



How Big Data Helps Manage Big Networks

Making Sense Of Machine Chatter

Tom Griffin

Director – Systems Engineering - EMEA

- Wikipedia
...big data is a collection of data sets so large and complex that it becomes difficult to process using on-hand data management tools
- Gartner
Its not about the size of the data, but what you do with it.
- Data is only as good as the decisions it helps us make

What Defines the Performance of Big Data – The 4 Vs!

- **Volume of Data**
 - Historical Storage
 - Rate of Incoming Data
- **Velocity of Analytics**
 - Real time analytics
 - Reporting
- **Variety**
 - Big Data is Any Data
 - Analyzing different data together yields better insights
- **Veracity**
 - Big Data is about decisions
 - You can't act on data you don't trust

The poster children of Big Data

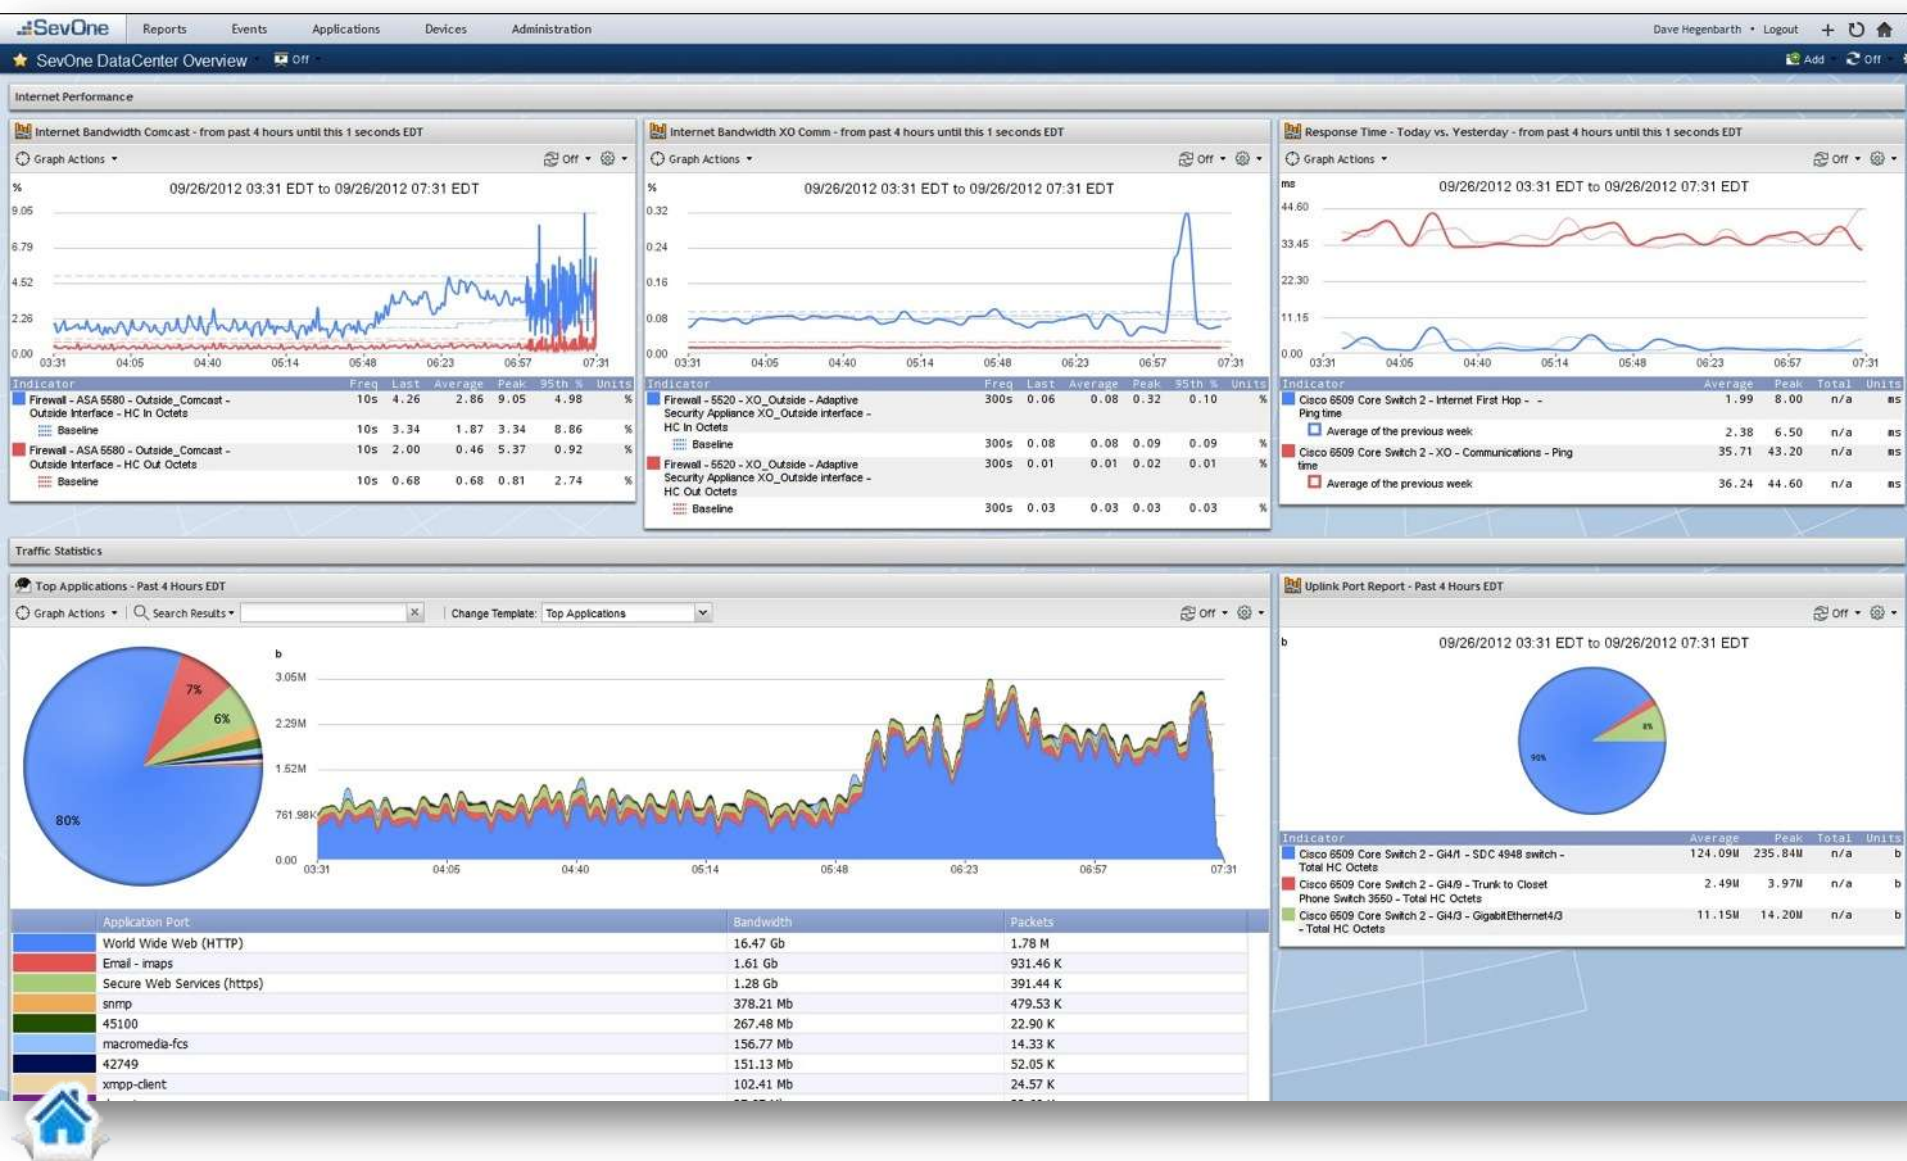
- Social Media
- Financial Trading
- Telecommunications

Big Data is connected directly to the industry's revenue models. Being able to extract more value from the data the organization is capable of generating is a distinct competitive advantage and determines the success.

