

# sFlow®

Data network visibility and control



“You can’t control what you can’t measure”  
*Tom DeMarco*



# Reason 1: Widely supported industry standard



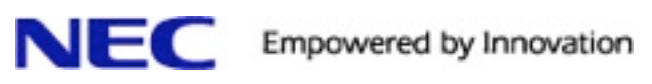
# Reason 1: Widely supported industry standard



BROCADE



ARISTA



AlaxaIA



Alcatel-Lucent





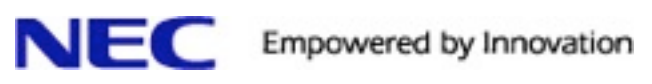
# Reason 1: Widely supported industry standard



BROCADE



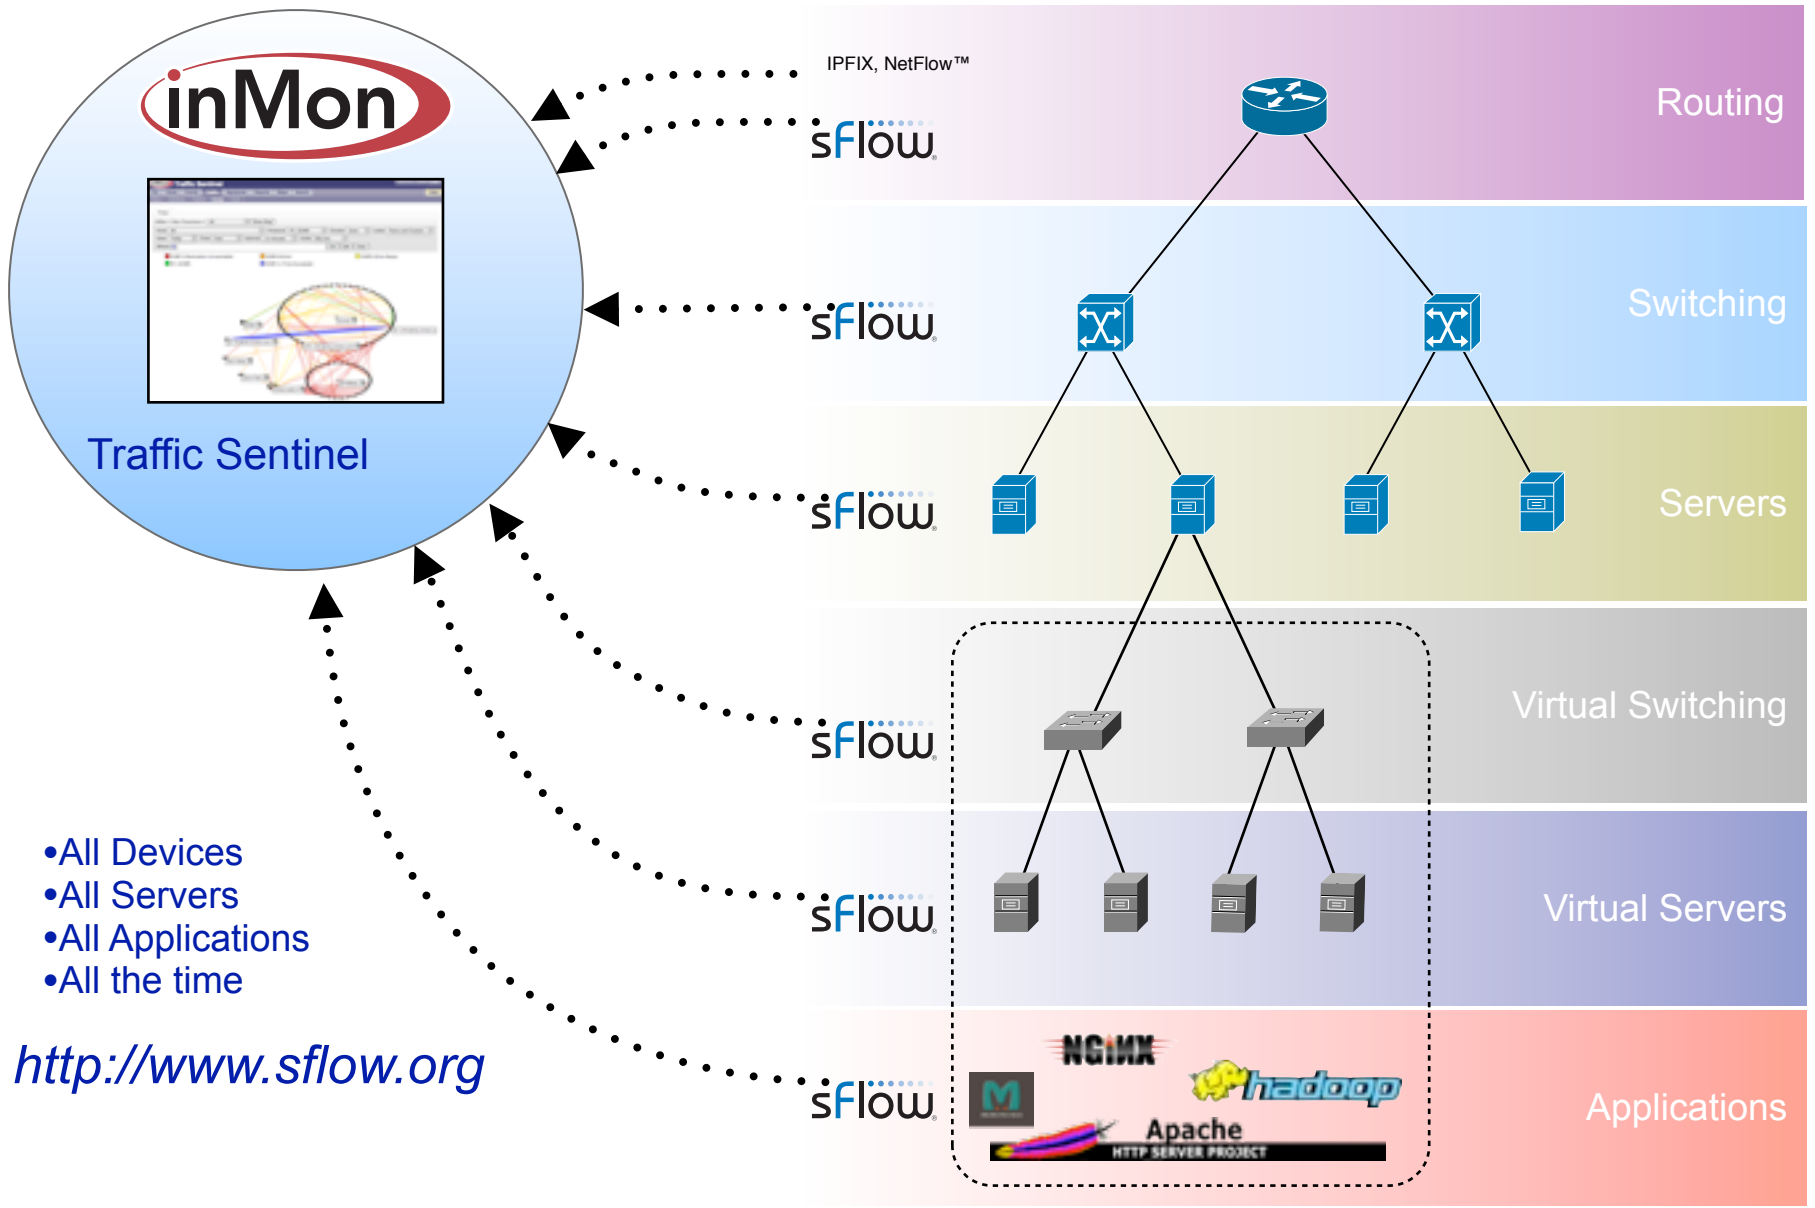
ARISTA



AlaxaIA



# Reason 2: Comprehensive



- All Devices
- All Servers
- All Applications
- All the time

<http://www.sflow.org>

# sFlow Overview: replaces counter polling

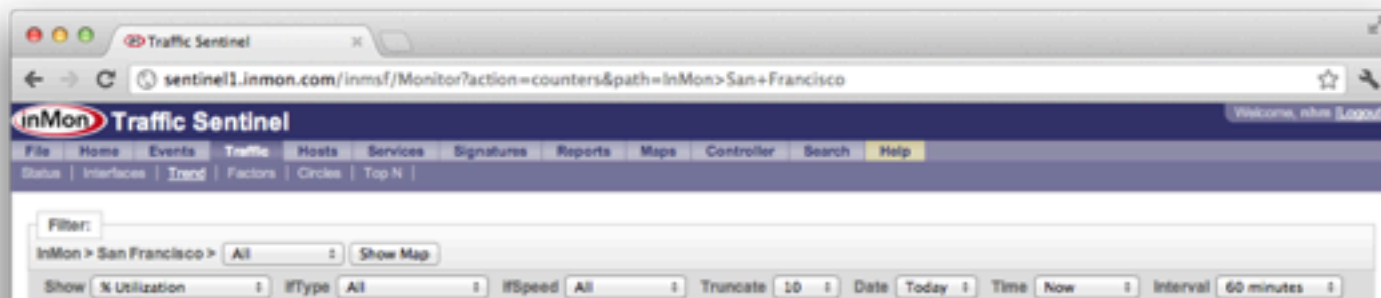
## *“De-synchronized, Parallel Push”*

- sFlow agent automatically pushes full set of SNMP ifTable counters<sup>1</sup>
  - Compared to SNMP polling, counter push results in 10-20x fewer packets on network, reduces CPU load on switch and on network management software (XDR is easier to encode/decode than SNMP)
  - Single sFlow collector can easily monitor 200,000 switch ports with 1 minute granularity. SNMP polling with 5 minute granularity requires 5-10 collectors.
1. ifIndex, ifType, ifSpeed, ifDirection, ifAdminStatus, ifOperStatus, ifInOctets, ifInUcastPkts, ifInMulticastPkts, ifInBroadcastPkts, ifInDiscards, ifInErrors, ifInUnknownProtos, ifOutOctets, ifOutUcastPkts, ifOutMulticastPkts, ifOutBroadcastPkts, ifOutDiscards, ifOutErrors, ifPromiscuousMode



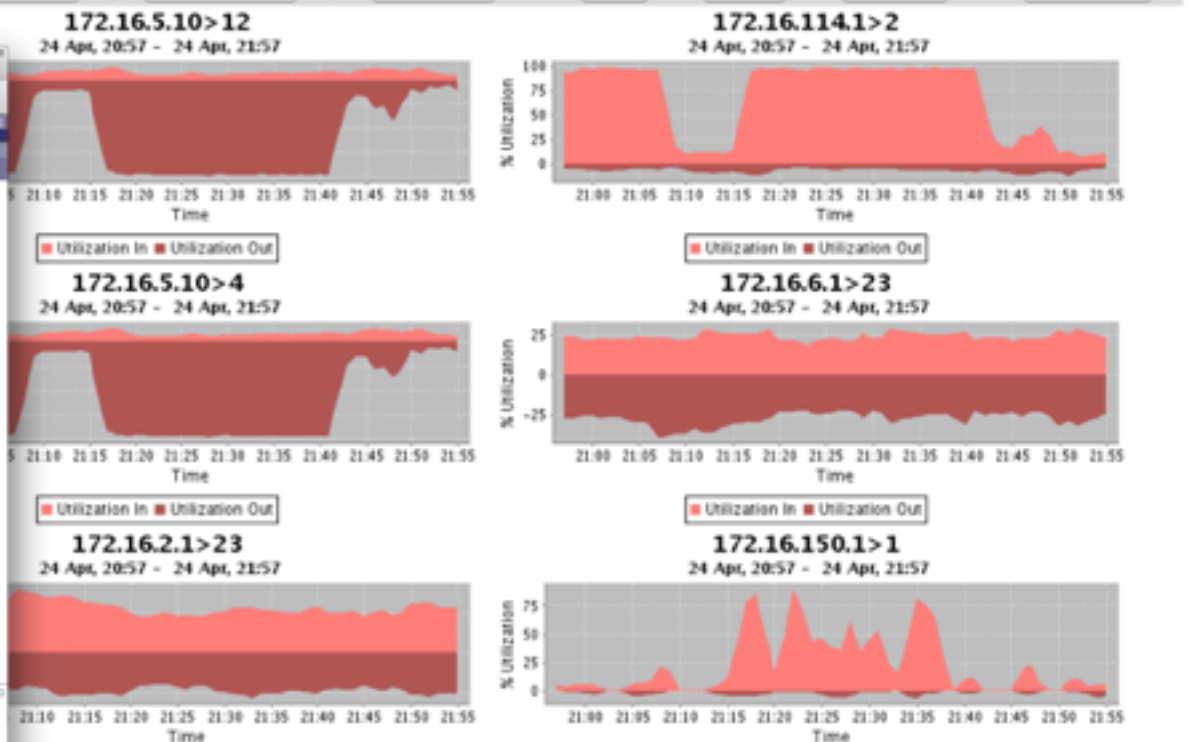
# Traffic Sentinel: Interface counters

- 200,000+ ports
- 1-minute granularity
- Thresholds/alerts
- Compare all interfaces

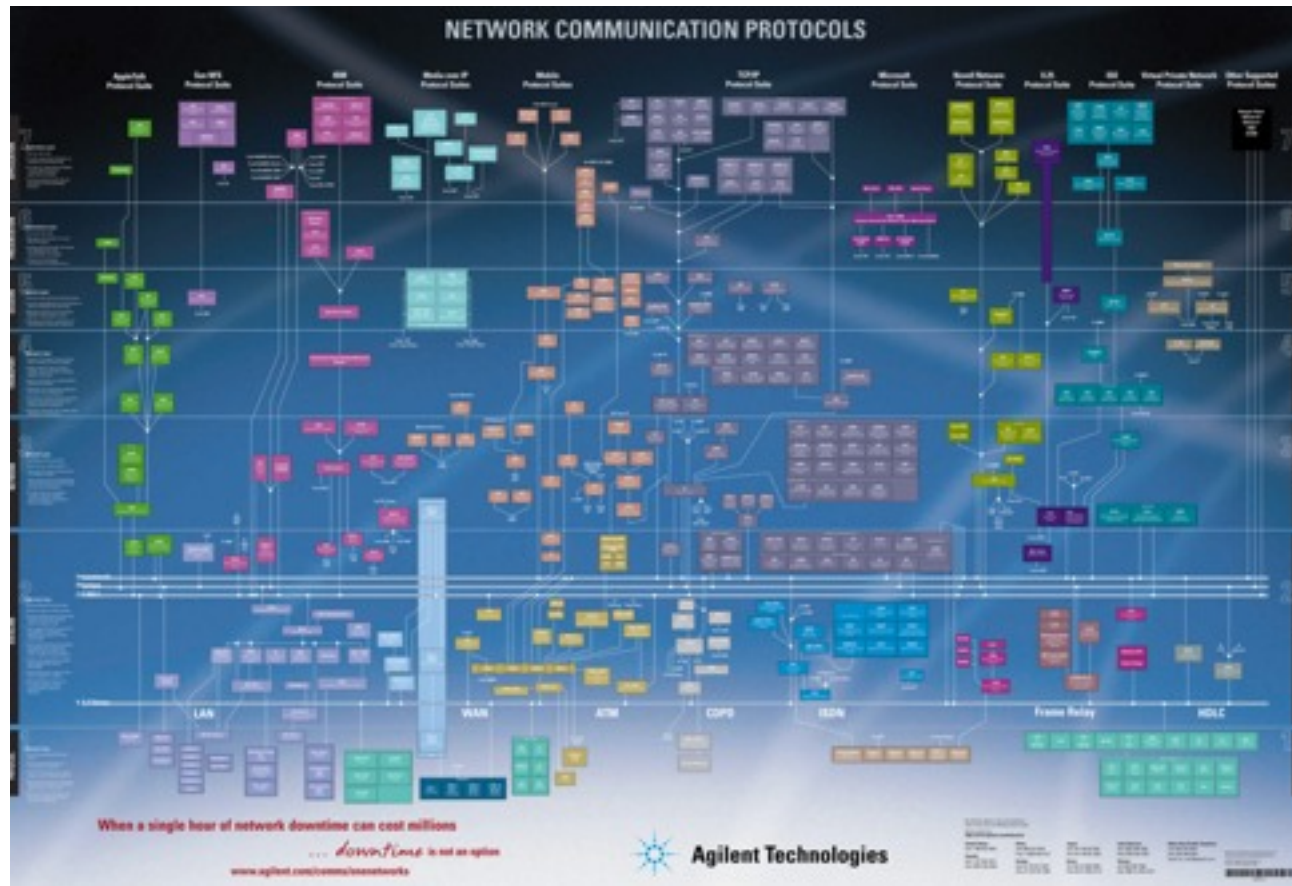


The screenshot shows a table of interface counters. The table has columns for 'Frames/sec', '% Utilization', 'Broadcasts/sec', 'Multicasts/sec', 'Errors/sec', 'Discards/sec', 'Agent', 'Interface', and 'IFSpeed'. The data is presented in a grid with green and red bars indicating utilization levels.

Frames/sec	% Utilization	Broadcasts/sec	Multicasts/sec	Errors/sec	Discards/sec	Agent	Interface	IFSpeed
45.04 K	26.16	231.3	42.07	0	0	172.16.5.1	22	1000/sec
43.06 K	25.21	236.3	46.37	0	12.6	172.16.6.1	23	1000/sec
18.74 K	10.86	79.75	36.73	0	0	172.16.6.1	25	1000/sec
18.36 K	10.44	753.3 m	5.9	0	0	172.16.96.200	23	1000/sec
18.3 K	10.35	80.13	36.9	0	0	172.16.96.200	1	1000/sec
16.03 K	14.94	700 m	3.217	0	0	172.16.96.12	6	1000/sec
13.25 K	13.28	103.3 m	3.25	0	0	172.16.96.12	6	1000/sec
11.208 K	11.52	523.3 m	016.7 m	0	0	172.16.96.223	16	10000/sec
30.76 K	0.807	62.15	41.5	0	0	172.16.2.1	302	2000/sec
29.9 K	0.80	48.95	40.87	0	0	172.16.96.12	290	2000/sec
62.37	0.311	0	0	0	0	172.16.5.10	12	1.5400/sec
0	0.199	0	0	0	0	172.16.5.10	4	1.5400/sec
51.83	0.155	0	0	0	0	172.16.114.1	2	1.5400/sec
57.48	1.94	0	0	0	0	172.16.5.10	16	1.5400/sec
0	7.8	0	0	0	0	172.16.5.10	8	1.5400/sec
76.3	7.585	863.3 m	12.017	0	0	172.16.147.2	68	1000/sec
56.73	7.182	0	0	0	0	172.16.136.1	1	1.5400/sec
67.18	5.989	0	47.8	0	0	172.16.96.200	30202	1100/sec
103.3	5.809	0	47.8	0	0	172.16.96.200	30211	1100/sec
104	4.301	0	47.8	0	0	172.16.96.200	30214	1100/sec



# sFlow Overview: monitors all protocols



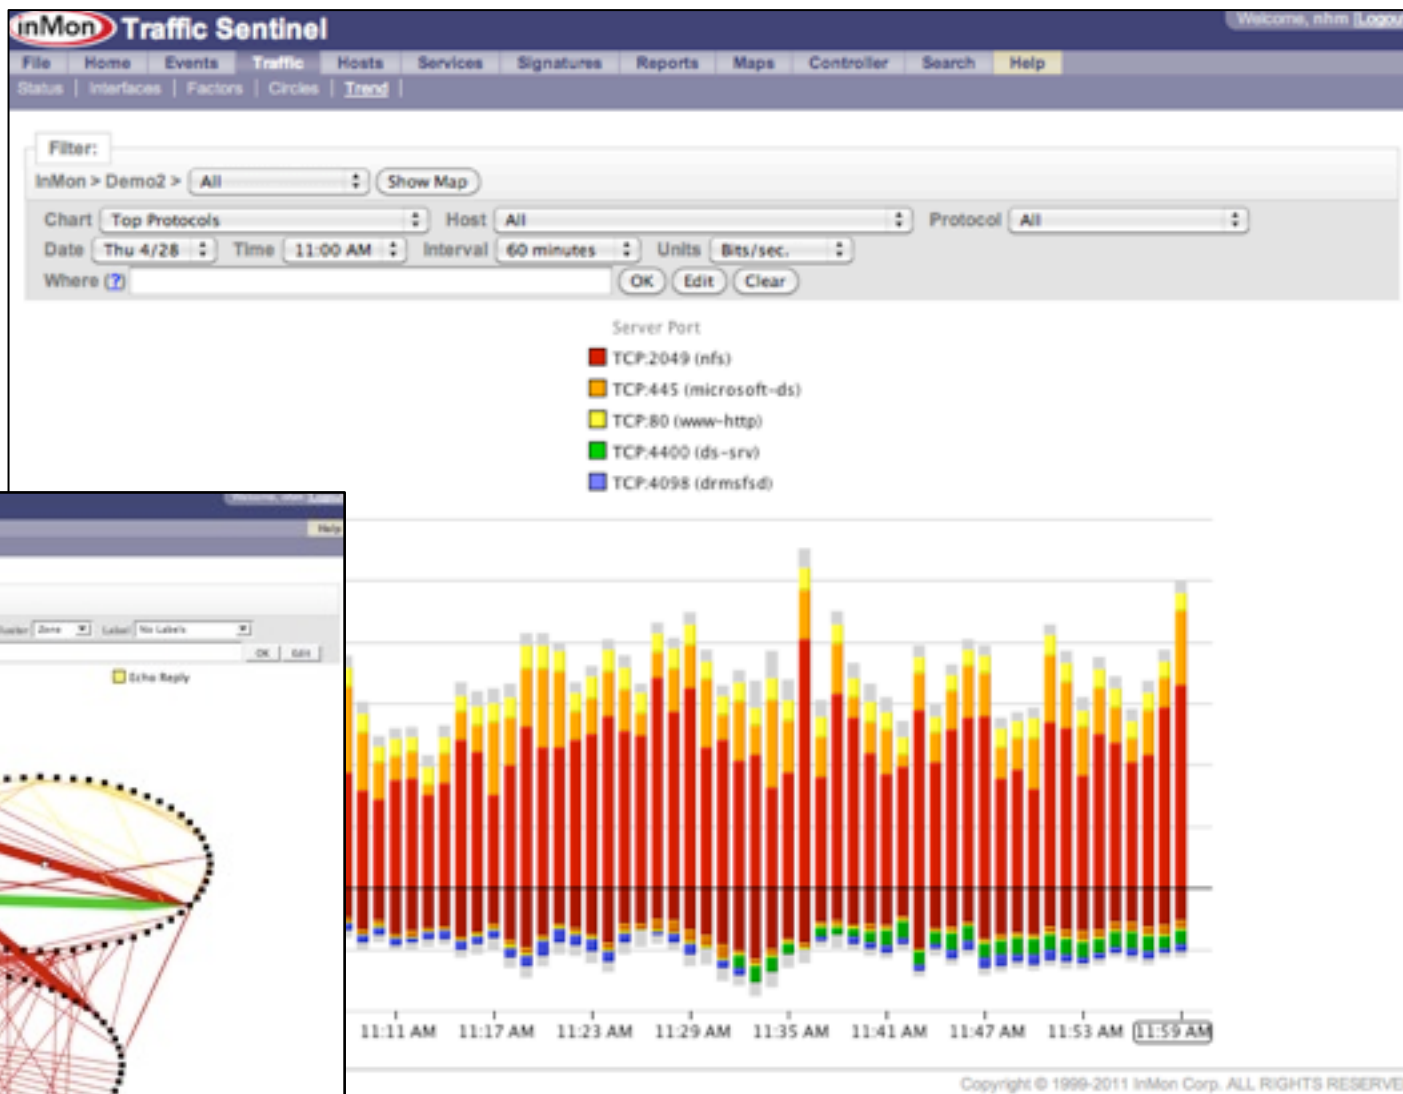
- Simple agents: packet headers sent to sFlow collector for decoding.
- Easier to add decodes to central collector than to every device in a multi-vendor network (e.g. IPv6, FCoE etc.)
- Captures complex layering (e.g. MAC/VLAN/MPLS/IPv4/IPv6): critical for tracing packet paths through network.





