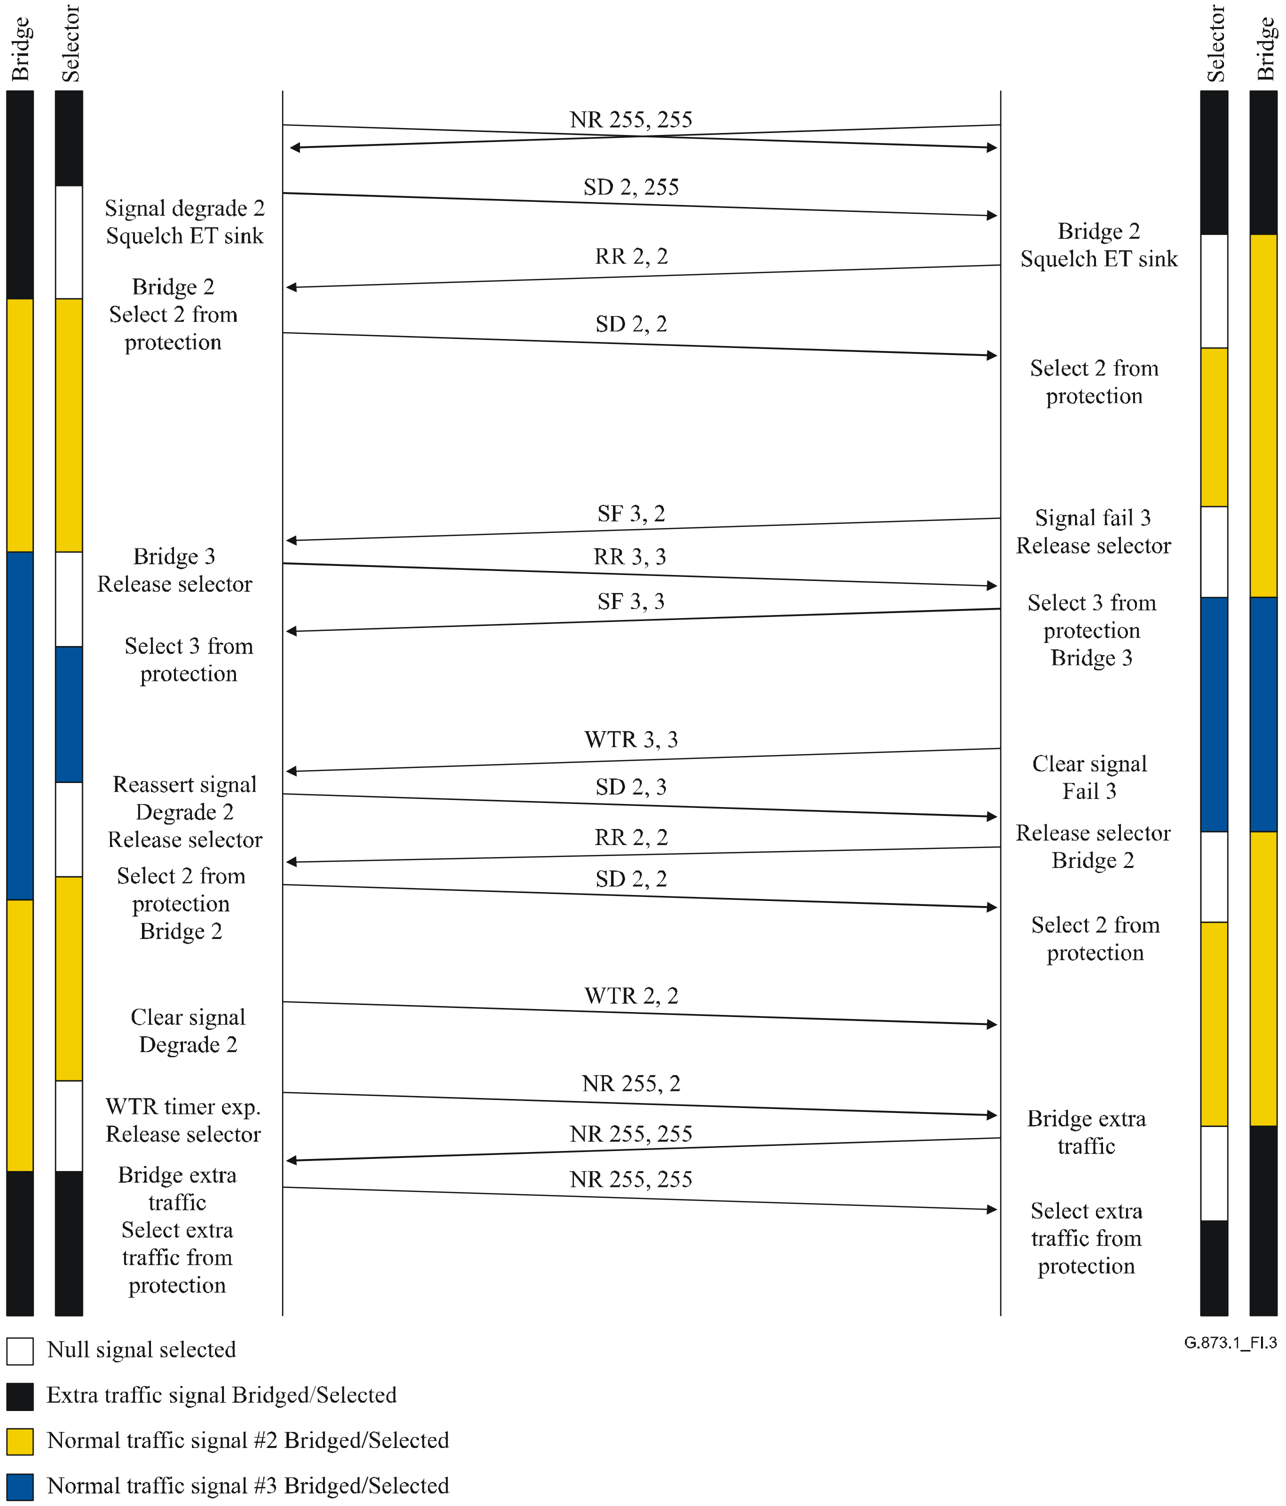

Figure I.3/G.873.1 – Example APS message flow for 1:n bidirectional switching

I.4 Exercise command operation

The Exercise command is a test that the far end will respond to an APS channel request in bidirectional switching without operating the selector. This command is low priority so as not to interfere with the actual operation of the protection group. It is only valid when the current request is NR or DNR, as it is lower in priority than all other requests.

Figures I.4, I.5, I.6 and I.7 give examples of operation of the Exercise command. In all cases, neither the requested nor bridged entity numbers are changed for the Exercise command. A successful response is receiving an "RR" with the same entity number. Note that having DNR answered with DNR provides a way to test that the Exercise command receives the appropriate RR response.

