



Core Information Model (CoreModel)

TR-512.A.11 Appendix – Resilience Examples

Version 1.4
November 2018

ONF Document Type: Technical Recommendation
ONF Document Name: Core Information Model version 1.4

Disclaimer

THIS SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Any marks and brands contained herein are the property of their respective owners.

Open Networking Foundation
1000 El Camino Real, Suite 100, Menlo Park, CA 94025
www.opennetworking.org

©2018 Open Networking Foundation. All rights reserved.

Open Networking Foundation, the ONF symbol, and OpenFlow are registered trademarks of the Open Networking Foundation, in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Important note

This Technical Recommendations has been approved by the Project TST, but has not been approved by the ONF board. This Technical Recommendation is an update to a previously released TR specification, but it has been approved under the ONF publishing guidelines for 'Informational' publications that allow Project technical steering teams (TSTs) to authorize publication of Informational documents. The designation of '-info' at the end of the document ID also reflects that the project team (not the ONF board) approved this TR.

.

Table of Contents

Disclaimer	2
Important note	2
Document History	6
1 Introduction	7
1.1 References.....	7
1.2 Definitions	7
1.3 Conventions	7
1.4 Viewing UML diagrams.....	7
1.5 Understanding the figures.....	7
1.6 Appendix Overview	7
2 Introduction to this Appendix document.....	8
2.1 Key to diagrams.....	8
3 Linear protection schemes.....	9
3.1 1?1 cases.....	9
3.2 1?1 open protection cases.....	17
3.3 1:N Cases	21
4 Mesh Network cases	25
4.1 N:1 with multicast nodal Cases.....	25
5 Ethernet Ring Protection [ITU-T G.8032]	26
5.1 The protection scheme	26
5.2 Relevant pieces of the resilience model for [ITU-T G.8032].....	32
5.3 Using the spec model to explain the alternative raw and encapsulation forms.....	33
5.4 Representing the block	36
5.5 Forming the ring.....	36
6 Other protected ring schemes	44
6.1 Network Wrapping	46
6.1.1 The scheme	46
6.1.2 The model applied	46
6.2 Network Steering	51
6.2.1 The scheme	51
6.2.2 The model applied	52
6.3 The model in detail for both Steering and Wrapping	57

List of Figures

Figure 2-1 Instance diagram key	8
Figure 3-1 Simple summary example of 1?1 cases (represented via partition)	9
Figure 3-2 Showing detail of a single ended view of 1+1 and 1:1 switches	10
Figure 3-3 Showing an emergent abstract controller in a 1:1 case	11
Figure 3-4 Showing a basic route based representation of the 1?1 protection scheme	12
Figure 3-5 Showing a two level route based representation of the 1?1 protection scheme	13
Figure 3-6 Showing the preferred route based representation of the 1?1 protection scheme	13
Figure 3-7 Route based representation of the 1?1 protection scheme showing C&SCs	14
Figure 3-8 Showing detail of a single ended view of 1+1 and 1:1 switches in a route context	15
Figure 3-9 Single ended view of 1:1 switches in a route context with peer C&SC coordination	16
Figure 3-10 Nodal controller peering in a route context.....	17
Figure 3-11 Simple summary example of open 1?1 cases.....	17
Figure 3-12 Showing an emergent abstract controller in an open 1?1	18
Figure 3-13 Simple summary example of open 1?1 cases showing route approach.....	19
Figure 3-14 Single ended view of open 1:1 switches in a route context with peer C&SC coordination	20
Figure 3-15 Nodal controller peering in a route context.....	20
Figure 3-16 Simple summary example of 1:N cases (represented via partition)	21
Figure 3-17 Showing detail of a single ended view of 1:N line system.....	22
Figure 3-18 Showing an emergent abstract controller in a 1:N case	23
Figure 3-19 Showing route based representation of the 1:N protection scheme	24
Figure 3-20 Showing Extra Traffic in a route based representation of the 1:N protection scheme	24
Figure 3-21 Showing independent two ended view of W1 route detail in a 1:N protection scheme.....	25
Figure 4-1 Showing a unidirectional N:1 scheme fragment.....	26
Figure 5-1 [ITU-T G.8032] Ring node control and signaling	27
Figure 5-2 [ITU-T G.8032] Rings showing traffic flow under normal and failure conditions (1)	28
Figure 5-3 [ITU-T G.8032] Rings showing traffic flow under normal and failure conditions (2)	28
Figure 5-4 Basic model for a ring node	29
Figure 5-5 Detailed node model focusing on signal flow	30
Figure 5-6 Detailed view of a control function from [ITU-T G.8032]	31
Figure 5-7 Ring node represented by C&SC encapsulations	32
Figure 5-8 Resilience model structure for G.8032	33
Figure 5-9 View of NE actively participating in two rings showing spec and encapsulation	33

Figure 5-10 Relationship between the two instance views shown via their related spec	35
Figure 5-11 Applying the "blocks" to the ring	36
Figure 5-12 The basic [ITU-T G.8032] ring	37
Figure 5-13 Two overlaid [ITU-T G.8032] Major Rings showing signaling only.....	38
Figure 5-14 Basic [ITU-T G.8032] Sub-Ring.....	38
Figure 5-15 Major Ring showing Traffic	39
Figure 5-16 Basic [ITU-T G.8032] Major Ring and Sub-Ring	40
Figure 5-17 [ITU-T G.8032] Major Ring and Sub-Ring showing traffic.....	40
Figure 5-18 [ITU-T G.8032] Major Ring and Sub-Ring showing only traffic	41
Figure 5-19 [ITU-T G.8032] Major Ring and Sub-Ring showing traffic with zero length link	41
Figure 5-20 [ITU-T G.8032] Major Ring & Sub-Ring showing only traffic with zero length link	42
Figure 5-21 [ITU-T G.8032] Major Rings in a mesh.....	42
Figure 5-22 MEPs and R-APS insertion [ITU-T G.8032]	43
Figure 5-23 MEPs and R-APS insertion without R-APS virtual Channel [ITU-T G.8032].....	44
Figure 6-1 Diagram key.....	45
Figure 6-2 The network	45
Figure 6-3 The network showing wrapping	46
Figure 6-4 Wrapping: NE X and NE Y (no failure in ring)	47
Figure 6-5 Wrapping: NE Z (no failure in ring)	48
Figure 6-6 Wrapping: NE Y with failure on port 2	49
Figure 6-7 Wrapping: NE X with failure on NE Y port 2	50
Figure 6-8 Wrapping: NE Z with failure on NE Y port 2	51
Figure 6-9 Network showing steering	51
Figure 6-10 Steering: NE Z (no failure in ring)	52
Figure 6-11 Steering: NE X (no failure in ring).....	53
Figure 6-12 Steering: NE Y (no failure in ring).....	53
Figure 6-13 Steering: NE W (no failure in ring).....	54
Figure 6-14 Steering: NE Z with failure on NE Y port 2	54
Figure 6-15 Steering: NE Y with failure on port 2 (same as no failure)	55
Figure 6-16 Steering: NE X with failure on NE Y port 2.....	56
Figure 6-17 Steering: NE W with failure on NE Y port 2.....	56

Document History

Version	Date	Description of Change
		Appendix material was not published prior to Version 1.3
1.3	September 2017	Version 1.3 [Published via wiki only]
1.3.1	January 2018	Addition of text related to approval status.
1.4	November 2018	No change.

1 Introduction

This document is an appendix of the addendum to the TR-512 ONF Core Information Model and forms part of the description of the ONF-CIM. For general overview material and references to the other parts refer to [TR-512.1](#).

1.1 References

For a full list of references see [TR-512.1](#).

1.2 Definitions

For a full list of definition see [TR-512.1](#).

1.3 Conventions

See [TR-512.1](#) for an explanation of:

- UML conventions
- Lifecycle Stereotypes
- Diagram symbol set

1.4 Viewing UML diagrams

Some of the UML diagrams are very dense. To view them either zoom (sometimes to 400%) or open the associated image file (and zoom appropriately) or open the corresponding UML diagram via Papyrus (for each figure with a UML diagram the UML model diagram name is provided under the figure or within the figure).

1.5 Understanding the figures

Figures showing fragments of the model using standard UML symbols and also figures illustrating application of the model are provided throughout this document. Many of the application-oriented figures also provide UML class diagrams for the corresponding model fragments (see [TR-512.1](#) for diagram symbol sets). All UML diagrams depict a subset of the relationships between the classes, such as inheritance (i.e. specialization), association relationships (such as aggregation and composition), and conditional features or capabilities. Some UML diagrams also show further details of the individual classes, such as their attributes and the data types used by the attributes.

1.6 Appendix Overview

This document is part of the Appendix to TR-512. An overview of the Appendix is provided in [TR-512.A.1](#).

3 Linear protection schemes²

3.1 1?1 cases

This section deals with basic 1+1 and 1:1 cases and shows how they can be represented. The abbreviation 1?1 has been used where the description is common between both 1+1 and 1:1.




Figure 3-1 Simple summary example of 1?1 cases (represented via partition)

The figure above shows a simple summary example of a 1?1 case in a basic network with three NEs. Clearly this can be generalized further to be in a rule form. A specific solution can include zero or more NEs on either path³. The end-end FC is partitioned into subordinate (i.e. is an aggregation of the subordinate parts via FcHasLowerLevelFcs). The scheme may involve signalling.

² The examples in this section were in TR-512.5 V1.2.

³ Note that there is work in progress to develop scheme specs that will provide a rule based view of the scheme.

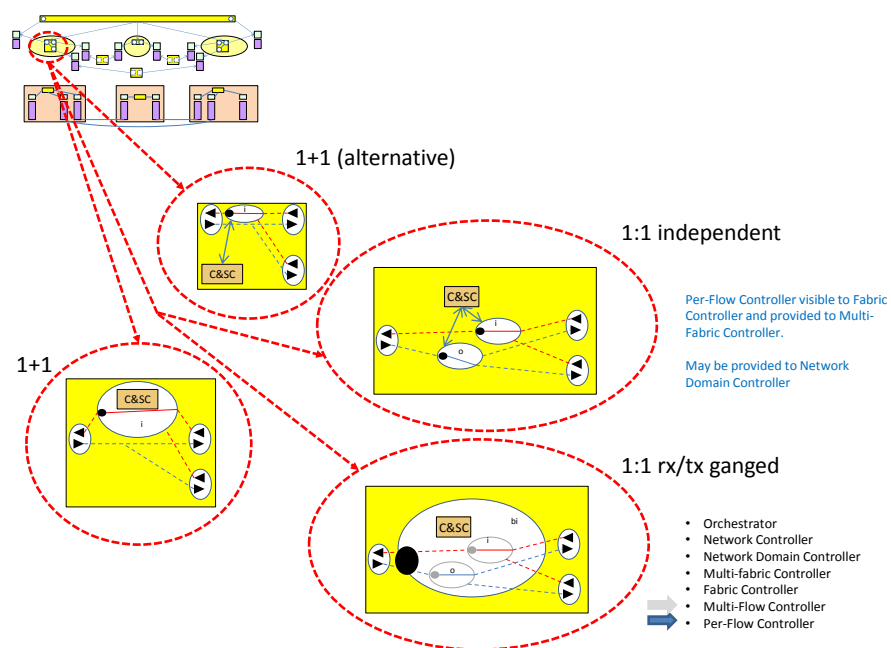


Figure 3-2 Showing detail of a single ended view of 1+1 and 1:1 switches

The figure above shows a nodal view and highlights ConfigurationAndSwitchControllers (C&SC) encapsulated in the FcSwitch in some cases and in FC in others. The encapsulation chosen depends upon the scope of control of the C&SC. The encapsulation is via FcSwitchCoordinatedByInternalControl when in the FcSwitch and FcSwitchesInFcCoordinatedBySwitchCoordinator when in the FC.

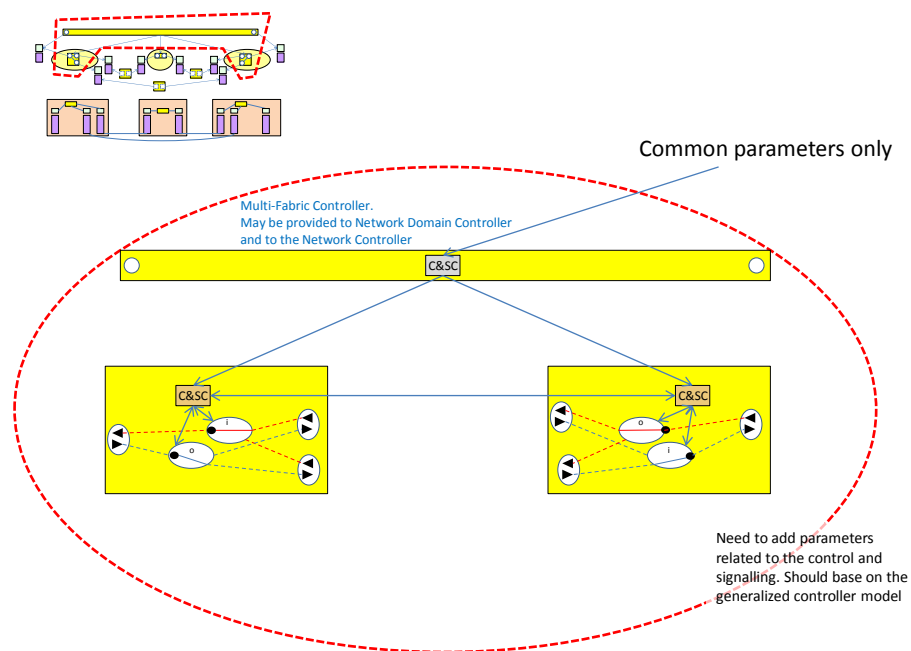


Figure 3-3 Showing an emergent abstract controller in a 1:1 case

The figure above shows a case of 1:1 independent switching (where the two directions of traffic are switched independently). The figure assumes that there is a distributed control solution (where the C&SCs in the FCs signal each other) and highlights an emergent C&SC which does not actually exist in the real control solution but which can be expressed to collect together parameters that should be set to the same value at both ends. In the network the coordination occurs through peer signaling. Above the network the SDN controller may realize the coordination⁴.

⁴ This recognition of levels of control from the most basis local two state switch controller through the various levels shown here and on two ring controllers and the SDN controller peer-hierarchy is a manifestation of and a validation of the concept of the Management-Control Continuum. Representation of the Management-Control Continuum will be further explored in the next release.

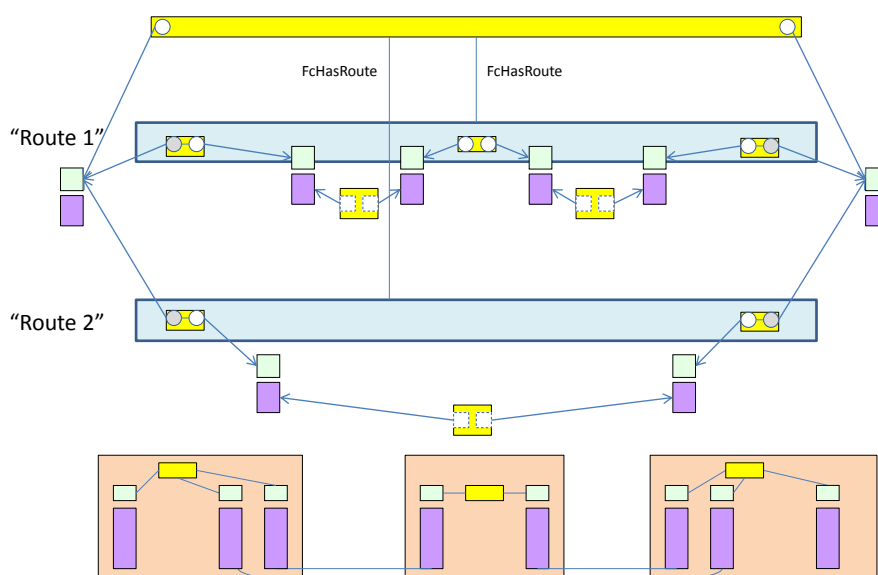


Figure 3-4 Showing a basic route based representation of the 1?1 protection scheme

The figure above shows an alternative representation of the 1?1 to that shown in Figure 3-1 Simple summary example of 1?1 cases on page 9. In the representation above two FcRoutes are used to represent the two alternative flows across the network. It should be noted that the FCs at the ends of each route are associated with the same LTPs and are only not conflicting because of the switches that they encapsulate (which when appropriately coordinated can ensure that only one FC is feeding the LTP at any time). The FcPort that can be switched off, i.e. be open, to ensure conflict can be avoided are depicted in grey (an output FcPort that can be switched off can share an LTP with another similar output FcPort and hence is called a sharing FcPort in this document). The FcSpec would identify the port via the switch configuration definition.

