INSP
Inter-Network Service Protection

Zehavit Alon

November 2010
V01

Nokia Siemens Networks

INSP – Inter-Network Service Protection

- Inter-Network Service Protection (INSP) is a mechanism that protects services over MEF-defined EIs (UNIs and ENNIs) in the interconnection zone (IZ) and guarantees service delivery between domains.
- A network domain can deploy any packet technology and protection mechanism.
- Combining INSP and intra-domain protection mechanisms can provide end-to-end service protection.

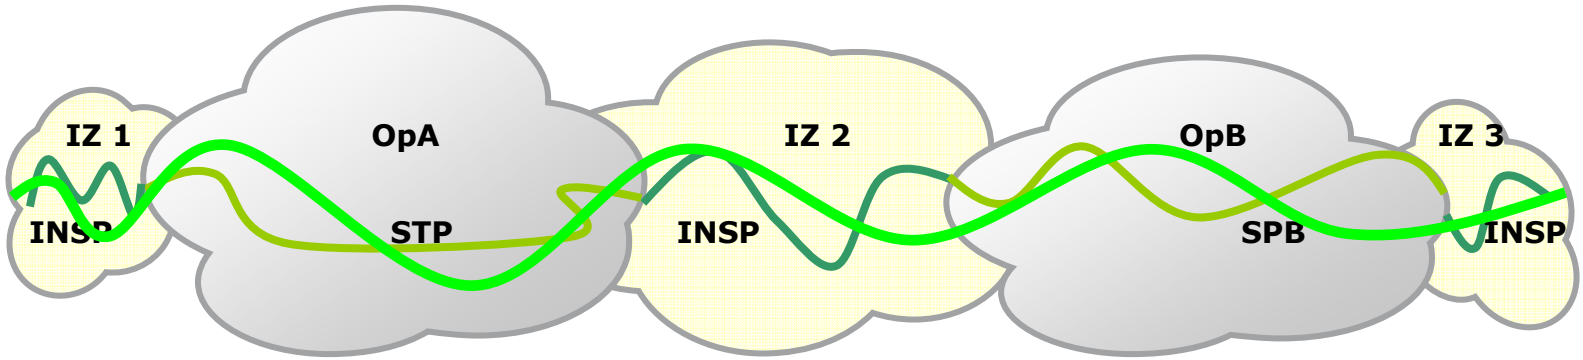

INSP components

INSP components comprise:

- Border nodes – responsible for conveying services to and from the interconnection zone
  - Border nodes protecting a service (or a service bundle) are grouped into a service portal (SP).
- External links – connect border nodes between peer SPs
- Internal links – connect border nodes inside an SP

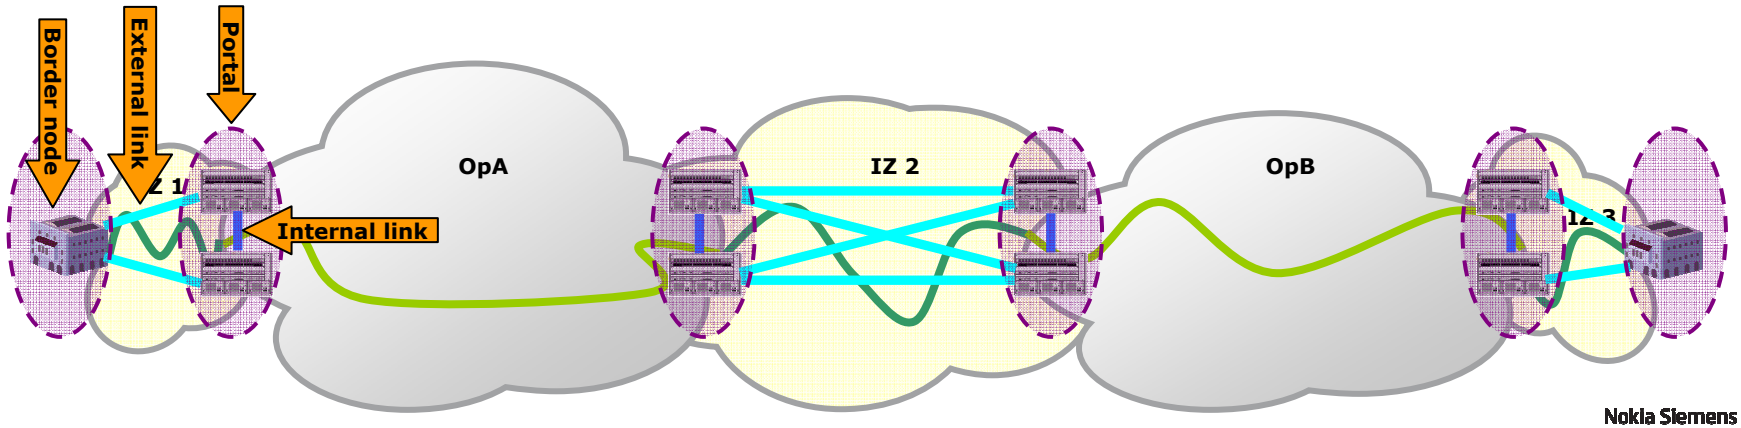

INSP location in IEEE architectural model

- The protection entity, known as a Virtual Protection Group (VPG), is a logical Bridge Port (as defined in 802.1Q) which is configured in each border node residing in a service portal.
- The VPG groups two or more ports (physical or logical (LAG)) which participate in the protection mechanism.
- A port in the group can be connected to an external link or an internal link, and is named external port or internal port respectively.
- The number of internal and external links determines the protection topology.

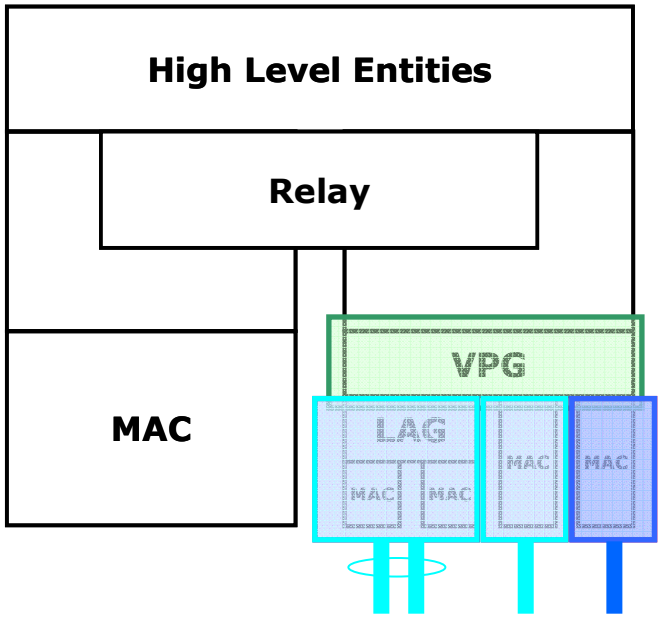

Full mesh topology

Service portal with two nodes. Each node comprises two external ports and one internal port.

