Trellis

An Open-Source, White-Box, SDN Based Leaf-Spine Fabric

Charles Chan, Ph.D.

Outline

- What is Trellis?
- Trellis features
- Why Trellis?
- Recent activities (in 2018)
- Deployments & use cases
- Ecosystem & roadmap

What is Trellis?

Production-ready
multi-purpose leaf-spine fabric
designed for NFV

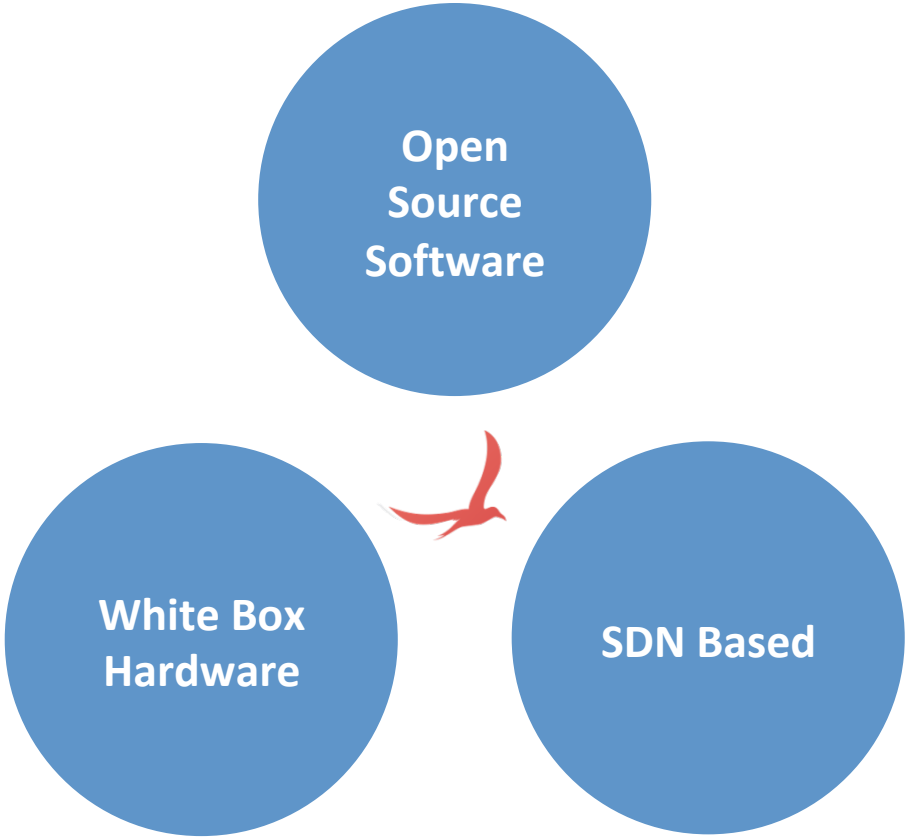

Trellis Full Topology

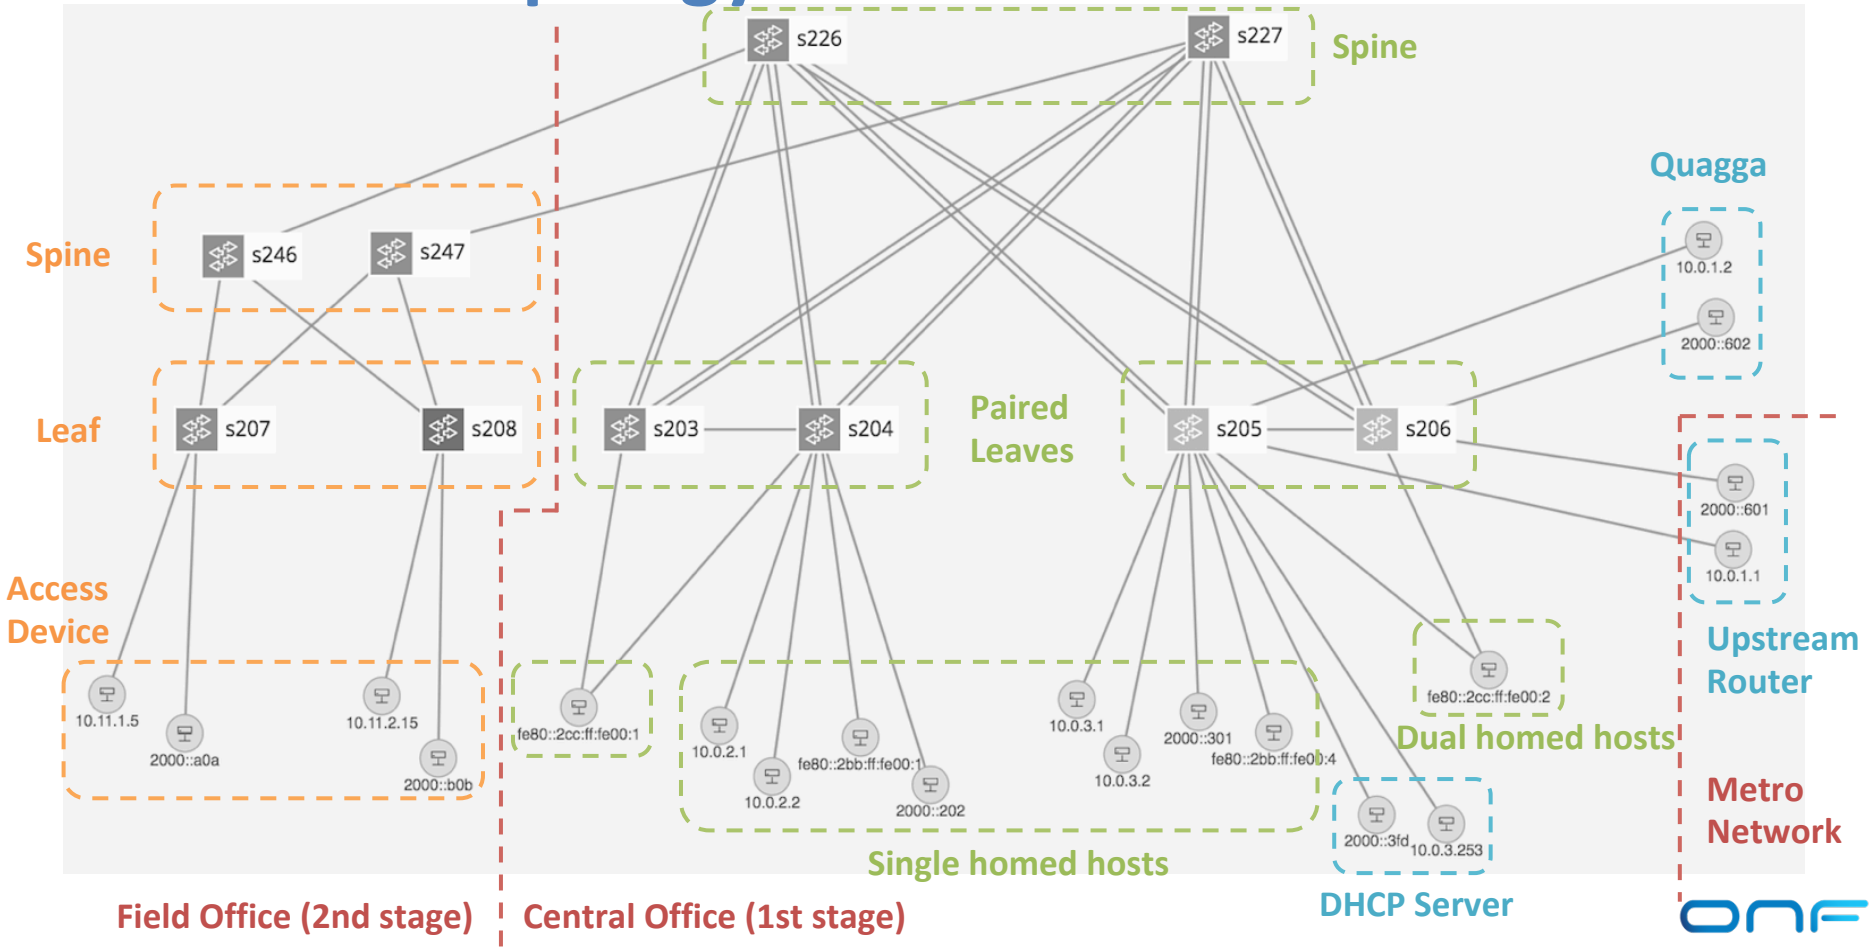

Way Too Complicated?

Trellis supports simpler configurations too!

Single Leaf

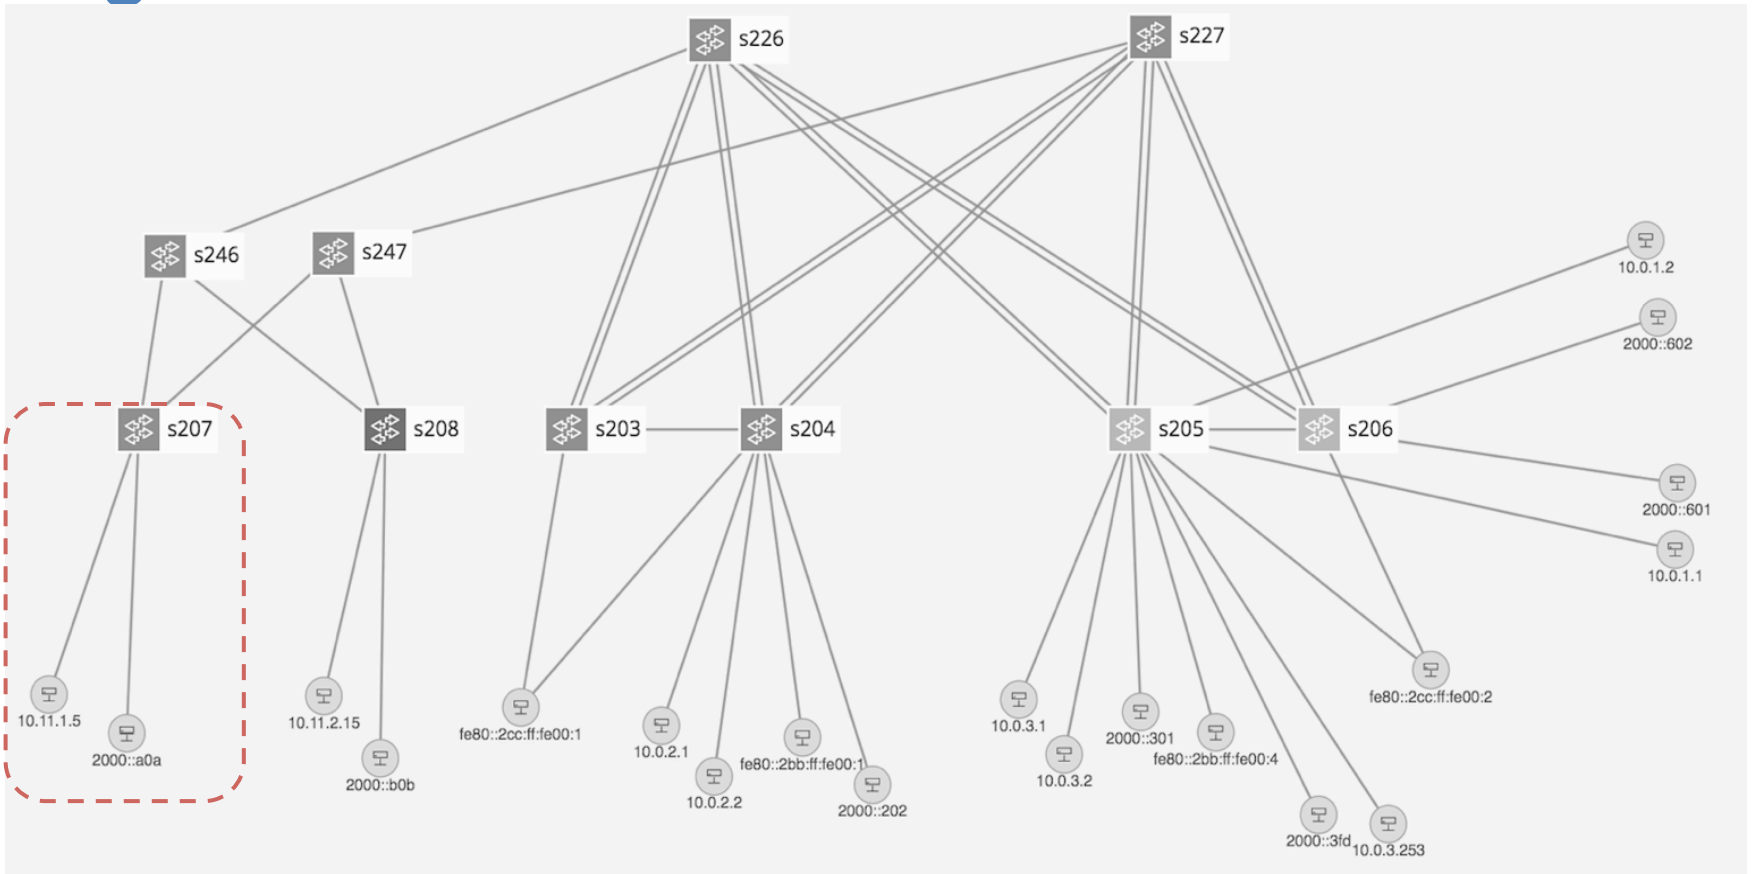
s226
s227
s246
s247
s207
s208
s203
s204
s205
s206
10.0.1.2
2000::602
2000::601
10.0.1.1
fe80::2cc:ff:fe00:2
10.11.1.5
2000::a0a
10.11.2.15
2000::b0b
fe80::2cc:ff:fe00:1
10.0.2.1
10.0.2.2
fe80::2bb:ff:fe00:1
2000::202
10.0.3.1
10.0.3.2
2000::301
fe80::2bb:ff:fe00:4
2000::3fd
10.0.3.253

