Why LACP for the NNI protocol?

Using LACP for the NNI problem

Rev. 1

Norman Finn

nfinn@cisco.com

Ethernet NNI

- This contribution is available at: new-nfinn-why-LACP-for-NNIl-0111-v01.pdf.
- See also the contributions: new-nfinn-LACP-vs-buffer-networks-1110-v1.pdf and new-haddock-Distributed-LAG-Models-1010-v2.pdf.
- The purpose of this contribution is to show why the 802.1AX Link Aggregation Control Protocol is a good choice for the Resilient Network Interface protocol.

Background

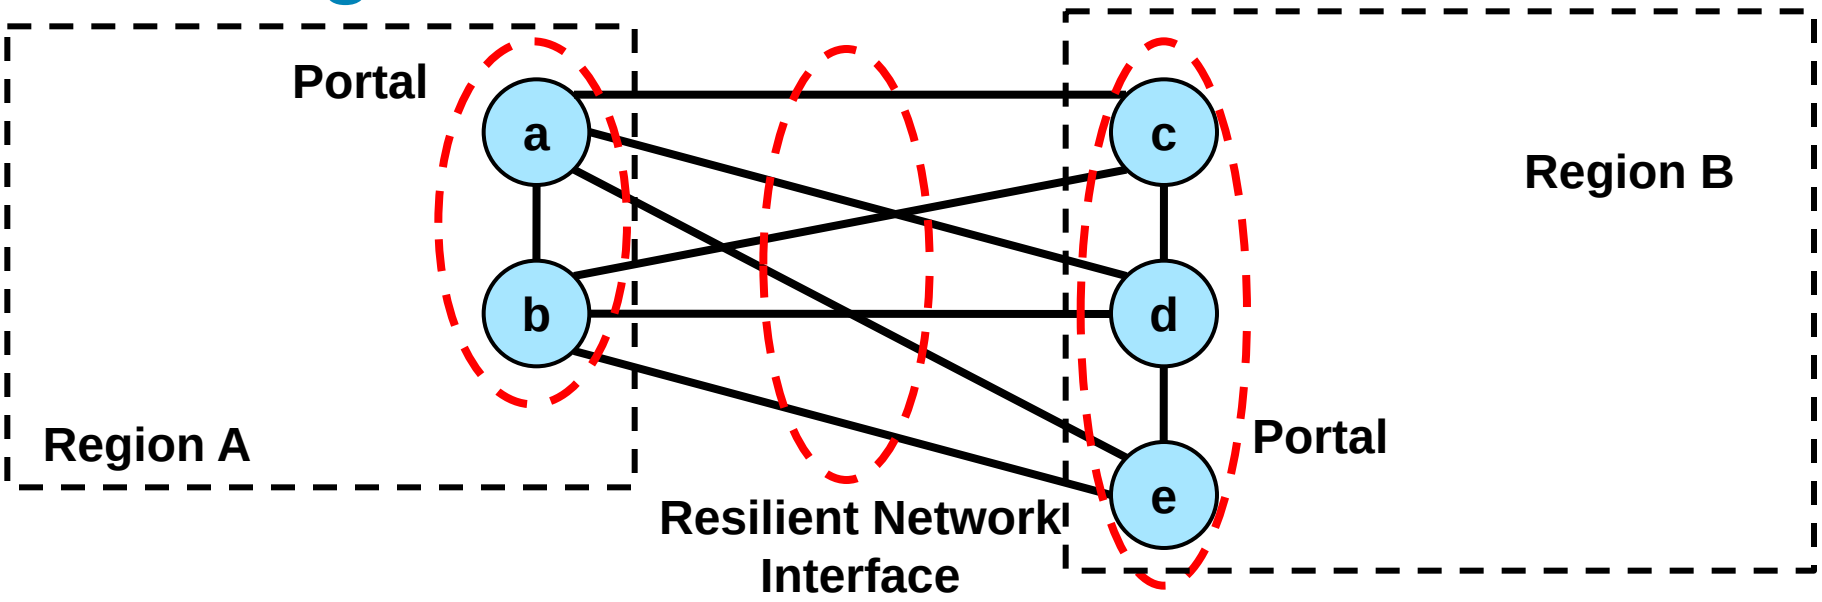

- As shown in new-nfinn-LACP-vs-buffer-networks-1110-v1.pdf, the basic architectural blocks, control flows, and data flows for a Resilient Network Interface are largely determined by the problem requirements, independently of the choice of a particular control protocol.

Protocol choices

As set out in new-nfinn-LACP-vs-buffer-networks-1110-v1.pdf, there are separate (but related) choices to be made for data encapsulation and for control protocol for a Resilient Network Interface.

- There appear to be two choices for the data plane:
  1. Data can carry an extra encapsulation across the RNI.
  2. Data can use the encapsulation of either attached network.
- There appear to be three choices for the control plane:
  1. Extensions to Link Aggregation Control Protocol (802.1AX)
  2. Extensions to Connectivity Fault Management (802.1ag)
  3. A new protocol invented for this purpose.

