

Trellis Introduction

Saurav Das, Charles Chan & Jono Hart

with contributions from many more ...

CORD Build 2017, San Jose January 14, 2018



Outline



• What is Trellis?

• Why Trellis?

• Trellis Roadmap

Trellis

Multi-purpose leaf-spine fabric designed for NFV



Trellis – Multi-purpose Leaf-Spine Fabric

____ Open Network Operating System



Disaggregation – Bare-metal + Open-Source



Leaf/Spine Switch Software Stack



OCP: Open Compute Project ONL: Open Network Linux ONIE: Open Network Install Environment BRCM: Broadcom Merchant Silicon ASICs OF-DPA: OpenFlow Datapath Abstraction

Fabric ASIC Pipeline^{*} (BRCM's OF-DPA)



Trellis Control Plane Redundancy



Spreading mastership over controller instances contributes to scale

Trellis Control Plane Redundancy



Spreading mastership over controller instances contributes to scale

Trellis Data Plane Redundancy



Trellis – CORD Network Infrastructure



Trellis Provides Common control over underlay & overlay networks, including

- Service Composition for Tenant Networks
- Distributed Virtual Routing
- Optimized Delivery of Multicast Traffic Streams

Trellis is the enabling Network Infrastructure for CORD

CORD Architecture





Outline



• What is Trellis?

• Why Trellis?

• Trellis Roadmap

Why Trellis?



• Trellis is designed for Service Providers & NFV

Trellis in CORD



Why Trellis?



• Trellis is designed for Service Providers & NFV

• SDN allows simpler/easier/optimized features

vRouter in Software?



<u>Issues</u>: Hairpinning Embedded control plane complicates scale-out

vRouter Operation



vRouter Operation



vRouter Operation

Trellis vRouter → Implemented as a big distributed router

→ Presents entire n/w infrastructure as a single router to outside world



Why Trellis?



Trellis is designed for Service Providers & NFV

- SDN allows simpler/easier/optimized features
- SDN + Programmable pipelines == New features

SDN + Programmable Pipelines



Why Trellis?



Trellis is designed for Service Providers & NFV

- SDN allows simpler/easier/optimized features
- SDN + Programmable pipelines == New features
- Open-source == SP ownership & customizability

Outline



• What is Trellis?

- Why Trellis?
- Trellis Roadmap

Trellis Roadmap vs. Trellis-CORD Roadmap 從

Trellis Features Used in CORD

- Bridging & VLANs (Access/Trunk/Native)
- IPv4 Unicast & Multicast Routing
- vRouter BGPv4 & Static Routing

Trellis Features To-Be-Used in CORD

- Dual-Homing
- IPv6
- DHCP L3 relay
- Pseudowires

Trellis Roadmap



Upcoming ONOS 1.12 release (Dec 2017)

- Dual-homing (released 1.11.1) fixes for known issues
- Pseudowire support
- Initial QoS support
- IPv6 additional features
 - DHCPv6 relay additional features (contributed by Nokia)
 - IPv6 Multicast (contributed by Nokia)
 - IPv6 Router Advertisement app (contributed by Infosys)
- Support for New ASICs & Bare-metal switches
 - Broadcom Qumran (QMX)
 - Cavium Xpliant
 - Quanta switches (QCT LY8)
 - Barefoot Tofino using P4 (not in 1.12 release)

Roadmap



See here for complete roadmap

https://docs.google.com/spreadsheets/d/1SmKxsZ9iHtLHoEtvsntsJ98tLf_cYT_qXuuO7eAUtt4/edit?usp=sharing

