


NVO3 Architecture
draft-narten-nvo3-arch-00.txt
IETF87 – Berlin
July 31, 2013

David Black, Jon Hudson, Larry
Kreeger, Marc Lasserre, Thomas
Narten

NVO3 Architecture Purpose

- Architecture identifies key system components (NVE, NVA, etc.) and how they fit together for an overall system
 - WG has discussed but not formally confirmed various decisions
- Components interact with another through well-defined interfaces
 - Interfaces between components represent “on-the-wire” protocols (i.e., potential IETF work areas)
- Internal implementation of component not IETF matter, so long as interface behavior maintained
 - Allows for independent evolution of individual components
- Architectural decisions lead to requirements, requirements feed directly into gap analysis

NVE and NVA Components



Data Plane Encapsulations

- Assertion: WG should not pick or bless one encapsulation
 - Multiple encaps exist today, deployments will have multiple encaps
- Implication: Architecture must support multiple encapsulations
 - NVEs should use common encapsulation and direct tunneling where possible
 - Traffic should flow through translating gateways when NVEs do not support same encapsulation
 - Should not require operator intervention – should just work
- Summary: Control plane must be aware of and support existence of different encapsulations on different NVEs
 - Impacts the control plane requirements

NVE-to-NVA Protocol

- Goal: NVEs should implement NVO3 functionality once, then not again
 - Many NVEs in a deployment – upgrading them will be difficult going forward
 - Future innovation/evolution will be within NVAs
 - SHOULD NOT require NVE upgrades
- Assertion: there will likely be a range of NVA types
 - Should hide details from NVE
- Implies need for well-defined NVE-to-NVA protocol with clear interface