The figure below shows an alternative, slightly more verbose, representation of the 1?1 protection using two levels of route whether the top level routes have FCs that have the same span and the end-end FC⁵

The FCs of the route are contained in the route via the RouteIsDescribedByFc composition association and hence are not members of an FD. The FCs are used to represent flow and are defined in terms of the LTPs they reference in the context of the Route. The FD if visible would still have the FCs as shown in Figure 3-1 Simple summary example of 1?1 cases.

⁵ This pattern of “decomposition” of the FC into two parallel FCs is also used when the FC is representing a Control Plane Call and when there is a need to combine two unidirectional FCs into a bidirectional FC. In these cases the decomposition takes place via the FcHasLowerLevelFcs association and the FCs are members of an FD via the FdContainsFc association.

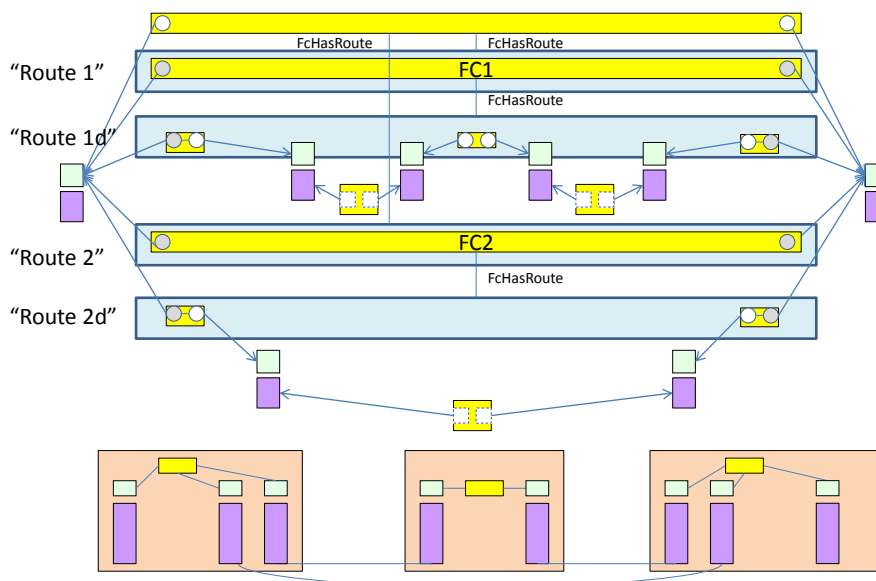


Figure 3-5 Showing a two level route based representation of the 1?1 protection scheme

The figure below shows the preferred route based representation which is a hybrid of the two where the FC of a route is described in terms of FCs via the `FcHasLowerLevelFc` such that the lower level (nodal) FCs are in the context of FDs via the `FdContainsFc` aggregation (a usual partition).

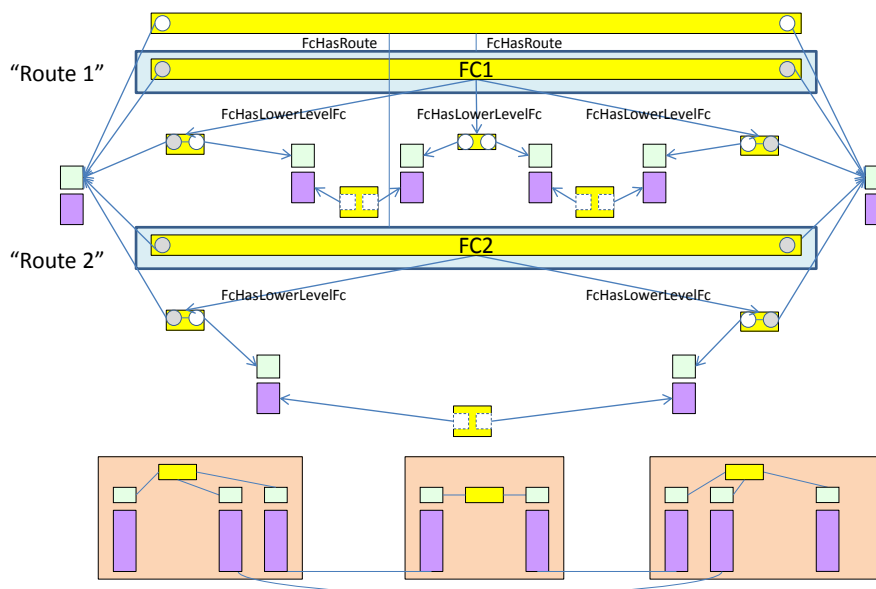


Figure 3-6 Showing the preferred route based representation of the 1?1 protection scheme

This approach for the 1?1 case, which may involve signalling, with a decomposition then partition is used in following diagrams.

The figure below shows the ConfigurationAndSwitchController (C&SC) positions and their associations (ControllerGovernsSubordinateController). The figure shows a number of potential emergent controllers as well as some real controllers assuming a distributed control scheme.

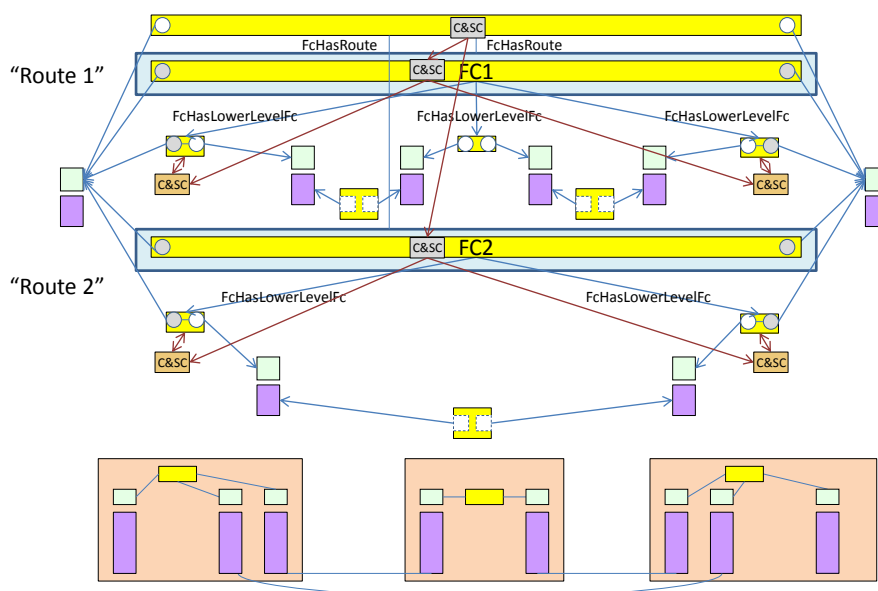


Figure 3-7 Route based representation of the 1?1 protection scheme showing C&SCs

The figure below shows a nodal view for one route and highlights ConfigurationAndSwitchControllers (C&SC) encapsulated in the FcSwitch in some cases and in FC in others. The encapsulation chosen depends upon the scope of control of the C&SC. The encapsulation is via FcSwitchCoordinatedByInternalControl when in the FcSwitch and FcSwitchesInFcCoordinatedBySwitchCoordinator when in the FC.

Some of the diagrams in the figure below (in dotted red ellipses) use a mixture of output and input switches (designated by "o" and "i" respectively).

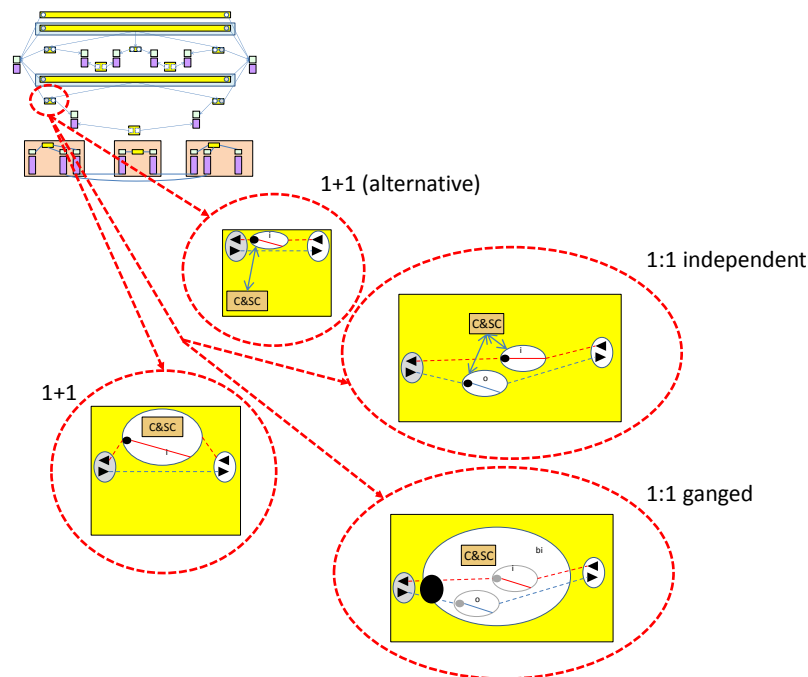


Figure 3-8 Showing detail of a single ended view of 1+1 and 1:1 switches in a route context

The figure below shows the interaction between the C&SCs of the FCs of the two routes. The interaction is via a balanced dual form of the `ControllerGovernsController`⁶ association used to indicate a peer relationship. Setting values for one controller will affect the values in the peer.

Rules for the effect need to be stated in the spec. If aspects of the peering can be disabled this would lead to attributes to control those aspects.

⁶ Note that this association is Experimental

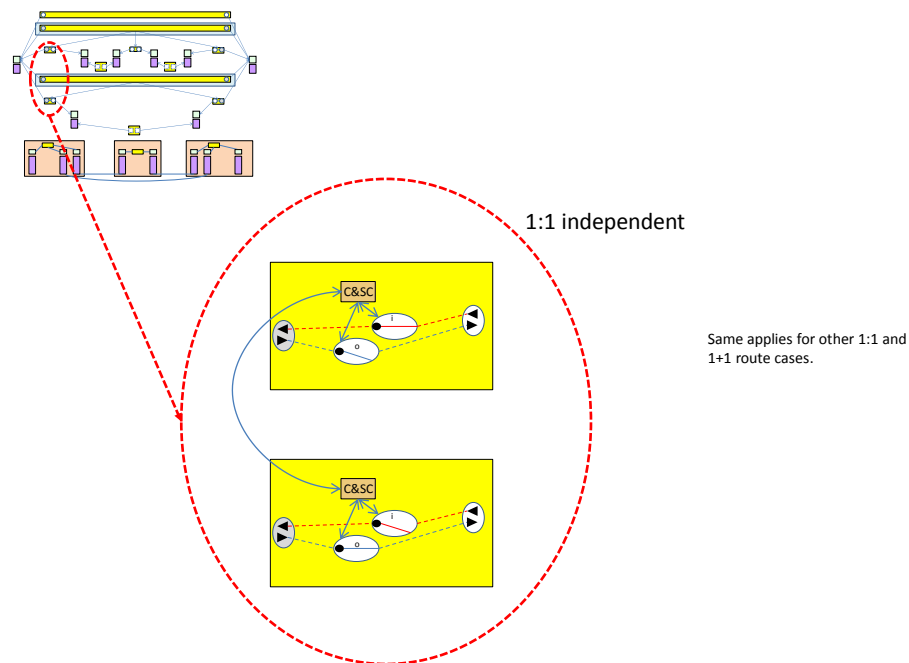


Figure 3-9 Single ended view of 1:1 switches in a route context with peer C&SC coordination

The figure below shows the controller peering between routes (emergent as the scheme is assumed to be a distributed control scheme) and also the emergent control in the end-end FC. It is proposed that the orientation convention is that input switch is preferred when ambiguous.




Figure 3-10 Nodal controller peering in a route context

3.2 1?1 open protection cases

The figures in this section are similar to those in the previous section.




Figure 3-11 Simple summary example of open 1?1 cases

The figure above shows a simple summary example of an open 1?1 case (e.g. where only one end of the recovery scheme is within the scope of the SDN controller) in a basic network with three NEs. Clearly this can be generalized further to be in a rule form.

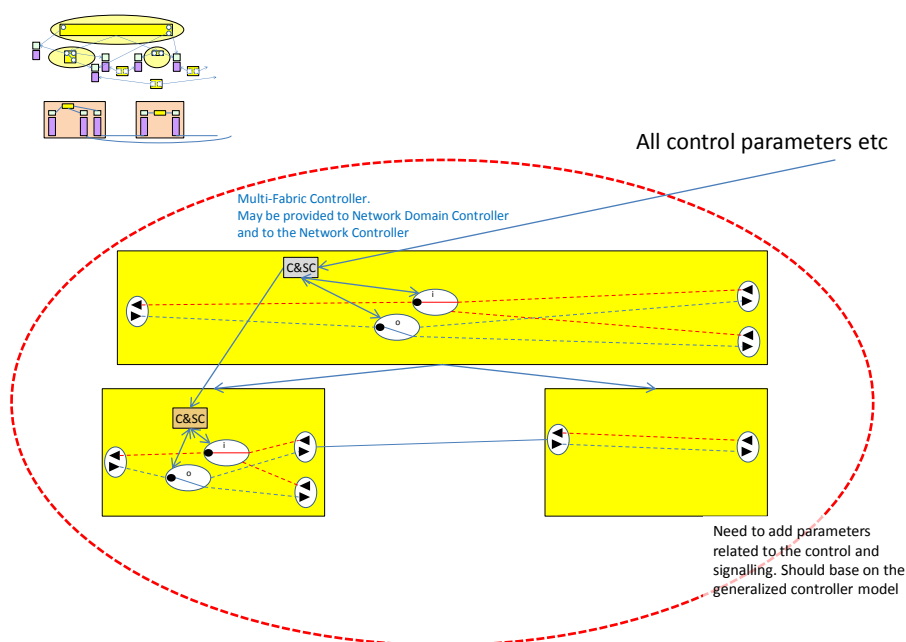


Figure 3-12 Showing an emergent abstract controller in an open 1?1

The figure above shows a case of 1:1 independent switching (where the two directions of traffic are switched independently). The figure assumes that there is a distributed control solution (where the C&SCs in the FCs signal each other) and highlights an emergent C&SC which does not actually exist in the real control solution but which can be expressed to collect together parameters that should be set to the same value at both ends. In the network the coordination occurs through peer signaling where the peer signaling is between C&SCs one of which is outside this view. Above the network the SDN controller may realize the coordination but to do this it will itself need to have communication with network peers (SDN controllers or other management-control entities) that control the off-network end(s) of the protection scheme.




Figure 3-13 Simple summary example of open 1?1 cases showing route approach

The figure below shows the preferred route based representation which is a hybrid of the two where the FC of a route is described in terms of FCs via the `FcHasLowerLevelFcs` such that the lower level (nodal) FCs are in the context of FDs via the `FdContainsFcs` aggregation (a usual partition).