⊞	Fabric	Features & Status 🛛 🛱 🖿				sauravdas@alumni.	stanford.edu 👻
	File Ed	it View Insert Format Data Tools Add-ons Help Last ed	dit was on May 6				Share Share
£		T \$ % .0 .00 123 - Arial - 12 - B		5 - <u>†</u> - ⊐ - 17 - 00 □ □ □ 7	ν - Σ -		*
Jx	# A	В	С	D	E	F	G
1	#	Feature	Use Case	Status	Notes	OF-DPA Dependency (Indicates version onwards)	
2	F1	L2 switching within a rack handled at leaf-switches	Platform	available in CORD 1.0 (Aug' 16)	Note: recently reported (Dec'16)by Johan BRCM CSP# 1117220 (unreferenced flood group causes implicit forwarding) - patch available for 2.0 GA	ofdpa-I.12.1.1 (Dec'15)	
3	F2	L3 forwarding across racks using ECMP hashing and MPLS segment routing.	Platform	available in CORD 1.0 (Aug' 16)		ofdpa-I.12.1.1 (Dec'15)	
4	F3	vRouter integration for using the fabric as the dataplane for vRouter, to interface with upstream router and provide reachability to publicly routable IP addresses.	Platform	available in CORD 1.0 (Aug' 16)		ofdpa-I.12.1.1 (Dec'15)	
5	F4	VLAN cross-connect feature to bridge QinQ packets between OLT I/O blades and vSG containers	R-CORD / AT&T	available in CORD 1.0 (Aug' 16)	Caveat: can only bridge based on outer VLAN, and can only bridge OLT to server in the same rack - requires PW support to overcome this limitation - see F15)	ofdpa-I.12.1.1 (Dec'15)	
6	F5	IPv4 multicast forwarding and pruning for IPTV streams (with vRouter integration) from upstream router to residential subscribers	R-CORD / AT&T	available in CORD 1.0 (Aug' 16)		ofdpa-I.12.1.1 (Dec'15)	
7	F6	Support for OF-DPA 3.0 EA0	Platform	available in ONOS 1.8 & CORD 2nd release (Dec' 16)	CORD-311, Workaround made for BRCM CSP#1100576 (cannot add flow to point to empty group)	OF-DPA 3.0 EA0 (Jul'16)	
8	F7	DHCP relay app support with fabric (includes ability to directly connect dhcp-server to fabric, or place it behind external router)	Platform / Comcast	available in ONOS master (Mar ' 17)	CORD-522	OF-DPA 3.0 EA3 (Nov'16)	
9	F8	Ability to statically configure and route a prefix to a configurable Next-Hop	Comcast	available in ONOS master (Mar ' 17)	CORD-523	OF-DPA 3.0 EA3 (Nov'16)	
10	F9	Support for 802.1x authentication (without blocking support- see F12)	R-CORD / Comcast	available in ONOS master (Mar ' 17)	CORD-527	OF-DPA 3.0 EA3 (Nov'16)	
11	F10	IPv6 support including ICMPv6 NDP, vRouter(BGP v6) and static routes	Platform / Comcast	available in ONOS master (Mar ' 17)	CORD-520	OF-DPA 3.0 EA4 (Jan'17)	
12	F17	Access/Trunk port configuration for bridging - includes the ability to route between vlans (v4 and v6), configure static routes, dhcp for tagged hosts, and native-vlan support on trunk ports	Platform / Comcast / M-CORD	available in ONOS master (Mar ' 17)	CORD-800, CORD-906	OF-DPA 3.0 EA4 (Jan'17)	
13	F20	Multiple Quagga support for vRouter HA preferably with ability to connect quaggas to leaf switches in different racks	Comcast / Platform	under development, scheduled for June 17	CORD-903,904,905		
14	F11	Support for Multi-homing of servers to dual ToRs, with changes in routing and bridging to accomodate hosts reachable via two ToRs	Platform / Comcast	under development, scheduled for June 17	CORD-256		
15	E12	Support to create and update ACLs for blocking traffic	Diatform / Comparet	scheduled for beyond June,	0080 534		

Summary



- What is Trellis?
 - Multi-purpose leaf-spine fabric
 - Built with → bare-metal hardware + open-source software + SDN
 - Overlay + Underlay when used in CORD
- Why Trellis?
 - Designed for SP use cases & NFV
 - SDN allows simpler, easier & more-optimized features
 - SDN + P4 = new features in hardware
 - Open-source gives SP's ownership & full customizability
- Trellis Roadmap
 - Different roadmaps for Trellis, and the use of Trellis in CORD



and more...