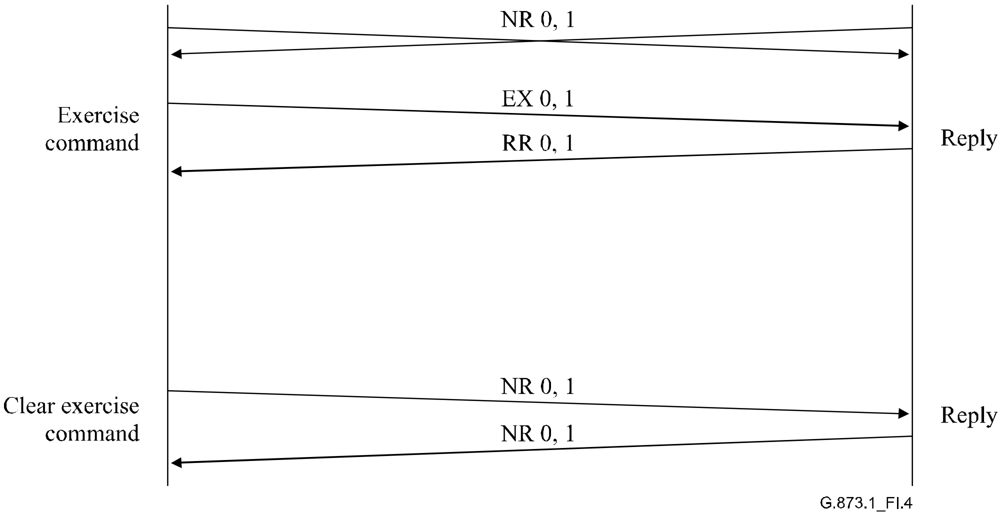

Figure I.4/G.873.1 – Example of exercise command from 1+1 NR state

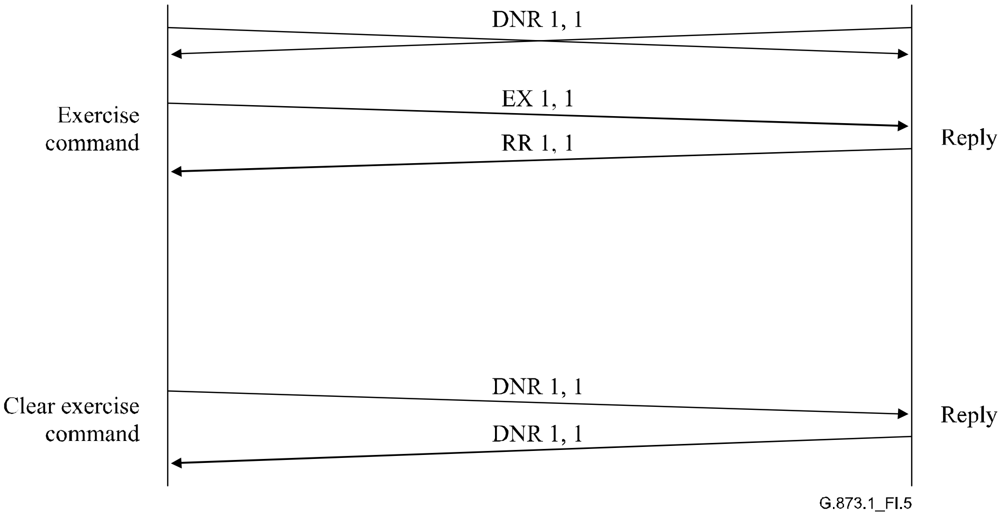

Figure I.5/G.873.1 – Example of exercise command from 1+1 DNR state

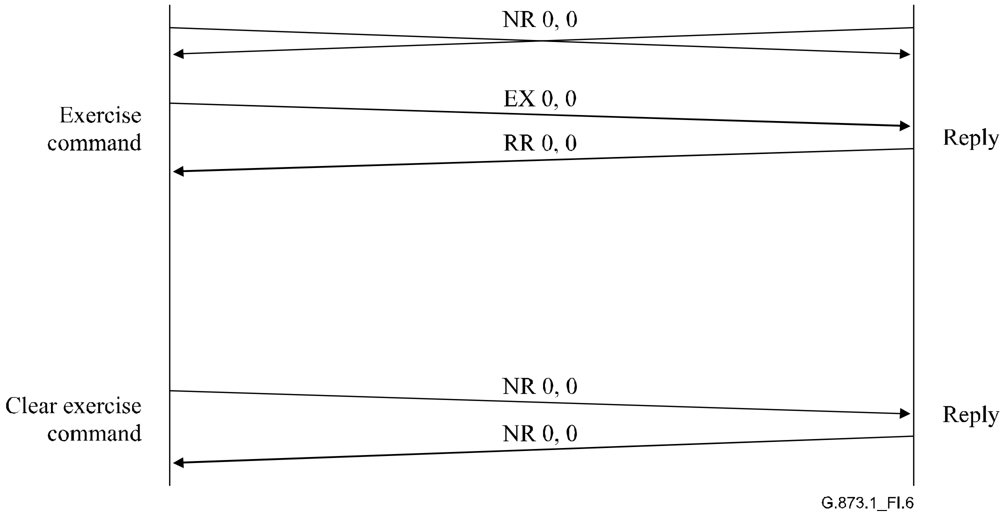

Figure I.6/G.873.1 – Example of exercise command from 1:n NR state without extra traffic

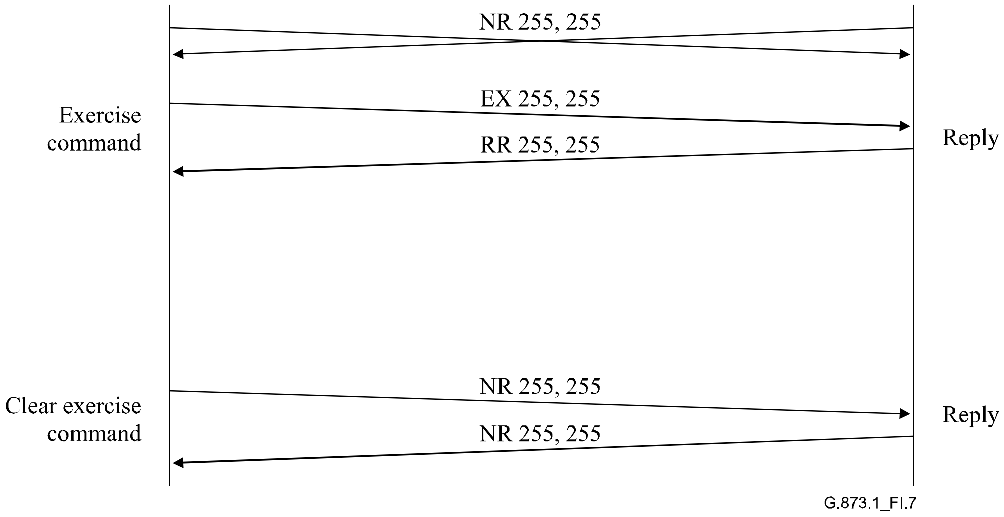

Figure I.7/G.873.1 – Example of exercise command from 1:n NR state with extra traffic

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
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	Telecommunication management, including TMN and network maintenance
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, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2006