The figure below shows the interaction between the C&SCs of the FCs of the two routes. The interaction is the same as discussed earlier for Figure 3-9 Single ended view of 1:1 switches in a route context with peer C&SC coordination on page 16.




Figure 3-14 Single ended view of open 1:1 switches in a route context with peer C&SC coordination

The figure below shows the controller peering between routes (emergent as the scheme is assumed to be a distributed control scheme) and also the emergent control in the end-end FC.




Figure 3-15 Nodal controller peering in a route context

3.3 1:N Cases

This section deals with basic 1:N cases and shows how they can be represented.

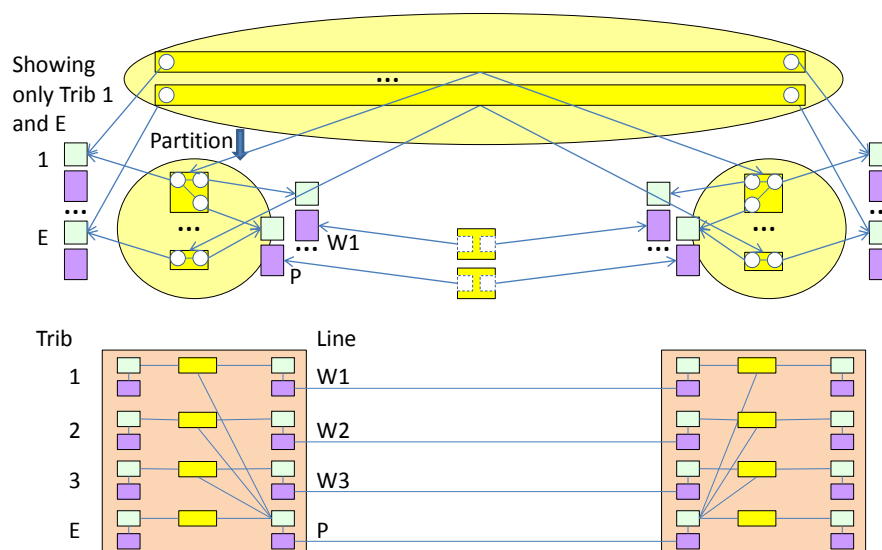


Figure 3-16 Simple summary example of 1:N cases (represented via partition)

The figure above shows a simple summary example of a 1:N case in a basic network. As shown in the detailed NE view at the bottom of the figure, the scheme provides protection to three traffic signals (1, 2 and 3) and also provides a lower grade path for "Extra Traffic" (E). The traffic signals 1, 2 and 3 normally each use a dedicated "Worker"⁷ paths (W1-W3 (numbered to match the traffic signal numbers)). The Protection path (P) provides an alternative for any one of the Workers. The "Extra Traffic", E, uses the protection path, P, when it is not needed to protect any of W1-W3. Clearly this can be generalized further to be in a rule form. A specific solution can include one more traffic paths.

The FC view shows only one of the traffic paths (1) and the "Extra Traffic" (E). The end-end FC representing the traffic path is partitioned into subordinate (i.e. is an aggregation of the subordinate parts via FcHasLowerLevelFcs) as is the "Extra Traffic" path (E). It should be noted that the nodal FC from E to P and the FC from 1 to W1 and P use the same LTP at P. The apparent conflict is resolved by the C&SC (not shown). The scheme will involve signalling.

⁷ The term "Worker" means normal path for particular traffic




Figure 3-17 Showing detail of a single ended view of 1:N line system

The figure above shows a nodal view and highlights ConfiguraionAndSwitchControllers (C&SC) encapsulated in the FcSwitch in some cases and in FC in others. The encapsulation chosen depends upon the scope of control of the C&SC. The encapsulation is via FcSwitchCoordinatedByInternalControl when in the FcSwitch and FcSwitchesInFcCoordinatedBySwitchCoordinator when in the FC.

In the case of 1:N with Extra Traffic it is necessary for the switch of the Extra Traffic to be coordinated with the switches for protection of the main traffic (and likewise for the switches of each of the main traffic signals to be coordinated). It is assumed here that there is a real C&SC that carries out that coordination. The C&SCs encapsulated in the FCs/FcSwitches are assumed subordinate and hence the ControllerGovernsController association is one way from the independent C&SC to the C&SCs in the FCs/FcSwitches.




Figure 3-18 Showing an emergent abstract controller in a 1:N case

The figure above shows a case of 1:N independent switching (where the two directions of traffic are switched independently). The figure assumes that there is a distributed control solution (where the C&SCs in the FCs signal each other) and highlights the emergent C&SCs (not all controllers are relevant for control and control may be scattered across the controllers⁸).

The abstract C&SCs can be expressed to collect together parameters that should be set to the same value at both ends or to some other complementary values for competing switches at the same end. In the network the coordination occurs through peer signaling. Above the network the SDN controller may realize the coordination⁹.

In the figure above the Extra Traffic has been switched off although only one direction of the Protection route is being used. It is assumed here that the Extra Traffic is bidirectional in nature and the loss one direction makes the signal useless (and hence both directions should be switched together).

⁸ At a later point this will be clarified and the C&SCs that are relevant for control will be highlighted. The scheme spec will define which C&SCs are the target for commands etc

⁹ This recognition of levels of control from the most basis local two state switch controller through the various levels shown here and on two ring controllers and the SDN controller peer-hierarchy is a manifestation of and a validation of the concept of the Management-Control Continuum. Representation of the Management-Control Continuum will be further explored in the next release.

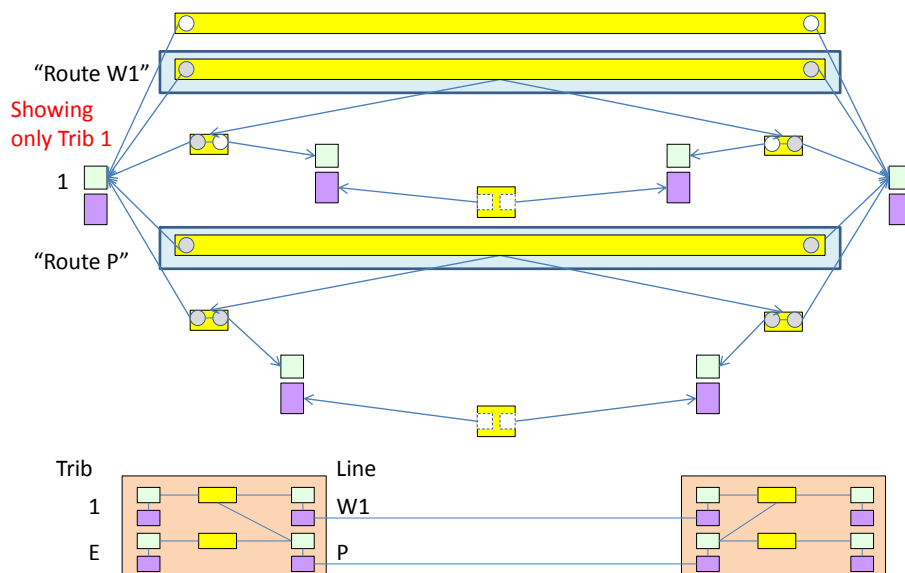


Figure 3-19 Showing route based representation of the 1:N protection scheme

The figure above shows a fragment of the route based representation. The figure detail only shows one traffic signal to avoid clutter. All traffic signals and the Extra Traffic are modeled with the same essential form. Extra Traffic is shown in the figure below.

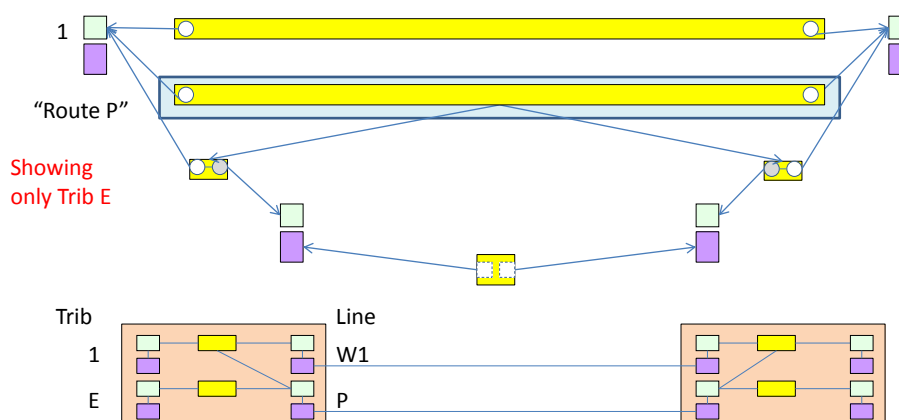


Figure 3-20 Showing Extra Traffic in a route based representation of the 1:N protection scheme

The figure below shows detail of C&SCs and switches for W1 in the 1:N scheme.




Figure 3-21 Showing independent two ended view of W1 route detail in a 1:N protection scheme

4 Mesh Network cases

4.1 N:1 with multicast nodal Cases

The figure below shows a fragment of an N:1 network wide scheme. It is assumed in the figure that I1 is an input from a network external to the one depicted and hence I1 is represented in the upper FC.

The scheme depicted is related to distribution unidirectional signal that has multiple resilient sources (from left to right in the figure). There is assumed to be a single network that has essentially one vast network FC representing the points in the network where the signals are originated and are terminated and used etc. The depiction of the upper FC is intentionally vague as the focus here is not on the representation of termination but solely on the protection scheme. This structure would apply to broadcast TV and to time synchronization. The timing network is represented in detail as described in [TR-512.2](#) and [TR-512.A.8](#).

Considering the protection scheme:

- There may be many inputs carrying the same signal (or an equivalent)
- There may be many outputs for this signal to be propagated to other places
- There may be monitoring or use of the signal at any switching point
- In the case of time synchronization there is some processing of the signal on transit that needs to be represented

The figure below only shows one node in detail. The small FC taking the inputs I1 to In feeds to a signal processing element represented by an LTP (grey) that then feeds O1 to Om via a single unidirectional multi-cast FC.

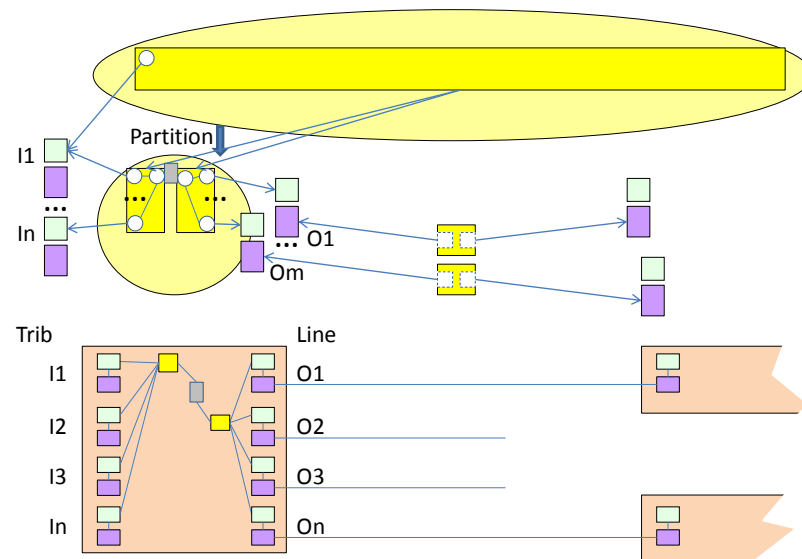


Figure 4-1 Showing a unidirectional N:1 scheme fragment

5 Ethernet Ring Protection [ITU-T G.8032]

5.1 The protection scheme

The [ITU-T G.8032] protection scheme is a network scheme that is built by constructing logical rings that are formed from assemblies of the basic nodal structure shown in the figure below.

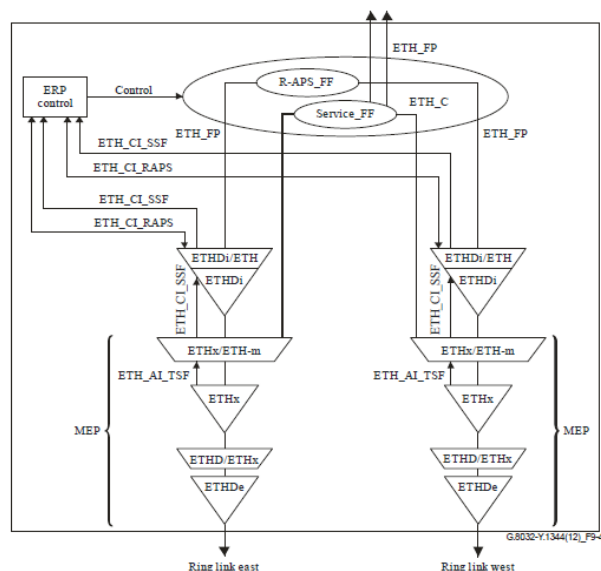


Figure 9-4 – MEPs and R-APS insertion function in Ethernet ring node (normal Ethernet ring node)

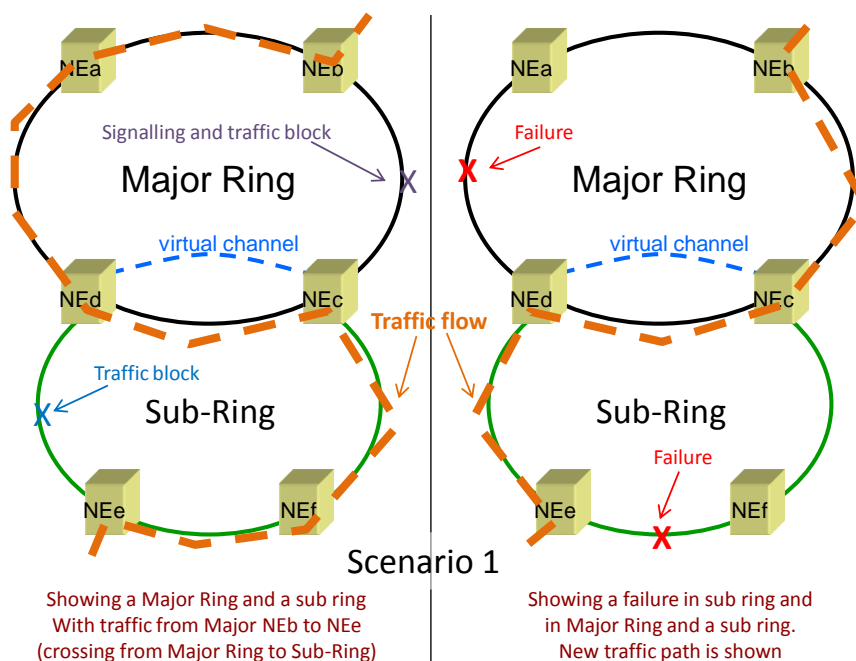
Figure extracted from the referenced ITU-T document

Figure 5-1 [ITU-T G.8032] Ring node control and signaling

A logical ring is defined by the flow of signalling messages that control the protected ring. The protected traffic and control messages use the same links between nodes. When traffic frame enters the protected ring it is sent both ways round the ring. Continuous circulation of Ethernet frames within the ring is prevented by blocks in both the control and the traffic forwarding paths at some point in the ring. When a failure occurs these blocks are moved to position over the failure thus maintaining the flow of both traffic and signalling.

The scope of the scheme can be extended beyond the ring by adding Sub-Rings. The Sub-Ring is a partial ring from a signalling perspective. When one or more Sub-Rings are present the complete ring is called a Major Ring. The Sub-Ring closes protection via the Major Ring.