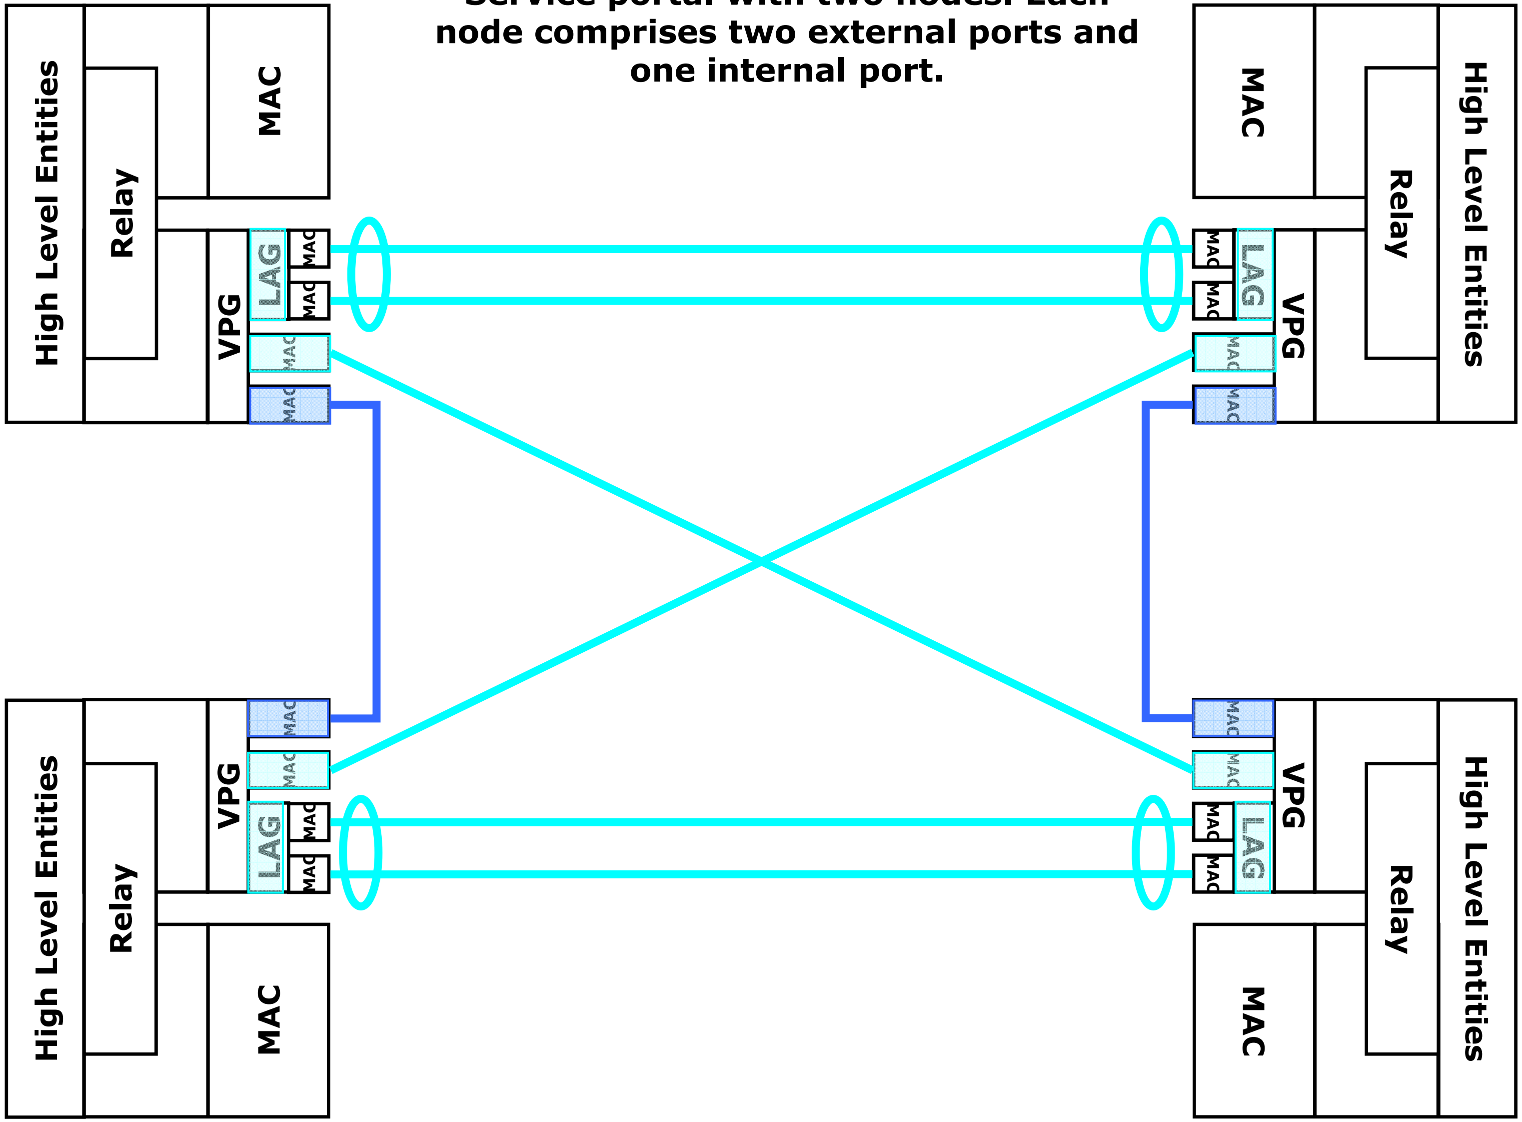

Hourglass topology

Service portal with two nodes. Each node comprises two external ports only (no internal ports).