# Traffic Sentinel: Traffic Breakdown

- MAC, VLAN, IP, IPv6, TCP, UDP, MPLS, TRILL, RTP etc. (over 100 fields)
- 1-minute granularity
- Thresholds/alerts
- Automatic de-duplication
- Subnet rollups



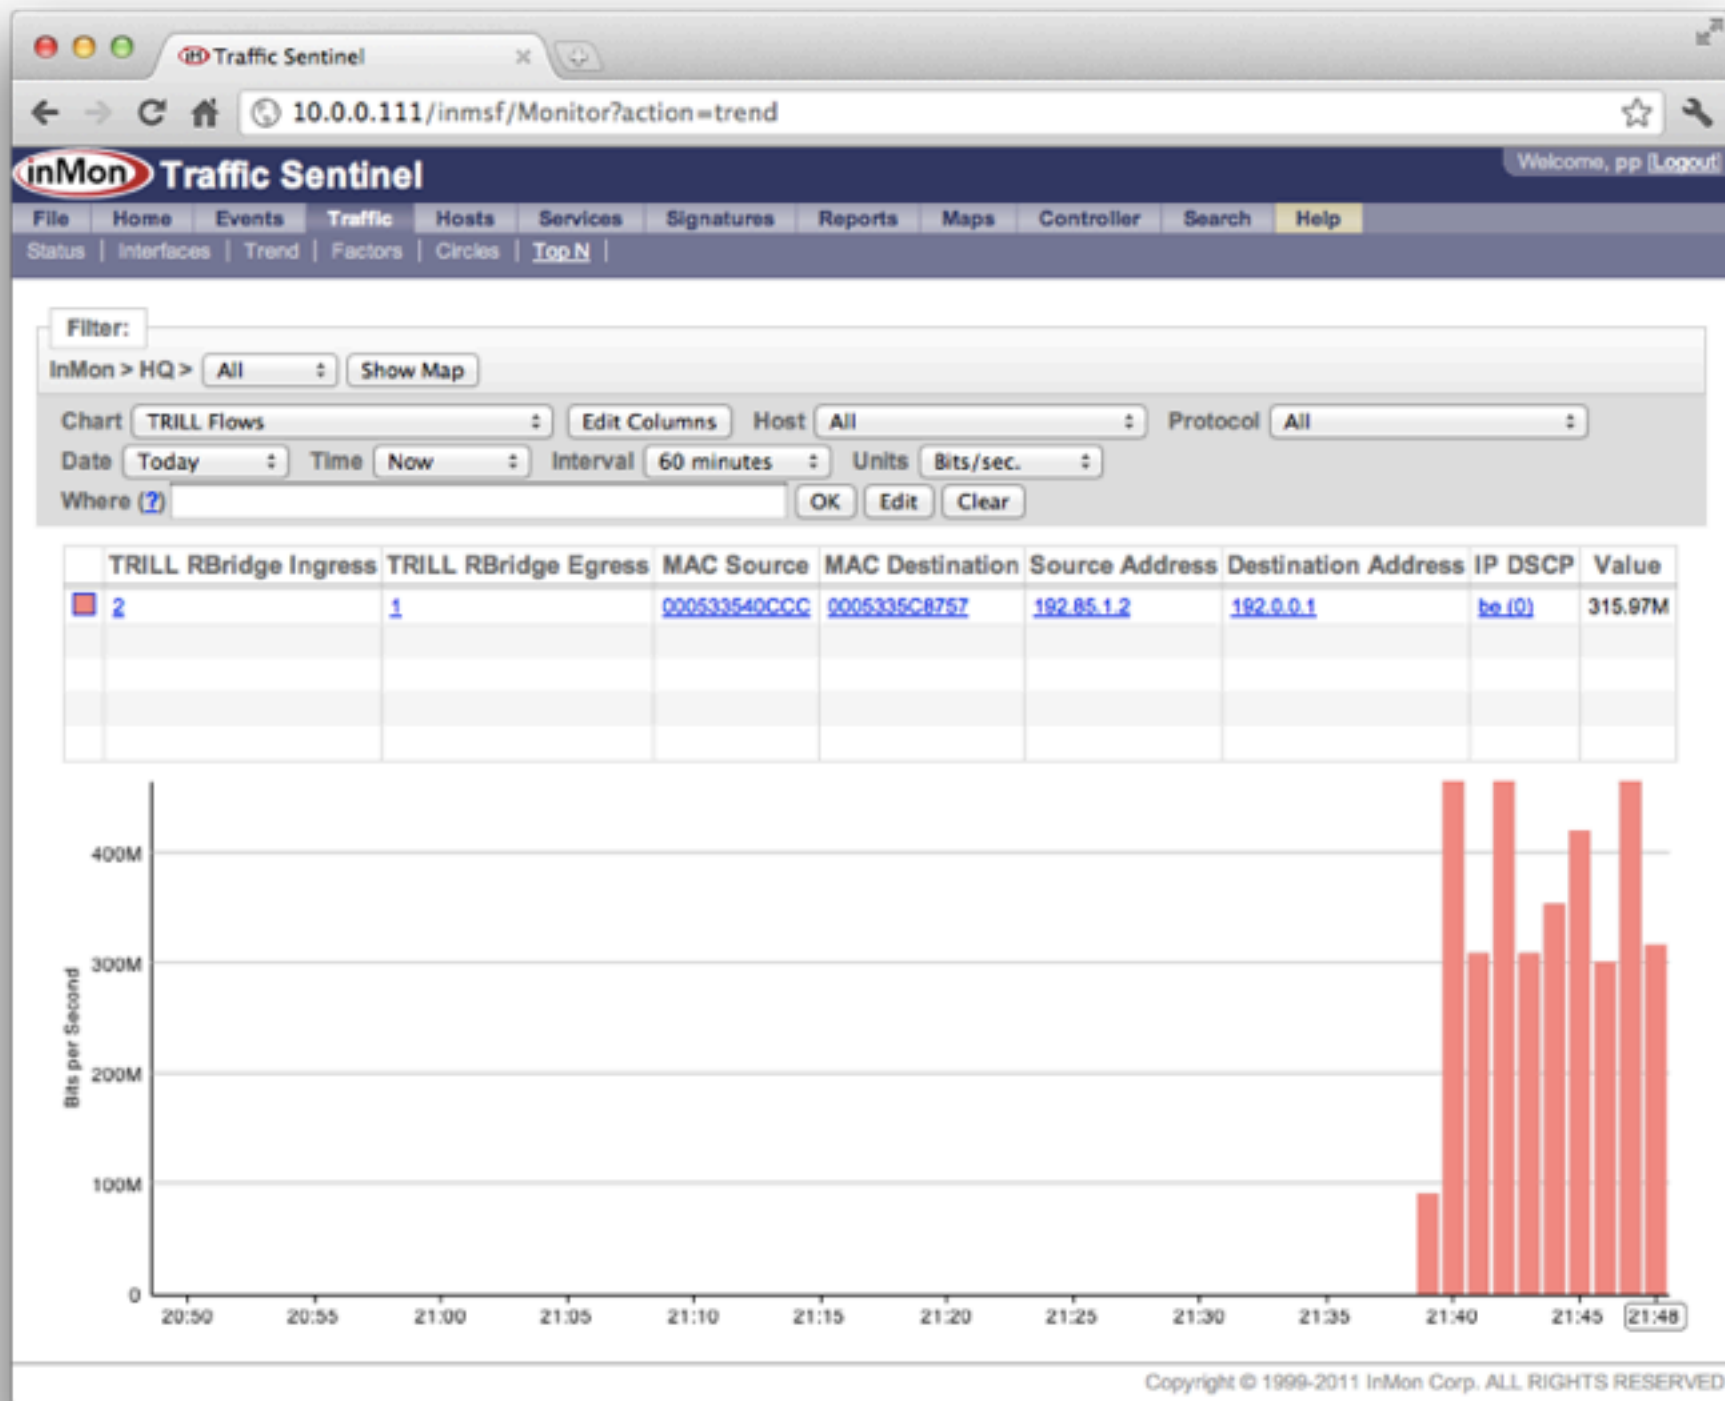
# TRILL Fabrics

```
+++++
|                               Outer Destination MAC Address                               |
| Outer Destination MAC Address | Outer Source MAC Address                             |
|                               Outer Source MAC Address                               |
| Ethertype = IEEE 802.1Q      | UP |C|      Outer VID                             | |
| Ethertype = TRILL           | V | Hop Limit |M| Reserved                         |
| Egress RBridge Address      | Ingress RBridge Address                             |
|                               Inner Destination MAC Address                          |
| Inner Destination MAC Address | InnerSource MAC Address                             |
|                               Inner Source MAC Address                              |
| Ethertype = IEEE 802.1Q      | UP |C|      Inner VID                             |
|                               Original Ethernet Payload                              |
|                               New FCS                                             |
+++++
```

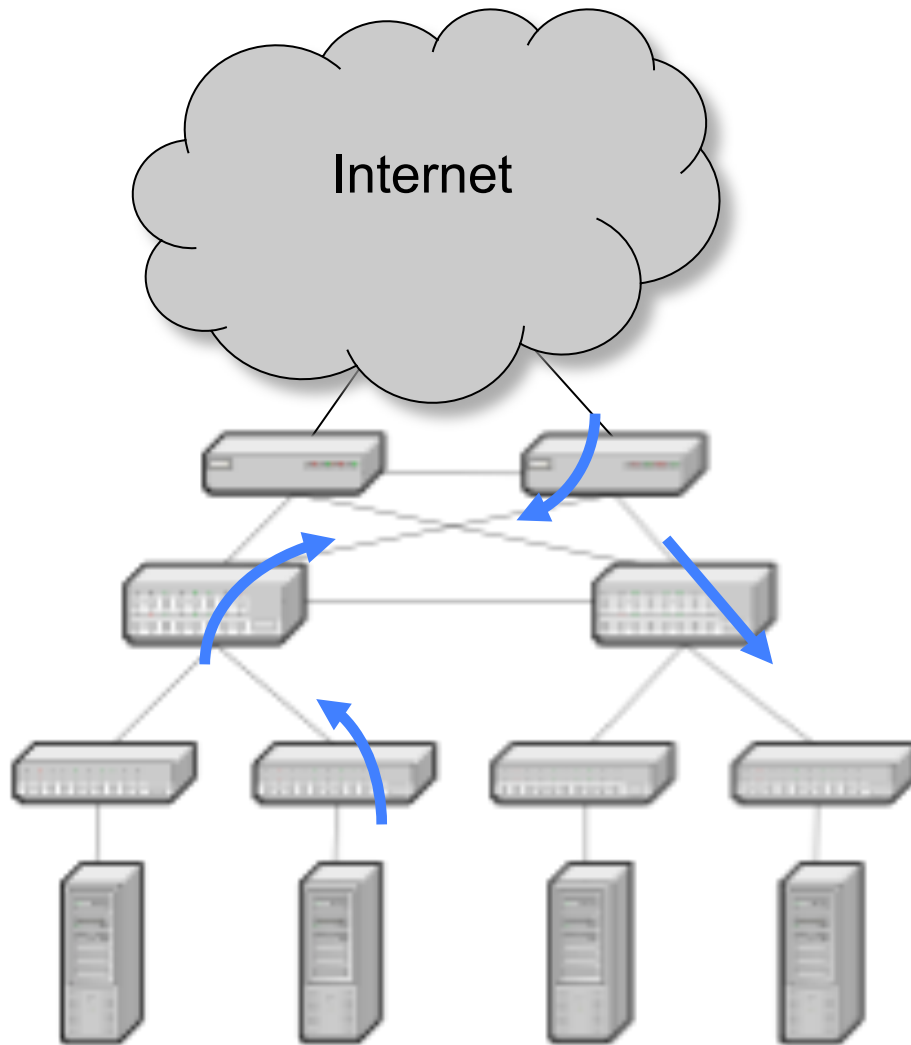
- sFlow on existing Ethernet switches captures the following TRILL fields
  - TRILL RBridge Addresses
  - Forwarding path
  - Hop count
  - Broadcast bit
- As well as inner/outer MAC addresses and encapsulated TCP/IP etc. data
- sFlow monitoring provides information about path utilization, applications using a path etc. Critical for load balancing and troubleshooting TRILL

-

# TRILL Fabrics



# sFlow Overview: captures packet path



- Each packet sample captures the forwarding path for the packet
- Threading together the paths provides a constantly updating picture of network topology and host locations
- The combination of forwarding table data and packet headers provides an integrated view of traffic. E.g. you can filter on forwarding attributes (VLAN, MPLS, route) and see traffic, or filter on traffic and identify forwarding paths.

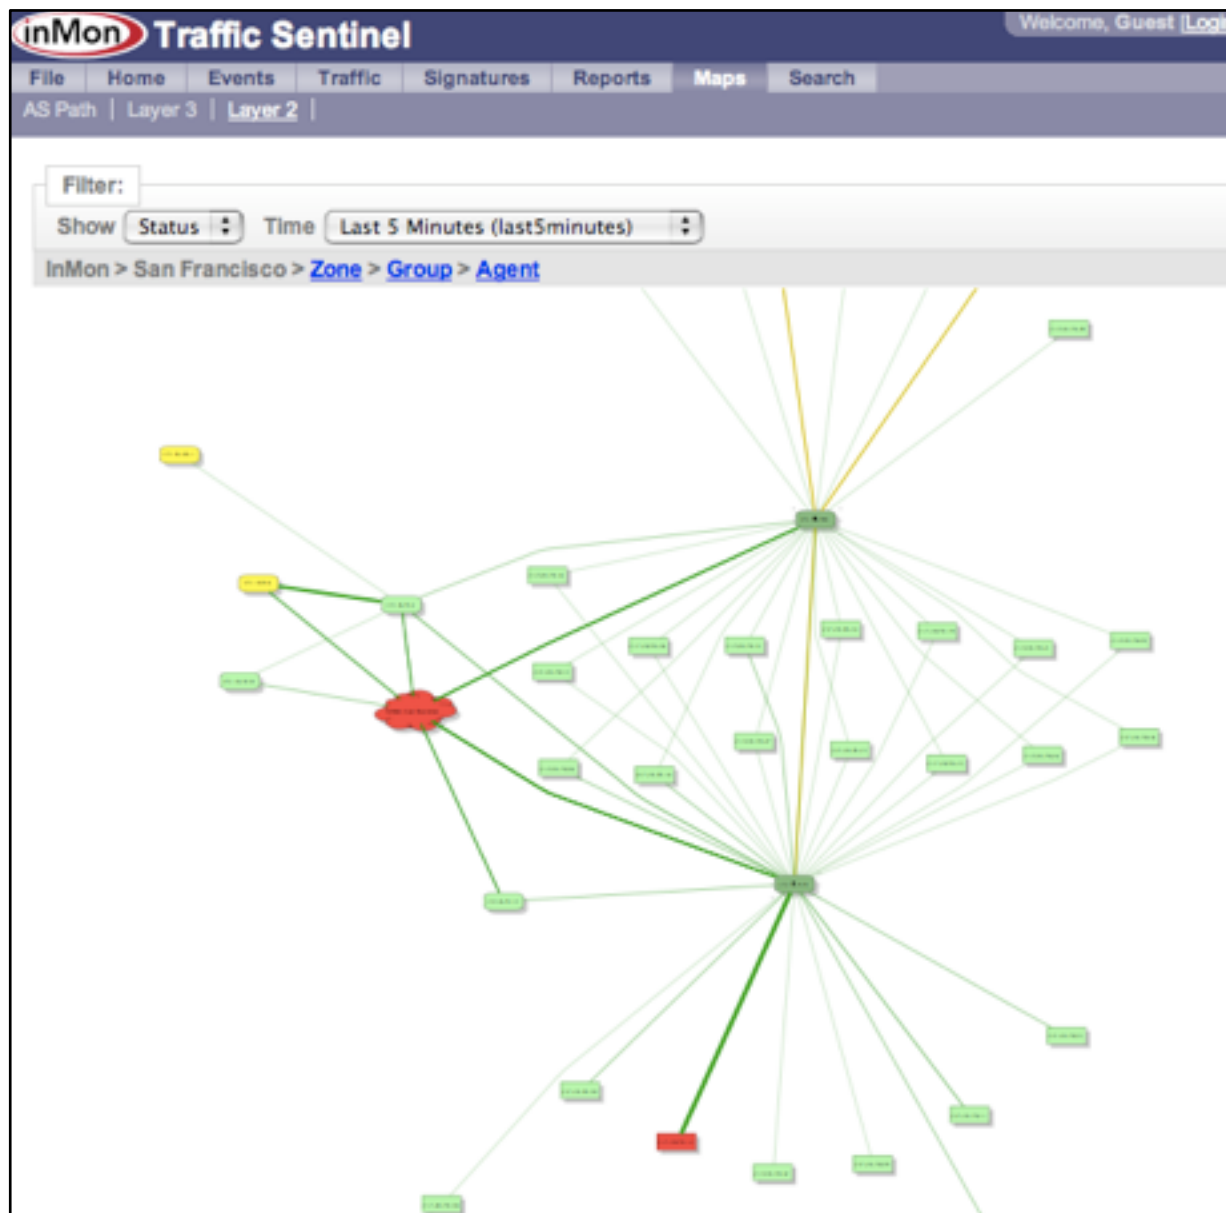


# Traffic Sentinel: Multivendor topology discovery

Uses:

- sFlow
- CDP
- FDP
- LLDP
- Spanning-tree
- Bridge-tables
- and more...

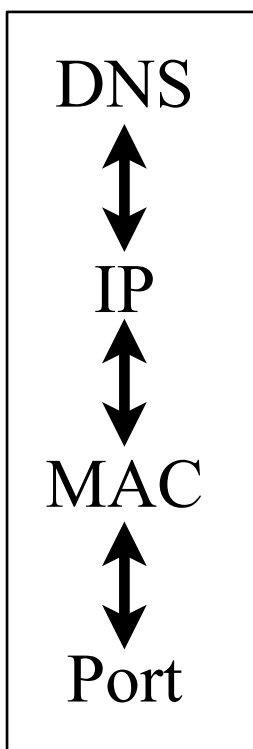
- Auto-layout
- Mouse-wheel zoom
- Show Status, Traffic (refreshed every minute)





# Traffic Sentinel: End-host location

Uses:  
sFlow  
SNMP



The screenshot shows the inMon Traffic Sentinel interface with the following navigation and filter options:

- Menu: File, Home, Events, Traffic, Signatures, Reports, Maps, Search
- Sub-menu: Status, Interfaces, Circles, Trend
- Filter: InMon > San Francisco > Embarcadero > Mission > 12-70.demo.inmon.com
- Show: Hosts, Interfaces, Connected Only

The main content area displays a table of host locations:

Interface	IfSpeed	Hosts
ethernet1/1	1Gb/sec	000480F548C2
ethernet2/1	1Gb/sec	172.16.239.85
ethernet2/2	1Gb/sec	95-118.demo.inmon.com
ethernet2/3	1Gb/sec	203-239.demo.inmon.com
ethernet2/5	1Gb/sec	31-238.demo.inmon.com
ethernet2/6	1Gb/sec	63-239.demo.inmon.com
ethernet2/7	1Gb/sec	48-44.demo.inmon.com
ethernet2/8	1Gb/sec	222-122.demo.inmon.com
ethernet2/11	1Gb/sec	160-118.demo.inmon.com
ethernet2/12	1Gb/sec	160-118.demo.inmon.com
ethernet2/13	1Gb/sec	4-239.demo.inmon.com
ethernet2/14	1Gb/sec	76-239.demo.inmon.com
ethernet2/15	1Gb/sec	41-238.demo.inmon.com

