A framework for defining an IEEE 802.1 NNI

Version 1

Norman Finn

Cisco Systems

References

- This presentation is: http://www.ieee802.org/1/files/public/docs2009/new-nfinn-nni-framework-1209-v01.pdf

Definitions

Definitions

- Network: The global interconnected set of links and nodes.
- Link: A point-to-point or multipoint-to-multipoint connection among two or more nodes.
- Node: A bridge or switch that is common to exactly two subnetworks, and belongs to exactly one portal.
- Subnetwork: A subset of the links and nodes of a network intended to provide connectivity for services among some number of portals.

Definitions

- Portal: A set of one or more nodes all of which belong to the same two subnetworks.
- UNI: A portal to a customer.
- NNI: A subnetwork with exactly two portals.
- Cloud: A subnetwork with two or more portals.
- Service: A path from terminus to terminus through the network carrying data for a customer.
- Terminus: The point, always in a node, at which a service transfers from one subnetwork to another.

Definitions

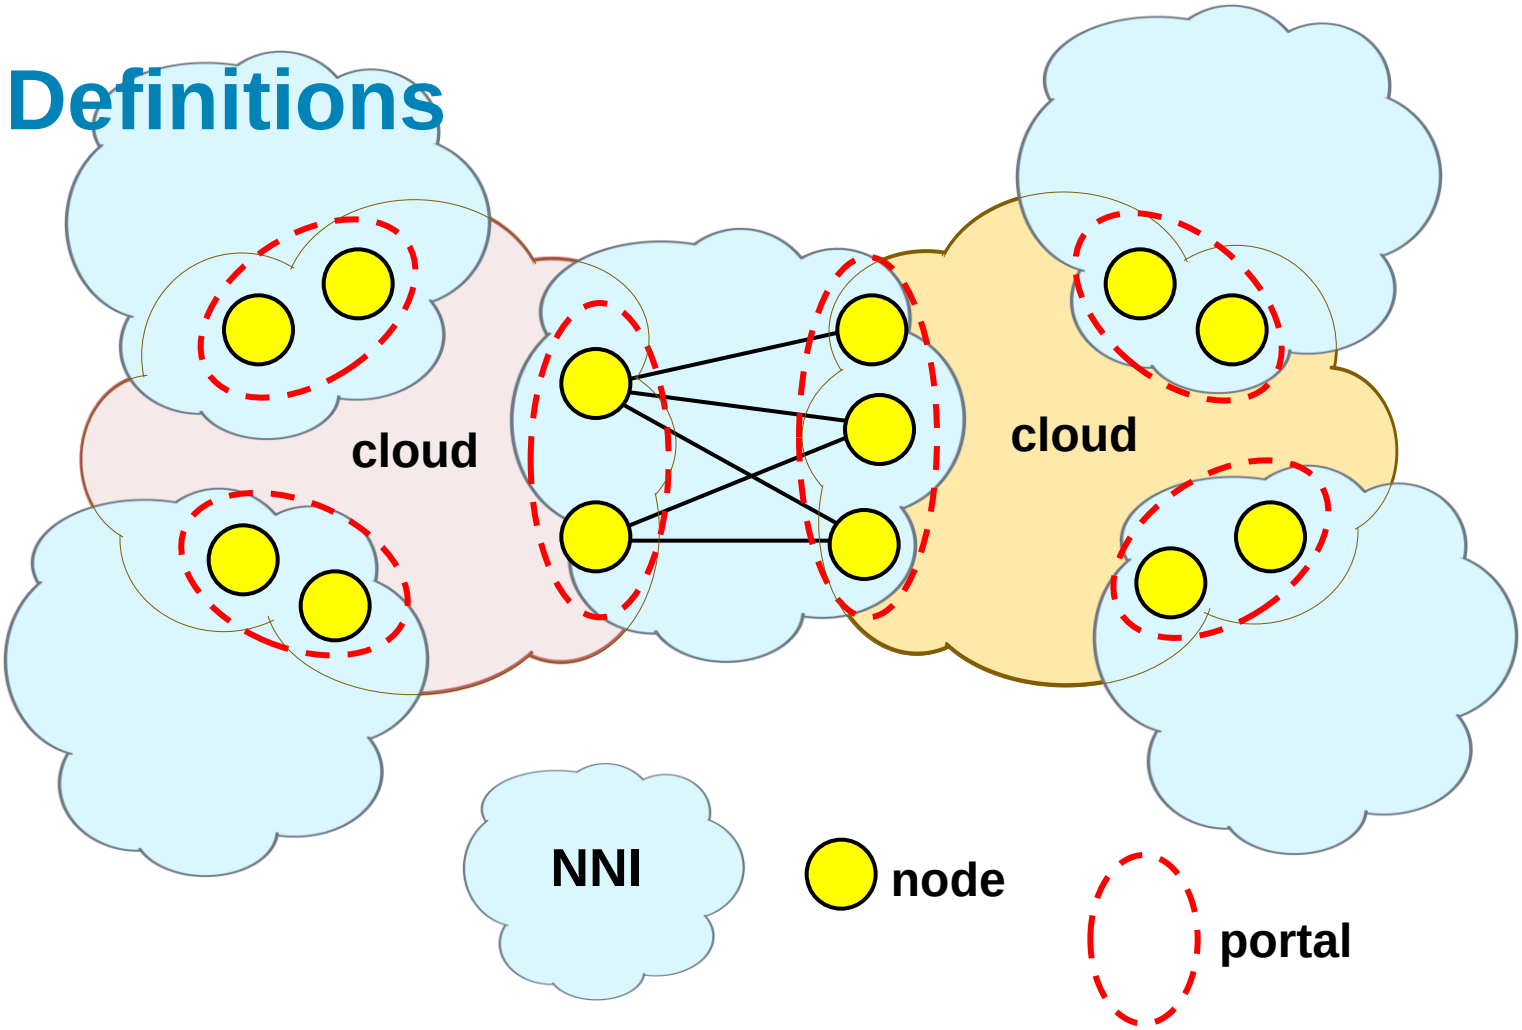

- Two clouds connected via an NNI, with other NNIs to four other clouds (not shown).

Principles

- A subnetwork provides connectivity for every service, with the ability to connect a single terminus in one portal to a terminus in any other portal (for a point-to-point service) or to a terminus in every other portal (for a multipoint-to-multipoint service).
- A portal connects two subnetworks. Every service that passes from one subnetwork to another passes through a portal. In exactly one of the nodes in that portal is a terminus for that service. The terminus is the point at which the services passes from one subnetwork's control to the other.
- A service is carried over links from node to node. A service may pass through a node without having a terminus. It may not pass from one subnetwork to another through a portal without a terminus, and it must not have more than one terminus in a single portal.

Principles

- The path of a service through a network is configured in terms of subnetworks and portals.
- The choice of which links are used to carry a service from terminus to terminus across a subnetwork is made by the routing protocol(s) used by that subnetwork, independently of any other subnetwork.
- The choice of which of the nodes in a portal the terminus for each service resides is made jointly by the routing protocol(s) of the two networks that share the nodes in the portal connecting them.

Connectivity and fault tolerance

- It is easy for a subnetwork to provide connectivity from one portal to another, if that subnetwork's routing protocol has the freedom to make terminus-to-node assignments. (Two 2-node portals with one link between the upper nodes and one link between the lower nodes.) For example, only two links are required to provide against a single failure between two portals.
- But, if every subnetwork is that simple, then when you connect those subnetworks together, the whole network can ensure connectivity for only a single failure in the whole network. There is no isolation of faults.
- In order to ensure connectivity against a single failure of a node or link, if the choice of terminus-to-node assignment is in the hands of the other subnetworks, requires 2^n links to interconnect n nodes.

