

Next Steps with draft-ietf-teas-ietf-network-slices

Adrian Farrel (Editor)
adrian@olddog.co.uk

Changes in -05

- Issue #1 Connectivity Matrices
 - Clarified definitions for P2P, P2MP, MP2P, MP2MP
 - Added a new matrix Any-to-Any (A2A)
 - Aims
 - Make the general case model tunnel connectivity
 - Support a VPN type of service
 - A2A is not based on tunnel connectivity, it is a routable full mesh
 - Discussion on next slide
 - Can we now handle every type of service you want to offer?
 - Does A2A make MP2MP and MP2P unnecessary?
 - Kiran's new question (SLOs in P2MP service)
- Issue #3 SLOs/SLEs per Sender/Receiver
 - Closely tied to Issue #1 (discussion on next slide)

The Types of Connectivity Matrix

- **P2P** One sending CE and one receiving CE
 - All traffic injected at the sending CE is intended to be received by the receiving CE. Like a private wire or a tunnel
 - The SLOs and SLEs apply at the sender (and implicitly at the receiver)
- **Bidirectional P2P** Two CEs: each may send to the other
 - Two sets of SLOs and SLEs: each applies to one of the CEs as a sender
- **P2MP** One sending CE and more than one receiving CE
 - All traffic from the sending CE is intended to be received by all the receiving CEs. This is like a P2MP tunnel or multi-access VLAN segment.
 - There is one set of SLOs and SLEs that apply at the sending CE (and implicitly at all receiving CEs)
- **MP2P** One receiving CE and (N - 1) sending CEs
 - All traffic injected at any sending CE is received by the single receiving CE. This is like a set of P2P connections all with a common receiver.
 - Each sending CE has its own set of SLOs and SLEs (the combination of those SLOs and SLEs gives the implicit SLOs and SLEs for the receiving CE)
- **MP2MP** Each of N CEs can be a sending CE: traffic is delivered to all of the other CEs
 - Each sending CE has its own set of SLOs and SLEs (the combination of those SLOs/SLEs gives the implicit SLOs/SLEs for each/all of the receiving CEs since each receiving CE is expect to receive all traffic from all/any sender.
- **A2A** Any sending CE may send to any one receiving CE or any set of receiving CEs
 - There is an implicit level of routing in this connectivity matrix that is not present in the other connectivity matrices
 - The matrix must determine to which receiving CEs to deliver each packet
 - The SLOs/SLEs apply to individual sending CEs and individual receiving CEs
 - There is no implicit linkage and a sending CE may be "disappointed" if the receiver is over-subscribed.

Changes in -05

- Issue #2 Connectivity Matrices per Slice
 - We wanted to allow an operator to choose
 - One slice per matrix
 - Multiple matrices per slice
- Issue #4 Service Definition
 - The definition of an *IETF Network Slice Service* was polished in -05
 - Main issues were about endpoints (see #5) and connectivity matrices per slice (see #2)
- Issue #5 Endpoints
 - Figure and description inserted per discussions on the list and at interim
 - Added 3.2.1 “Ancillary CEs” to embody “service functions” – traffic sources/synchs within the provider’s network
- Issue #6 Realization Process
 - Figure and description inserted per discussions on the list and at interim
 - Modifications per emails from Med, Joel, and John D
- Issue #7 Workflow
 - No change made per agreement to not include any further explanation
- Editorials
 - Fix “customer” not “consumer”
 - Fix “IETF Network Slice Service”

Further Issues for Further Resolution

- Proposed for the next revision...
 - Wholesale editorial pass on the document
 - There are some very stale sections unchanged from -00
 - There is a degree of duplication
 - Tidy up remaining language around “end points”, “NSE”, “CE”, Service Demarcation Point
 - Usual spelling and grammar
 - Converging with RFC 8309 (next slide)
 - Clarifying “technology-agnostic” (further slide)
 - Editorials from Med, John D
- Further work
 - Worked examples of how construct a few simple services
 - Any other issues

NBI/SBI

- This terminology gets confusing
 - It is context-specific (your northbound is my southbound)
- Med proposed to align with RFC 8309 and reference L2SM/L3SM
 - Pretty simple idea
 - Same terminology for all “IETF Services”
- Is the SBI based on a Network Model or a Network Configuration Model?
 - Compare with L2NM and L3NM

