

Considerations on Network Virtualization and Slicing

draft-arkko-arch-virtualization

*Jari Arkko
Jeff Tantsura
Joel Halpern
Balázs Varga*

Reasons for Thinking about This Topic

- Growing importance of various virtualized solutions
- Work on data models
- 5G “slice” concept and associated IETF BOFs and proposals
- Frustrations on trying to talk about a topic in a top-down fashion, without enough linkage to what already exists

Some Observations to Start With

- What terminology we use shouldn't (necessarily) affect what new technology we need
 - 5G slicing, for instance, should be able to rely on many existing (or under development) VPN, traffic engineering, data model, etc. works
- The big role of software — many (most) virtualisation can be achieved without touching the protocols
- Protocols vs. data models vs. concepts

Virtualization Tech @ IETF

- Virtualization does not generally affect TCP/IP or applications
- Some exceptions to this, when assumptions made somewhere are broken due to virtualization, leading to a need to add information to application protocols:
 - E.g., early HTTP versions assumed that 1 server = 1 website
 - With virtual hosting, modern HTTP versions carry intended web site name inside the protocol

Virtualization Tech @ IETF

- Lower layers can carry virtualisation-related information (e.g. VLAN tags in Ethernet frames)
 - Some IETF-defined support protocols (e.g., AAA) and tunnelling protocols can also indicate desired domains
- End-point VPN technologies, e.g., IPsec-based VPNs
- Provider-based VPNs and associated technical components — MPLS, L2VPN, L3VN, LDP, BGP, ...
- Traffic engineering, e.g., TEAS WG works on enhancements to traffic-engineering capabilities for MPLS and GMPLS networks

Virtualization Tech @ IETF

- Tunnels — e.g., GRE, LISP
- Service chaining — SFC WG's NSH
- Management frameworks — e.g., NETCONF, YANG
- Data models — e.g., L2SM, L3SM

Architectural Observations

- Trend: Increasing role of software
- Trend: Centralization of functions, which makes things easy. But:
 - Keeping operations up when networks partition or center is down
 - Avoiding the loss of control in the network edges
- Tailored vs. general-purpose networking; what are the economics of special-purpose treatment and QoS?
- Think about data model layering! E.g., service vs. network/device data models.
- General over specific — does it make sense for IETF to do general designs or designs for someone's specific requirements at specific time?

Possible New Work

- Data models
- Managing heterogeneous technologies, e.g., across SDN and traditionally built networks, or manage both general-purpose and very technology-specific virtualization parameters (e.g., VPNs + 5G radio)
- New underlying VPN or API capabilities — e.g., ability to specify "statistical" rather than hard performance parameters
- Cross-domain management and VPNs — but these are very hard
- Some terminology and conceptual alignment across industry would also be useful, e.g., to know what words to use in SLAs...

Next Steps

- Connect the top-down and bottom-up approaches to thinking about virtualization and slicing
- Determine what concrete things should be done
- Divide work to existing or new WGs...
- Data model development for virtual networks seems as it should be a continued topic for IETF work
- Not only 5G slicing -driven issue, also core to what IETF outputs are. And we need to talk about how our own various pieces fit together (service models... traffic engineering... detnetworking...)
- ...