Coupled protection

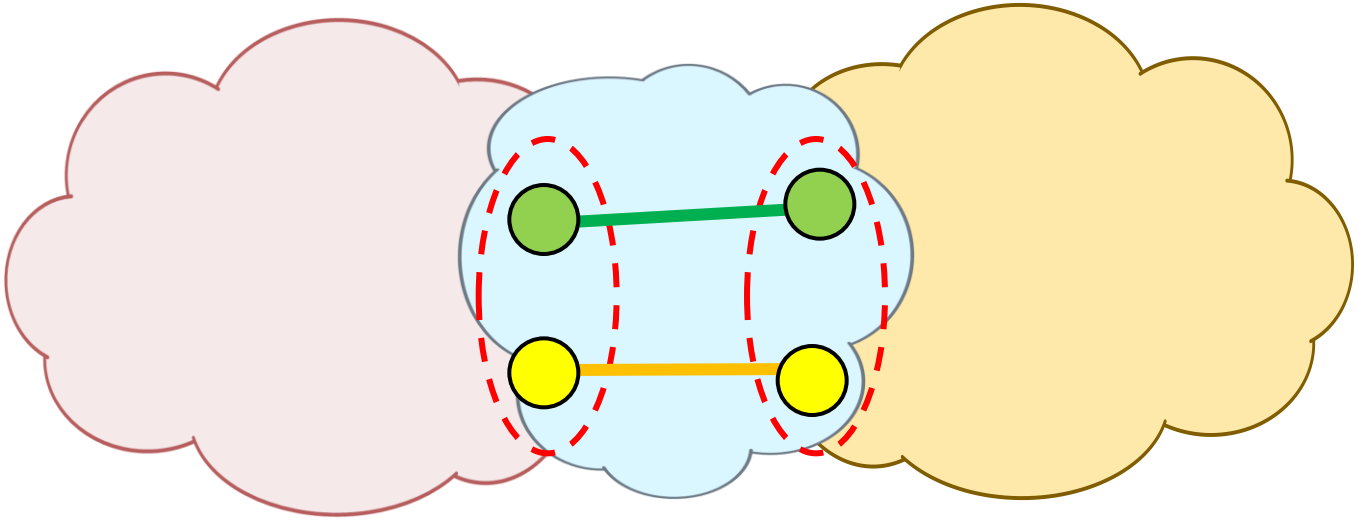

- Simple NNI – four nodes, two links.
- It DOES guarantee to connect left portal to right portal, even if one link or one node fails.
- But, it couples failures.