With sFlow, host locations can be updated within 60 seconds



## Simple agents



# sFlow<sup>®</sup> Architecture

Simple agents



# sFlow<sup>®</sup> Architecture

Simple agents



Smart collector

# sFlow<sup>®</sup> Architecture

Simple agents

Easy to implement



Smart collector

# sFlow<sup>®</sup> Architecture

Simple agents

Easy to implement

Embedded, wire-speed



Smart collector

# sFlow<sup>®</sup> Architecture

Simple agents

Easy to implement

Embedded, wire-speed

Low cost



Smart collector

# sFlow<sup>®</sup> Architecture

Simple agents

Easy to implement

Embedded, wire-speed

Low cost

Counter-push



Smart collector

# sFlow<sup>®</sup> Architecture

Simple agents

Easy to implement

Embedded, wire-speed

Low cost

Counter-push

+ packet/transaction sampling



Smart collector

# sFlow<sup>®</sup> Architecture

Simple agents

Easy to implement

Embedded, wire-speed

Low cost

Counter-push

+ packet/transaction sampling



Smart collector

Network-wide,  
integrated,  
visibility and control

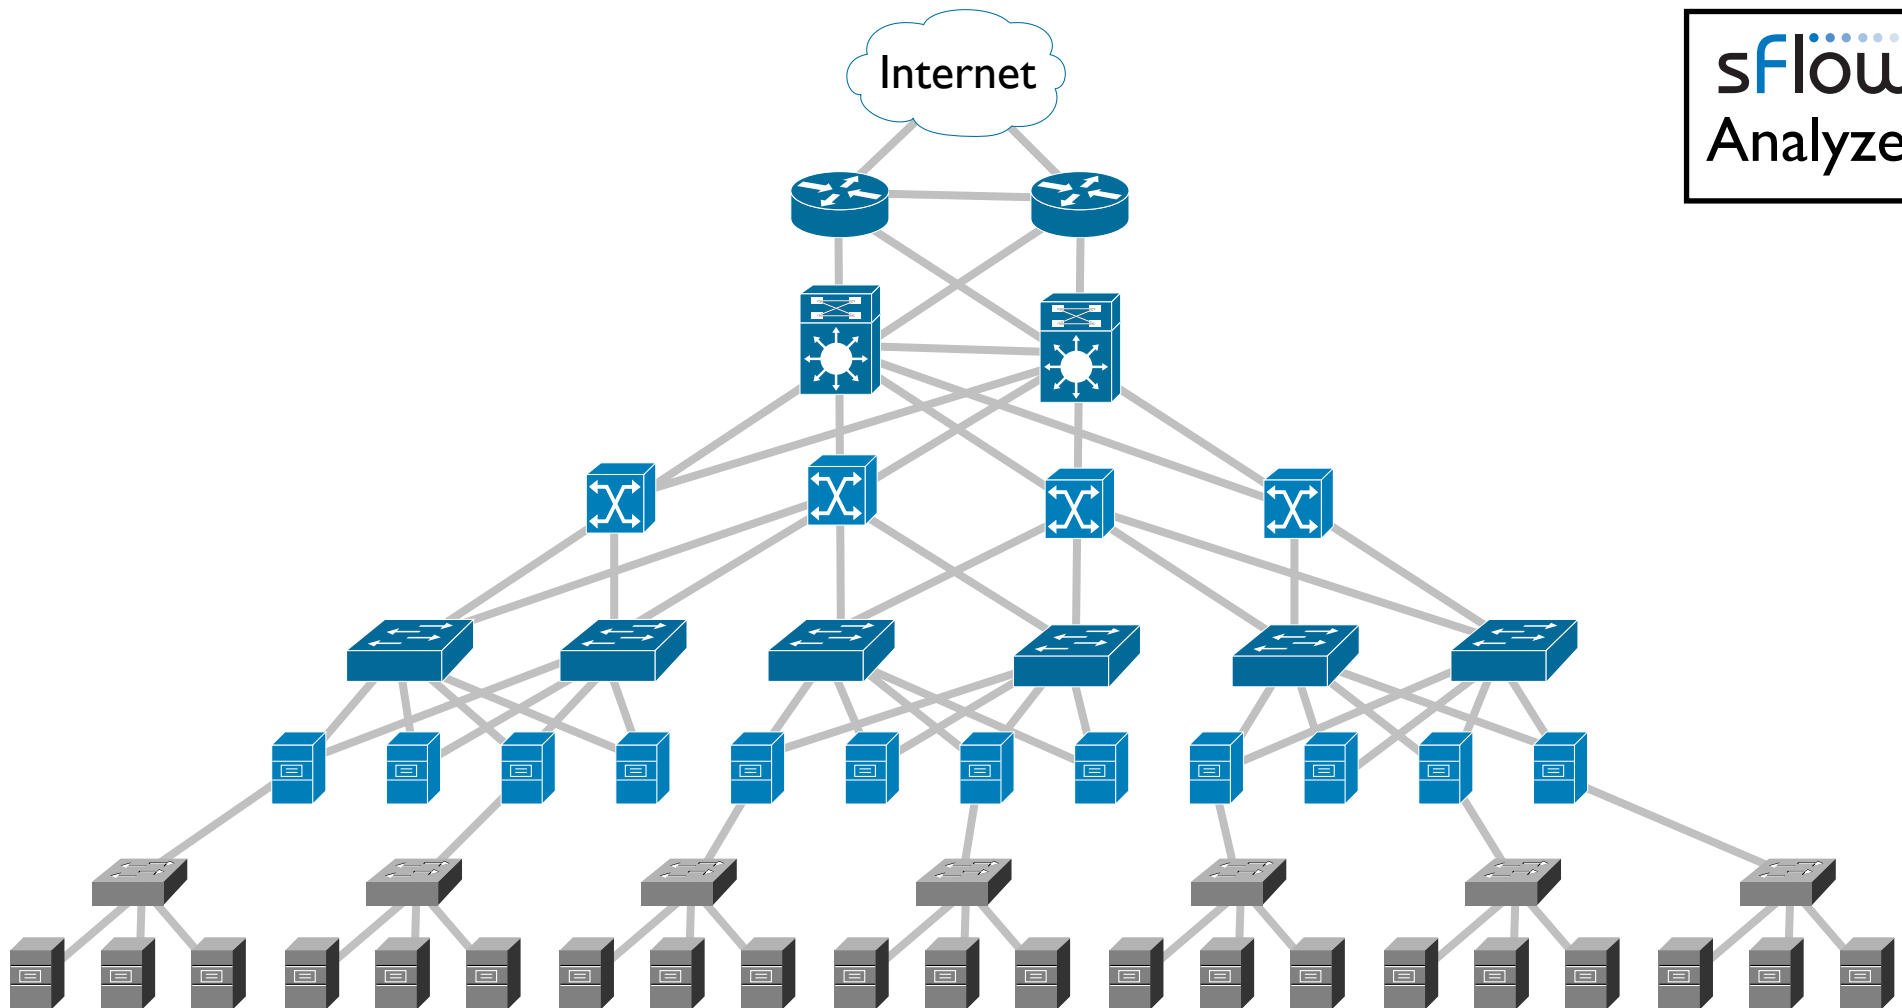


# sFlow Architecture

Simple agents



Smart collector

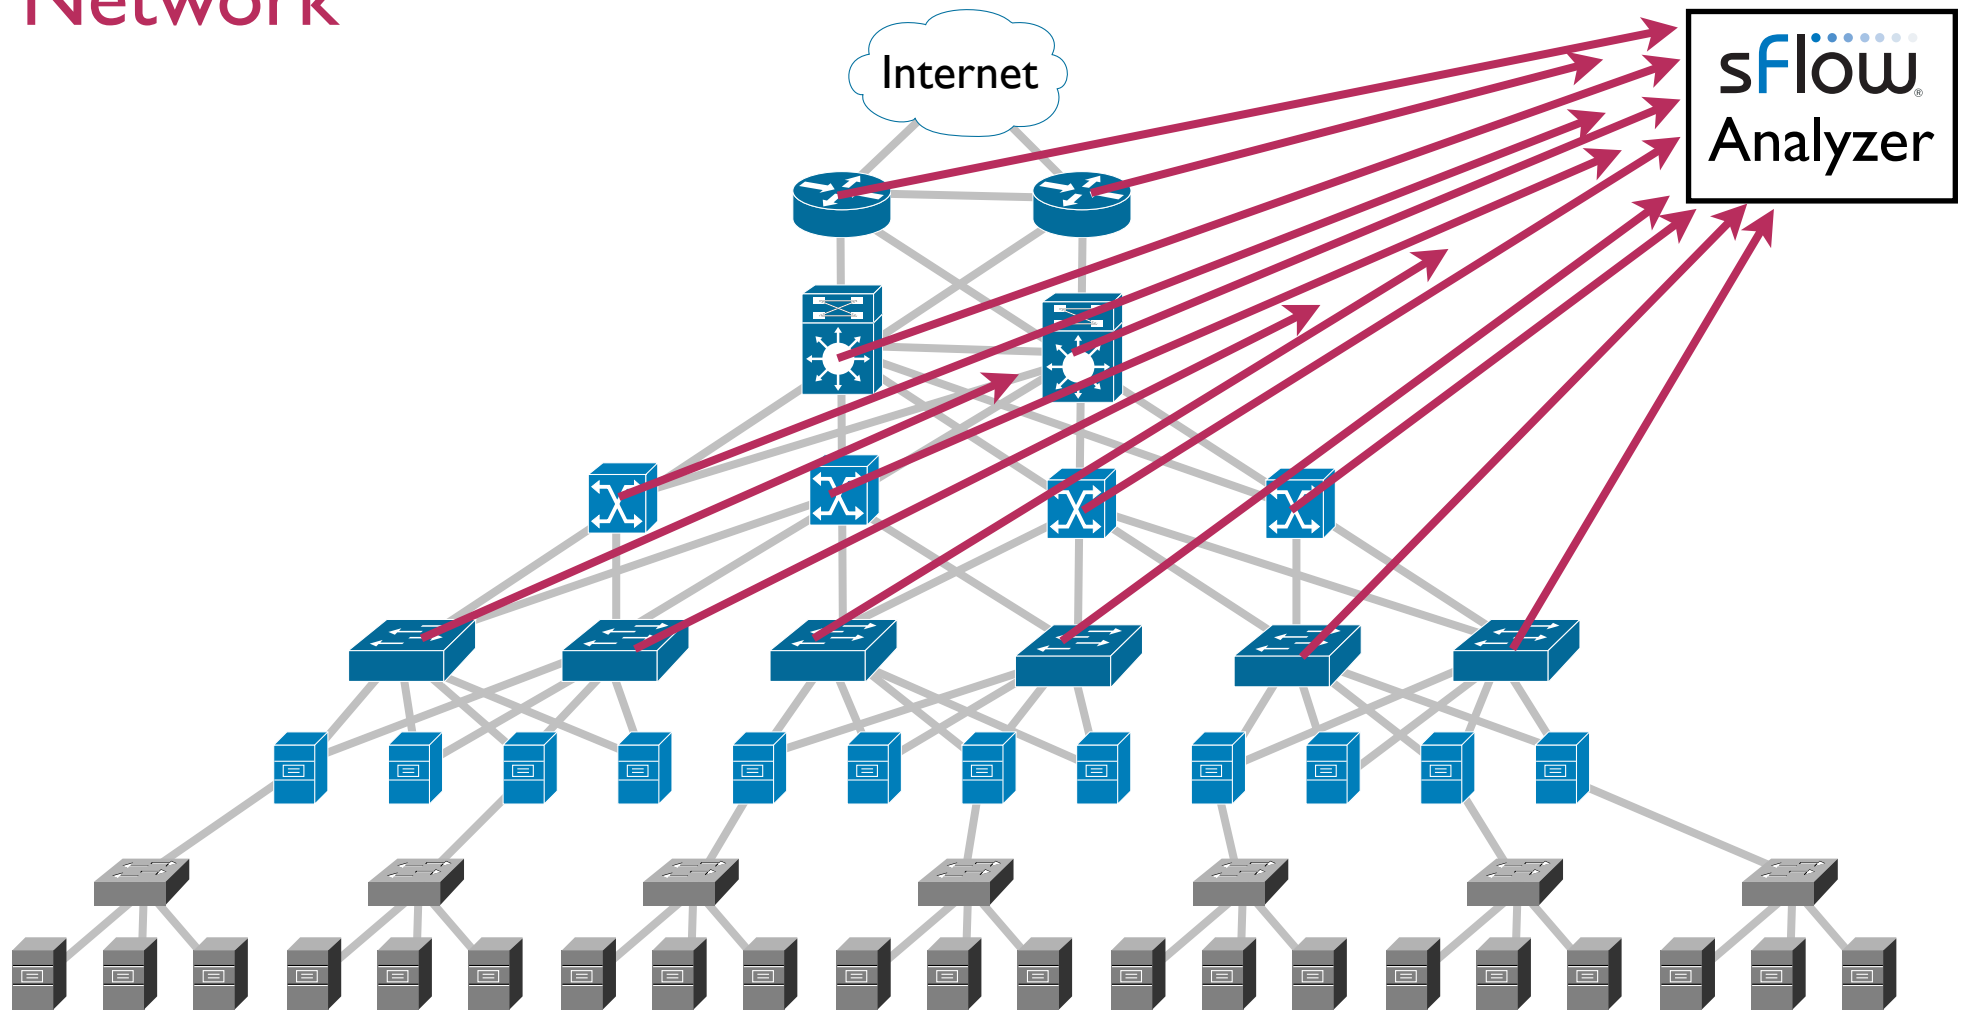


# sFlow<sup>®</sup> Architecture

Simple agents  
Network



Smart collector



# sFlow Architecture

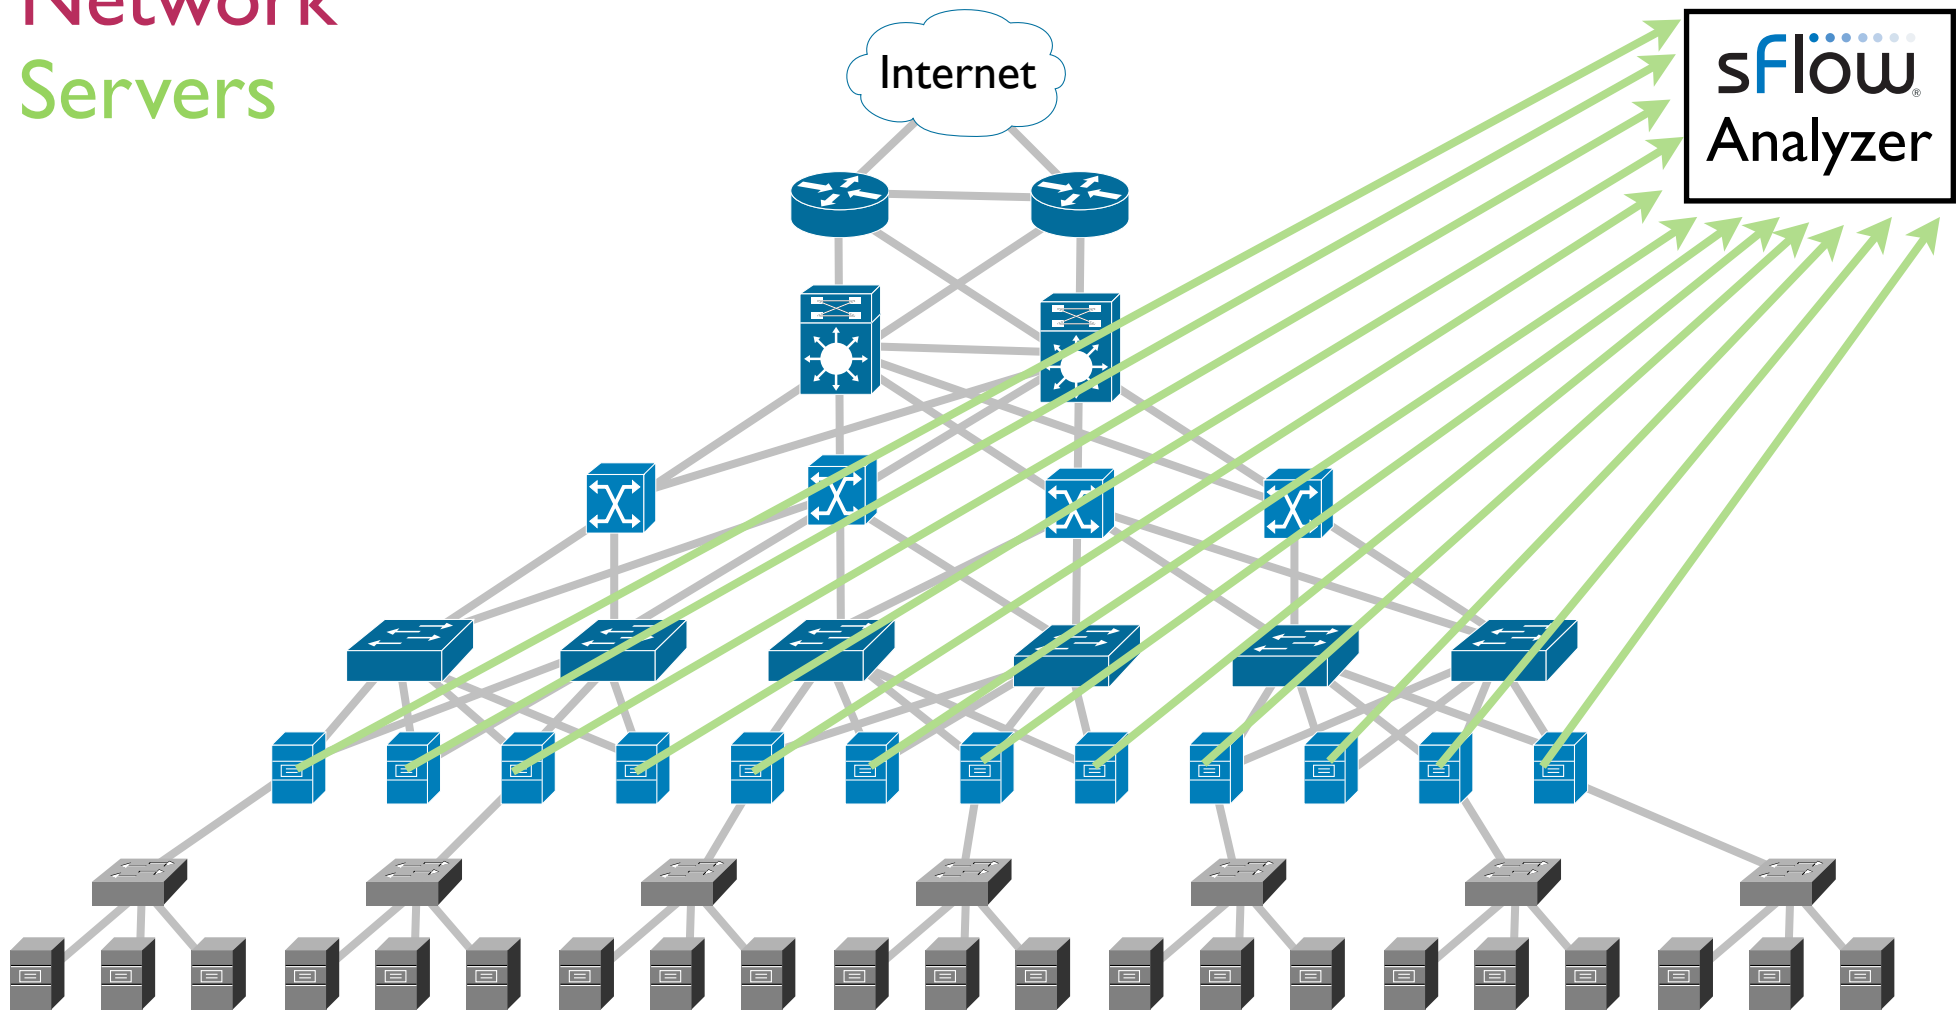
Simple agents

Network

Servers



Smart collector



# sFlow Architecture

Simple agents

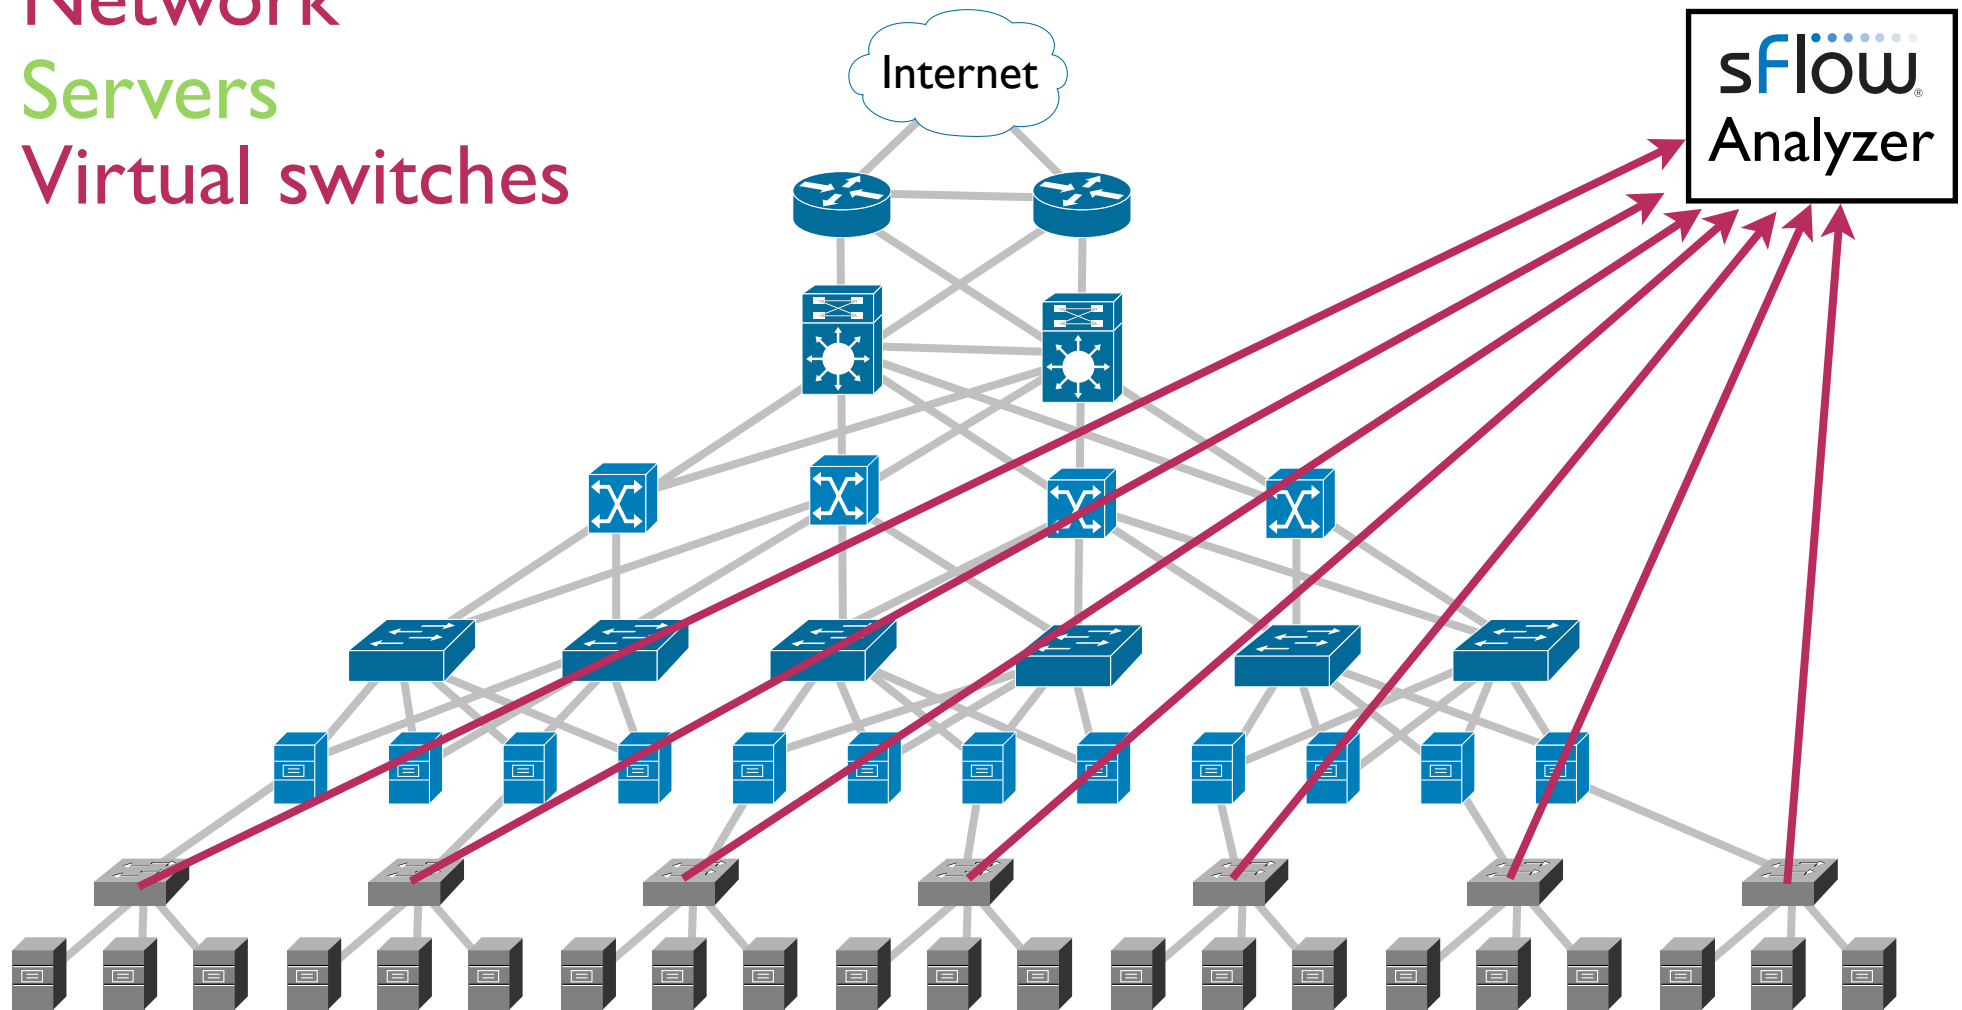
Network

Servers

Virtual switches



Smart collector



# sFlow<sup>®</sup> Architecture

Simple agents



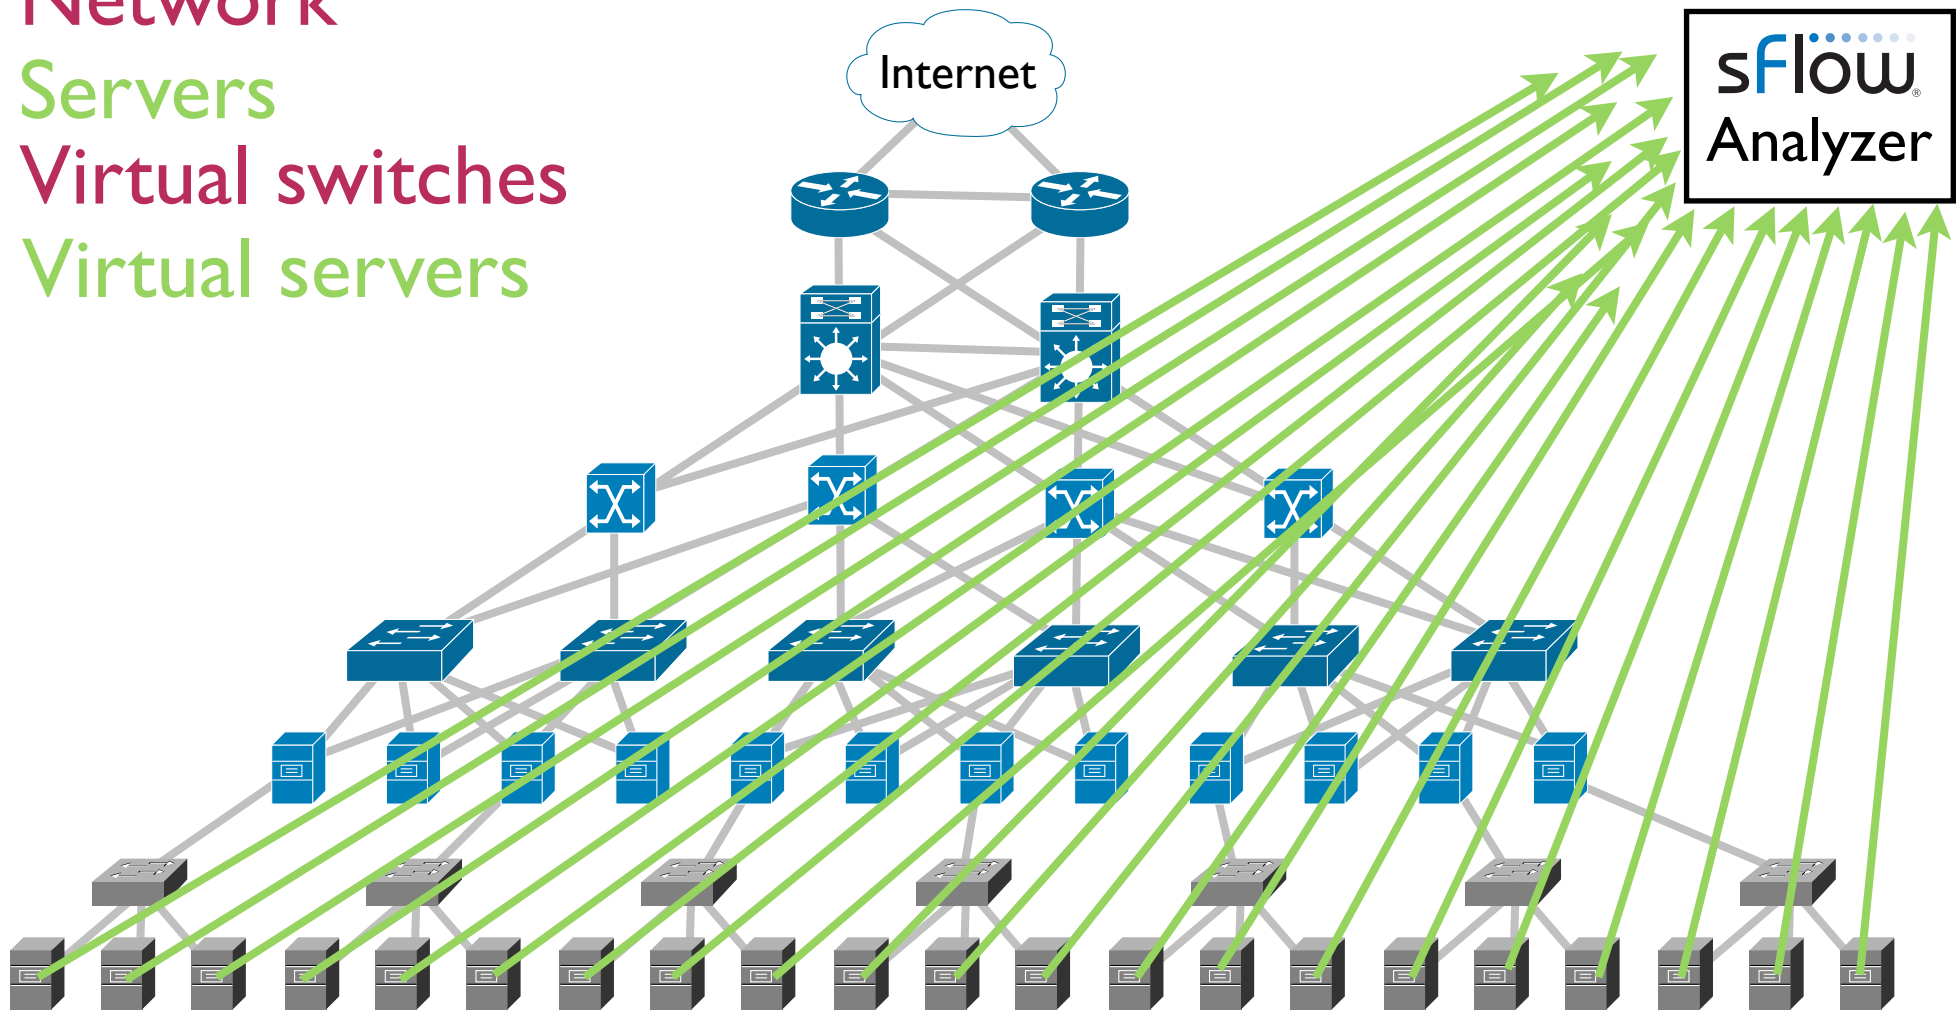
Smart collector

Network

Servers

Virtual switches

Virtual servers



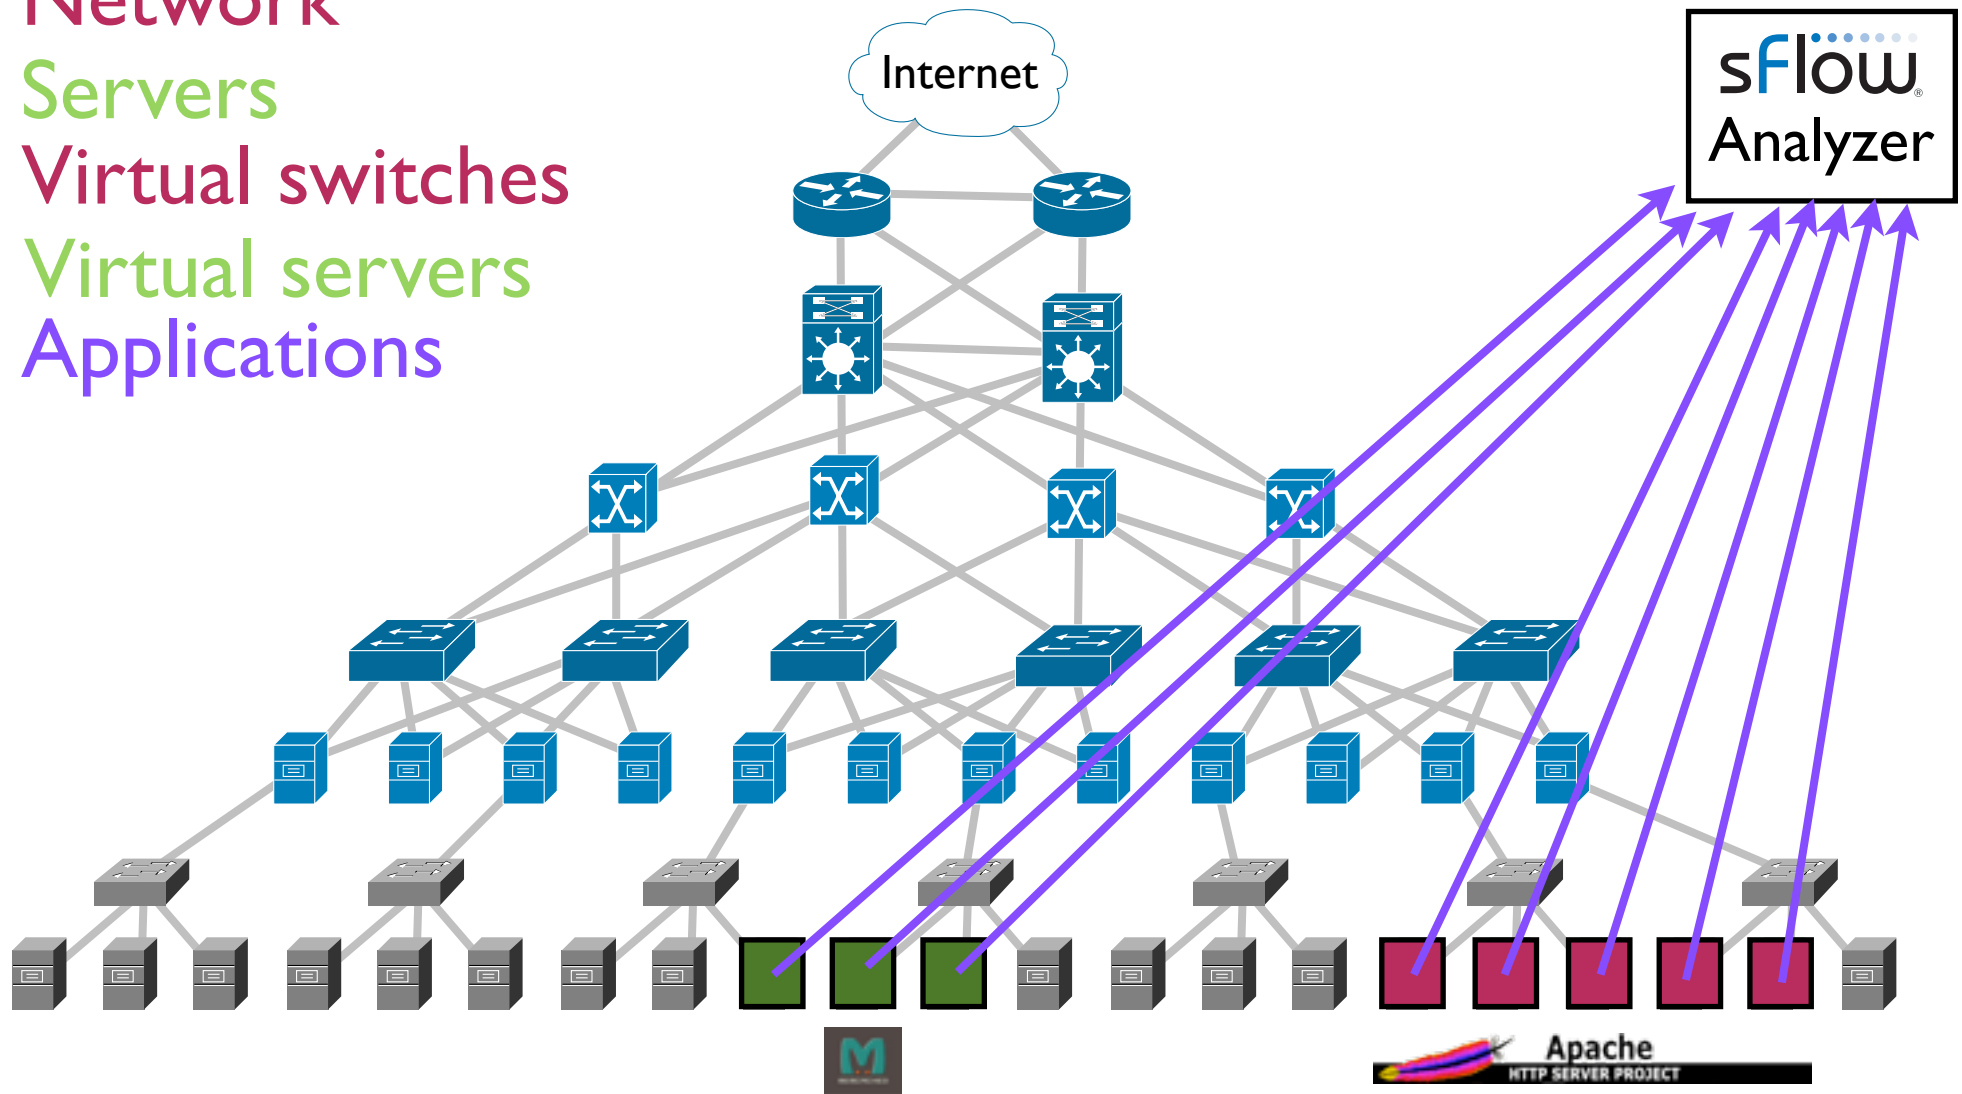
# sFlow Architecture

Simple agents



Smart collector

- Network
- Servers
- Virtual switches
- Virtual servers
- Applications



# sFlow<sup>®</sup> Architecture

Simple agents



Smart collector

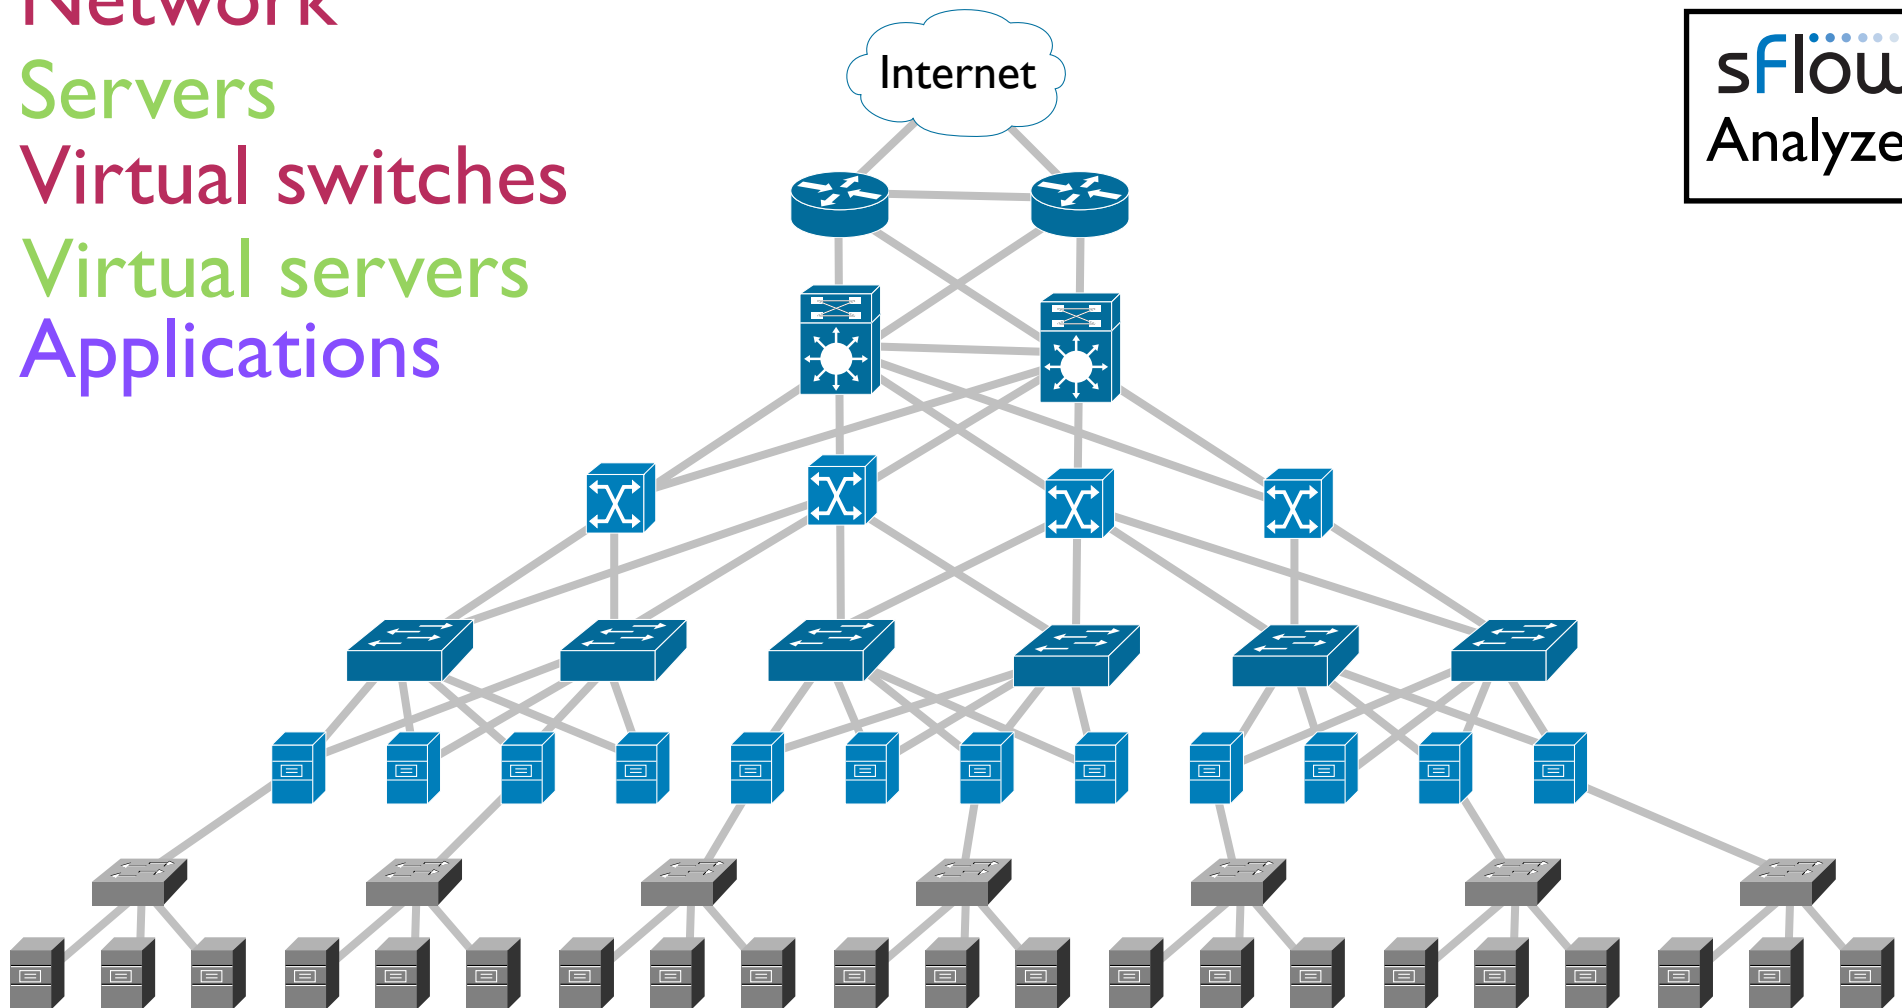
Network

Servers

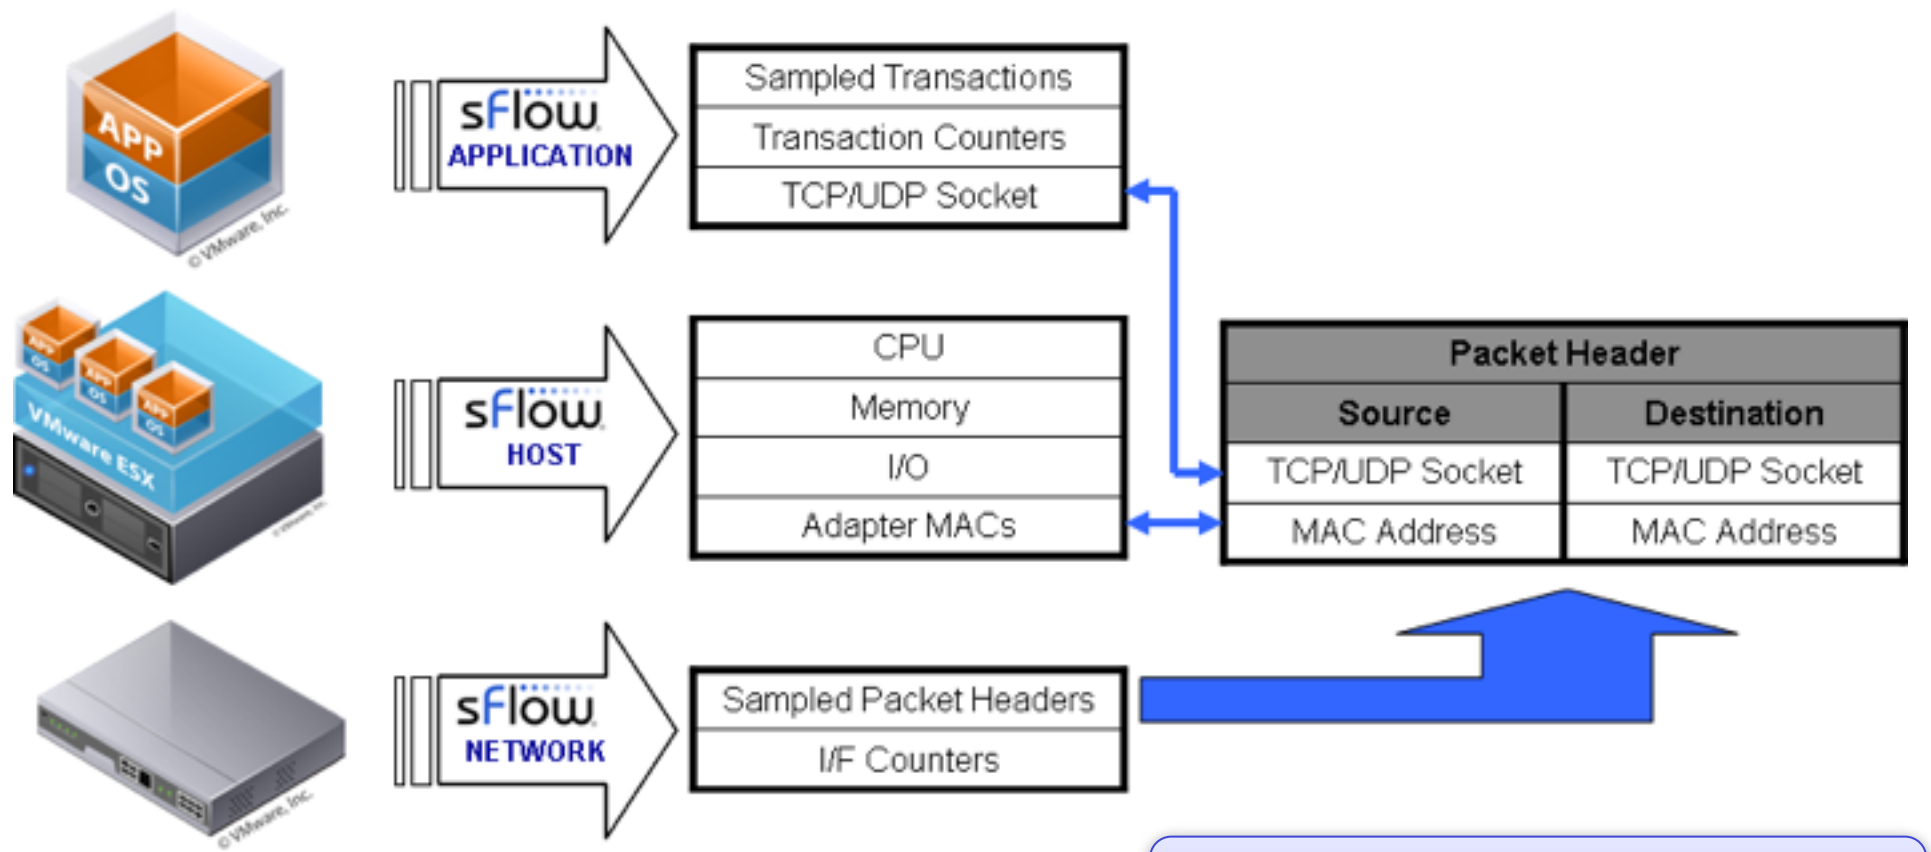
Virtual switches

Virtual servers

Applications



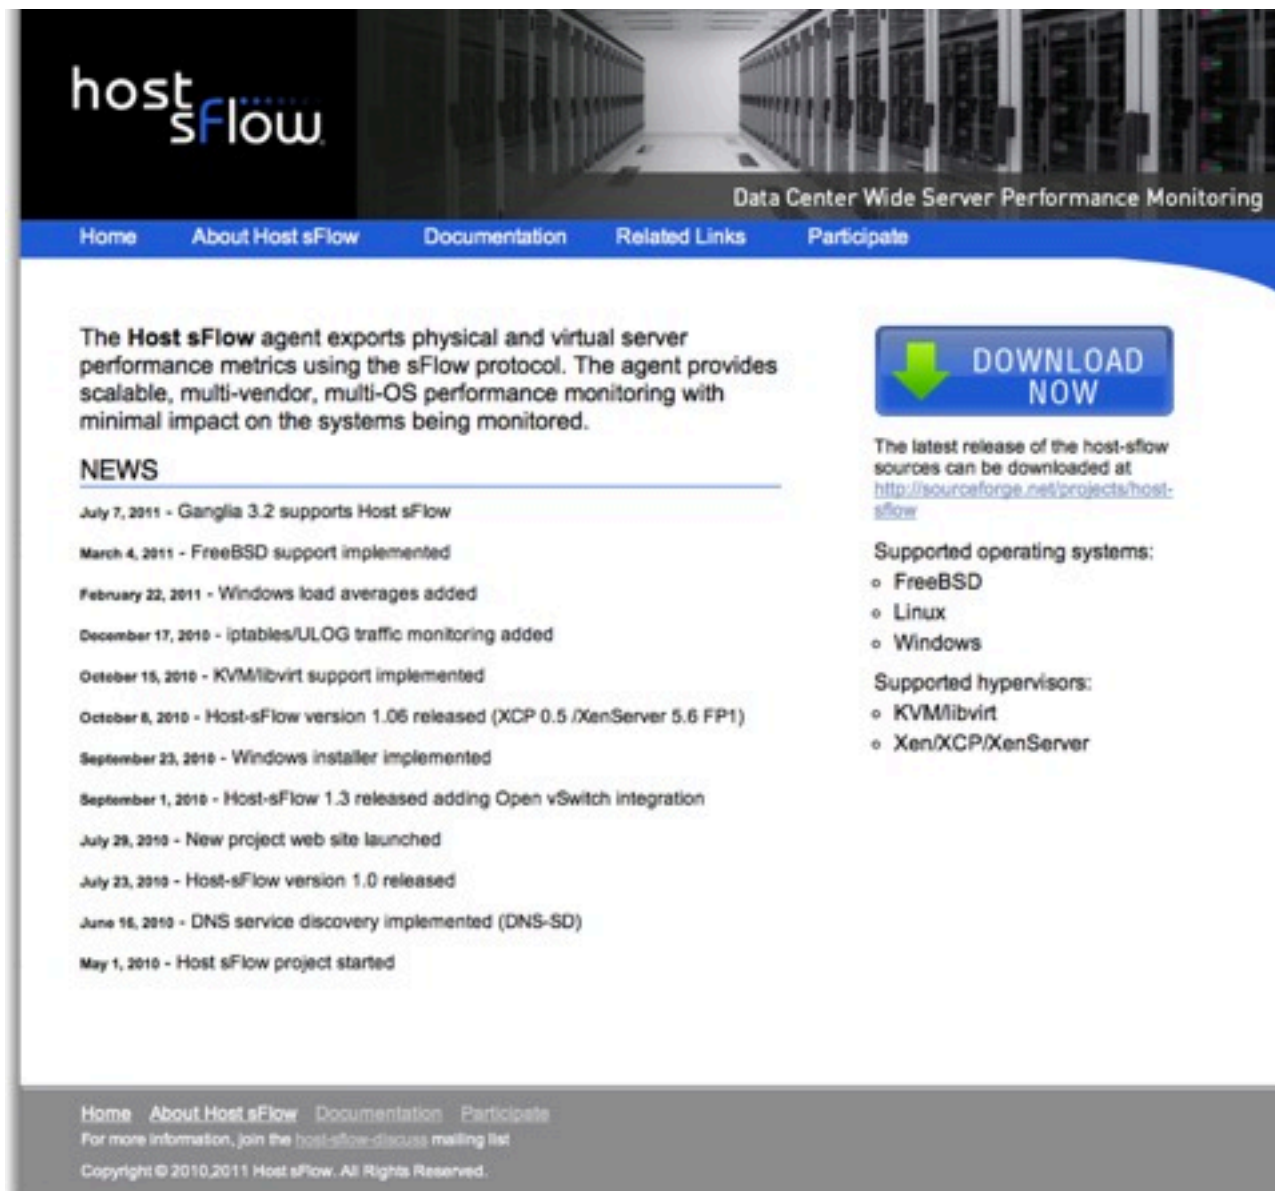
# Cross-layer correlation: Application, Host and Network



Standard measurements from different sources designed to be joinable at the collector

e.g. application response time increase correlated directly to congestion on network path