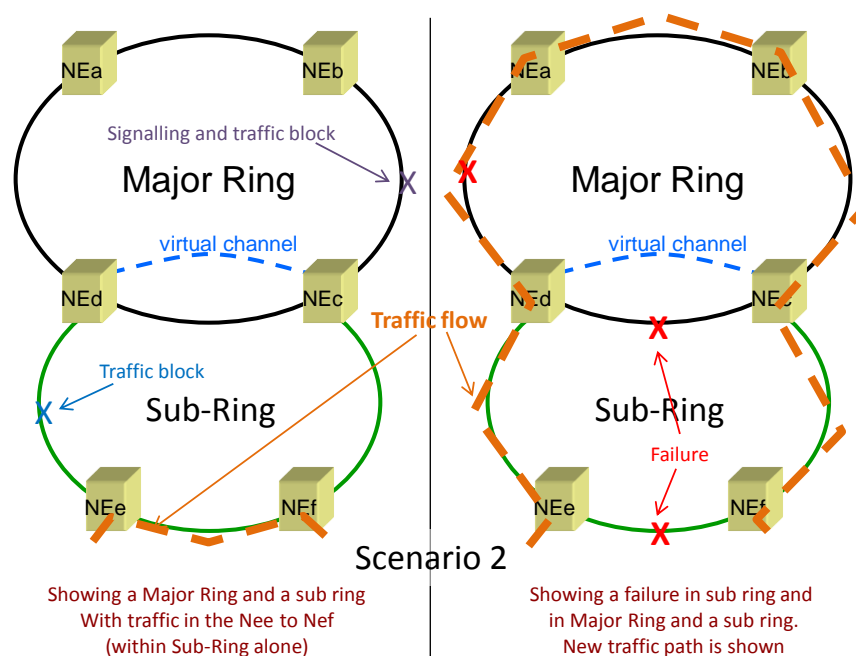
The two figures below show a view of [ITU-T G.8032] protection highlighting major ring and sub-ring configurations and examples ring behaviour on failure.



Based on Fig. 7-31/G.8052 (Control process types)

Figure 5-2 [ITU-T G.8032] Rings showing traffic flow under normal and failure conditions (1)

The figure above shows traffic that flows between a node in the Major Ring and a node in the Sub-Ring. The figure below shows how the major ring offers protection to traffic that is between nodes in the Sub-Ring



Based on Fig. 7-31/G.8052 (Control process types)

Figure 5-3 [ITU-T G.8032] Rings showing traffic flow under normal and failure conditions (2)

In a major ring that has no failure present the traffic block is typically positioned in the same place as the control signalling block (although fundamentally it need not be the case). In a sub-ring only the traffic needs to be blocked. Each active node in a ring has a controller for the ring to ensure the block behaviour. The signalling may transit nodes that are not involved in the scheme so long as the traffic transits the same nodes with no drop opportunity.

The protection scheme supports multiple overlaid logical rings (see Figure 5-13 Two overlaid [ITU-T G.8032] Major Rings on page 38). The protection scheme uses reserved MAC addresses for the ring controllers. All controllers on any specific ring have the same (reserved) MAC address. Part of the MAC address is the Ring ID. All controllers for any case of ring ID have the same MAC address and it is this MAC ring that essentially defines the boundary of an individual protection control domain. A single VID can be used to carry signalling (the R-APS VID) for multiple different rings (where each ring has a different ring IDs). Each traffic ring has one of more dedicated VIDs. Multiple traffic VIDs may be controlled by a single protection control domain.

The Controller controls a subset of the FCs passing through the NE (see Base Model in Figure 5-9 View of NE actively participating in two rings showing spec and encapsulation on page 33). The blocking of the signalling is per MAC address whereas the blocking of traffic is per VID.

Signalling goes both clockwise and anticlockwise. Signalling traffic uses the normal multi-cast drop and continue forwarding behaviour.

In the ring node model shown in the figure below there is a clear split between VID termination and MAC termination to handle the different multiplicity (one VID and n MAC terminations).

Note that in an implementation the Control functionality may be shared between rings but in the model it is the logical function that is considered and from that perspective the control functions are completely separate/disjoint.

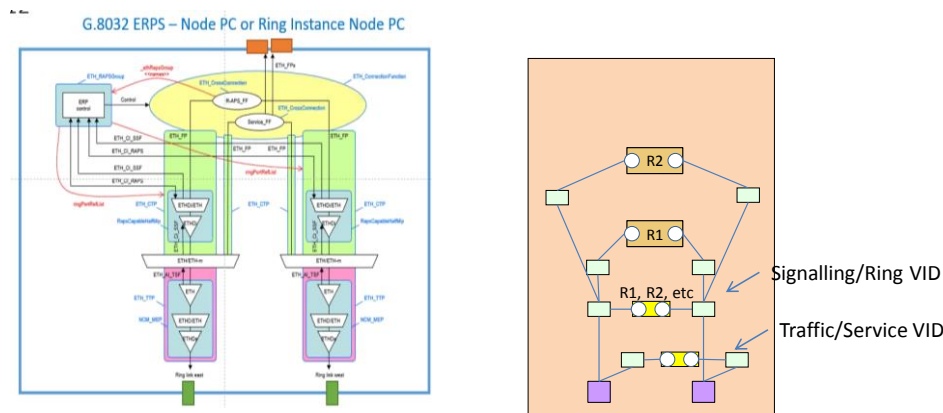


Figure 5-4 Basic model for a ring node

The model reflects the following key features

- Propagation of signalling is by drop and continue
 - The effect in the Ring is that messages do not terminate and hence there is a signalling block
- The ring is defined by the RingId in the last byte of the MAC Address
- The signalling FC is at the VID scope and hence it is a filter¹⁰ in the LTP not a switch in the FC that does the ring block
 - The block on the signalling VID on a per MAC basis
 - Not all rings using the same VID need to be blocked at the same node
 - Not all rings using the same VID need to be co-routed
- The signalling VID block is constructed by provisioning MAC to propagate (the RingX signalling block is formed by not configuring the MAC address (this is a detail))
- The signal block and the traffic block are assumed to be at the same node for a ring (this appears to be a minor simplification).
- The traffic block could be represented as as switch in the FC or an attribute in the LTP

The figure below show expanded detail in the LTPs that deals with signaling.




Figure 5-5 Detailed node model focusing on signal flow

In the figure above the FC supports the signalling information flow (i.e. it is no different to the traffic forwarding function).

¹⁰ A filter in the LTP is essentially a switch in an FC encapsulated in the LTP

The figure above shows a single C&SC. There would be a C&SC for each ring. The C&SC is a relatively complex function as shown in the figure below but dealing with the detail in the controller appears not relevant. In a future form that detail could (should) be laid out in a scheme spec (as it is invariant) such that in future a smart controller could interpret failure modes etc.

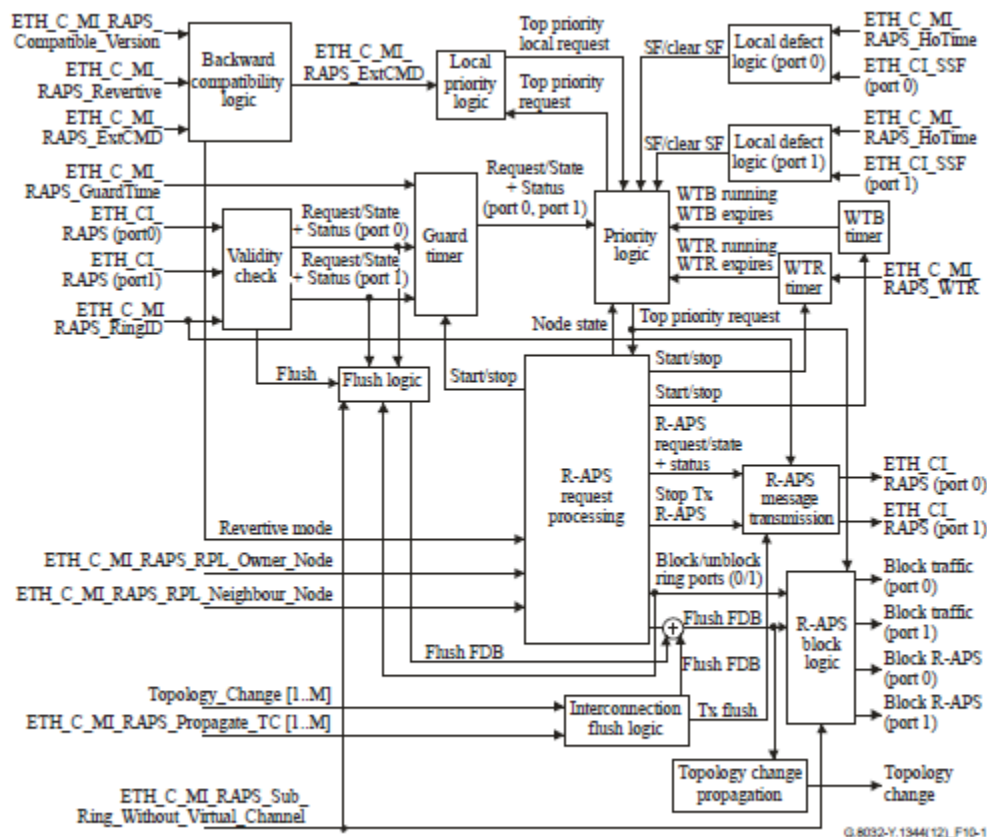


Figure 10-1 – Breakdown of the ERP control process

Figure 5-6 Detailed view of a control function from [ITU-T G.8032]

There are several different views of the ring control that can be derived from the protection model:

- The raw functionality can be exposed per logical ring node as shown in Figure 5-5 Detailed node model focusing on signal flow on page 30
 - The C&SCs for all nodes in the ring can be considered as a single emergent C&SC (see Figure 5-12 The basic [ITU-T G.8032] ring on page 37). This allows "ring wide" parameters to be provisioned across the ring instead of per node.
- The functionality, VID through signalling traffic and all C&SCs related to the signalling VID can be encapsulated in a larger C&SC where that element would represent the control of all rings (see the figure below). This encapsulation obscures some details. For example;

- If there is a ring that simply passes through the node then that Ring ID is not apparent other than via configuration of the C&SC ports.
- Where there is an intermediate NE that does not engage in the protection scheme but does pass the signalling there would be a normal FC such that the representation is not uniform from NE to NE.
- Both the control and the traffic VIDs could be encapsulated in a C&SC
 - The model supports this view as any VIDs for the FD can be allocated to the C&SC as the FCs are created. The traffic VIDs are not directly owned by the C&SC other than when the traffic FCs are created

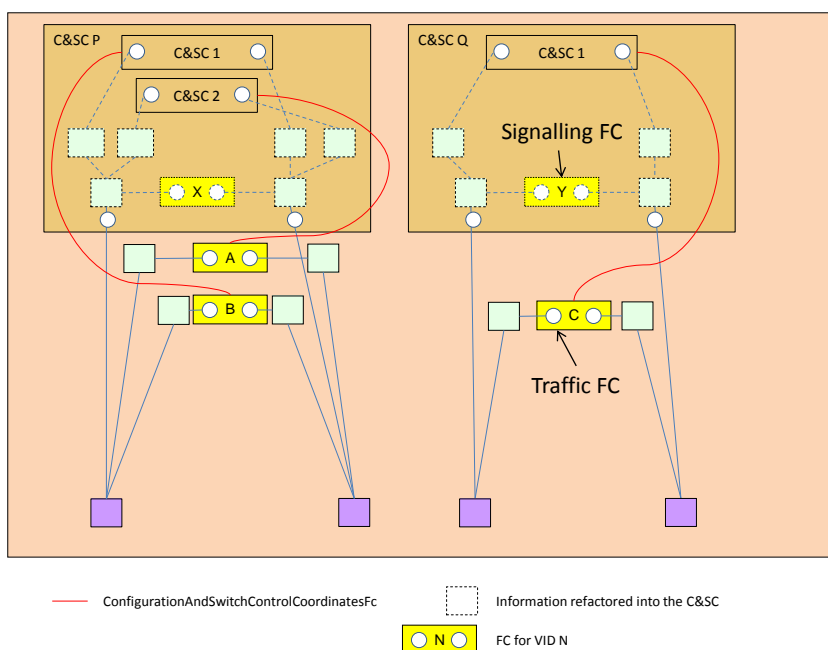
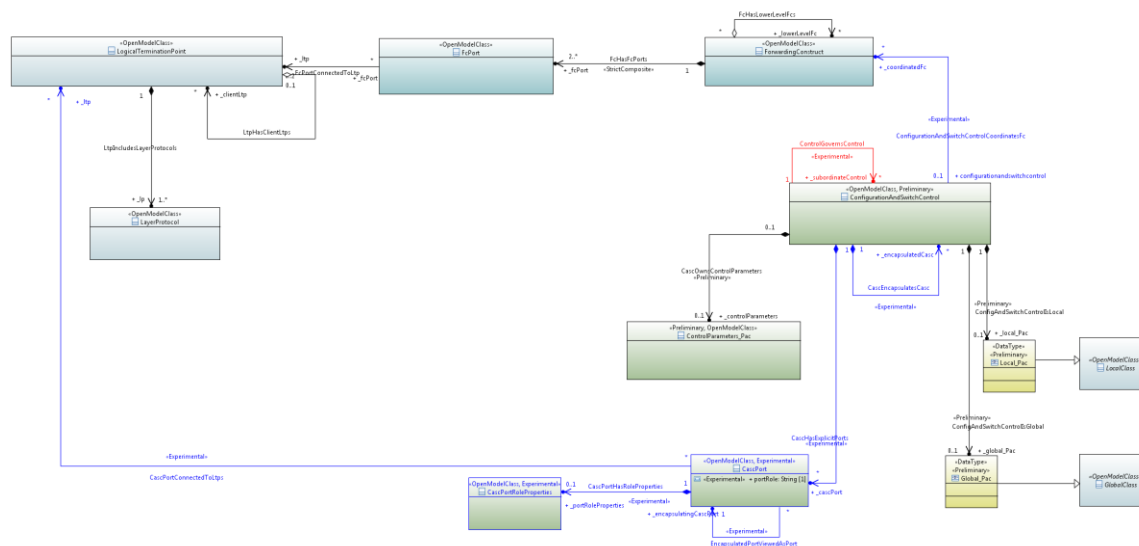


Figure 5-7 Ring node represented by C&SC encapsulations

5.2 Relevant pieces of the resilience model for [ITU-T G.8032]

The figure below highlights the key classes and associations used to represent [ITU-T G.8032].



CoreModel diagram: Resilience-G.8032-Pattern

Figure 5-8 Resilience model structure for G.8032

5.3 Using the spec model to explain the alternative raw and encapsulation forms

The [ITU-T G.8032] ring is can be defined by a scheme spec. In a full model the scheme spec mechanism would apply as shown below to represent the basic model, an encapsulation form and the relationship between the encapsulation form and the basic model.




Figure 5-9 View of NE actively participating in two rings showing spec and encapsulation

The "Base model" diagram shows an example layout, in an instance diagram form, of a single [ITU-T G.8032] node that is involved in two protected rings (R1 and R2).

The "Base scheme spec sketch..." diagram shows a detailed representation of the nodal aspects of the [ITU-T G.8032] protection scheme spec (see [TR-512.7](#) for more details on scheme specs). The elements of the spec are created from the ONF CIM using the "Prune and Refactor" (P&R) approach. This supports the construction of several cases of any class from the model with corresponding associations from the ONF CIM. The spec is augmented with rules that constrain the creation of instances of entities of the scheme so as to abide by the scheme definition. For example "no more than two ports per ring id" on the signaling FC.

The "C&SC Encapsulated..." diagram shows the results of a second P&R stage where the scheme spec is taken and embedded in a C&SC shell. This intermediate step provides an aspect of the mapping of the raw scheme to the "Abstracted C&SC spec". As the FCs and LTP cannot be embedded in the C&SC the model is somewhat of a hybrid but it allows the next step of construction.