Coupled protection

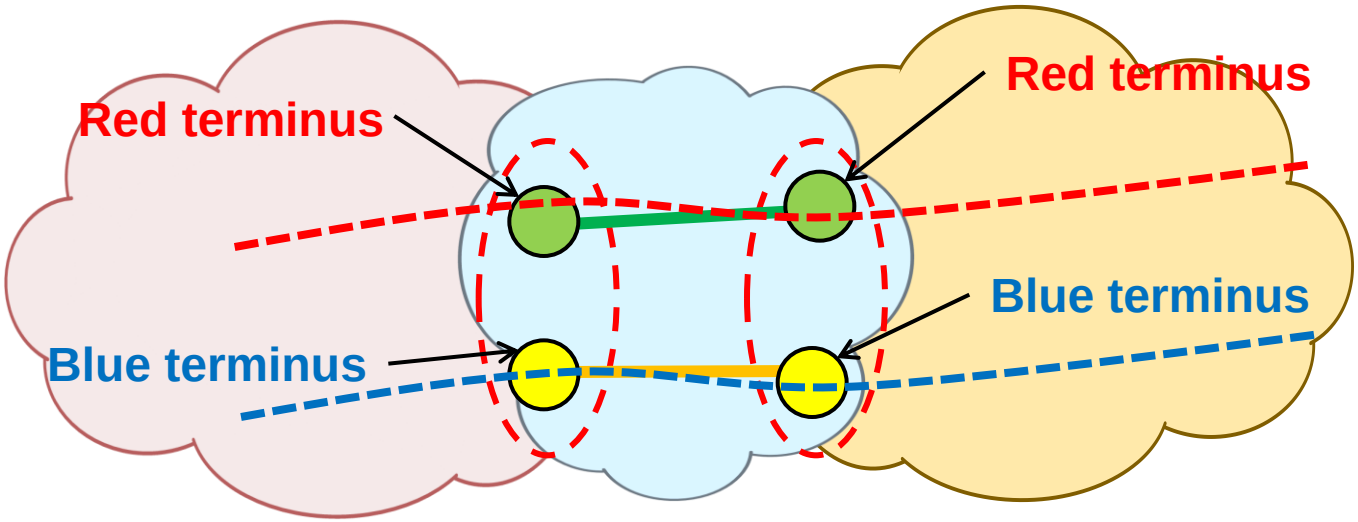

- Each service (red or blue) is guaranteed connected from left portal to right portal, even if the green link fails.
- BUT, if the green link fails, the Red terminus must change in both the left and right clouds.
- Moving the red terminus in the right cloud forces an error recovery condition on the right cloud.

Coupled protection

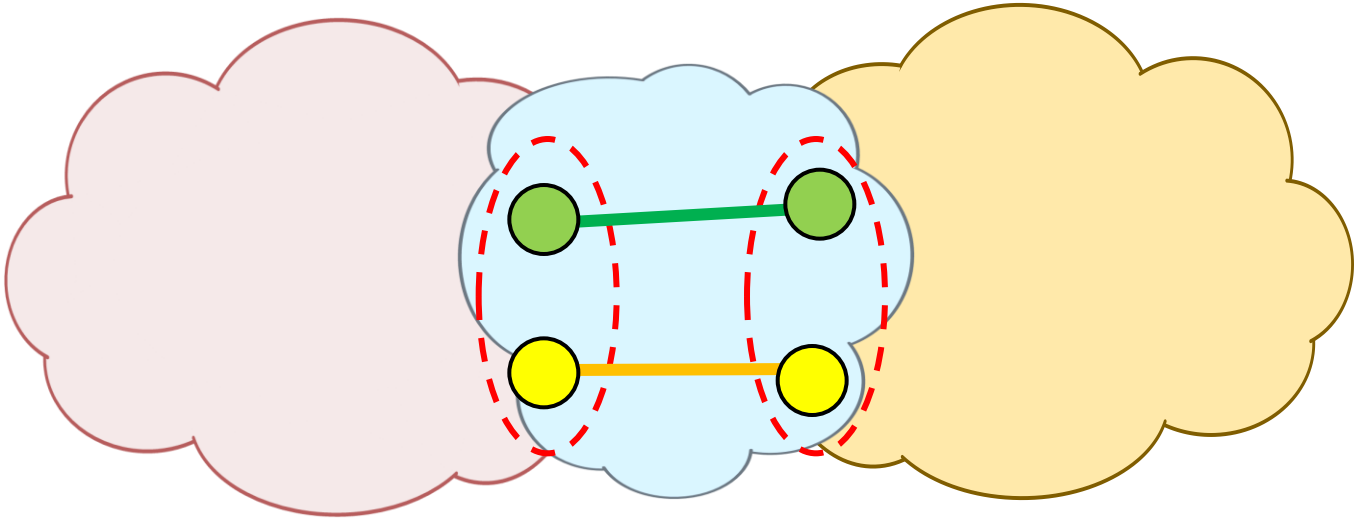

- So, this simple-minded NNI is inadequate.

