sFlow®

Data Network visibility and control

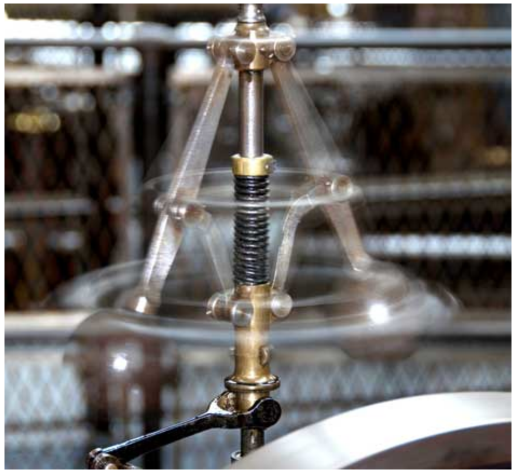

"You can't control what you can't measure"
Tom DeMarco

InMon Corp.

- Inventors of sFlow
- Leading supplier of software solutions that take advantage of embedded traffic monitoring
- Partner with switch and router vendors to deliver effective traffic management solutions
- Based in San Francisco, California
- Worldwide customer deployments
  - Enterprise, Education, Government, Media, ISP
  - Global presence through reseller network

inMon sFlow

Example customers

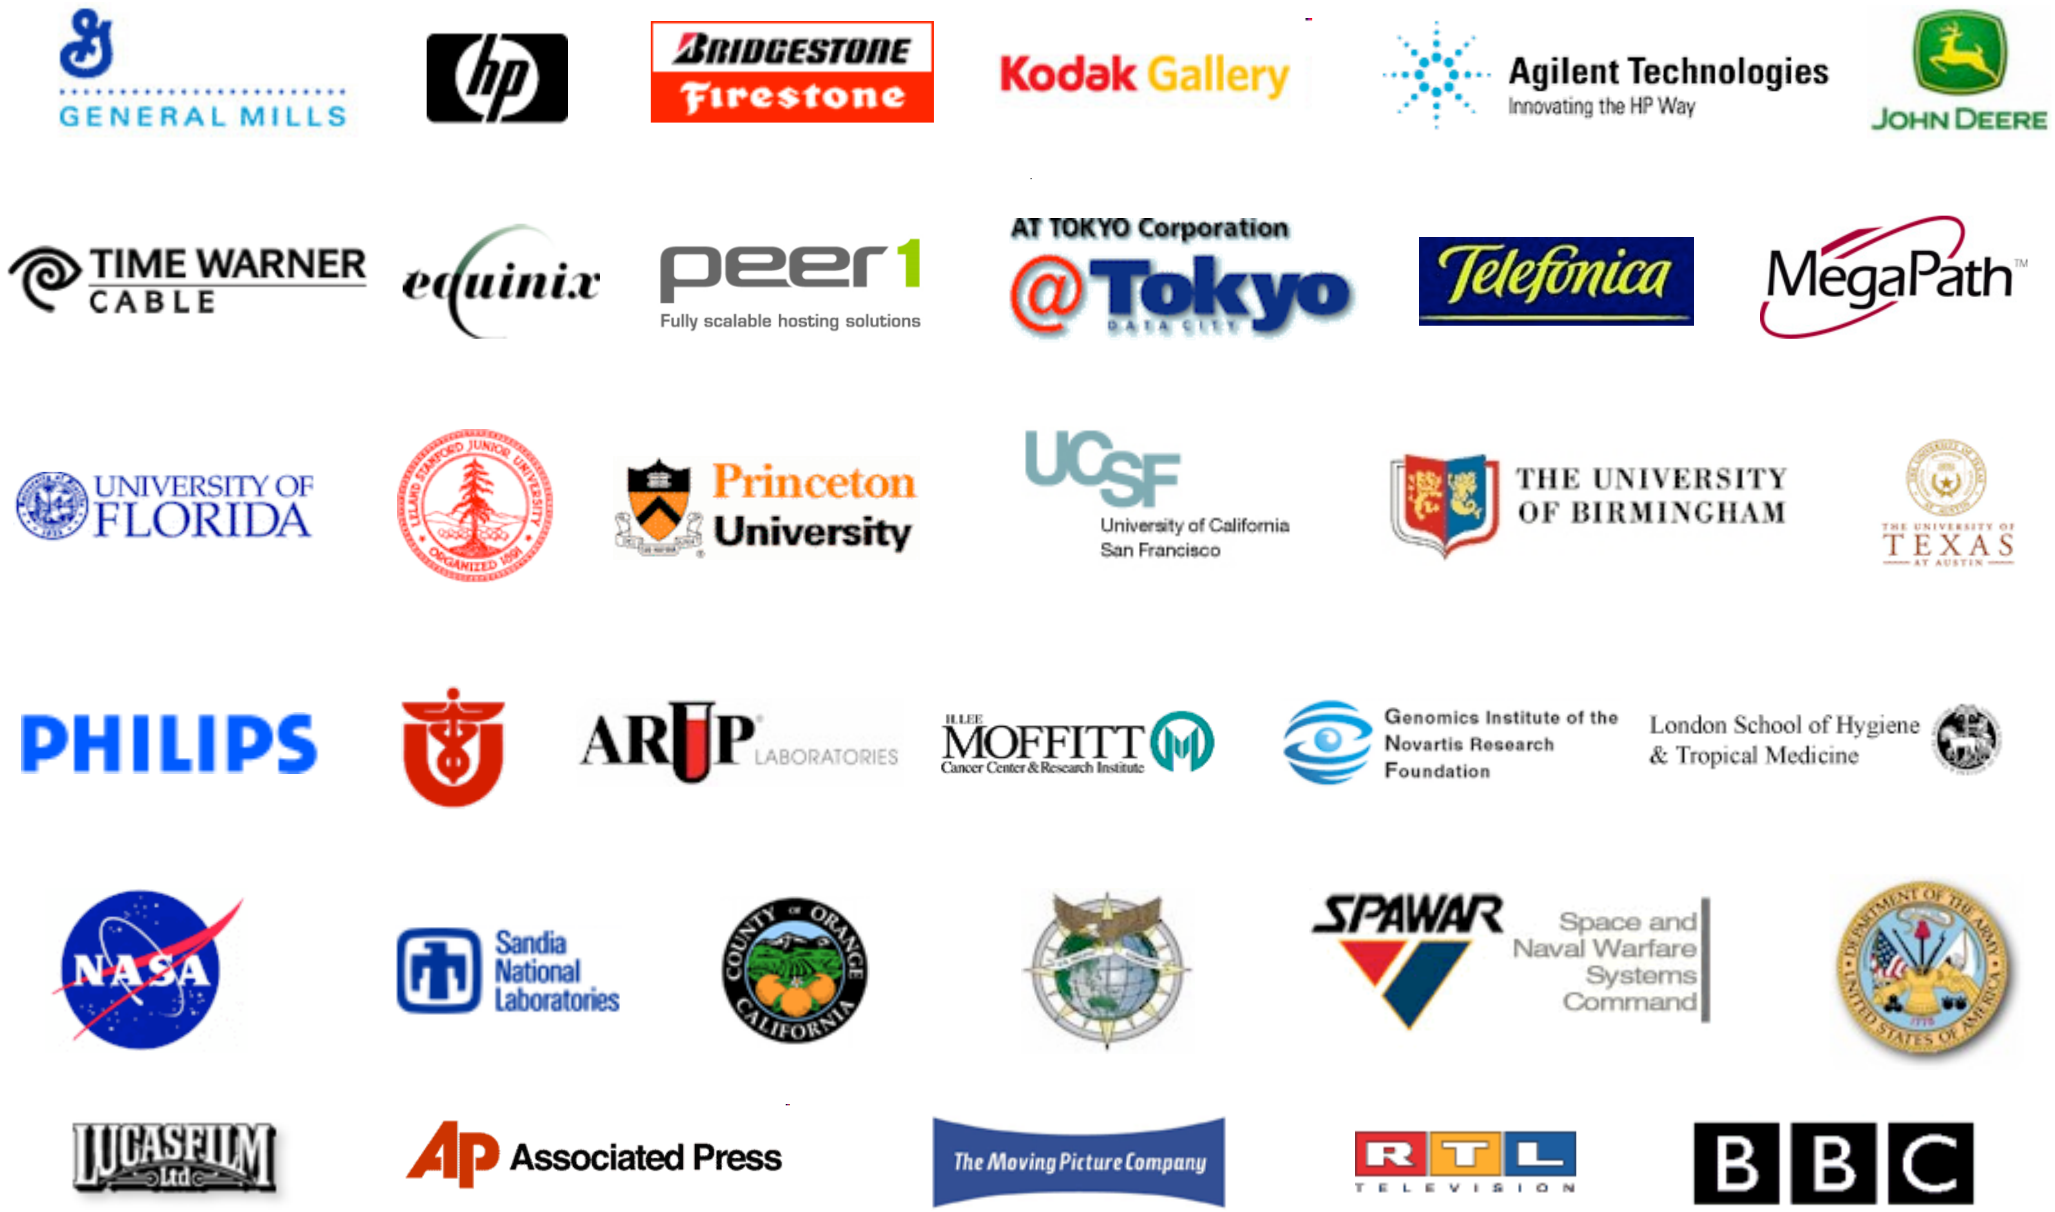

Not Monitoring is Expensive

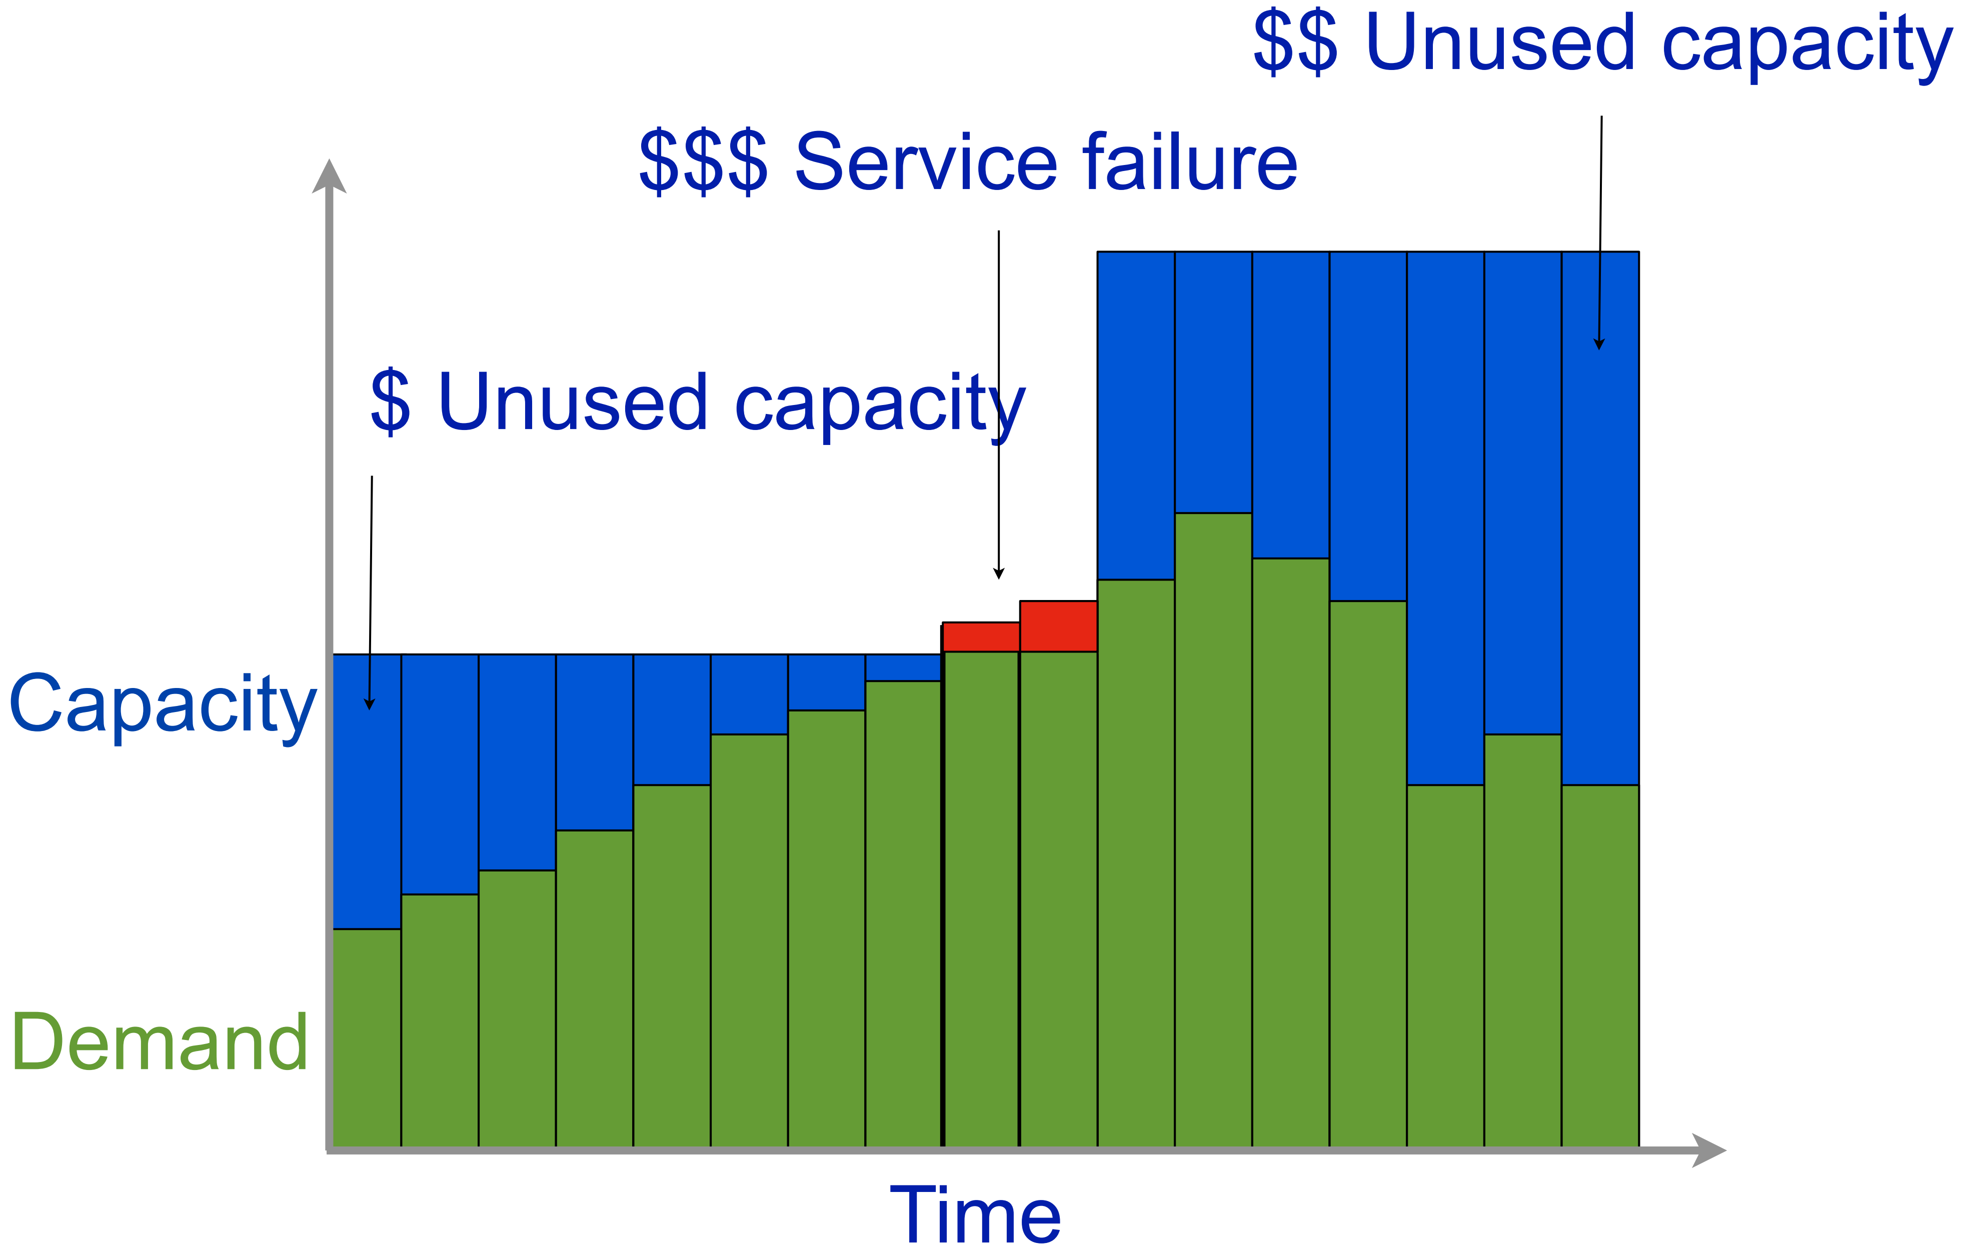

$$ Unused capacity

$$$ Service failure

$ Unused capacity

Capacity

Demand

Time

Measurement Saves Money

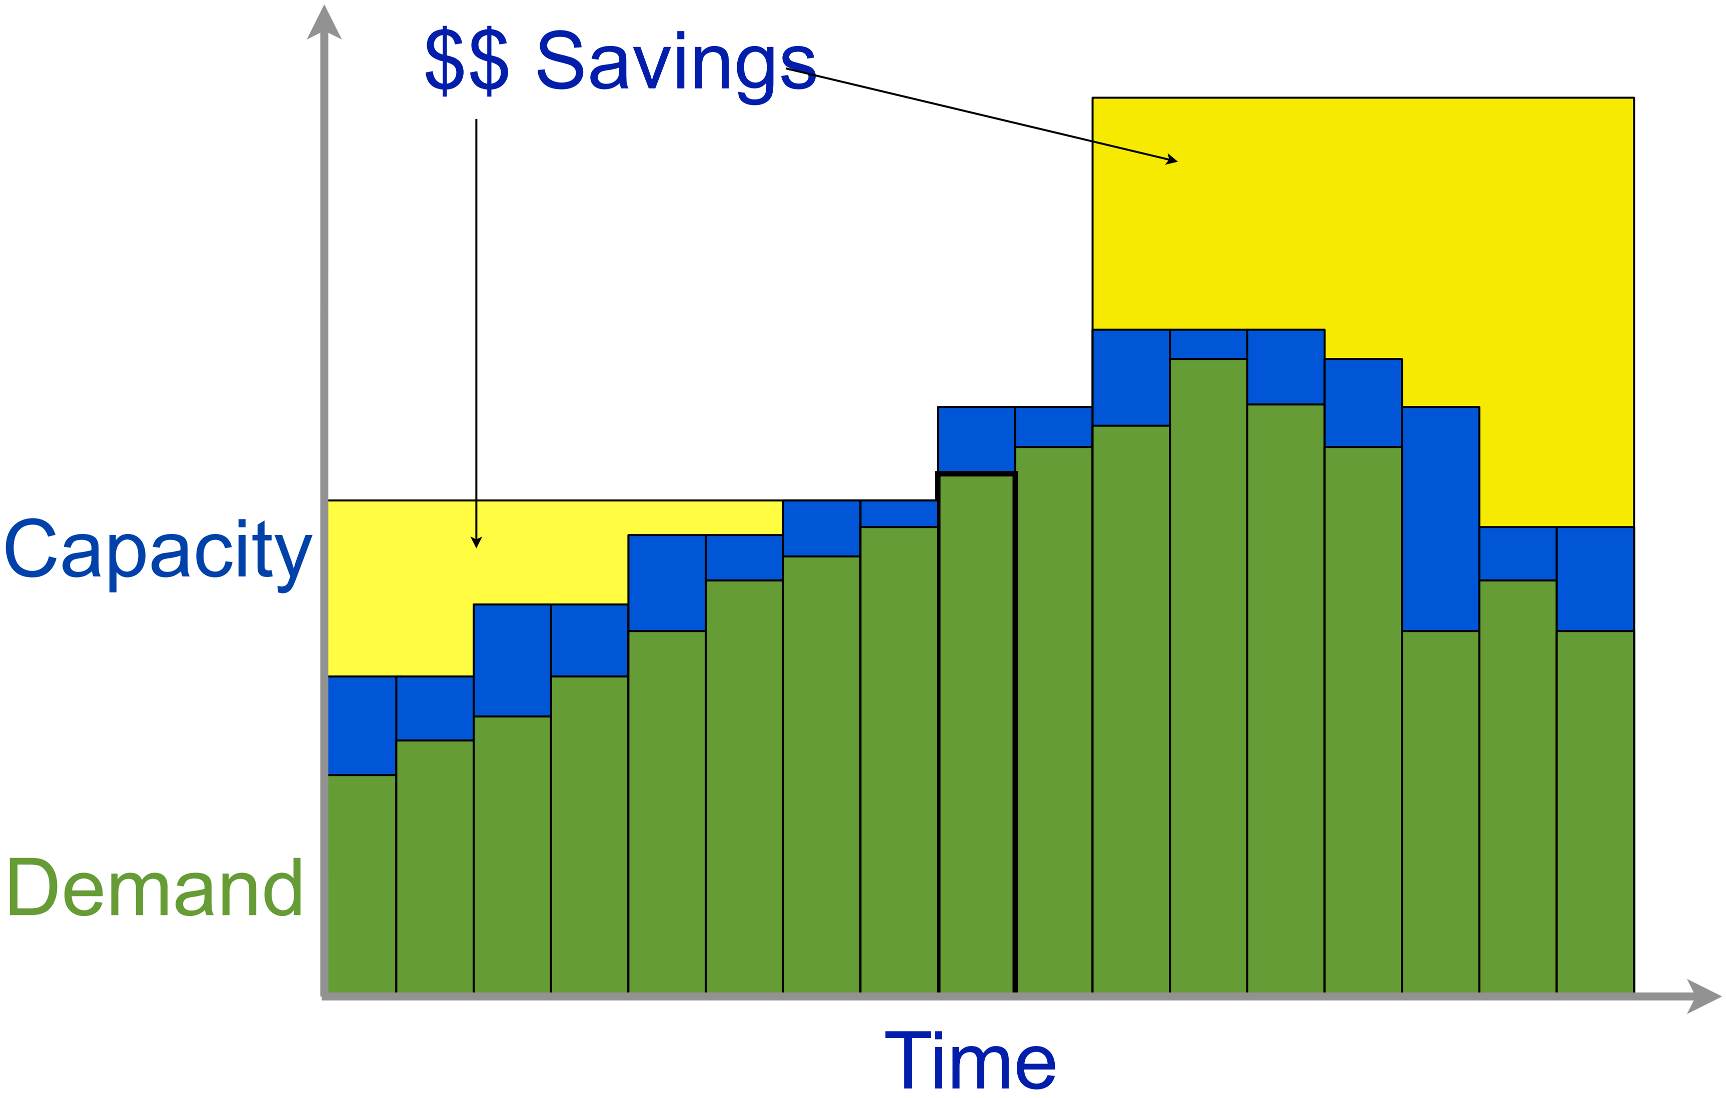

Network Visibility: sFlow

http://www.sflow.org

Always on

- Continuous monitoring of every port
- Robust under all conditions

Complete visibility

- All devices = L2-L7 flows end-end
- Detailed real-time and historical data

Cost effective

- Embedded in every port

Scalable

- Measures traffic flows on all ports
- Effective even at 100Gbs speeds
- Does not impact network performance

hp

BROCADE

IBM

DELL

H3C
IToIP Solutions Expert

Juniper NETWORKS

BLADE
NETWORK TECHNOLOGIES

3com

ARISTA

FORCE10

HITACHI
Inspire the Next

NEC Empowered by Innovation

AlaxalA

Alcatel·Lucent

D-Link
Building Networks for People

extreme networks
GO BEYOND

Allied Telesis

enterasys
Secure Networks™

FORTINET

VYATTA
Open Networking

CITRIX | XenServer™

Xen Cloud Platform

Standard-export: many collectors...