The screenshot shows the website for host-sflow, a project for Data Center Wide Server Performance Monitoring. The page features a navigation menu with links for Home, About Host sFlow, Documentation, Related Links, and Participate. The main content area includes a description of the Host sFlow agent, a 'NEWS' section with a list of updates from May 2010 to July 2011, and a 'DOWNLOAD NOW' button with a green arrow icon. Below the button, there is a link to the sourceforge project page and lists of supported operating systems (FreeBSD, Linux, Windows) and hypervisors (KVMlibvirt, Xen/XCP/XenServer). The footer contains a secondary navigation menu, a mailing list link, and copyright information for 2010 and 2011.

host sFlow  
Data Center Wide Server Performance Monitoring

Home About Host sFlow Documentation Related Links Participate

The **Host sFlow** agent exports physical and virtual server performance metrics using the sFlow protocol. The agent provides scalable, multi-vendor, multi-OS performance monitoring with minimal impact on the systems being monitored.

**NEWS**

- July 7, 2011 - Ganglia 3.2 supports Host sFlow
- March 4, 2011 - FreeBSD support implemented
- February 22, 2011 - Windows load averages added
- December 17, 2010 - iptables/ULOG traffic monitoring added
- October 15, 2010 - KVMlibvirt support implemented
- October 8, 2010 - Host-sFlow version 1.06 released (XCP 0.5 /XenServer 5.6 FP1)
- September 23, 2010 - Windows installer implemented
- September 1, 2010 - Host-sFlow 1.3 released adding Open vSwitch integration
- July 29, 2010 - New project web site launched
- July 23, 2010 - Host-sFlow version 1.0 released
- June 16, 2010 - DNS service discovery implemented (DNS-SD)
- May 1, 2010 - Host sFlow project started

**DOWNLOAD NOW**

The latest release of the host-sflow sources can be downloaded at <http://sourceforge.net/projects/host-sflow>

Supported operating systems:

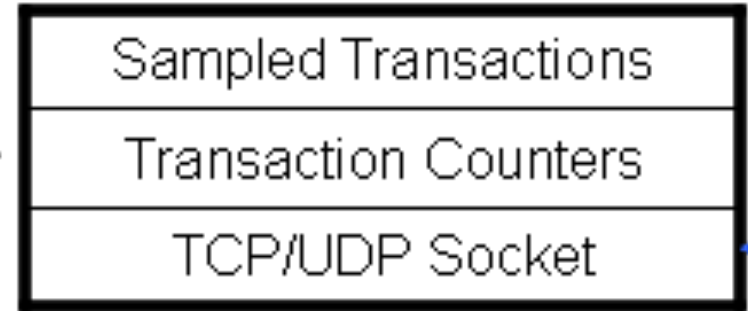
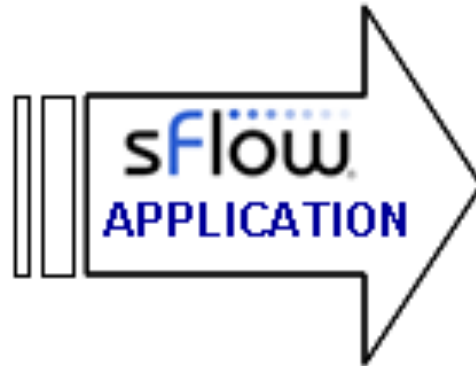
- o FreeBSD
- o Linux
- o Windows

Supported hypervisors:

- o KVMlibvirt
- o Xen/XCP/XenServer

Home About Host sFlow Documentation Participate  
For more information, join the [host-sflow-discuss](#) mailing list  
Copyright © 2010,2011 Host sFlow. All Rights Reserved.