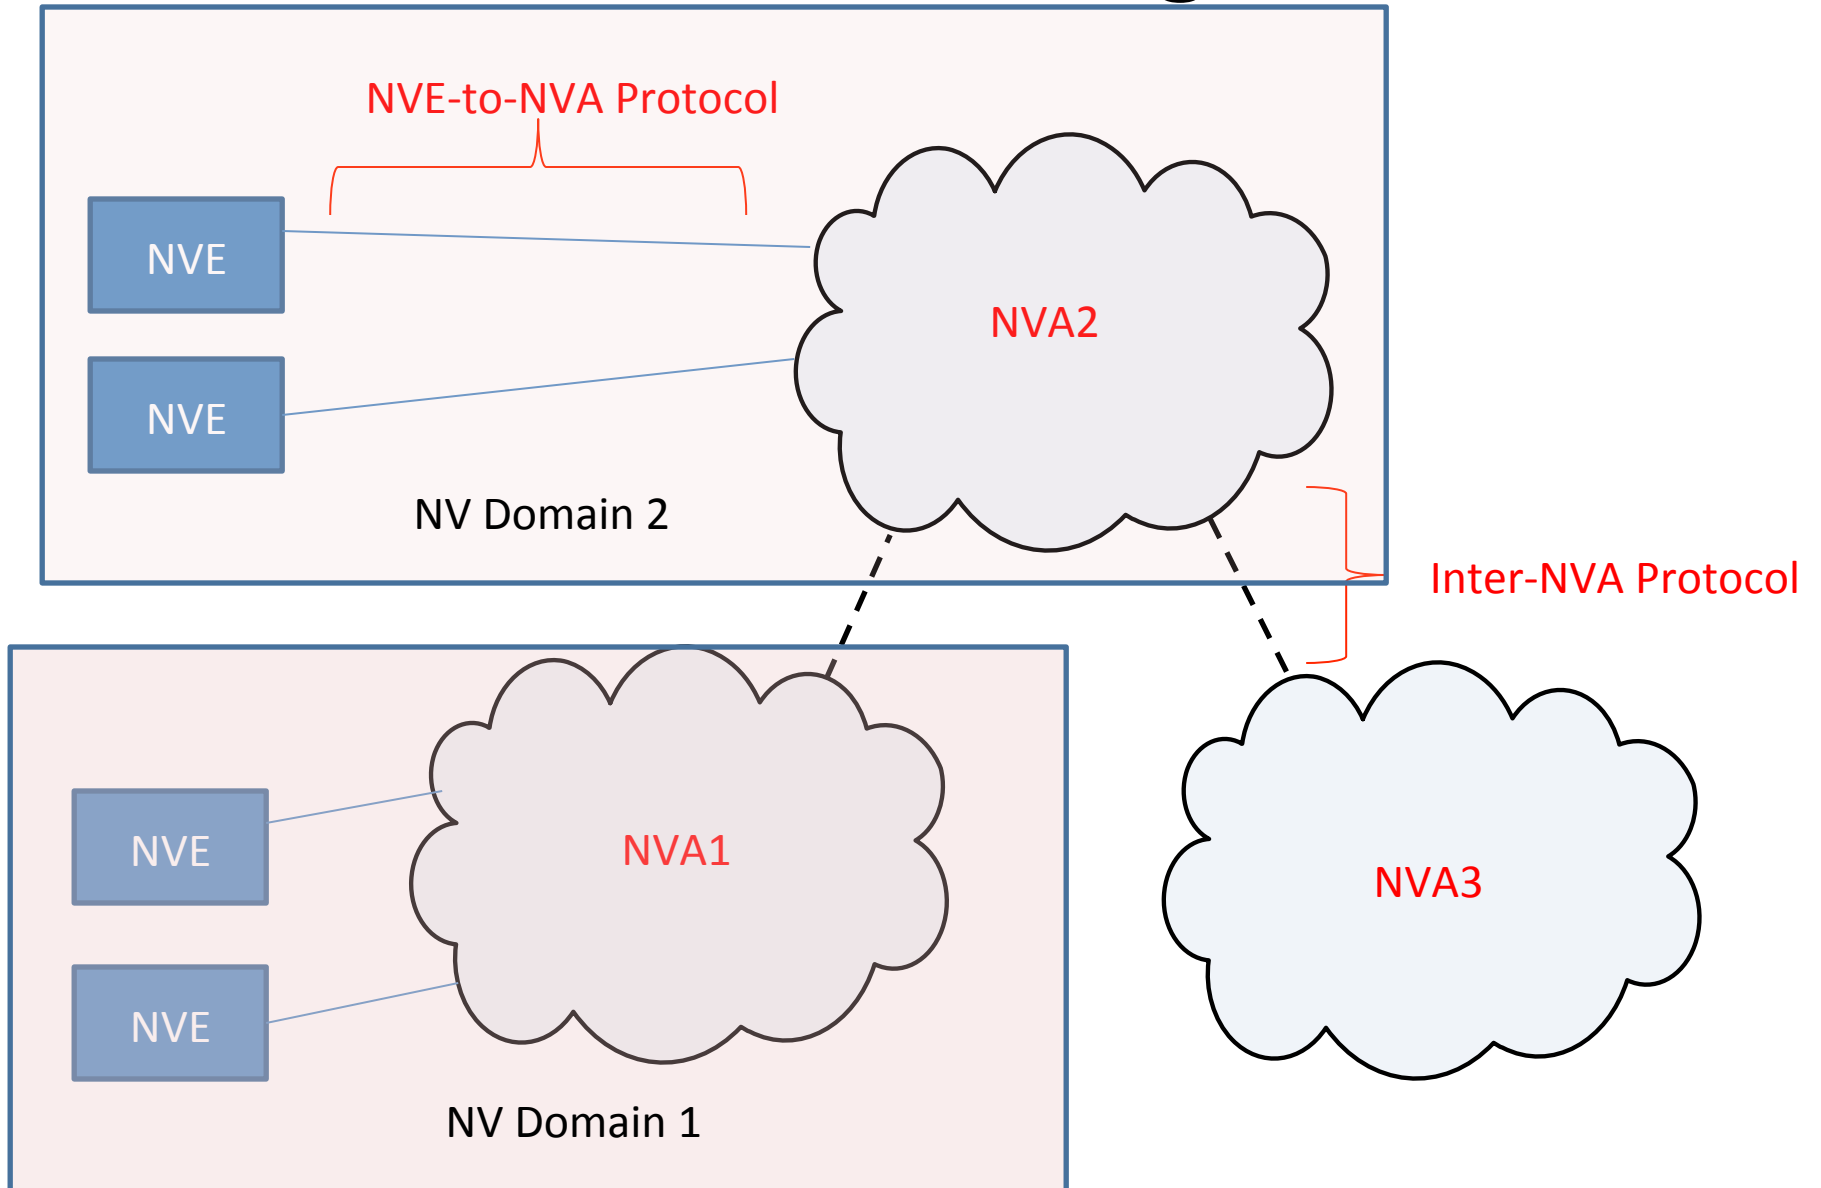
Internal NVA Organization

- Reliability requirement implies:
 - Distributed implementation (e.g, IGP/BGP like), or
 - Use of clustering technology
- NVA-internal architecture/implementation is important, but does not necessarily require IETF standardization
 - BGP extensions (if needed) would be IETF activity
 - Development of database clustering approach (likely) not appropriate for IETF standardization

NVA Federations

- NV Domain – Administrative construct: NVA, NVEs, virtual networks
- NV Region: Two or more NV Domains that share information about virtual networks, to allow VN to span multiple NV domains
- NVAs will need to share information with each other
 - On a per virtual network basis
 - Under policy/configuration control
- Federation of NVAs implies
 - Well-defined/clean interface between NVAs
 - On-the-wire protocol between NVAs
 - Assertion: potential are for IETF work

NV Domains & Regions



Push vs. Pull

- We've had much discussion (at abstract level) about whether “push” or “pull” is better
- “all push” and “all pull” are two ends of a spectrum
 - Neither is what we are likely to see in practice
- Architecture should recognize that both will need to be supported
 - Specific NVA solutions will define where on spectrum a particular NVA approach will lie
 - NVE should support range of models

Questions?