



Building data centres on open standards

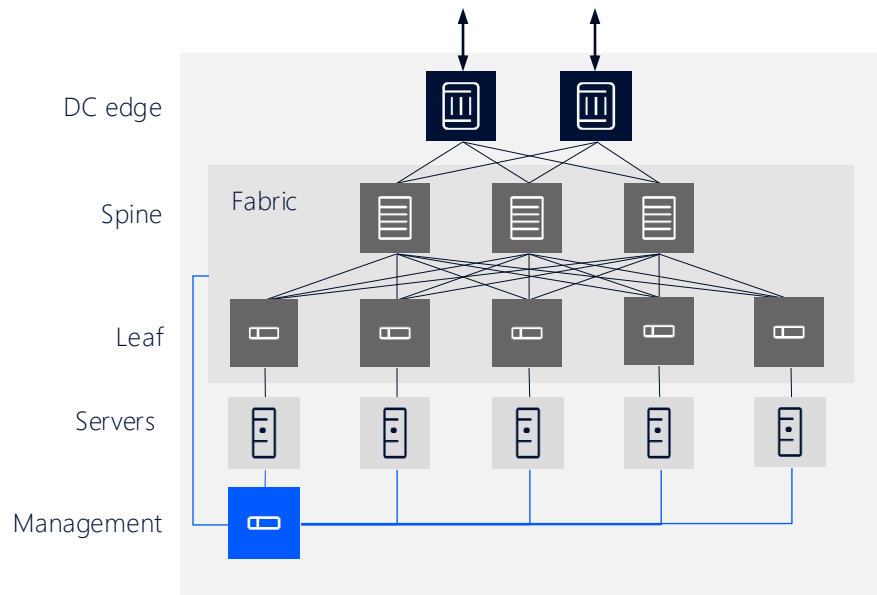
Bastien Claeys

Regional Product Manager at Nokia

NOKIA

Data center network architectures

The industry has converged



Non-blocking fabrics



- IP and EVPN fabrics
- DC gateway or border leaf derivatives
- Collapsed core for edge DC
- Scale via super spines / pods

ASICs tailored per use case



- Range of different ASICs on the market
- Key properties: latency, programmability, port speed & density, feature set


OOB management



- Merchant silicon
- 1G/10G port speeds

The myth of the ultimate switch

Why one vendor never fits all (forever) in data center networking



- Deciding factors when selecting a platform
 - Port density, feature set, ease of operation/automation, cost
- Using equipment from a single vendor seems practical
 - Train for a single network operating system, easy to support and maintain
- Requirements evolve over time, e.g. data centre renewal
 - Sticking with the same vendor can be costly ... and change require time and effort

Breaking free

Escaping vendor lock-in with multivendor networks


- Being locked into one vendor reduces flexibility but not only
 - Operators are more sensitive to supply chain delays and high cost
 - Massive blast radius in case of bug in the network stack implementation
- There's a need for a multivendor approach - how do we manage this efficiently ?



One network, many vendors

Standardisation as the key to multivendor networks

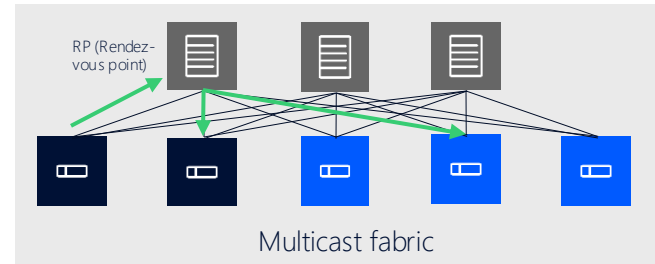
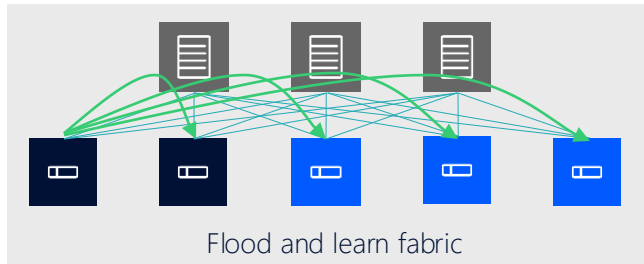
- Proprietary solutions can be effective, as a vendor has then full control on its solution design. However, it doesn't have to be this way.
- Today, underlays can be deployed in many ways using OSPF, IS-IS or BGP. Overlay networks are typically deployed with BGP-EVPN. There are mature solutions for designing data centres that simply works.
- Opting for standard practices is a way for operators to remain flexible and agile when it comes to using equipment from different vendors
- How do we get there ?



Moving away from proprietary solutions

Running VxLAN in Data Centres


- Early 2010's : first VxLAN-based fabrics
- Flood & learn approach
- Remote VTEPs statically assigned
- Traffic flooded to all VTEPs without service knowledge
- Vendor-specific solution and not scalable ☹️
- 1-2 years later : VxLAN-based fabrics running IP multicast
- 1 multicast group per service
- VTEPs subscribe to relevant groups
- Traffic is only sent to other members
- Better traffic replication than flood & learn
- Vendor-specific solution and multicast is needed ☹️
- Then, emergence of HW VTEP with SDN controllers
- Decoupled control plane and data plane
- Scalable but proprietary solution ☹️



Ethernet VPN : the standard for modern networks

Ensuring seamless connectivity across vendors

- VPN services were traditionally delivered using different technologies depending on the service type: BGP/LDP for VPLS and VPWS, MP-BGP/MPLS for IP VPNs, and BGP/PIM for multicast VPNs.
- Ethernet VPN (EVPN) : standardized BGP-based control plane solution that enables scalable, efficient, and flexible Layer 2 and Layer 3 network virtualization across data centers and service provider networks
- EVPN is today a key component of a standardised architecture and an IETF workgroup is taking care of future improvements.