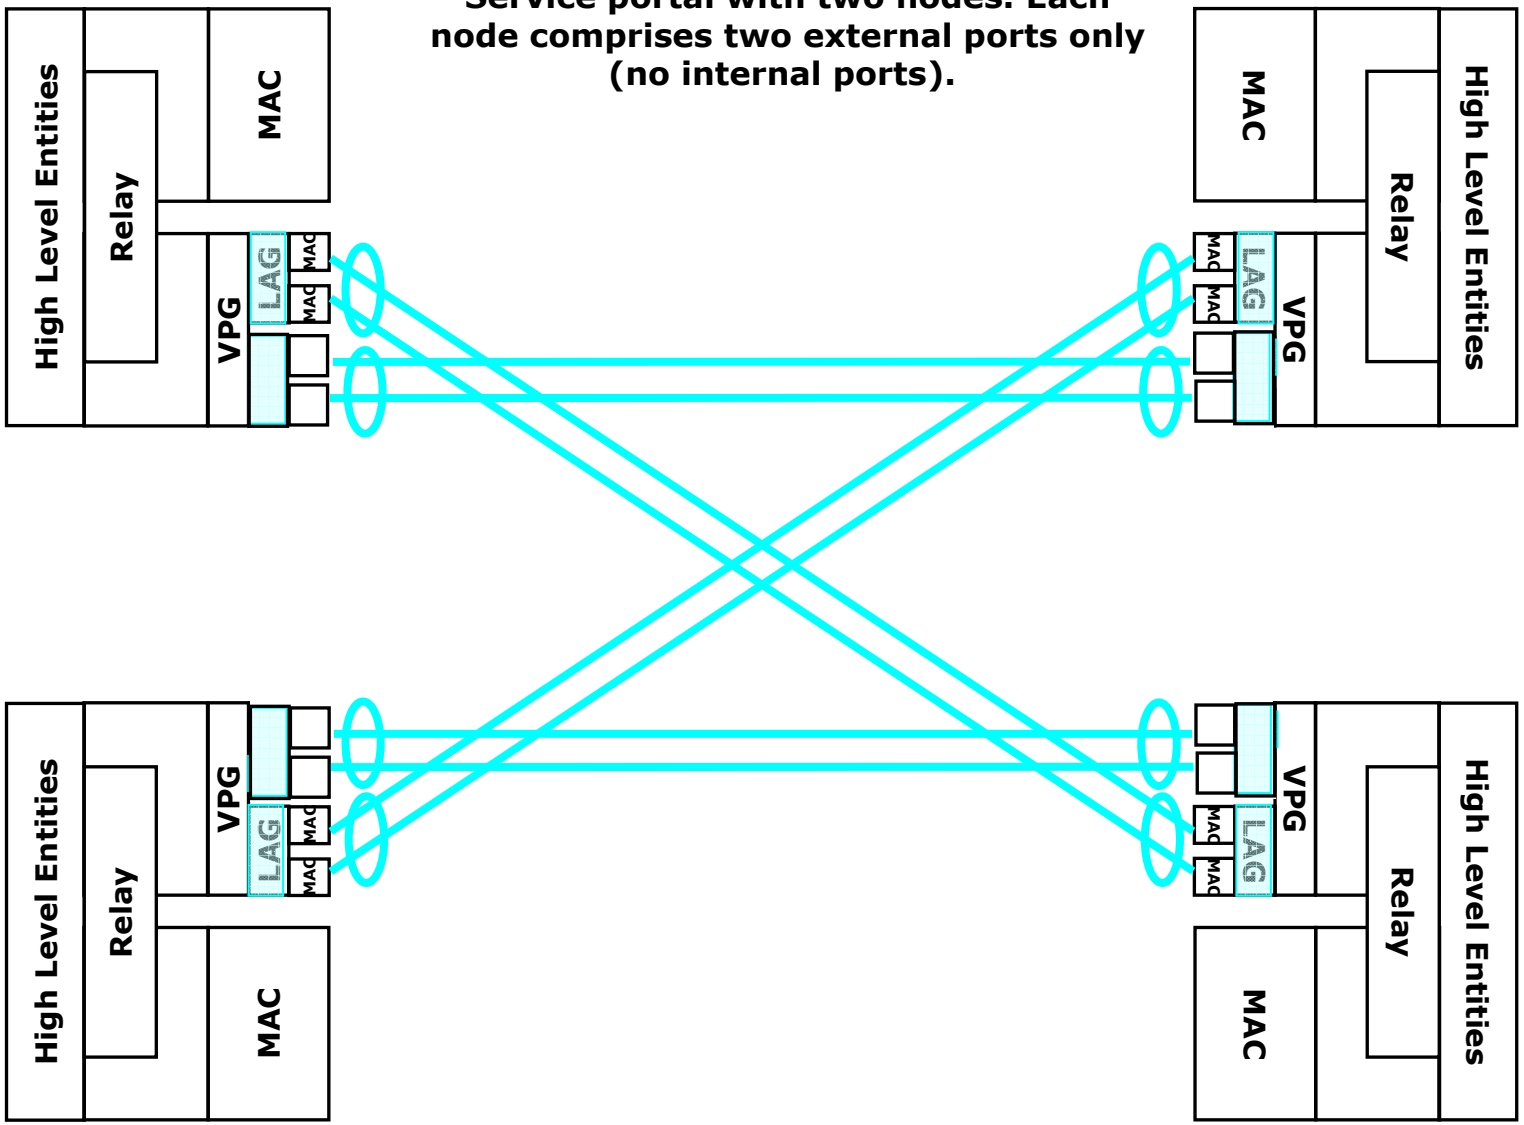

Ring topology

Service portal with two nodes.
Each node comprises one external port and one internal port.

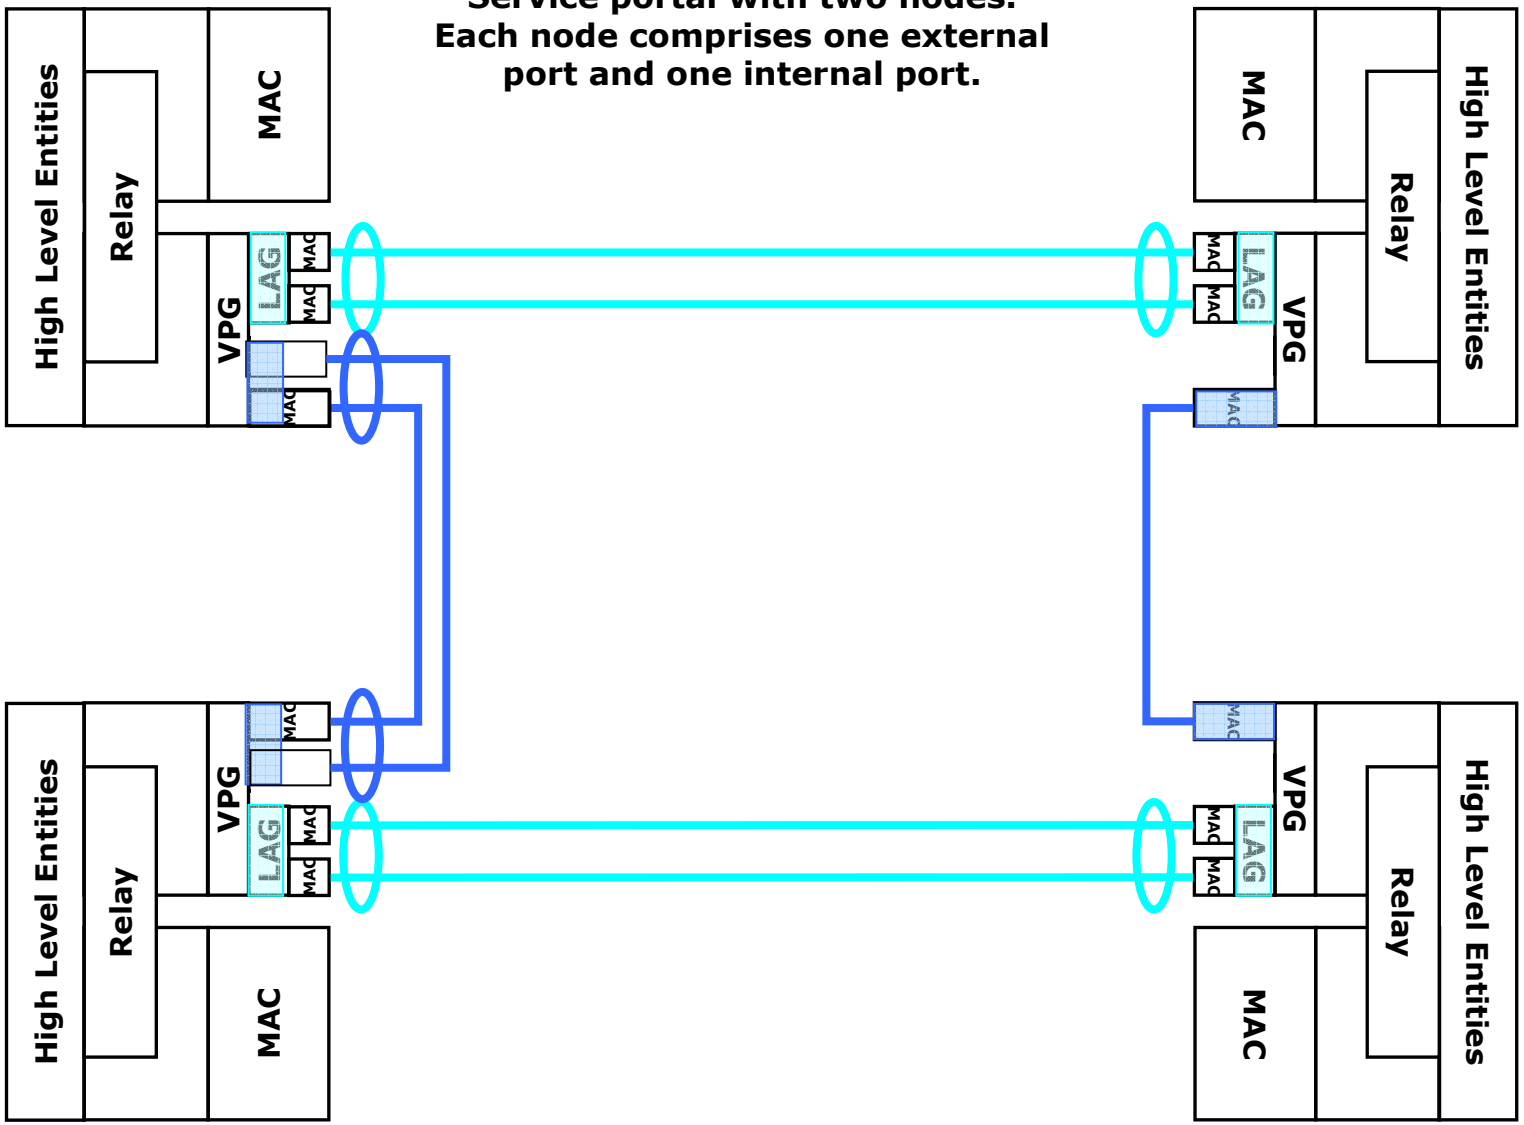

Full mesh topology

Service portal with three nodes.

