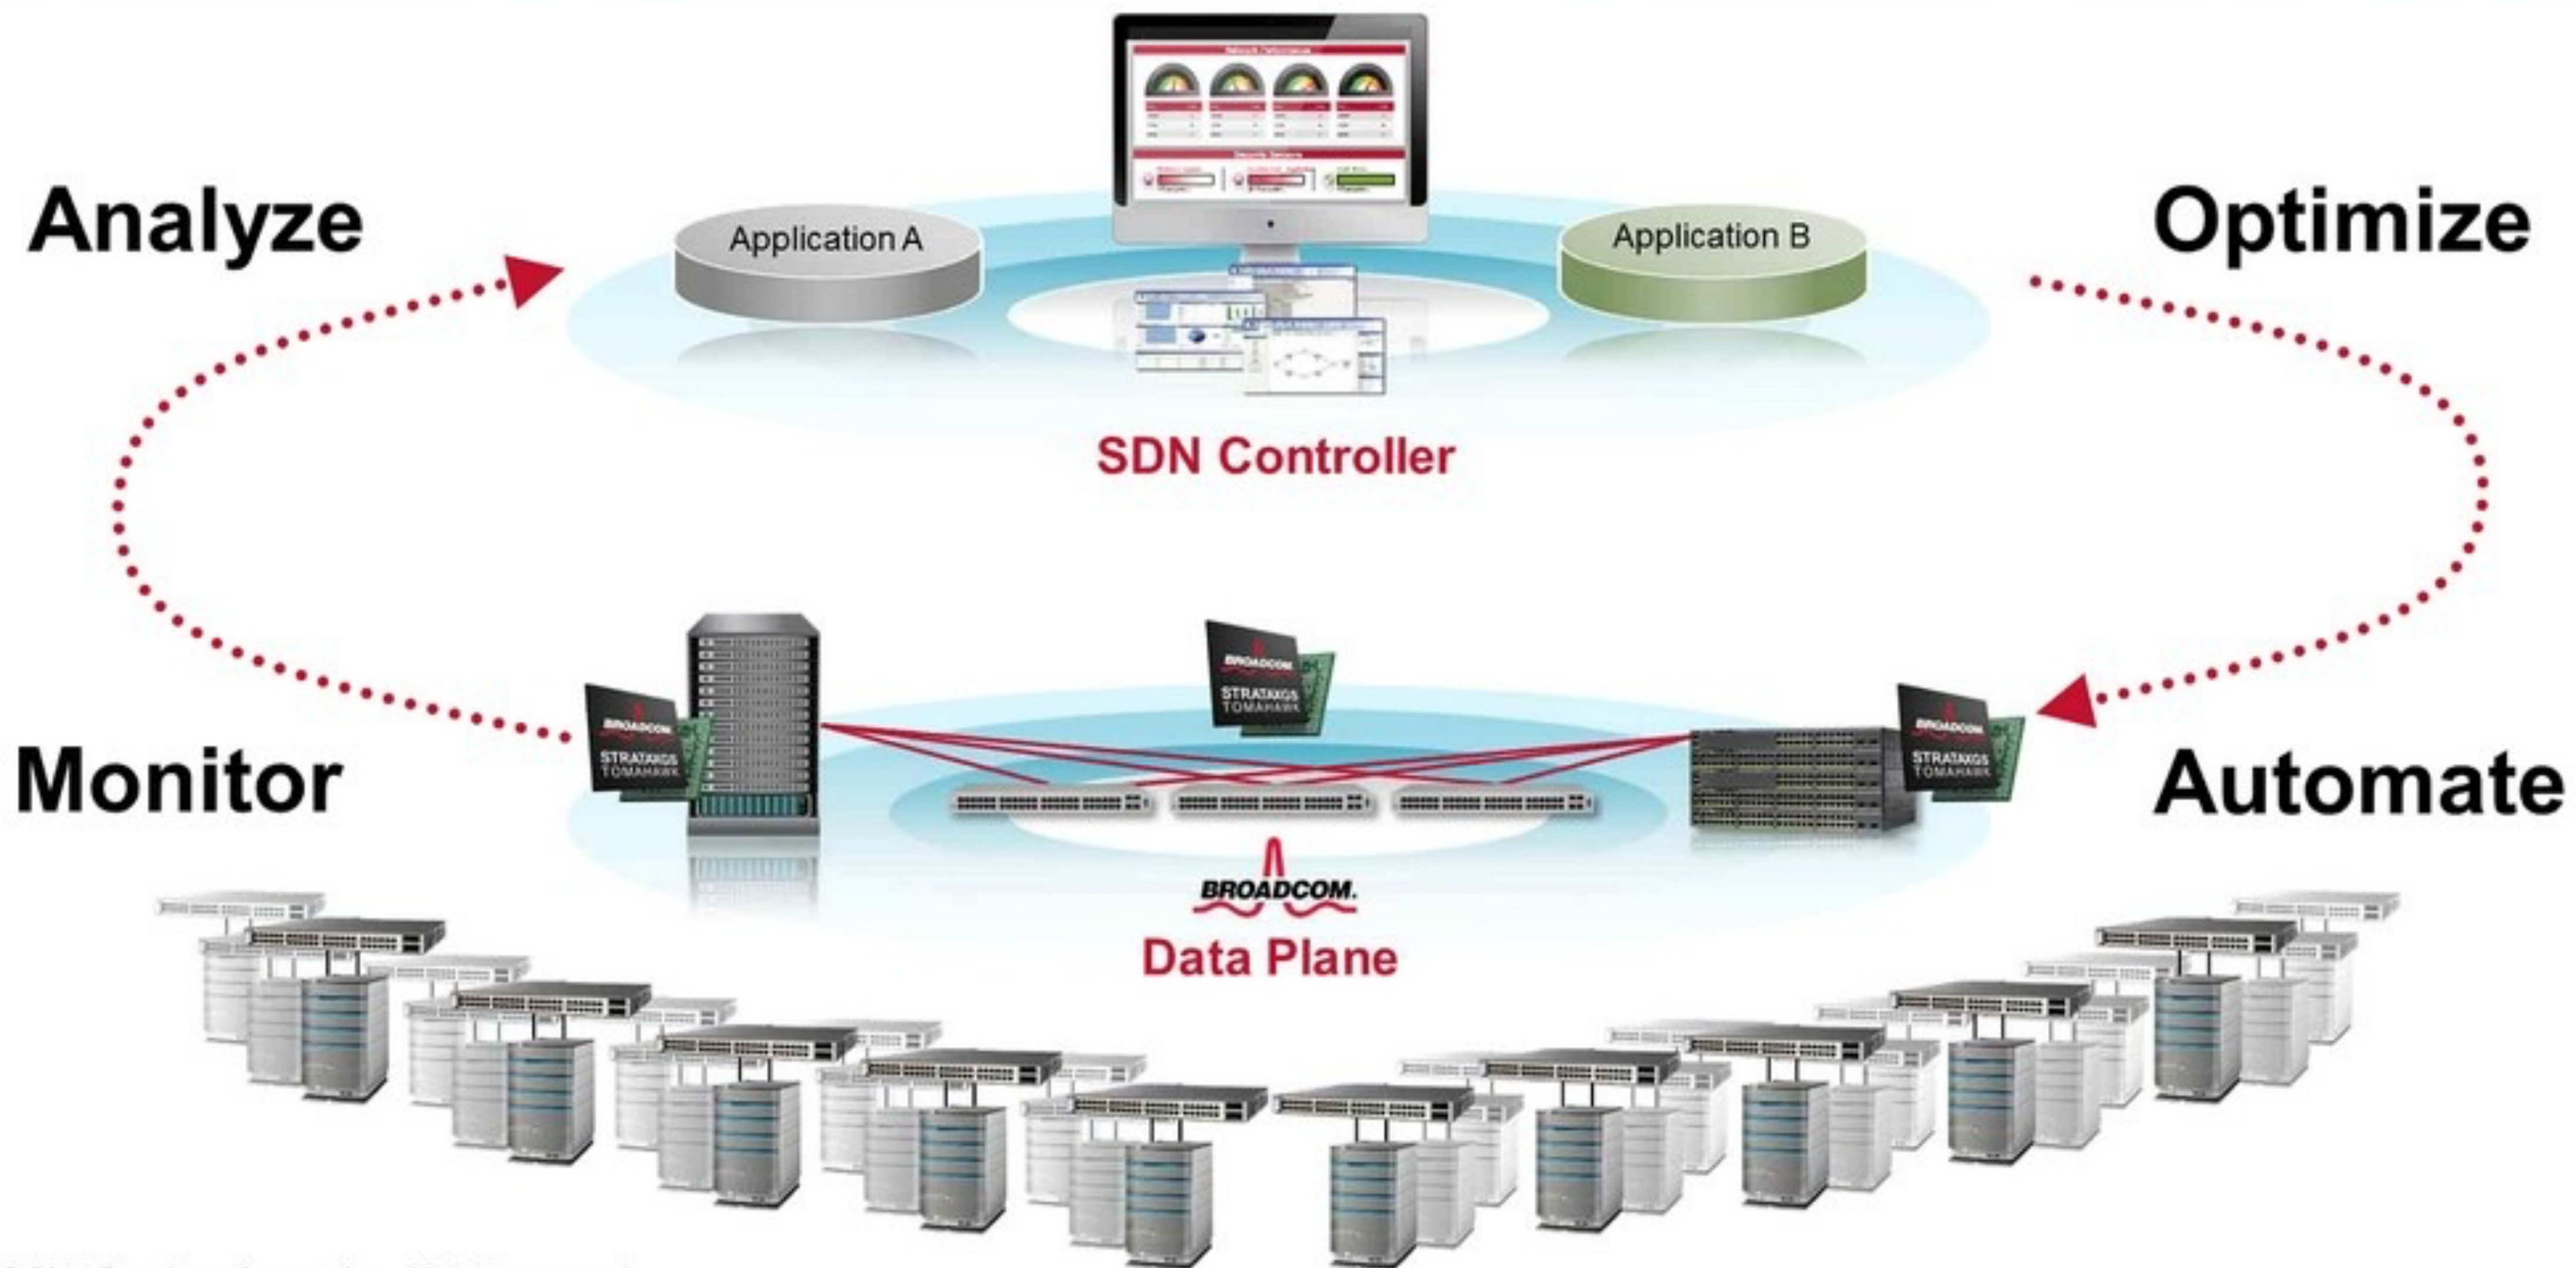


New OVS instrumentation features aimed at real-time monitoring of virtual networks

Open vSwitch 2015 Fall Conference

Peter Phaal
InMon Corp.
November 2015

SDN IS ABOUT NETWORK VISIBILITY AND CONTROL




sFlow: Embedded Infrastructure Monitoring

- Industry standard measurement technology integrated in merchant silicon, white box switches, and branded switches from over 40 vendors (sFlow.org)
- Supported in Open vSwitch since v0.99.1 (25 Jan 2010)
- Open source Host sFlow agent for hosts and hypervisors (sFlow.net)
- Periodic export of standard counters
- Asynchronous export of random packet samples

Comprehensive low latency monitoring required for closed-loop control

Traffic visibility and control with sFlow,
Peter Phaal, InMon, [OVS 2014 Fall Conference](#)

New OVS Specific sFlow Counter Metrics



- Datapath hits
- Datapath misses
- Datapath lost
- Datapath flows
- OVS CPU user
- OVS CPU system
- OVS memory

sFlow agent sends OVS metrics periodically sent with standard interface counters

Managing Open vSwitch across a large heterogeneous fleet,
Chad Norgan, Rackspace, [OVS 2014 Fall Conference](#)

sFlow Packet Flow Sampling

- **Packet Flow** A Packet Flow is defined as the path or trajectory that a packet takes through a Network Device (i.e. the path that a packet takes as it is received on one interface, is subject to a switching/routing decision and is then sent on another interface.
- **Packet Flow Sampling** Packet Flow Sampling refers to the random selection of a fraction of the Packet Flows observed at a DataSource.