Single Leaf Pair

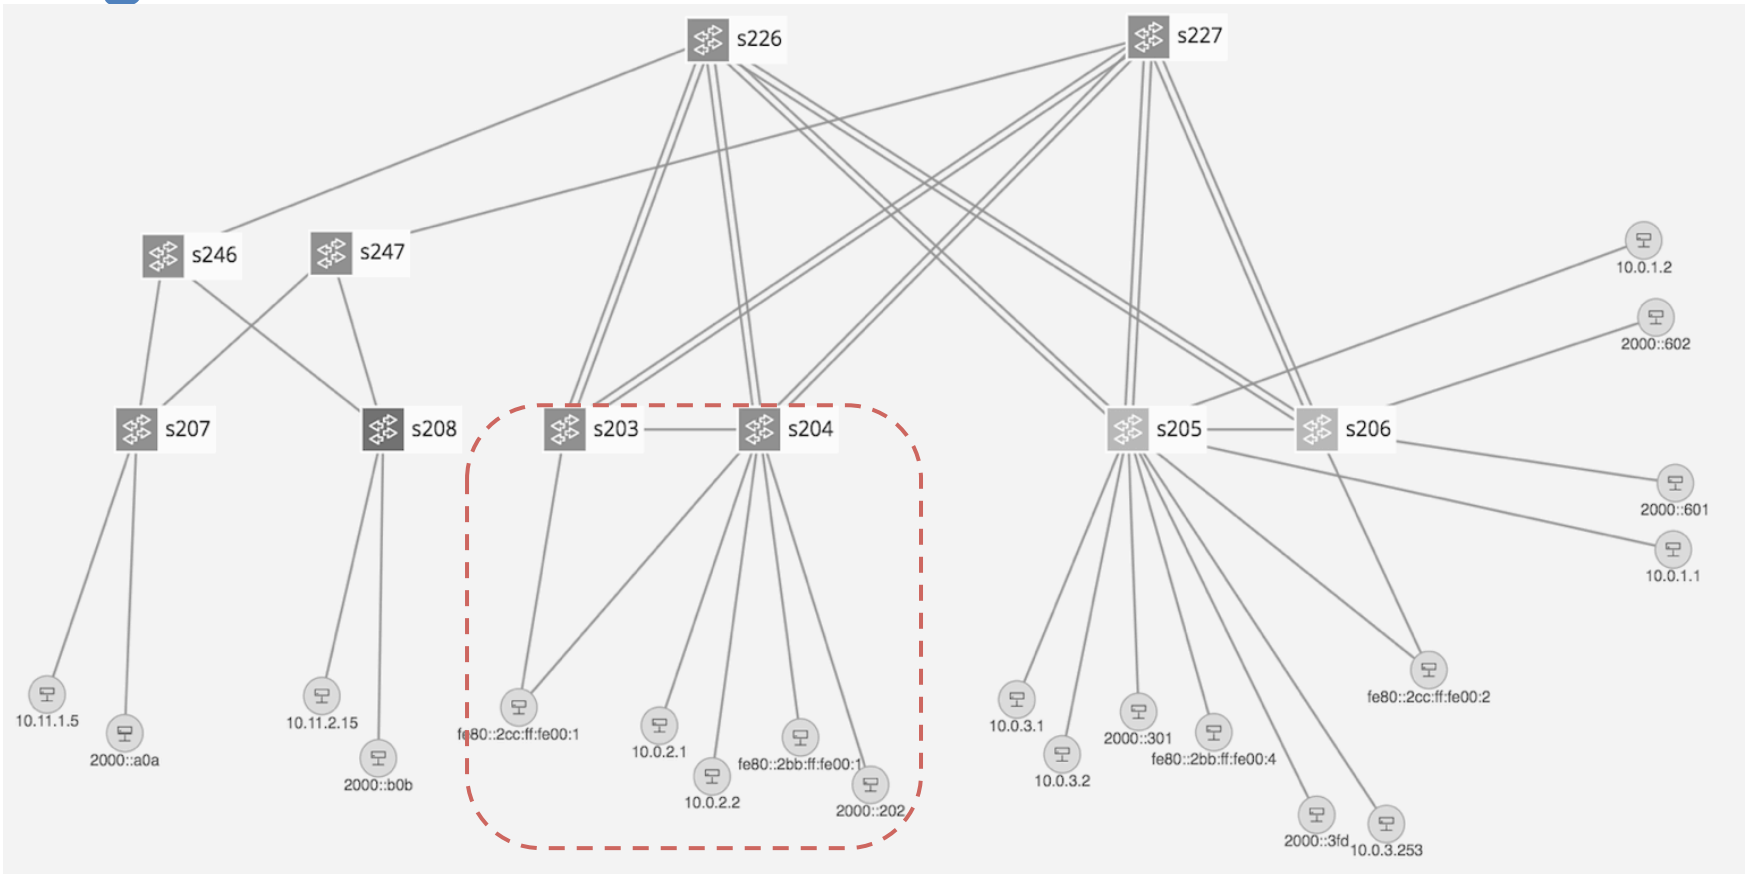

Single Leaf Pair with Upstream

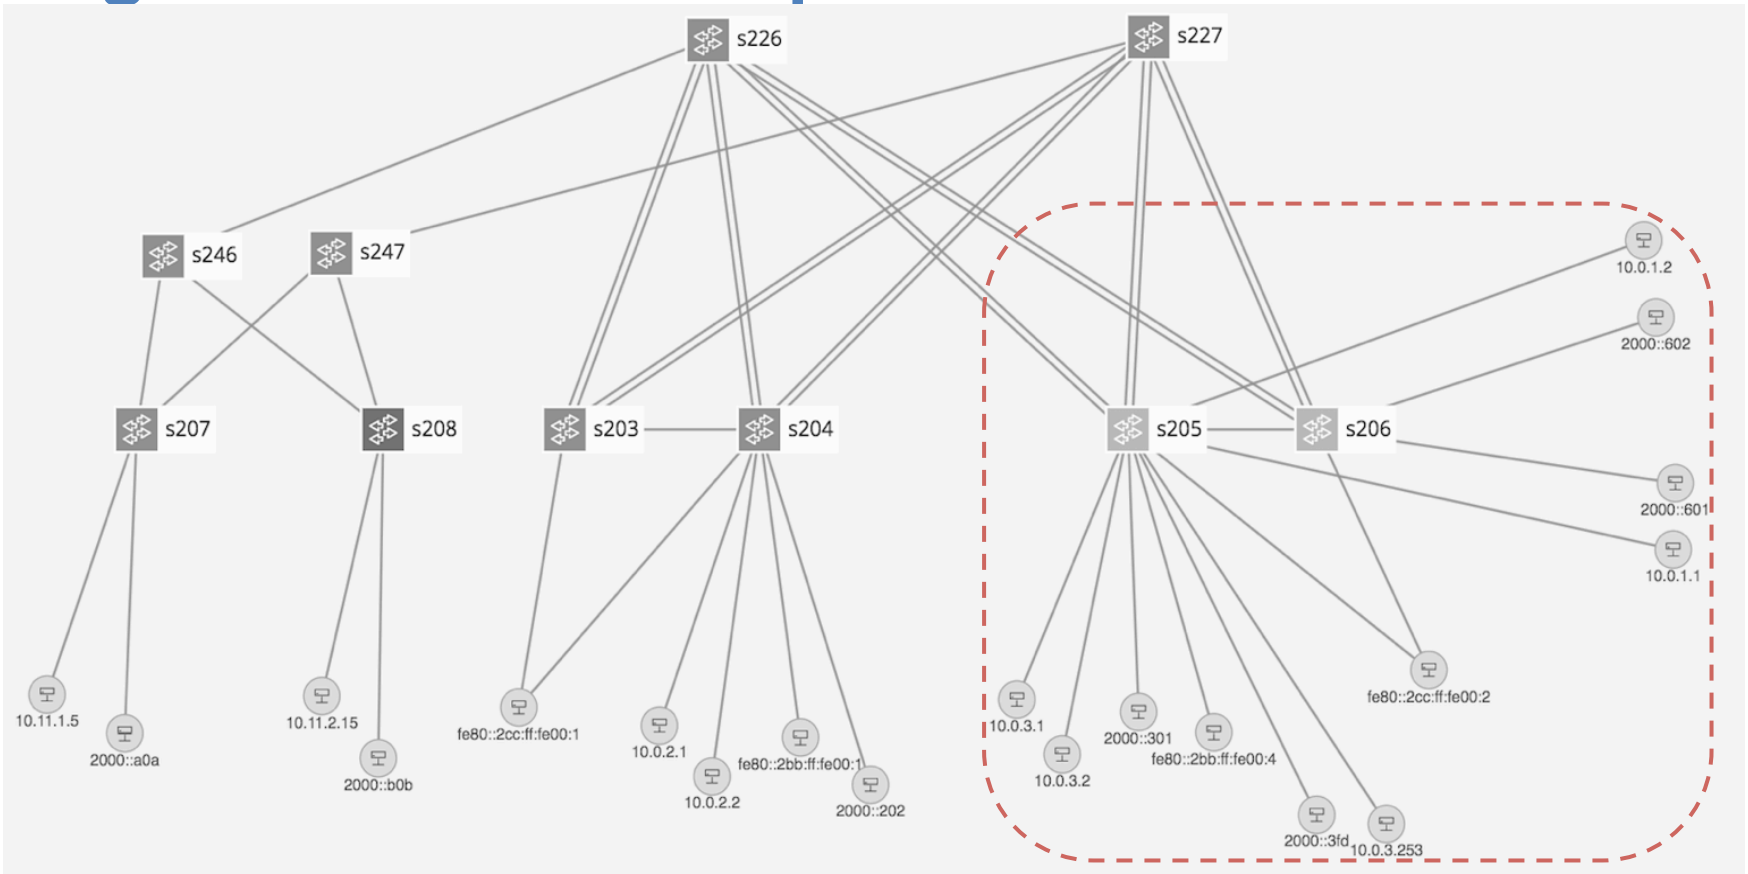

Single Stage

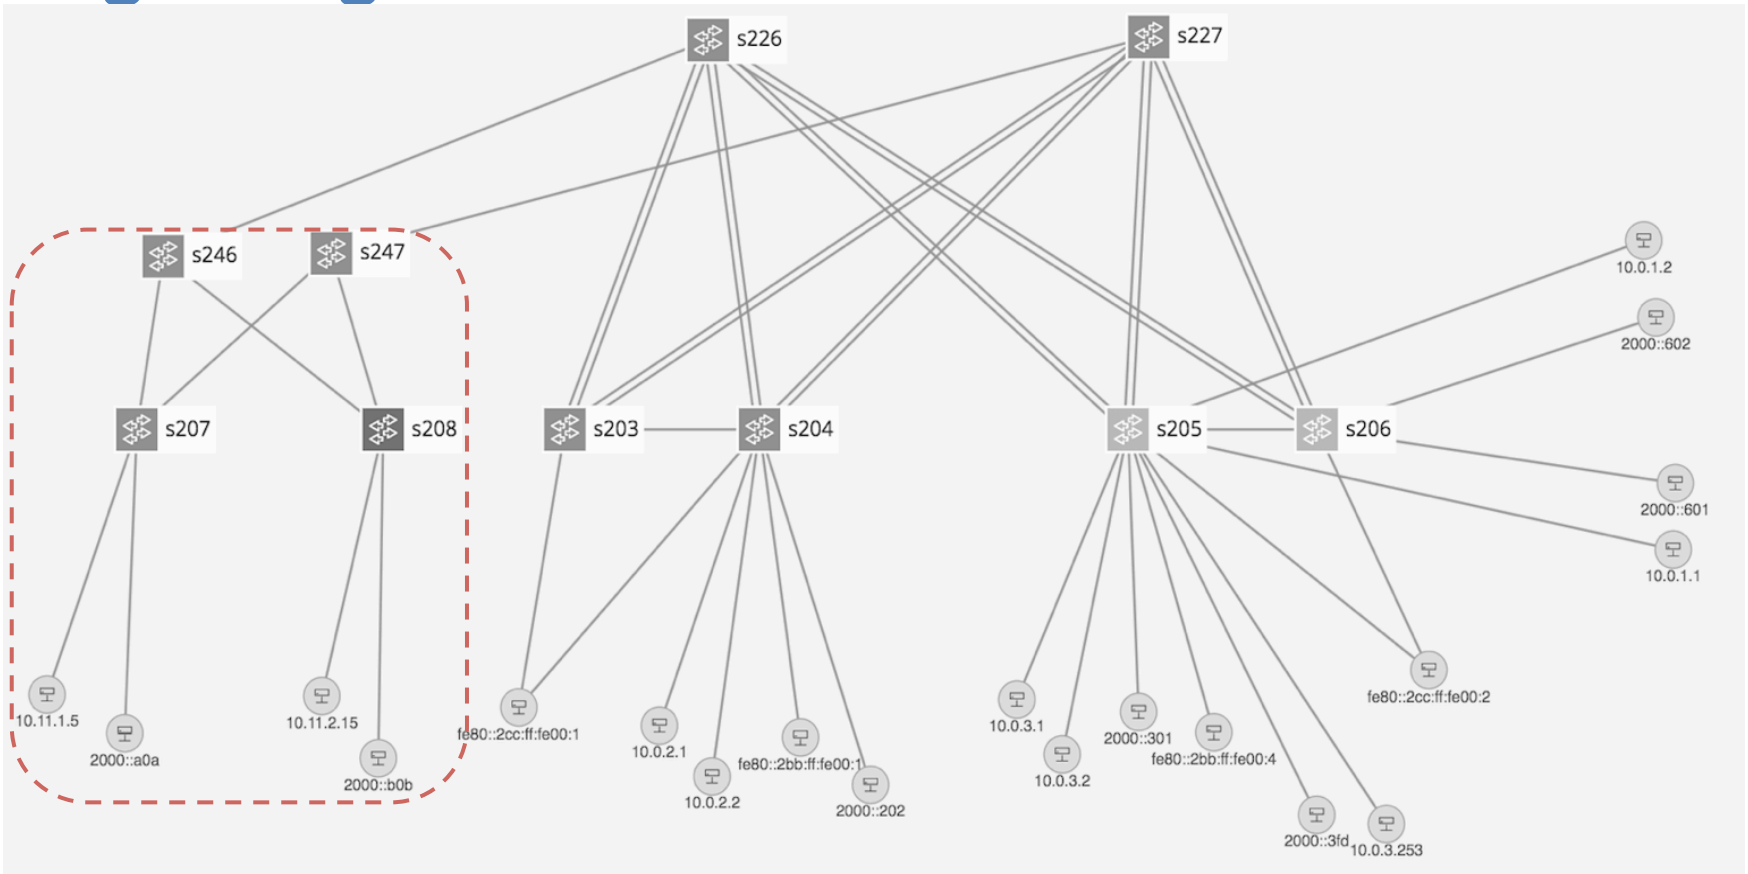

Single Stage with Leaf Pairs & Upstream

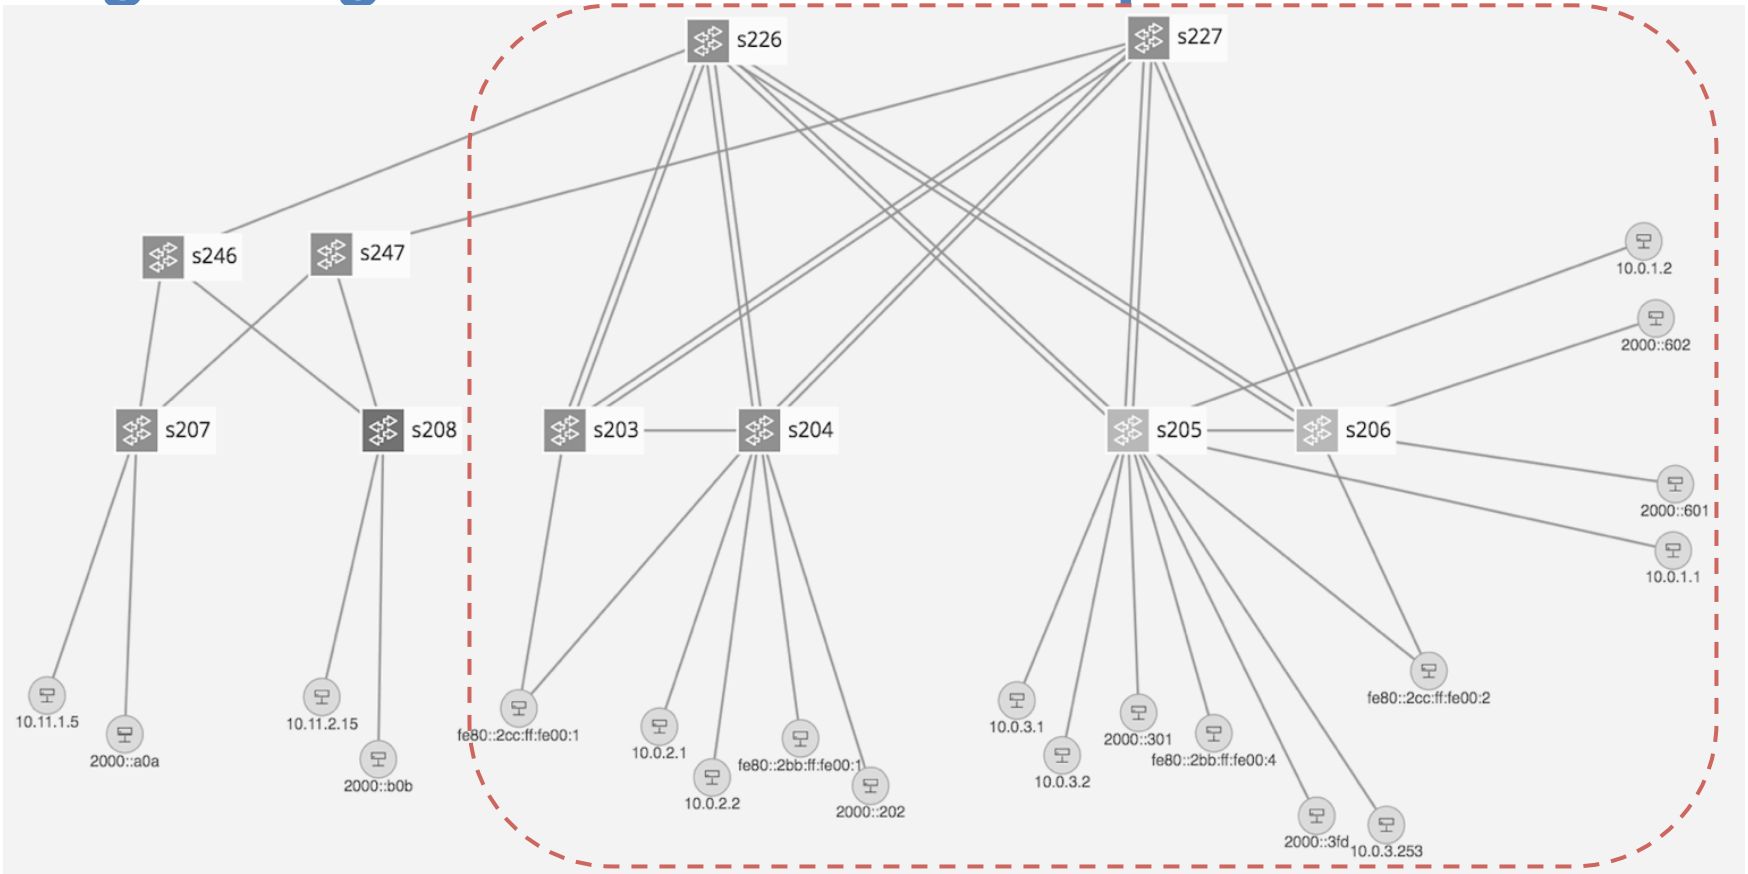

Trellis in Production

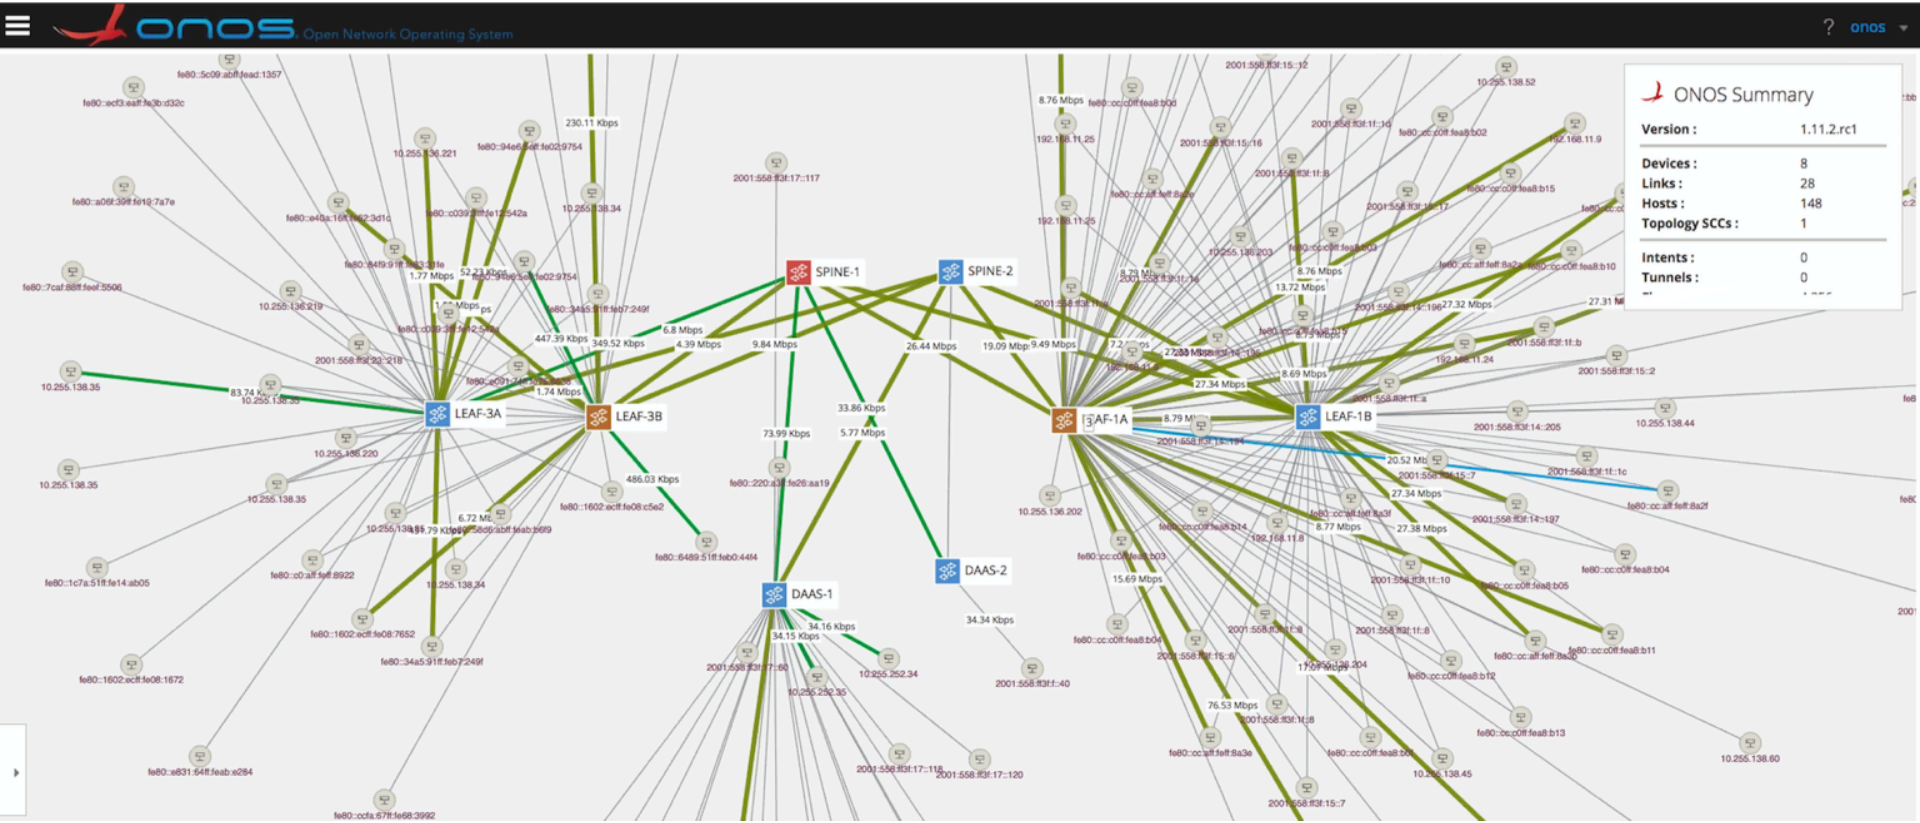

Trellis Features

- **Bridging** with Access & Trunk VLANs (within a rack)
- **Routing** (inter-rack)
  - IPv4 & IPv6 Unicast routing with MPLS Segment-Routing
  - IPv4 & IPv6 Multicast routing
- **Dual-homing** for compute-nodes and external routers
- **Multi-stage** fabrics (2 layers of spines)
- **vRouter** - entire fabric behaves as a single router
  - BGP (v4/v6) support for external connectivity
  - Static routes, route blackholing
- **DHCP L3 relay** (v4/v6)
- **MPLS Pseudowire**
- **Q-in-Q termination**
- **T3** (Trellis Troubleshooting Tool), **onos-diags**