Coupled protection

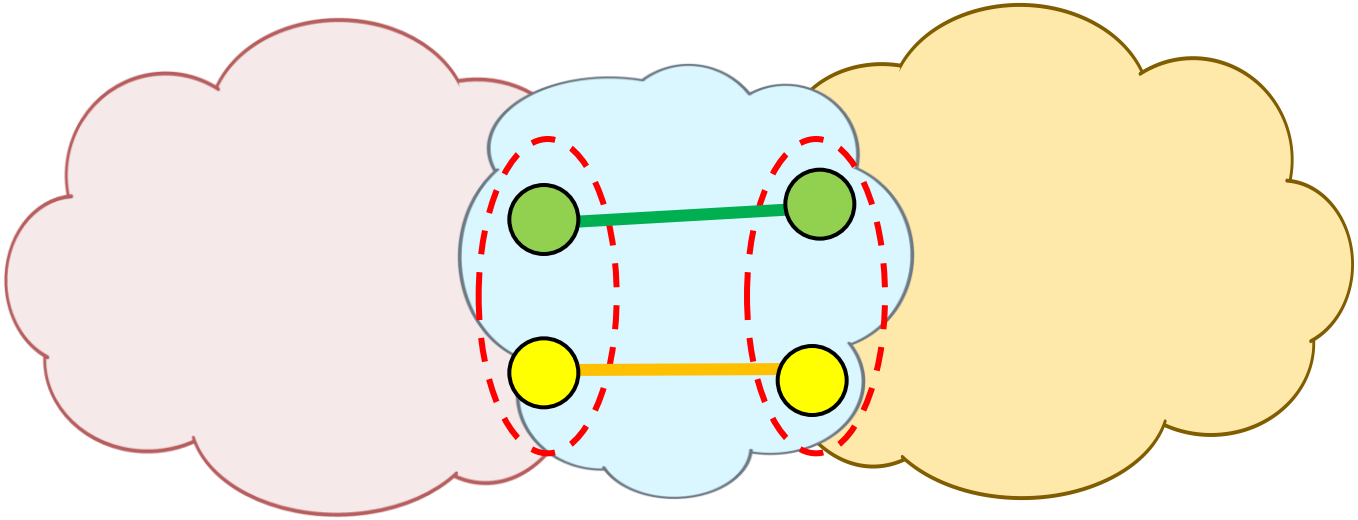

- So, **this simple-minded NNI is inadequate**.

Coupled protection

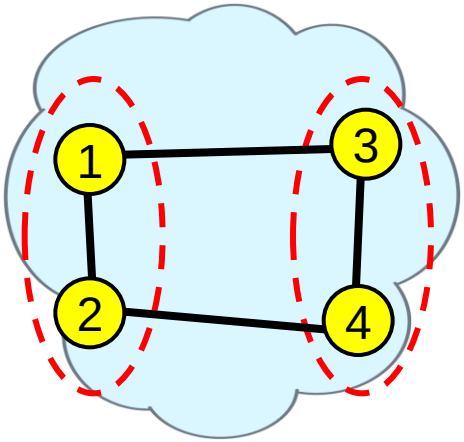

- But, **this interconnect**, along with **eight protected segments** (1-3, 1-2-4-3, 1-3-4, 1-2-4, 2-4, 2-1-3-4, 2-1-3, and 2-4-3) can provide protection against any link failure **without triggering an unnecessary change of terminus** in ether cloud.
- Note that four nodes with two between-cloud links requires eight segments and two real or virtual intra-could links (1-2 and 3-4) available for use solely by NNI.