February 2015 : RFC 7432 BGP MPLS-Based Ethernet VPN Ethernet VPN

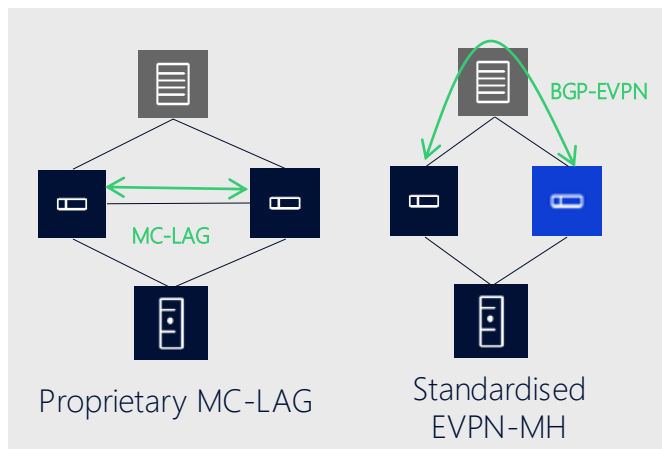
...

August 2024 : RFC 9625 EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding

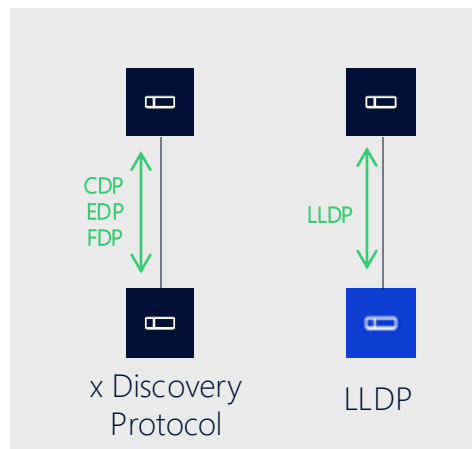
Moving away from proprietary solutions (2)

EVPN and beyond

- Multihoming has for long been provided through MC-LAG
- EVPN Multihoming is a standardised solution that works between platforms from different vendors



- In 1994, Cisco Discovery Protocol was introduced
- LLDP is now the industry approach, compatible with practically all NOS and OS




- Multi-Source Agreements (MSA) make sure form factors are designed following the same guidelines
- While not being official, those agreements ensure interoperability between equipment
- Examples
 - Optical transceiver form factors (SFP, QSFP, ...)
 - Optical transmission (CLR4, CWDM4, ...)


EVPN-VxLAN fabrics

Best practices and repeatable blueprints

- With EVPN control plane properly defined by RFCs, one can easily deploy this protocol on different silicon as long as the routing protocol stack follows standards
- For instance, multiple ways to deploy a fabric underlay (IS-IS, OSPF, BGP), but sticking with a standard implementation allows to bring diversity
- Tuning of specific parameters to reach the best configuration possible, such as BGP timers to reach faster convergence for instance
- It's a fact - vendors define blueprints that are easily reproducible



Example of fabric underlay design



Example of a fabric overlay design

EVPN in the real world

Validation across vendors


- Standardisation is important and vendors know it : some of them publish validated designs that operators can execute
- If following standards is important, so is validation. Interoperability tests make sure equipment from different vendors are compatible, e.g. EANTC.
- Interop tests driven by vendors and involving operators to define test scenarios or evaluate results
- Use cases as close as possible to the field to confirm compatibility and inspire confidence in each vendor's network protocol implementation.



Programmable interfaces

Why model-driven matter so much

- Model-driven interfaces relying on abstraction via data models
- Transactional approach to ensure predictable changes
- Benefits
 - Validation, rollback
 - Favours automation
 - Easily integrated as part of a feedback loop
- Effective on CLI, even more powerful through management interfaces
 - NETCONF or gNMI



Programmable interfaces

gRPC framework for network equipment

gRPC
HTTP/2 - Protobuf

gNMI - gRPC Network
Management Interface

Management protocol

Configuration
modification, state
retrieval etc.

gNSI- gRPC Network
Security Interface

Security infrastructure
services

AAA, certificates and
key management

gNOI - gRPC Network
Operations Interface

Operational commands

Reboot, upgrade, file
management, ping etc.

gRIBI - gRPC Routing
Information Base
Interface

Route injection

Programmatic route
installation into a
router's RIB

MD + standardised interfaces




Automation
CI/CD integration

Multivendor fabric management

Bringing different vendors together

- Multivendor environments are the most cost effective and robust way to build data centres
- Fabric management is key to make this possible
- Emergence of intent-based controllers emphasise this aspect. They manage switches using standardised interfaces and network protocols.



Key takeaways

Vendor lock-in exposes customers to supply chain delays and higher costs

Open standards enable robust multi-vendor interoperability

Model-driven interfaces for config and state are critical for modern DC networks

Automation and observability are day-0 requirements

NOKIA