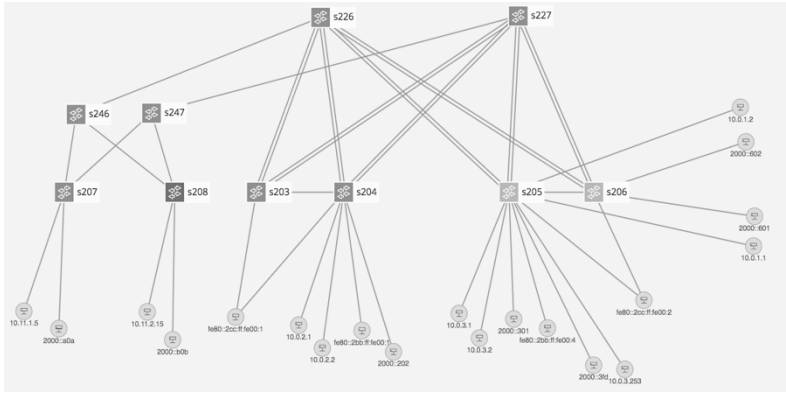

Trellis Features / Bridging & Routing

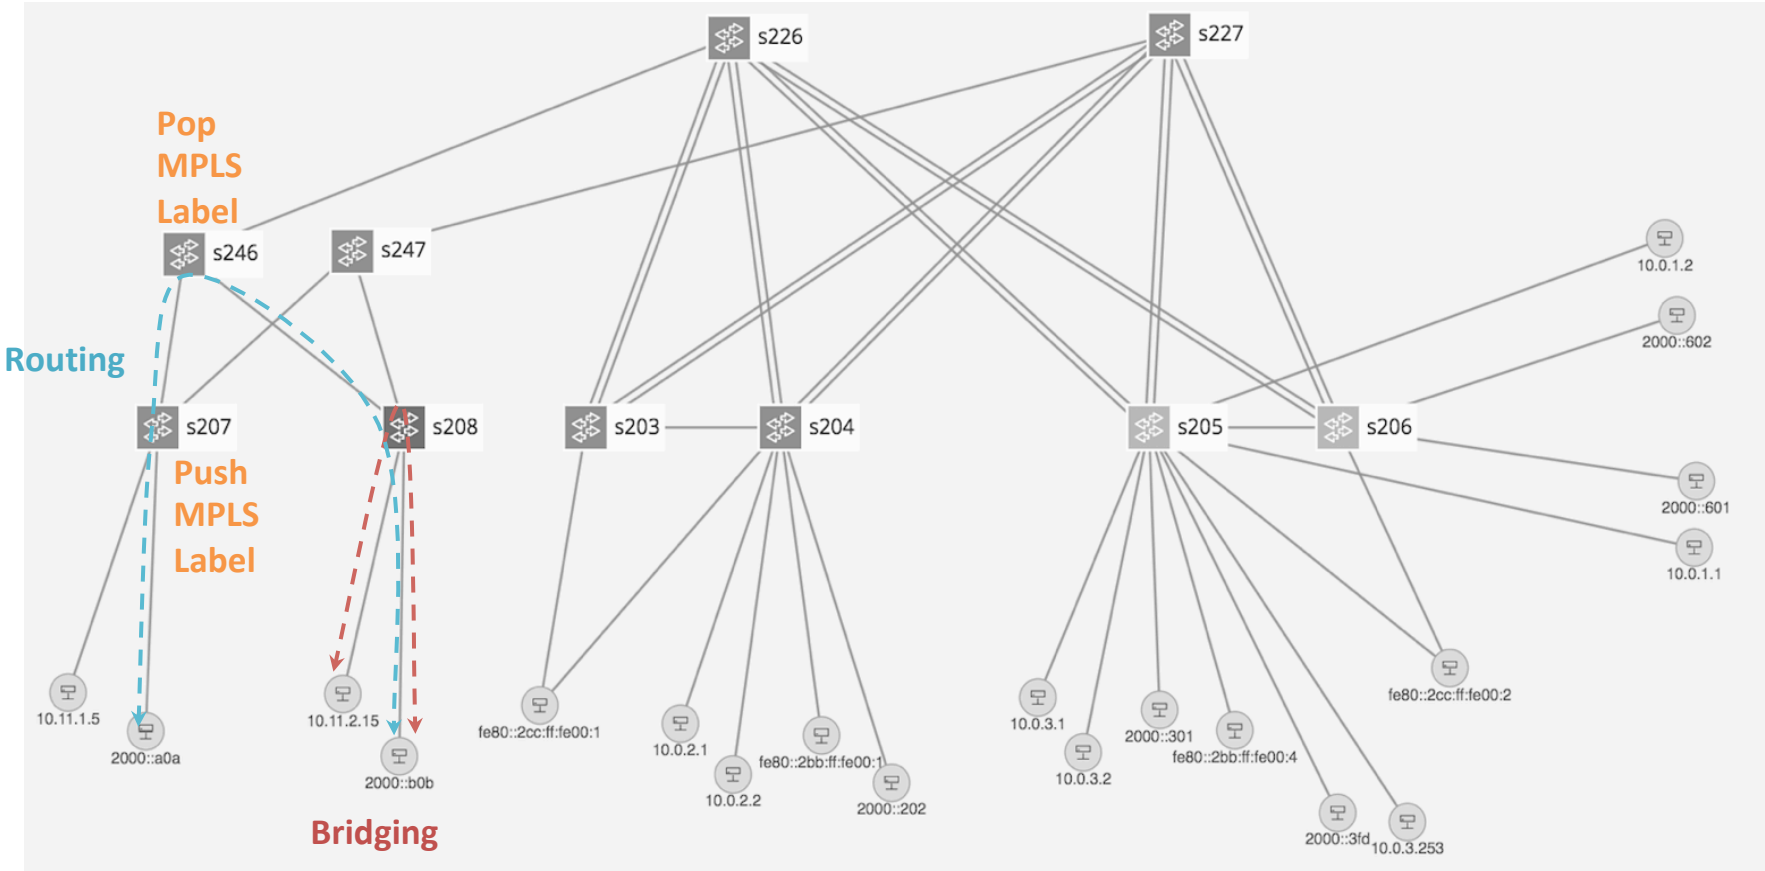

Trellis Features / Multicast (1/3)

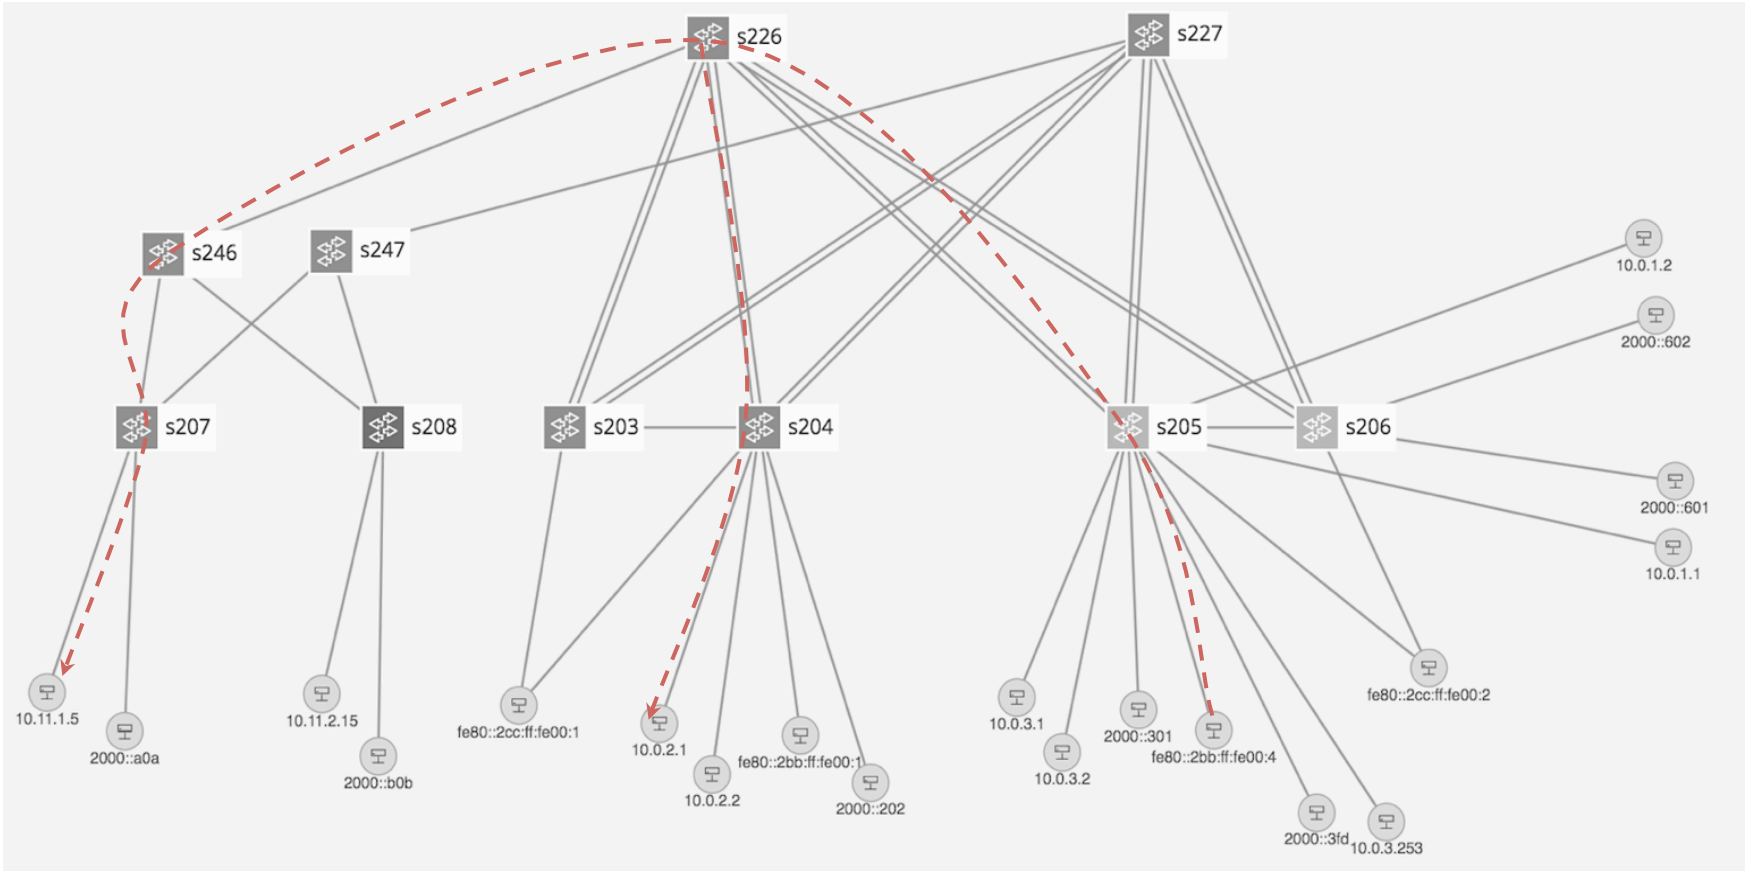

Single-homed source, single-homed sink

Trellis Features / Multicast (2/3)

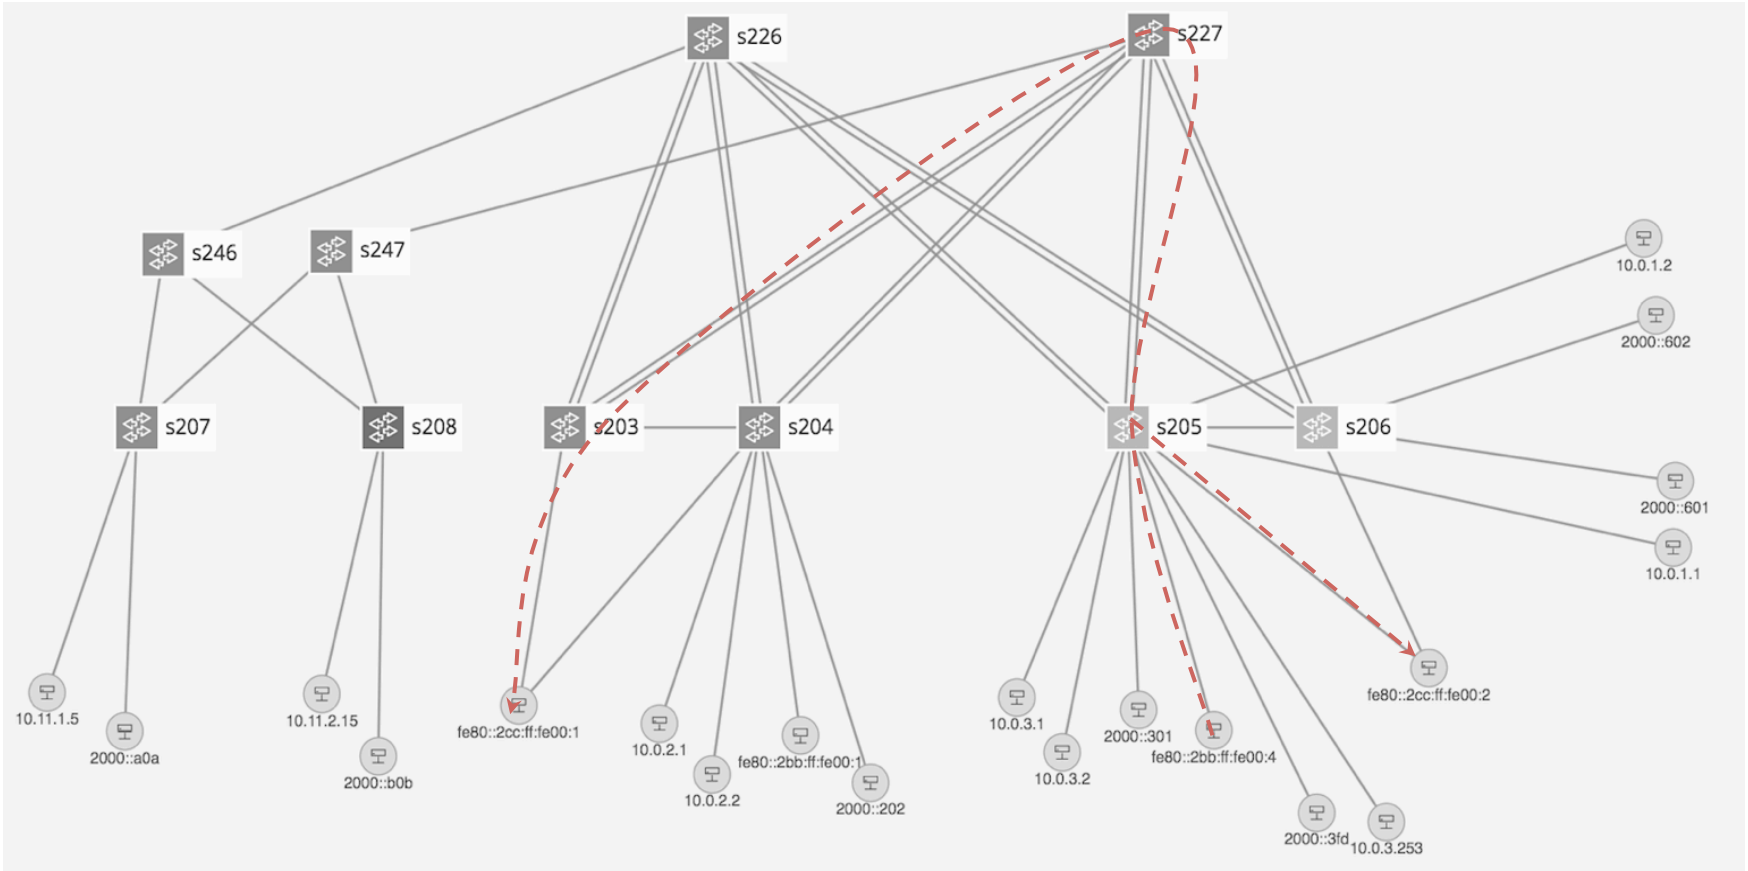

Single-homed source, dual-homed sink

Trellis Features / Multicast (3/3)