Packet Header + Forwarding Information (FIB)

http://sflow.org/sflow_version_5.txt

Capture Actions with Packet

Kernel patch (upstreamed)

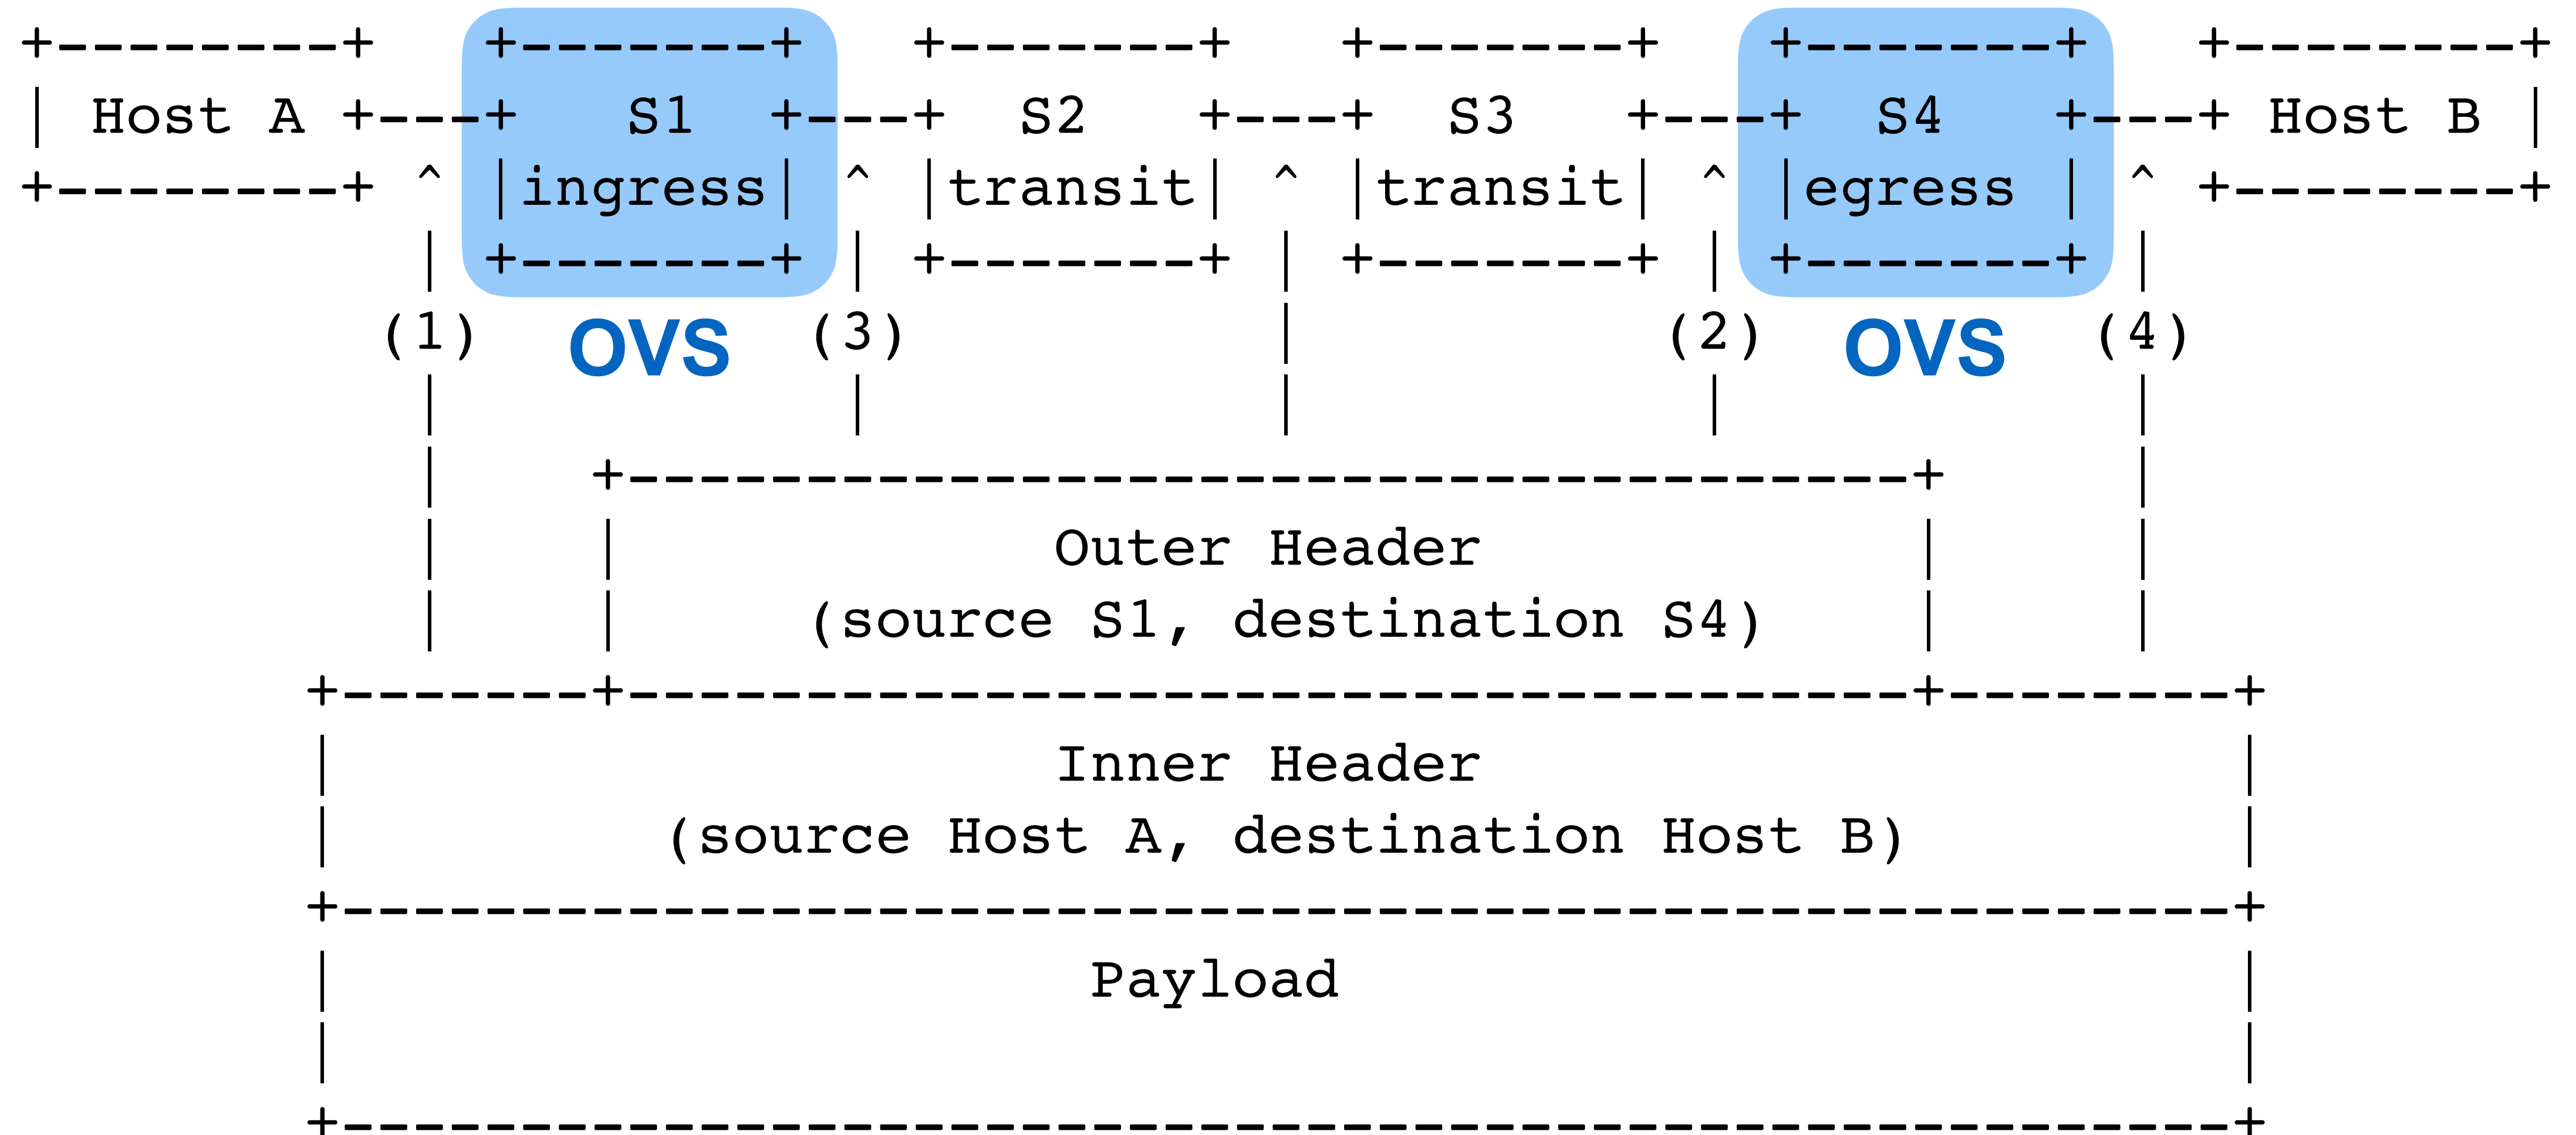
```
- return output_userspace(dp, skb, key, a);  
+ return output_userspace(dp, skb, key, a, actions, actions_len);
```

Fall back to user space rule cache when running on older Kernel

Features:

- Future proof - all actions available to user space agent, no longer limited by information that can be encoded in cookie
- Accurate - packet and actions fused together in datapath

sFlow Tunnel Structures



Tunnel Structures

- sFlow Tunnel Structures extension defines metadata to describe L2, L3, L4 tunnel encap / decap actions and tunnel attributes
- sFlow Version 5 already included MPLS and VLAN stack encapsulations

OVS user space sFlow agent extended to parse captured actions, populate sFlow tunnel structures, and export tunnel structures with sampled packet header