Machine to Machine (m2m) Data Is Growing

- Thousands of devices generate millions of metrics (50 billion connected devices by 2021)
 - Temperature
 - Bandwidth
 - Power
 - RFID
- **Size of Data**
 - Variety and Nature of Data
 - Poll vs. Stream
- **Time to Respond**
 - When is it two minutes too late

IT Performance Monitoring



Traffic Statistics

Top Applications - Past 4 Hours EDT

Application	Port	Bandwidth	Packets
World Wide Web (HTTP)		16.47 Gb	1.78 M
Email - imaps		1.61 Gb	931.46 K
Secure Web Services (https)		1.28 Gb	391.44 K
snmp		378.21 Mb	479.53 K
45100		267.48 Mb	22.90 K
macromedia-fcs		156.77 Mb	14.33 K
42749		151.13 Mb	52.05 K
xmpp-client		102.41 Mb	24.57 K

Uplink Port Report - Past 4 Hours EDT

Indicator	Average	Peak	Total	Units
Cisco 8509 Core Switch 2 - Gi4/9 - SDC 4948 switch - Total HC Octets	124.09M	235.84M	n/a	b
Cisco 8509 Core Switch 2 - Gi4/9 - Trunk to Closet Phone Switch 3550 - Total HC Octets	2.49M	3.97M	n/a	b
Cisco 8509 Core Switch 2 - Gi4/3 - GigabitEthernet4/3 - Total HC Octets	11.15M	14.20M	n/a	b

Application Monitoring

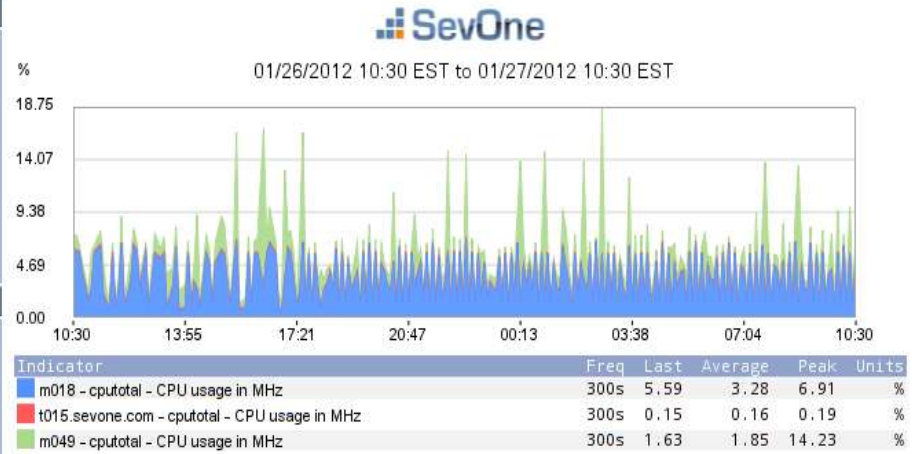
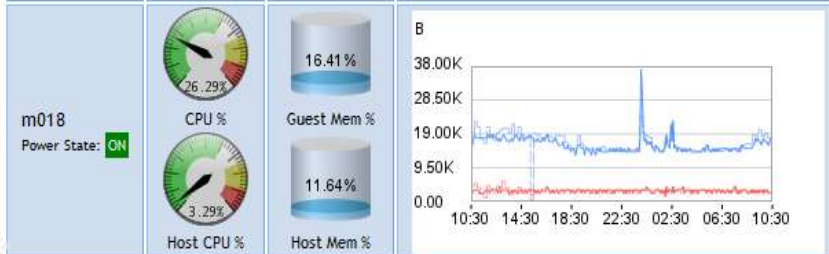