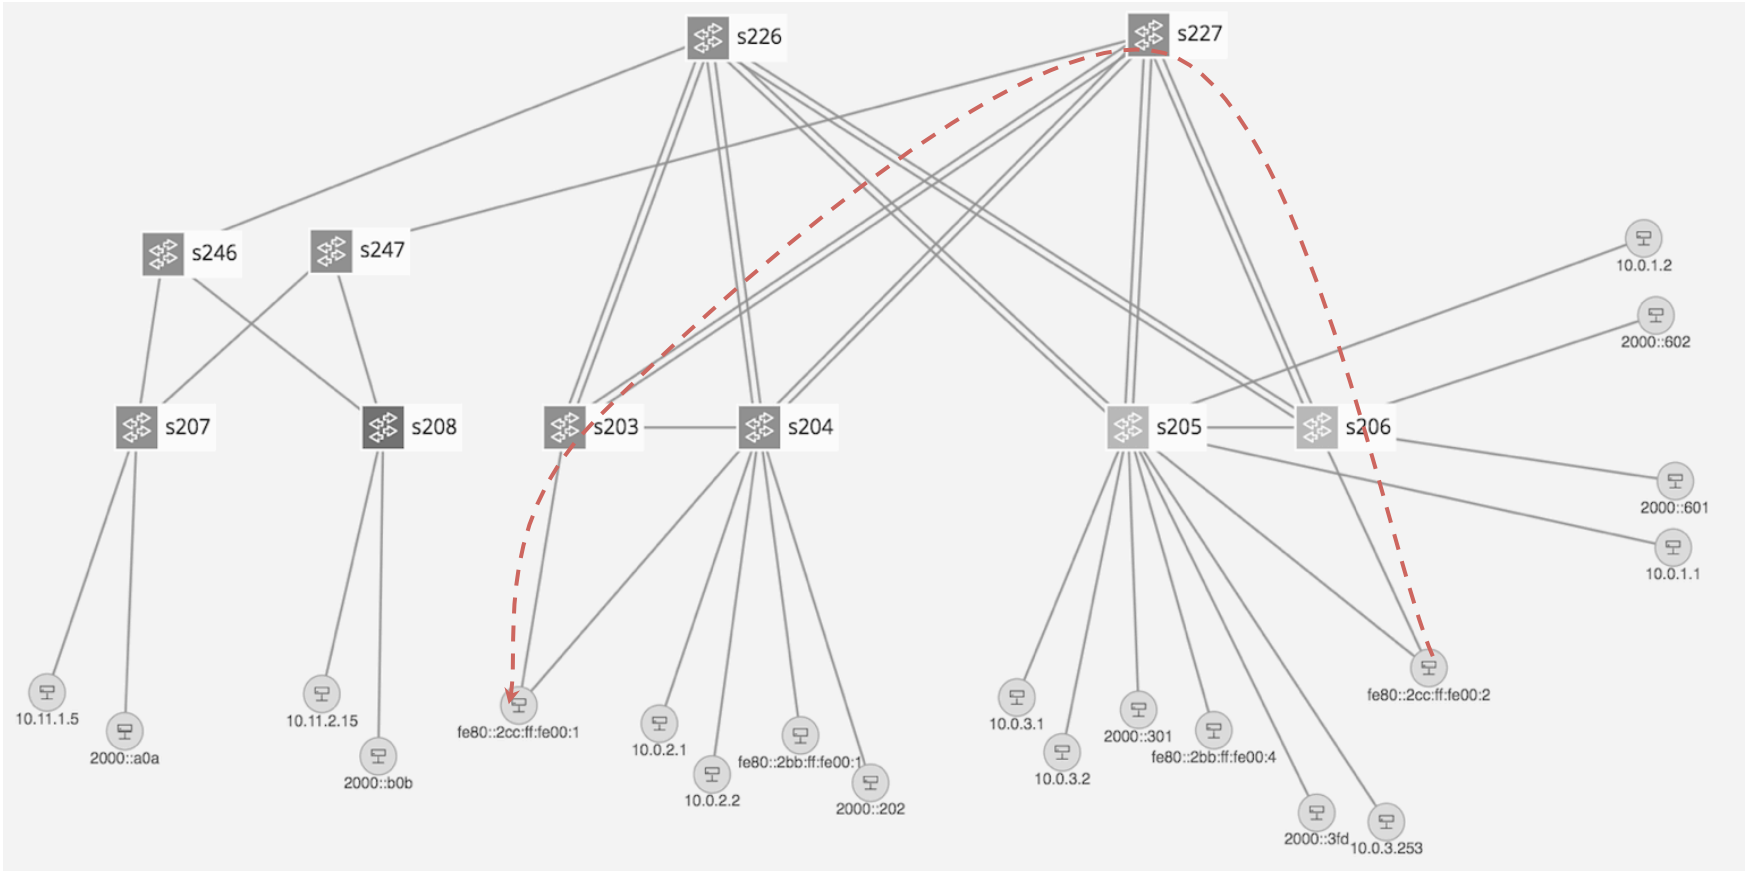

Dual-homed source, dual-homed sink

Trellis Features / Dual-Homing

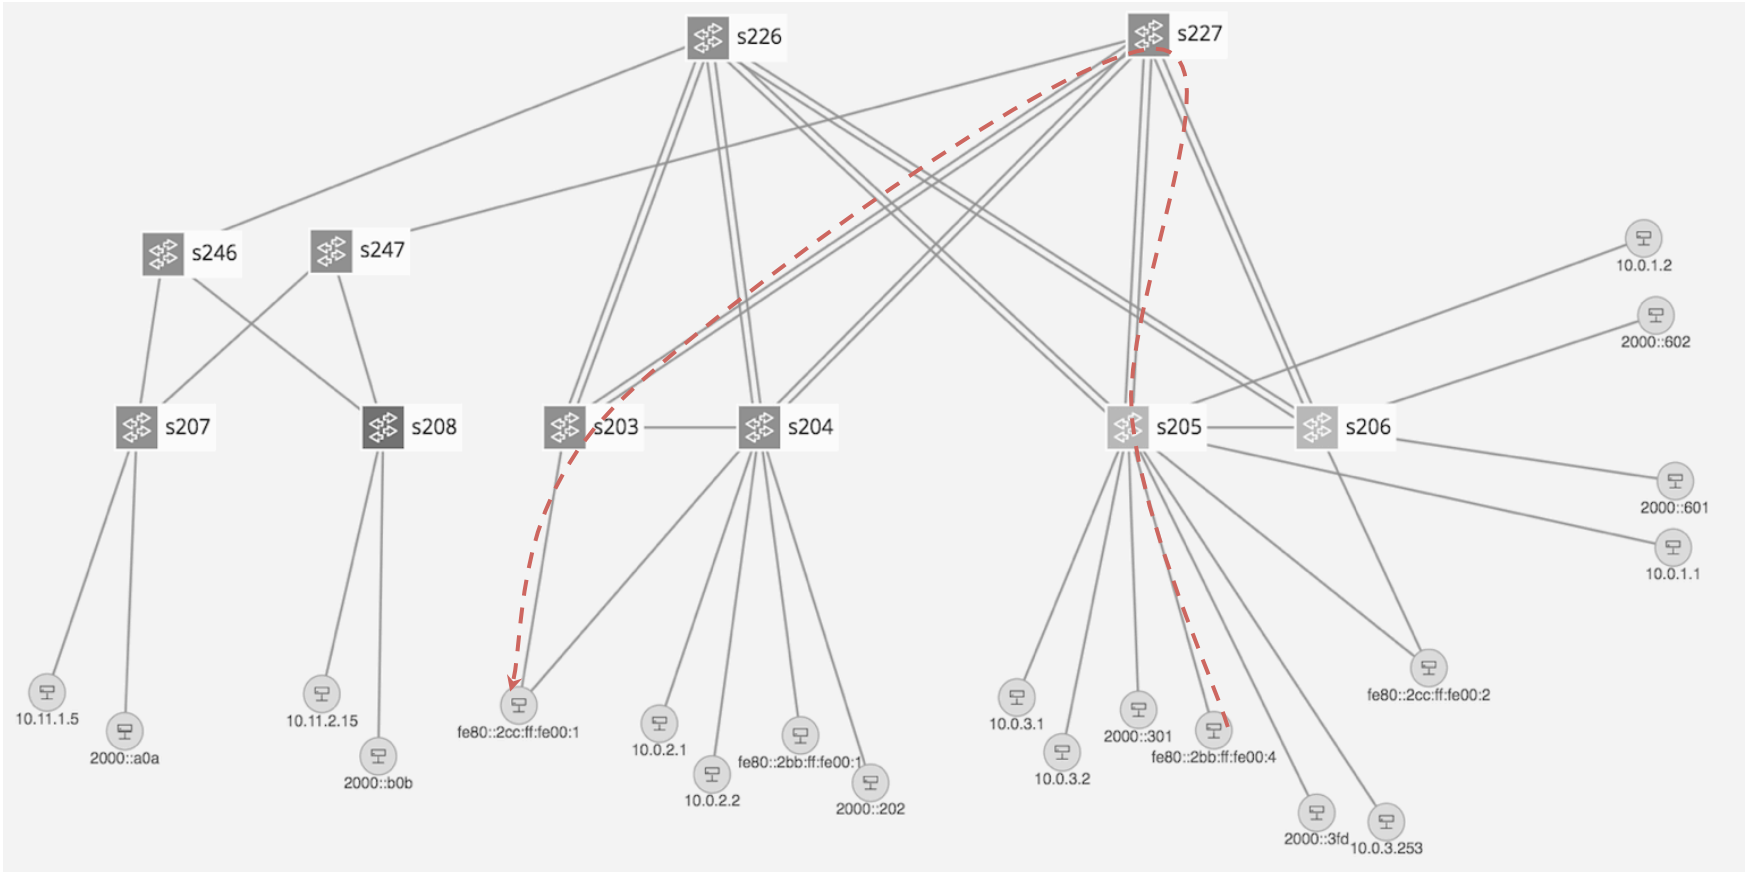

Trellis Features / Dual-Homing / Failure (1/2)

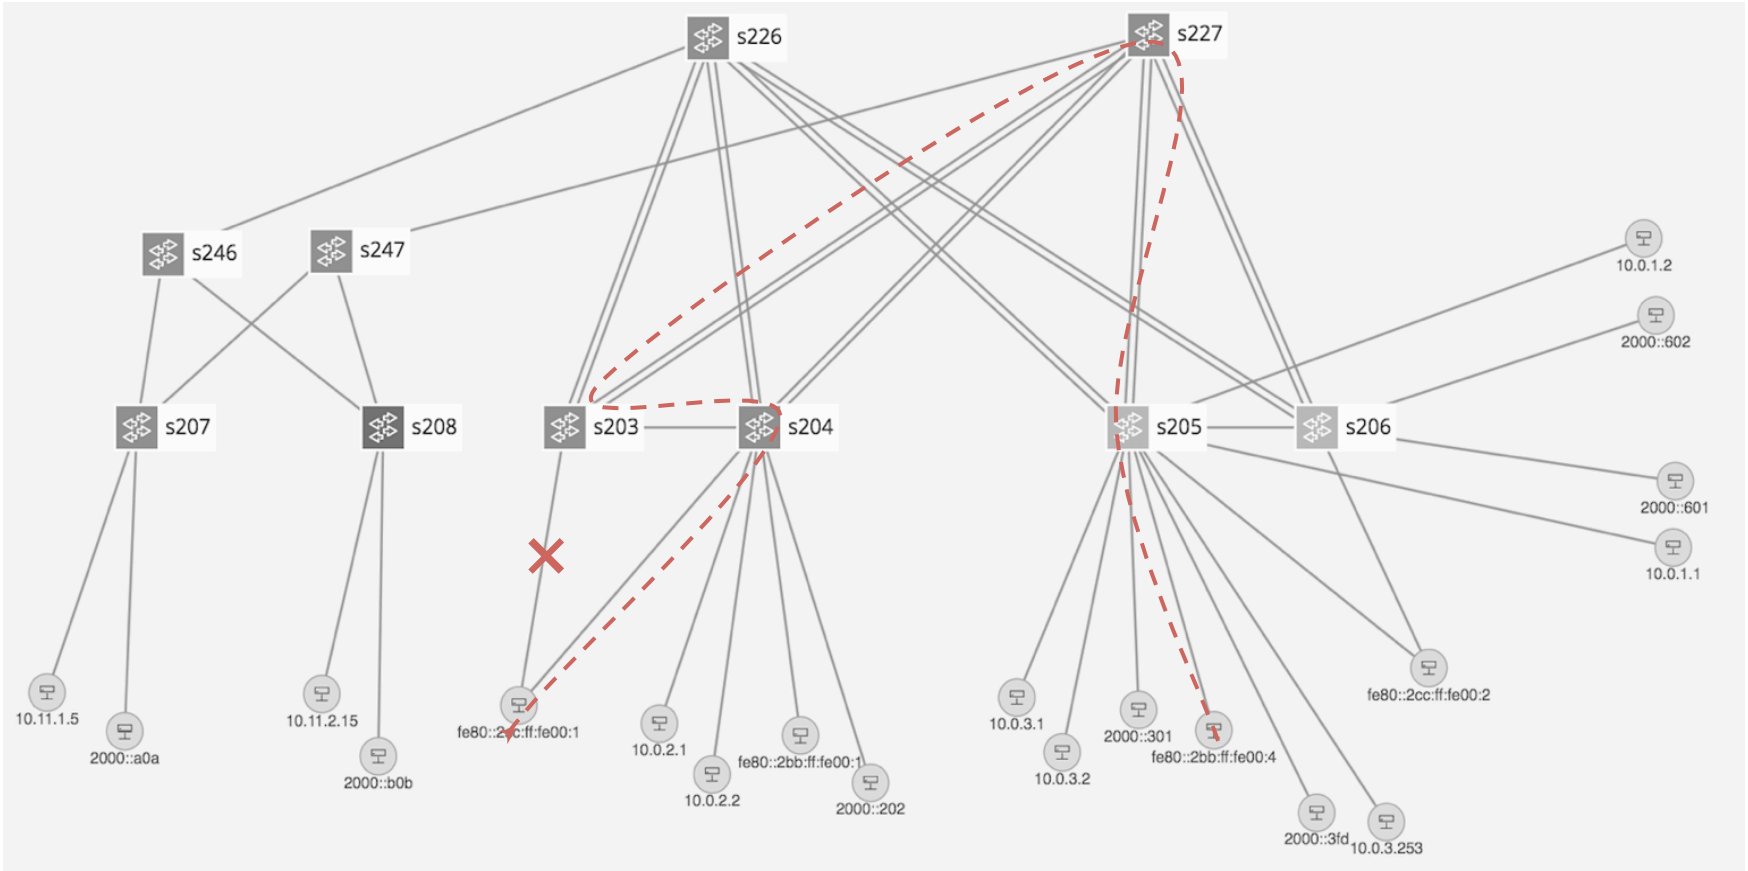

Pair link is only used to recover local failure

Trellis Features / Dual-Homing / Failure (2/2)