Packet Header

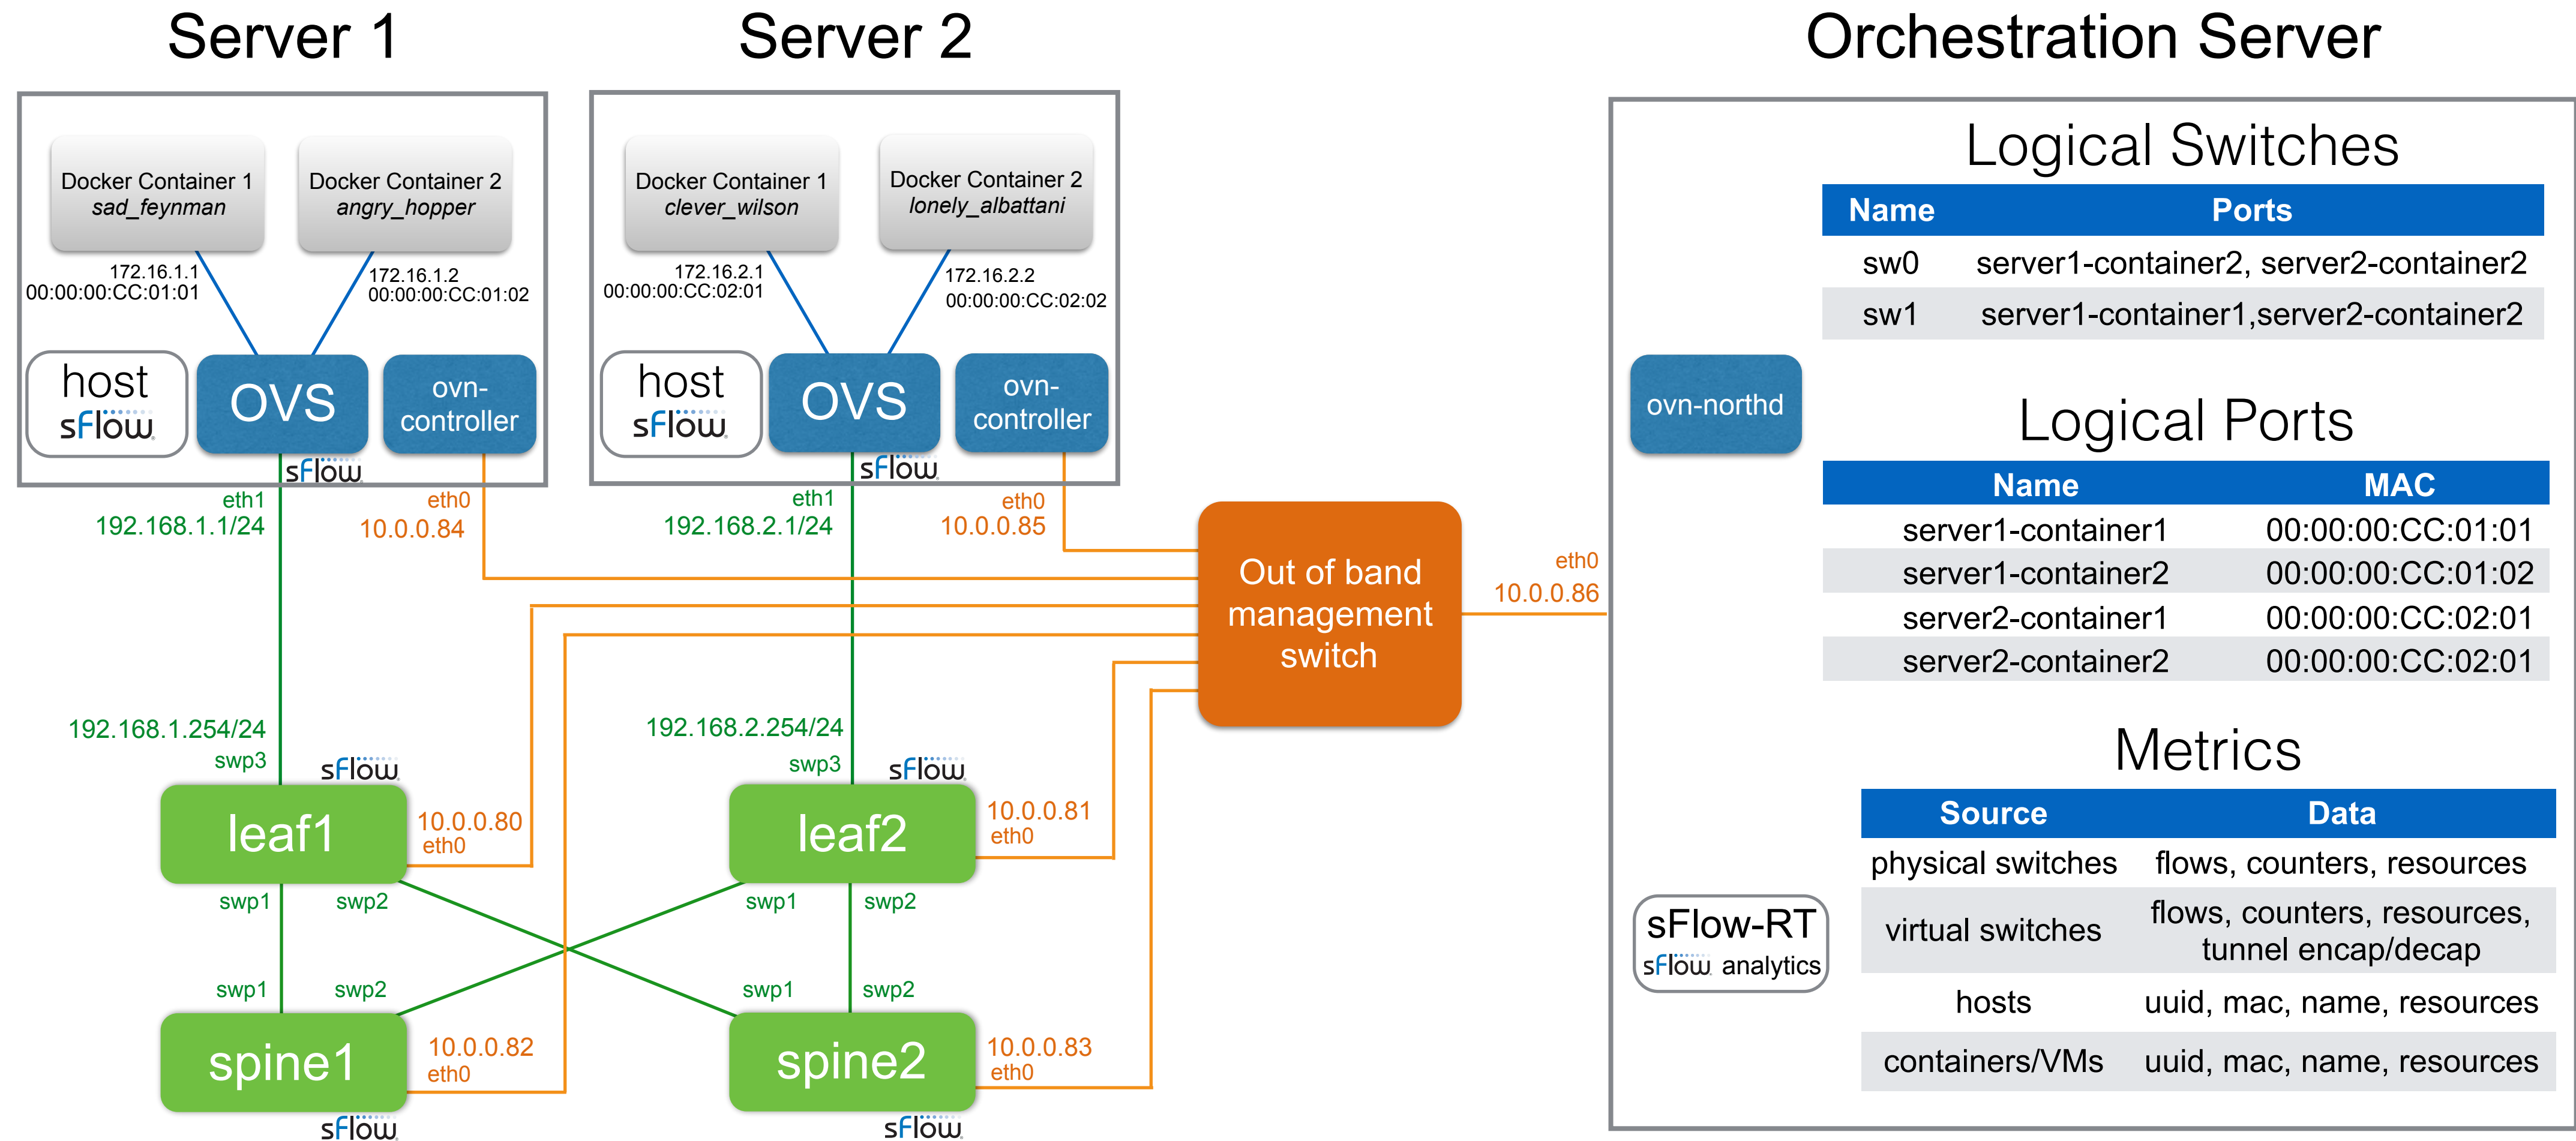
```
flowSampleType HEADER
headerProtocol 1
sampledPacketSize 1070
strippedBytes 4
headerLen 128
headerBytes 00-00-00-CC-02-02-00-00-00-CC-01-02-08-00-45-00-04-1C-15-6F-40-00-40-11-C6-3D-AC-10-01-02-
AC-10-02-02-BF-5B-13-89-04-08-5F-CD-00-00-08-48-56-35-49-7A-00-08-5C-89-00-00-00-00-00-00-00-01-00-00-
13-89-00-00-04-00-00-98-96-80-FF-FF-F8-30-36-37-38-39-30-31-32-33-34-35-36-37-38-39-30-31-32-33-34-35-
36-37-38-39-30-31-32-33-34-35-36-37-38-39-30-31-32-33-34-35
```