The "Abstracted C&SC spec" diagram shows the results of a third P&R stage where the properties of the LTPs (including association ends) are merged into the corresponding C&SC ports as are the port properties of the FC and the FC itself (the FC itself has no relevant properties).

In this case the "Base scheme spec..." is further backed up by a more detailed set of specs for the C&SC for [ITU-T G.8032]. At this point the spec for [ITU-T G.8032] is not documented in a machine interpretable form. The longer term intention is that it would be machine interpretable.

As a consequence of the above steps there is a formal path from the "Abstracted C&SC spec" to the definition of the detailed underlying mechanism. As a consequence the representation from an implementation that uses the "Abstract C&SC spec" form can be transformed in a running solution to a view that follows the "Base scheme spec.." using a machine interpretable definition of the transformation.

As the detailed set of specs are moved to machine interpretable forms an advanced controller will have the information to fully interpret the protection scheme and its data.

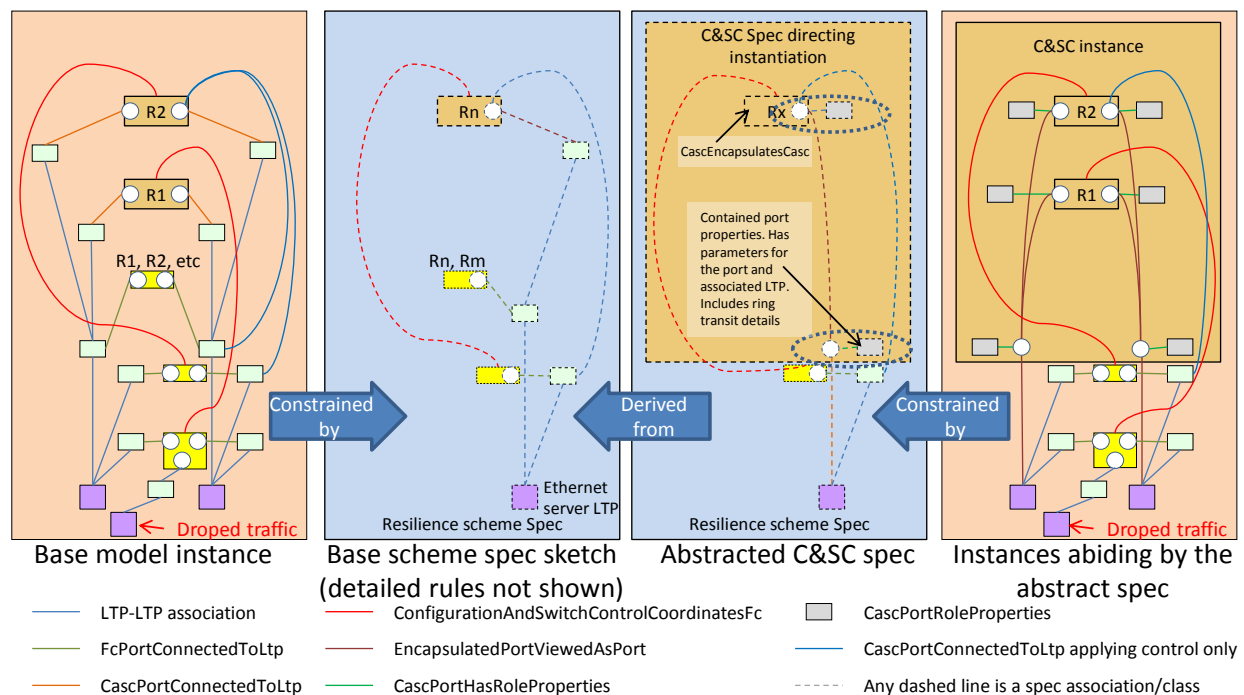


Figure 5-10 Relationship between the two instance views shown via their related spec

The figure above show the relationship between instance sketches and the corresponding specs and highlights the relationship between the specs.

The scheme specs and the rigorous relationship between those enable a controller to interpret and expand a compact form. If any other compact forms are chosen they should be rigorously related

5.4 Representing the block

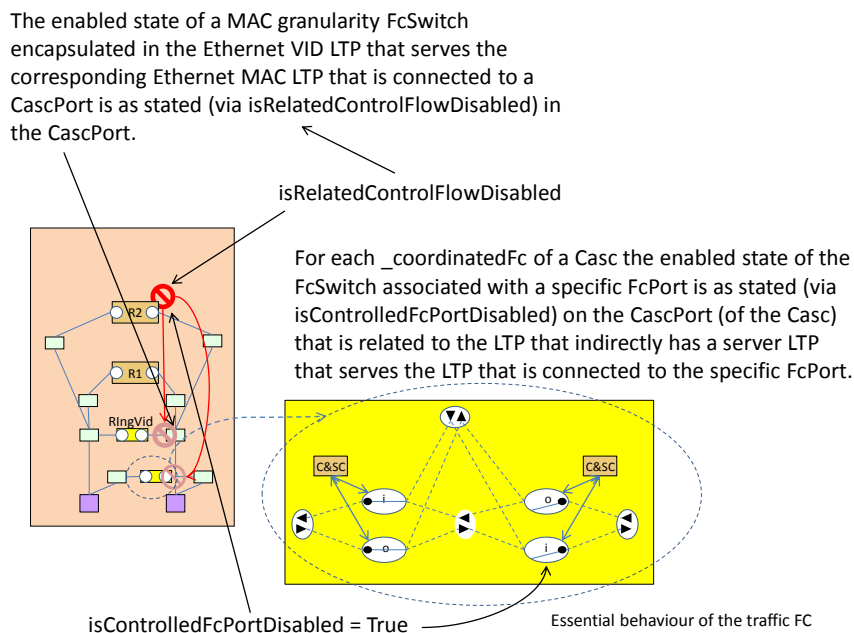


Figure 5-11 Applying the "blocks" to the ring

The CascPort supports both the sending of signaling through and/or application of control to the associated LTP and/or the gathering of monitoring information from the associated LTP. The controls can be applied directly to the associated LTP and/or indirectly to an appropriately deterministically related LTP peer or server to the associated LTP as described by the scheme spec and illustrated in the figure above. The same applies to the gathering of monitoring information.

Considering [ITU-T G.8032] protection as an example the control parameter related to the "isRelatedControlFlowDisabled" property of the port applies also to the indirectly related LTP dealing with the control signal and the "isControlledFcPortDisabled" property of the port applies specifically to the port of the controlled FC as explained by the scheme spec.

In addition the scheme spec will indicate whether the actual state of each individual controlled FC can be determined directly from the FC or whether only the aggregate state is available. Clearly the former may cause performance issues in an implementation if hundreds of FCs are controlled and switched together especially if notifications are sent for changes in every one independently.

5.5 Forming the ring

The figure below shows an example of the protection scheme where there is only one ring set up (controllers and signaling) and there is no traffic applied to the ring. As noted previously the ring is emergent from the signaling and nodal control arrangement. The ring can be represented by a superior C&SC that groups the nodal C&SCs for the ring.

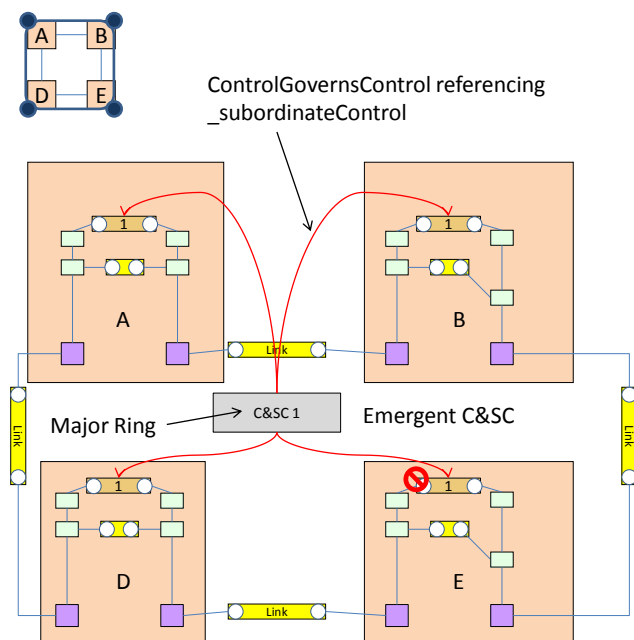


Figure 5-12 The basic [ITU-T G.8032] ring

Whilst in this simple case there also appears to be a control forwarding ring formed from the FCs, in more complex cases with various overlaid rings the control FC can become a complex structure where the VID is used by multiple rings that are not co-routed. This is because it is at the MAC level that the ring appears rather than at the VID level.

The following figure emphasizes the point as it shows two overlaid rings that are not co-routed. The figure could represent a deployment situation where protected private networks are being supported and two customers of the provider have private networks that happen to have some sites in common.

In this case there is no need to pass traffic between the two private networks. If traffic were to need to be passed traffic from the two rings could only be interconnected at one point to prevent loops. In this configuration there would be a single point of failure. To enable protection without a single point of failure an alternative configuration is constructed that uses Sub-Rings. Sub-Rings are shown in later figures.

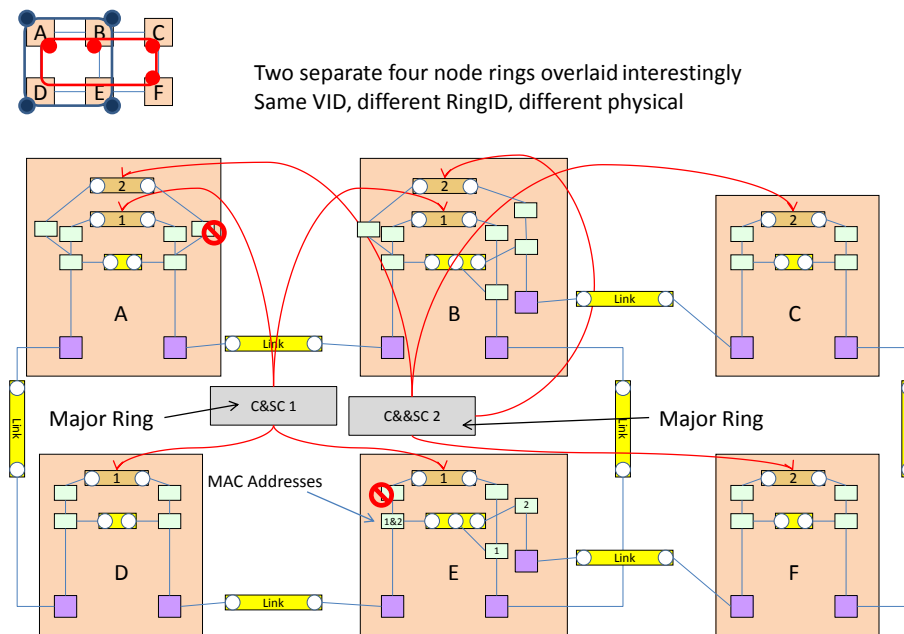


Figure 5-13 Two overlaid [ITU-T G.8032] Major Rings showing signaling only

As can be seen from the figure the two rings share a single VID for signaling but are not co-routed. Ring 2 simply transits through NE D and NE E (there is no controller present that deals specifically with Ring 2 but the signaling VID is set up to allow Ring 2 MAC to pass).

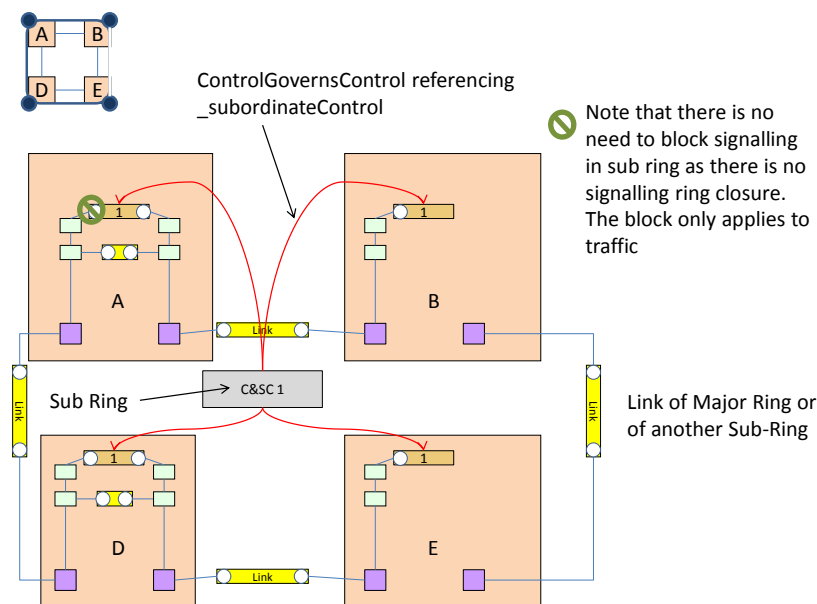


Figure 5-14 Basic [ITU-T G.8032] Sub-Ring

The figure above shows a sub-ring configuration. As described earlier, the Sub-Ring closes protection through a Major Ring (or another Sub-Ring where the Sub-Ring configuration is attached to at least one Major Ring). A Major Ring may support many Sub-Rings.

The figure below shows a Major Ring with one protected Traffic VID between ports on NE A and NE E. In the figure the block is such that the traffic will flow via D.

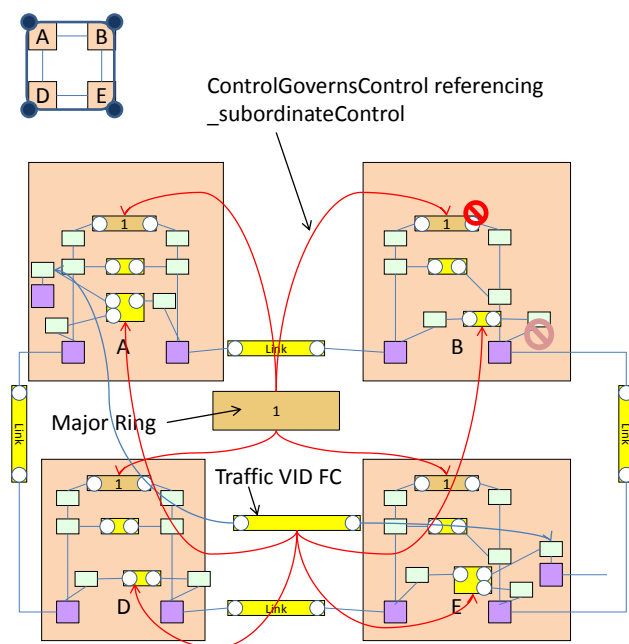


Figure 5-15 Major Ring showing Traffic

The figure below shows a Sub-Ring and associated a Major Ring from a signaling perspective. As the figure only shows signaling the two rings appear to have no association.