# sFlow-APPLICATION



## Examples:

**NFS/CIFS transactions** (file path, bytes, response-time, socket)

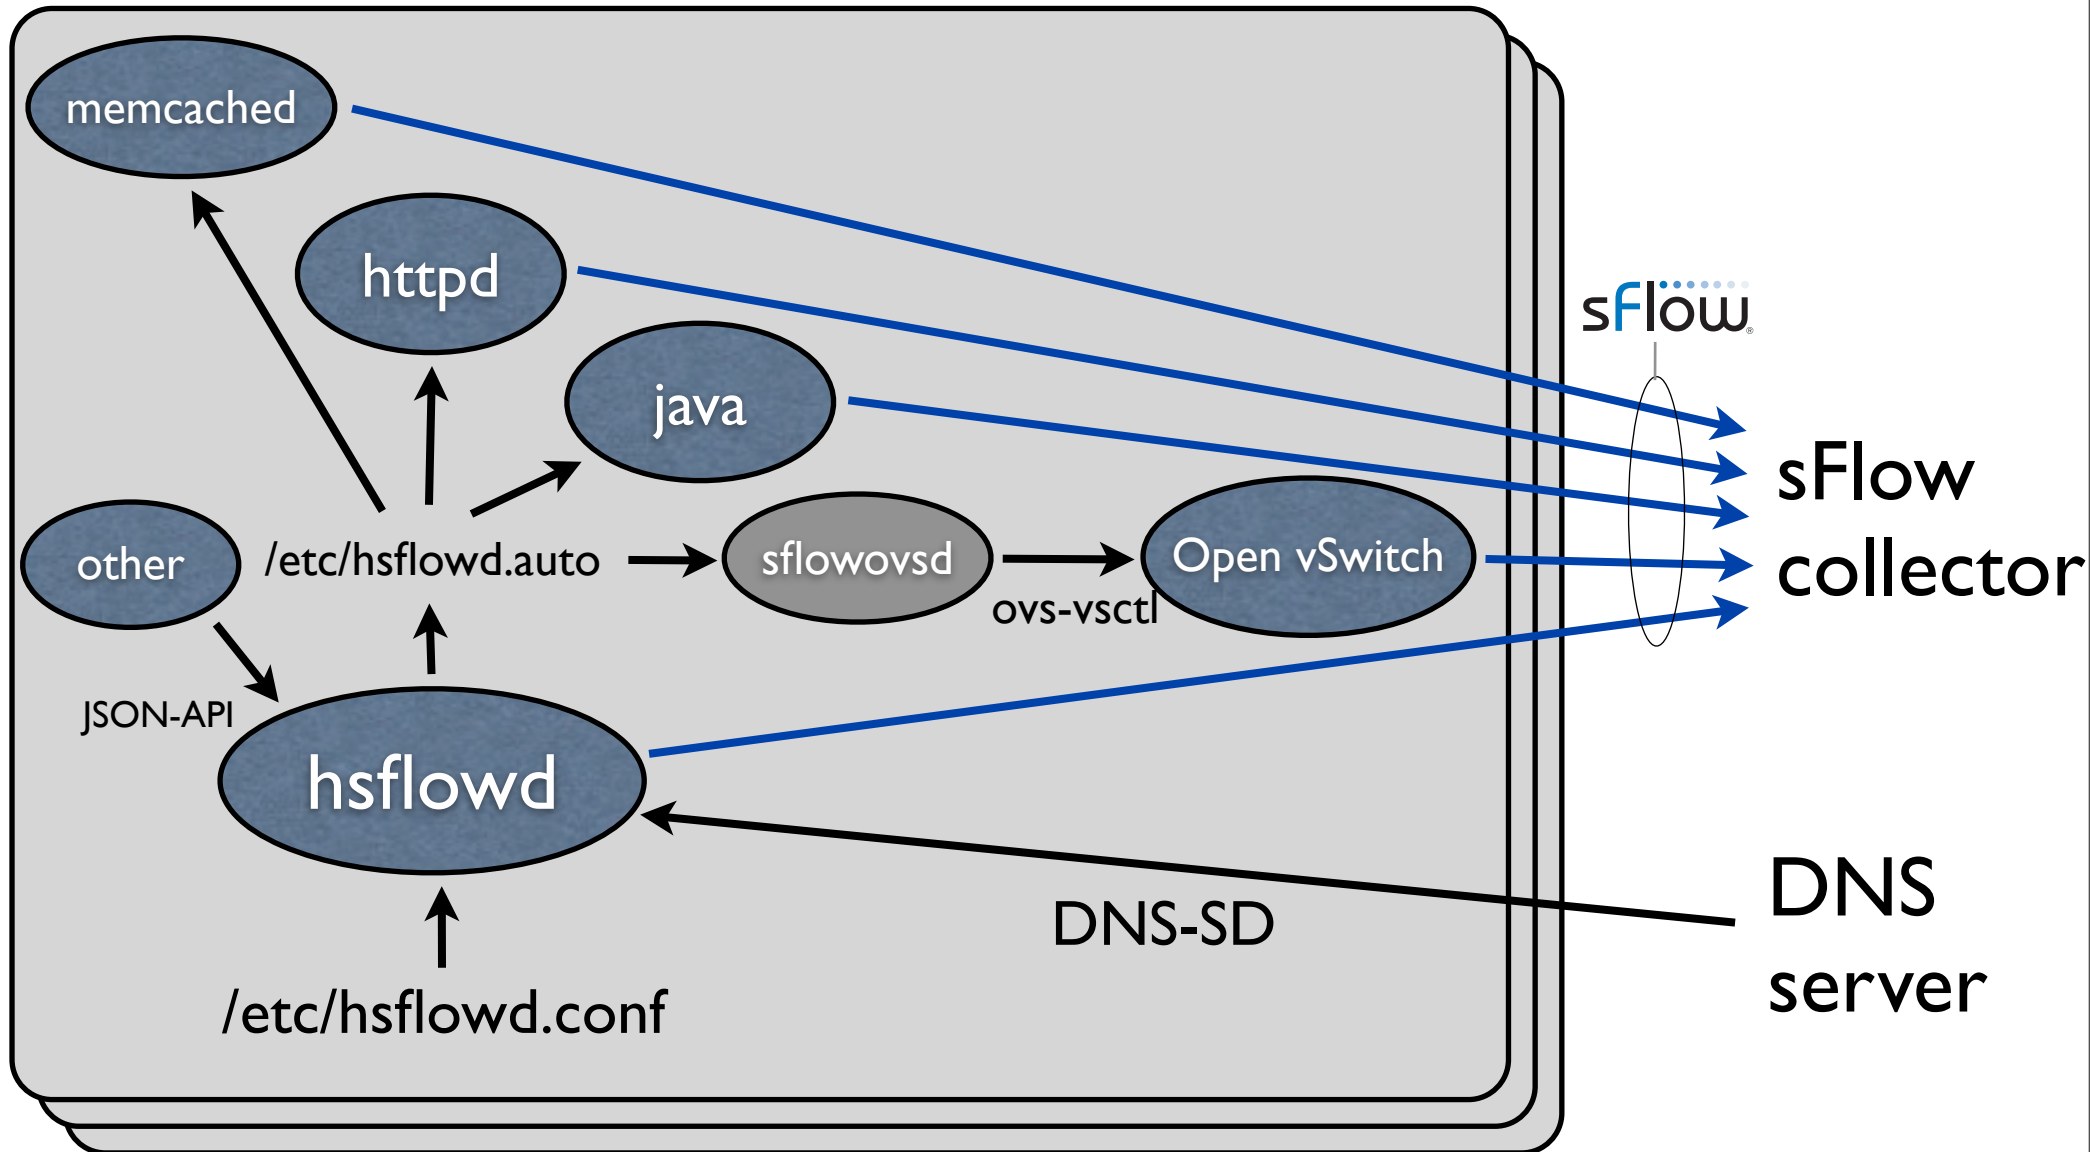
**HTTP requests** (URL, user-agent, mime-type, bytes, response-time, socket)

**Memcached lookups** (key, value-bytes, hit/miss, socket)

**Database queries** (query#, response-time, socket)

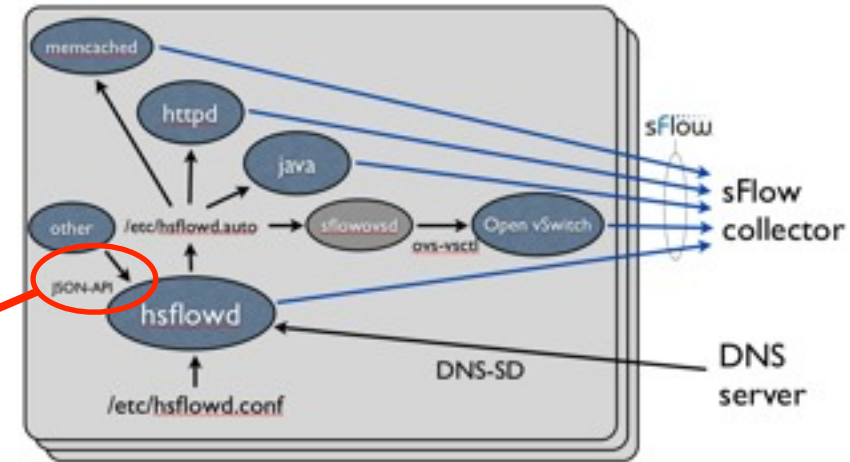
Application layer measurements much more valuable when correlated with performance of every component in underlying infrastructure!

# Host sFlow distributed agent



# Host sFlow - JSON API

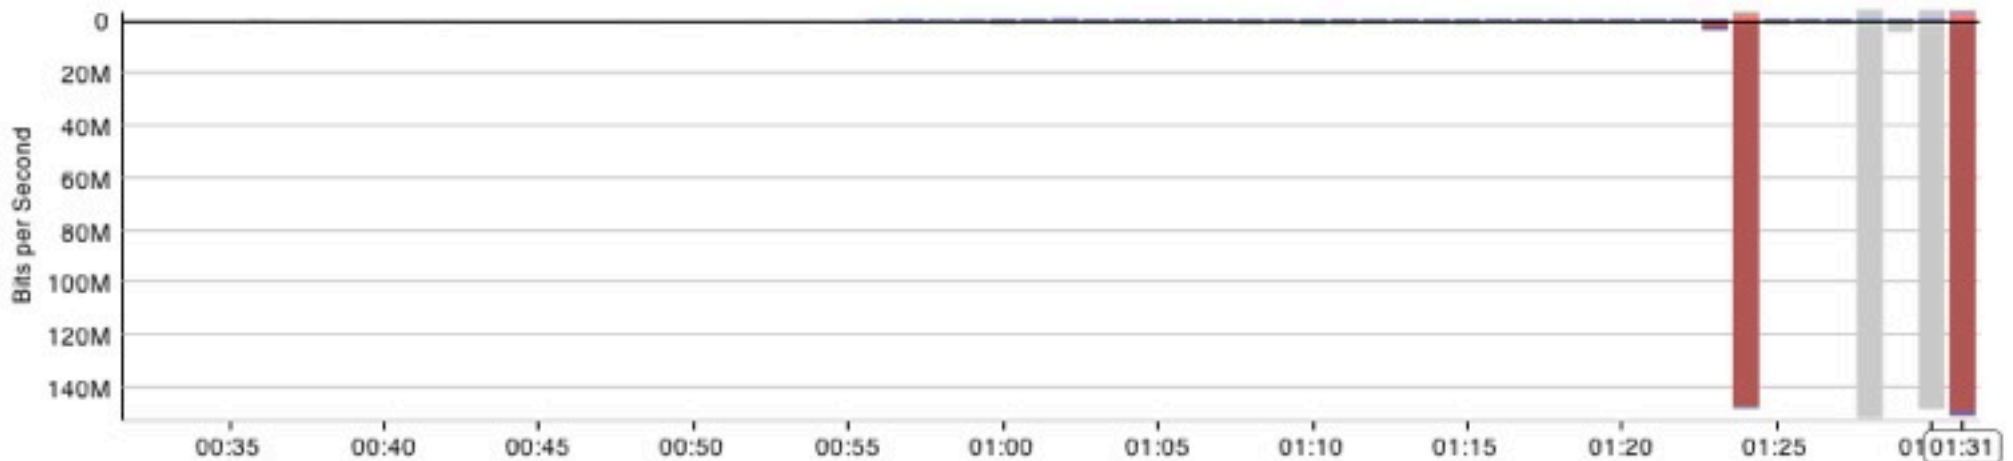
```
{
  "flow_sample": {
    "app_name": "myapp",
    "sampling_rate": 100
  },
  "app_operation": {
    "operation": "task.start",
    "attributes": "id=123&user=root",
    "status_descr": "OK",
    "status": 0,
    "req_bytes": 43,
    "resp_bytes": 234,
    "uS": 2000
  },
  "app_initiator": {"actor": "123"},
  "app_target": {"actor": "231"},
  "extended_socket_ipv4": {
    "protocol": 6,
    "local_ip": "10.0.0.1",
    "remote_ip": "10.0.0.23",
    "local_port": 123,
    "remote_port": 43032
  }
}
```



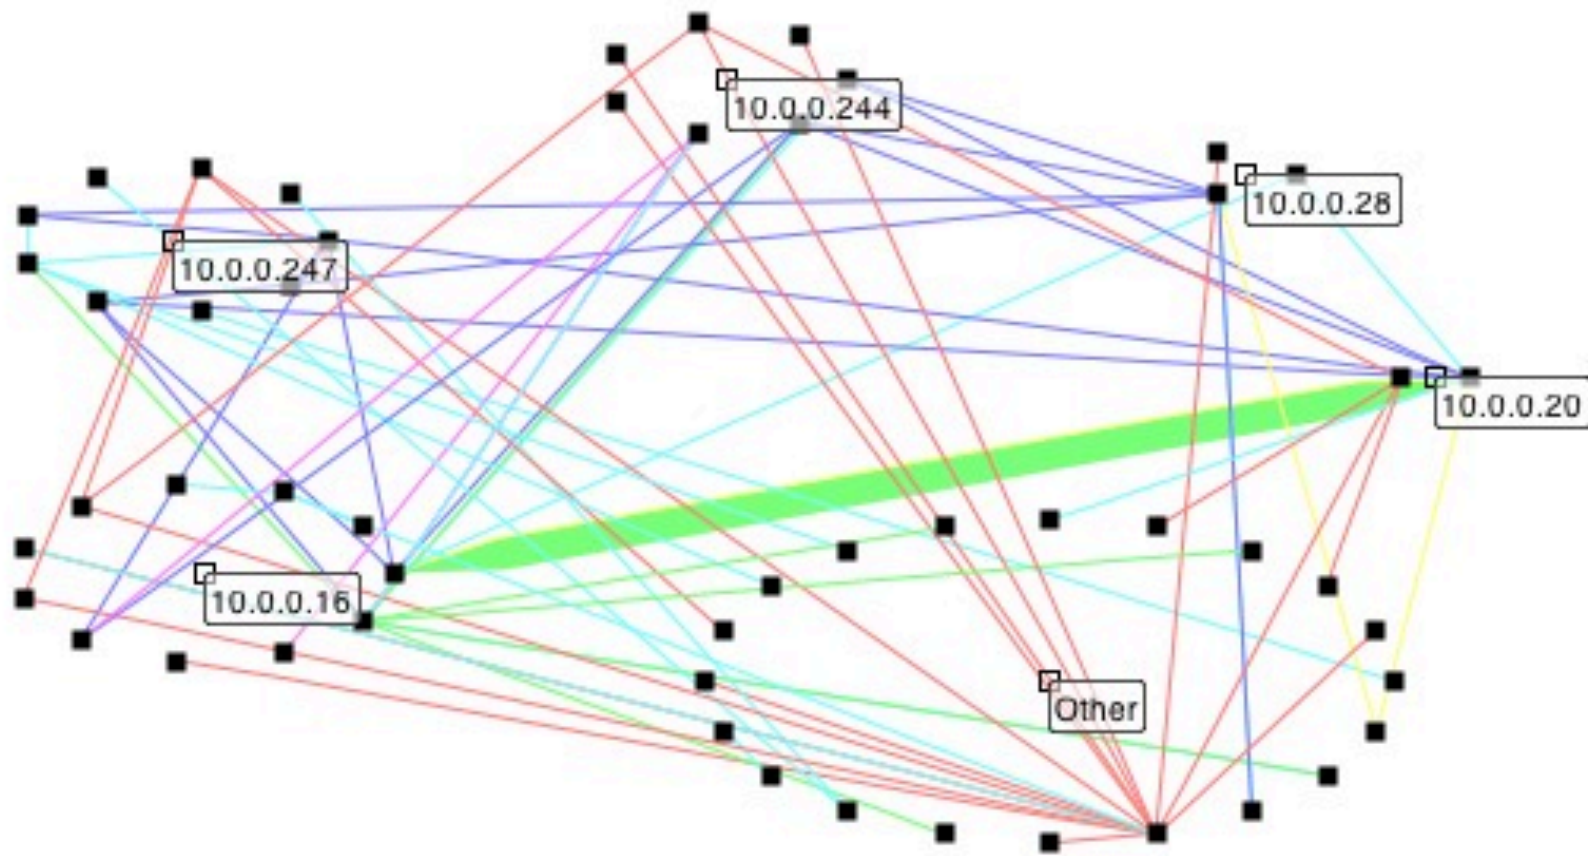
- UDP msg to hsflowd on localhost:36343
- most fields **optional**
- sampling in app. or hsflowd
- counters in app. or hsflowd
- hsflowd sends binary sFlow-APPLICATION feed to configured collectors.

# XenMotion bandwidth

	Server Address	Server Port	Client Address	Bytes From Server	Bytes To Server	Bytes
■	<a href="#">xenserver2 (10.0.0.16)</a>	<a href="#">TCP:80 (www-http)</a>	<a href="#">xenserver1 (10.0.0.20)</a>	2.73M	149.54M	152.26M
■	<a href="#">ganglia (10.0.0.112)</a>	<a href="#">TCP:11211 (memcache)</a>	<a href="#">xenvm1 (10.0.0.150)</a>	773.75K	1.07M	1.85M
■	<a href="#">openfiler (10.0.0.18)</a>	<a href="#">TCP:3260 (iscsi-target)</a>	<a href="#">xenserver1 (10.0.0.20)</a>	251.93K	21.21K	273.14K
■	<a href="#">openfiler (10.0.0.18)</a>	<a href="#">TCP:3260 (iscsi-target)</a>	<a href="#">xenserver2 (10.0.0.16)</a>	167.92K	72.85K	240.77K
■	<a href="#">openfiler (10.0.0.18)</a>	<a href="#">TCP:2049 (nfs)</a>	<a href="#">xenserver1 (10.0.0.20)</a>	0	211.39K	211.39K

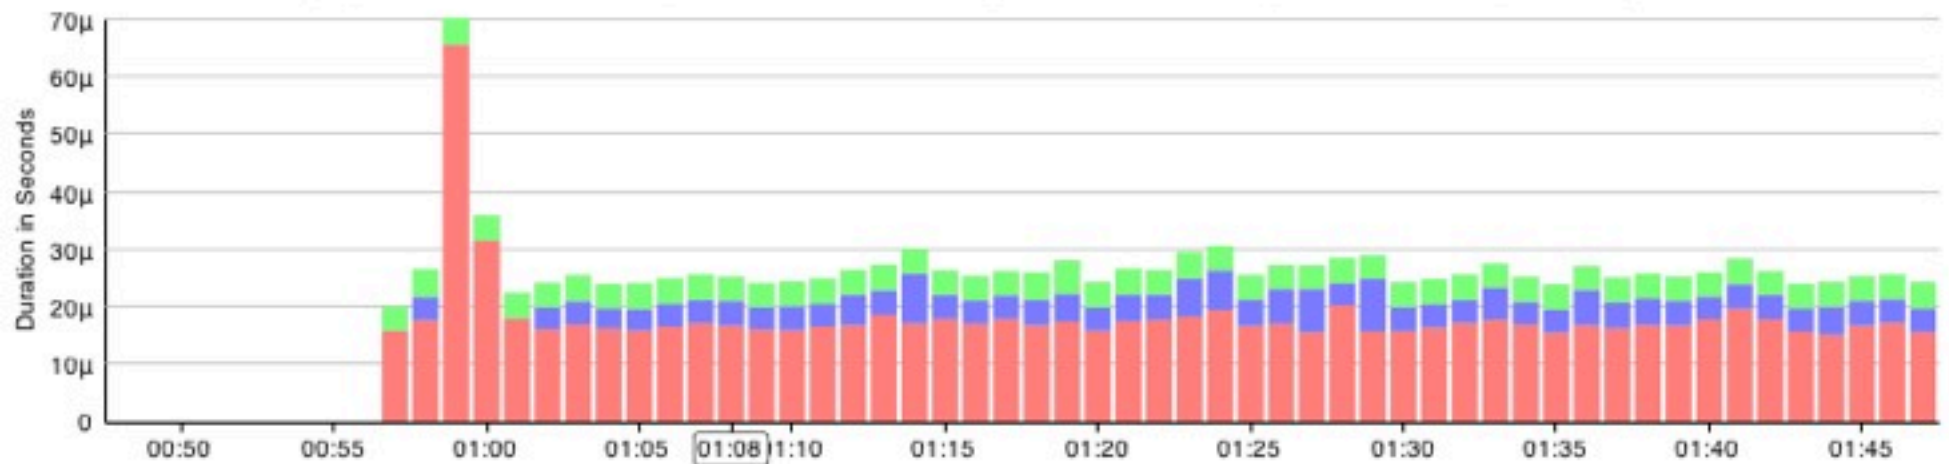


# XenMotion bandwidth



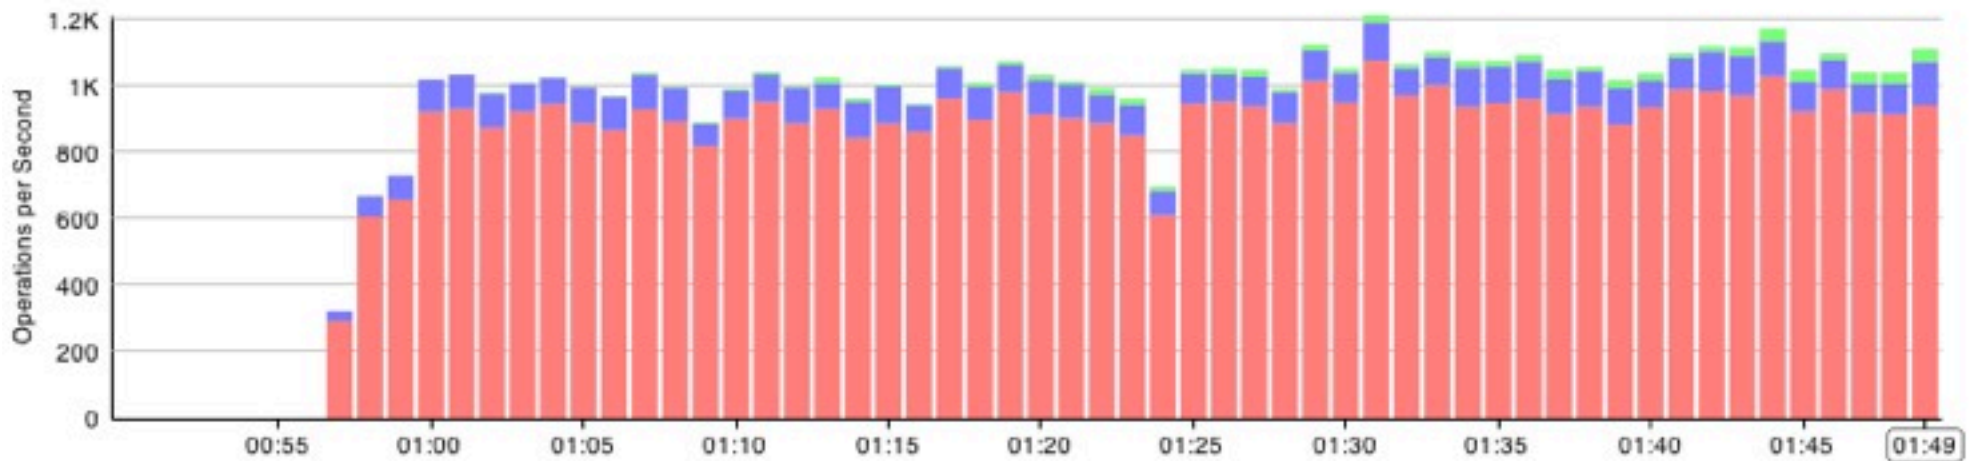
# Application throughput and response time

	Memcache Protocol	Memcache Command	Memcache Status	Memcache Key	Duration
■	ASCII	SET	STORED	brutis	16.70μ
■	ASCII	GET	EXISTS	brutis	4.25μ
■	ASCII	GET	NOT_FOUND	brutis	4.21μ