Coupled protection

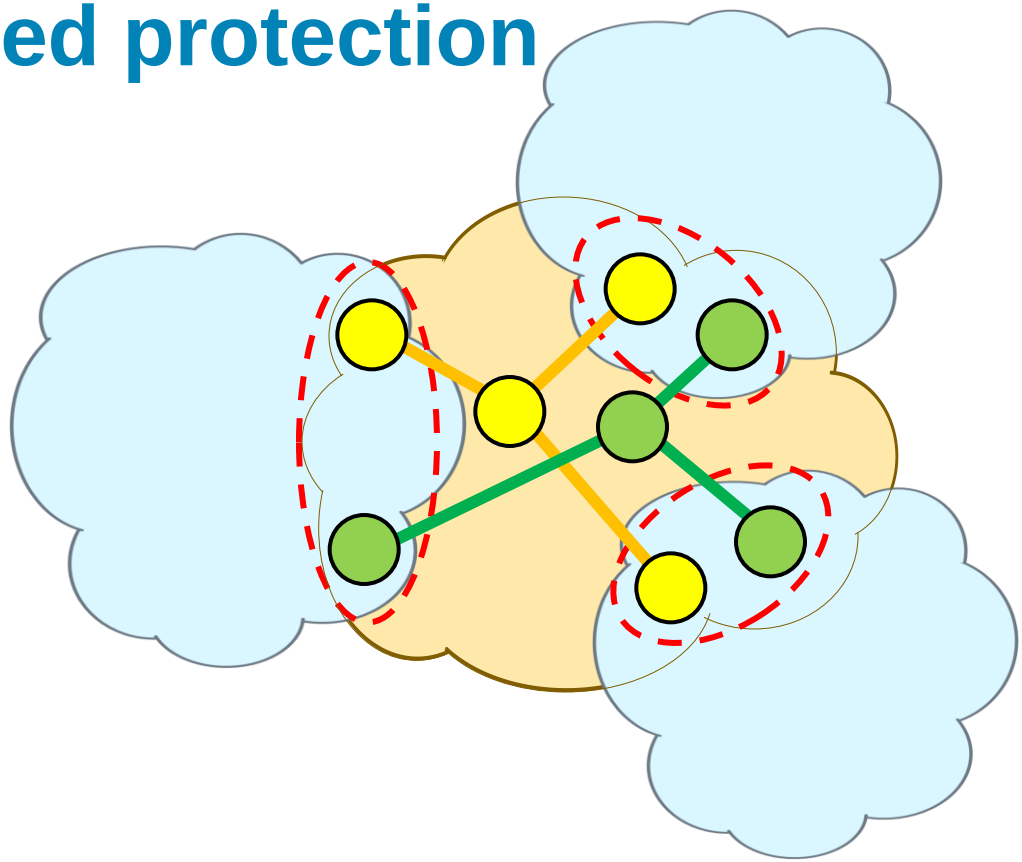

- Two separate redundant multi-node nets.
- If a link/node fails, its whole green or yellow net fails.
- Failure of one Portal's node affects all portals.