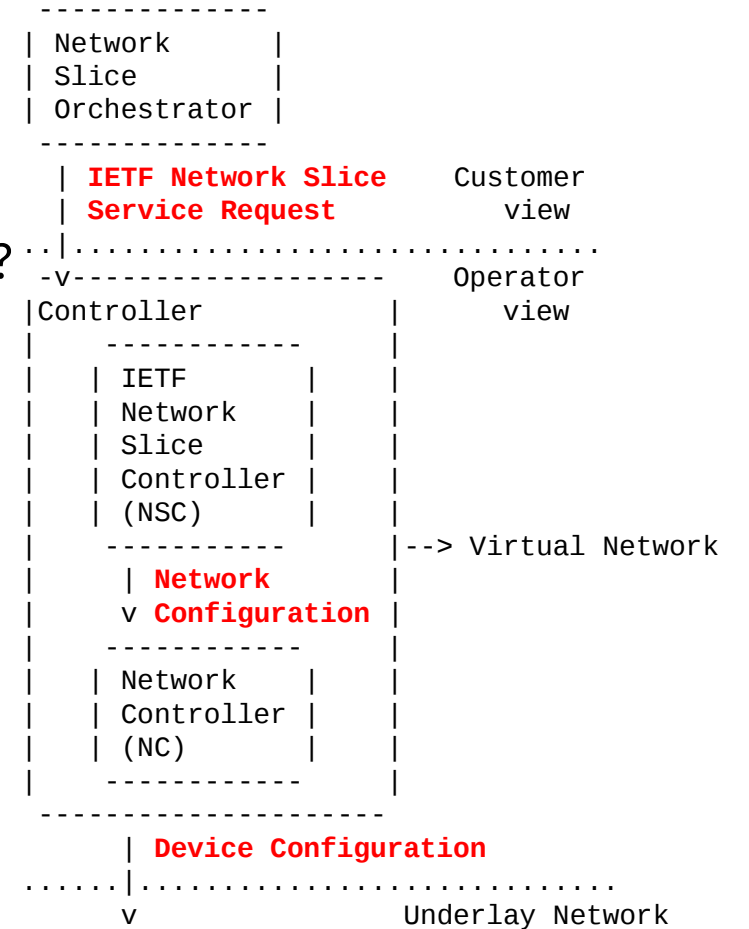
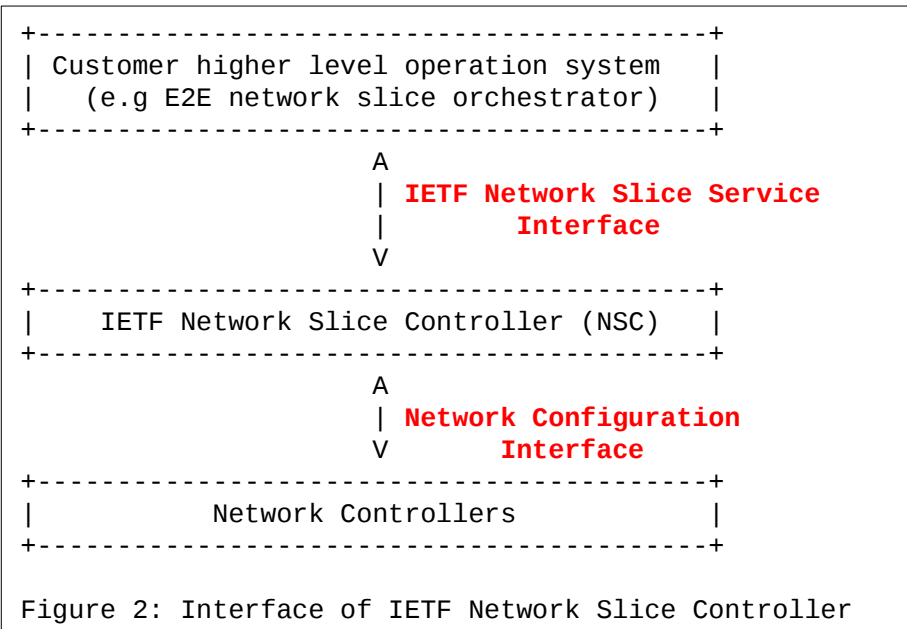


Figure 3: Interface of IETF Network Slice Management Architecture

What is Technology-Agnostic/Specific

- What we know:
 - How the service is provided is not the business of the customer
 - The same service can be provided in multiple different ways
 - The same service can be provided over multiple different technologies
- But:
 - The traffic supplied by the customer is of a specific technology
 - Maybe packet type or encoding of an input stream (such as Attachment Circuit)
 - Thus, the service has a technology-specific aspect
 - The Access Circuits are of a specific technology
 - Thus, if the AC is part of the service, there is a technology-specific aspect
 - The AC may itself be sliced
- Simply clarify this by stating all of these points?
 - Or remove discussion of technology agnosticism?
- “The service is agnostic to the technology in the underlay network”

Raise your other issues here

- Or send mail to the list

The plan...

- -06 as described
 - I would plan to do this in November
- Pre-last call major review
 - All interested parties give the document a thorough reading
 - Maybe run this through to the end of December
- -07 to catch all issues raised
 - First couple of weeks in January
- Possibly ready for WG last call by end of January