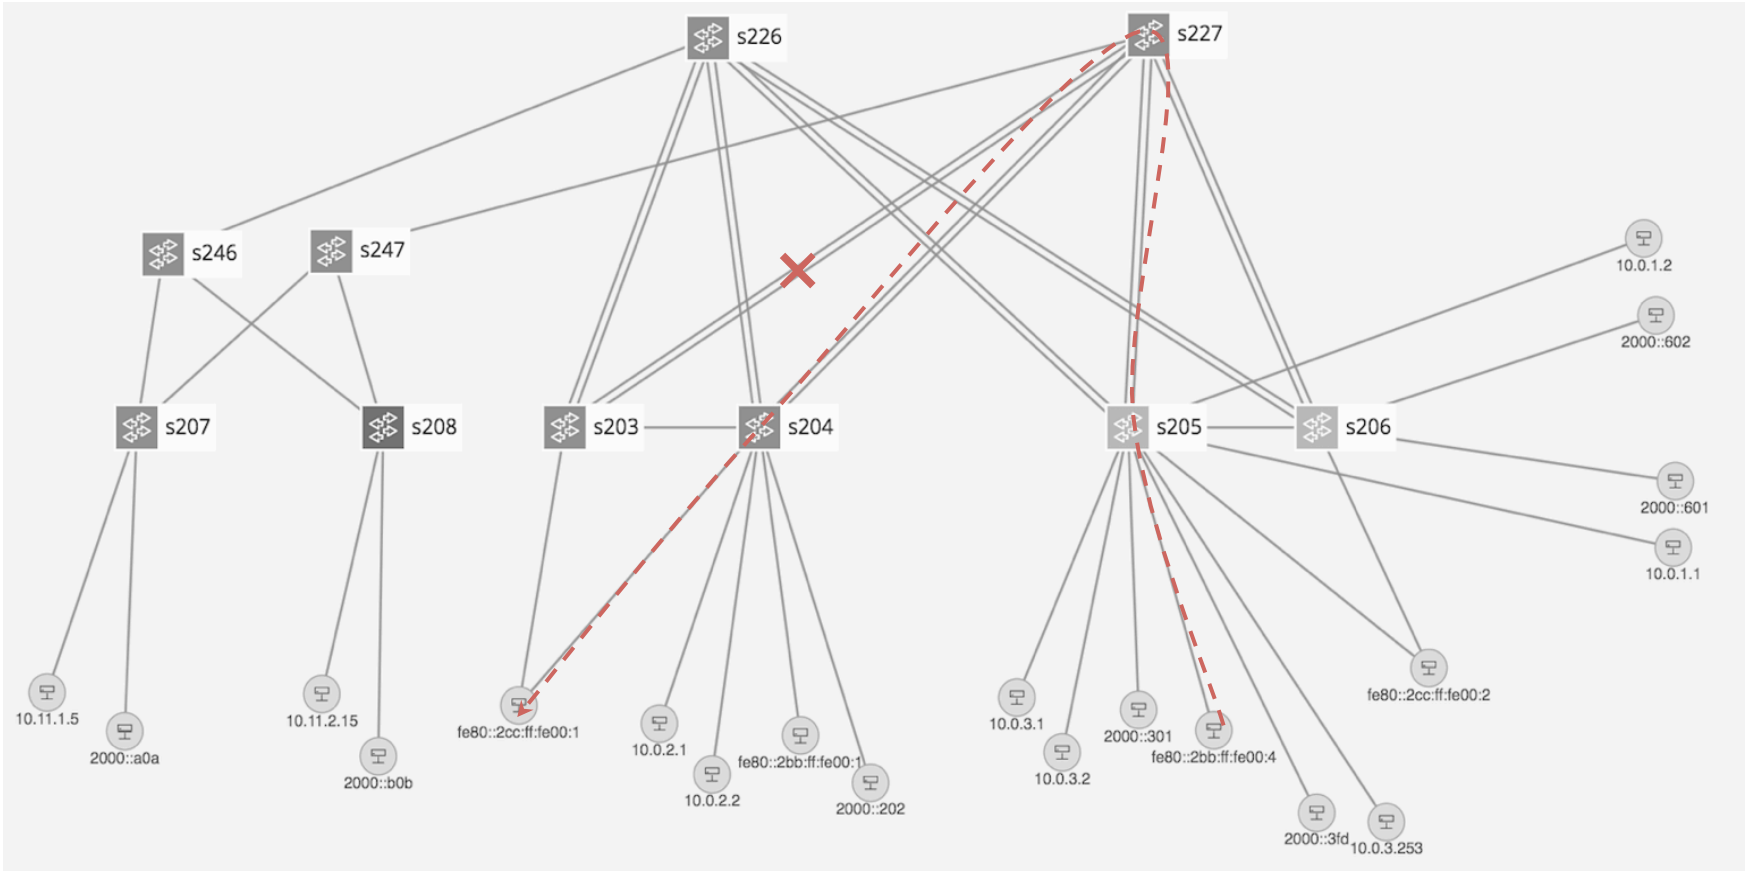

Trellis Features / Multi-Stage

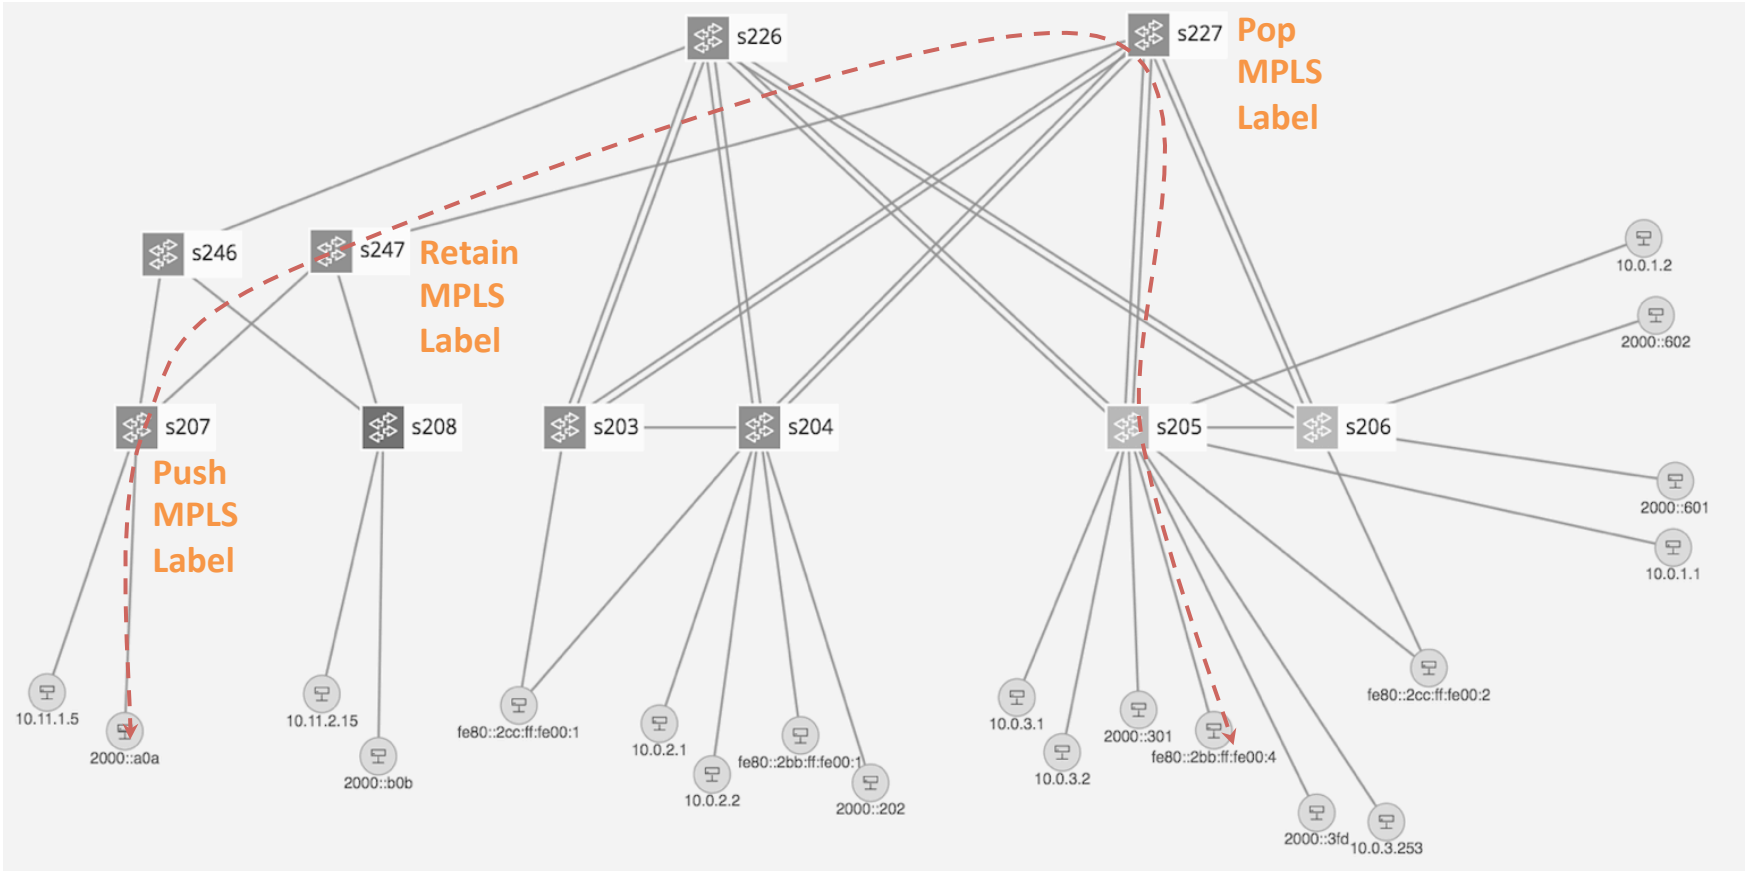

Trellis Features / vRouter

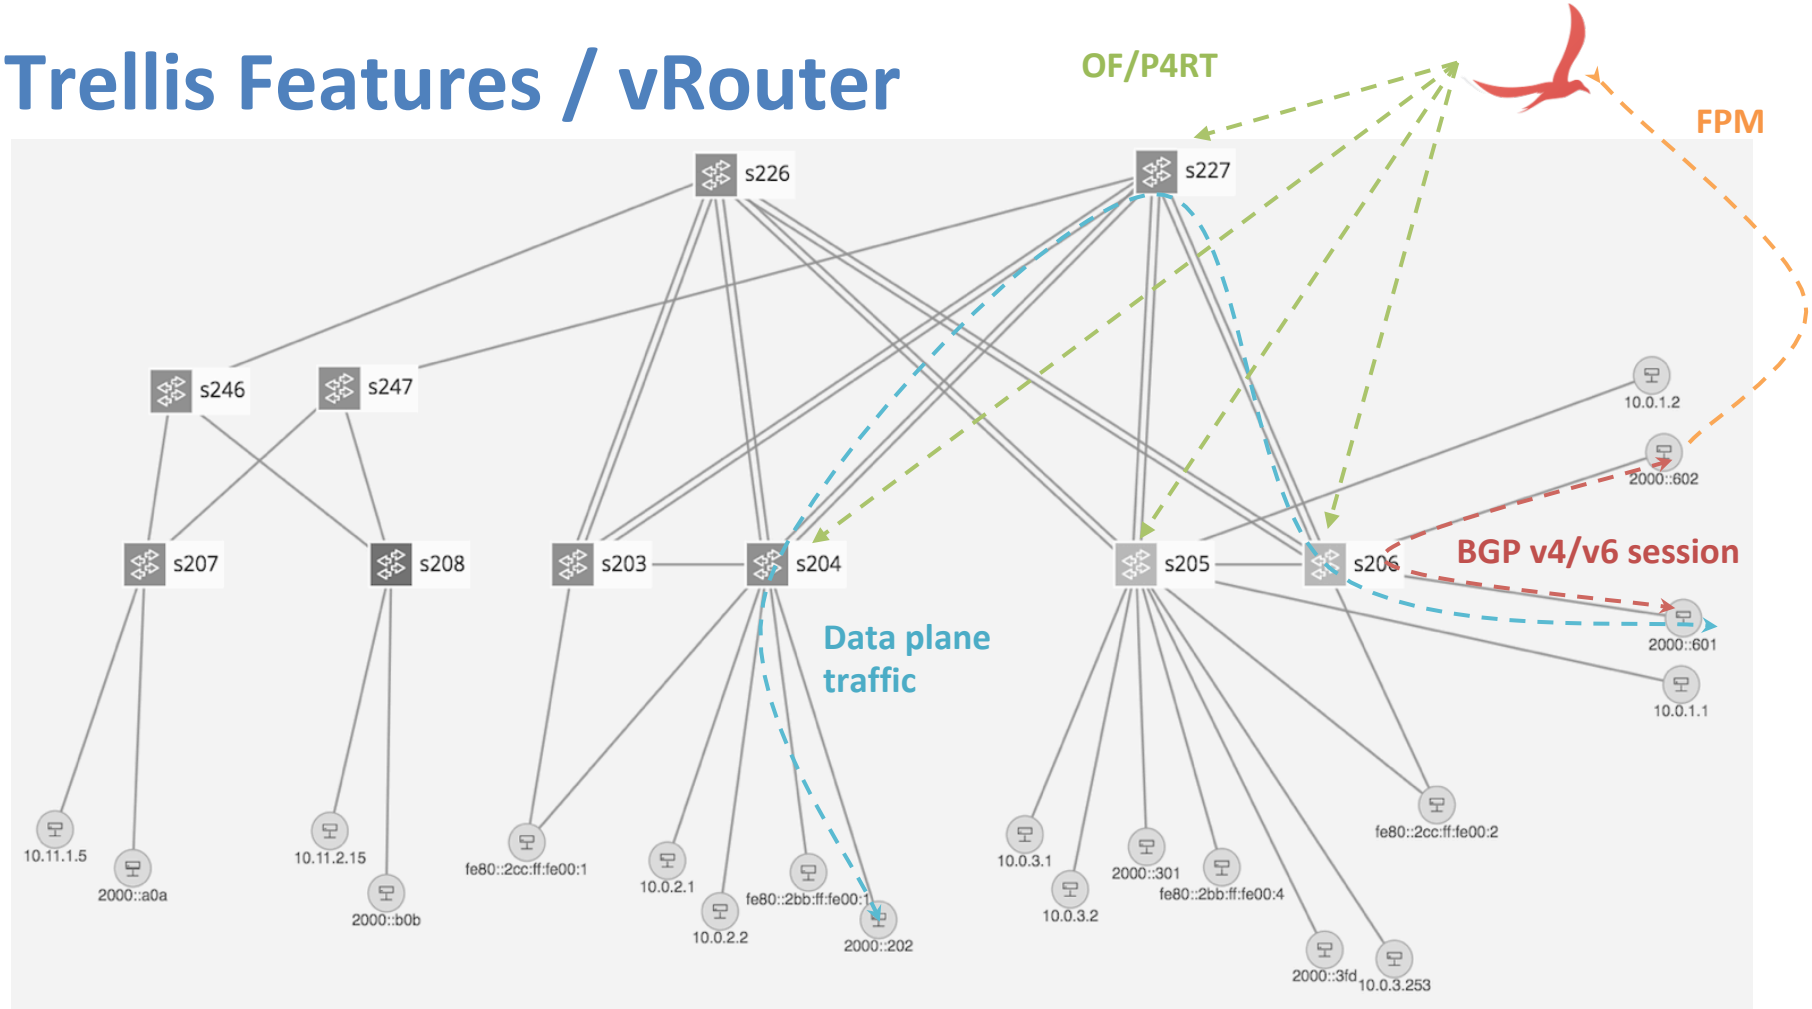

vRouter

- Control / data separation
- Entire fabric as a big router
- Also supports static routes and route blackholing