Coupled protection

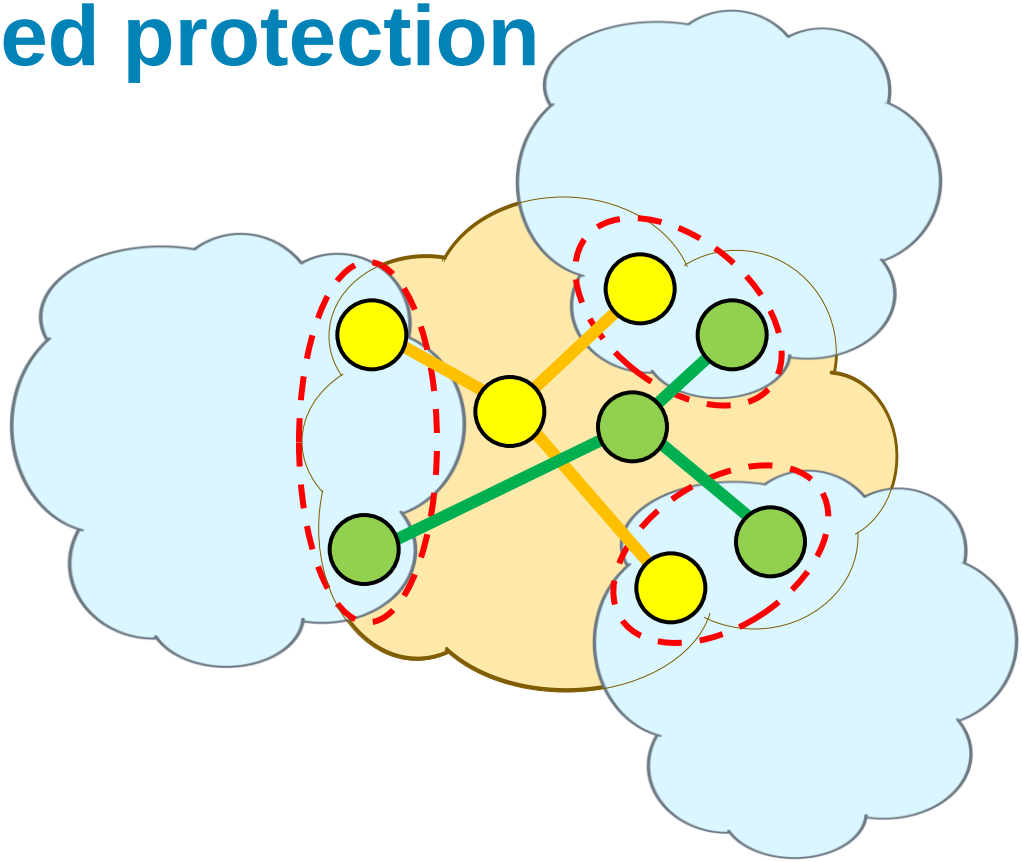

- Proofing a 3-portal, 2-nodes per portal network against coupling protection faults requires **vastly more** segments than the simple NNI.

Connectivity and fault tolerance

- That is, a subnetwork with two portals, each of two nodes, requires four real or virtual links among the nodes to guarantee connectivity if it cannot choose in which node each terminus is placed.
- A subnetwork with four two-node portals requires 16 real or virtual links among those nodes.
- A typical carrier network can have very many portals to customers and/or other providers. Its ability to provide connectivity can be greatly simplified if it can make the terminus-to-node assignment.

NNI

- Two issues force us to divide subnetworks into two classes:
  - The fact that two differenet networks' routing protocols must jointly make the decision for assigning termini to nodes suggests that an $O(p^2)$ interoperation problem among p routing protocols can be reduced to an $O(p)$ problem by creating a standard subnetwork with a standard routing protocol, called an "NNI."
  - There are massive differences in the scale of the problem of ensuring fault-tolerant connectivity in a subnetwork with two portals, and providing those features in a subnetwork with more than two portals, if control of the assignment of termini to nodes is surrendered.

NNIs and clouds

- Therefore, we must have cloud subnetworks, that:
  - Have any number of portals;
  - Run any routing protocol; and
  - Decide in which node of the portal each service's terminus resides; and
- NNI subnetworks, that:
  - Have exactly two portals;
  - Run the IEEE standard NNI routing protocol; and
  - Leave the assignment of terminus to node in the hands of the routing protocol of the cloud subnetwork that shares control of its nodes.

NNI protocol requirements

Among the requirements for an NNI protocol, therefore, we must list:

- The NNI controls a network consisting of nodes and links.
- All of the nodes belong to portals, so control of every node is shared between the NNI and the routing protocol of the neighboring cloud.
- The links and nodes may be virtual or real, but for the purposes of the NNI, can be considered as if they were physical. That is, no intermediate switches or bridges between the nodes need be considered.

NNI protocol requirements

- The ability to react to changes in the availability of the links and nodes in the NNI subnetwork is obvious.
- What is not so obvious is that the NNI protocol's interaction with its two neighbor cloud subnetworks' routing protocols, at least in terms of reaction to faults, is limited to the cloud protocols' terminus-to-node assignment decisions.
- The NNI protocol must react to terminus reassignments with the same alacrity that it reacts to link and node failures.
- **In short, clouds get to propagate faults to NNIs as terminus reassignments, but NNIs don't get this privilege.**