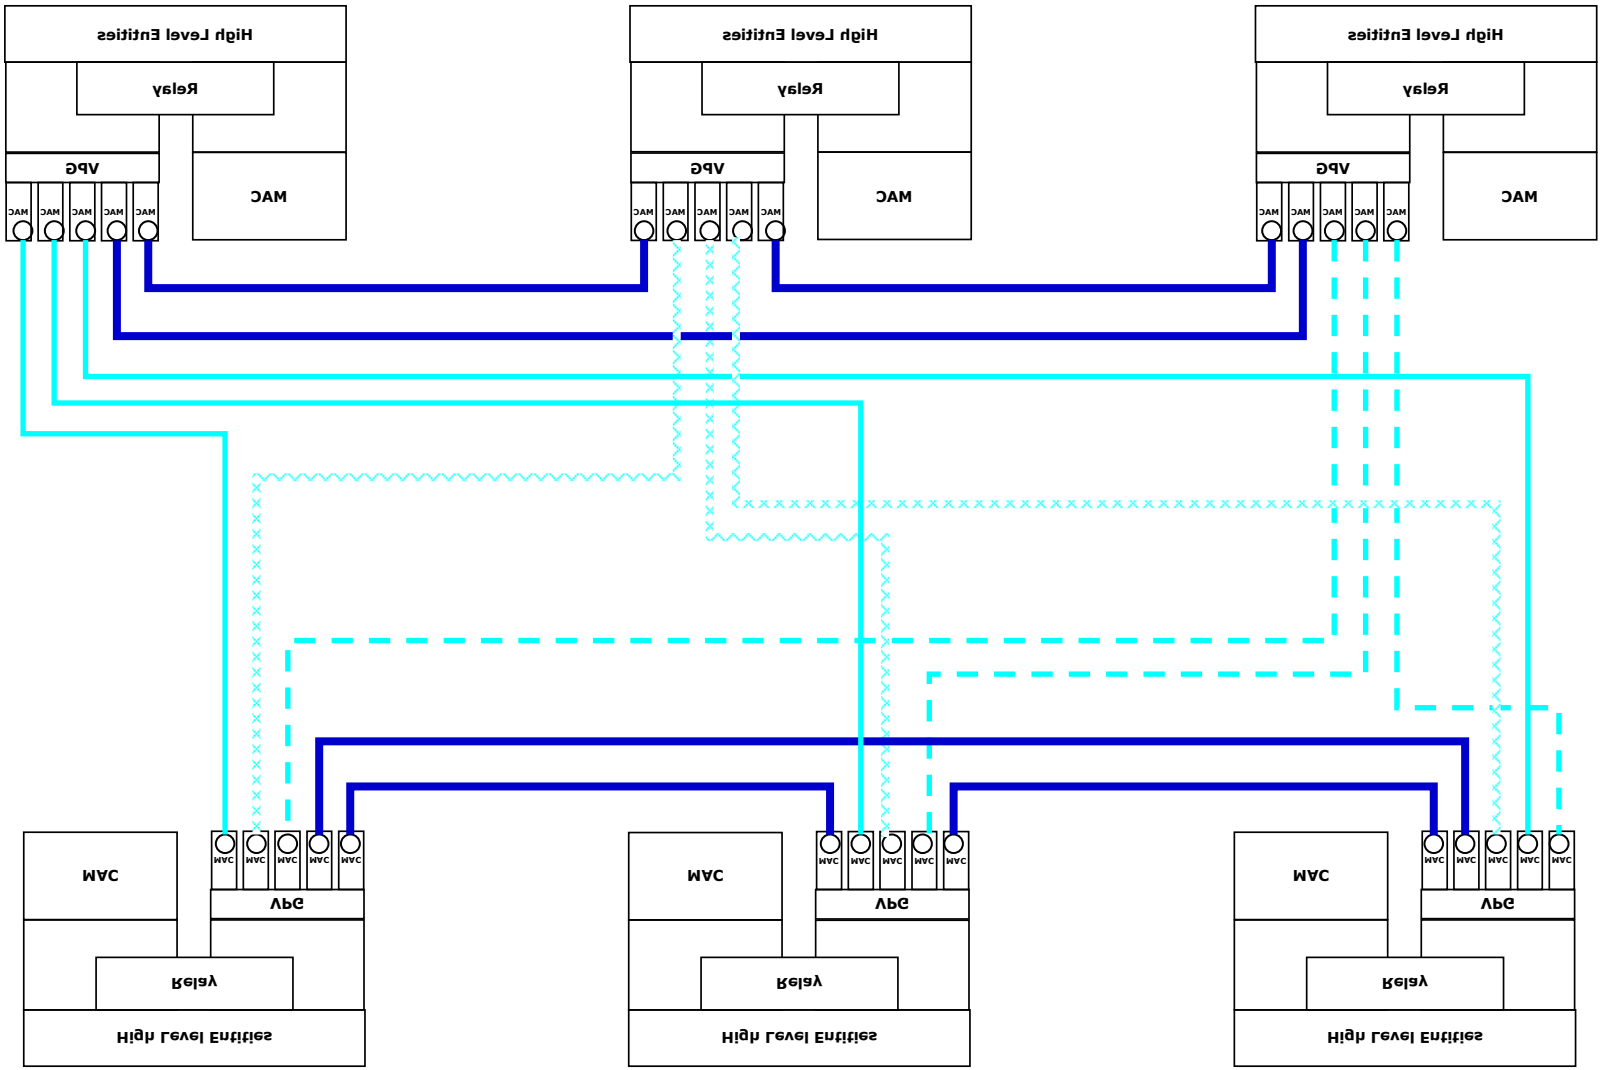

VPG functionality

- The VPG port belongs to the member set of all VLANs that it supports
- Each protected VLAN is configured on at least two underlying ports, at least one of them external.
- VLAN traffic is relayed from the VPG to one of the potential underlying ports (configured with that VLAN), and from all underlying ports to the VPG
- The port selection algorithm is executed on the VPG.

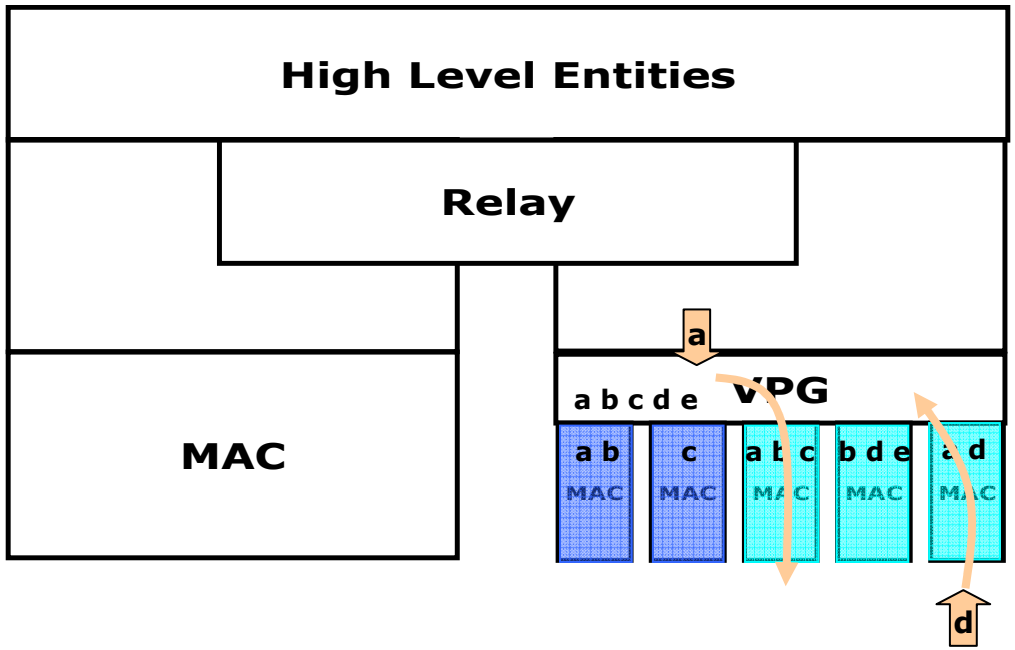

VPG functionality (cont'd)

A port in a VPG can accept or reject service traffic, and can be in one of the following states:

- Standby – drops service traffic
- Active – relays service traffic to the VPG

The VPG can have a maximum of two active ports assigned to each VLAN, one of them external.