# Application throughput and response time

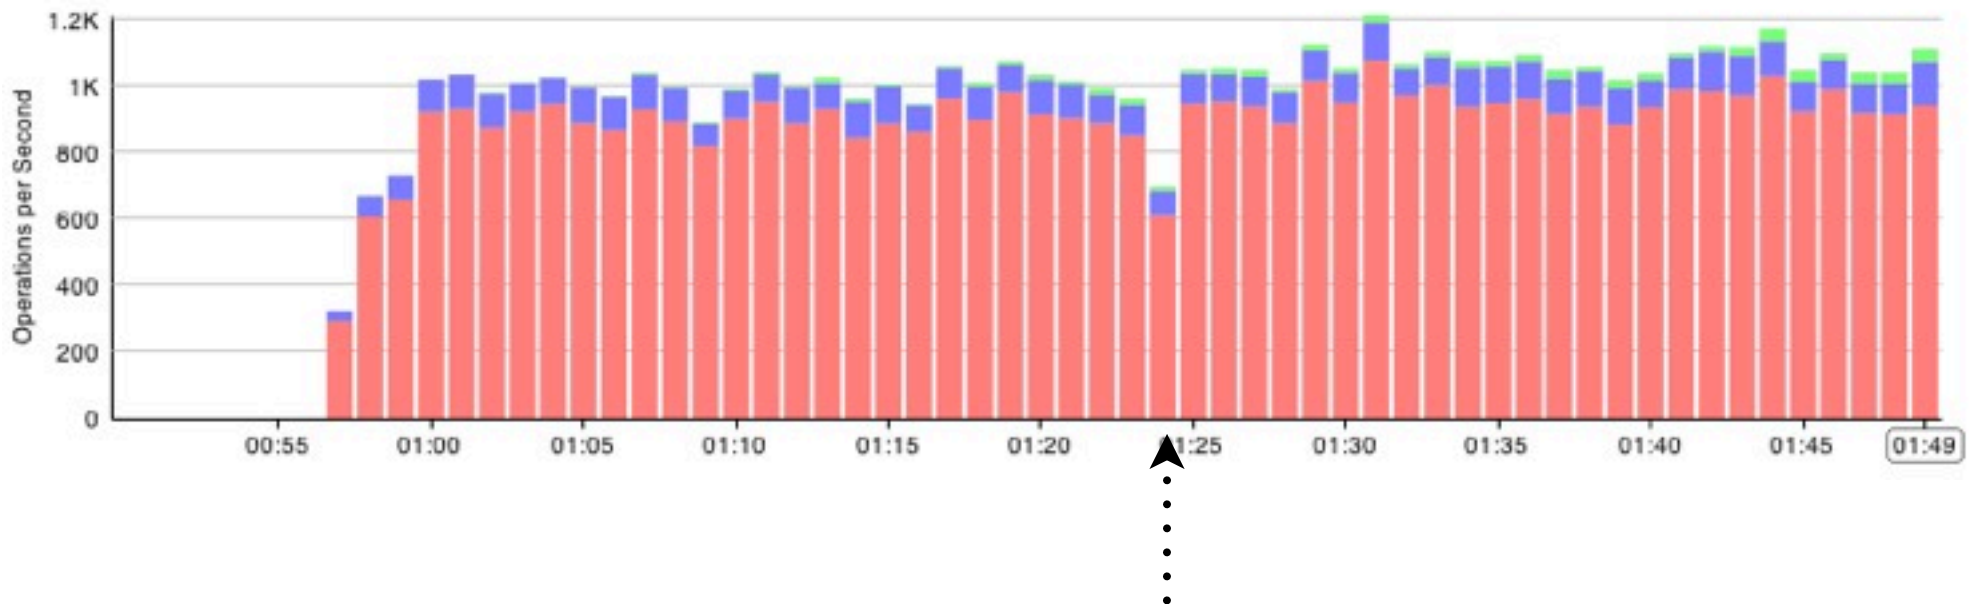
	Memcache Protocol	Memcache Command	Memcache Status	Memcache Key	Operations
■	ASCII	GET	NOT_FOUND	brutis	938.33
■	ASCII	SET	STORED	brutis	130.00
■	ASCII	GET	EXISTS	brutis	37.50





# Application throughput and response time

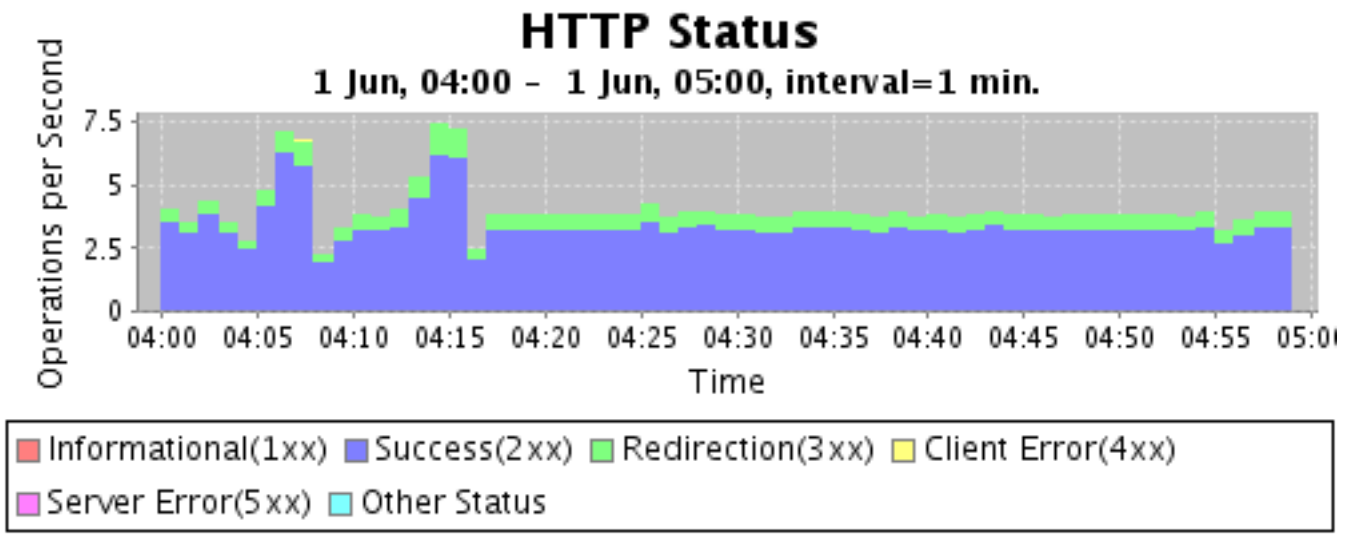
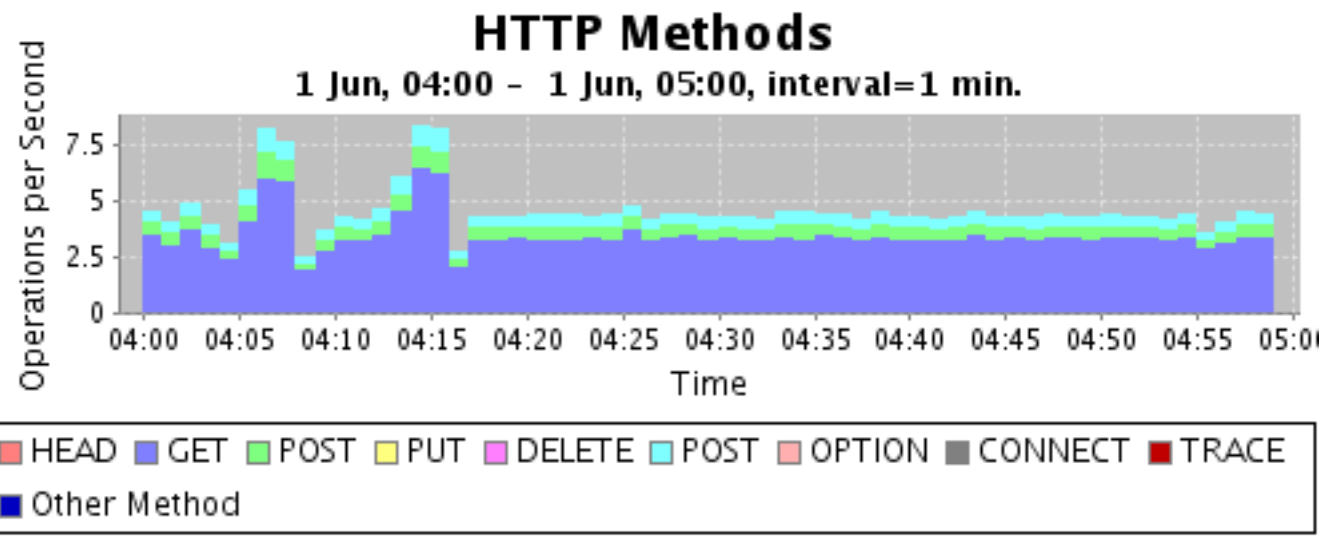
	Memcache Protocol	Memcache Command	Memcache Status	Memcache Key	Operations
■	ASCII	GET	NOT_FOUND	brutis	938.33
■	ASCII	SET	STORED	brutis	130.00
■	ASCII	GET	EXISTS	brutis	37.50



Drop in throughput  
during XenMotion of  
Memcache client



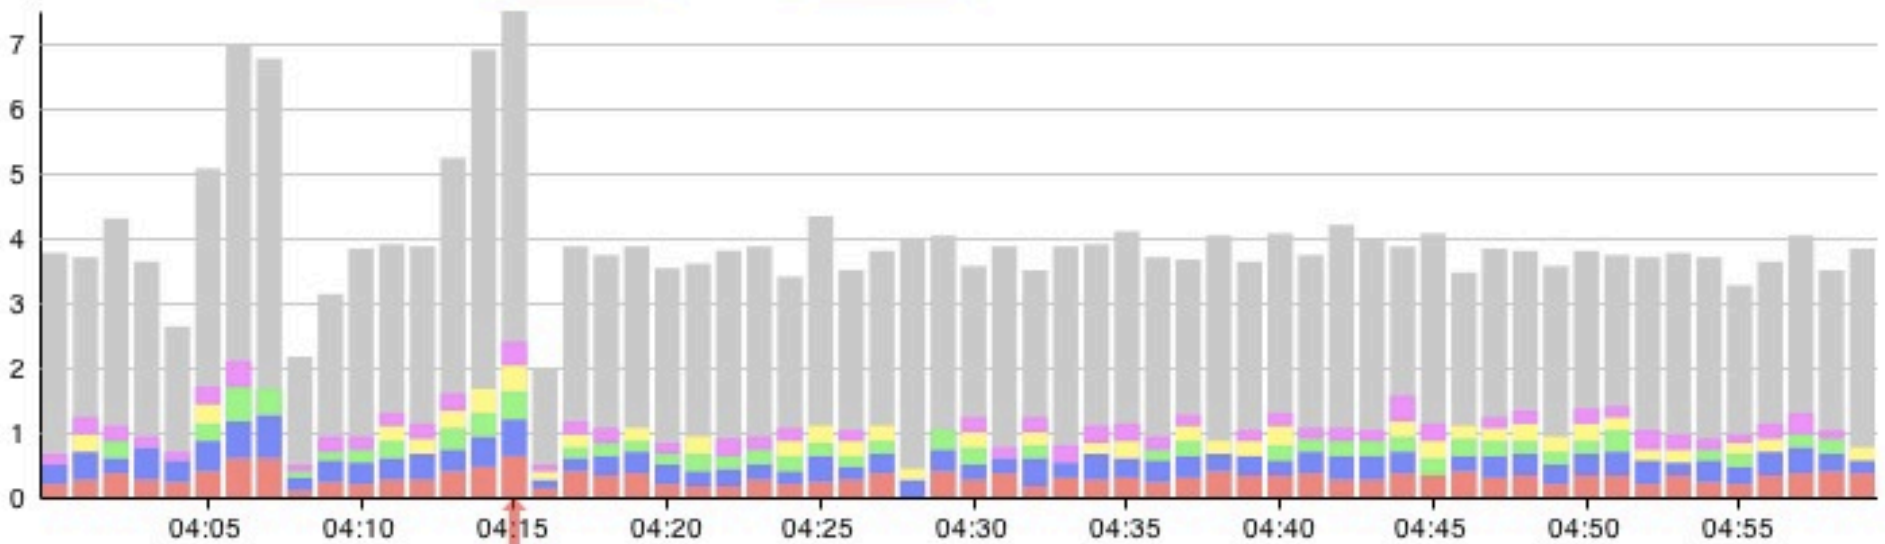
# sFlow-APPLICATION example: transactions





# sFlow-APPLICATION example: transaction detail

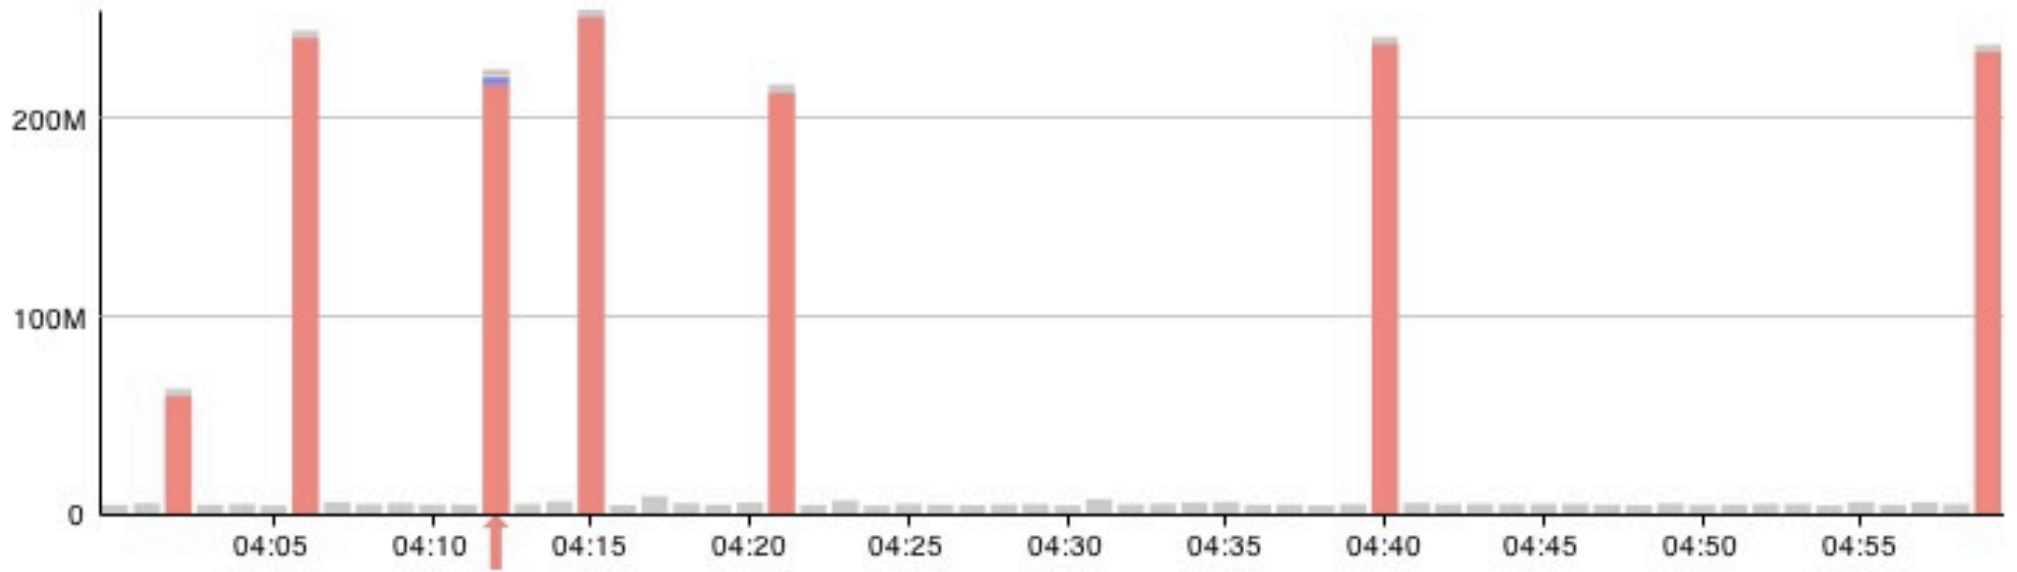
	URI Host	URI Path	Operations
	<a href="http://goku.brightcove.com">goku.brightcove.com</a>	<a href="/1pix.gif">/1pix.gif</a>	0.63
	<a href="http://www.miitv.com">www.miitv.com</a>	<a href="/">/</a>	0.57
	<a href="http://www.miitv.com">www.miitv.com</a>	<a href="/module/setup.php">/module/setup.php</a>	0.43
	<a href="http://c.brightcove.com">c.brightcove.com</a>	<a href="/services/messagebroker/amf">/services/messagebroker/amf</a>	0.40
	<a href="http://www.miitv.com">www.miitv.com</a>	<a href="/adspace.php">/adspace.php</a>	0.37





# sFlow-APPLICATION example: latency

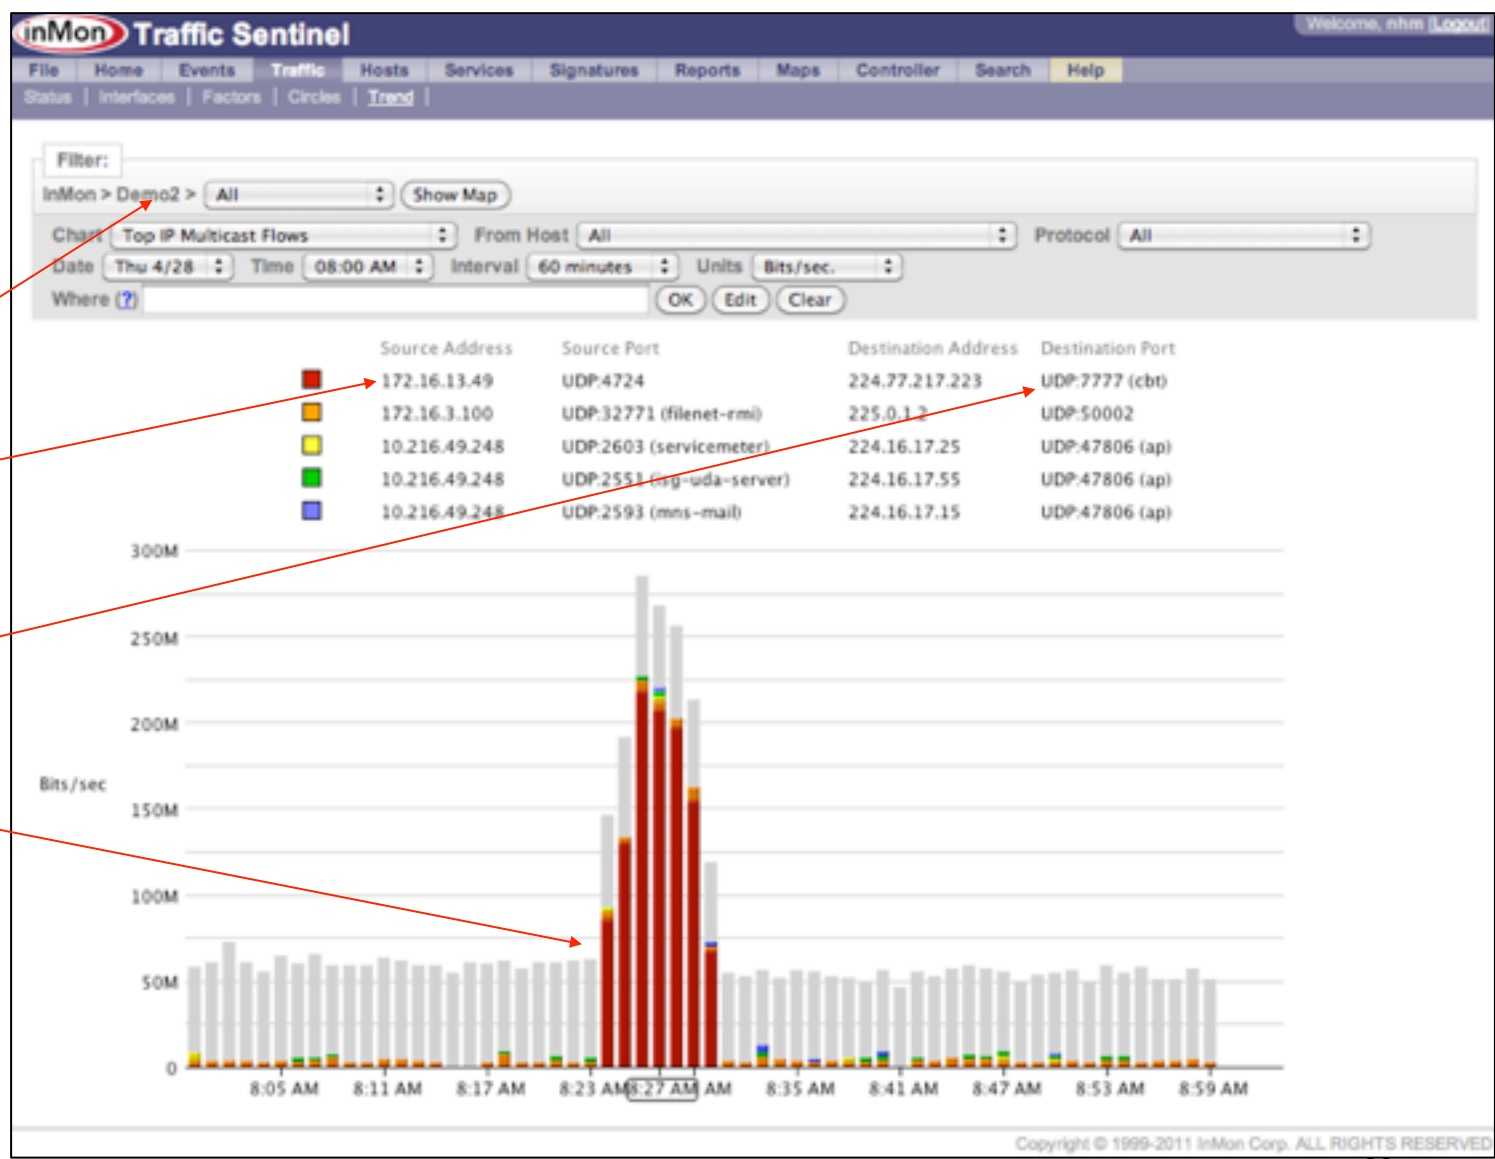
	URI Host	URI Path	Duration
	<a href="http://talkgadget.google.com">talkgadget.google.com</a>	<a href="/talkgadget/channel/bind">/talkgadget/channel/bind</a>	216.39M
	<a href="http://ds.addthis.com">ds.addthis.com</a>	<a href="/red/psi/sites/www.mitv.com/p.json">/red/psi/sites/www.mitv.com/p.json</a>	3.47M
	<a href="http://c1142172.cdn.cloudfiles.rackspacecloud.com">c1142172.cdn.cloudfiles.rackspacecloud.com</a>	<a href="/100x75/j0c08ui84r.jpg">/100x75/j0c08ui84r.jpg</a>	1.13M
	<a href="http://c1142172.cdn.cloudfiles.rackspacecloud.com">c1142172.cdn.cloudfiles.rackspacecloud.com</a>	<a href="/100x75/eghxushmko.jpg">/100x75/eghxushmko.jpg</a>	1.12M
	<a href="http://c1142172.cdn.cloudfiles.rackspacecloud.com">c1142172.cdn.cloudfiles.rackspacecloud.com</a>	<a href="/100x75/gdsussbdf4.jpg">/100x75/gdsussbdf4.jpg</a>	1.12M