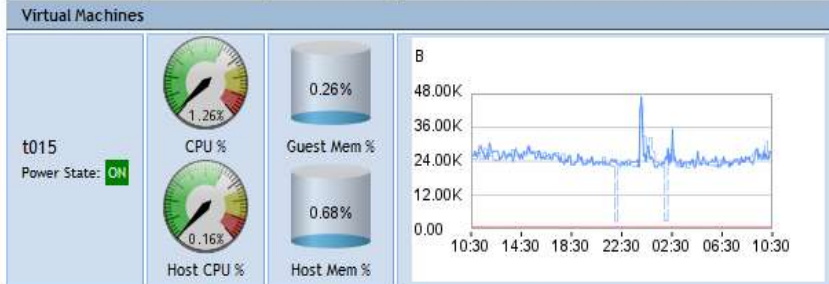
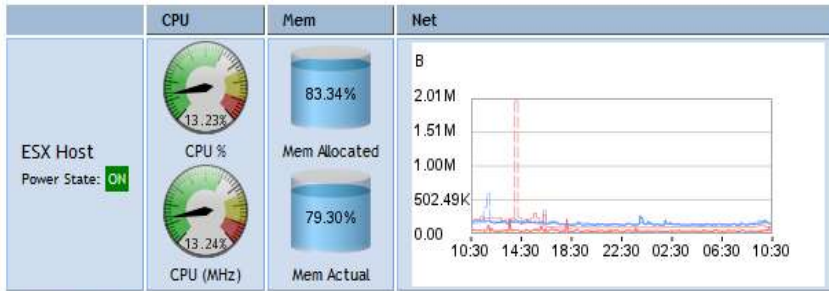