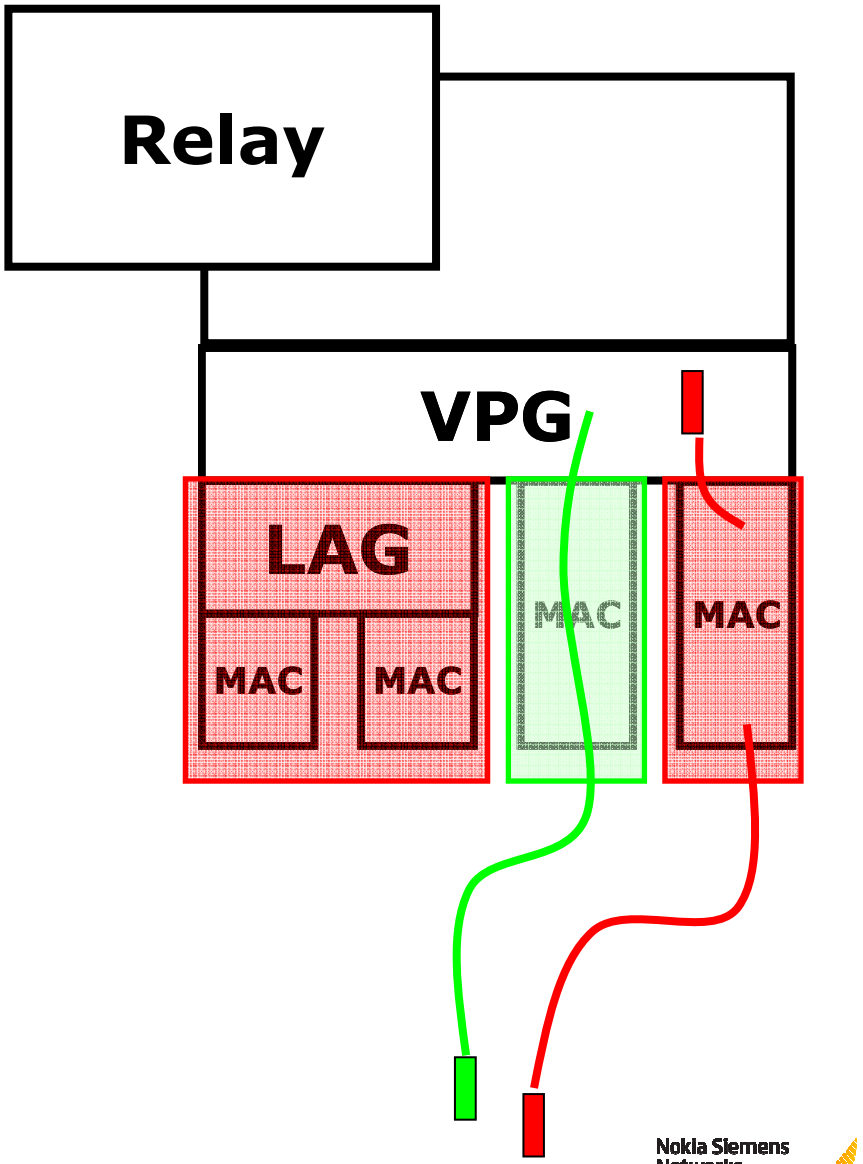

VPG functionality (cont'd)

A VPG can accept, reject or tunnel service traffic, and can be in one of the following states:

- Down – does not relay packets between bridge ports. Received packets are dropped

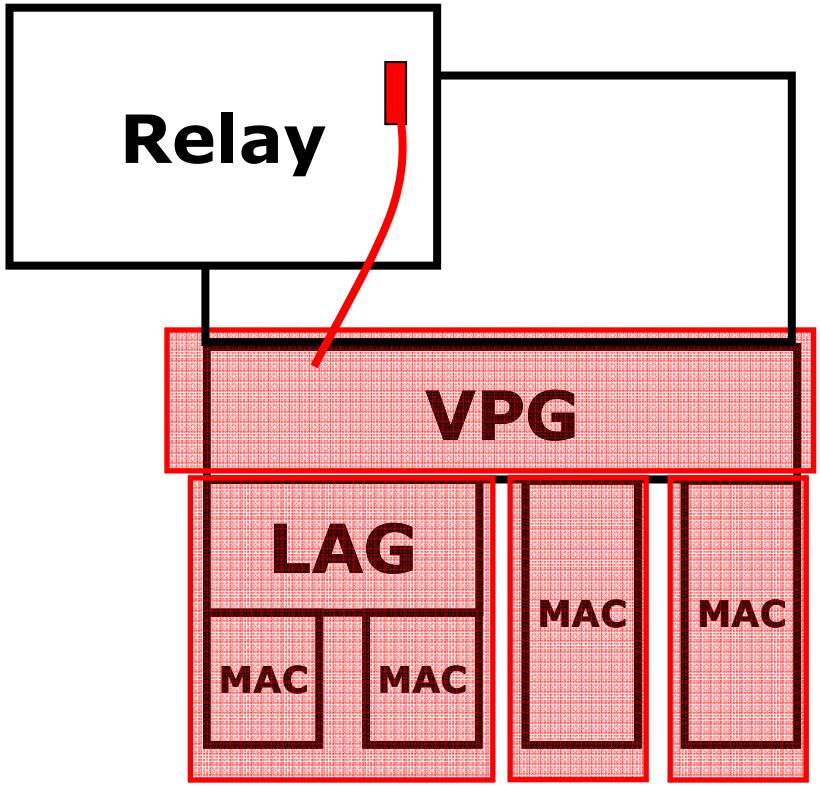

VPG functionality (cont'd)

A VPG can accept, reject or tunnel service traffic, and can be in one of the following states:

- Down – does not relay packets between bridge ports. Received packets are dropped
- Initialized – relays packets between bridge ports. (However, all the underlying ports are in the standby state)

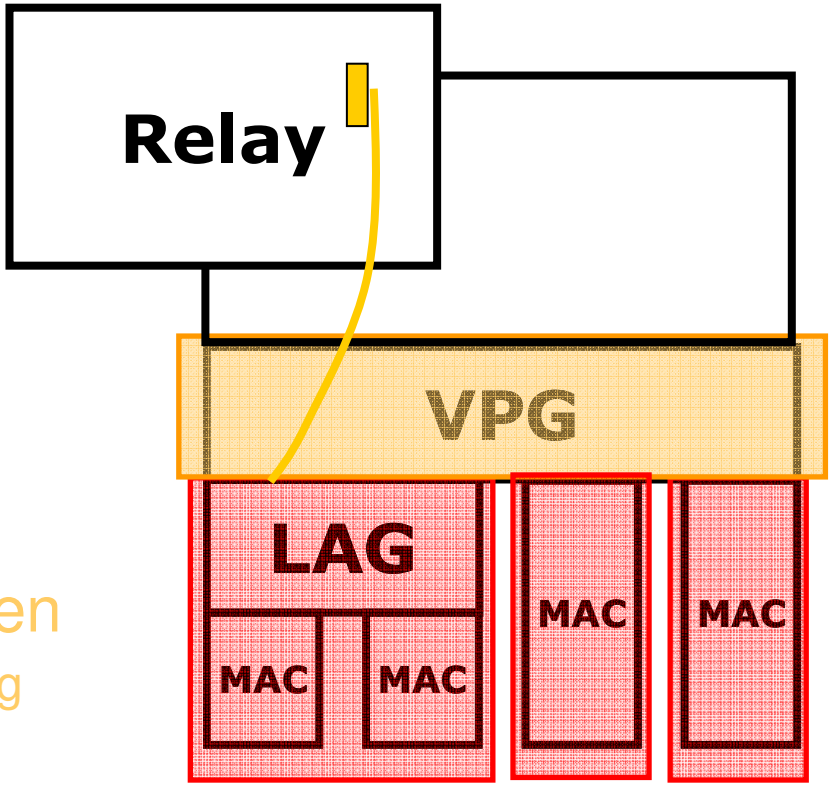

VPG functionality (cont'd)