+

Decapsulated Input Tunnel

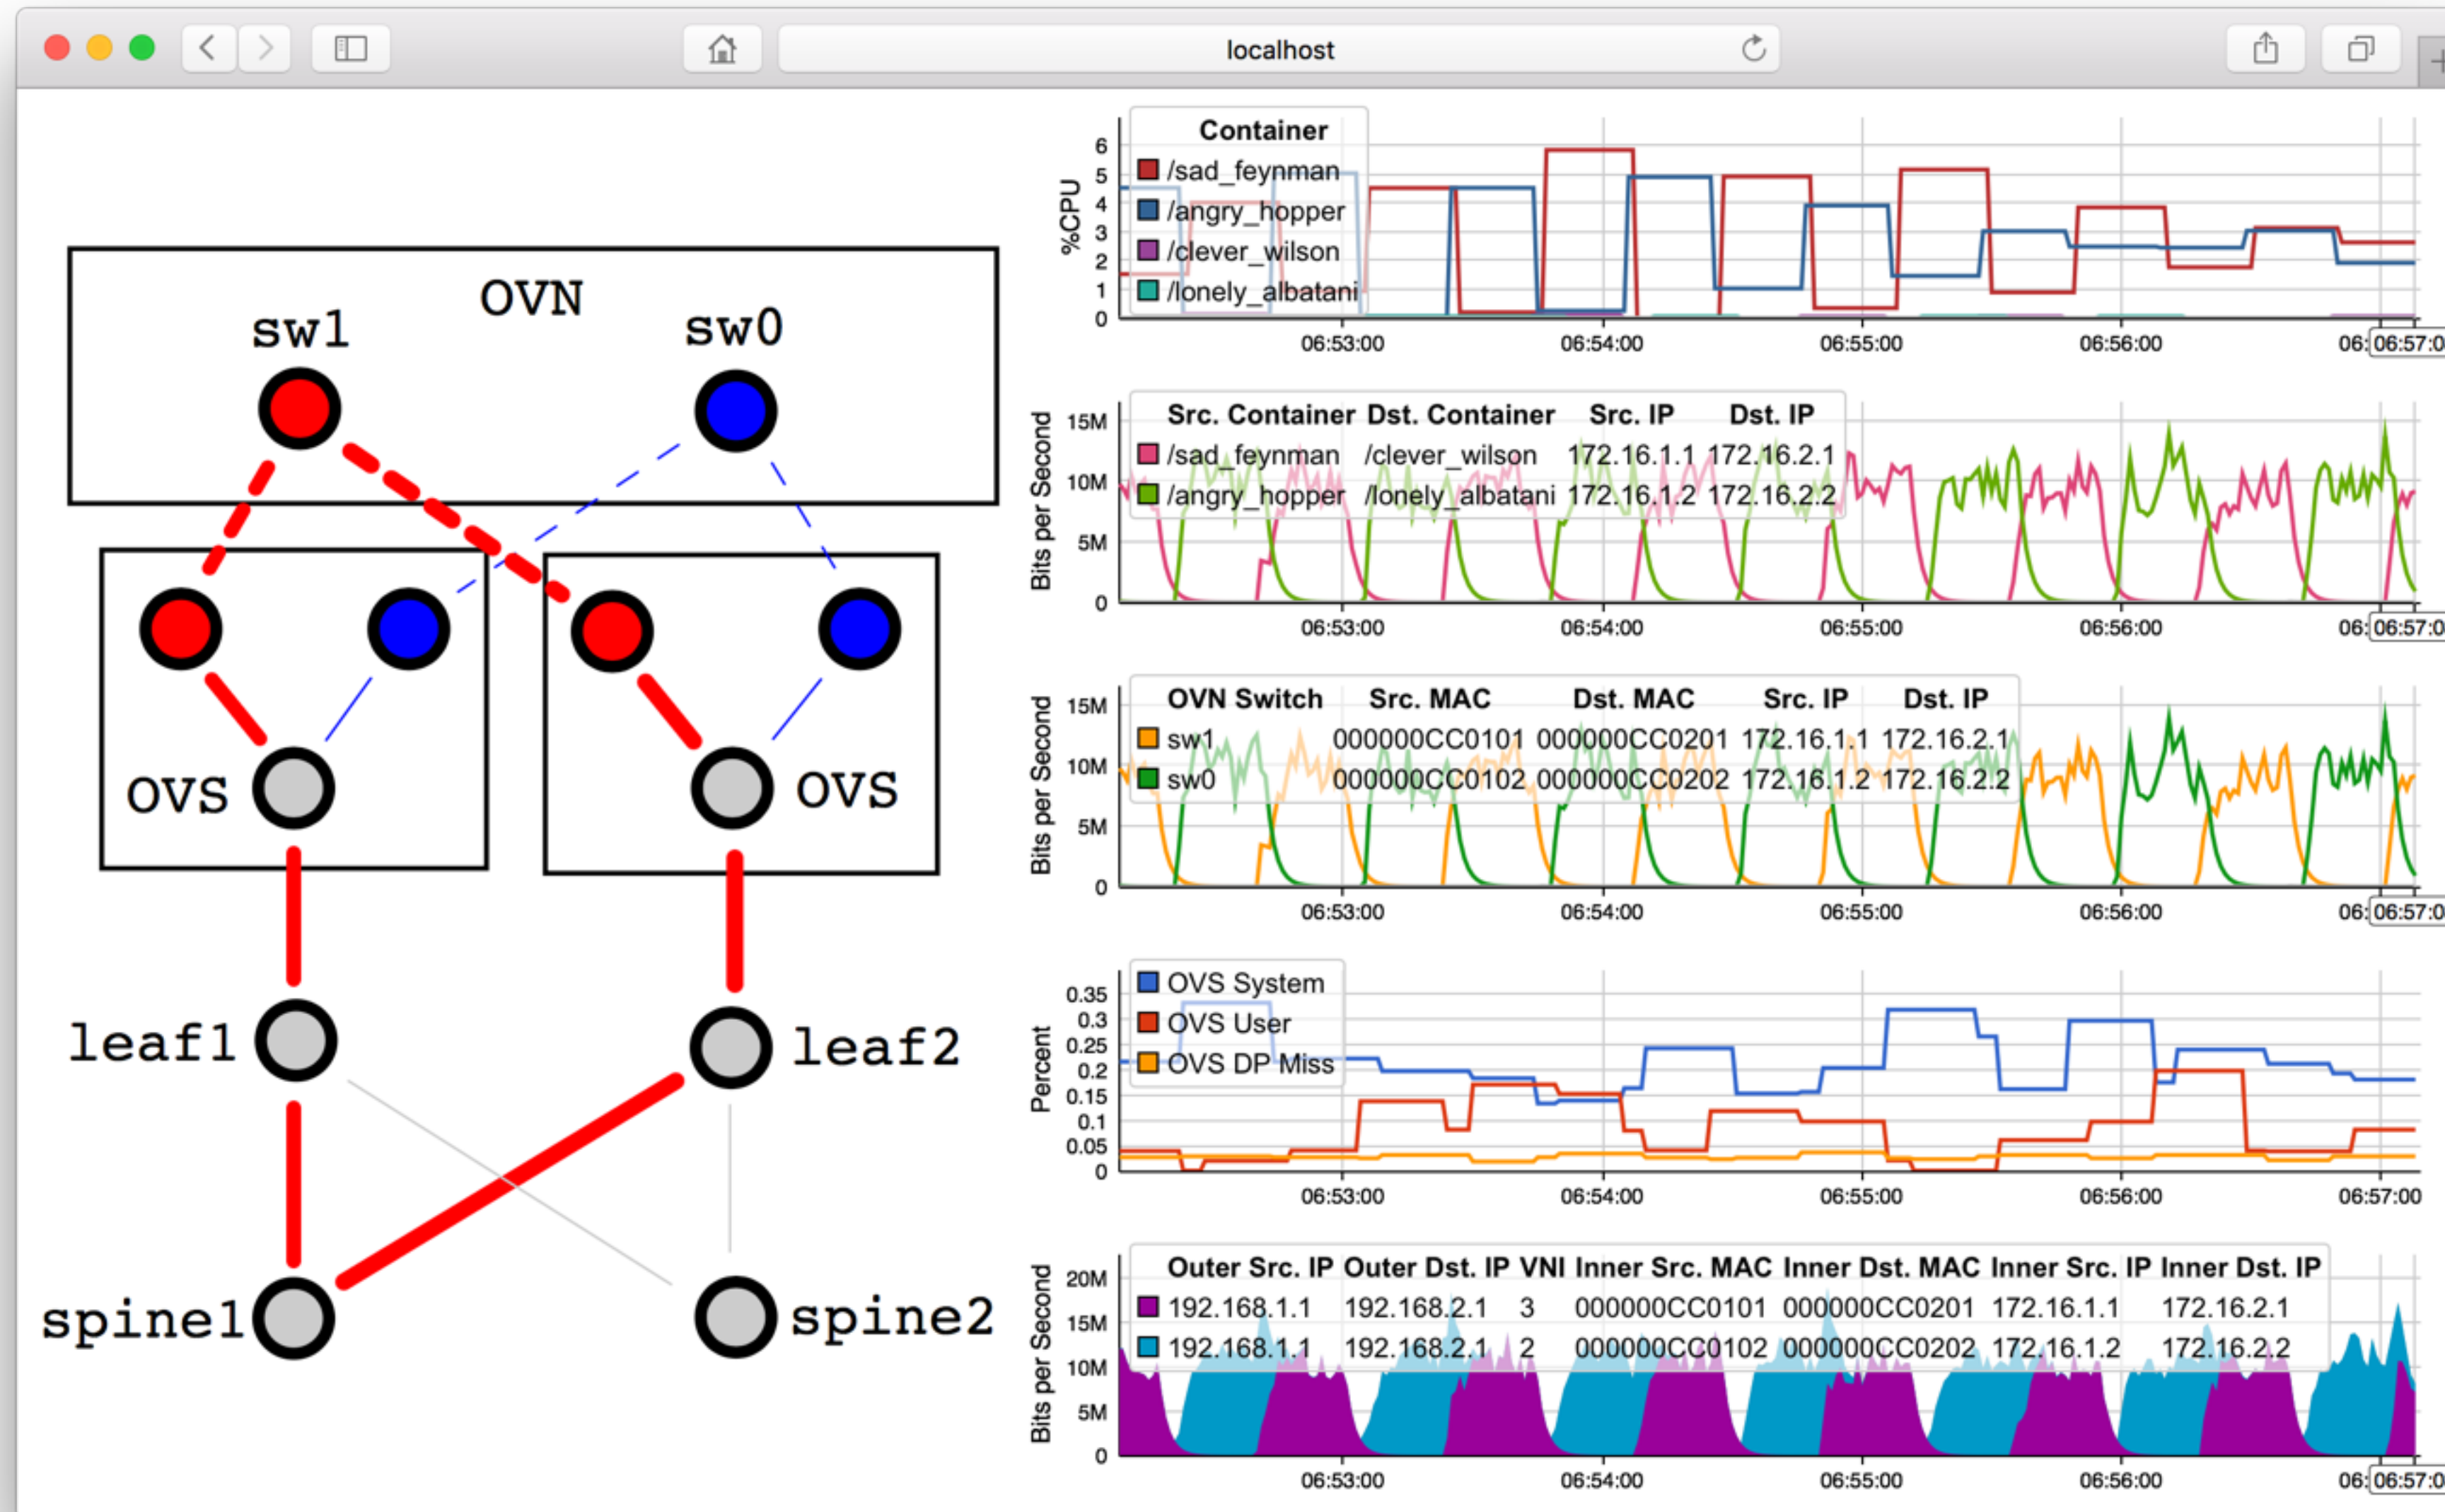
```
flowBlock_tag 0:1030
extendedType in_VNI
in_VNI 2
flowBlock_tag 0:1024
flowSampleType tunnel_ipv4_in_IPV4
tunnel_ipv4_in_sampledPacketSize 0
tunnel_ipv4_in_IPSize 0
tunnel_ipv4_in_srcIP 192.168.1.1
tunnel_ipv4_in_dstIP 192.168.2.1
tunnel_ipv4_in_IPProtocol 17
tunnel_ipv4_in_IPTOS 0
tunnel_ipv4_in_UDPsrcPort 54024
tunnel_ipv4_in_UDPdstPort 49431
```

