Resilient Network Interconnect: Arguments for using Link Aggregation

Stephen Haddock
Extreme Networks
January 11, 2011

Introduction

- Over the last 18 months many presentations have been given investigating various aspects of implementing resilient connectivity between two independently administered networks: a Resilient Network-Network Interface.
- One of the proposed solutions for this is based on 802.1AX Link Aggregation.
- On a conference call in January 2011 Norm presented new-nfinn-why-LACP-for-NNIl-0111-v01.pdf which listed some reasons for using LACP as a base for a resilient NNI.
- This presentation simply emphasizes some of the points Norm made.

Control Plane Options

- Norm mentioned three possible choices as a starting point from which to develop a Resilient NNI control plane:
  - LACP
  - CFM
  - A new protocol
- It has been pointed out that any protocol that provides a fully-connected loop-free topology between the gateways for each VLAN in each network would also work:
  - RSTP/MSTP
  - SPB
  - G.8032 (single or multi-ring)

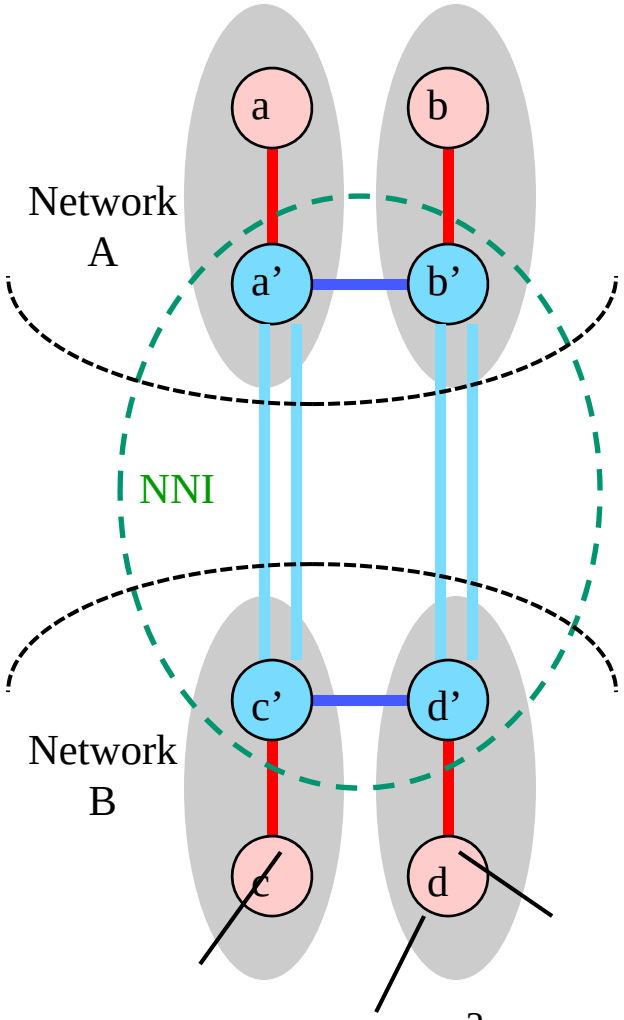

Why use Link Aggregation?

- Norm provided several reasons for using LACP as the starting point for a Resilient NNI.
- I would like to emphasize three very pragmatic issues:
  1. A highly desirable aspect of a solution based on Link Aggregation would be to maintain backwards compatibility with current Link Aggregation implementations.
  2. Many people (customer and vendors) already use proprietary distributed Link Aggregation solutions.
  3. Many people will continue to use distributed Link Aggregation solutions.

Backwards compatibility

- Supports “dual-homed” configurations for a NNI or, more likely, a UNI where the single system in one network simply runs the current standard LACP.
- If the new standard is backwards compatible with a single system running LACP, then it will be compatible with any multi-system implementation that looks like a single LACP system.
  - This includes all the current proprietary implementations of a distributed Link Aggregation (e.g. Split Multi Link Trunking (SMLT), Multi Chassis Link Aggregation (MLAG), etc.)
  - This includes implementations that make two or more systems appear to be a single bridge (stacking, virtual chassis, etc.)
- This is a very powerful migration story for getting a new standard deployed in real networks.

Current uses of LAG for NNI

1. The only currently specified protection mechanism in the Metro Ethernet Forum UNI and ENNI technical specifications is the current Link Aggregation standard.
   - This only provides protection against link failures, not node failures.
   - Extending Link Aggregation to protect against node failures is an obvious step.
2. Many actual network deployments currently use proprietary versions of a distributed Link Aggregation for a Resilient NNI solution.

- Both of these demonstrate a very high probability of market acceptance for a Resilient NNI protocol based on Link Aggregation.

Continued use of resilient LAG solutions

- Given the current use of proprietary distributed LAG solutions for NNIs, it is highly likely that this will continue in the future.
  - If we choose to develop a new protocol, this is our competition.
- Networks will continue to be designed with distributed LAG solutions in non-NNI applications.
  - If we choose to develop a new protocol, we are trying to split the market.
- There will be continued market demand for vendors to implement distributed LAG solutions.
  - If we choose to develop a new protocol we are asking vendors to develop, maintain and support a second solution addressing the same problem, and we are asking customers to learn (basic familiarization, trouble-shooting, etc.) a second solution addressing the same problem.

Conclusion

- Very pragmatic considerations lead to a conclusion that a Resilient NNI standard building on Link Aggregation has a very high probability of success, whereas developing a new protocol to compete with Link Aggregation has a very low probability of market acceptance.