A VPG can accept, reject or tunnel service traffic, and can be in one of the following states:

- Down – does not relay packets between bridge ports. Received packets are dropped
- Initialized – relays packets between bridge ports. (However, all the underlying ports are in the standby state)
- Up – relays packets between Bridge Ports. Has one active underlying port (can be internal or external)

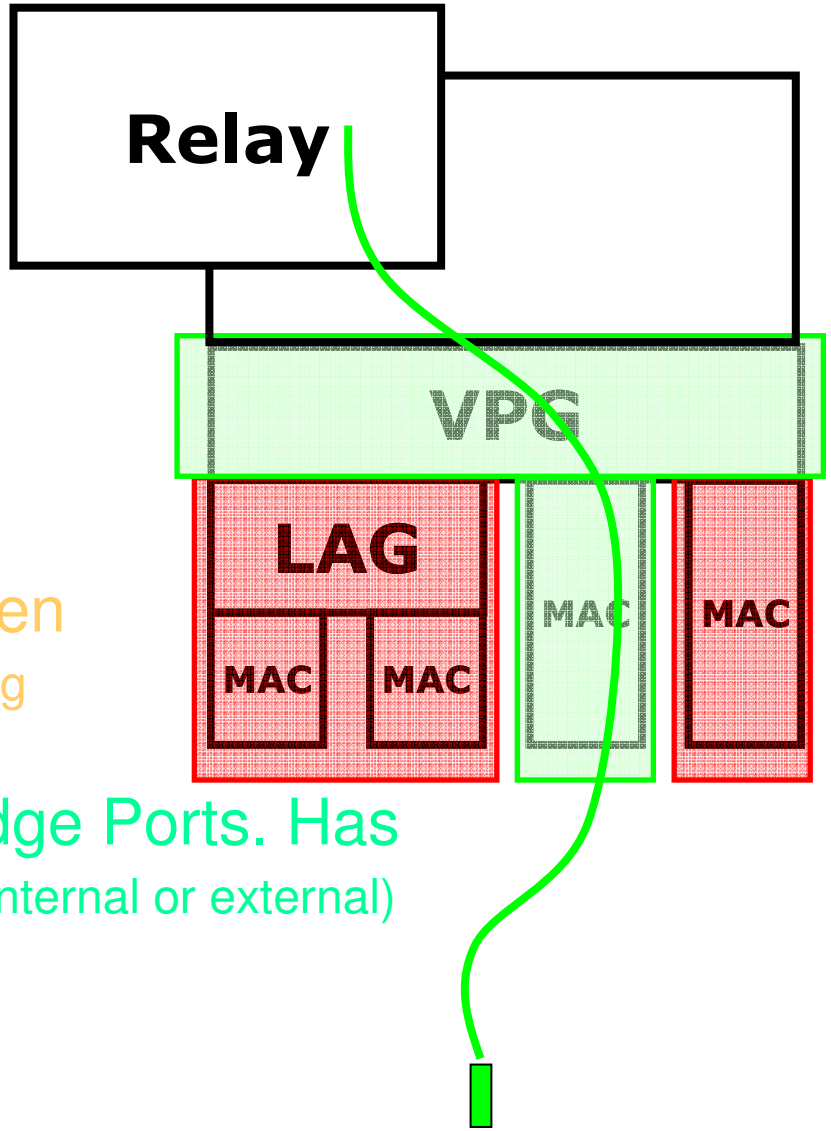

VPG functionality (cont'd)

A VPG can accept, reject or tunnel service traffic, and can be in one of the following states:

- Down – does not relay packets between bridge ports. Received packets are dropped
- Initialized – relays packets between bridge ports. (However, all the underlying ports are in the standby state)
- Up – relays packets between Bridge Ports. Has one active underlying port (can be internal or external)
- Tunnel – does not relay traffic between Bridge Ports. Relays traffic between two active underlying ports (one external and one internal)

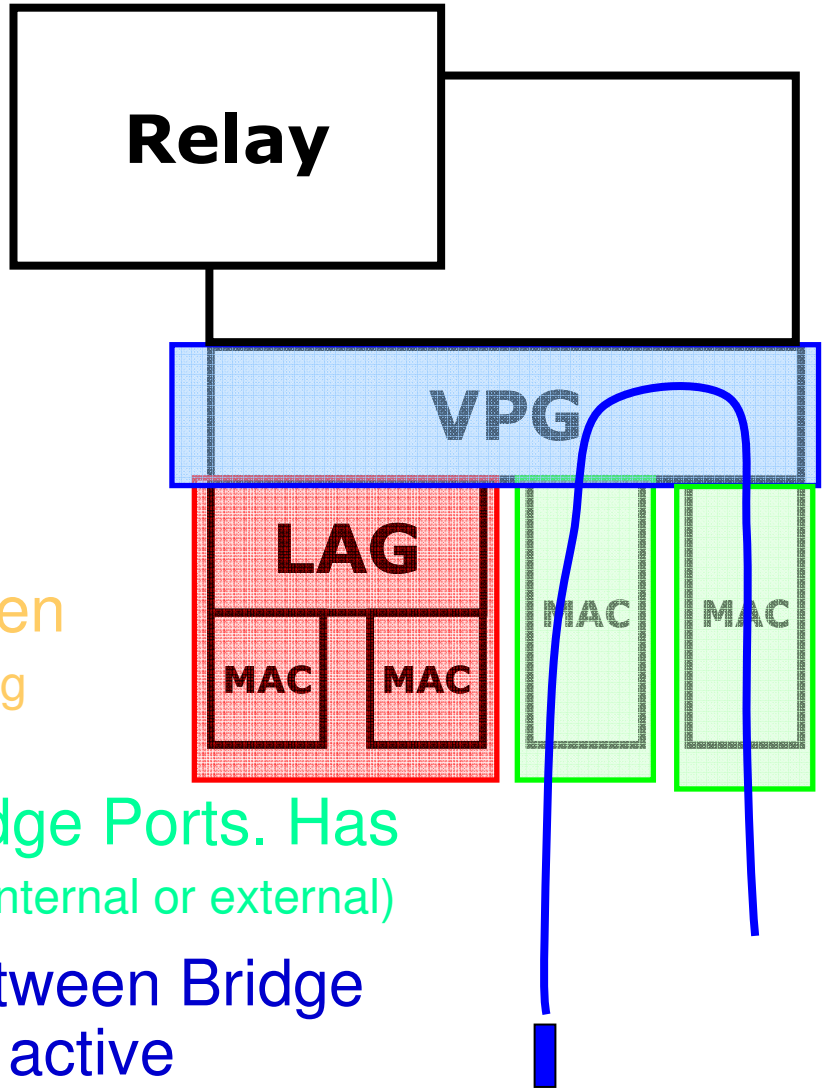

INSP protocol

- The INSP PDU contains the VPG state and the port state of the port from which it was sent
- INSP PDUs are sent by every underlying port
- The VPG may change its state according to messages received on its ports