More than 30
commercial sFlow
collector implementations
http://sflow.org/products/collectors.php

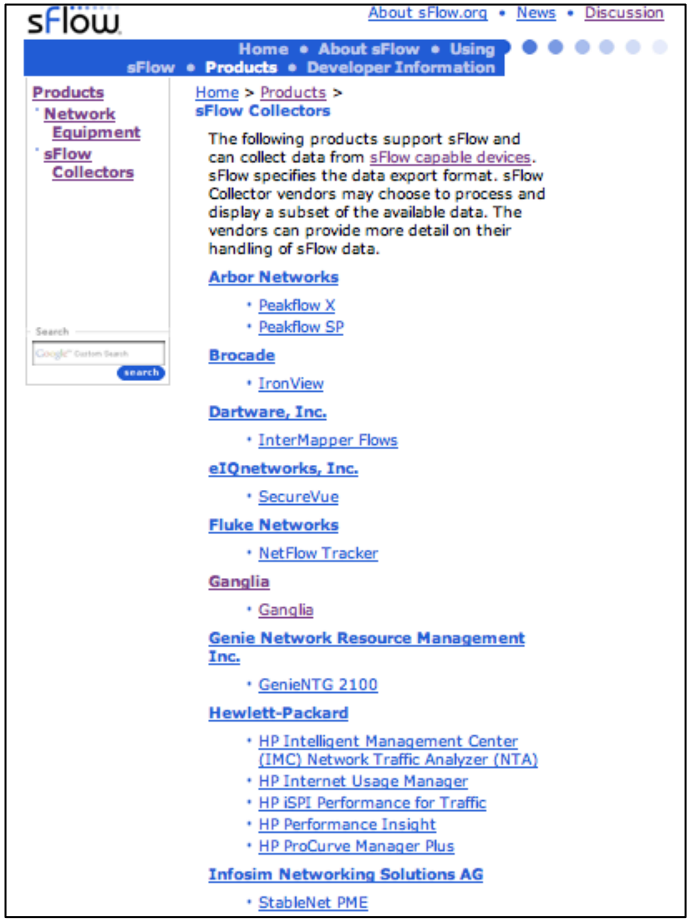

About sFlow.org • News • Discussion

Home • About sFlow • Using sFlow • Products • Developer Information

Products
- Network Equipment
- sFlow Collectors

Search

search

Home > Products >
sFlow Collectors

The following products support sFlow and can collect data from sFlow capable devices. sFlow specifies the data export format. sFlow Collector vendors may choose to process and display a subset of the available data. The vendors can provide more detail on their handling of sFlow data.

Arbor Networks
- Peakflow X
- Peakflow SP

Brocade
- IronView

Dartware, Inc.
- InterMapper Flows

eIQnetworks, Inc.
- SecureVue

Fluke Networks
- NetFlow Tracker

Ganglia
- Ganglia

Genie Network Resource Management Inc.
- GenieNTG 2100

Hewlett-Packard
- HP Intelligent Management Center (IMC) Network Traffic Analyzer (NTA)
- HP Internet Usage Manager
- HP iSPI Performance for Traffic
- HP Performance Insight
- HP ProCurve Manager Plus

Infosim Networking Solutions AG
- StableNet PME

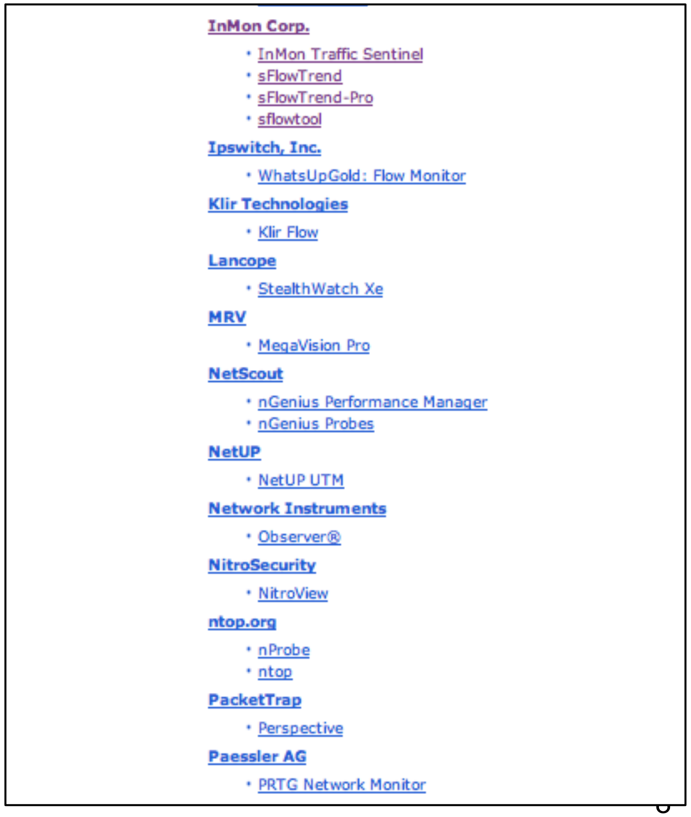

InMon Corp.
- InMon Traffic Sentinel
- sFlowTrend
- sFlowTrend-Pro
- sflowtool

Ipswitch, Inc.
- WhatsUpGold: Flow Monitor

Klir Technologies
- Klir Flow

Lancope
- StealthWatch Xe

MRV
- MegaVision Pro

NetScout
- nGenius Performance Manager
- nGenius Probes

NetUP
- NetUP UTM

Network Instruments
- Observer®

NitroSecurity
- NitroView

ntop.org
- nProbe
- ntop

PacketTrap
- Perspective

Paessler AG
- PRTG Network Monitor

sFlow replaces counter polling

- sFlow agent automatically pushes full set of SNMP ifTable counters[1]
- 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, 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

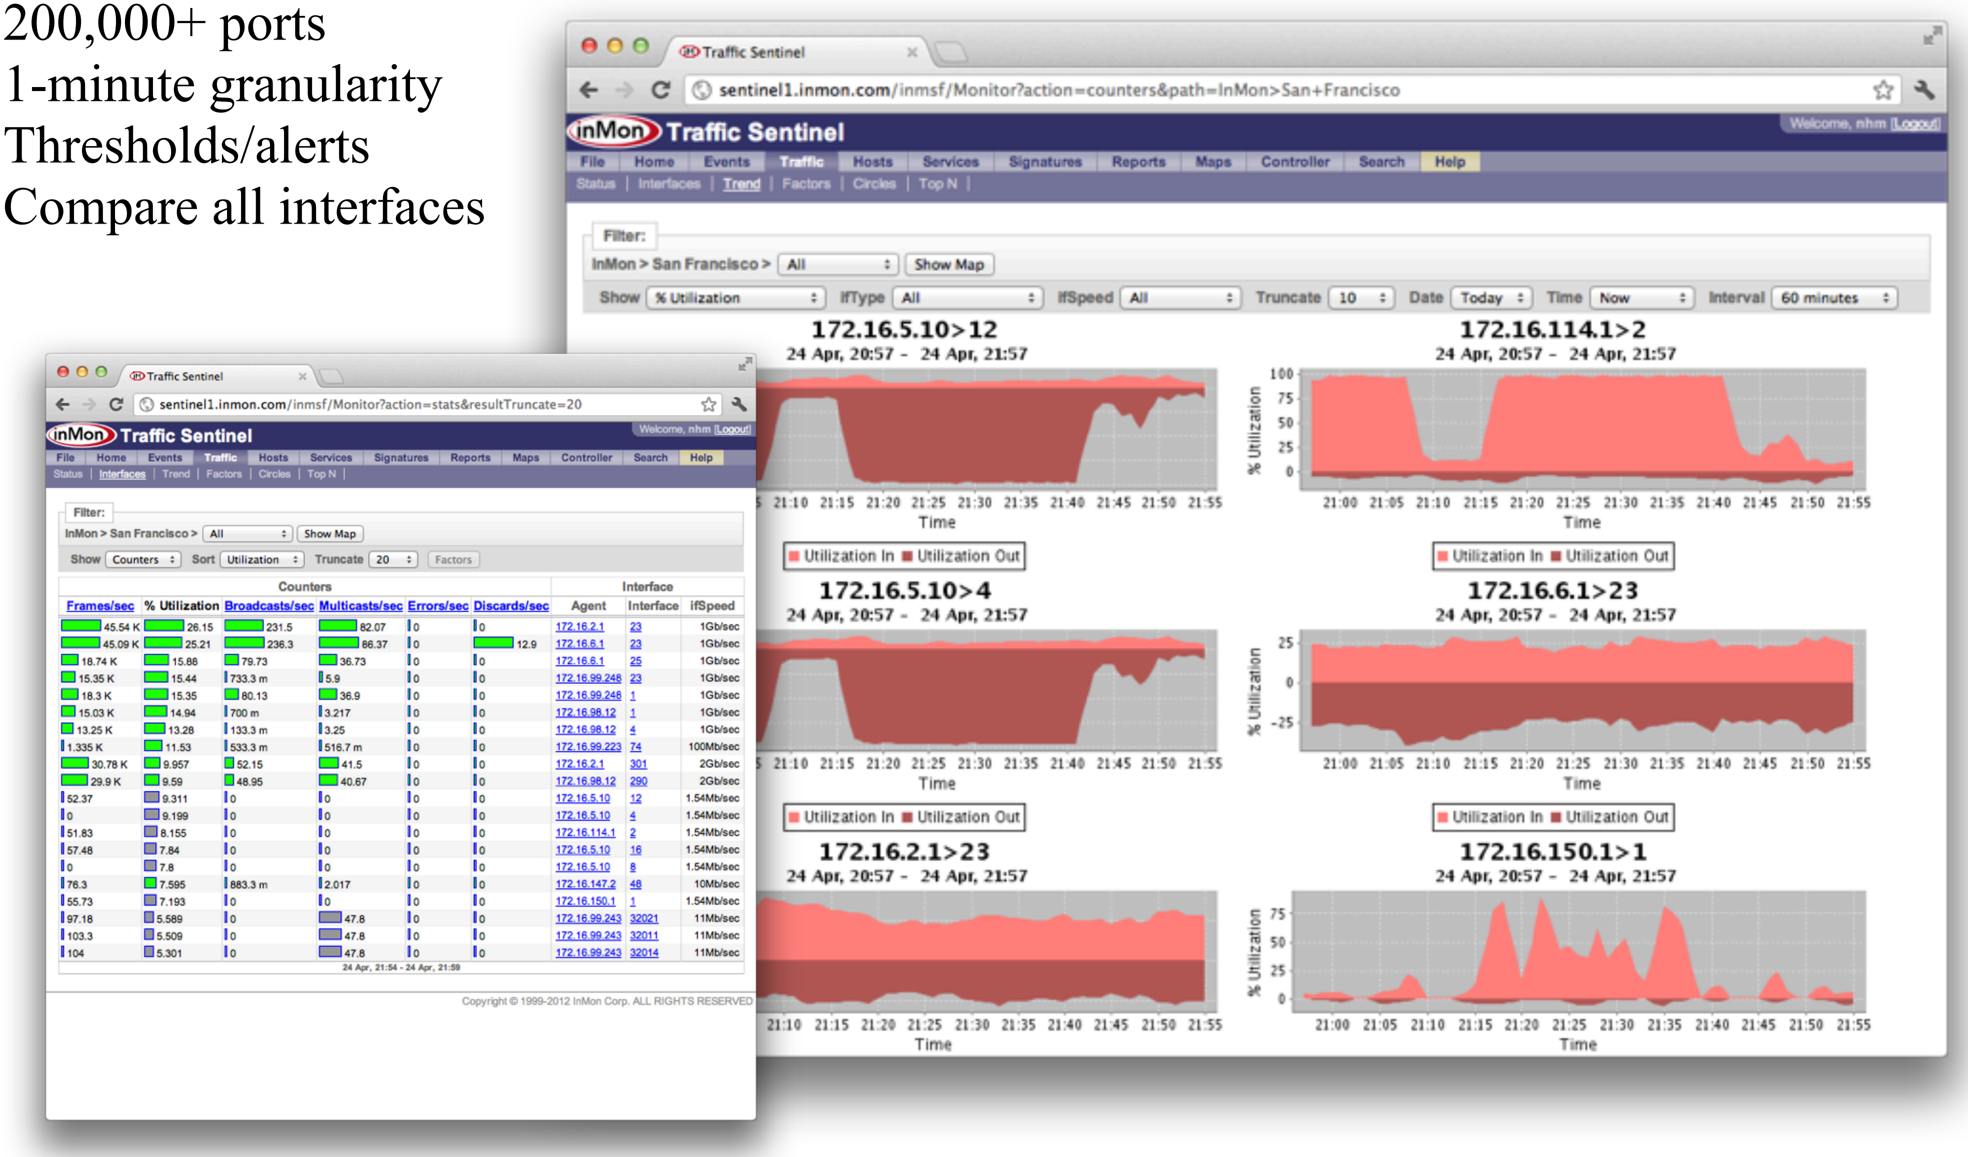

Traffic Sentinel
sentinel1.inmon.com/inmsf/Monitor?action=counters&path=InMon>San+Francisco
inMon Traffic Sentinel
Welcome, nhm [Logout]
File | Home | Events | Traffic | Hosts | Services | Signatures | Reports | Maps | Controller | Search | Help
Status | Interfaces | Trend | Factors | Circles | Top N
Filter:
Show Map
172.16.5.10>12
24 Apr, 20:57 - 24 Apr, 21:57
172.16.114.1>2
24 Apr, 20:57 - 24 Apr, 21:57
172.16.5.10>4
24 Apr, 20:57 - 24 Apr, 21:57
172.16.6.1>23
24 Apr, 20:57 - 24 Apr, 21:57
172.16.2.1>23
24 Apr, 20:57 - 24 Apr, 21:57
172.16.150.1>1
24 Apr, 20:57 - 24 Apr, 21:57
% Utilization
100
75
50
25
0
-25
Time
21:00 21:05 21:10 21:15 21:20 21:25 21:30 21:35 21:40 21:45 21:50 21:55
Utilization In
Utilization Out
sentinel1.inmon.com/inmsf/Monitor?action=stats&resultTruncate=20
Show Map
Factors

InMon > San Francisco > All

Show % Utilization | IfType All | IfSpeed All | Truncate 10 | Date Today | Time Now | Interval 60 minutes

InMon > San Francisco > All

Show Counters | Sort Utilization | Truncate 20

Counters						Interface		
Frames/sec	% Utilization	Broadcasts/sec	Multicasts/sec	Errors/sec	Discards/sec	Agent	Interface	ifSpeed
45.54 K	26.15	231.5	82.07	0	0	172.16.2.1	23	1Gb/sec
45.09 K	25.21	236.3	86.37	0	12.9	172.16.6.1	23	1Gb/sec
18.74 K	15.88	79.73	36.73	0	0	172.16.6.1	25	1Gb/sec
15.35 K	15.44	733.3 m	5.9	0	0	172.16.99.248	23	1Gb/sec
18.3 K	15.35	80.13	36.9	0	0	172.16.99.248	1	1Gb/sec
15.03 K	14.94	700 m	3.217	0	0	172.16.98.12	1	1Gb/sec
13.25 K	13.28	133.3 m	3.25	0	0	172.16.98.12	4	1Gb/sec
1.335 K	11.53	533.3 m	516.7 m	0	0	172.16.99.223	74	100Mb/sec
30.78 K	9.957	52.15	41.5	0	0	172.16.2.1	301	2Gb/sec
29.9 K	9.59	48.95	40.67	0	0	172.16.98.12	290	2Gb/sec
52.37	9.311	0	0	0	0	172.16.5.10	12	1.54Mb/sec
0	9.199	0	0	0	0	172.16.5.10	4	1.54Mb/sec
51.83	8.155	0	0	0	0	172.16.114.1	2	1.54Mb/sec
57.48	7.84	0	0	0	0	172.16.5.10	16	1.54Mb/sec
0	7.8	0	0	0	0	172.16.5.10	8	1.54Mb/sec
76.3	7.595	883.3 m	2.017	0	0	172.16.147.2	48	10Mb/sec
55.73	7.193	0	0	0	0	172.16.150.1	1	1.54Mb/sec
97.18	5.589	0	47.8	0	0	172.16.99.243	32021	11Mb/sec
103.3	5.509	0	47.8	0	0	172.16.99.243	32011	11Mb/sec
104	5.301	0	47.8	0	0	172.16.99.243	32014	11Mb/sec

24 Apr, 21:54 - 24 Apr, 21:59

sFlow monitors all protocols

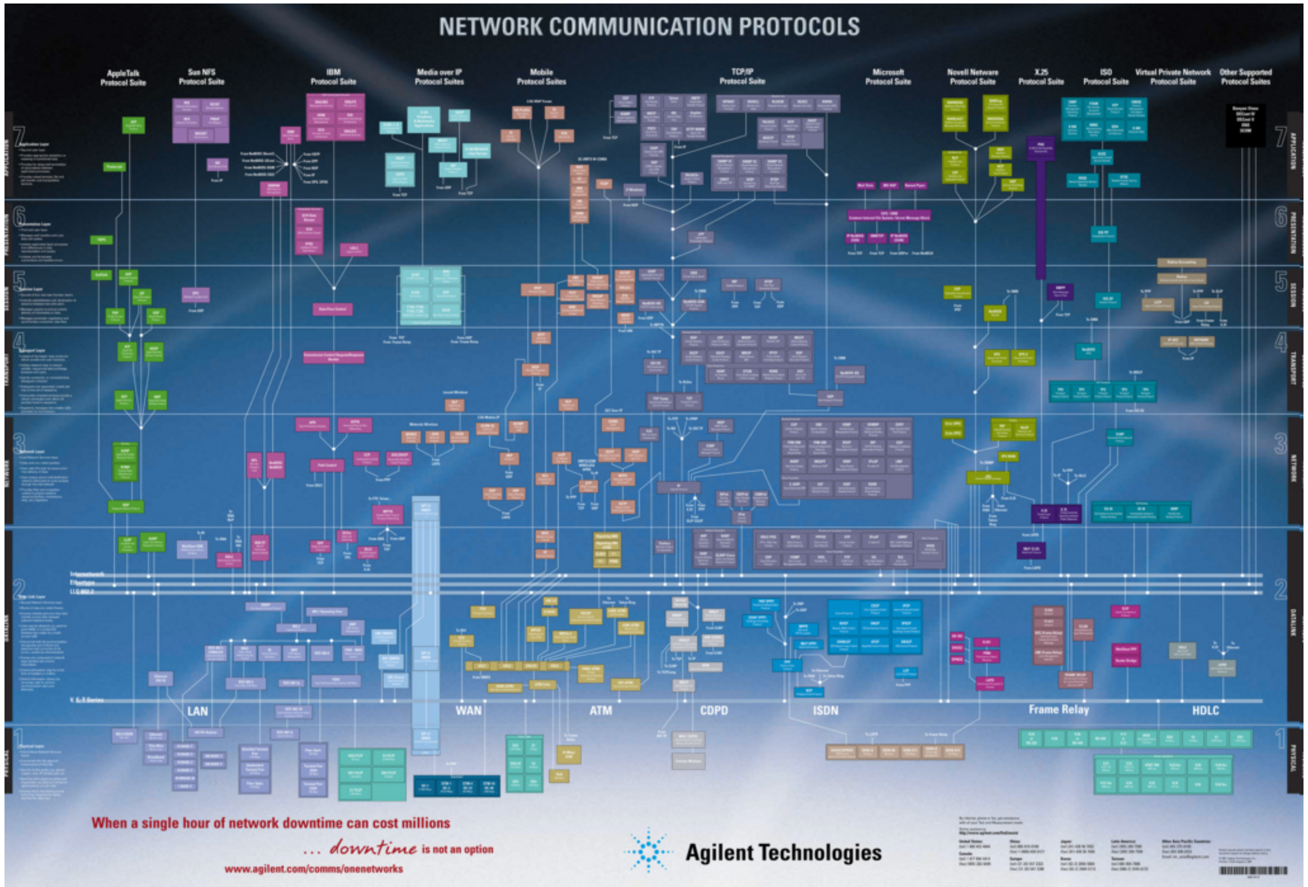
NETWORK COMMUNICATION PROTOCOLS
AppleTalk Protocol Suite
Sun NFS Protocol Suite
IBM Protocol Suite
Media over IP Protocol Suites
Mobile Protocol Suites
TCP/IP Protocol Suite
Microsoft Protocol Suite
Novell Netware Protocol Suite
X.25 Protocol Suite
ISO Protocol Suite
Virtual Private Network Protocol Suite
Other Supported Protocol Suites
7
6
5
4
3
2
1
LAN
WAN
ATM
CDPD
ISDN
Frame Relay
HDLC
When a single hour of network downtime can cost millions
... downtime is not an option
www.agilent.com/comms/onenetworks
Agilent Technologies