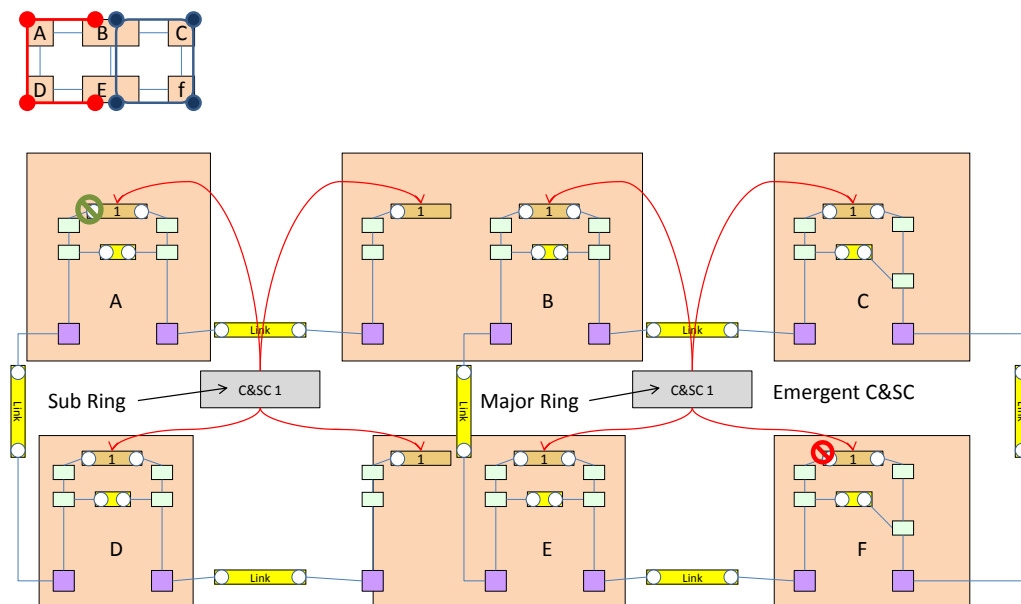


Figure 5-16 Basic [ITU-T G.8032] Major Ring and Sub-Ring

The figure below shows a single traffic VID between ports on NE A and NE F. In NE B and NE E there are three-way FCs that provide a broadcast that enables the protection scheme. The figure shows the necessary traffic blocks in NE A and NE F.

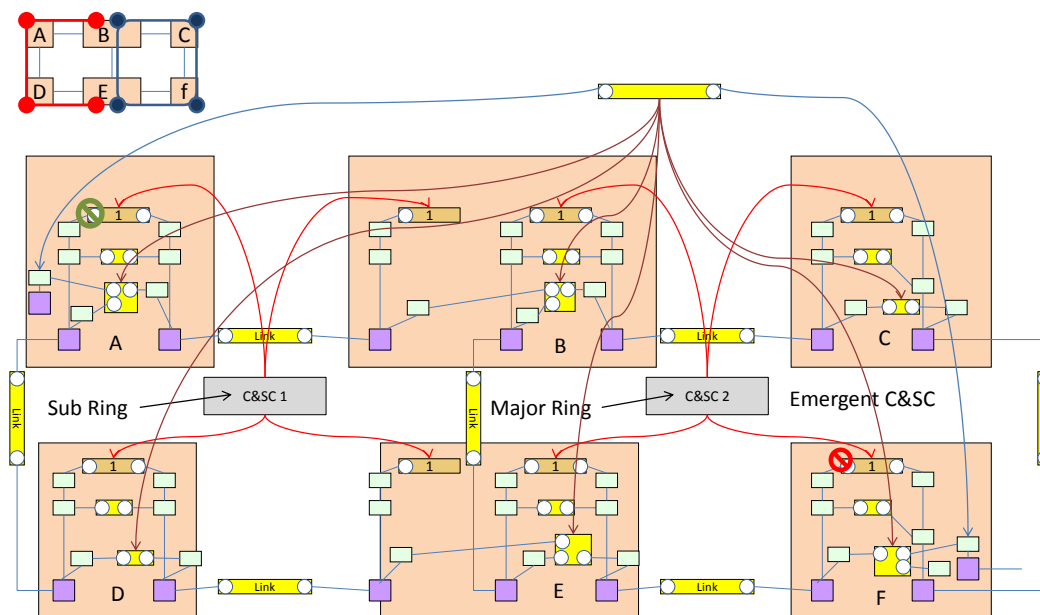


Figure 5-17 [ITU-T G.8032] Major Ring and Sub-Ring showing traffic

To reduce the clutter the figure below shows only the traffic.




Figure 5-18 [ITU-T G.8032] Major Ring and Sub-Ring showing only traffic

In the approach shown in the two figures above the FCs in NE B and NE E are partly controlled by C&SC 1 and partly controlled by C&SC 2. The figure below shows an alternative layout of traffic where there are dedicated FCs per control domain which are interconnected via a zero length link which is an artificial termination construct that represents the boundary of the control domains.

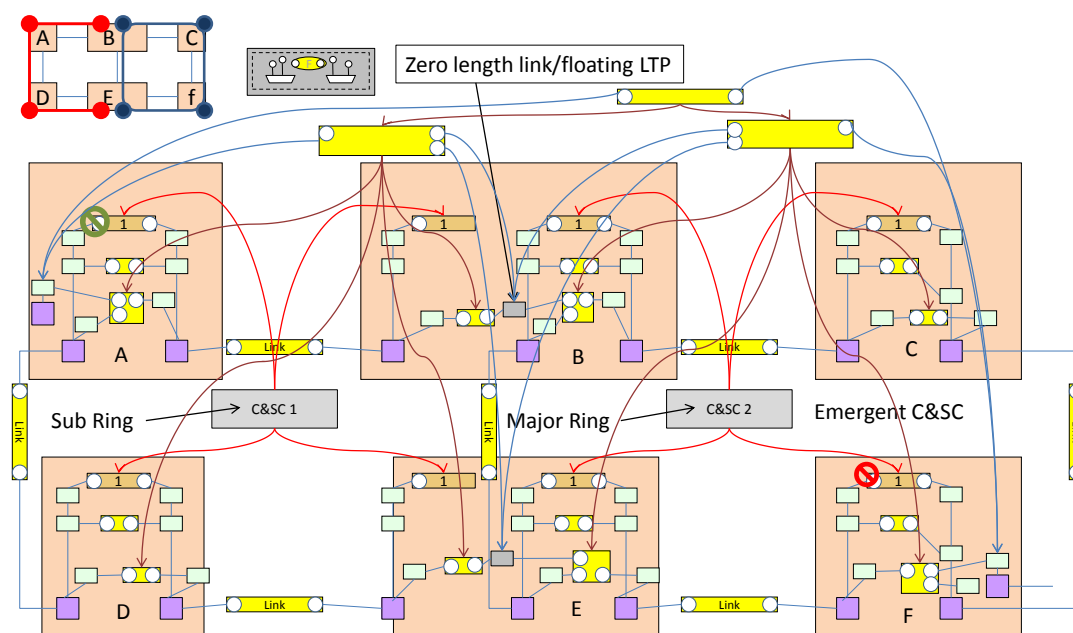


Figure 5-19 [ITU-T G.8032] Major Ring and Sub-Ring showing traffic with zero length link

A zero length link can be added per traffic VID. T

To reduce clutter the figure below shows only the traffic.

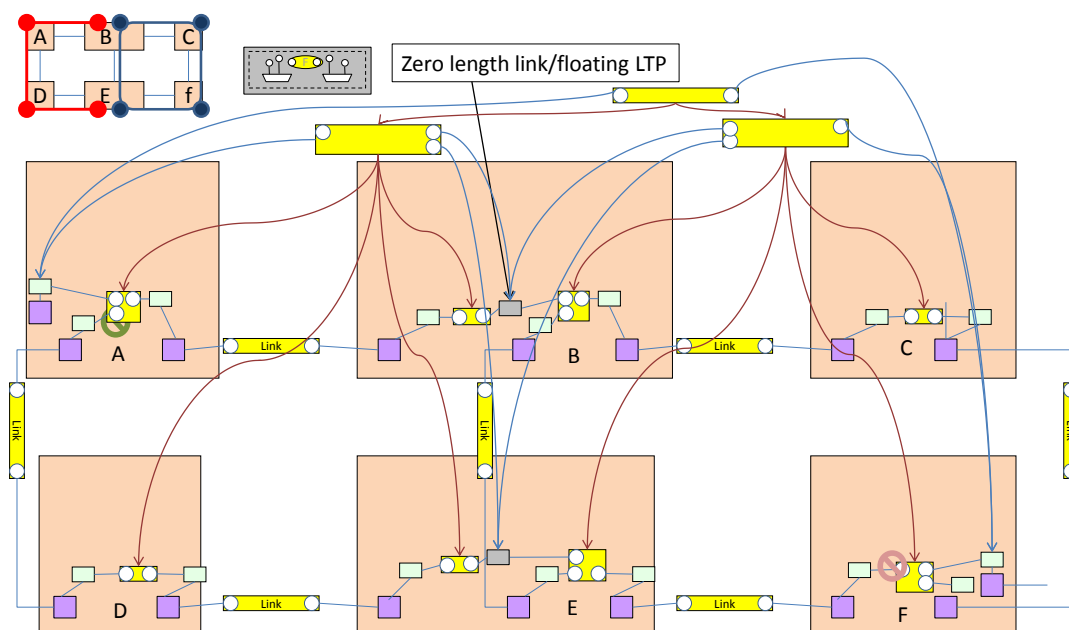


Figure 5-20 [ITU-T G.8032] Major Ring & Sub-Ring showing only traffic with zero length link

The approach using a zero length link adds complexity to the traffic path model but does allow a representation of control isolation.

The figure below show arrangements of rings in a mesh network. Each numbered Ring can use the same signalling VID so long as the Ring IDs are different for each Ring at an intersecting Node. In the example if the number is the ring ID then the VID can be the same. If the rings all have the same Ring ID then the number represents the VID.




Figure 5-21 [ITU-T G.8032] Major Rings in a mesh

More complex cases including those shown below, although not detailed here, are also covered by the model¹¹.

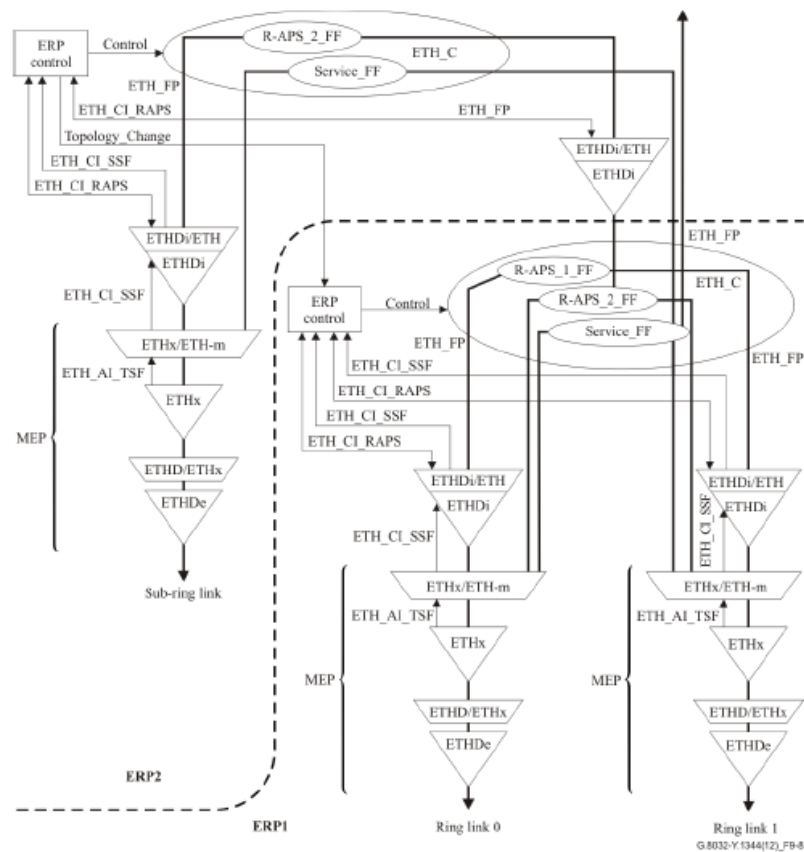


Figure 9-8 – MEPs and R-APS insertion function in an interconnection node with an R-APS virtual channel (different R-APS VIDs)

Figure 5-22 MEPs and R-APS insertion [ITU-T G.8032]

¹¹ Further work will be carried out in a later release to show these and other more complex cases in detail.

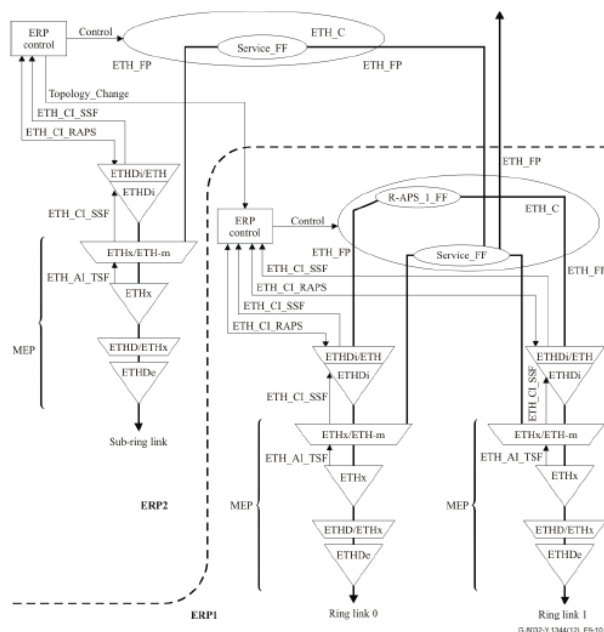


Figure 9-10 – MEPs and R-APS insertion function in a sub-ring interconnection node without an R-APS virtual channel (for a sub-ring connected to a major ring)

Figure 5-23 MEPs and R-APS insertion without R-APS virtual Channel [ITU-T G.8032]

In the figure above, it is not clear how the two Service_FF blocks are joined. This appears to be FC to FC (achievable with the current model), but there may be more complex behaviour that is implied by the figure. This requires further study.

6 Other protected ring schemes

This section details some other ring schemes. The following diagram key applies.

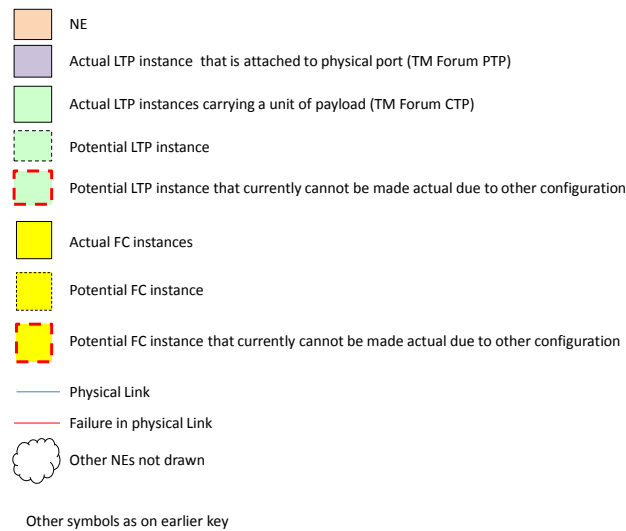


Figure 6-1 Diagram key

The figure below shows the basic network used to explain several ring based resilience schemes showing where a break will occur.




Figure 6-2 The network

- The network technology¹² is such that there are 8 channels of capacity on each link where 4 channels are available for traffic and 4 for protection.
- A single traffic signal could use just a single channel, could use two channels or could use all four channels
 - In the two channel case any available channels from the 4 can be used to make the capacity, i.e. the channels do not need to be adjacent
 - Different channels can be used on different links in the ring
 - Hence blocking is simply on capacity not channel
- The signals are numbered 1-4 for the single channel signal (B1) and 1-2 for the two channel signal (B2)

¹² This is not a real network technology. It has been contrived to make the case easier to express.

6.1 Network Wrapping

6.1.1 The scheme

The figure below shows a view of the network wrapping scheme.




Figure 6-3 The network showing wrapping

- A signal is passing from port 3 node W to port 3 node Z
- When a link Y-Z fails the traffic is routed back round the ring from the break on the corresponding protection capacity B2.P1
- Traffic can be monitored at intermediate points
- The following figures only show the 2 channel and 4 channel traffic (B2 and B4 respectively)
- To simplify the figures:
 - The same channel is maintained throughout the ring for both normal path and protection such that B2.1 must use B2.P1
 - No extra traffic is shown
- B1.n and B1.Pn is not shown. There is no B1 traffic in the ring
- Somewhere in the cloud there is a B2.2 service and a B4.1 service that require protection hence in all NEs shown there will be a B2.P2 and B4.P1 opportunity enabled.
 - If there was no B2.2 connection anywhere in the ring the B2.P2 would not be required.
 - If there was no B4.1 connection anywhere in the ring B4.P1 would not be required.