Demo: VirtualBox, Cumulux VX, OVN, Docker






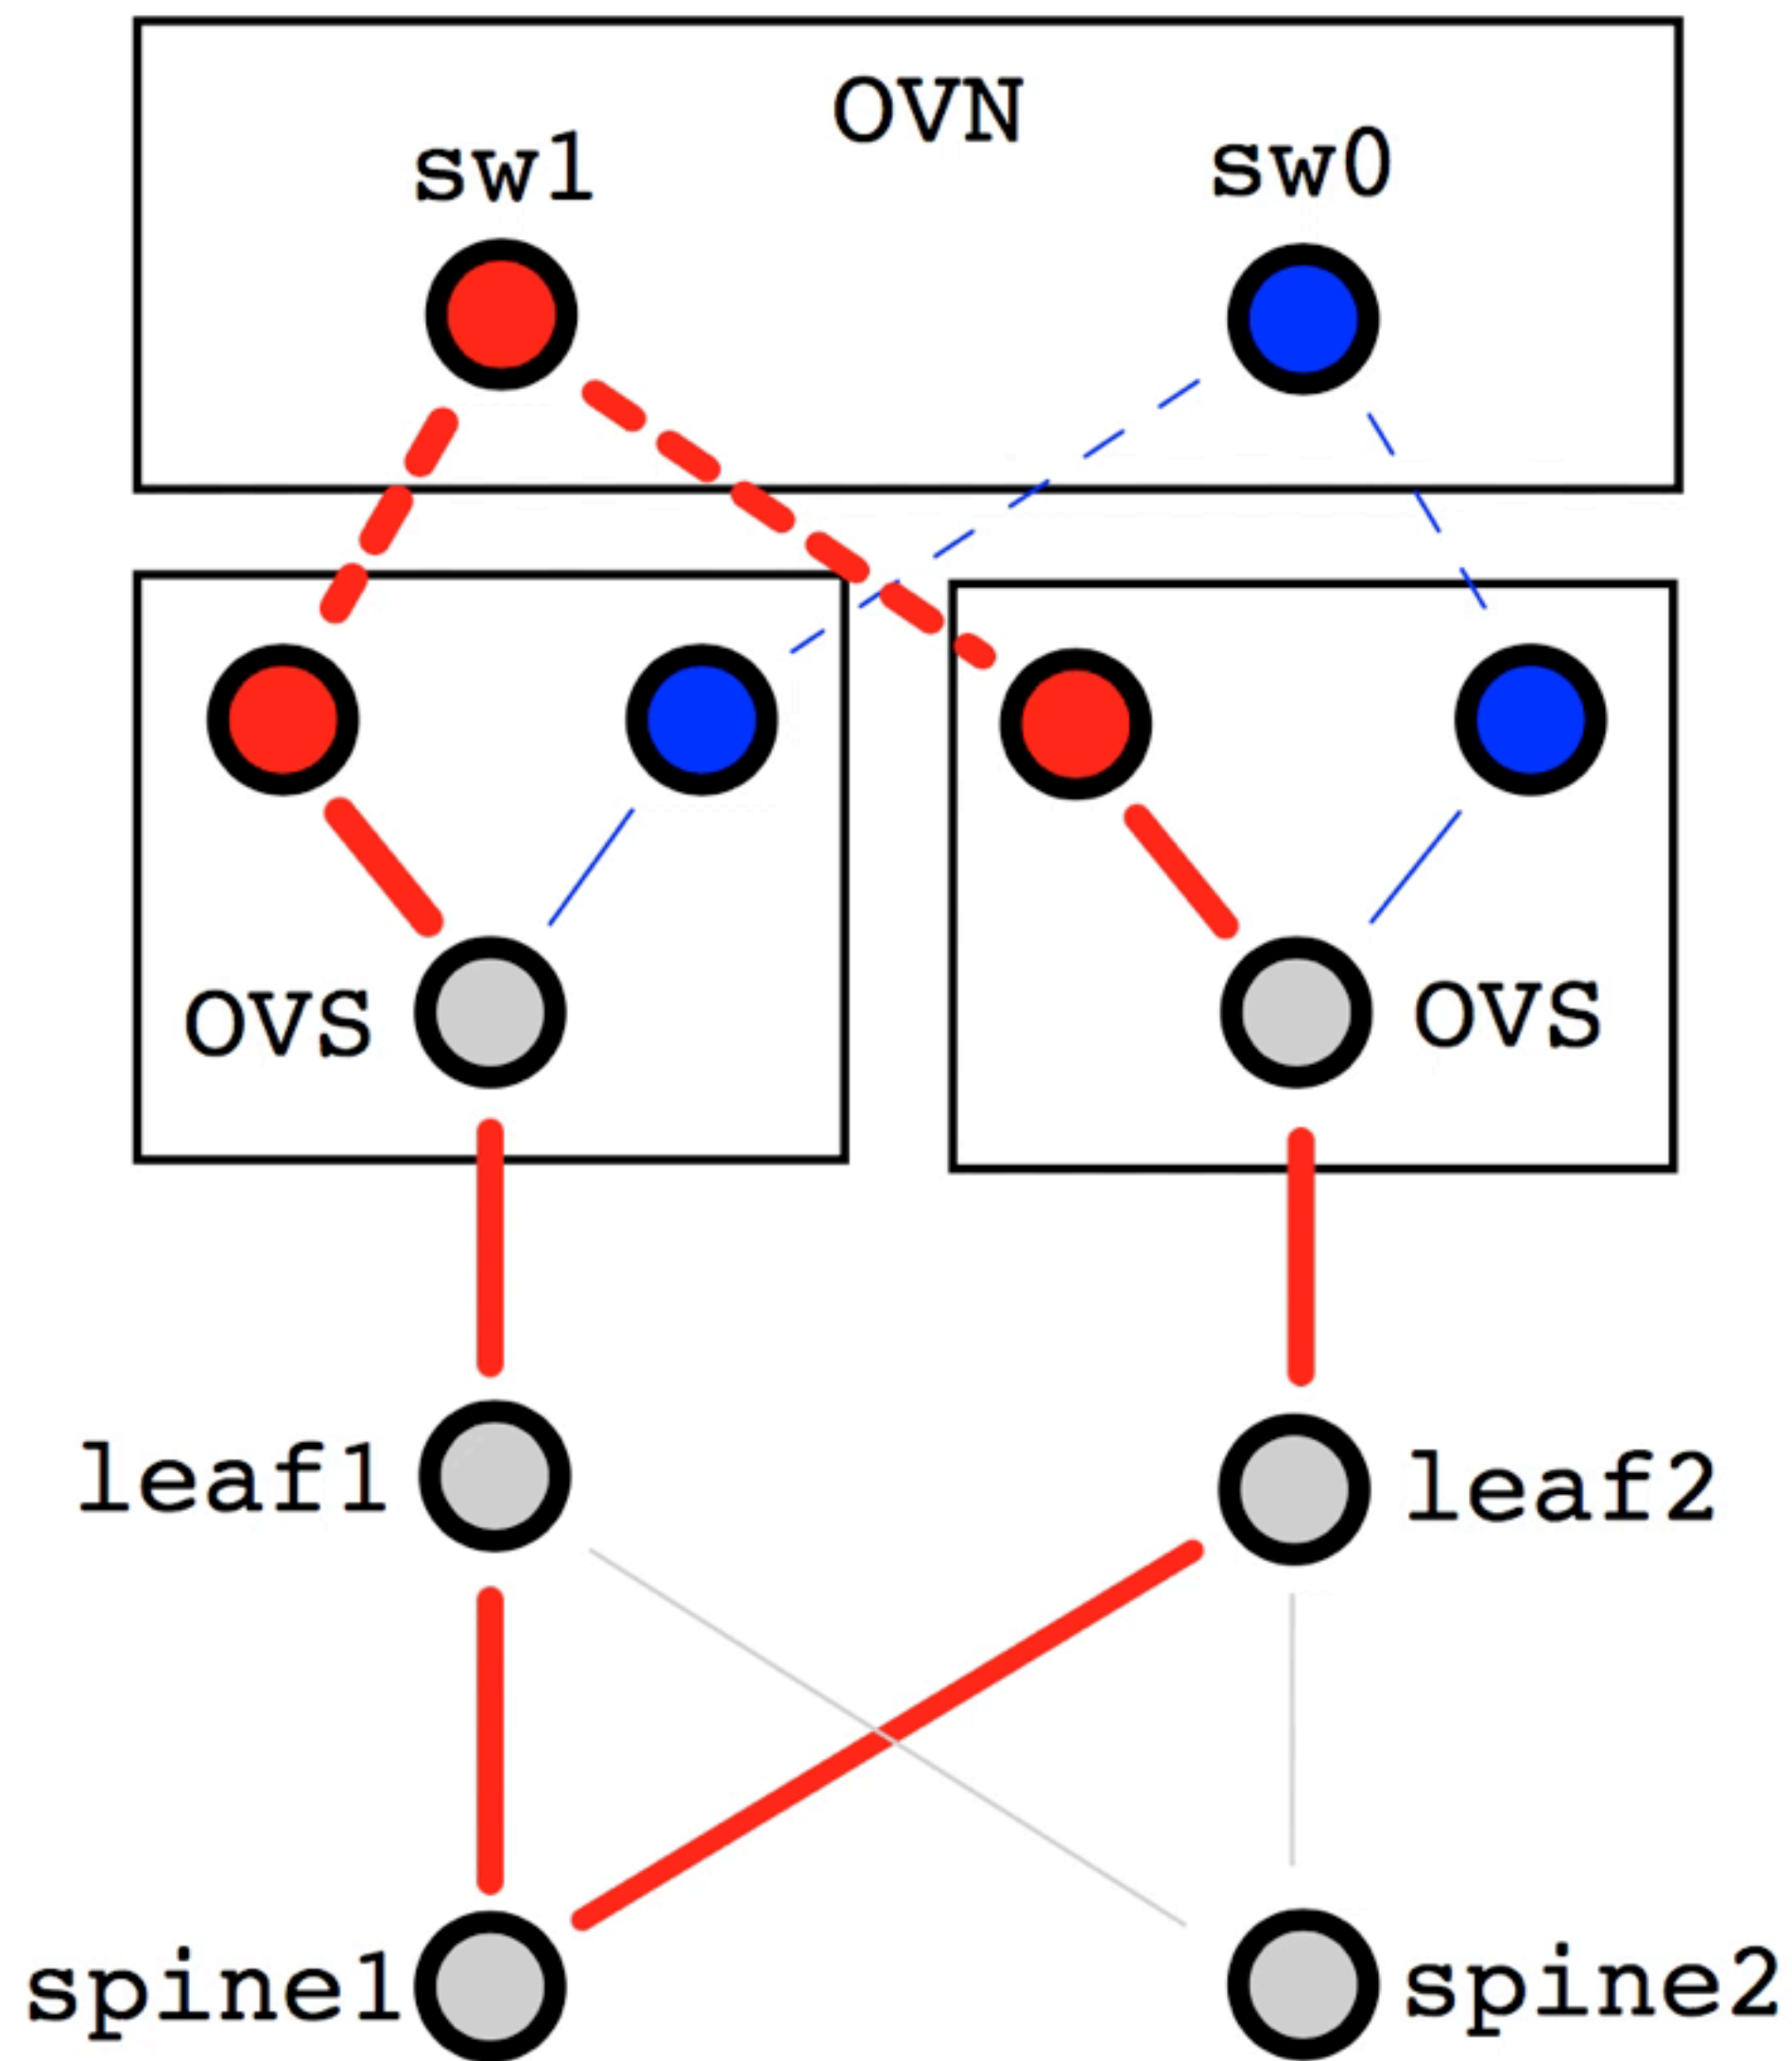


OVS part of ecosystem: How is sFlow combined to provide comprehensive visibility?