- 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

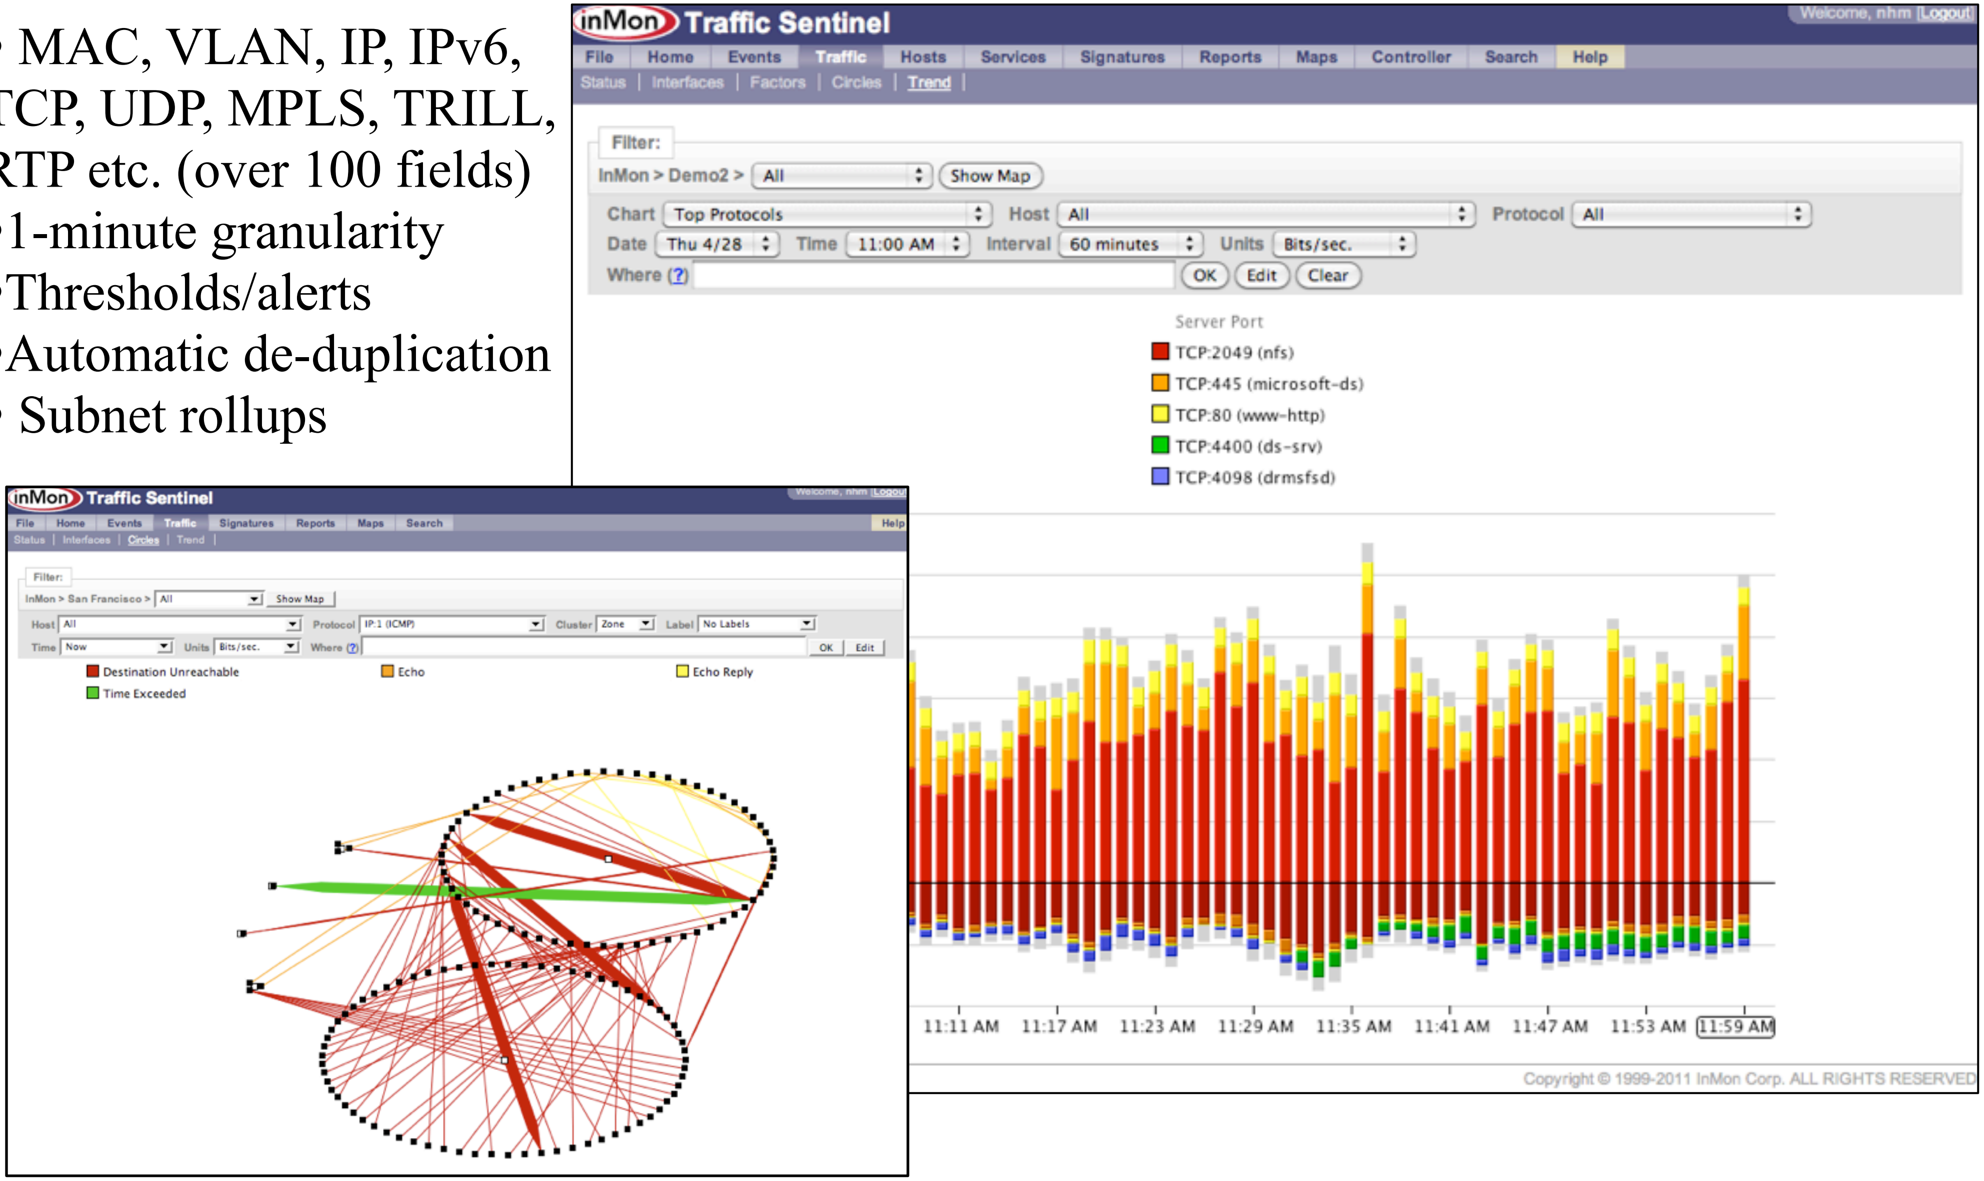

Filter:

InMon > Demo2 > All

Chart: Top Protocols | Host: All | Protocol: All

Date: Thu 4/28 | Time: 11:00 AM | Interval: 60 minutes | Units: Bits/sec.

Where (?)

Filter:

InMon > San Francisco > All

Host: All | Protocol: IP:1 (ICMP) | Cluster: Zone | Label: No Labels

Time: Now | Units: Bits/sec. | Where (?)

sFlow captures packet path

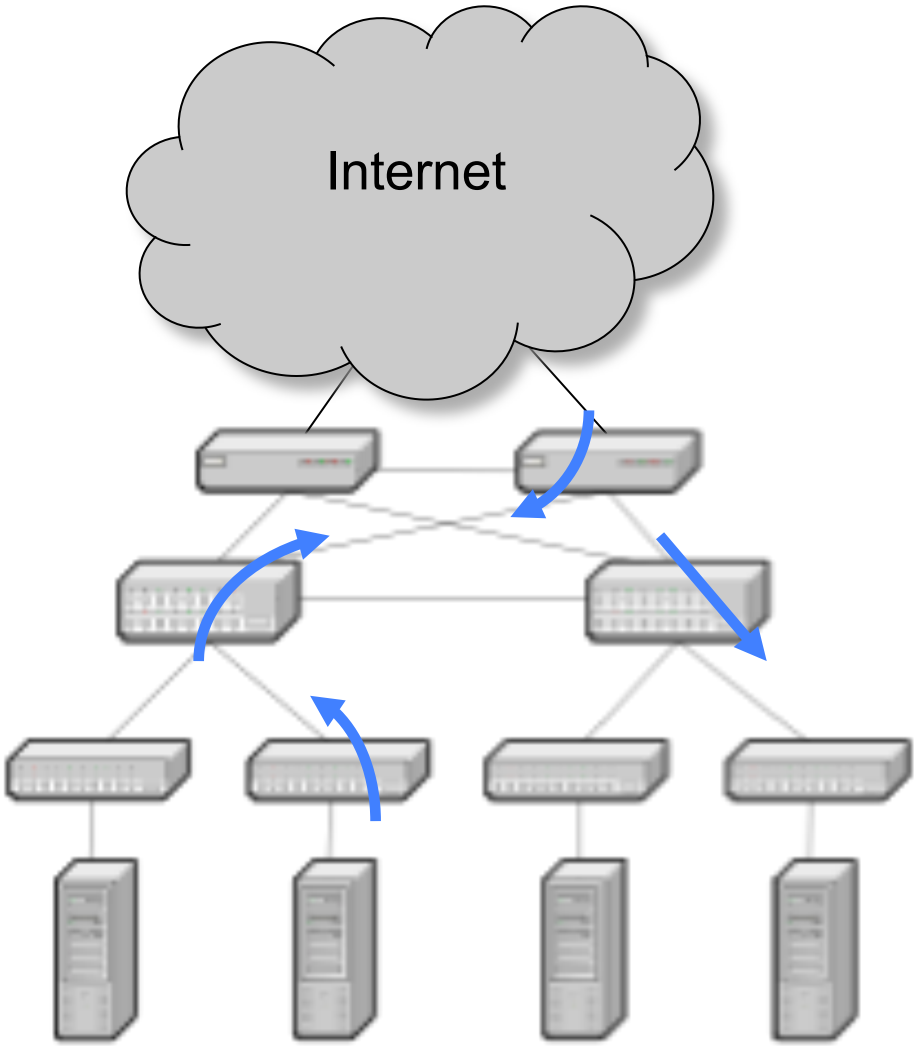

- Each packet sample captures the forwarding path for the packet
- Threading together the paths provides a constantly updating picture of network topology and host locations
- The combination of forwarding table data and packet headers provides an integrated view of traffic. E.g. you can filter on forwarding 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...

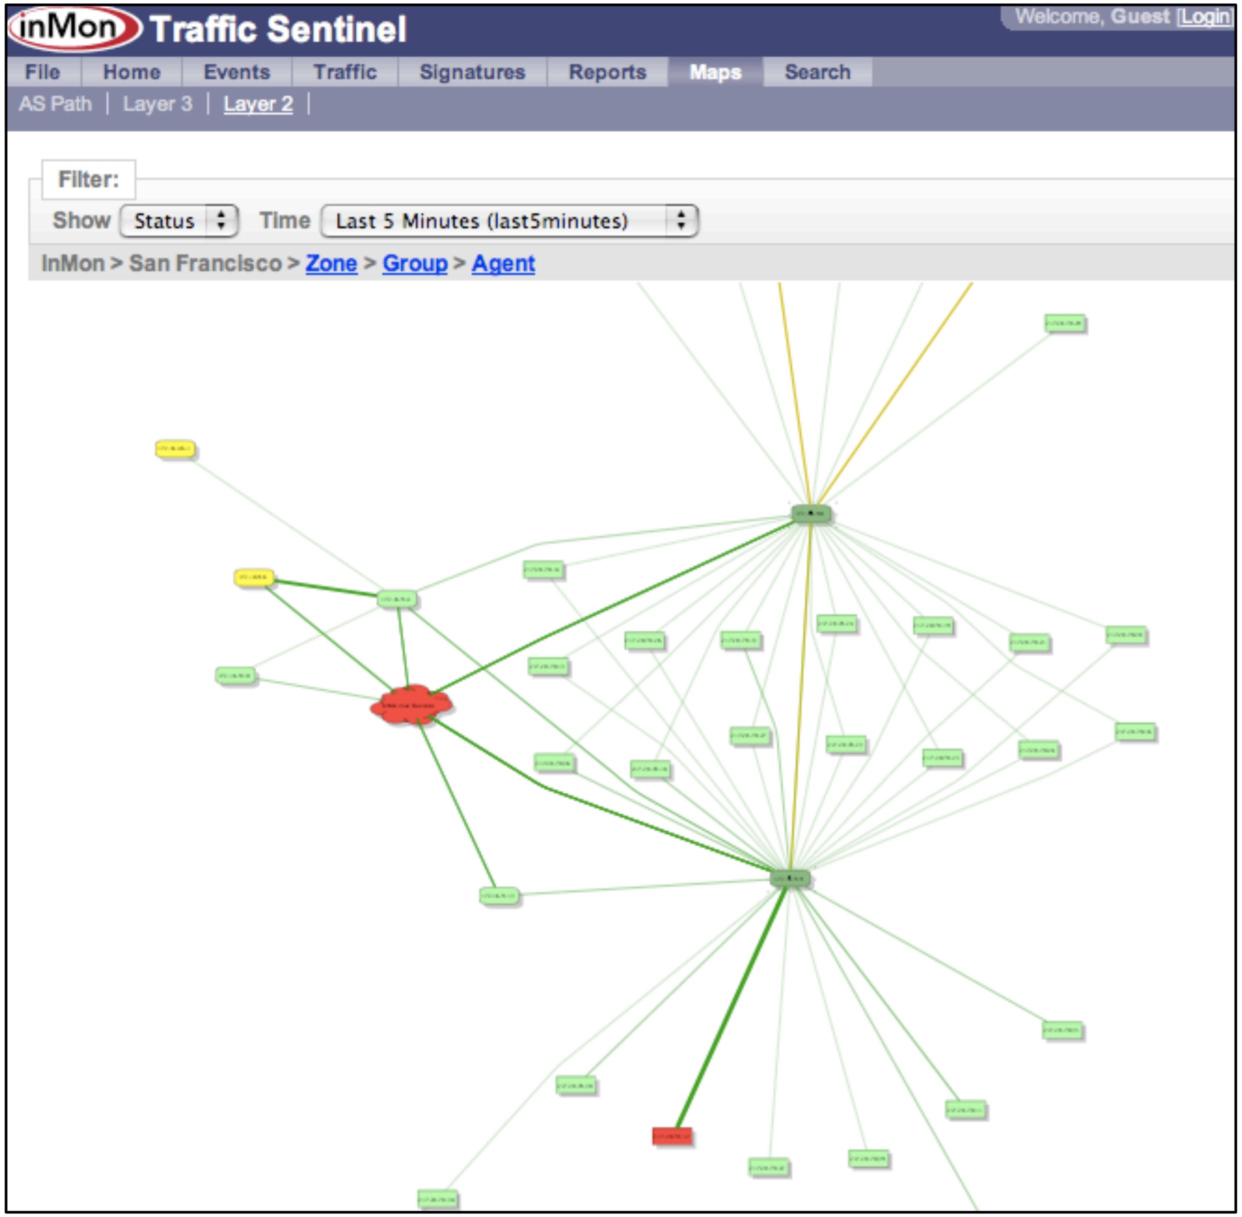

- Auto-layout
- Mouse-wheel zoom
- Show Status,Traffic

(refreshed every minute)

03/08/07

Traffic Sentinel: End-host location

Uses:
sFlow
SNMP

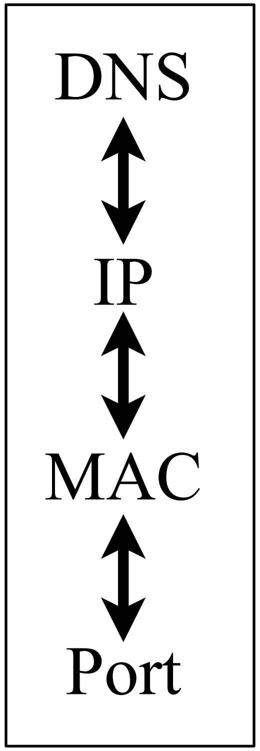
DNS
IP
MAC
Port

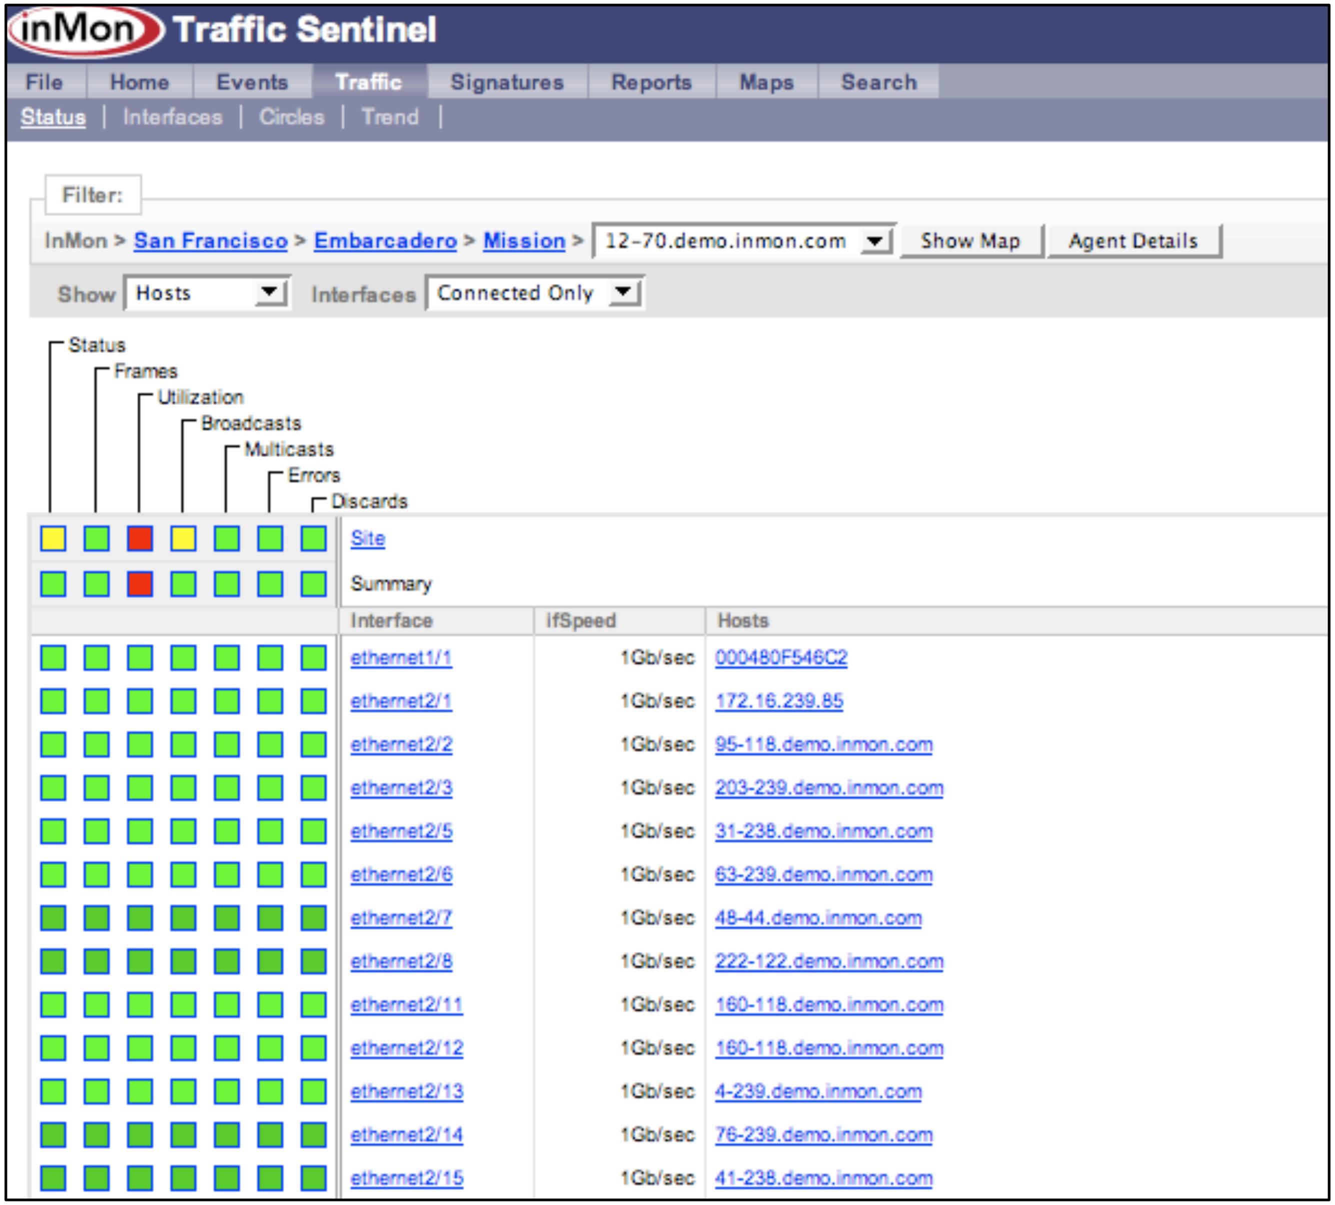

Filter:

InMon > San Francisco > Embarcadero > Mission > 12-70.demo.inmon.com

Show: Hosts   Interfaces: Connected Only

Status	Frames	Utilization	Broadcasts	Multicasts	Errors	Discards	Interface	ifSpeed	Hosts
🟨	🟩	🟥	🟨	🟩	🟩	🟩	Site		
🟩	🟩	🟥	🟩	🟩	🟩	🟩	Summary		
🟩	🟩	🟩	🟩	🟩	🟩	🟩	ethernet1/1	1Gb/sec	000480F546C2
🟩	🟩	🟩	🟩	🟩	🟩	🟩	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

- Configuring sFlow on all ports of a Brocade switch:

```
config> int e 1/1 to 1/48
interface> sflow forwarding
config> sflow destination 192.168.4.5
config> sflow sample 512
config> sflow polling-interval 30
config> sflow enable
```

Cross-layer correlation: Application, Host and Network

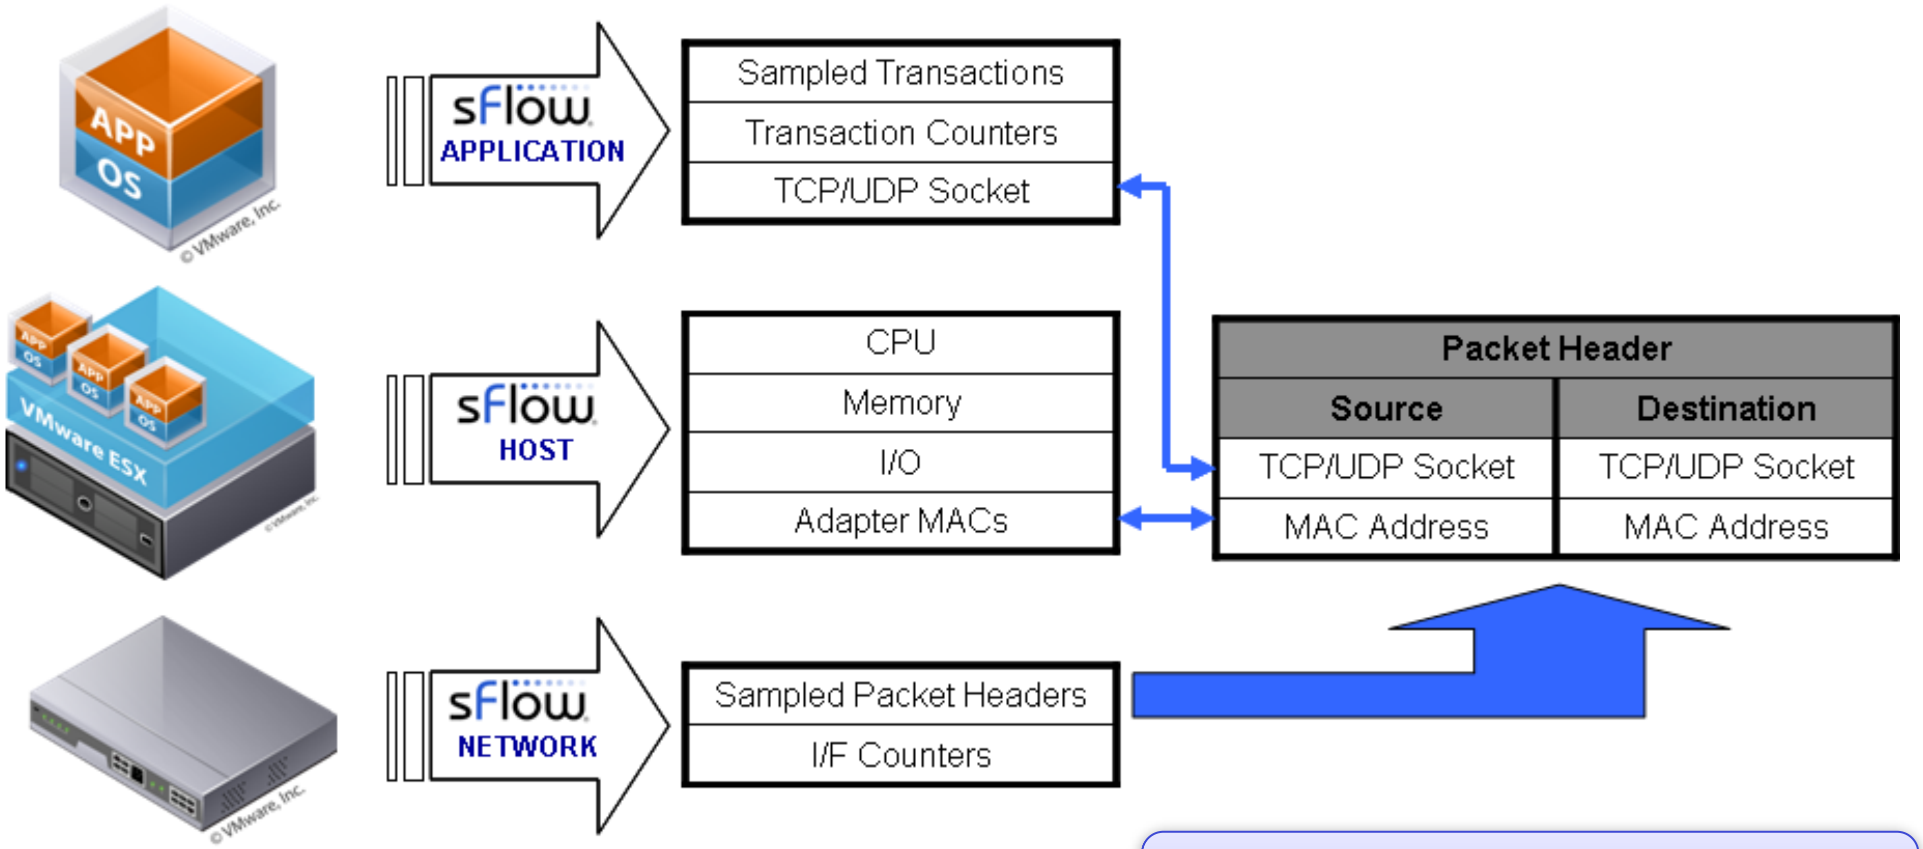

Sampled Transactions
Transaction Counters
TCP/UDP Socket

CPU
Memory
I/O
Adapter MACs

Sampled Packet Headers
I/F Counters

Packet Header	
Source	Destination
TCP/UDP Socket	TCP/UDP Socket
MAC Address	MAC Address

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

Simple Agents - sFlow-HOST
(http://host-sflow.sf.net)

- Portable, open-source (Windows, Linux, Solaris, BSD, Xen, KVM …)
- Ultra-light (100kB, 0.0% cpu)
- Zero licensing costs (but support available)
- Zero-config option (DNS-SD)
- Secure (push-only - does not listen for instructions)
- Scalable (still with 1-minute granularity)
- Standard (no vendor lock-in)

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 - KVM/libvirt 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

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

Supported operating systems:
- FreeBSD
- Linux
- Windows

Supported hypervisors:
- KVM/libvirt
- Xen/XCP/XenServer

For more information, join the host-sflow-discuss mailing list

Copyright © 2010,2011 Host sFlow. All Rights Reserved.

Developed in collaboration with Data-center and Supercomputing experts.

Simple Agents - sFlow-APPLICATION

(http://host-sflow.sf.net/relatedlinks.php)

- Application-layer sFlow agents for:
  - Java
  - Apache
  - NGINX
  - Memcached
  - node.js
  - Tomcat
  - Hadoop
  - PCoIP
- Generic API for other apps (e.g. SDSC “Rocks”)

Developed in collaboration with Data-center and Supercomputing

Why Monitor Everything?

1. Troubleshooting - always have context

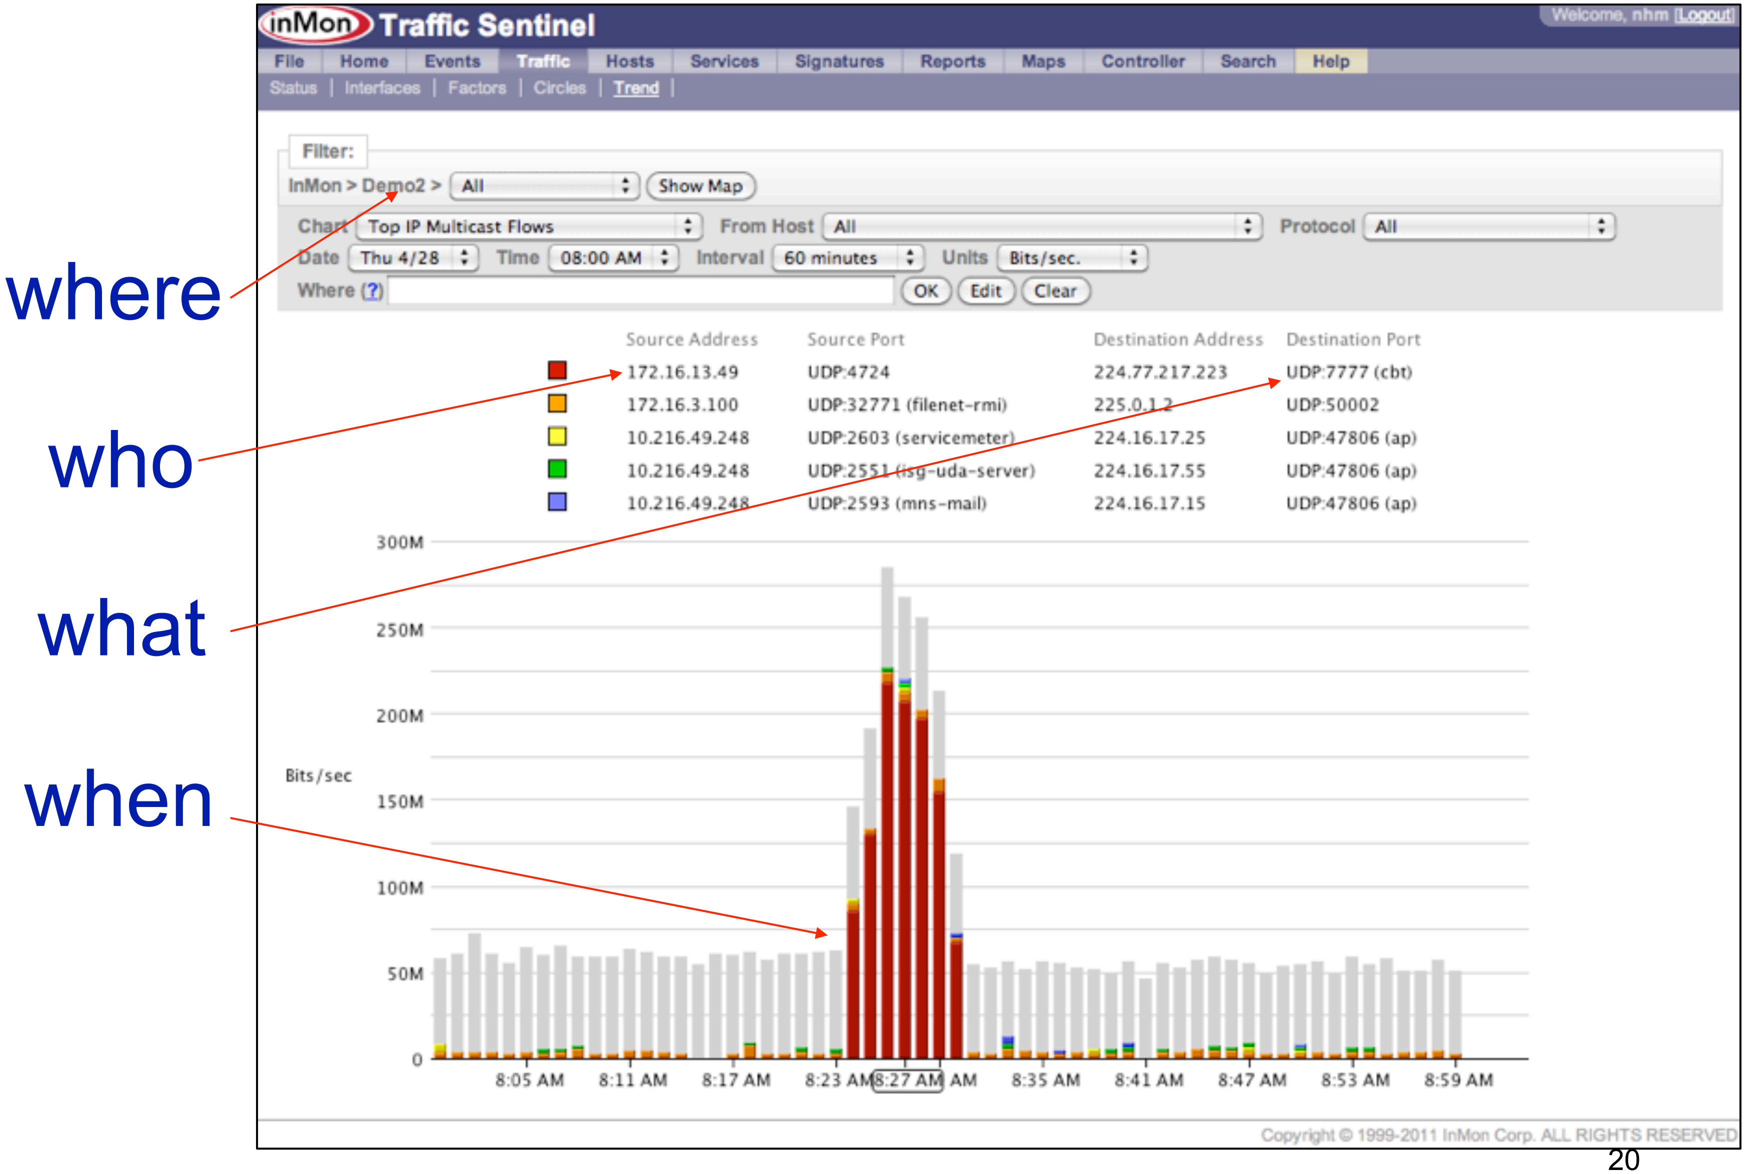

Filter:

InMon > Demo2 > All

Chart: Top IP Multicast Flows | From Host: All | Protocol: All

Date: Thu 4/28 | Time: 08:00 AM | Interval: 60 minutes | Units: Bits/sec.

Where (?)

	Source Address	Source Port	Destination Address	Destination Port
🟥	172.16.13.49	UDP:4724	224.77.217.223	UDP:7777 (cbt)
🟧	172.16.3.100	UDP:32771 (filenet-rmi)	225.0.1.2	UDP:50002
🟨	10.216.49.248	UDP:2603 (servicemeter)	224.16.17.25	UDP:47806 (ap)
🟩	10.216.49.248	UDP:2551 (isg-uda-server)	224.16.17.55	UDP:47806 (ap)
🟦	10.216.49.248	UDP:2593 (mns-mail)	224.16.17.15	UDP:47806 (ap)

where

who

what

when

Why Monitor Everything?

1. Troubleshooting - always have context

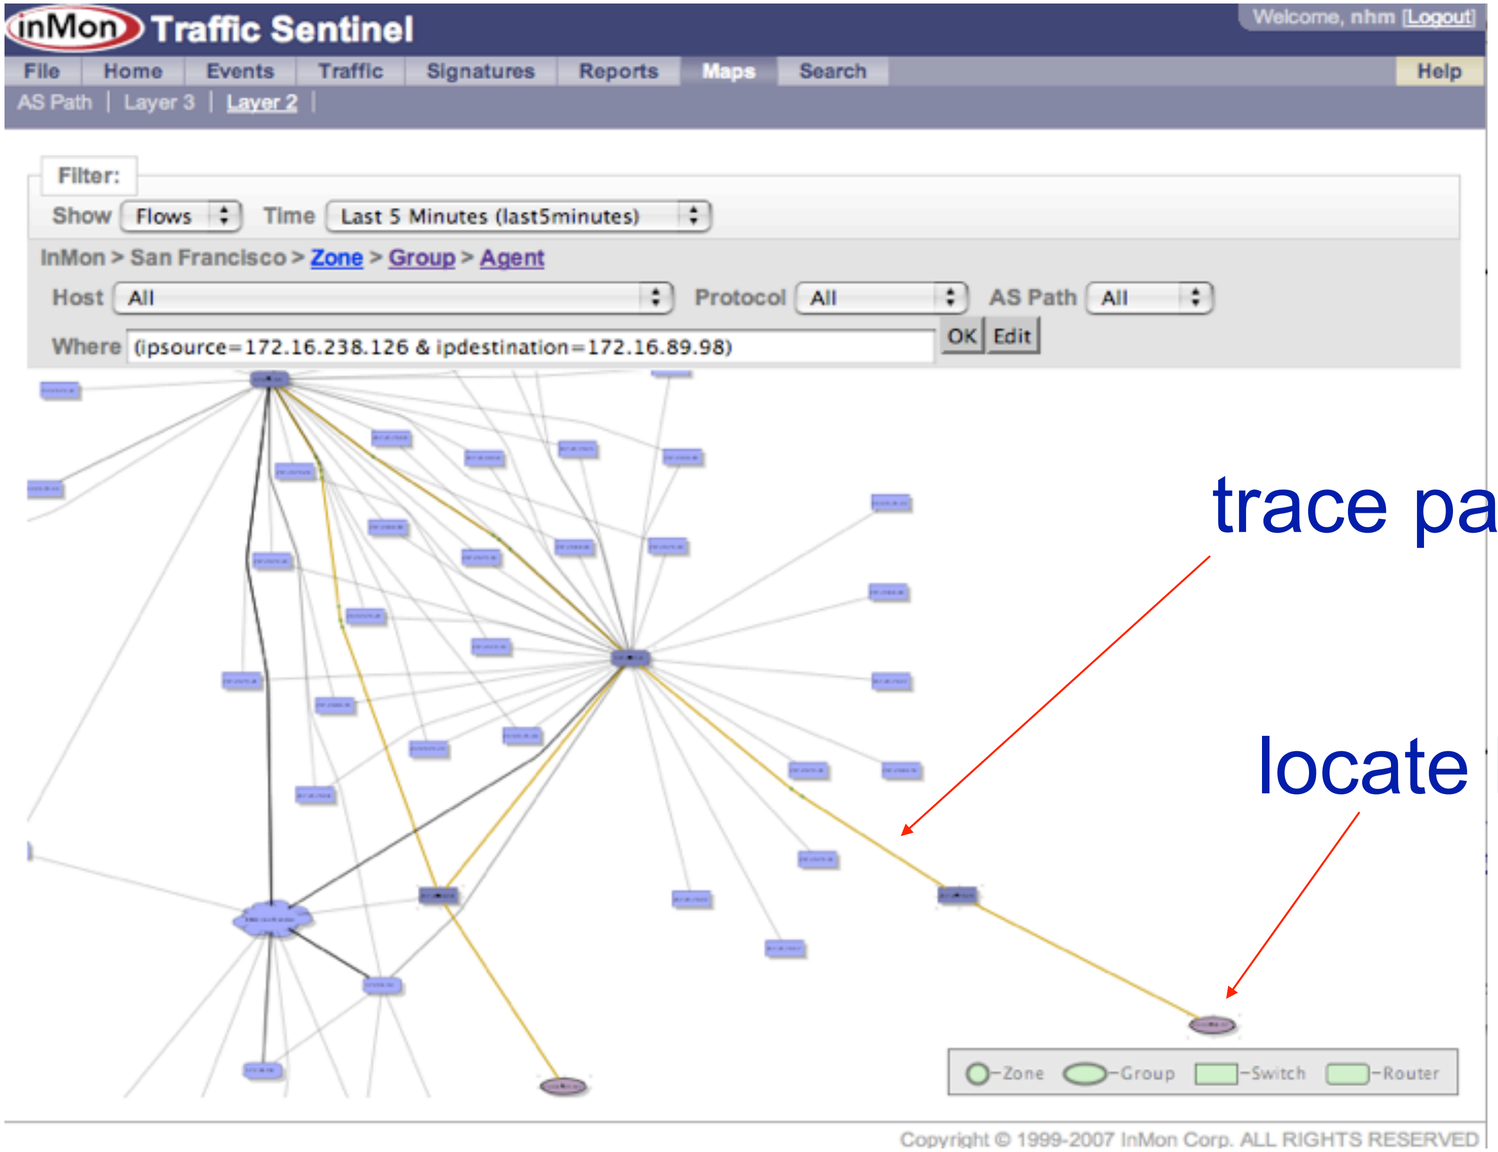

Filter:

Show: Flows
Time: Last 5 Minutes (last5minutes)

InMon > San Francisco > Zone > Group > Agent

Host: All
Protocol: All
AS Path: All

Where: (ipsource=172.16.238.126 & ipdestination=172.16.89.98)

Copyright © 1999-2007 InMon Corp. ALL RIGHTS RESERVED

trace path

locate hosts

Why Monitor Everything?