Protocol choices

- Because the only proposals for using CFM have (so far) also required an additional encapsulation across the RNI, and because this added encapsulation has issues when connecting 802.1ad (Q-in-Q) networks, this control plane choice has not been popular in the regular RNI Webex conferences.
- This leaves two choices:
  1. Extensions to Link Aggregation Control Protocol (802.1AX), perhaps something similar to those proposed by Norm Finn and Steve Haddock.
  2. A new protocol invented for this purpose, perhaps something similar to that proposed by Zehavit Alon.
- Both of these choices can use the same data plane.

Why LACP?

Compatibility

- Similar objectives of RNI and Link Aggregation.
  - Both utilize multiple physical links to provide a single logical link between exactly two entities.
  - Both leave it to higher layers to resolve a situation where multiple logical links are present that cannot be aggregated.
  - Both can be used to connect any kind of device, whether switch, end station, router, or other device.
  - Both must offer fast recovery from failures. (Restoration of server can take longer, for both.)
- What is different:
  - In the RNI, each of the two connected entities may consist of multiple physical switches. (This drives the requirements for extensions to LACP.)

Special case advantageous to LACP

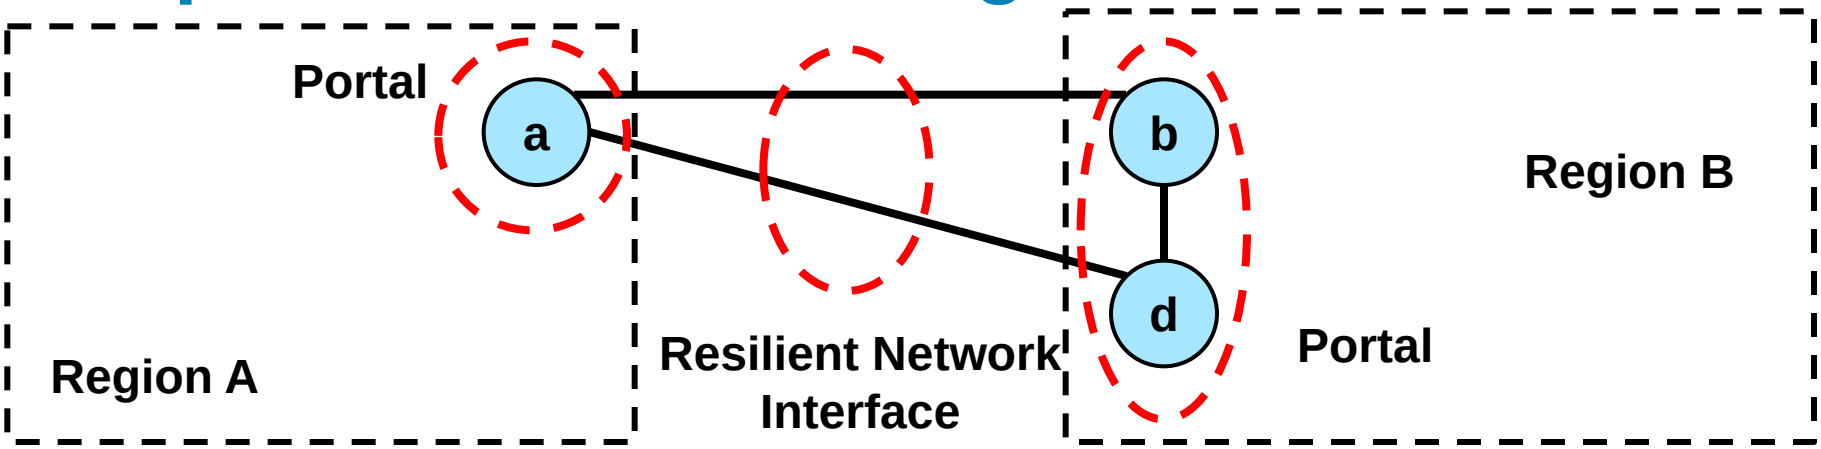

- Consider the case where one of the two connected networks has a single node in its Portal. (Perhaps this one node is the *only* node in Region A.)
- A single device (Node a, in this case) already knows a protocol for making multiple physical links behave like a single logical link – LACP!
- An RNI based on LACP that supports this scenario will have much greater applicability than a new protocol.

Current use

- Implementations by a number of vendors confirm that an extended LACP makes a satisfactory RNI.

Level of standardization

- Everything that is required to implement an RNI between two networks must be standardized.

Limitations on the RNI standard

- The links connecting the physical nodes of a Portal (intra-Portal links) are very simple if they support a single RNI.
- These links get much more complex, both in the data plane and in the control plane, if a single intra-Portal link supports multiple RNIs.
- Therefore, standardization of support by a single intra-Portal link of multiple RNIs should be deferred.