Live Demo



<https://github.com/pphaal/ovs-2015>
(includes PCAP file to replay data from test network)

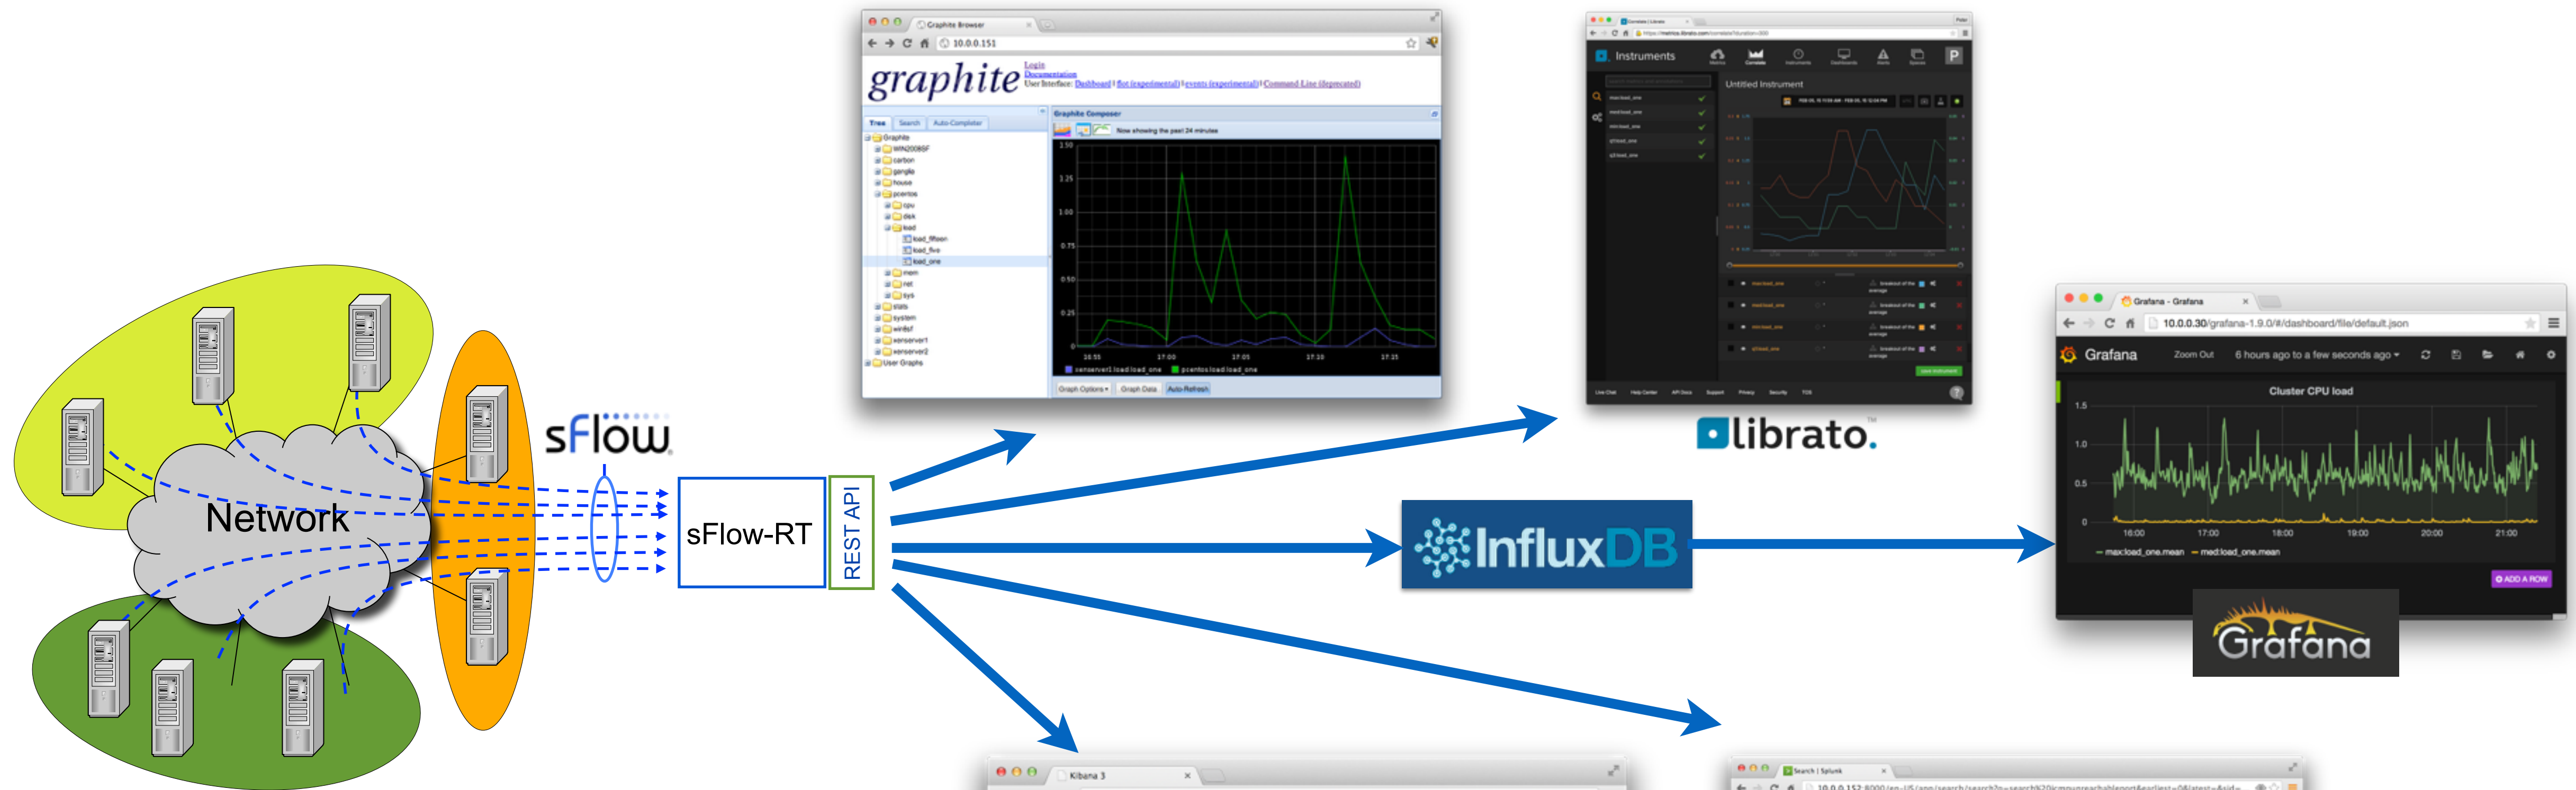




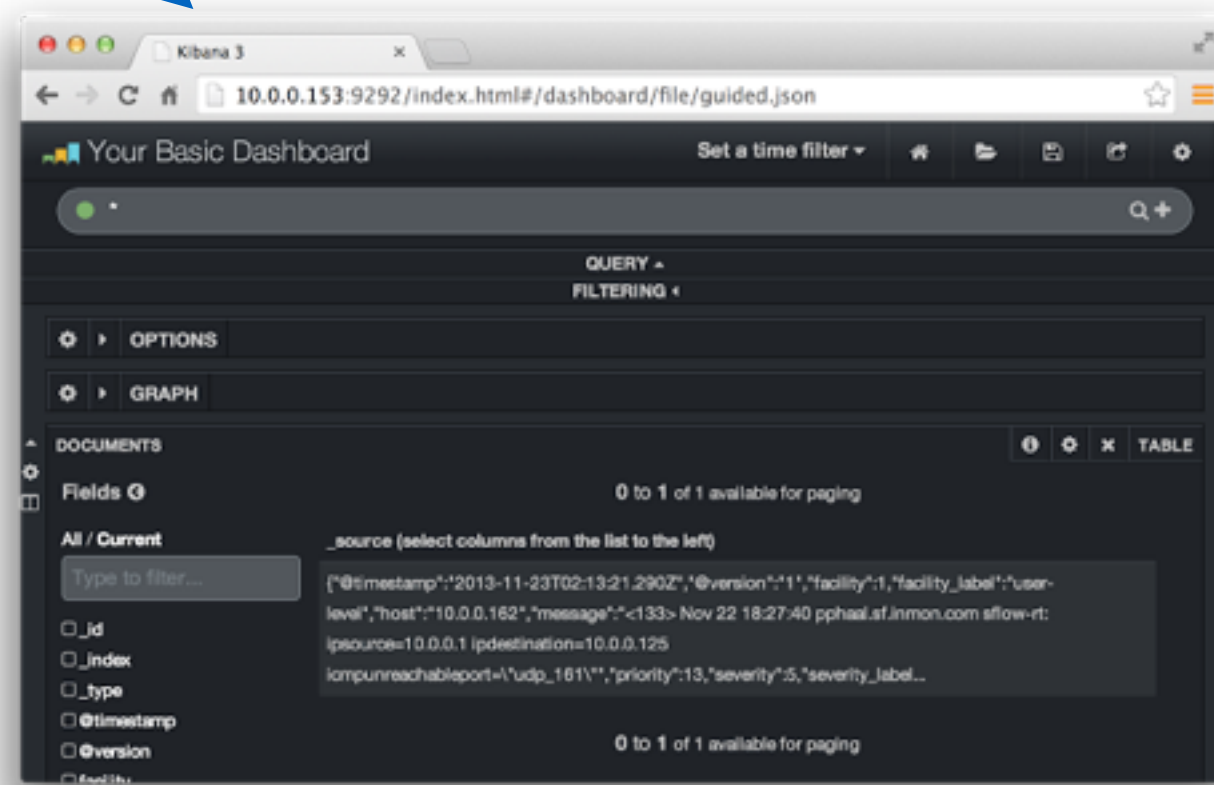
- SC15: The International Conference for High Performance Computing, Networking, Storage and Analytics, Austin, Nov 15th - 19th
- “SCinet brings to life a very high-capacity network that supports the revolutionary applications and experiments that are a hallmark of the SC conference”
- Live sFlow driven “weathermap”

<http://inmon.sc15.org/sflow-rt/app/sc15-weather/html/>

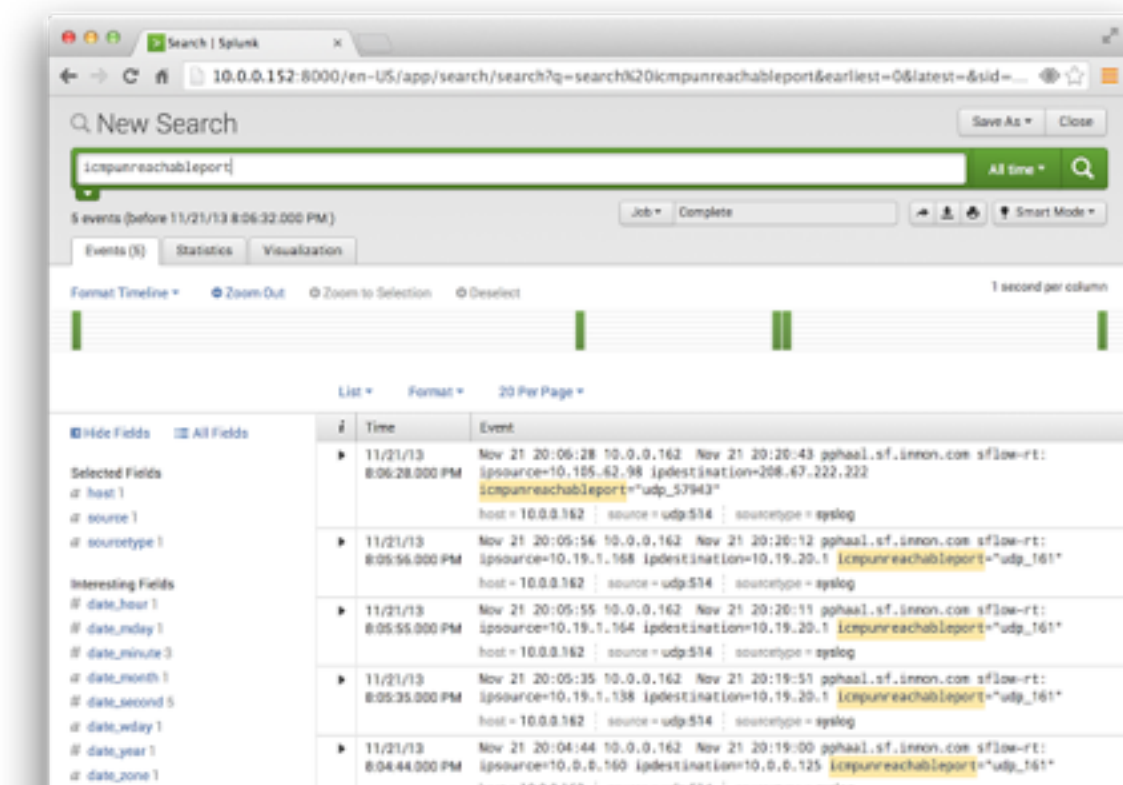
Visibility for DevOps tools



- Streaming filtering and summarization reduces data volume and increases scalability of backend tools
- Streaming flow analytics to generate application metrics