2. Put Network and Server teams on same page

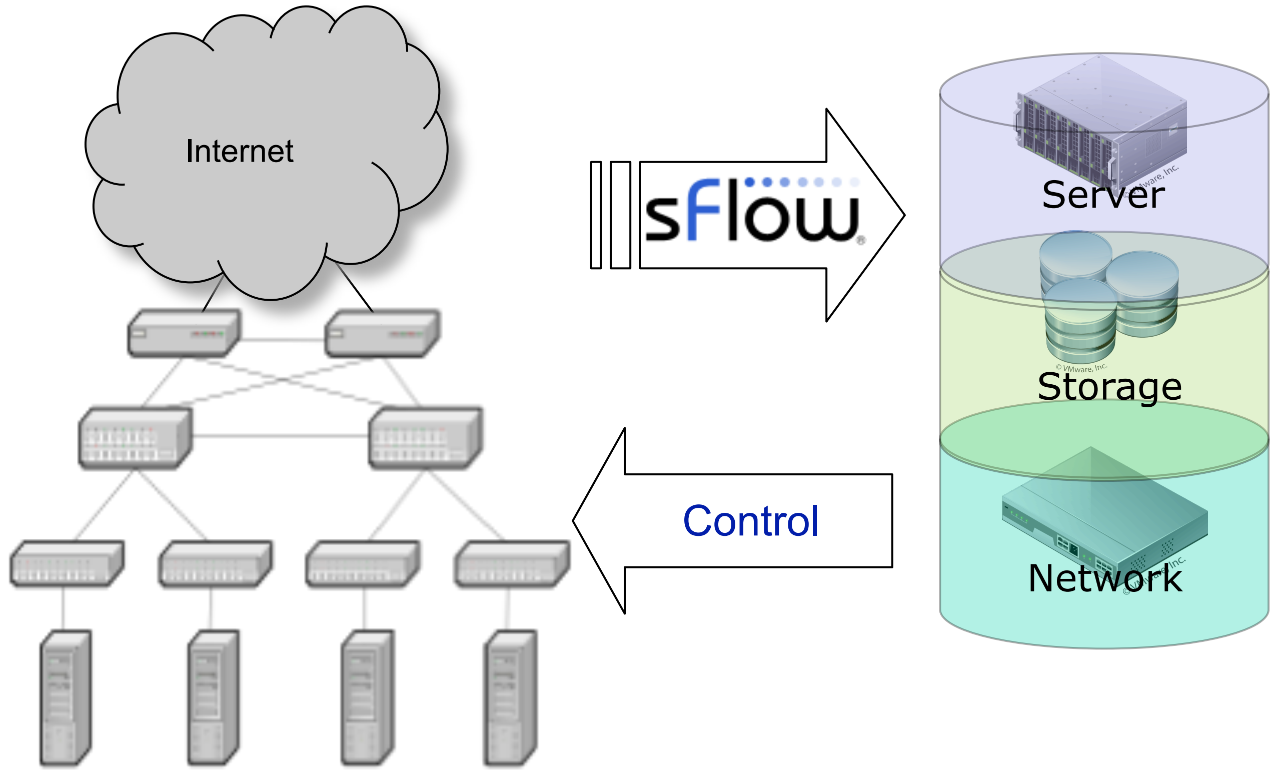

Control

3. Full “Observability” required for automated control

Example

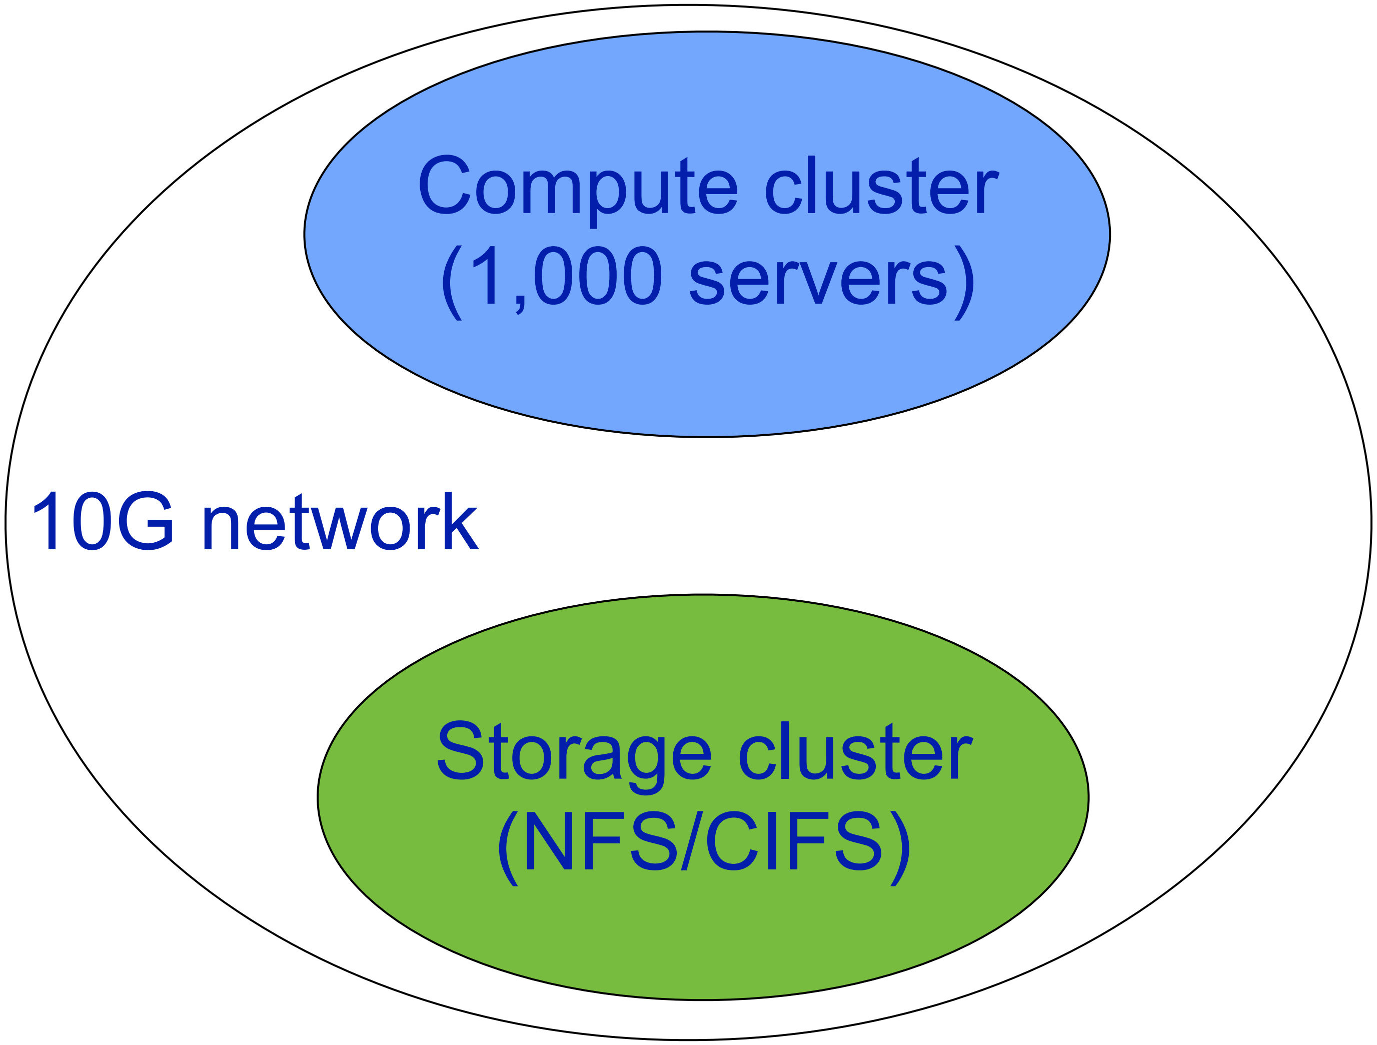
Compute cluster
(1,000 servers)
10G network
Storage cluster
(NFS/CIFS)