Trellis Features / DHCP L3 Relay

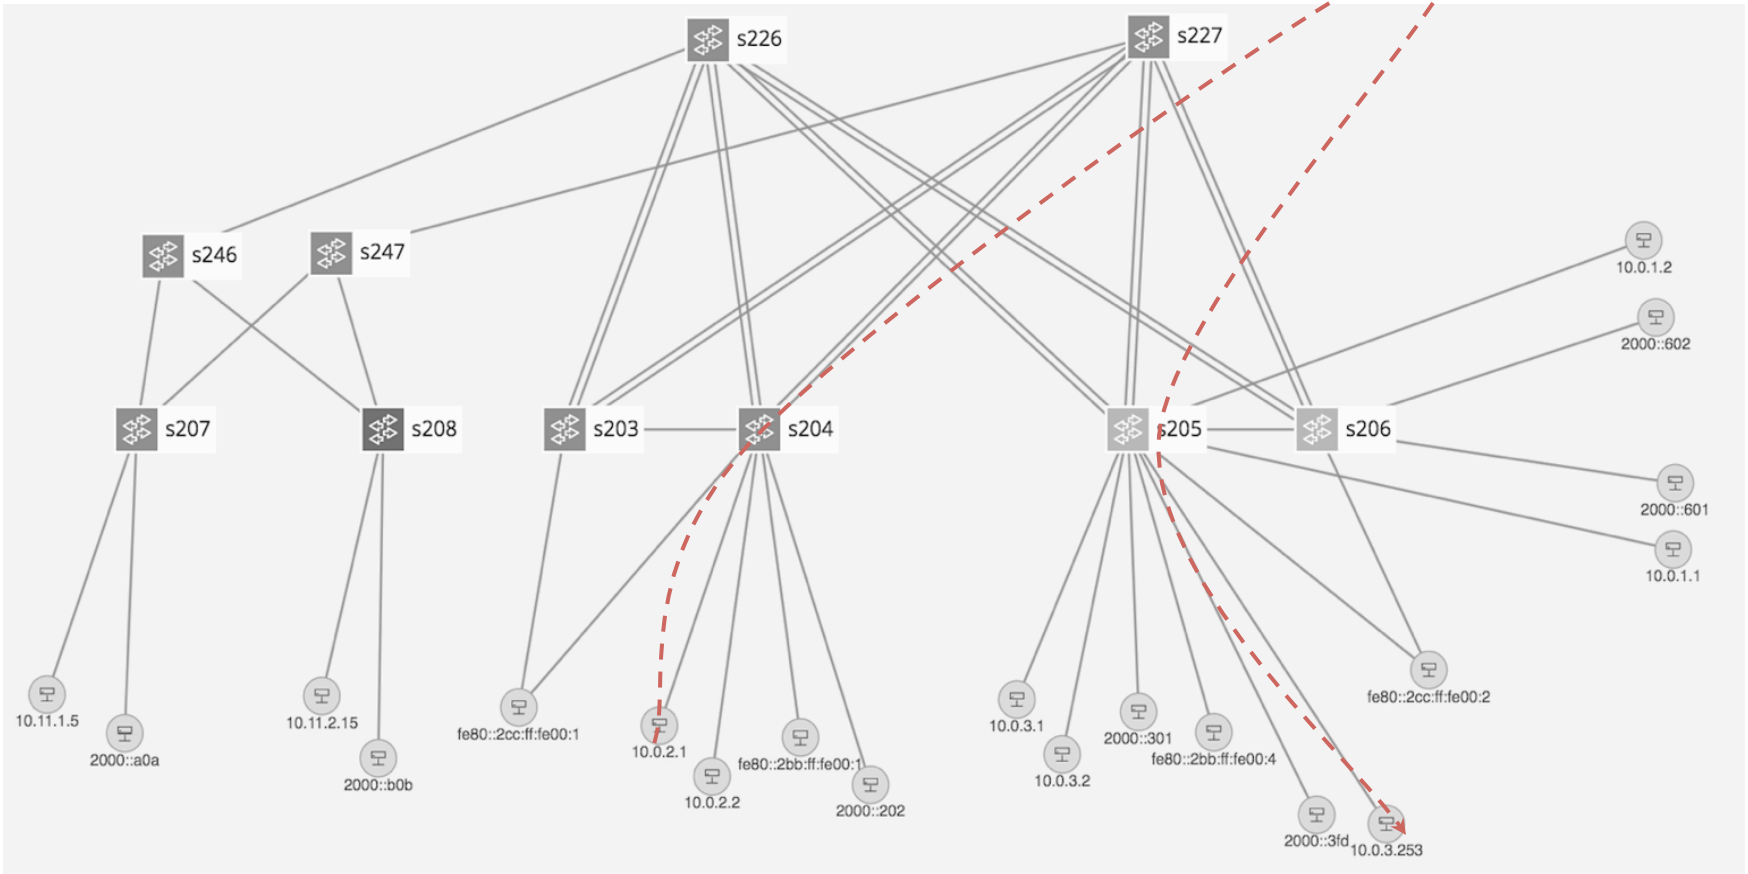

Trellis Features / Pseudowire

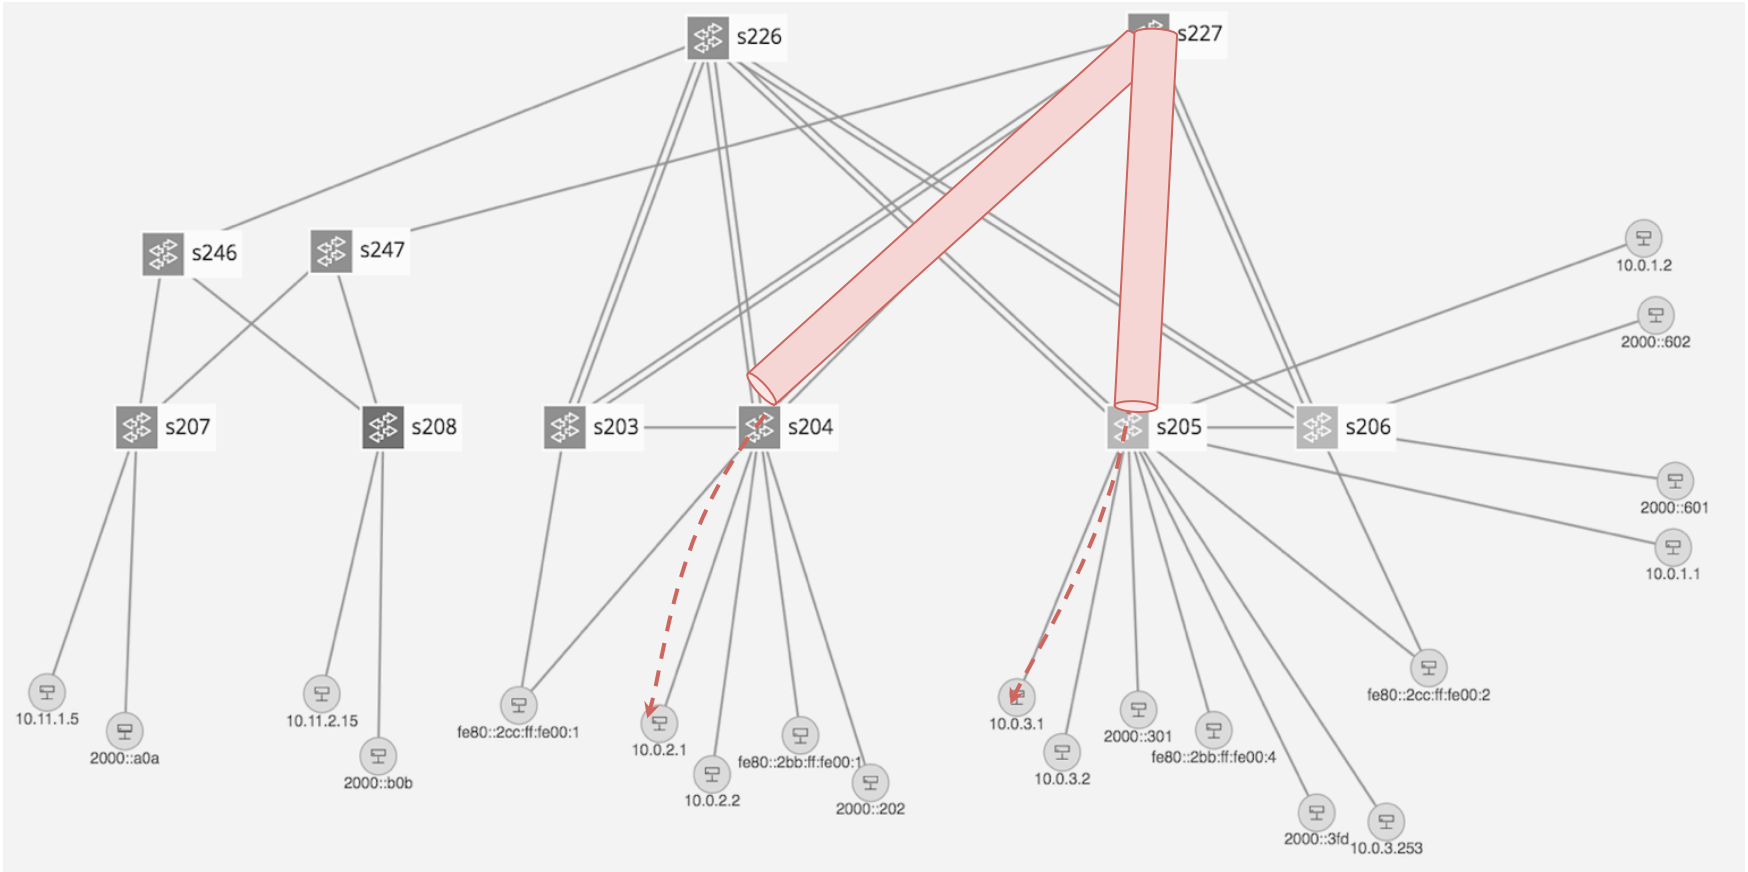

Trellis Features / Q-in-Q Termination

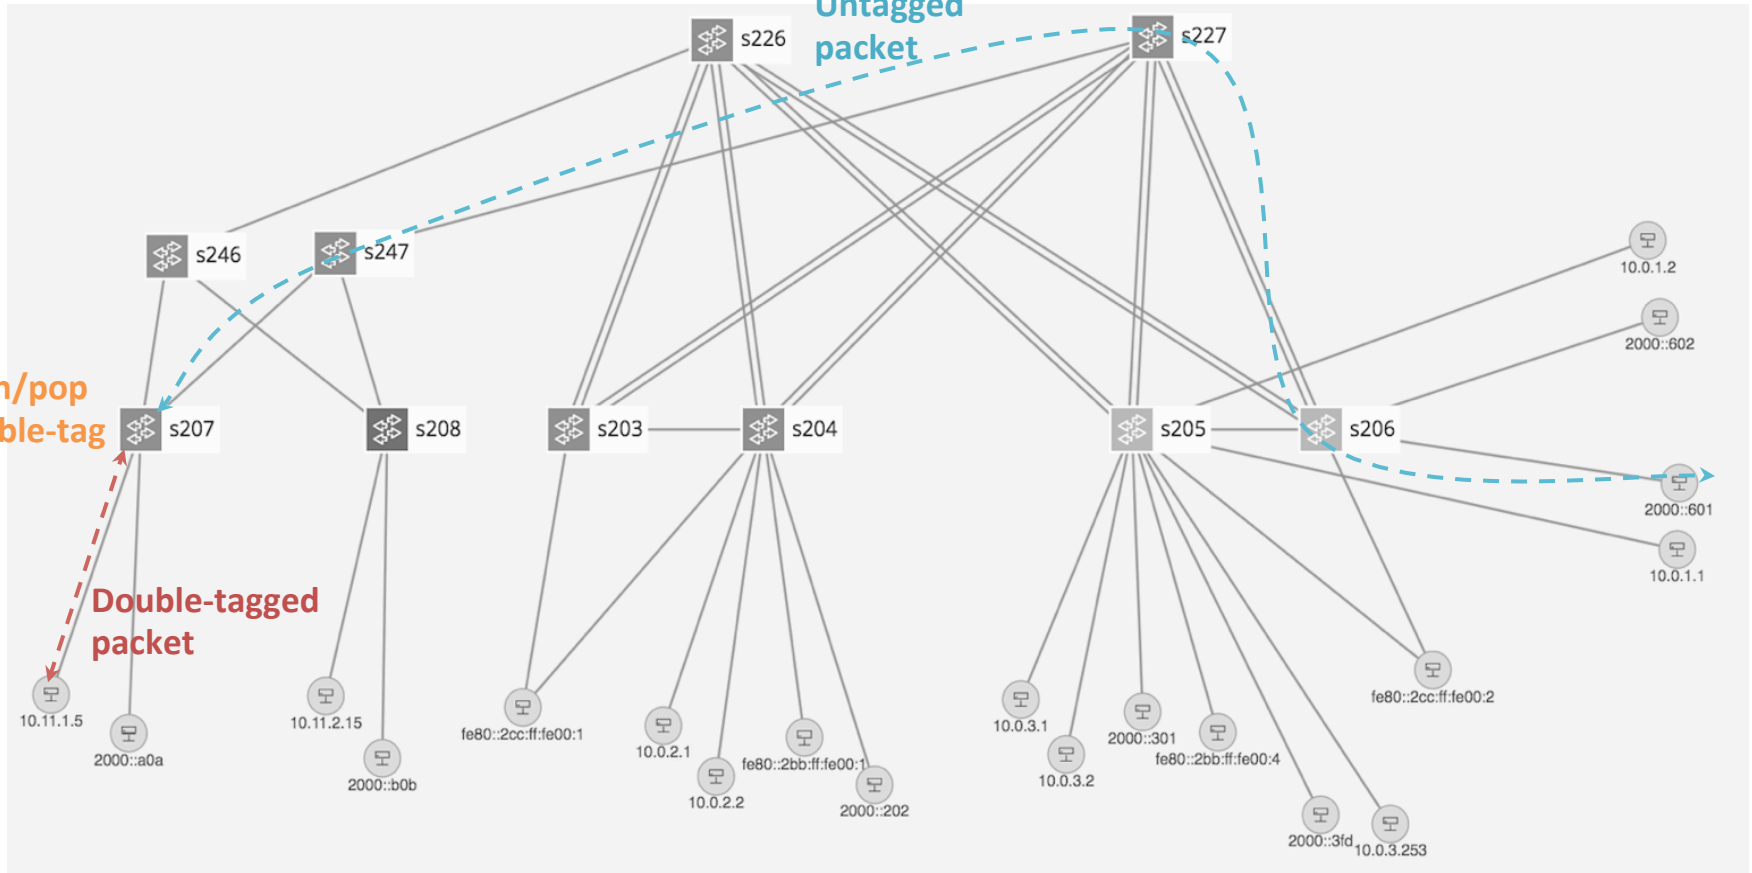

Works with single switch too

Various ASIC/Vendor Support

- Broadcom **Qumran, Tomahawk, Trident2** switches from EdgeCore, QCT, Delta, Inventec (WIP)
- Preliminary support P4-based **Tofino** switches from EdgeCore, Delta, Inventec
- Preliminary support for Mellanox **Spectrum** switches
- Preliminary support for Cavium **Xpliant** switches