# Why Monitor Everything?

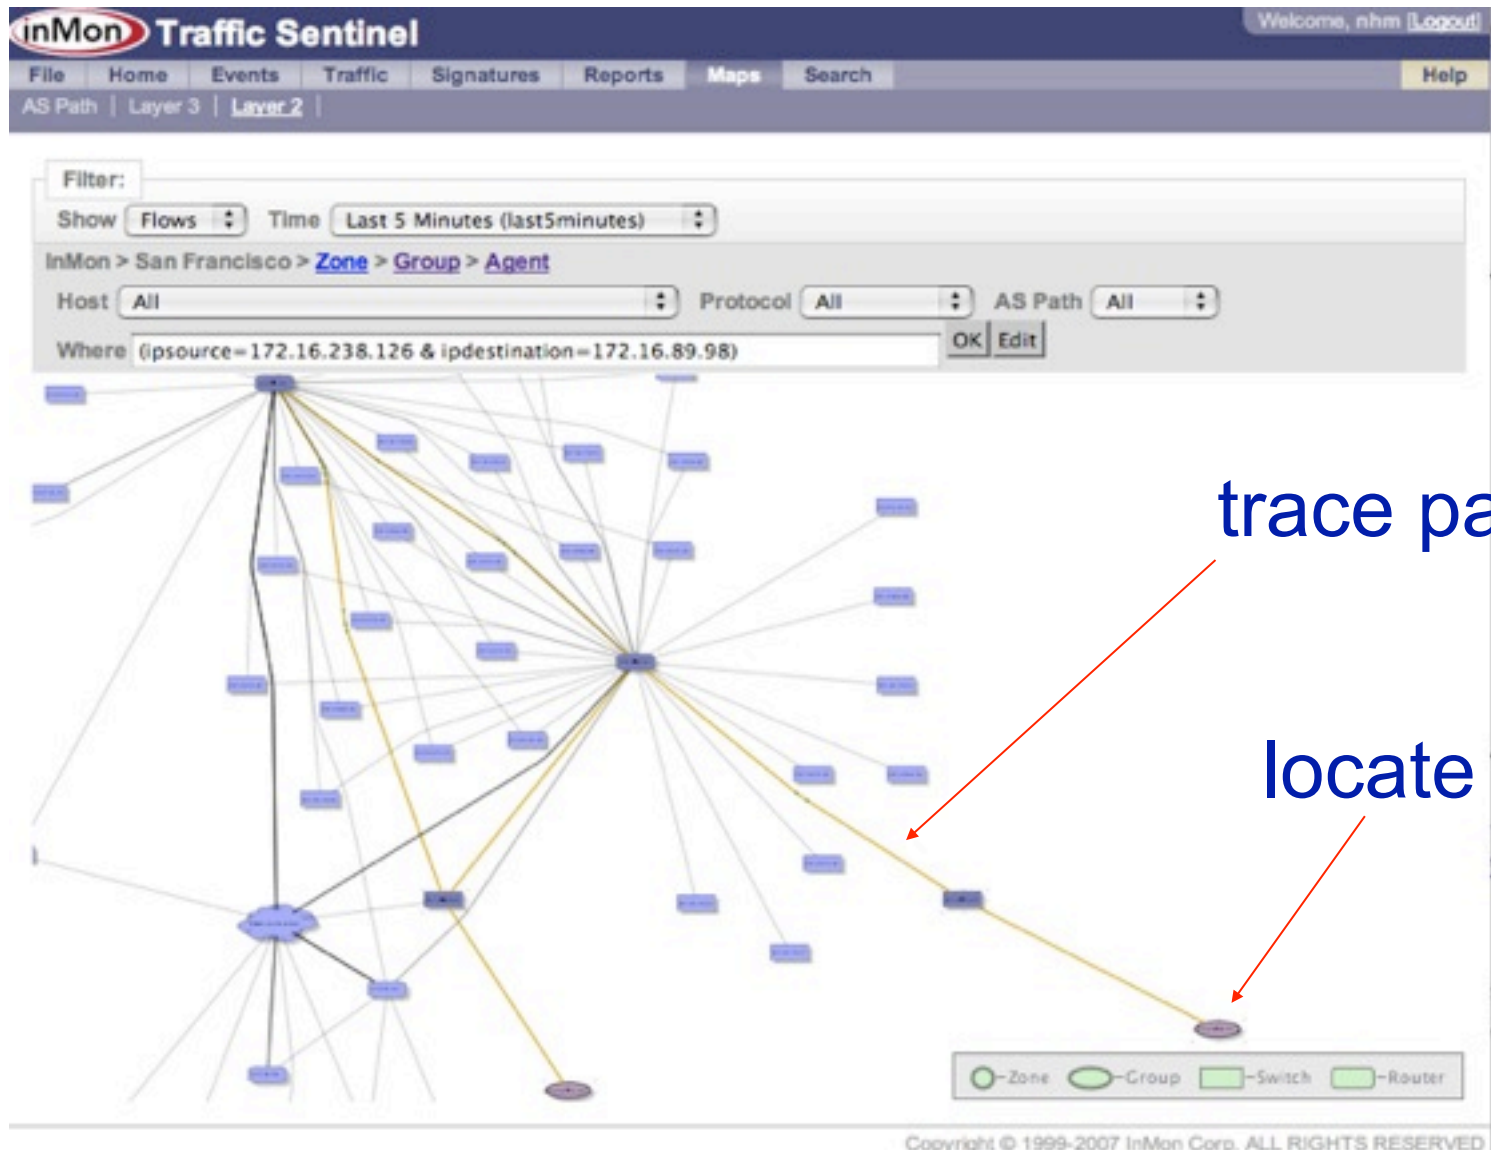
## 1. Troubleshooting - always have context

where  
who  
what  
when



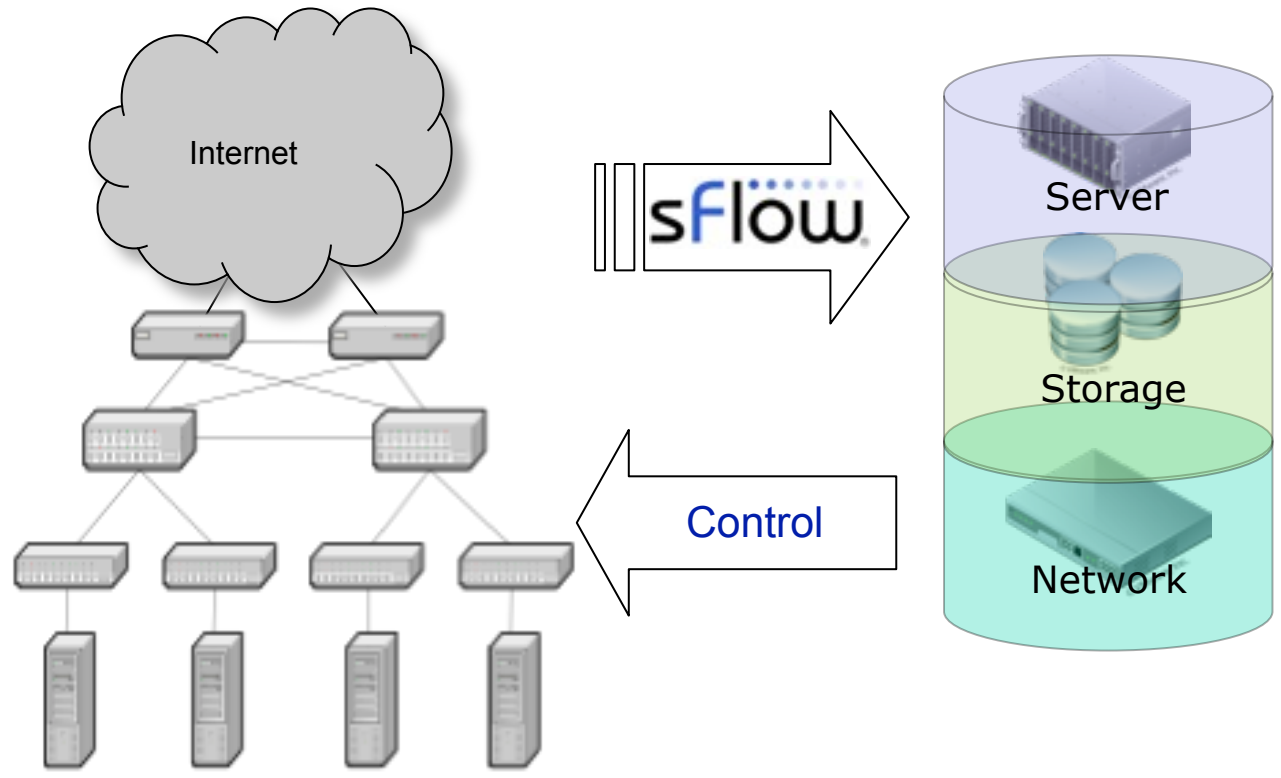
# Why Monitor Everything?

## 1. Troubleshooting - always have context



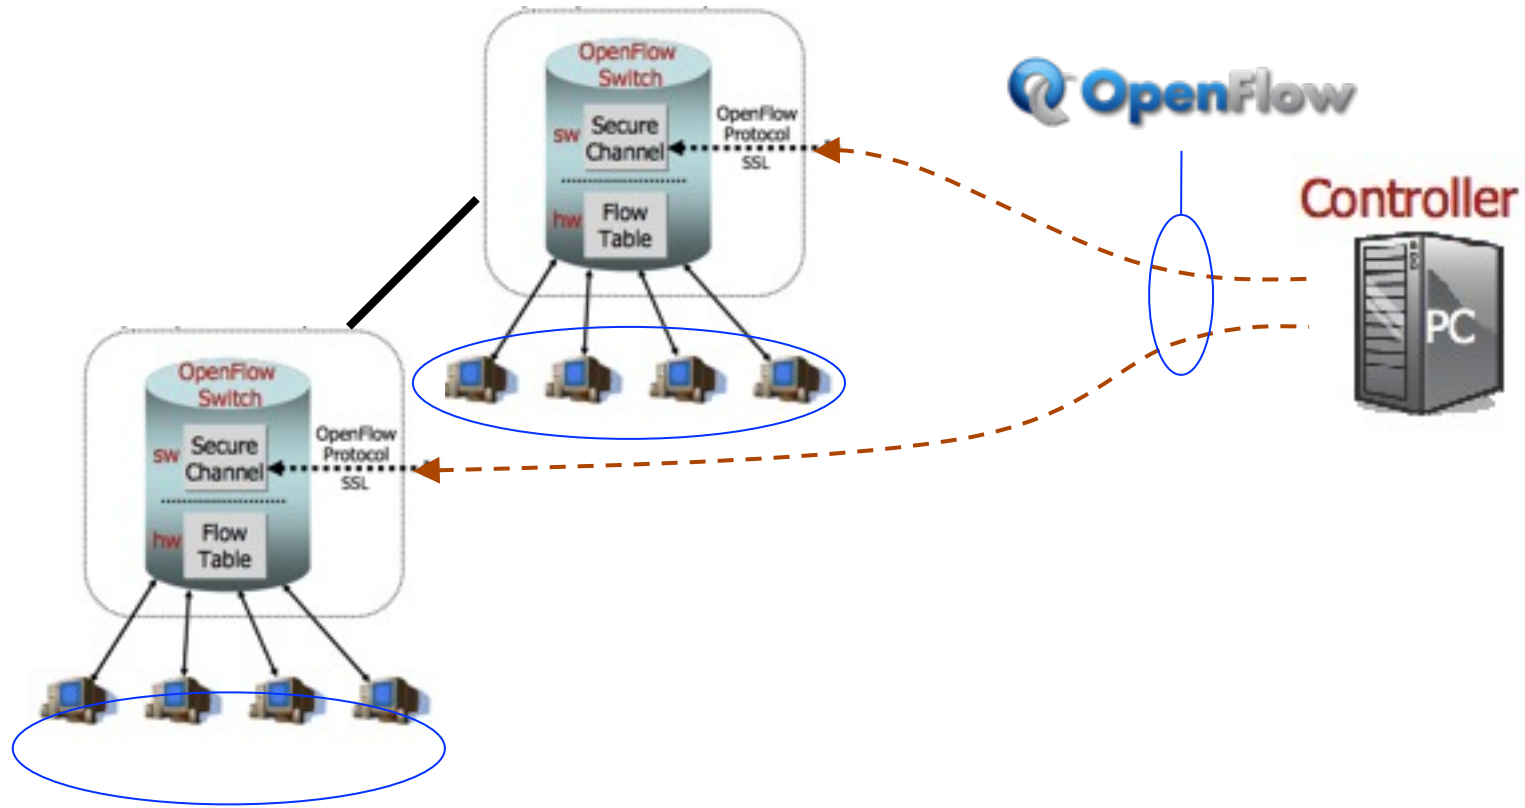
# Why Monitor Everything?

## 2. Put Network and Server teams on same page



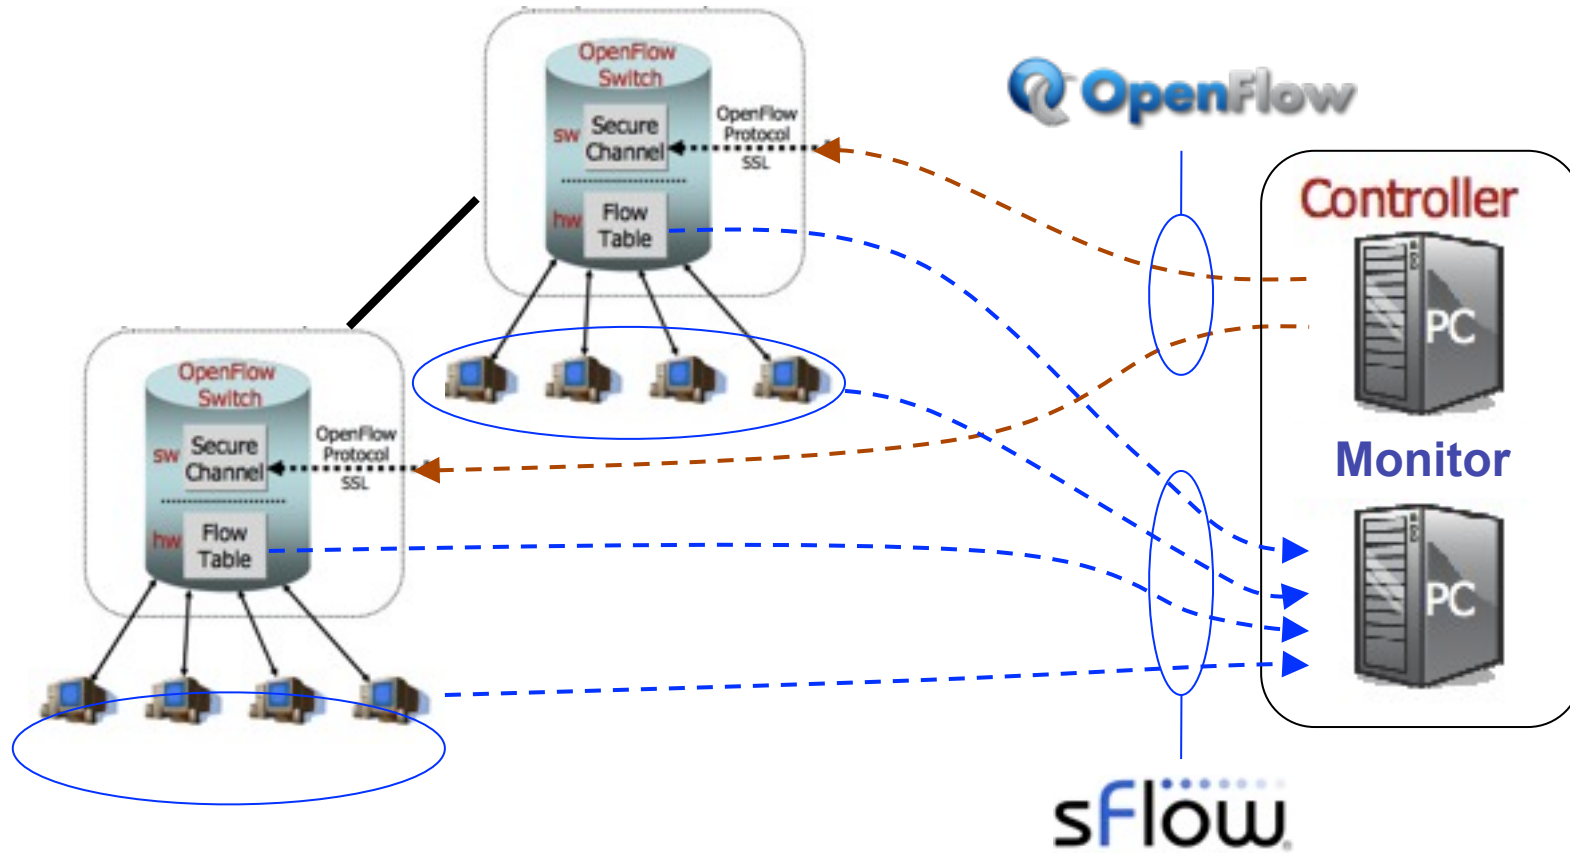
**3. Full “Observability” required for automated control**

# OpenFlow + sFlow = Adaptive Control

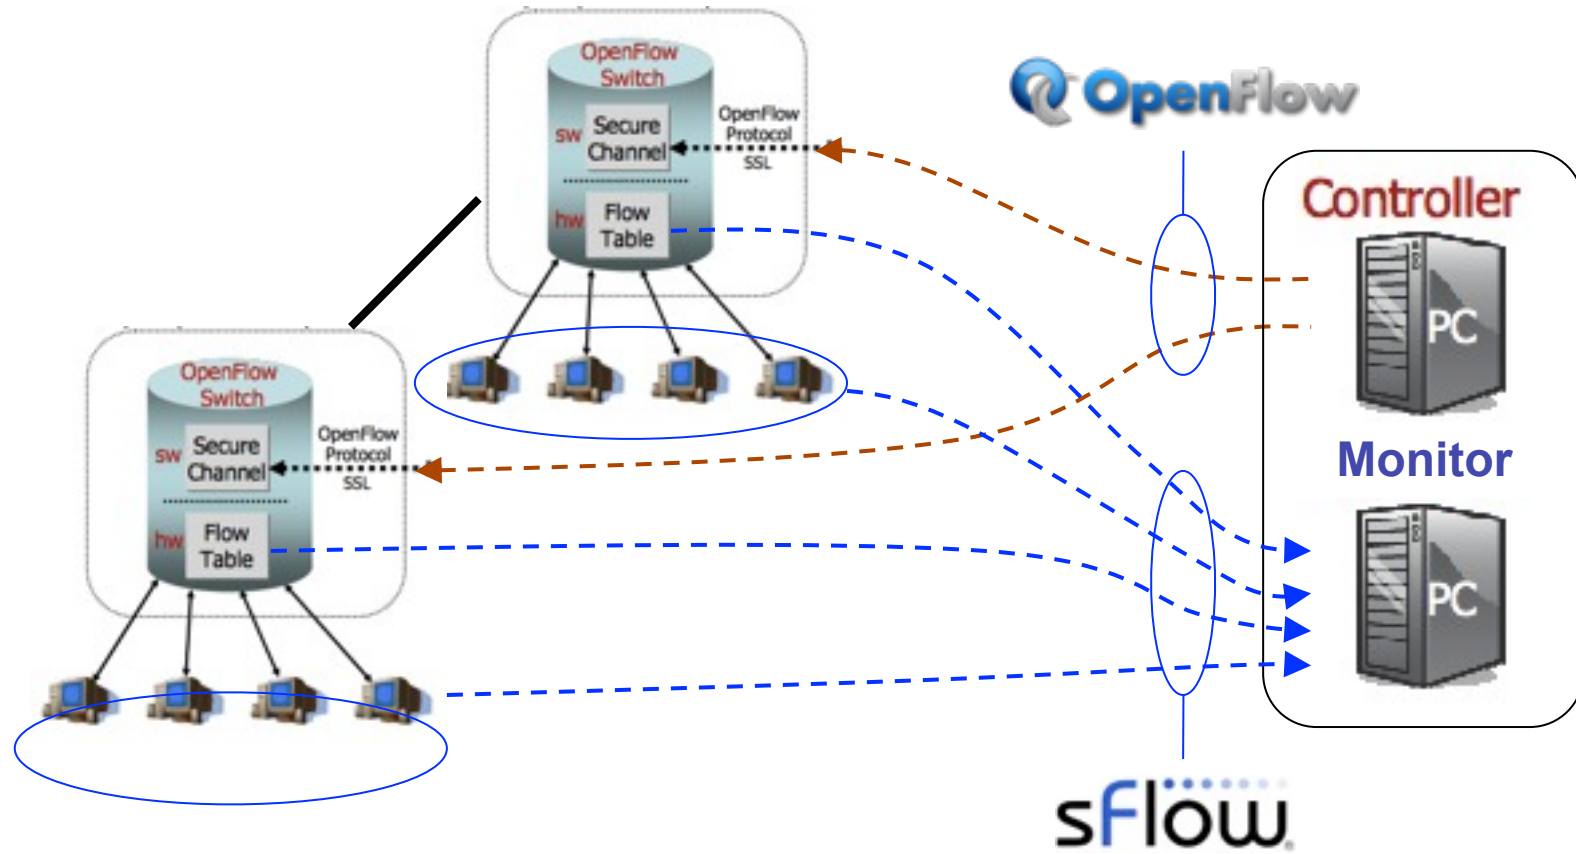




# OpenFlow + sFlow = Adaptive Control

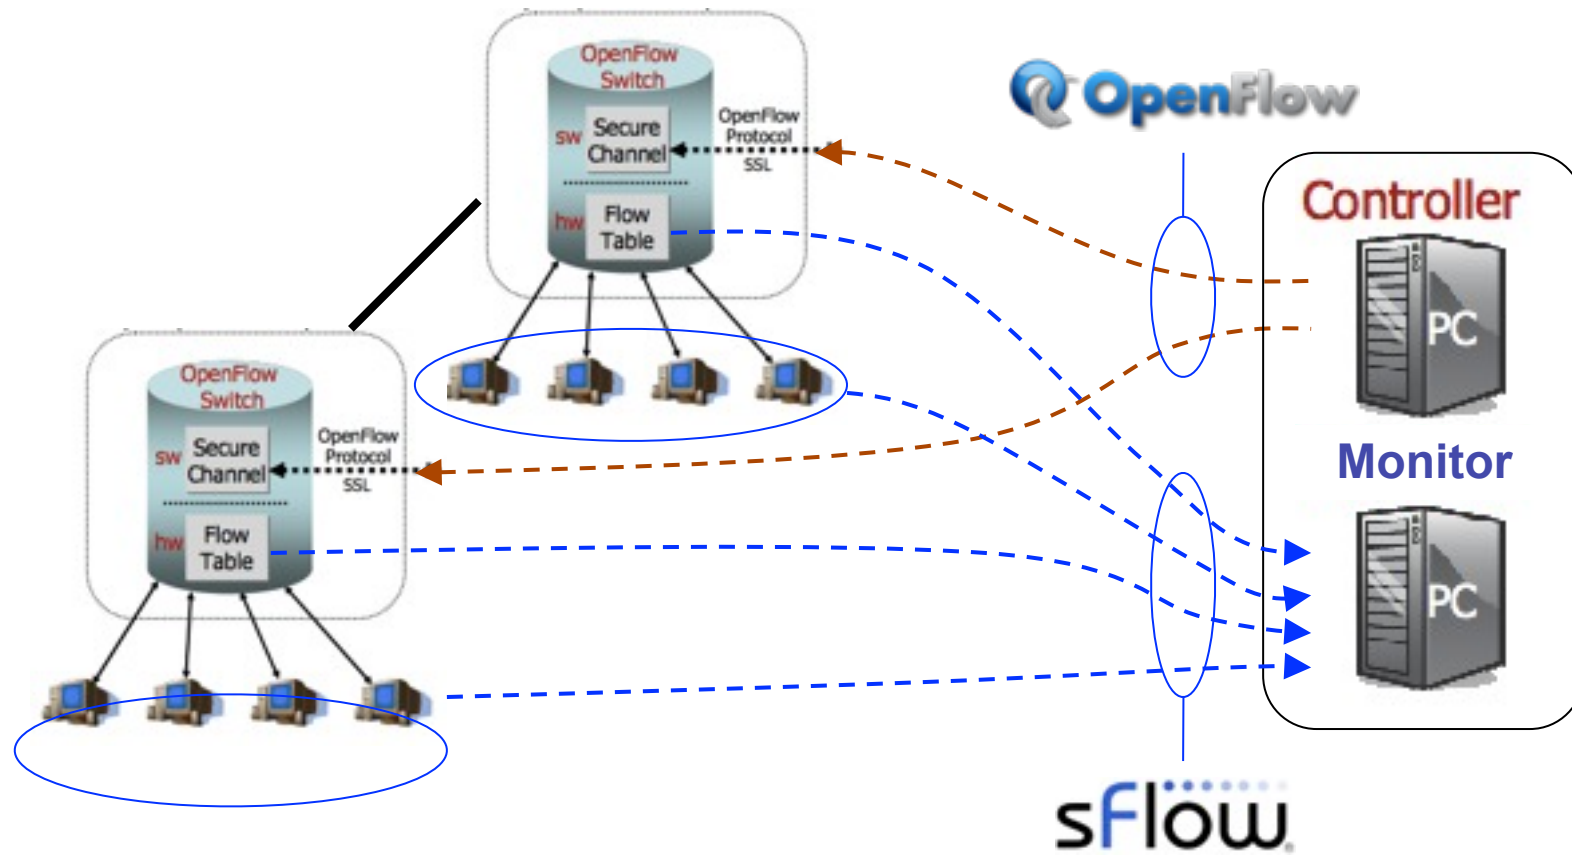


# OpenFlow + sFlow = Adaptive Control

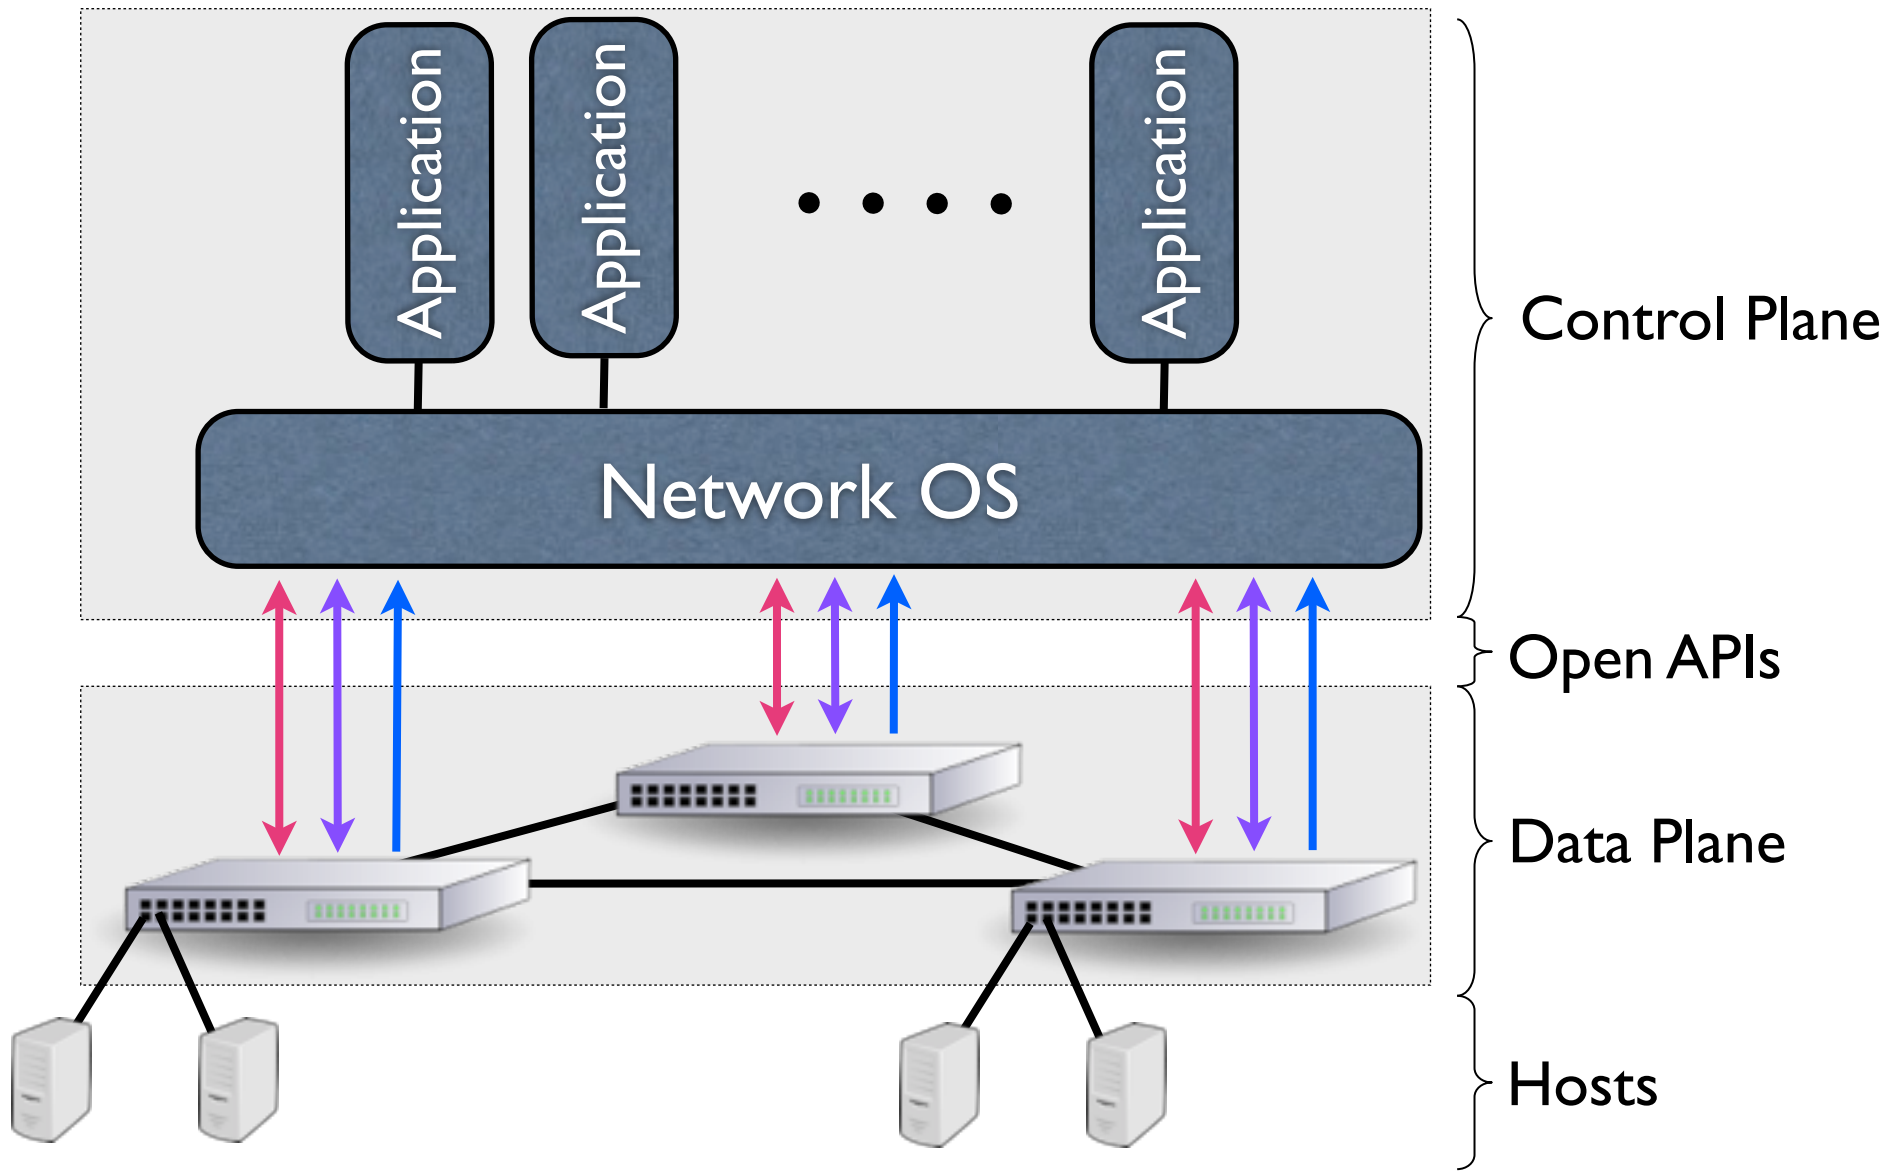


- Detailed, low-latency measurements from sFlow allows OpenFlow controller to adapt network to changing traffic patterns (load balancing, DDoS mitigation etc.).