Cluster performance

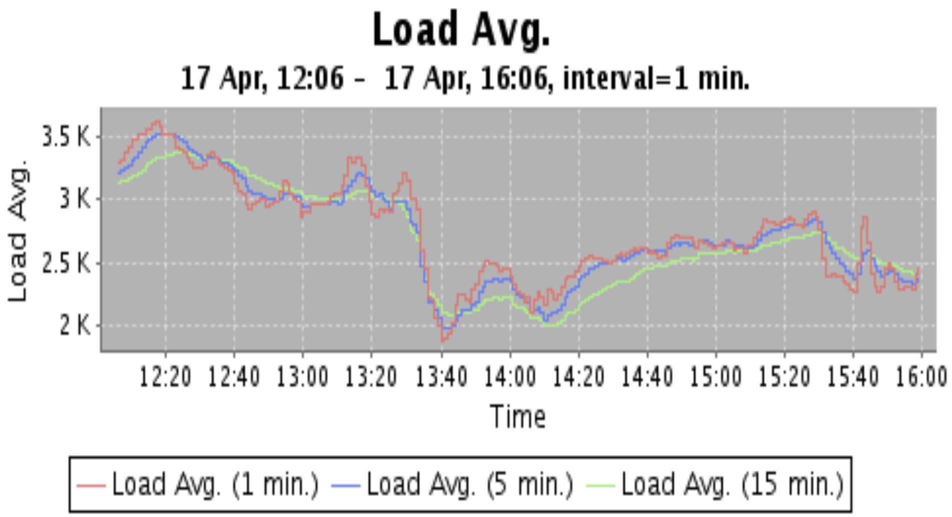

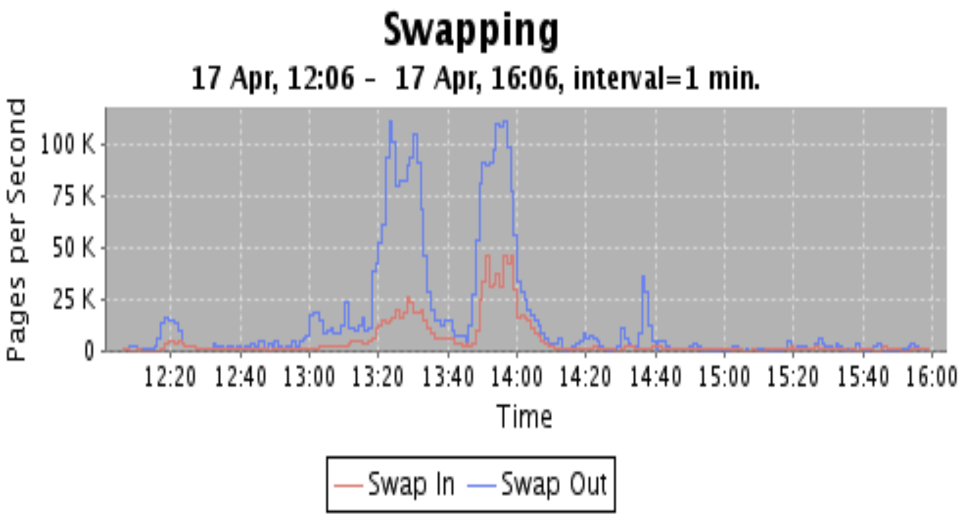

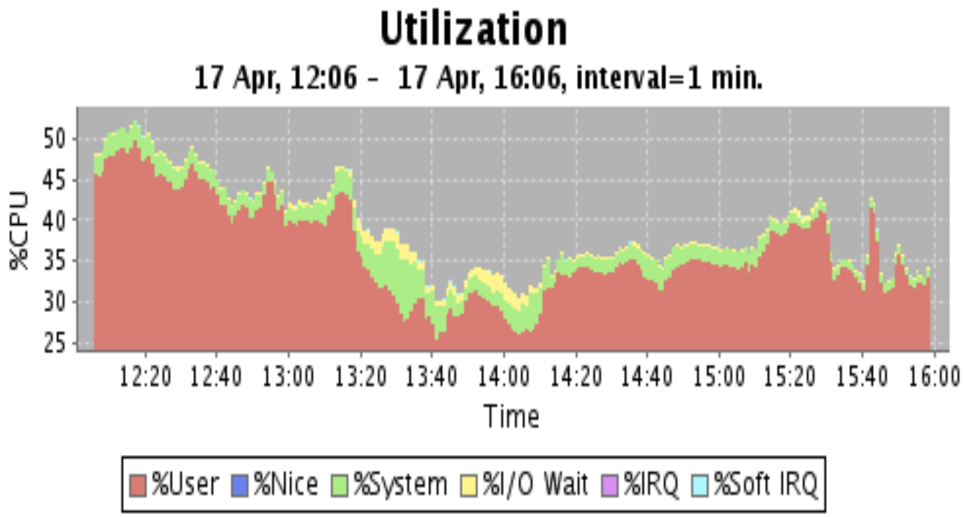

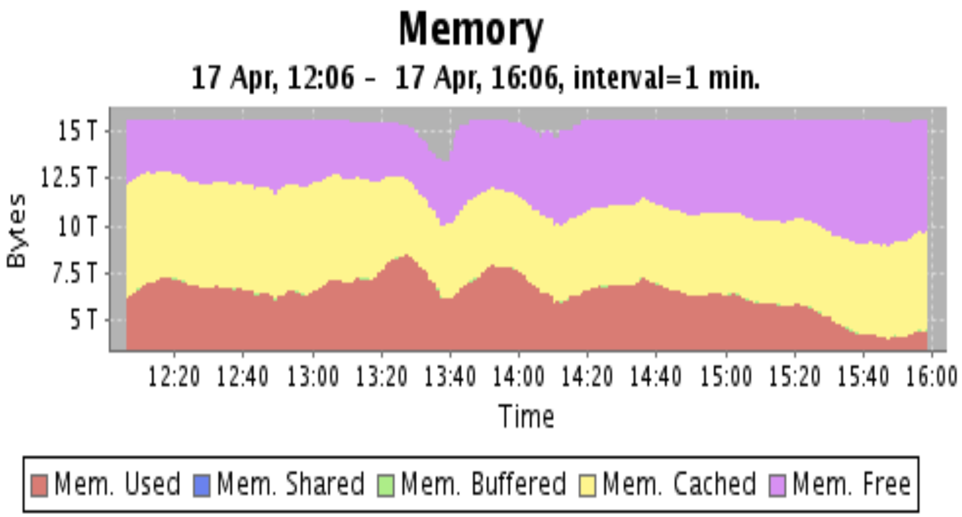

Cluster traffic

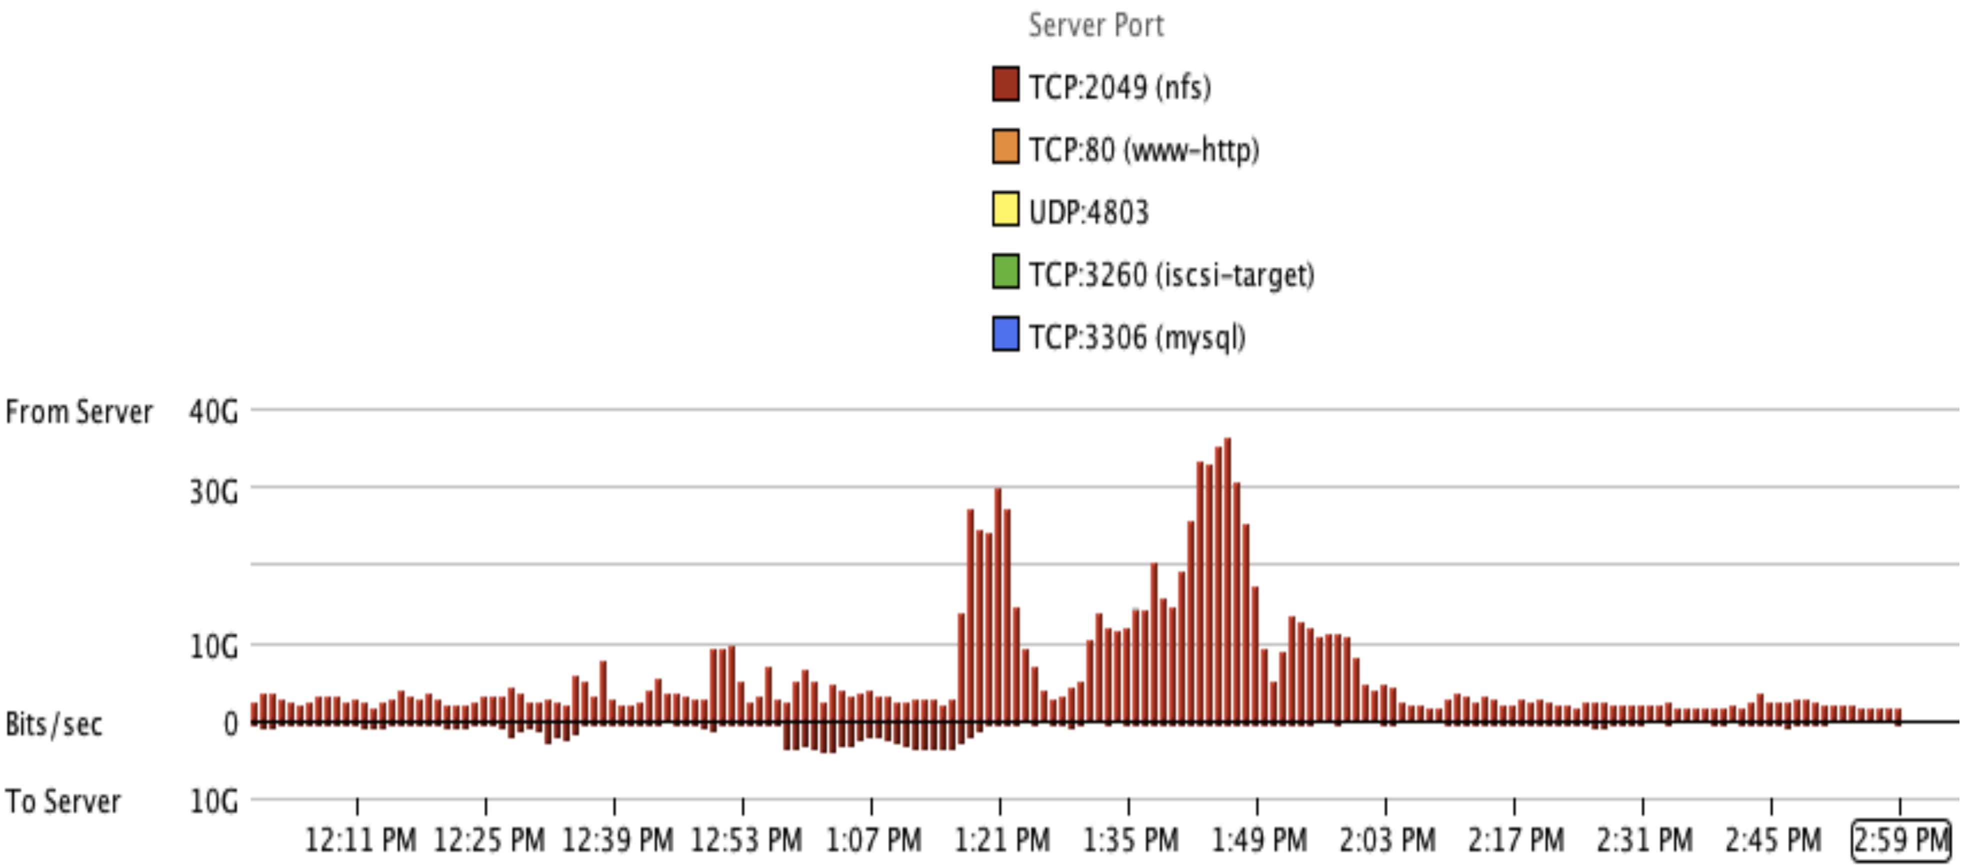

Top servers

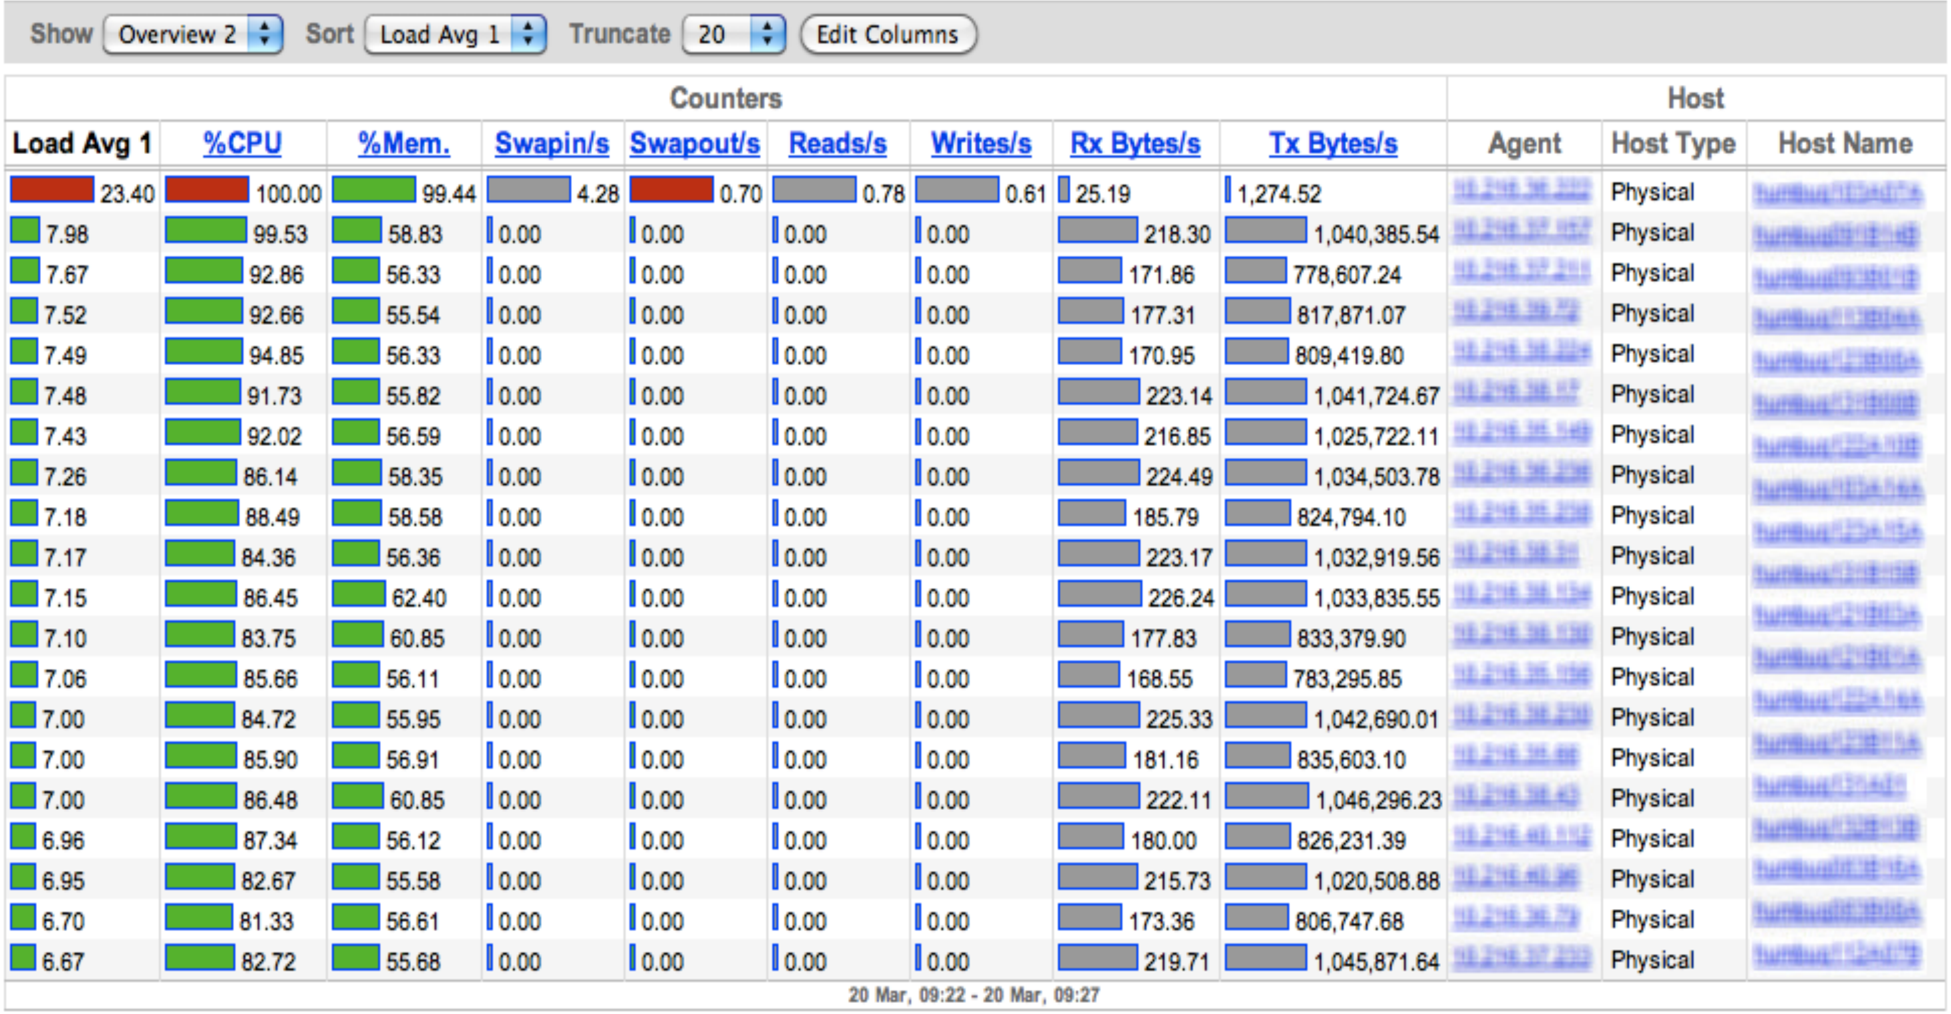

Show: Overview 2 · Sort: Load Avg 1 · Truncate: 20 · Edit Columns

Counters									Host		
Load Avg 1	%CPU	%Mem.	Swapin/s	Swapout/s	Reads/s	Writes/s	Rx Bytes/s	Tx Bytes/s	Agent	Host Type	Host Name
23.40	100.00	99.44	4.28	0.70	0.78	0.61	25.19	1,274.52	[redacted]	Physical	[redacted]
7.98	99.53	58.83	0.00	0.00	0.00	0.00	218.30	1,040,385.54	[redacted]	Physical	[redacted]
7.67	92.86	56.33	0.00	0.00	0.00	0.00	171.86	778,607.24	[redacted]	Physical	[redacted]
7.52	92.66	55.54	0.00	0.00	0.00	0.00	177.31	817,871.07	[redacted]	Physical	[redacted]
7.49	94.85	56.33	0.00	0.00	0.00	0.00	170.95	809,419.80	[redacted]	Physical	[redacted]
7.48	91.73	55.82	0.00	0.00	0.00	0.00	223.14	1,041,724.67	[redacted]	Physical	[redacted]
7.43	92.02	56.59	0.00	0.00	0.00	0.00	216.85	1,025,722.11	[redacted]	Physical	[redacted]
7.26	86.14	58.35	0.00	0.00	0.00	0.00	224.49	1,034,503.78	[redacted]	Physical	[redacted]
7.18	88.49	58.58	0.00	0.00	0.00	0.00	185.79	824,794.10	[redacted]	Physical	[redacted]
7.17	84.36	56.36	0.00	0.00	0.00	0.00	223.17	1,032,919.56	[redacted]	Physical	[redacted]
7.15	86.45	62.40	0.00	0.00	0.00	0.00	226.24	1,033,835.55	[redacted]	Physical	[redacted]
7.10	83.75	60.85	0.00	0.00	0.00	0.00	177.83	833,379.90	[redacted]	Physical	[redacted]
7.06	85.66	56.11	0.00	0.00	0.00	0.00	168.55	783,295.85	[redacted]	Physical	[redacted]
7.00	84.72	55.95	0.00	0.00	0.00	0.00	225.33	1,042,690.01	[redacted]	Physical	[redacted]
7.00	85.90	56.91	0.00	0.00	0.00	0.00	181.16	835,603.10	[redacted]	Physical	[redacted]
7.00	86.48	60.85	0.00	0.00	0.00	0.00	222.11	1,046,296.23	[redacted]	Physical	[redacted]
6.96	87.34	56.12	0.00	0.00	0.00	0.00	180.00	826,231.39	[redacted]	Physical	[redacted]
6.95	82.67	55.58	0.00	0.00	0.00	0.00	215.73	1,020,508.88	[redacted]	Physical	[redacted]
6.70	81.33	56.61	0.00	0.00	0.00	0.00	173.36	806,747.68	[redacted]	Physical	[redacted]
6.67	82.72	55.68	0.00	0.00	0.00	0.00	219.71	1,045,871.64	[redacted]	Physical	[redacted]

20 Mar, 09:22 - 20 Mar, 09:27

Individual server

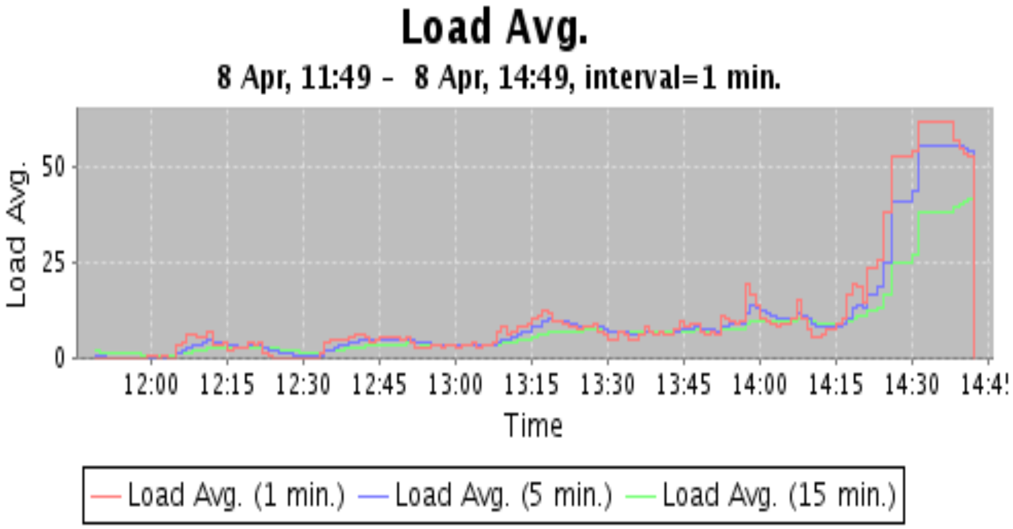

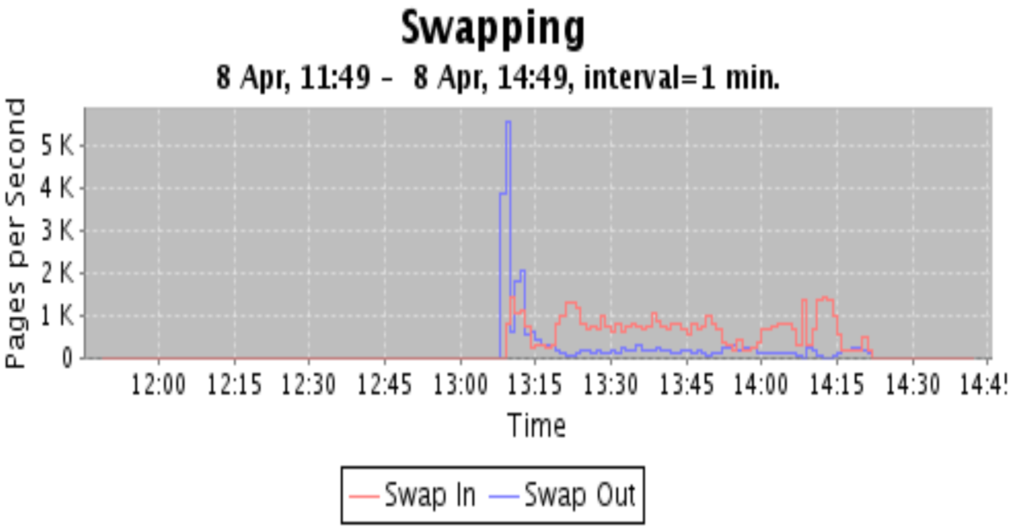

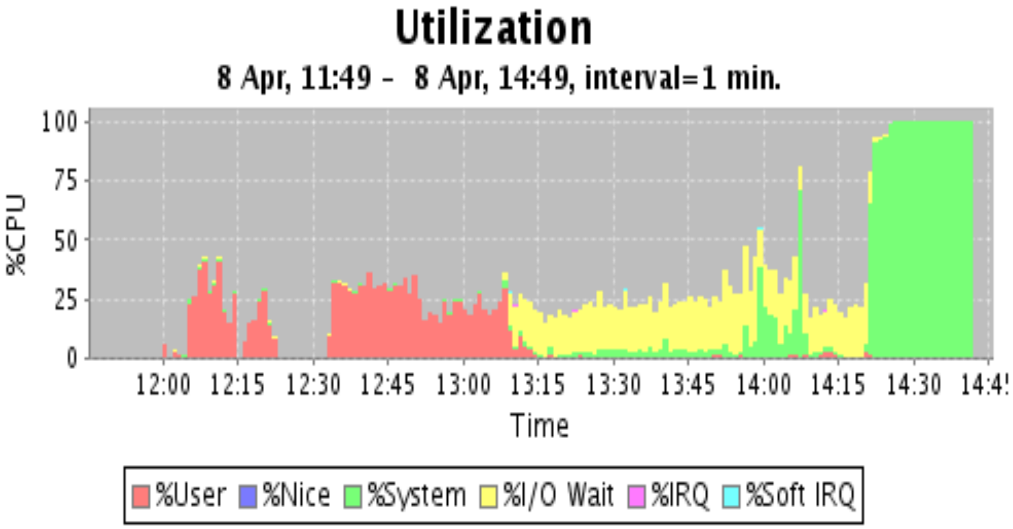

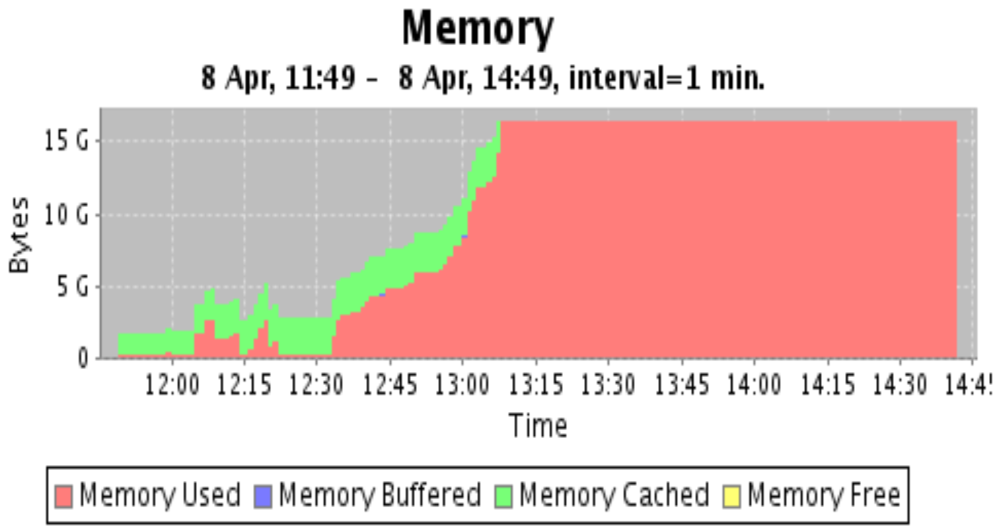

Application dependency map

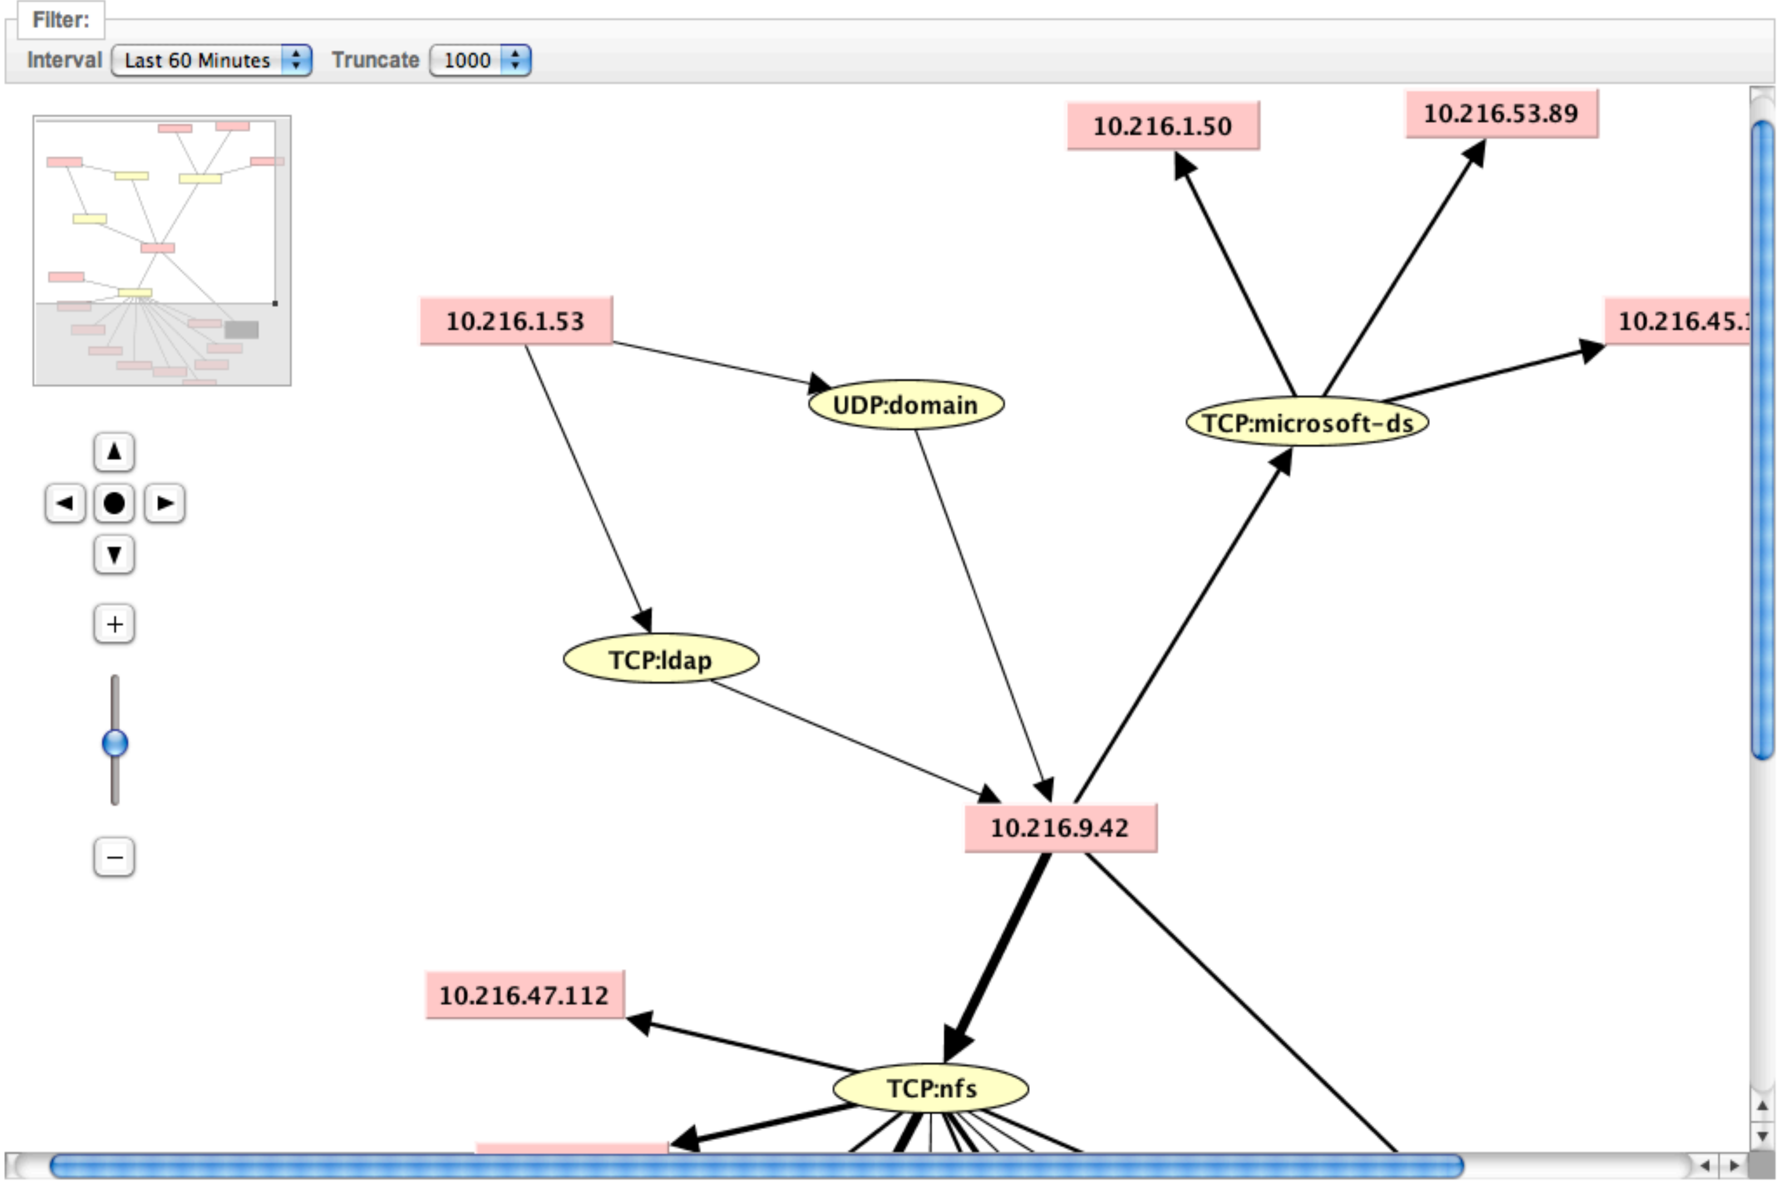

Filter:

Interval: Last 60 Minutes

Truncate: 1000

Network Visibility: sFlow

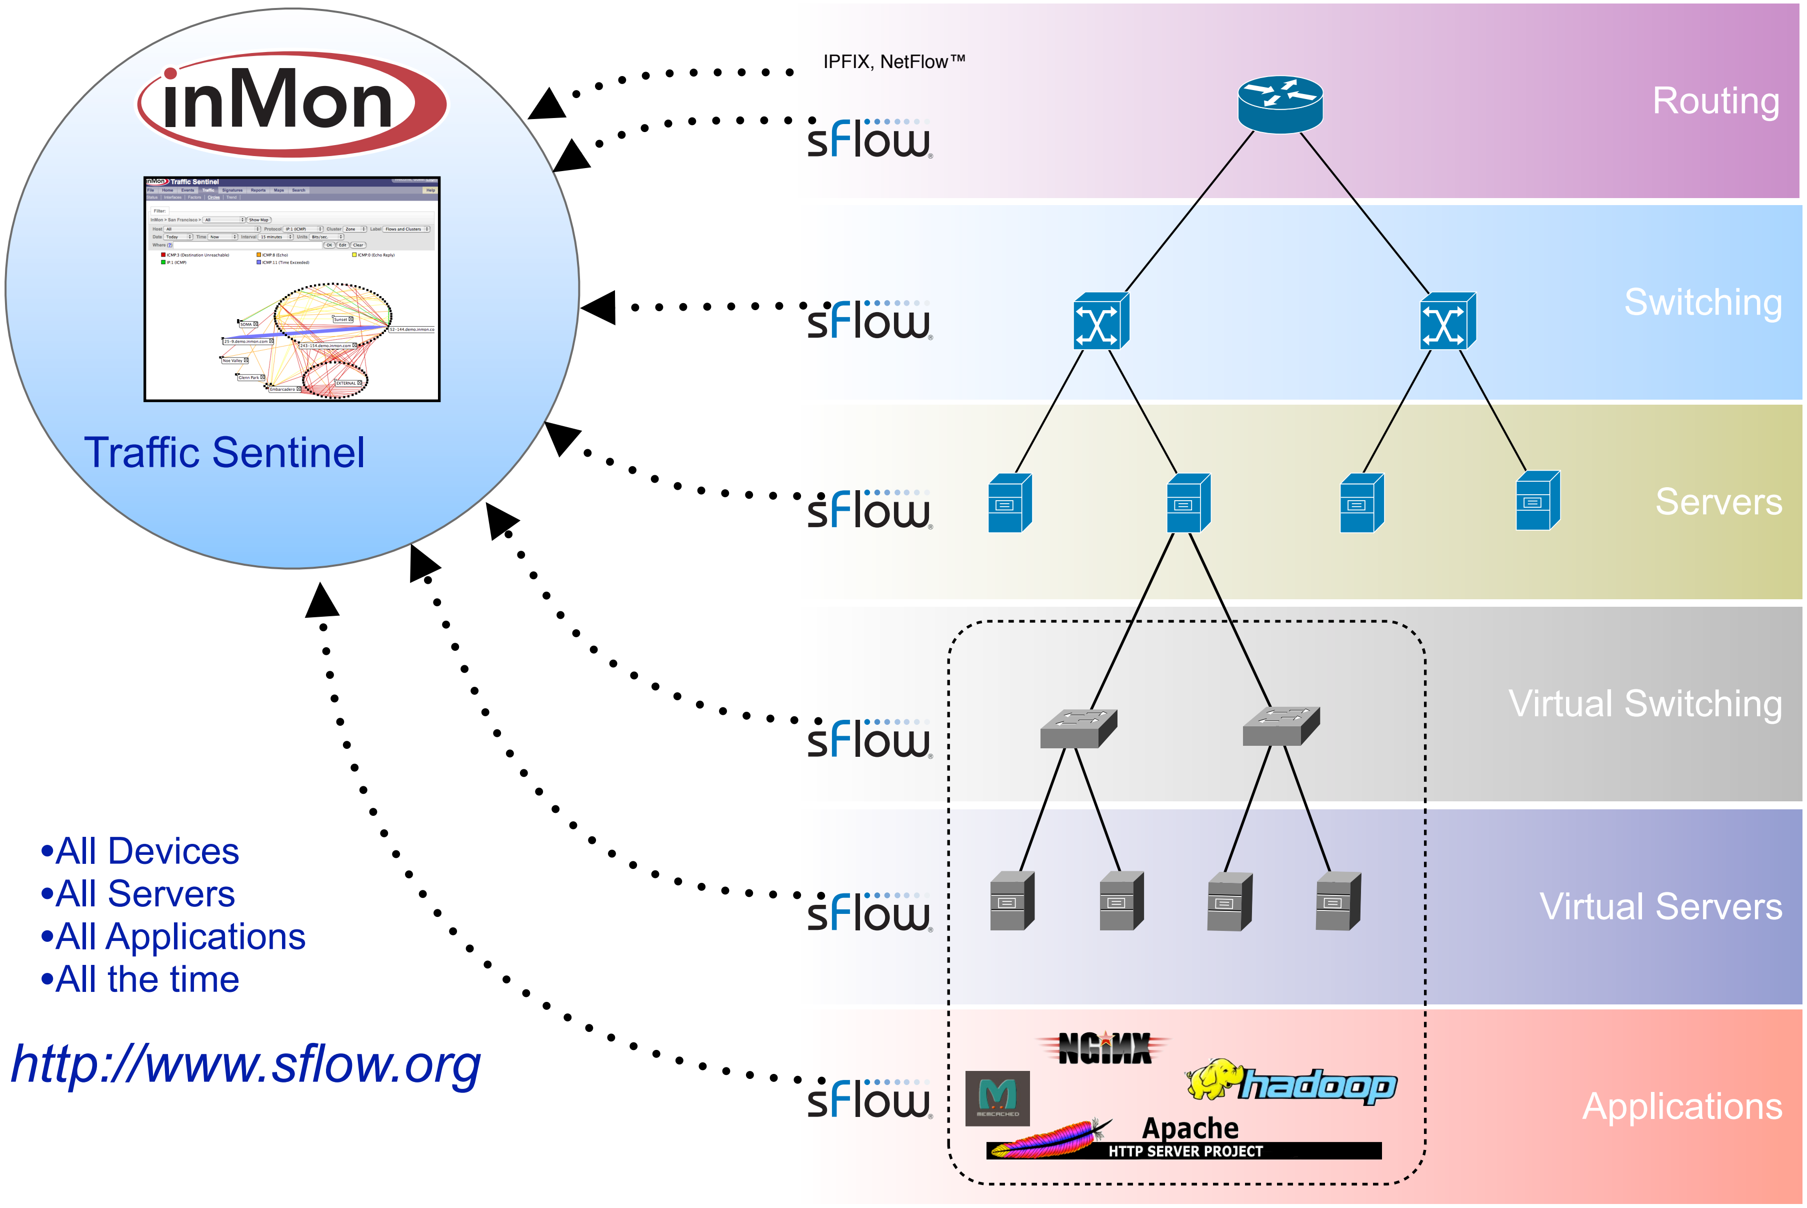

- •All Devices
- •All Servers
- •All Applications
- •All the time

http://www.sflow.org

Network Visibility: VDI

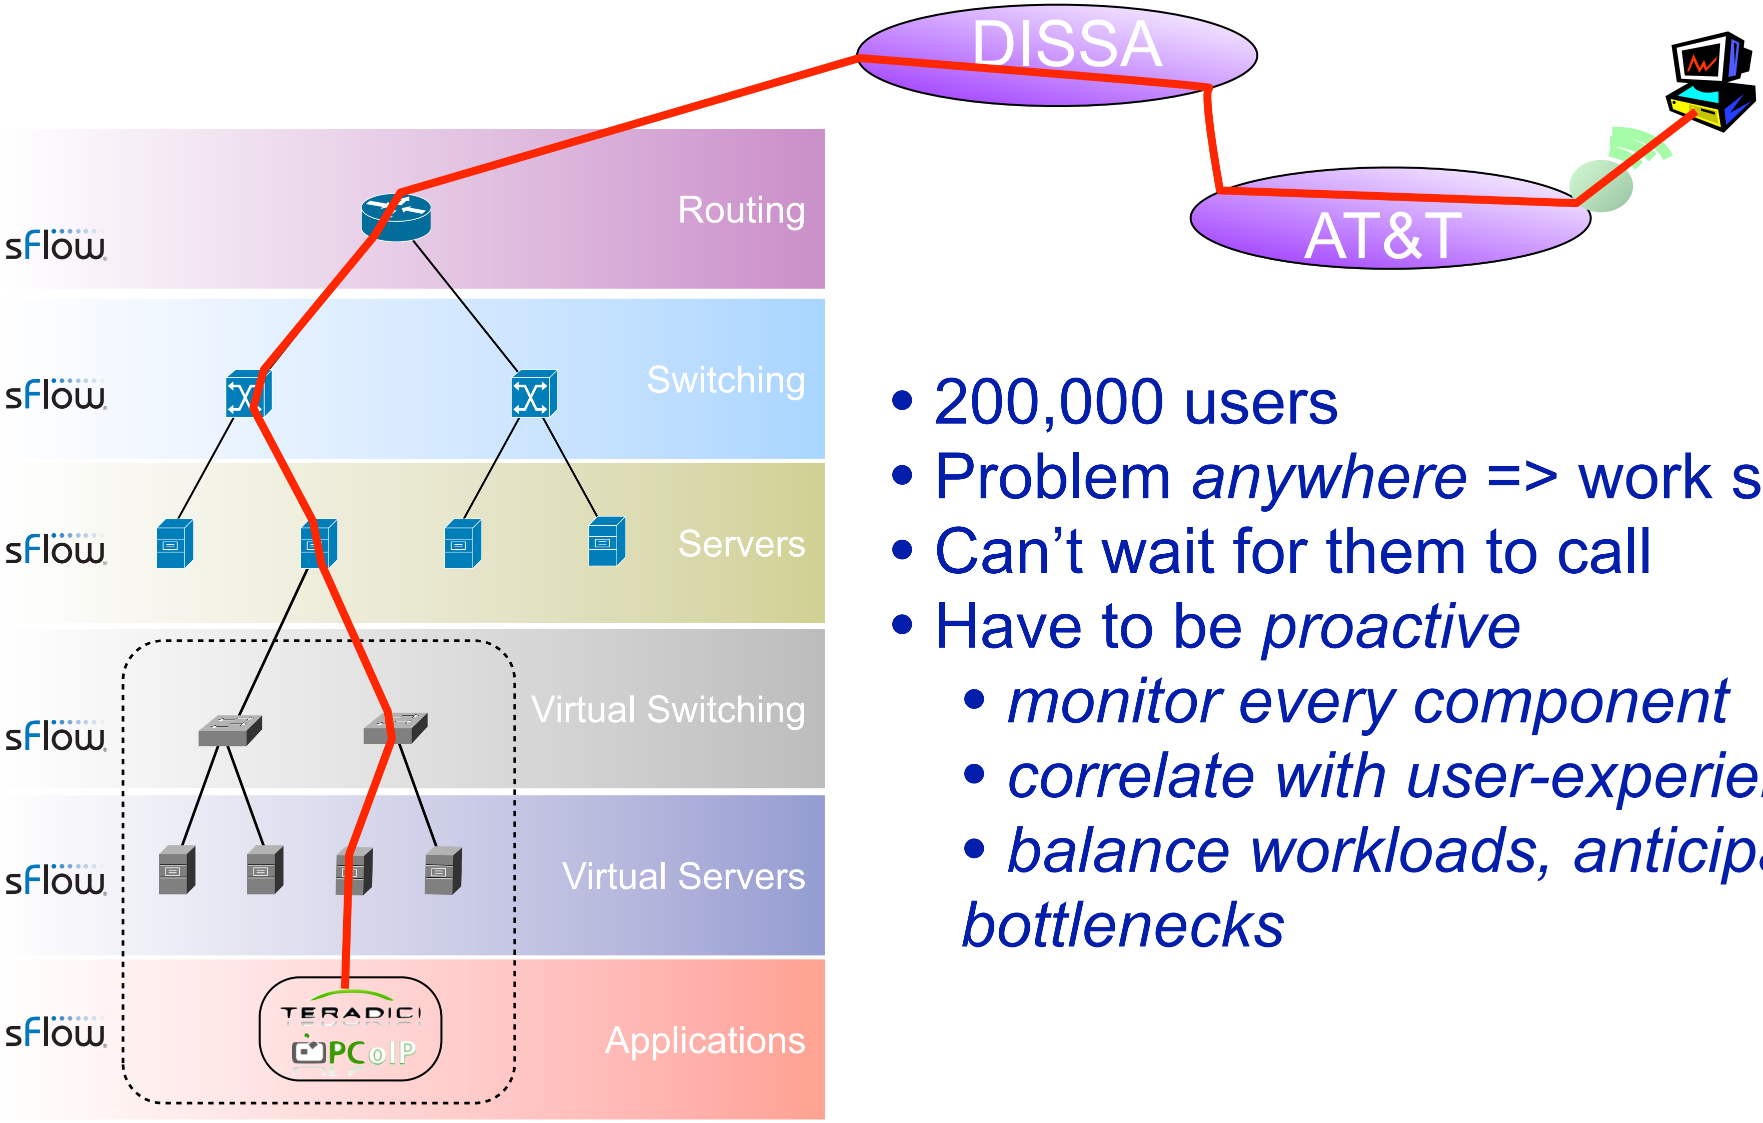

- 200,000 users
- Problem *anywhere* => work stops
- Can’t wait for them to call
- Have to be *proactive*
  - *monitor every component*
  - *correlate with user-experience*
  - *balance workloads, anticipate bottlenecks*

vSwitch sFlow

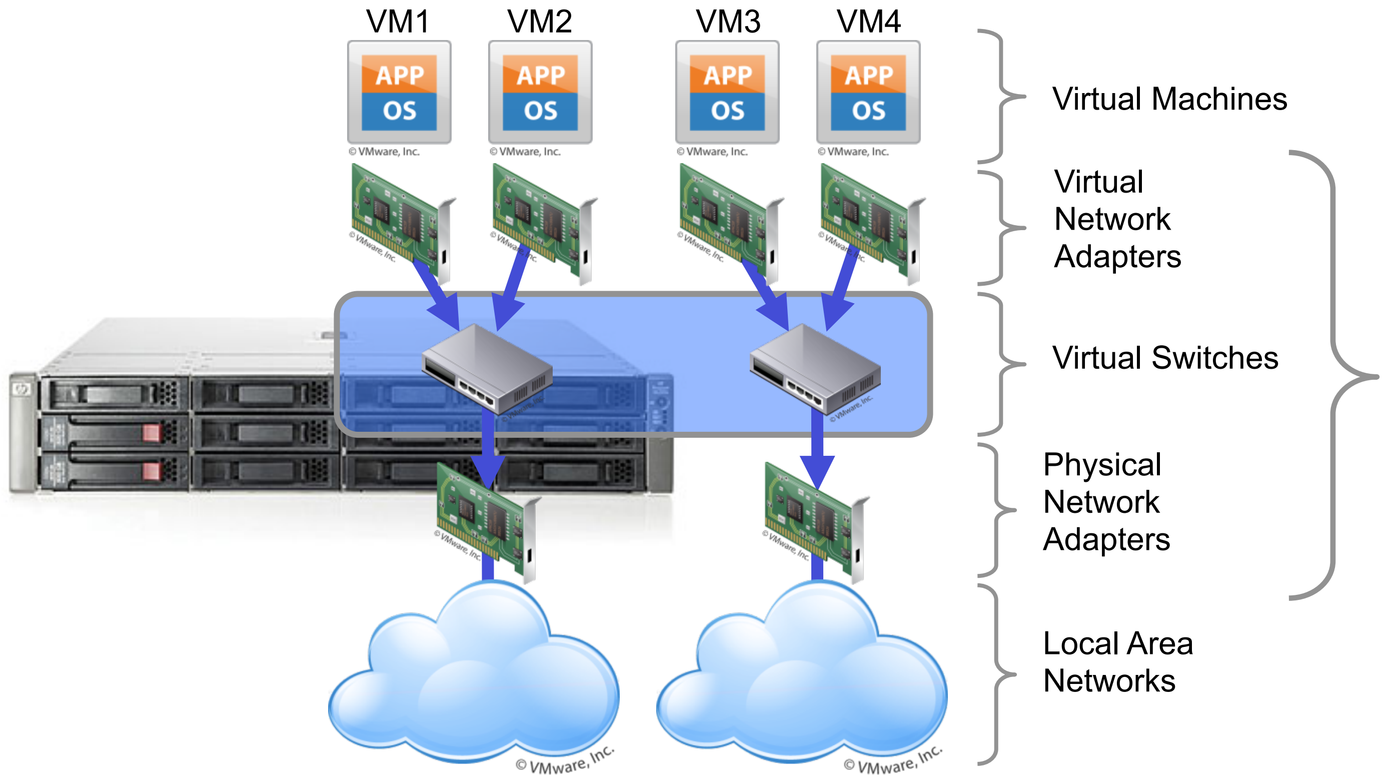

sFlow implemented in virtual switches extends visibility to virtual machines

- •Visibility into vSwitch critical, 20-40 times more vNICs that pNICs
- •Inter-VM traffic only visibility to vSwitch
- •sFlow in vSwitch unifies physical and virtual LAN management
- •Open vSwitch delivers sFlow (and OpenFlow) in open source virtualization stacks Xen/XenServer/KVM