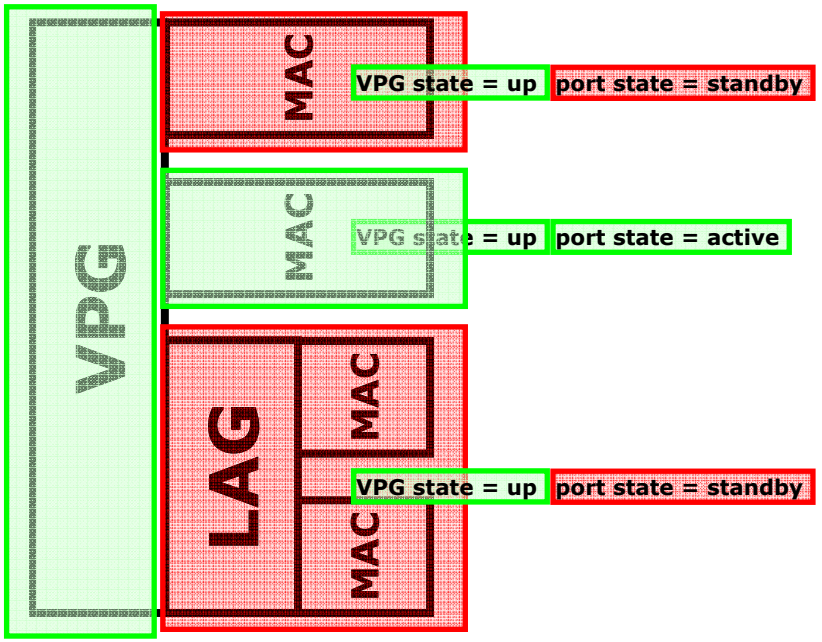

INSP functionality

- Service traffic is received from the network and relayed to the IZ, and vice versa, by a single selected border node in the service portal; this node is known as the service gateway (SG).
- The VPG on the SG selects one underlying port to transmit and receive VLAN traffic. A port is selected according to the following ordered rules (where rule # 1 represents the port of choice):
  1. The external port directly connected to the peer SG, otherwise
  2. One of the available external ports, otherwise
  3. One of the available internal ports
- If co-routing is required, the VPGs belonging to the peer SGs decide jointly which link is to be used.

Assumptions

- It is assumed that all border nodes in a service portal are leaves, i.e. a border node will not receive the same packet from the domain on more than one port. The INSP mechanism ensures that the packet is sent to the IZ by one border node only (as it may be received from the domain by more than one border node).
- Moreover, it is assumed that a packet from the domain can potentially reach all the border nodes in the service portal. This means that if the SG fails, another border node in the service portal can become the SG for that service, since it can be reached by the service traffic.

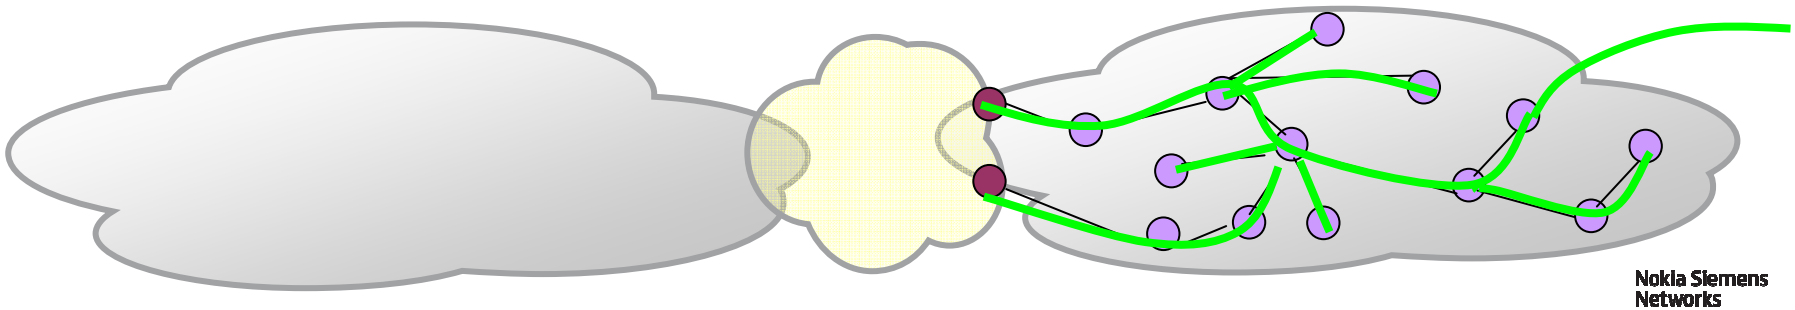

INSP advantages

- Provides a standard means of service protection over external interfaces (UNIs and ENNIs) with sub 50ms recovery time
- Avoids packet duplication
- Avoids packet misordering
- One border node can participate in several service portals
- Guarantees the prevention of fault propagation beyond the IZ in all topologies that support this functionality
- Scalable, as all the protected VLANs are monitored by one message
- All nodes participating in INSP are managed and controlled in a similar way; defining a management and control model is consistent and not problematic.

INSP advantages (cont'd)

- All services can be supported by one dedicated, internal link without encapsulation. They can also be supported on a shared link with one encapsulation type for all services (to distinguish them from domain traffic) – there is no need for a dedicated, internal link for each and every service.
- One protocol is used on all links – external and internal. As the same protocol is used inside and outside the portal, the mechanism allows interoperability, since it is capable of operating in a mixed environment (does not assume that a single vendor will be used).
- Enables the gradual upgrade of border nodes in a portal.
- Enables the replacement of a failed node with any standard bridge.

INSP advantages (cont'd)

- Intra-domain protocols, mechanisms and configurations should not be modified to support the INSP. All protocols work as usual and do not require modification.
- Supports arbitrary topologies including mesh (full and partial) and ring
- Support physical ports as well as LAG, so that service BW can be increased as the service grows
- Provides simple network expansion through the addition of a domain as service protection using INSP only requires the configuration of border nodes.
- Can be designed as part of the CFM suite with a new message type or a TLV on CCM

Thank You

zehavit.alon@nsn.com

Backup Slides

Possible entities states and messages

Node state which can be

- Down
- Initialized
- Up
- Tunnel

Port state which can be

- Active
- Standby

Node State	Port State	Message	symbol	Functionality
Up	Active	Active	A	Node is SG using this port
Up	Standby	Operational	O	Node is SG using another port
Initialized	Standby	Initialized	I	Node is SG but there is no active port yet
Tunnel	Active	Tunnel	T	Node is not SG. It uses this port and another port to tunnel traffic between internal and external ports inside the IZ
Tunnel	Standby	Bypass	B	Node is not an SG. It uses two other ports to tunnel traffic between ports
Down	Standby	Standby	S	Not an SG
		Down	D	No connectivity on this port
		Absent	Ab	The port is not configured, the link is not present
Reset Request		Reset	R	Signals to the adjacent node to move to idle. (This is only needed for the hourglass configuration, where there is no connectivity inside the service portal)

INSP node states and messages

This table summarizes the possible states in the state machine. The following are defined: node states, port states and the message each port sends over its link.