6.1.2 The model applied

The following figures show the LTP and FC configurations for nodes in the ring under normal and failure condition.

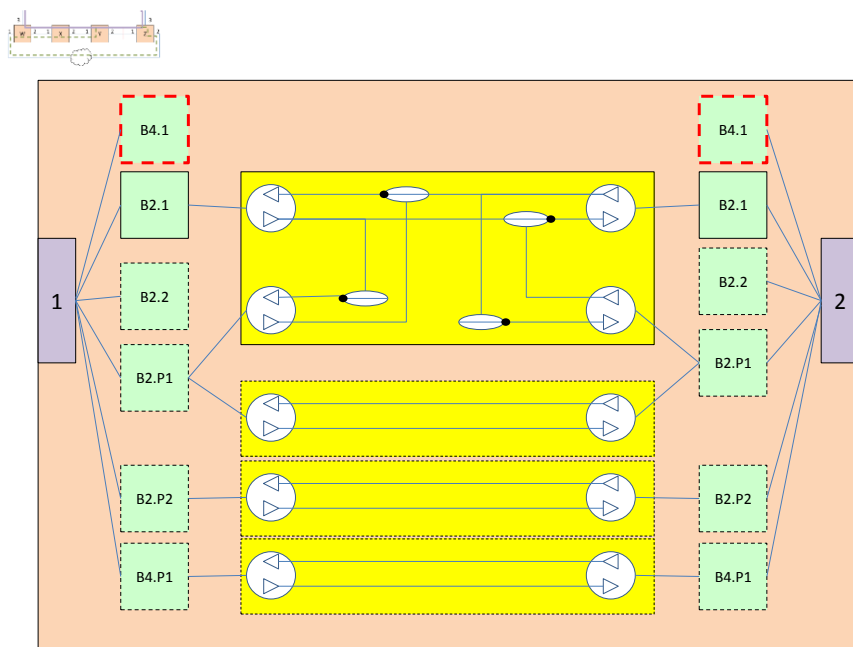


Figure 6-4 Wrapping: NE X and NE Y (no failure in ring)

The figure above shows the configurations of NE X and Y with the switches set to normal position. There is an actual FC allowing signal to flow between B2.1 on port 1 and B2.1 on port 2. There is potential for Signal to flow between B2.P1 on port 1 and B2.P1 on port 2 etc. and hence potential FCs are shown. Because B2.1 is used on port 1 then B4.1 cannot be used etc.

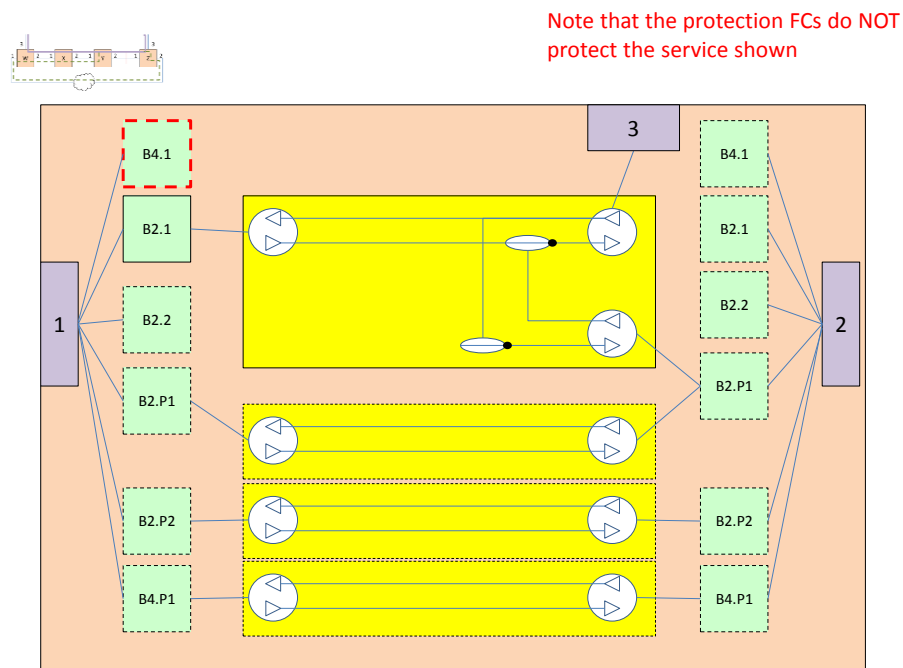


Figure 6-5 Wrapping: NE Z (no failure in ring)

The figure above shows the configurations of NE Z with the switches set to normal position. There is an actual FC allowing signal to flow between B2.1 on port 1 and port 3 (in this case B2.1 on port 2 is not used). There is potential for Signal to flow between B2.P1 on port 1 and B2.P1 on port 2 etc. and hence potential FCs are shown.

Note that NE W has essentially the same configuration although port 2 is used for normal signal flow and the protection faces port 1 not port 2.

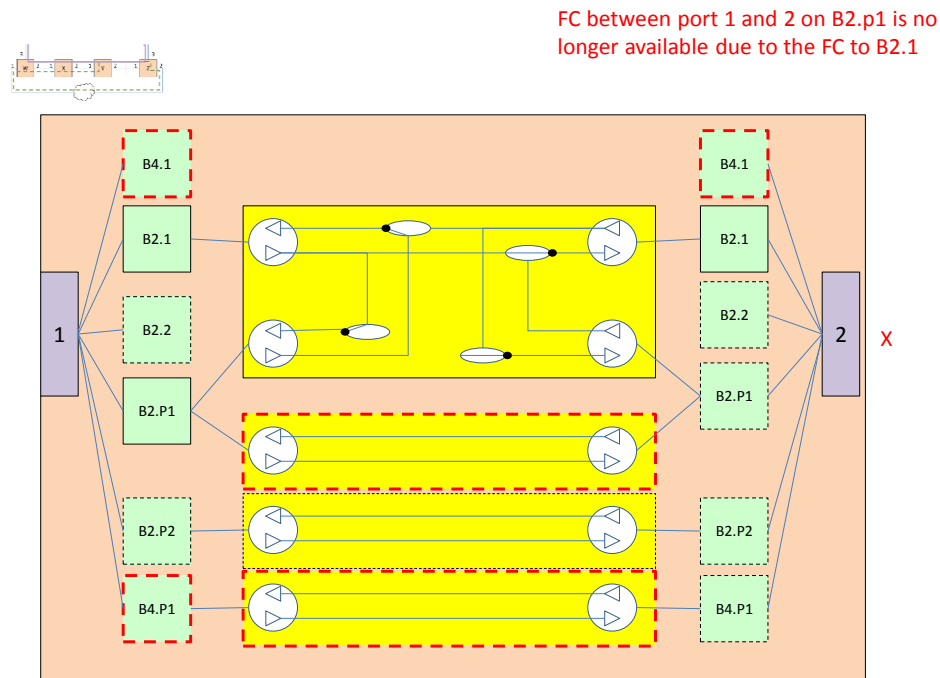


Figure 6-6 Wrapping: NE Y with failure on port 2

The figure above shows the configurations of NE Y with a failure on port 2 such that the switches on the left have been set to select B2.P1 on port 1. The FC allows signal to flow between B2.1 on port 1 and B2.P1 on port 1 such that the signal is wrapped back to port 1. There is no longer a potential for Signal to flow between B2.P1 on port 1 and B2.P1 on port 2 as B2.P1 on port 1 is now used for protection and hence potential FCs are shown as not usable. Because B2.P1 is used on port 1, B4.P1 cannot be used on port 1 and hence the FC between B4.P1 on port 1 and B4.P1 on port 2 cannot be used. B4.P1 on port 2 could be used but there is no other use shown.

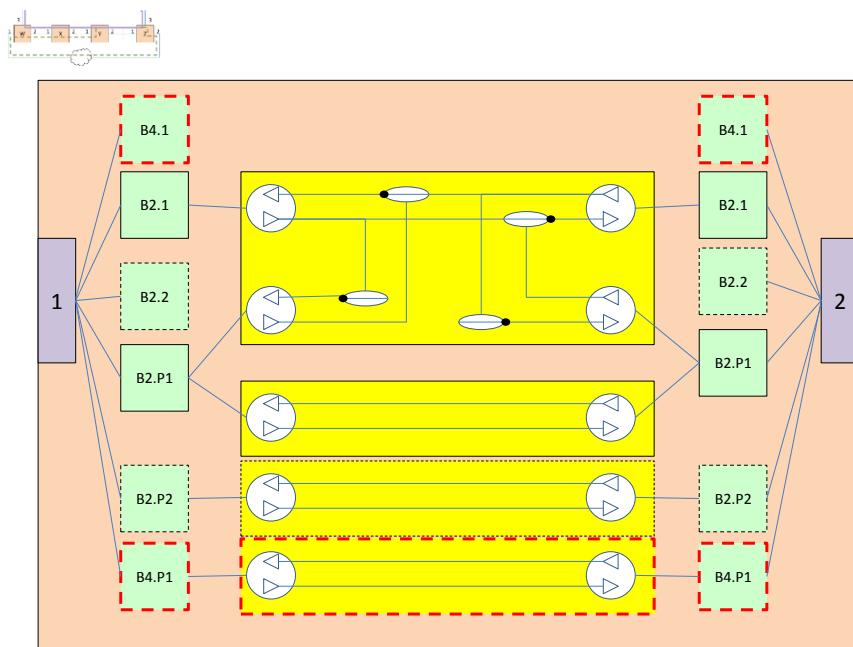


Figure 6-7 Wrapping: NE X with failure on NE Y port 2

The figure above shows the configurations of NE X for the failure on NE Y shown in the previous figure. There is an actual FC allowing signal to flow between B2.1 on port 1 and B2.1 on port 2. There is now also an actual FC shown between B2.P1 on port 1 and B2.P1 on port 2 that carries the protection traffic due to the wrap in NE Y shown in the previous figure. This actual FC is allowed due to the switch positions in the FC between B2.1 on port 1 and B2.1 on port 2 which do NOT select the B2.P1 channels on port 1 and port 2. Because of the B2.P1 usage on both port 1 and port 2 the B4.P1 FC and LTPs on both port 1 and port 2 cannot be used. The potential FC between B2.P2 on port 1 and B2.P2 on port 2 is still available. The assumption is that somewhere else in the ring B2.P2 is being used (otherwise there would be no need to commit the potential).

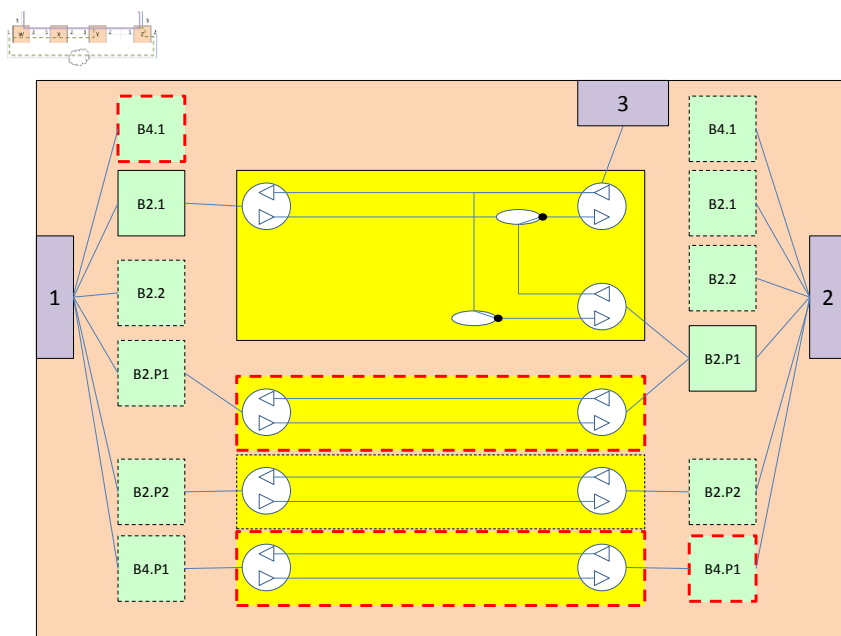


Figure 6-8 Wrapping: NE Z with failure on NE Y port 2

The figure above shows the case for NE Z where the port 3 signal is wrapped onto B2.P1 on port 2. The figure does not show a failure on port 1. The assumption here is that the failure is only detected in NE Y. The assumption is that the scheme does bidirectional switching. It is possible that signal can still flow from port 2 on NE Y to port 1 on NE Z but as the protection scheme is bidirectional it switches both directions of traffic and hence the unidirectionally viable link from NE Y to NE Z is not used.

NE W does not need to switch to protect the signal as NE Y and NE Z perform the protection function in this case. In general, for the wrapping scheme, the NEs either side of the failure perform the protection function.

6.2 Network Steering

6.2.1 The scheme




Figure 6-9 Network showing steering

- A signal is passing from port 3 node X to port 3 node Z
- When a link Y-Z fails the traffic is routed back round the ring from origin on corresponding protection capacity B2.P1
- Traffic can be monitored at intermediate points
- The following figures only show the 2 channel and 4 channel traffic (B2 and B4 respectively)

6.2.2 The model applied

The following figures show the LTP and FC configurations for nodes in the ring under normal and failure condition. The explanations for the figures is similar to that in the previous subsection and hence only differences have been highlighted.

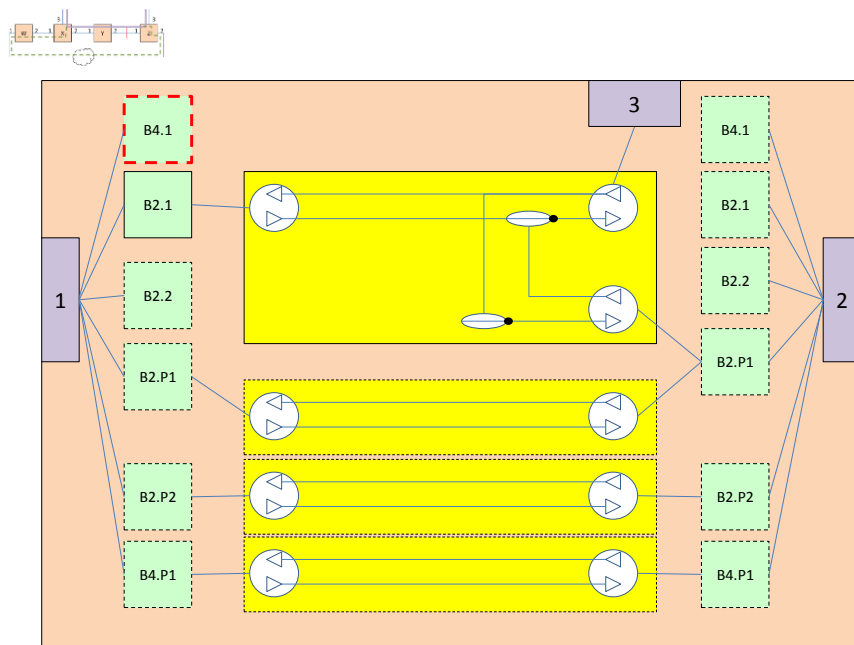


Figure 6-10 Steering: NE Z (no failure in ring)

Essentially as for the wrapping case.