sFlow monitoring of vSwitch: traffic flows

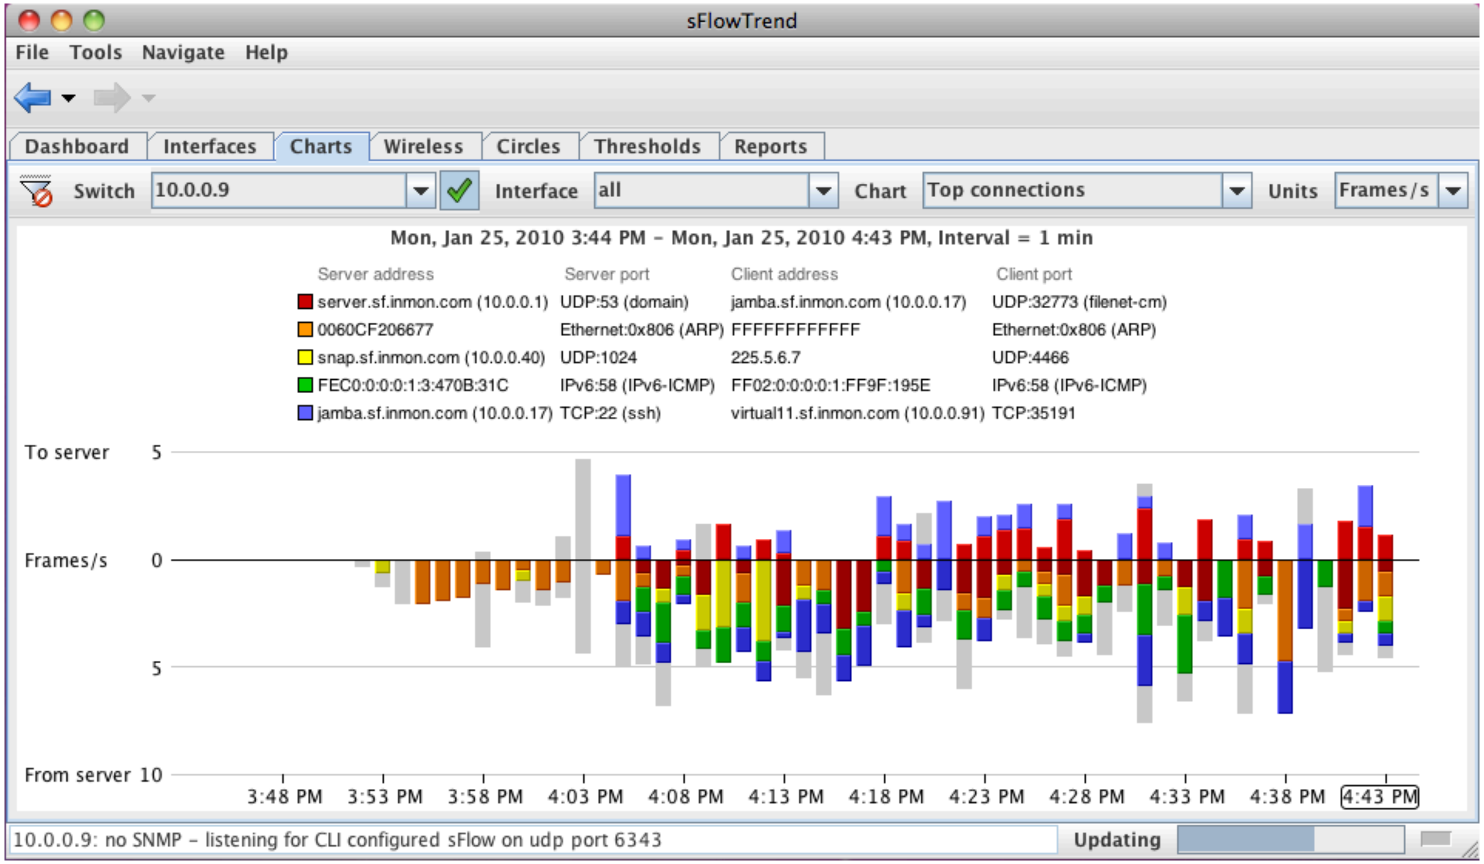

Switch: 10.0.0.9 Interface: all Chart: Top connections Units: Frames/s

Mon, Jan 25, 2010 3:44 PM – Mon, Jan 25, 2010 4:43 PM, Interval = 1 min

Server address	Server port	Client address	Client port
🟥 server.sf.inmon.com (10.0.0.1)	UDP:53 (domain)	jamba.sf.inmon.com (10.0.0.17)	UDP:32773 (filenet-cm)
🟧 0060CF206677	Ethernet:0x806 (ARP)	FFFFFFFFFFFF	Ethernet:0x806 (ARP)
🟨 snap.sf.inmon.com (10.0.0.40)	UDP:1024	225.5.6.7	UDP:4466
🟩 FEC0:0:0:0:1:3:470B:31C	IPv6:58 (IPv6-ICMP)	FF02:0:0:0:0:1:FF9F:195E	IPv6:58 (IPv6-ICMP)
🟦 jamba.sf.inmon.com (10.0.0.17)	TCP:22 (ssh)	virtual11.sf.inmon.com (10.0.0.91)	TCP:35191

10.0.0.9: no SNMP – listening for CLI configured sFlow on udp port 6343

- Visibility into all traffic
  - VM-VM,
  - VM to any other host
  - Layer 2
  - TCP/UDP
  - IPv6
- Data for managing switched traffic
  - VLANs
  - Layer 2 priorities

sFlow monitoring of vSwitch: interface counters

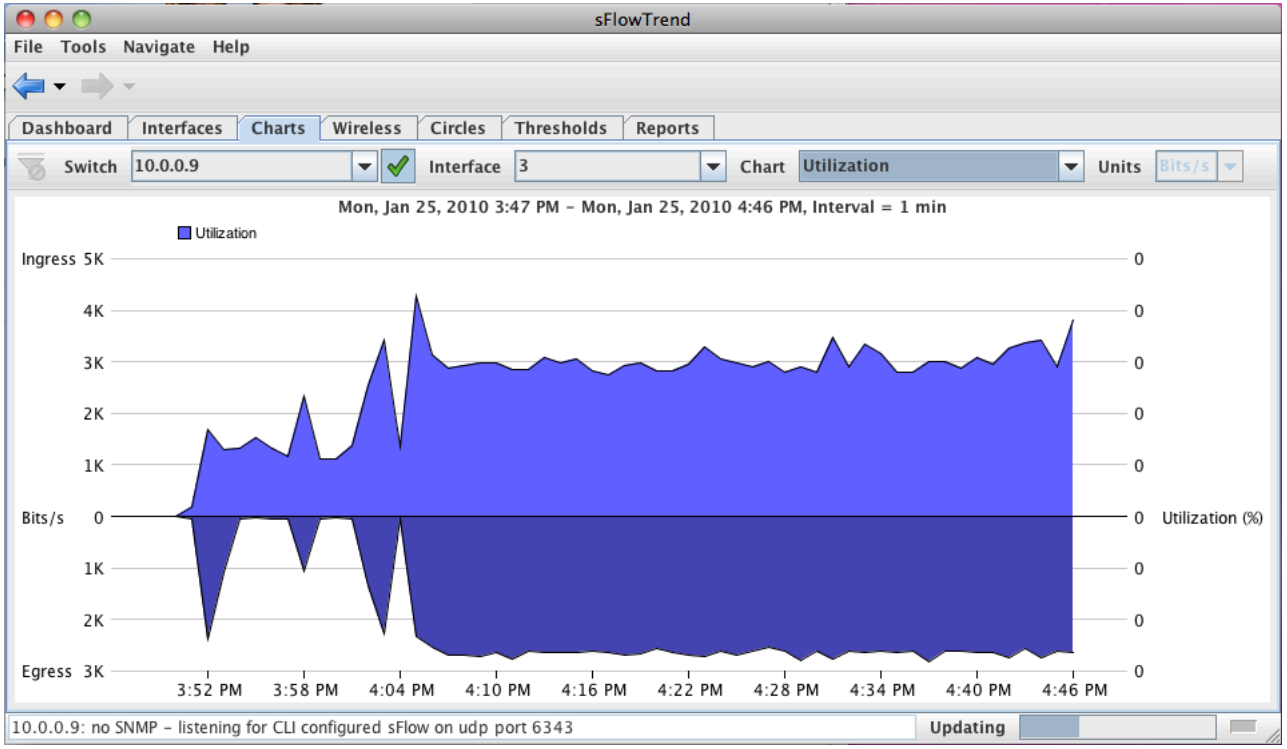

Switch 10.0.0.9 Interface 3 Chart Utilization Units Bits/s

Mon, Jan 25, 2010 3:47 PM – Mon, Jan 25, 2010 4:46 PM, Interval = 1 min

10.0.0.9: no SNMP – listening for CLI configured sFlow on udp port 6343

- Trending interface utilization a staple of network management
- sFlow is only realistic way to monitor virtual interface counters
  - sFlow counter export is scalable to support 200,000+ virtual ports (20-40 VMs per physical server makes this scalability a practical requirement in environments with 5,000 - 10,000 physical ports)

Traffic Sentinel in the NOC: Remote-site Management

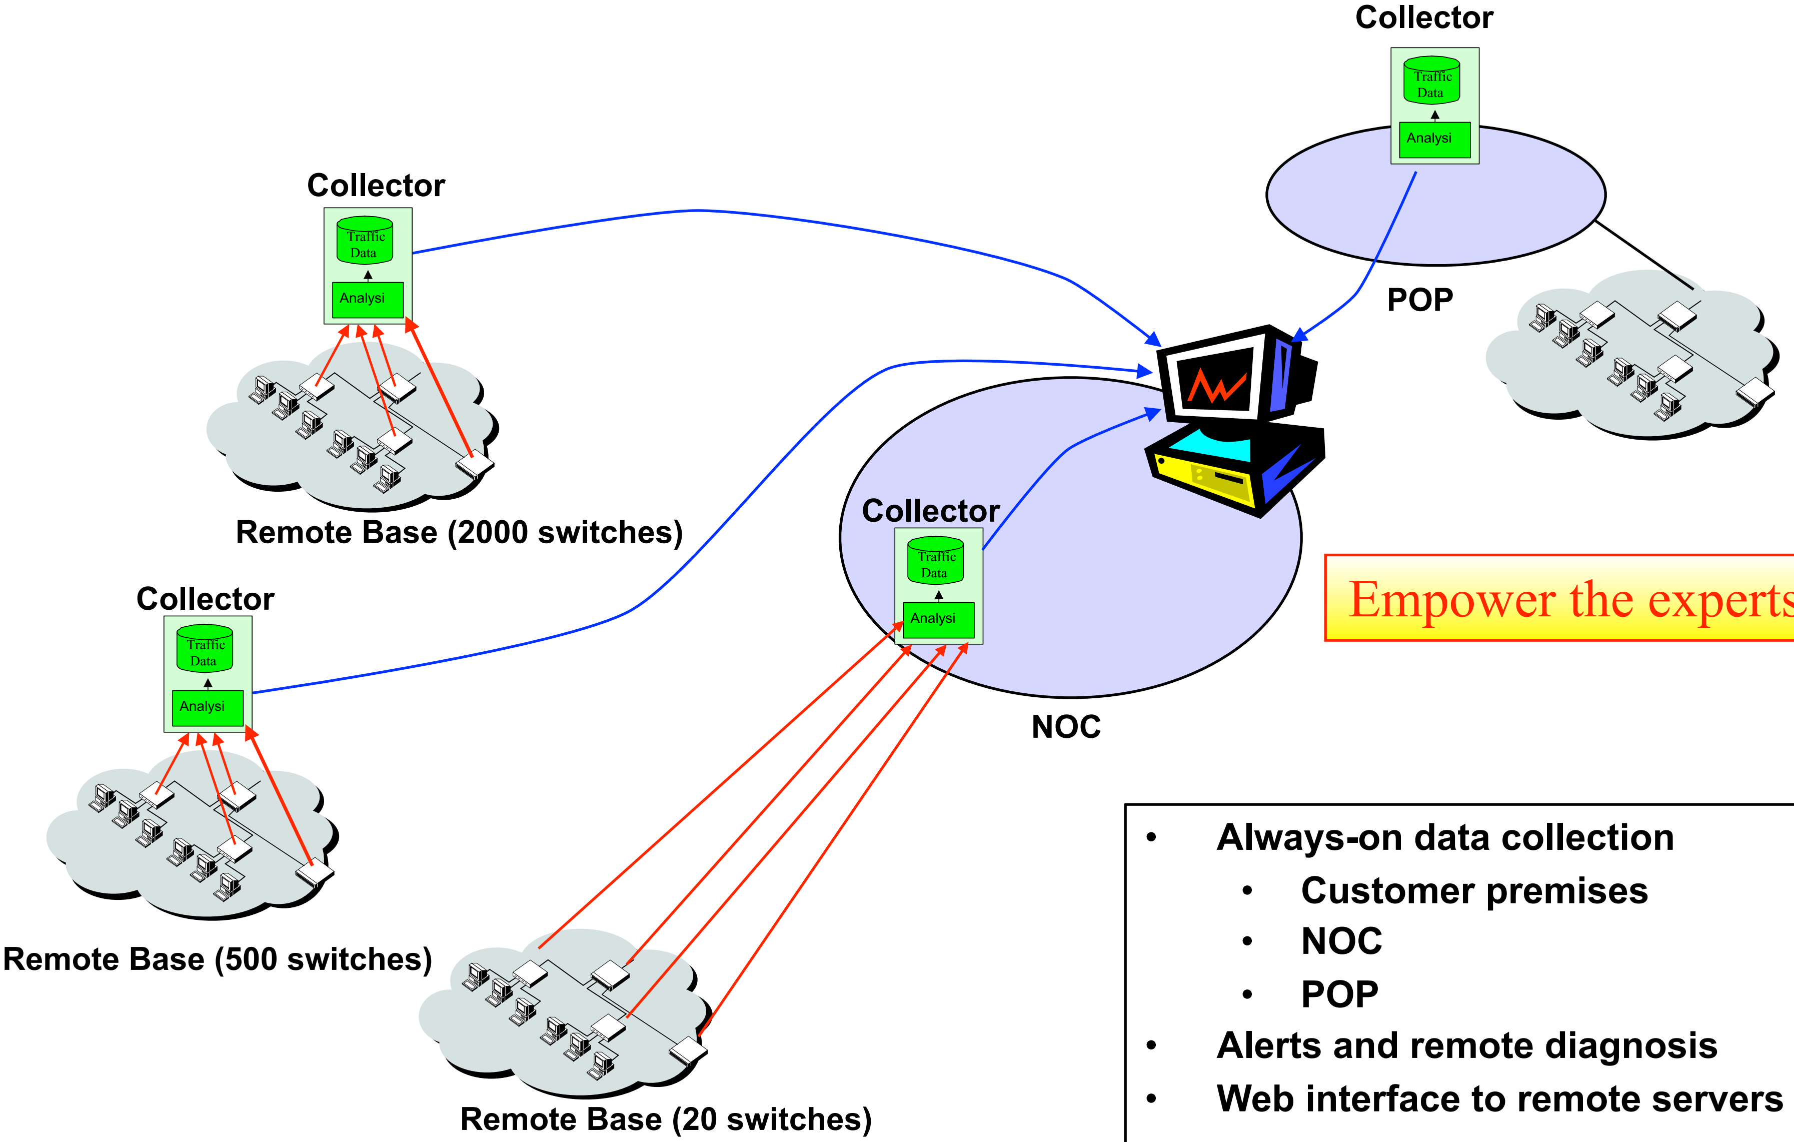

Empower the experts!

- Always-on data collection
  - Customer premises
  - NOC
  - POP
- Alerts and remote diagnosis
- Web interface to remote servers

sFlow vs NetFlow

- NetFlow gives partial visibility
  - Monitors routed L3 traffic only
  - Insufficient detail for effective LAN management
  - Significantly impacts switch/router performance
    - Cisco recommends monitoring of key interfaces only
  - Complex configuration
- sFlow can and should be enabled everywhere
  - sFlow monitors all traffic L2-L7, switched and routed, core to VM
  - Detailed data supports many applications
  - Negligible switch and network performance impact
  - Simple configuration
  - Collects interface counters too
  - Cross-layer measurements (e.g. MAC<->IP, GRE Tunneling)
  - Server and application monitoring too

More Information

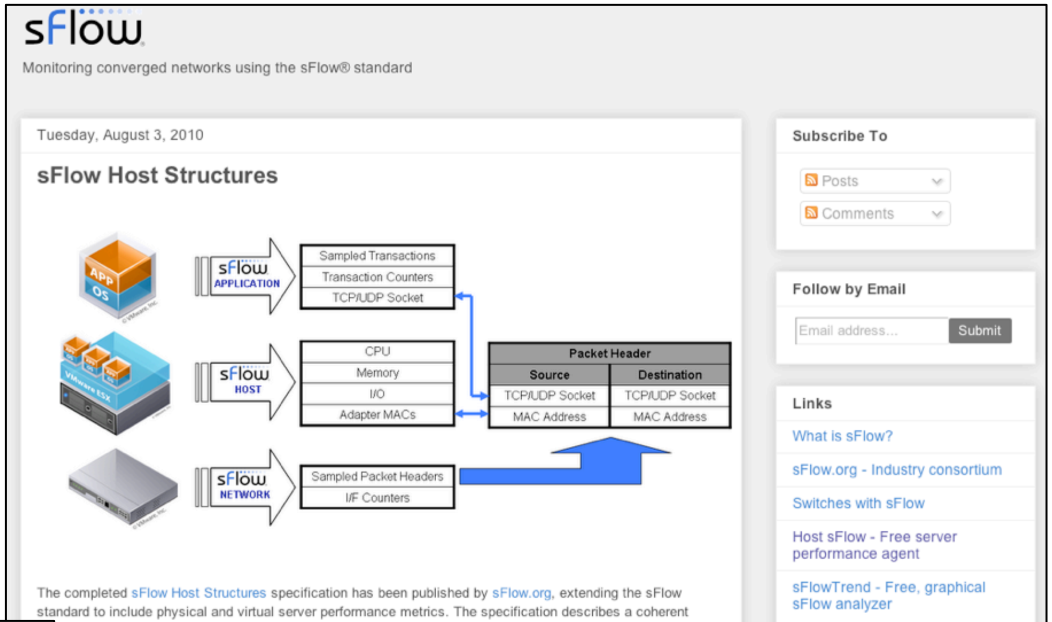

sFlow

Monitoring converged networks using the sFlow® standard

Tuesday, August 3, 2010

sFlow Host Structures

The completed sFlow Host Structures specification has been published by sFlow.org, extending the sFlow standard to include physical and virtual server performance metrics. The specification describes a coherent mework that builds on the sFlow metrics exported by most switch vendors, linking network, server and plication performance monitoring to provide an integrated picture of performance.

e diagram above shows how the packet header information exported by network devices is used to link

Subscribe To

Posts

Comments

Follow by Email

Email address... Submit

Links

- What is sFlow?
- sFlow.org - Industry consortium
- Switches with sFlow
- Host sFlow - Free server performance agent
- sFlowTrend - Free, graphical sFlow analyzer

Popular Posts

inmon.com

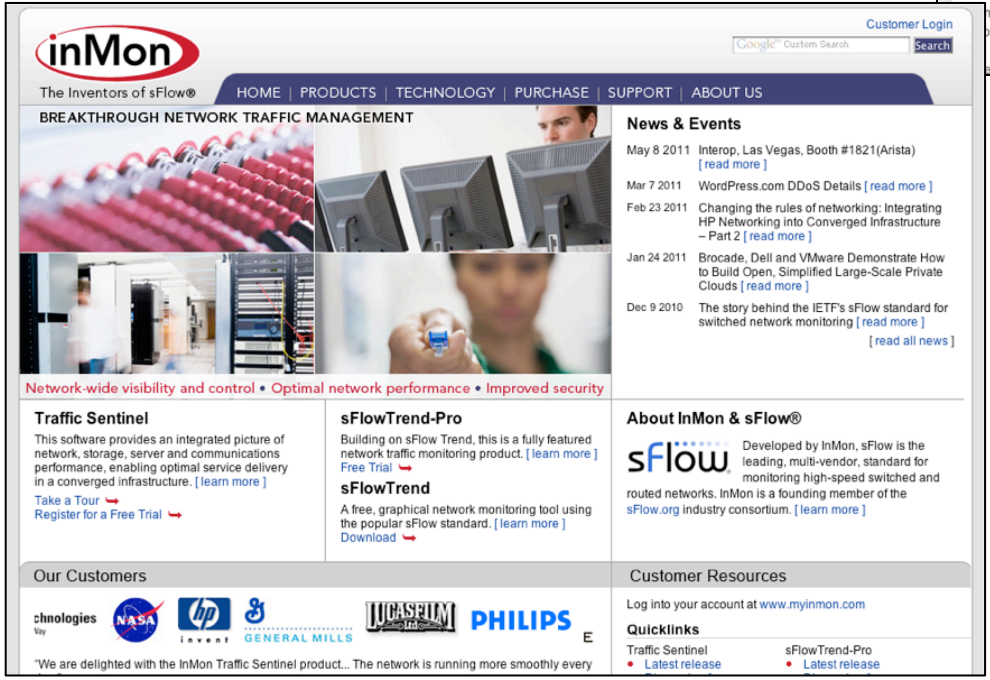

Customer Login

Google™ Custom Search Search

inMon

The Inventors of sFlow®

HOME | PRODUCTS | TECHNOLOGY | PURCHASE | SUPPORT | ABOUT US

BREAKTHROUGH NETWORK TRAFFIC MANAGEMENT

Network-wide visibility and control • Optimal network performance • Improved security

News & Events

May 8 2011	Interop, Las Vegas, Booth #1821(Arista) [read more]
Mar 7 2011	WordPress.com DDoS Details [read more]
Feb 23 2011	Changing the rules of networking: Integrating HP Networking into Converged Infrastructure – Part 2 [read more]
Jan 24 2011	Brocade, Dell and VMware Demonstrate How to Build Open, Simplified Large-Scale Private Clouds [read more]
Dec 9 2010	The story behind the IETF's sFlow standard for switched network monitoring [read more]

[read all news]

Traffic Sentinel

This software provides an integrated picture of network, storage, server and communications performance, enabling optimal service delivery in a converged infrastructure. [learn more]

Take a Tour ➥
Register for a Free Trial ➥

sFlowTrend-Pro

Building on sFlow Trend, this is a fully featured network traffic monitoring product. [learn more]

Free Trial ➥

sFlowTrend

A free, graphical network monitoring tool using the popular sFlow standard. [learn more]

Download ➥

About InMon & sFlow®

sFlow Developed by InMon, sFlow is the leading, multi-vendor, standard for monitoring high-speed switched and routed networks. InMon is a founding member of the sFlow.org industry consortium. [learn more]

Our Customers

"We are delighted with the InMon Traffic Sentinel product... The network is running more smoothly every

Customer Resources

Log into your account at www.myinmon.com

Quicklinks

Traffic Sentinel	sFlowTrend-Pro
• Latest release	• Latest release

blog.sflow.com