Deployments & Use Cases

- Trellis in Comcast
  - Utilizes almost all Trellis features
  - Trellis is installed and monitored by Kubernetes
  - Integrates with in-house VNFs, logging, telemetry and alarm systems
- Trellis in CORD / SEBA
  - Utilizes bridging, routing, multicast, cross connect, dual homing
  - L2 load balance (WIP)

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

Recent Activities (in 2018)

Production Readiness

December 2017 - November 2018

Support Comcast

- **Support Comcast design/dev/QA/ops teams** - issue analysis; root-cause;
- **Recommend best-practices/ training**
- Design discussions for new features & architectural improvements
- **Daily scrum** (-May '18)
- Documentation

Support Other teams

- **Broadcom**
- **Nokia**
- **Harmonic**

Deliver New Features

- **Pseudo wires** for in-band control
- Routing in **H-Agg** based topologies
- **Multicast** improvements
- **Dual-homing** improvements
- **DHCP v4/v6 Relay**
- **IPv6 Router Adv.**
- **ISSU** architectural discussions/progress
- Other small features

ONOS Stability & Scale

- Focus on stability of ONOS distributed stores (**9 releases of Atomix in 4 months**)
- Scale investigation ongoing

Tooling

- **T3** - Trellis Troubleshooting Tool
- **onos-diags** - Diagnostics collection tool
- **Mininet scripts** - Software-emulated test environment

QA & Bug Fixes

- Developing automated feature tests (**220 new tests in the 4 months**)
- Extending framework for hardware based tests
- **180 tickets resolved**

Stratum/P4 Integration

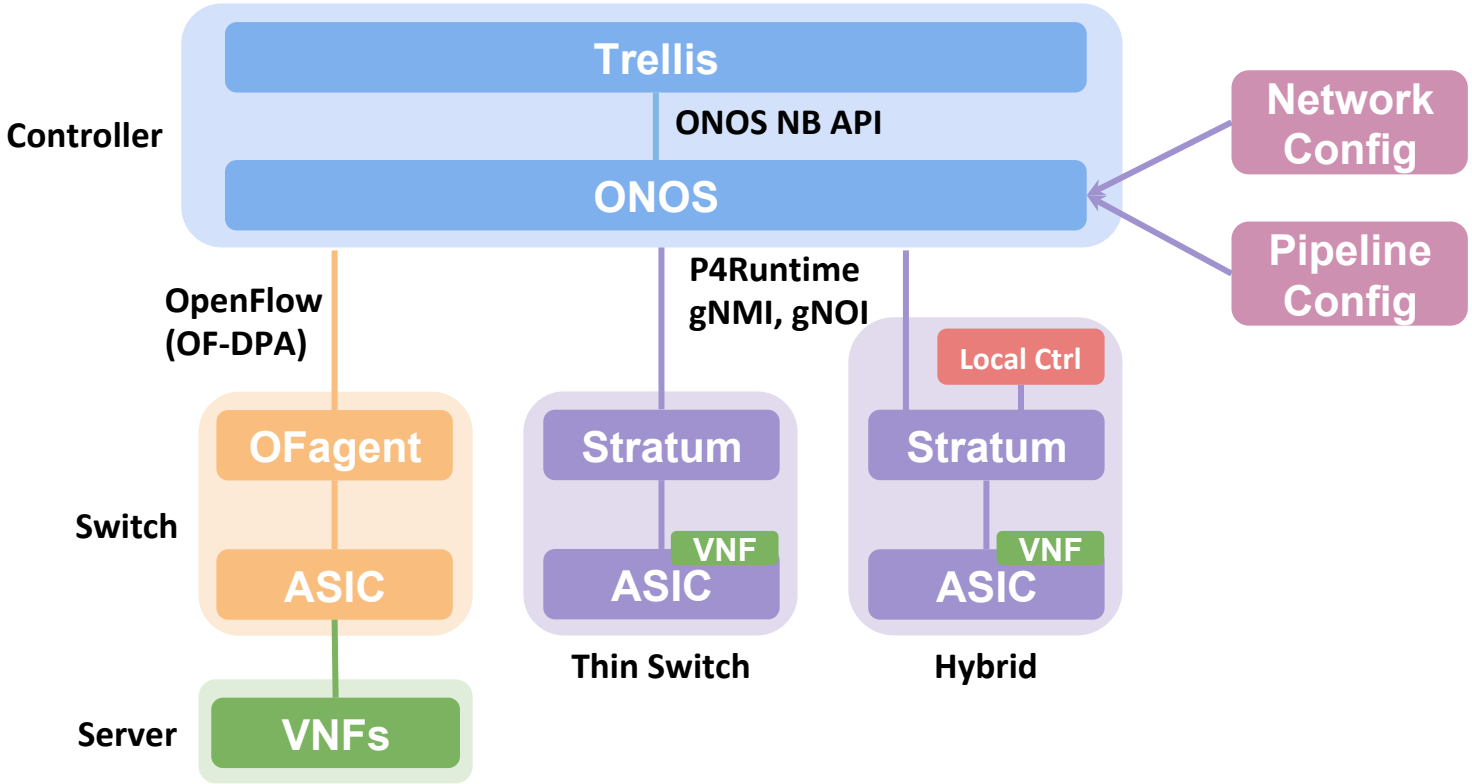

Ecosystem & Roadmap

Distributed DevOps Model

- **Analyze issues** reported by Harmonic in their setups
- **Code review** patches they submit to ONOS
- Design discussions for new features
- Bringing up ONF pod housed in San Jose

- **Report issues found; ~75 cases**
- Create automated tests to reproduce
- **Validate fixes** / releases
- Help with issues in their setup
- **Design discussions** for new features
- Daily scrum

- Report issues found in hw
- Validate T2 versions of switch software EdgeCore builds for us

- Bringing up ONF pod hosted by Flex
- QA collaboration

- Deliver **features**; meet ops needs
- Deliver **stability**, scale, perf, tools
- Support Comcast design/dev/QA teams – issue analysis; root-cause; **recommend best-practices/ training**
- **Design discussions** for new features
- Daily scrum
- **Documentation**

- **Design discussions** for apps/features created by Nokia
- **Code review** submitted patches (40 changesets; 4-5 patchsets/change)
- Validate some features; report issues


AT&T


COMCAST


NTT


Edge-corE
NETWORKS


Türk
Telekom


T


Fiber


Telefónica


CAVIUM


radisys


P4


BROADCOM


harmonic


ciena


flex


BAREFOOT
NETWORKS


NOKIA


QCT


Infosys


OPEN
Compute Project

and more…


ONF

Roadmap

- Scale & Performance improvements
- Dual homing for Access nodes (like OLTs)
- In Service Software Upgrades (ISSU)

- Stratum/P4 integration
- BNG features (e.g. PPPoE termination, hierachical QoS)
- 5G user plane features

Next-Gen SDN Demo

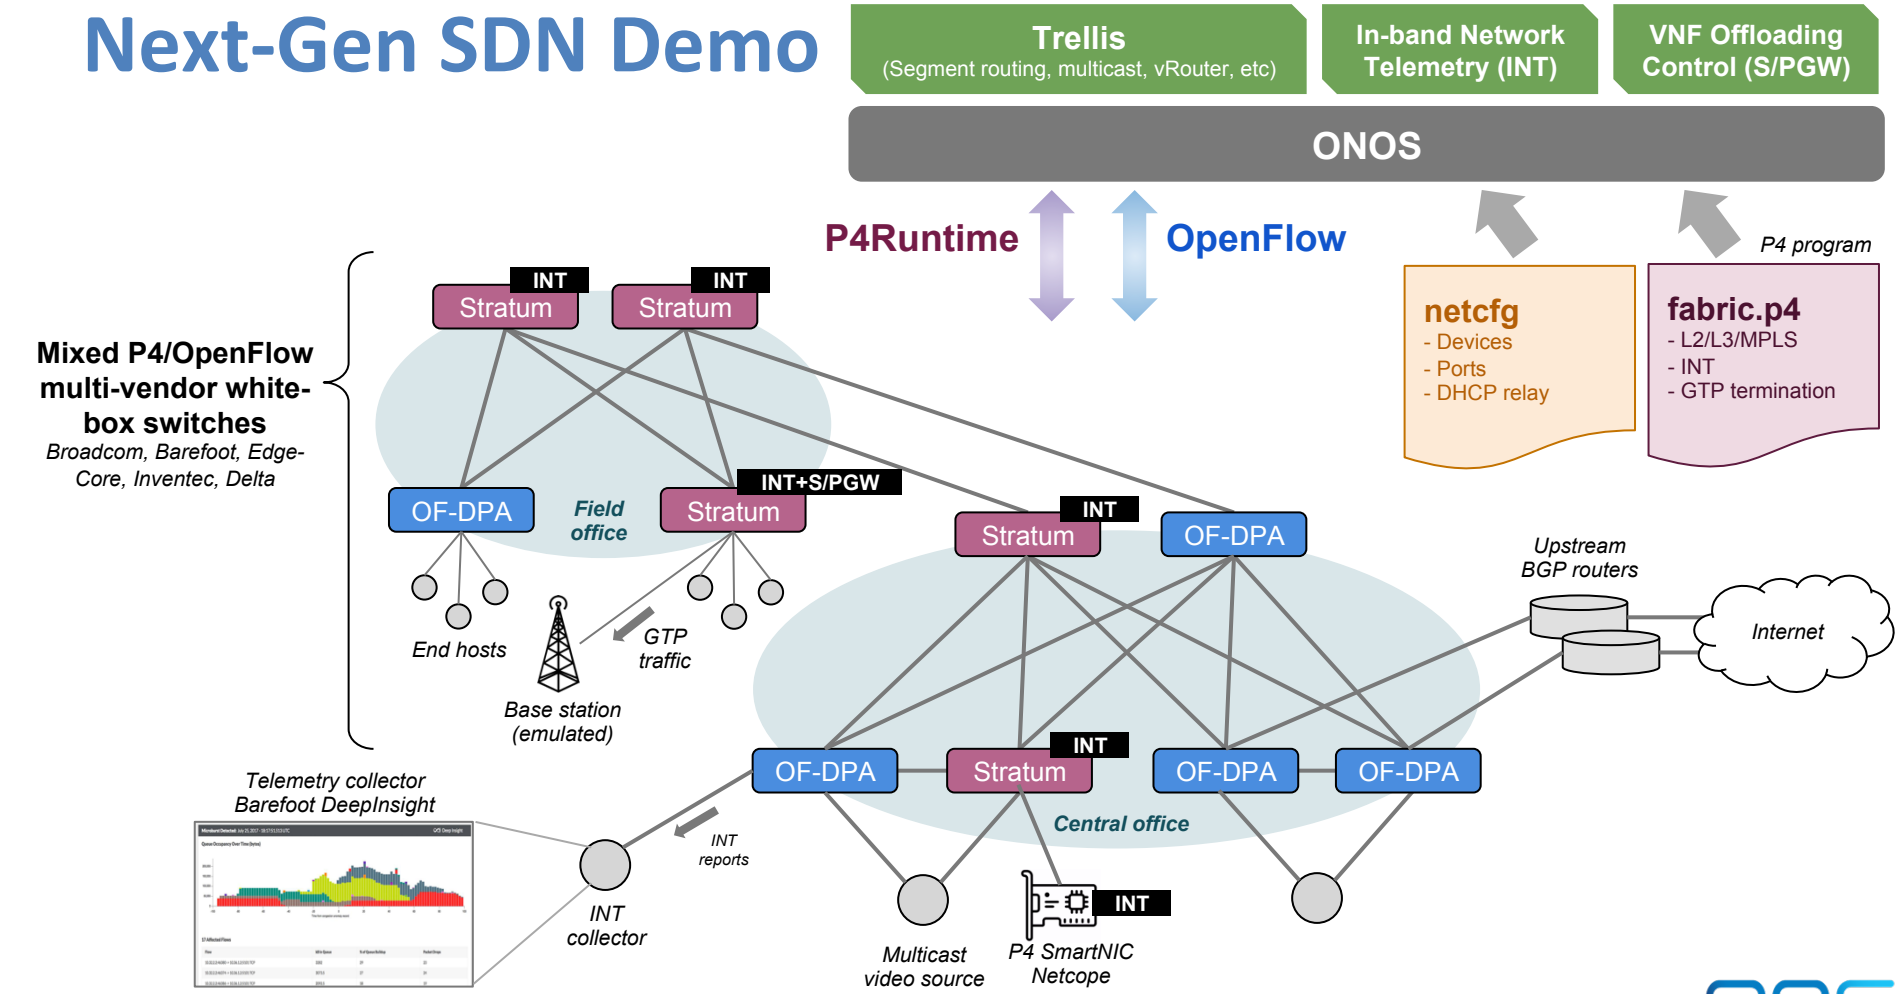

Charles Chan
charles@opennetworking.org