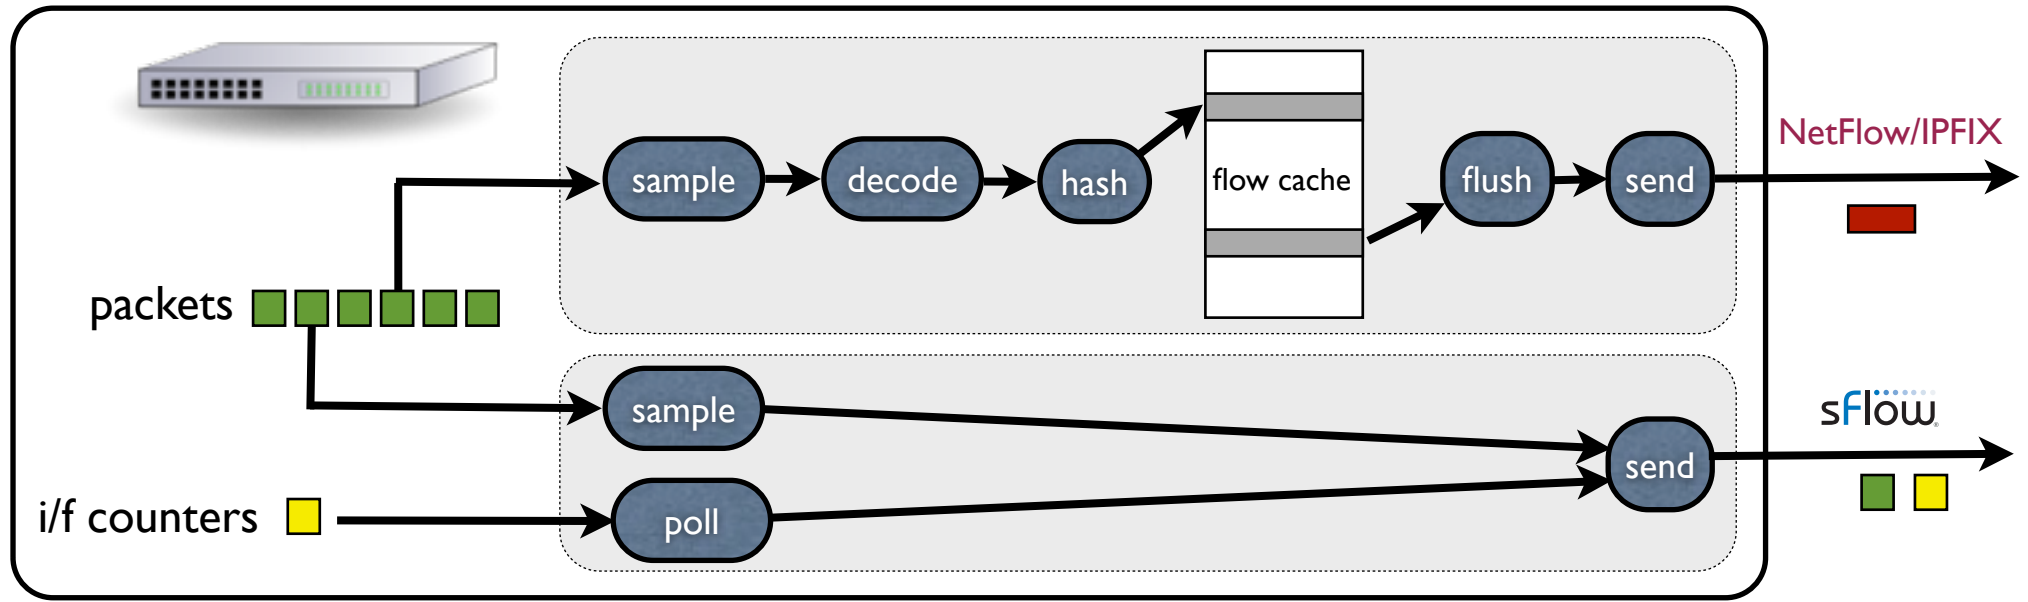
# OpenFlow + sFlow = Adaptive Control



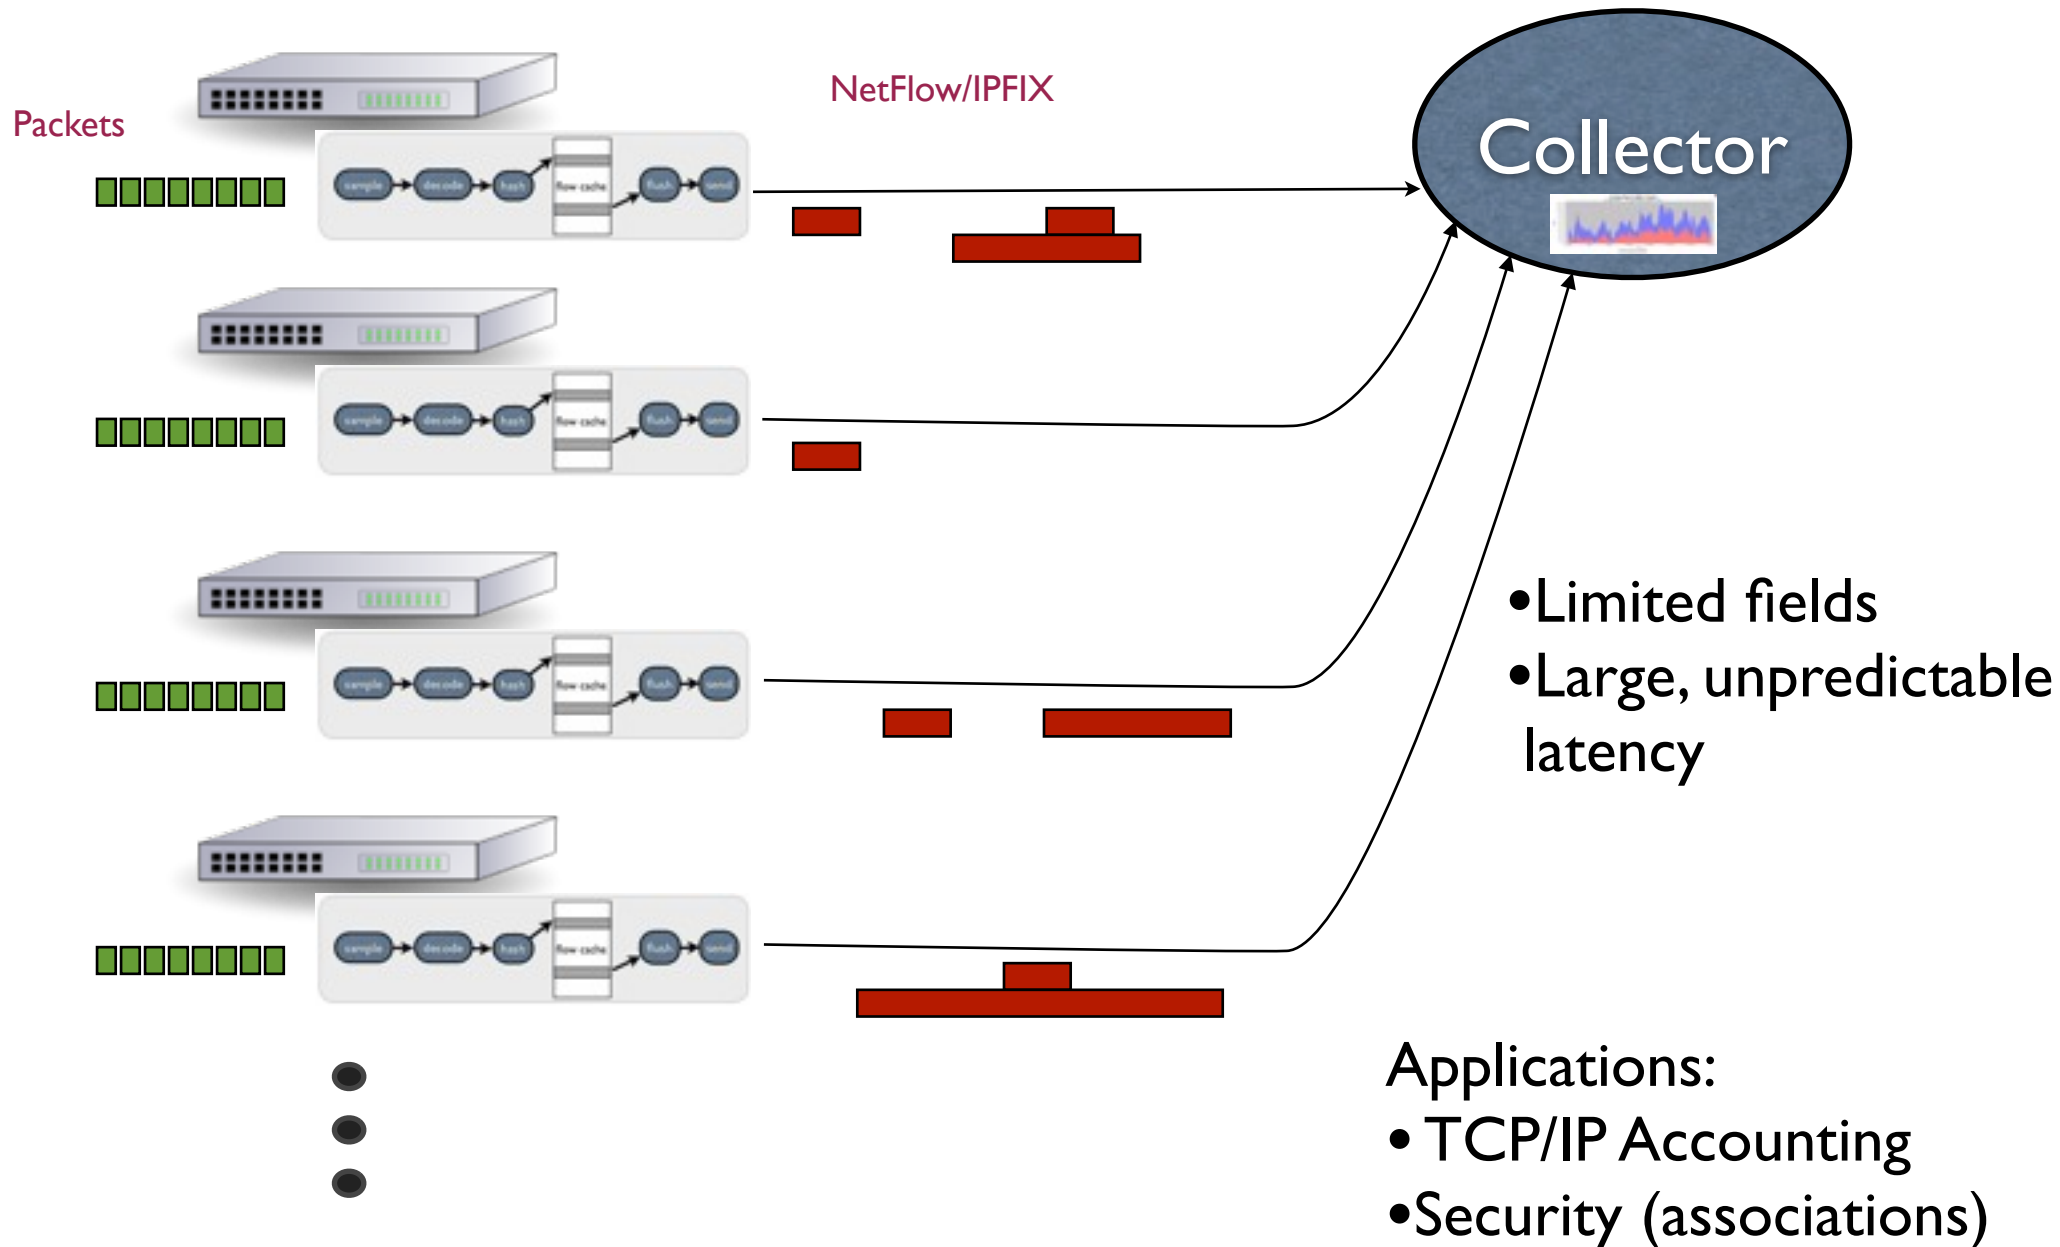
- Detailed, low-latency measurements from sFlow allows OpenFlow controller to adapt network to changing traffic patterns (load balancing, DDoS mitigation etc.).
- OpenFlow can be optimized for efficiency (e.g. by using wildcards), sFlow provides visibility to detect and manage large flows.



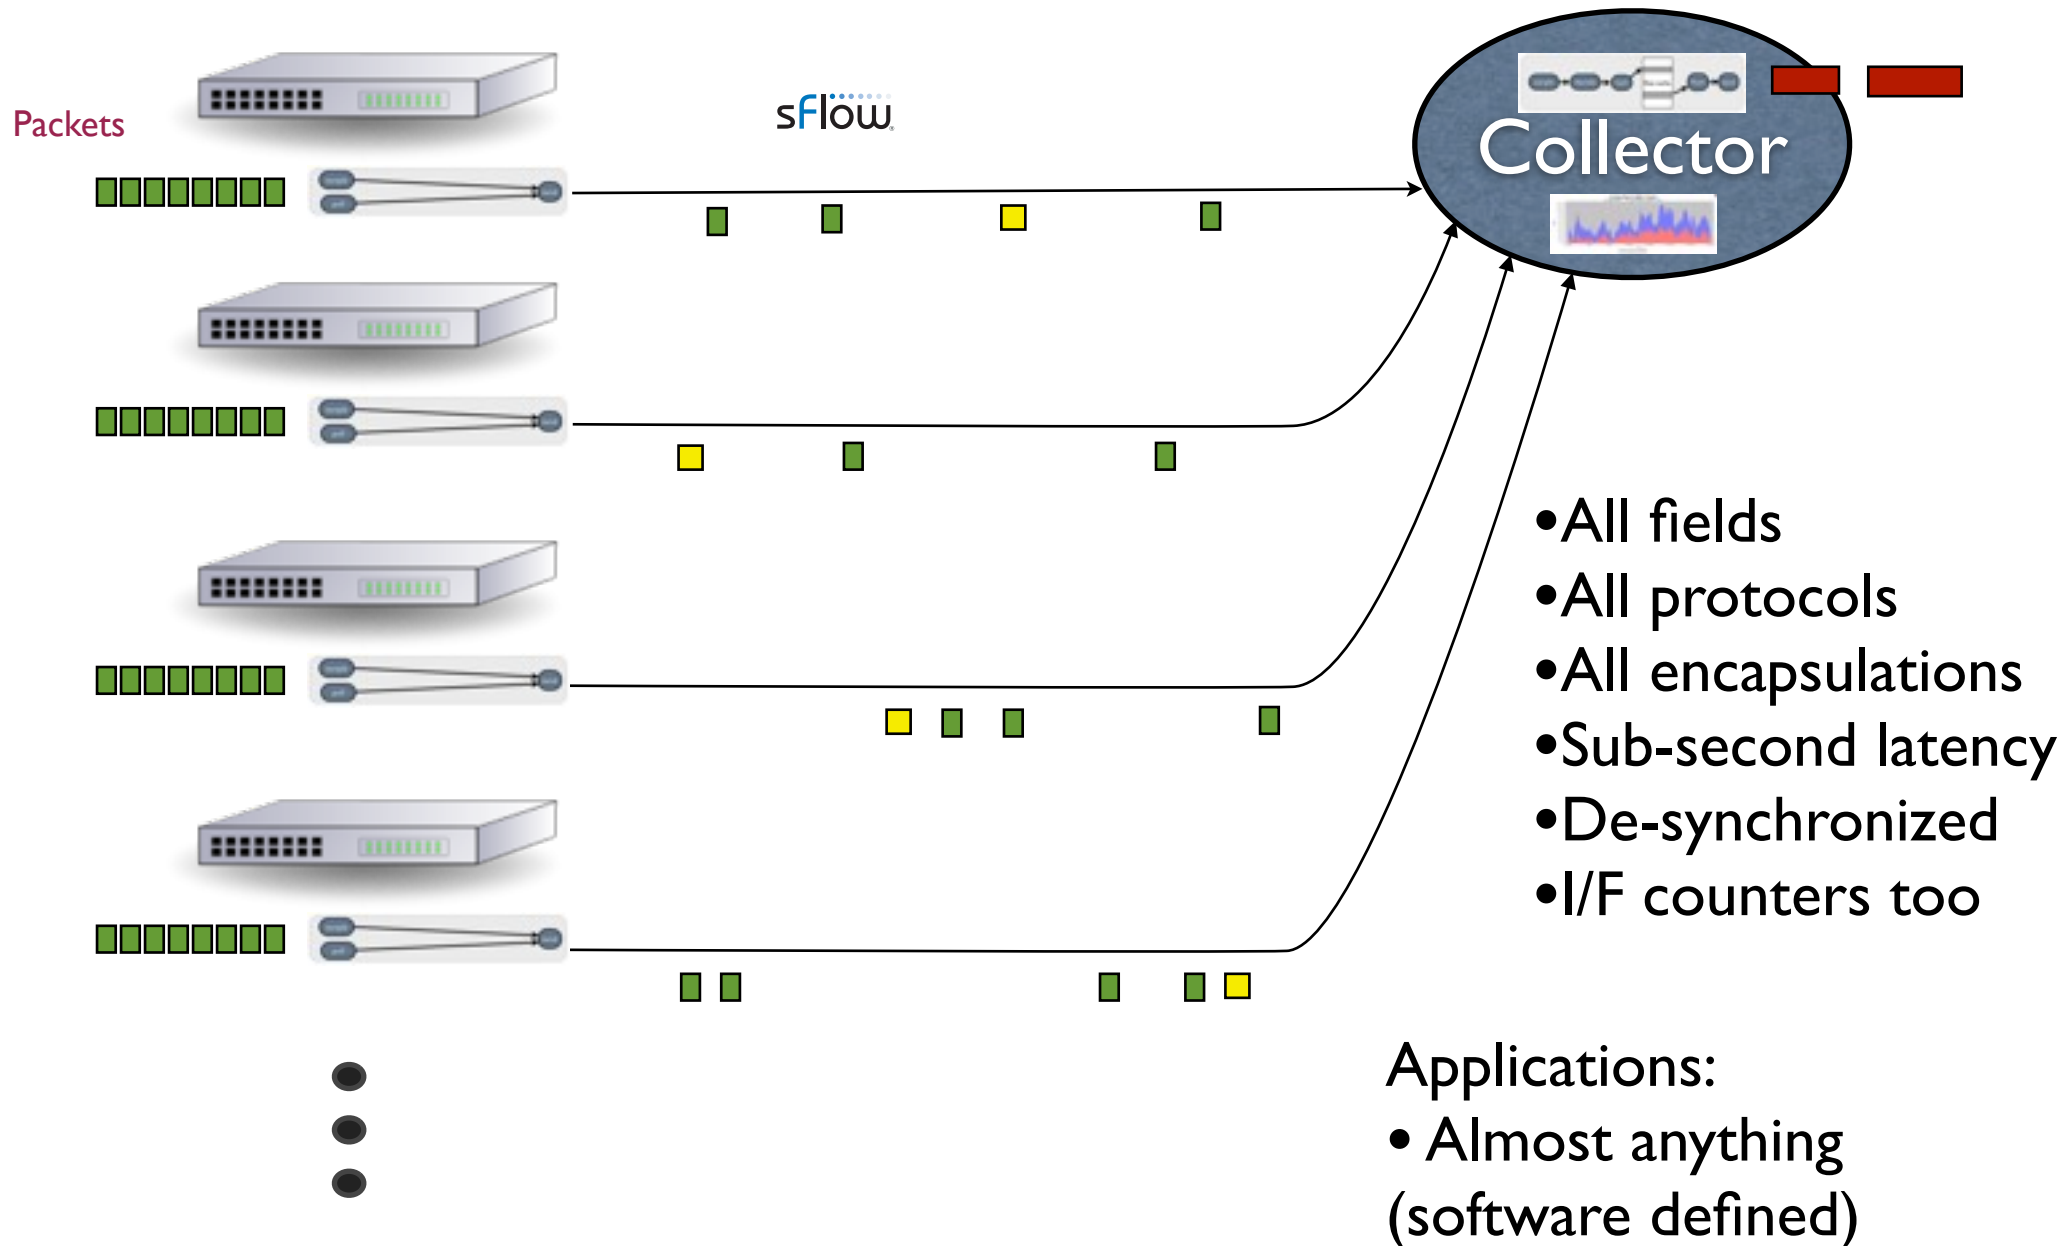
# sFlow vs NetFlow/IPFIX



# sFlow vs NetFlow/IPFIX : system-wide NetFlow/IPFIX



# sFlow vs NetFlow/IPFIX : system-wide sFlow

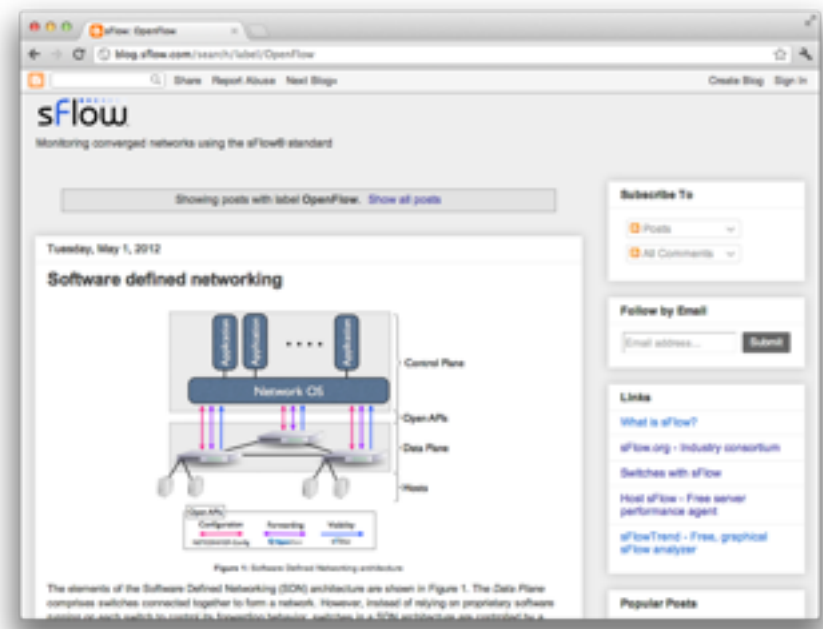




# More Information



[host-sflow.sourceforge.net](http://host-sflow.sourceforge.net)



[blog.sflow.com](http://blog.sflow.com)



[sflow.org](http://sflow.org)