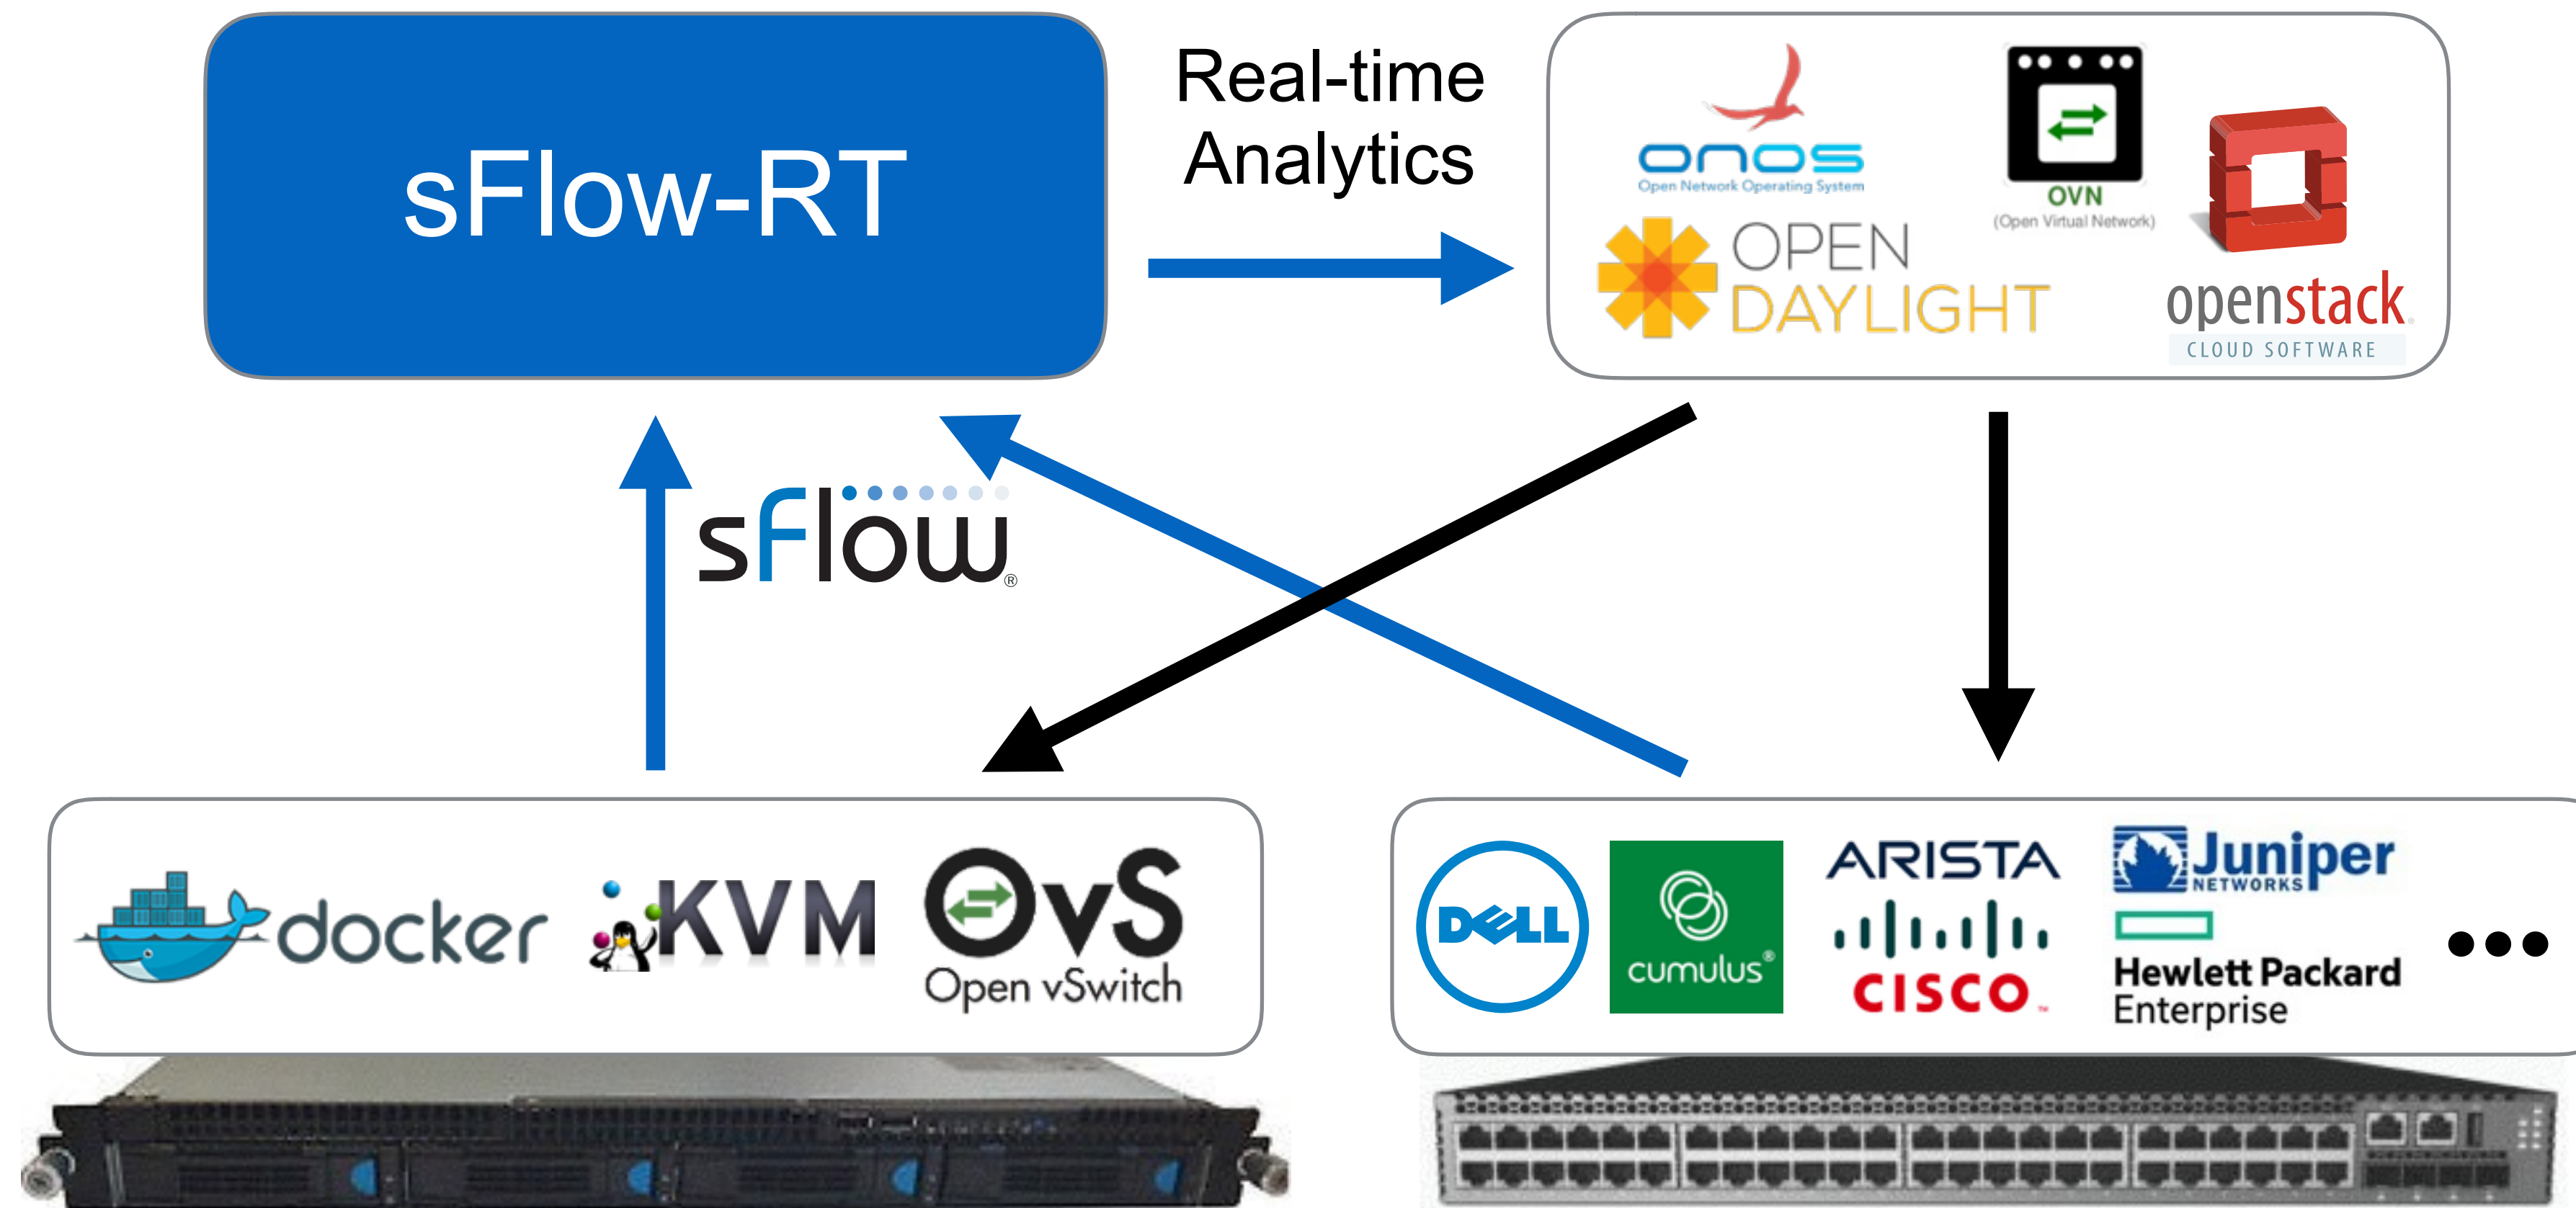


Logstash



splunk>

Feedback Control in Cloud Stacks

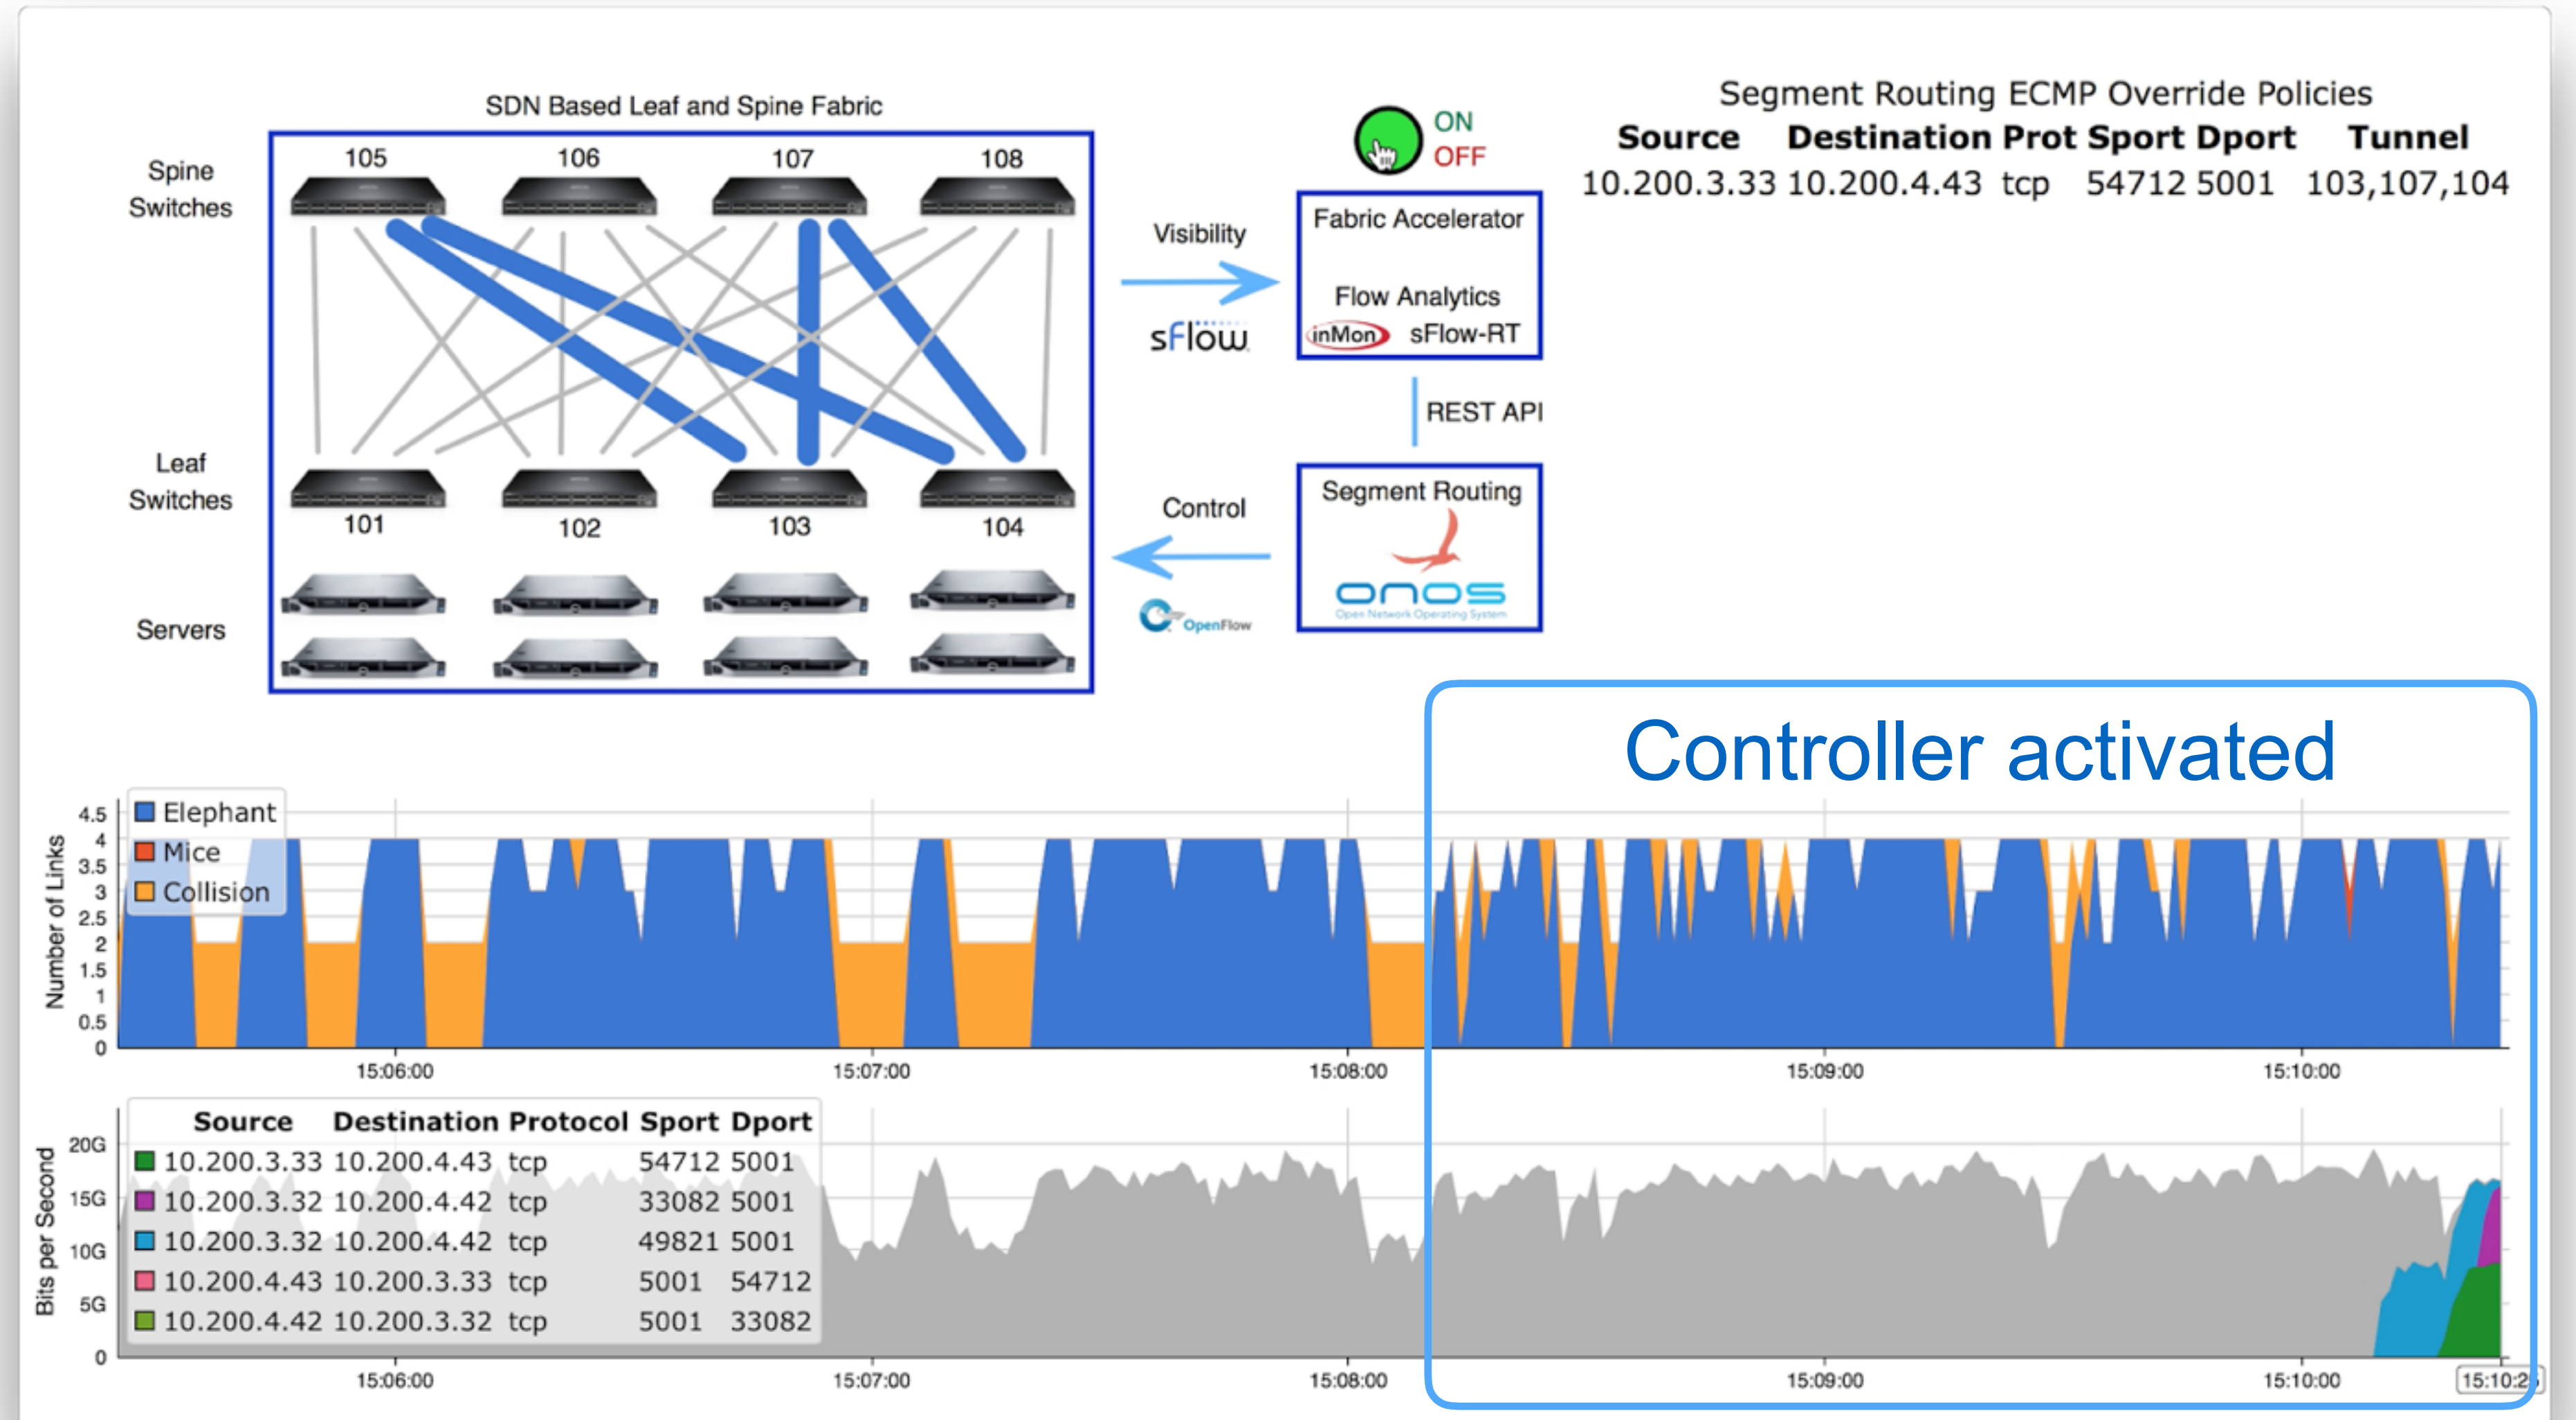


“Typically the resource that is most scarce is the network.”

Amin Vahdat, ONS2015 Keynote

<http://blog.sflow.com/2015/06/optimizing-software-defined-data-center.html>

ONS 2015: Leaf and Spine Traffic Engineering Using Segment Routing and SDN



<http://blog.sflow.com/2015/06/leaf-and-spine-traffic-engineering.html>
(Emulated in Mininet before deploying on physical network)

With visibility into underlay, OVS can push MPLS labels and load balance traffic

Additional use cases

- Bandwidth on demand: Dynamically provision optical circuits
<http://blog.sflow.com/2014/09/sdn-control-of-hybrid-packet-optical.html>
- Reduce latency: QoS marking of Elephant flows at network edge
<http://blog.sflow.com/2015/01/hybrid-openflow-ecmp-testbed.html>
- Increase availability: Detect and filter DDoS flood attacks
<http://blog.sflow.com/2014/03/ons2014-sdn-idol-finalist-demonstrations.html>
- Increase available bandwidth and reduce latency: Co-locate heavily communicating VMs / containers
- Increase available CPU: Auto scaling service pools
- Increase available bandwidth and reduce latency: Optimize network paths for NFV / service chaining

Final Comments

- Open vSwitch is part of a comprehensive cloud infrastructure ecosystem that embeds standard sFlow instrumentation
- Combining sFlow measurements from physical and virtual infrastructure provides end-to-end visibility into resources needed to deliver virtual network services.
- Measurement based feedback is key to optimizing cloud infrastructure

“If you can not measure it, you can not improve it.” Lord Kelvin

Questions?

Detailed OVN/OVS Packet Path