Forwarding Status of / State	Node State	External Port (x)	Message sent by External Port (x)	Internal Port (y)	Message sent by Internal Port (y)	Message sent by all other ports
IDLE (Not SG and do not convey traffic)	down	standby	S (standby)	standby	S (standby)	S (standby)
INIT (SG but do not convey traffic)	init (SG)	standby	I (init)	standby	I (init)	I (init)
ExtP(x) (SG using External Port x)	up (SG)	active	A (active)	standby	O (operational)	O (operational)
IntP(y) (SG using Internal port y)	up (SG)	standby	O (operational)	active	A (active)	O (operational)
TNL (x,y) (Not SG. tunnel between Internal port X and External port Y)	tunnel	active	T (tunnel)	active	T (tunnel)	B (bypass)
Reset (only in control node in hour glass configuration)	down	standby	R	standby	S (standby)	S (standby)

Control node state machine

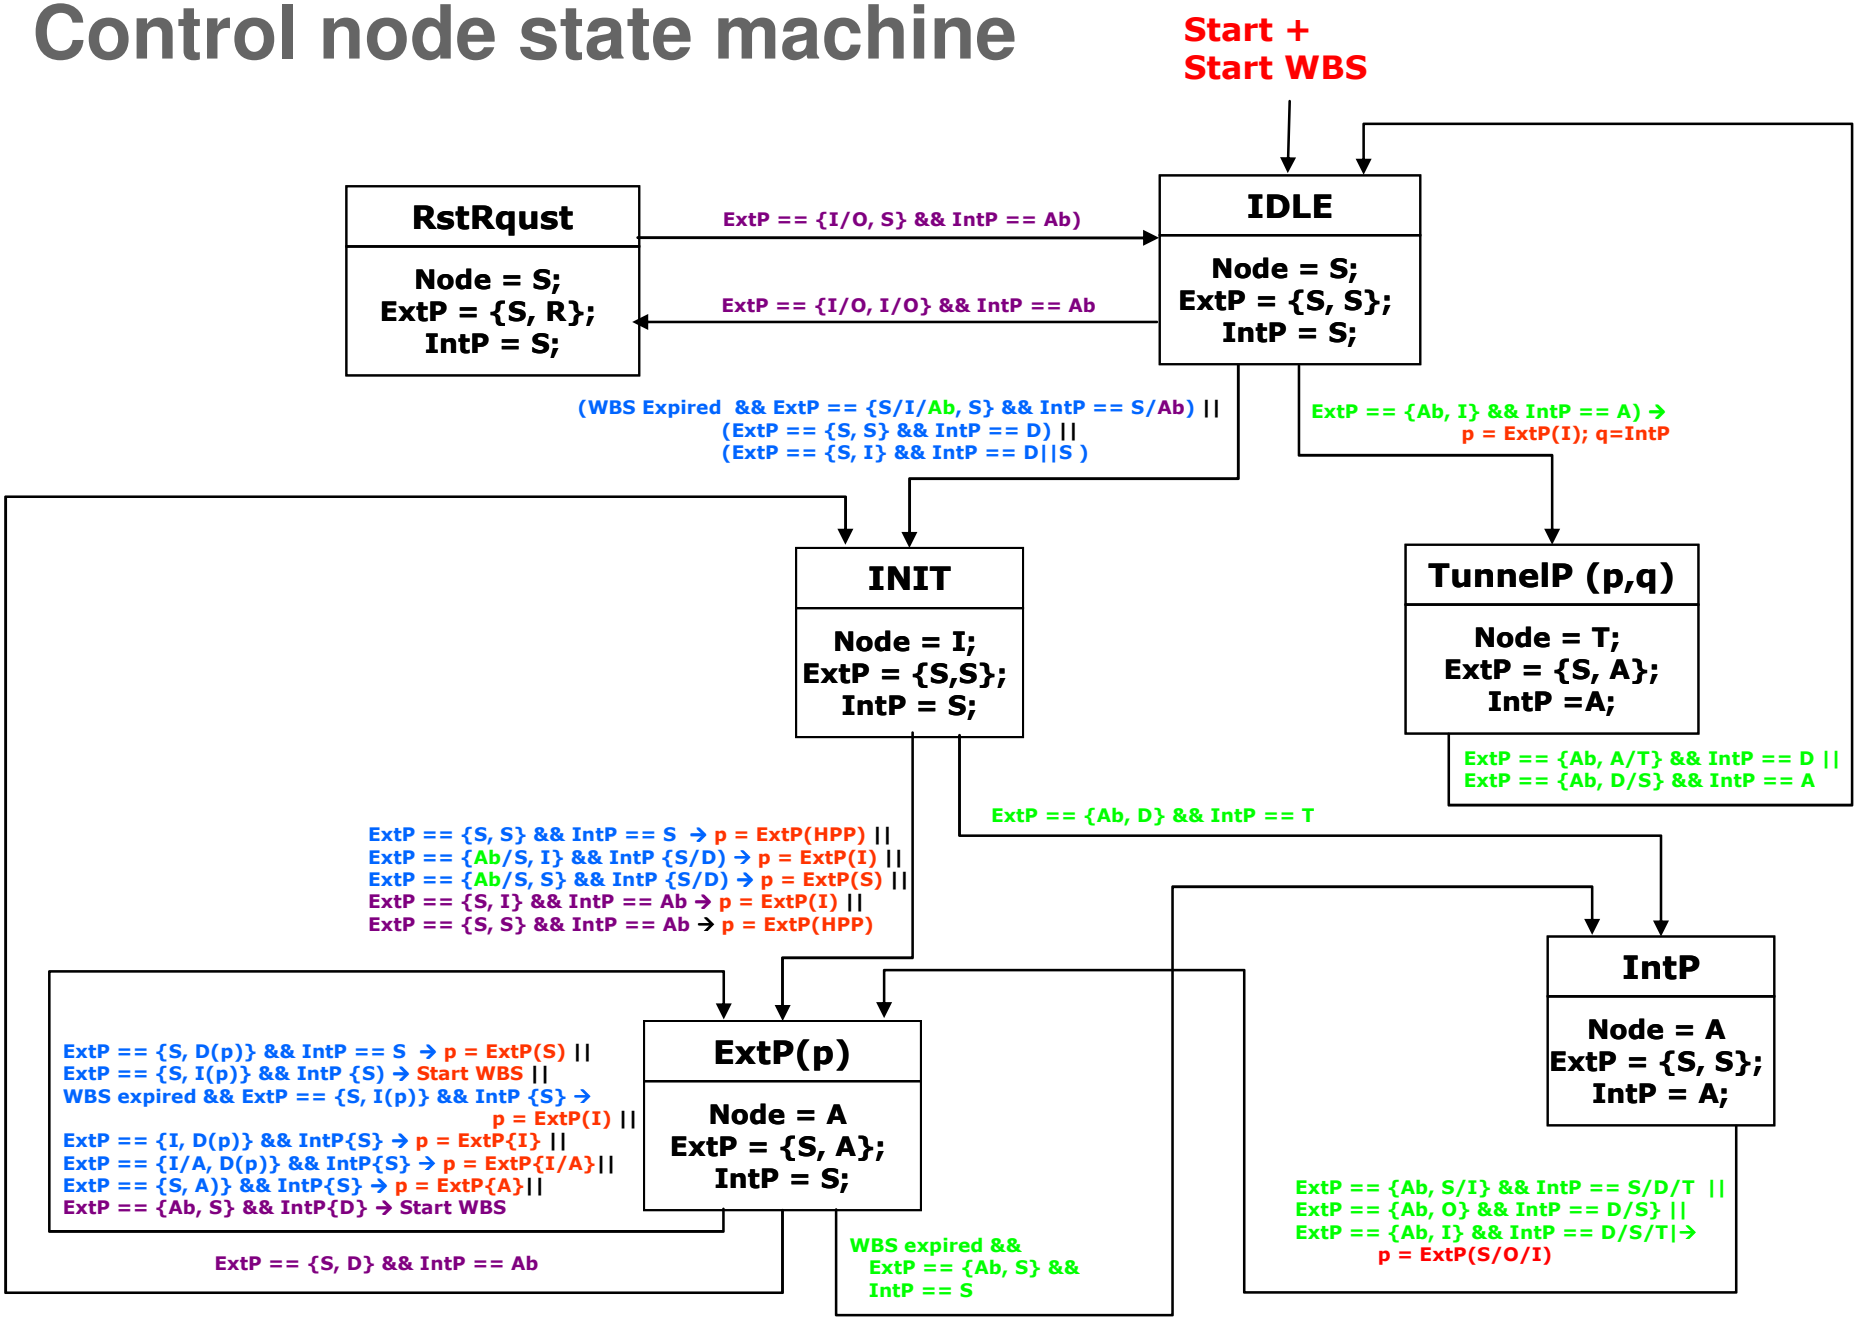

Slave node state machine