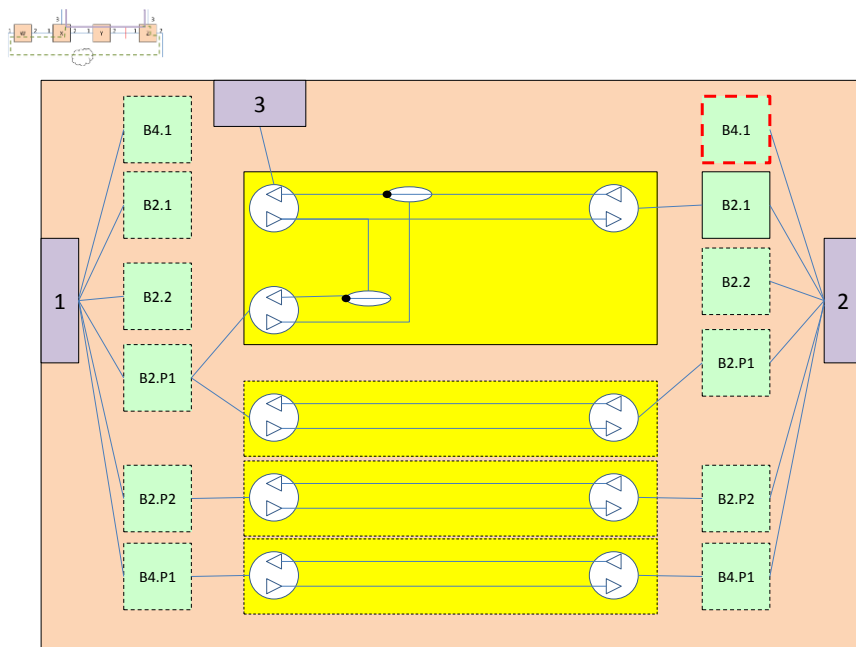


Figure 6-11 Steering: NE X (no failure in ring)

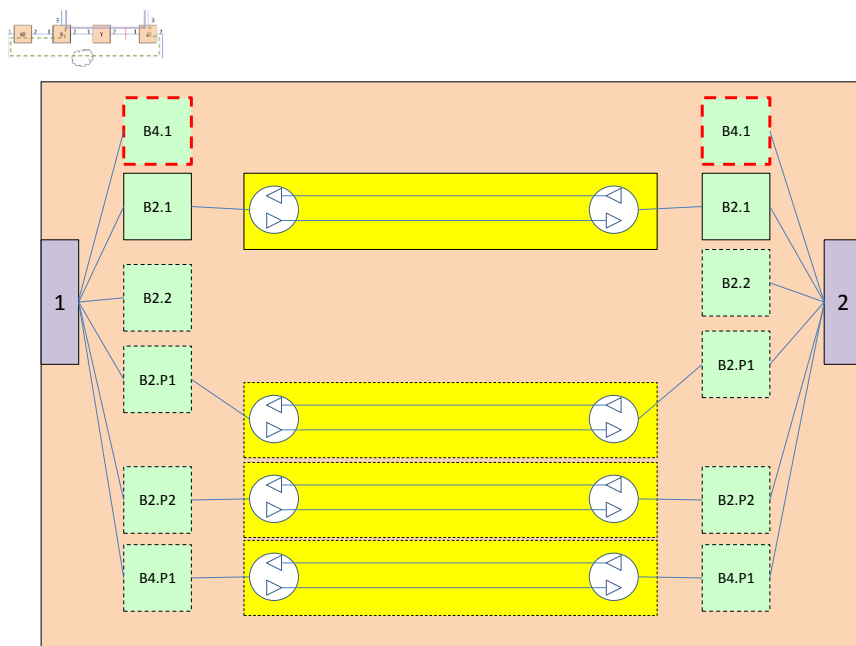


Figure 6-12 Steering: NE Y (no failure in ring)

Note that unlike the wrapping scheme, NE Y has no switching capability enabled, as in this scheme, switching is only performed at the entry points to the ring.

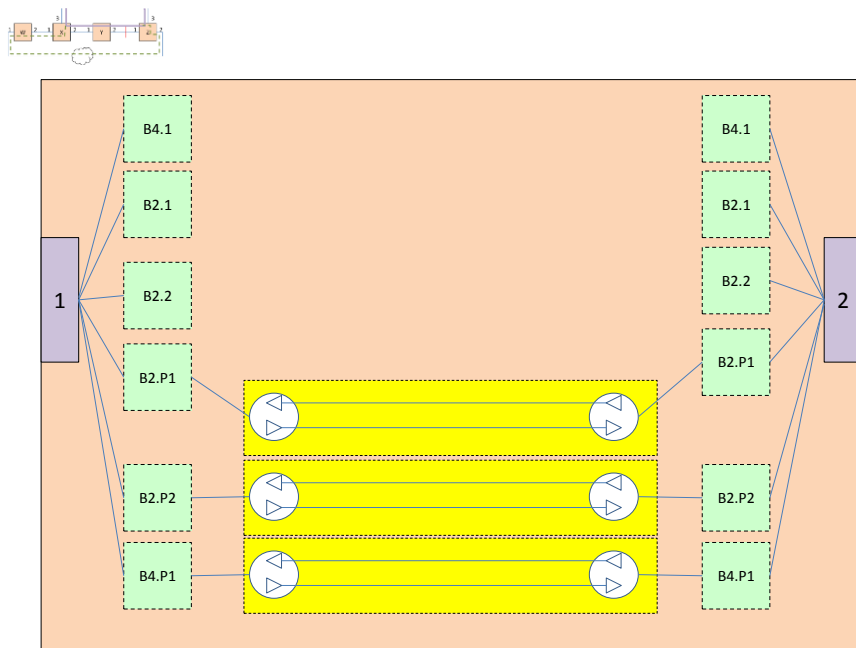


Figure 6-13 Steering: NE W (no failure in ring)

NE W does not participate in the main path of the signal.

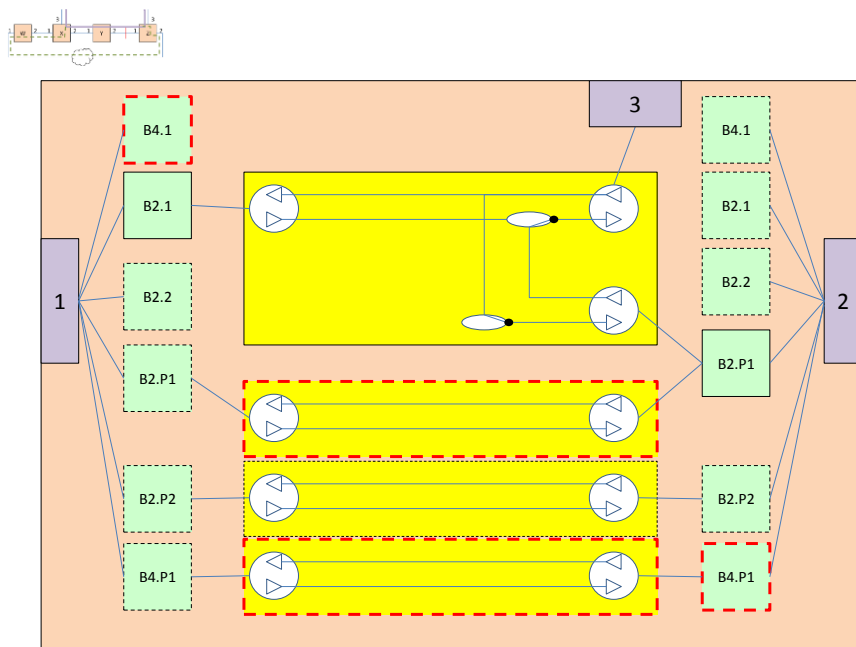


Figure 6-14 Steering: NE Z with failure on NE Y port 2

The figure above is the same as for the wrapping case because the failure is between NE Y port 2 and NE Z port 1 and NE Z is also the entry point for the signal.

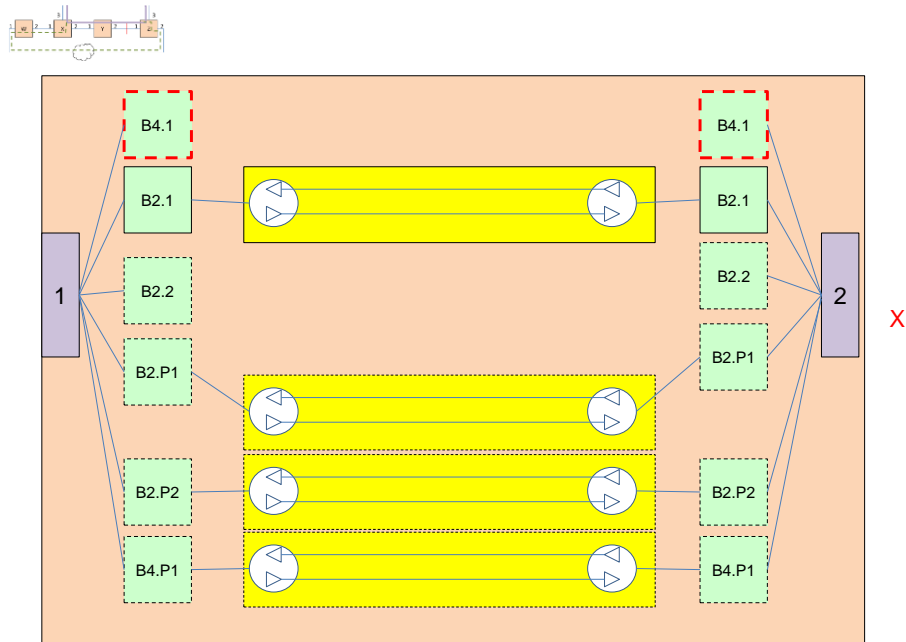


Figure 6-15 Steering: NE Y with failure on port 2 (same as no failure)

No change takes place in NE Y when the failure occurs as the responsibility to protect is at the NEs where the signal is applied to the protection scheme. For the wrapping scheme NE Y was involved in the protection.

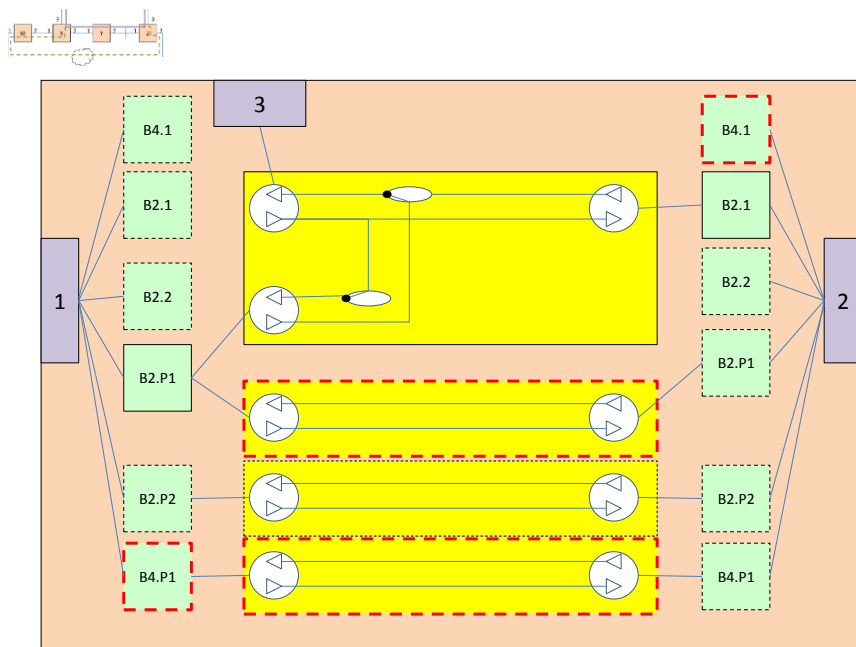


Figure 6-16 Steering: NE X with failure on NE Y port 2

The figure above show the switching occurring at the point of entry of the signal to the scheme. There is no failure either side of NE X so in the wrapping scheme NE X would NOT switch.

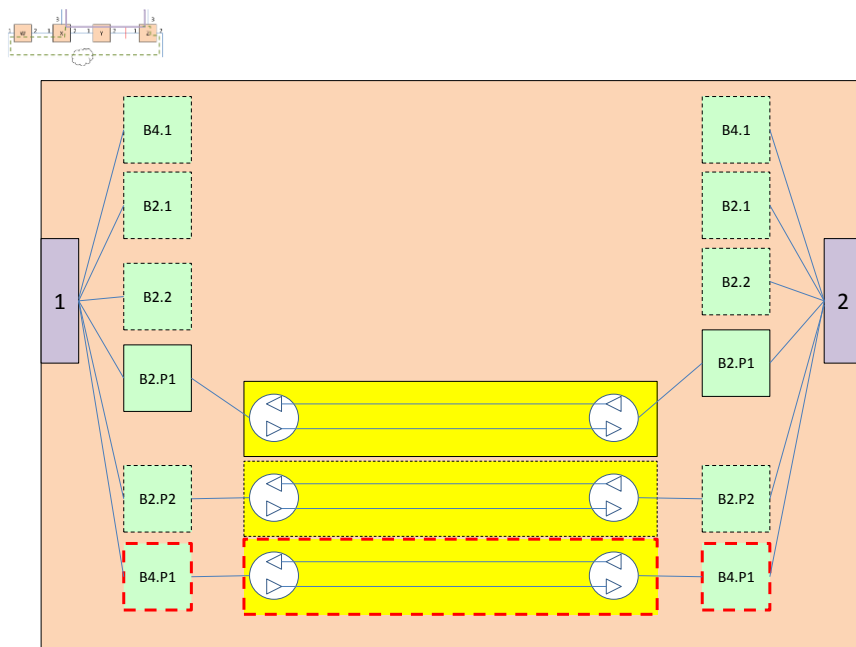


Figure 6-17 Steering: NE W with failure on NE Y port 2

NE W enables signal to pass through supporting the new path resulting from the steering at NE X and NE Z.

6.3 The model in detail for both Steering and Wrapping

- Three methods depending upon details of the actual scheme... FCs are
 1. "created" as potential and then activated when protection requires
 - The controller requests that particular FCs are changed from potential to actual (this potentially involves the controller communicating with all NEs when a protection/restoration action is required)
 2. not present until protection requires but are known to be potential through a specification
 - The controller coordinates the creation/deletion of FCs based upon the scheme description in the scheme spec. (this potentially involves the controller communicating with all NEs when a protection/restoration action is required)
 3. created as actual (rather than potential) with a switch disabling them and are switched on when protection requires
 - The controller requests the change of switch state
 4. The LTPs approach could be the same as the FC approach but there are some hybrids possible
 - The LTPs could be not present even if the FC is until selected by a switch
- Regardless of the approach from the methods described, use one or more specs to explain what can exist and what needs to be created when to form the correct behaviour
 - The spec can remove the need to report/notify entities
 - A composite notification could be designed to inform of a complex configuration change if defined in a spec
 - Potential LTPs and FCs can be considered as "partially created" in that a query on the live system could return them as instances and when they become active this could be considered as a state change rather than a creation (as the rule is indeed known by the NE)
 - A hybrid (with a switch) of potential+off and actual+on could be considered
 - When an LTP is disabled it is potential and when enabled it is actual
 - When LTPs and FCs are gathered /notified, only enabled FCs and LTPs would be reported
 - States may be
 - Actual
 - Potential
 - Potential – disallowed
 - Creation v state change when spec is provided...
 - Seems that much of the CTP/FC would be known from the spec so a simple state change is all that is required
 - Spec could identify group notifications that indicate a change of state of many classes or a batch notification could be provided
 - Challenges: Potential misalignment between spec and reality

- Note that the S&SC is really just a Controller
- The actual state **MUST** be available, the question is how much potential should be reported and how much should be in the spec. The feeling at this point is that the potentials should **NOT** be reported other than via the spec.

End of Document