All agentless monitoring

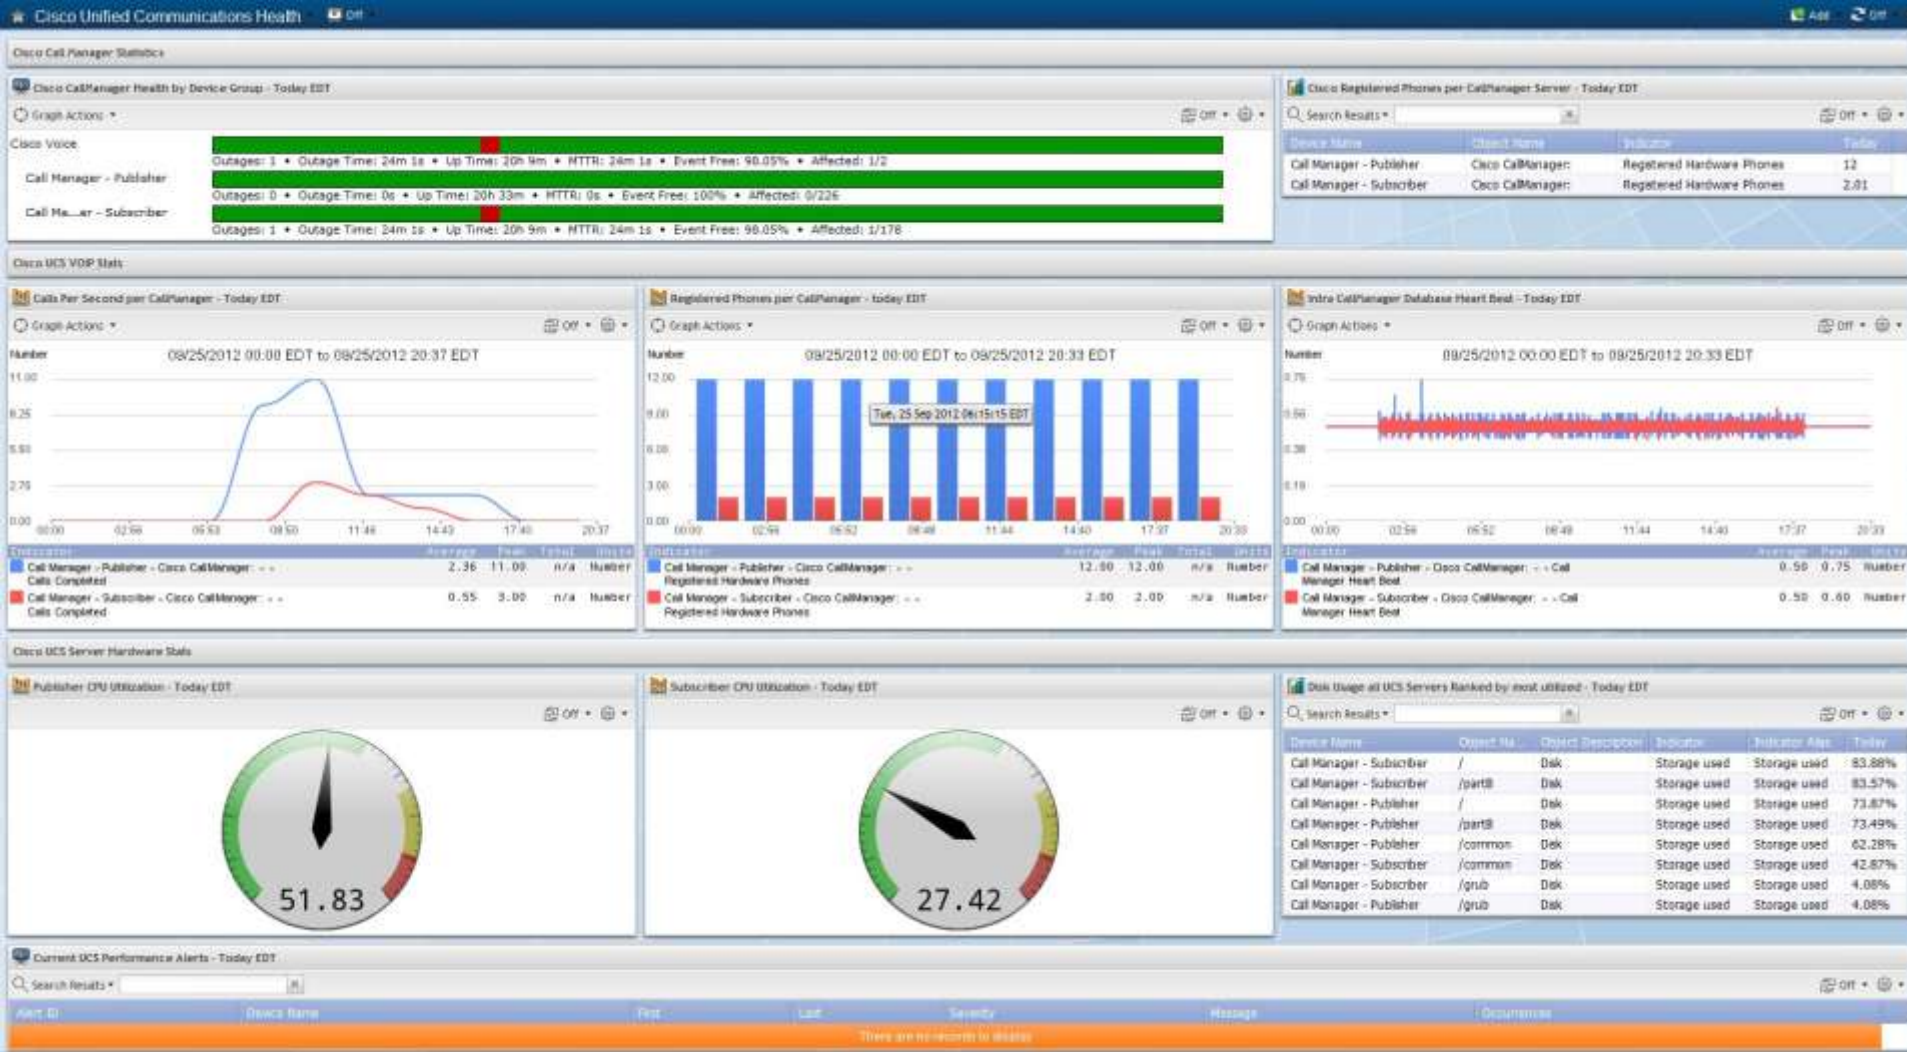


Virtual Infrastructure Monitoring

VMware Host Status

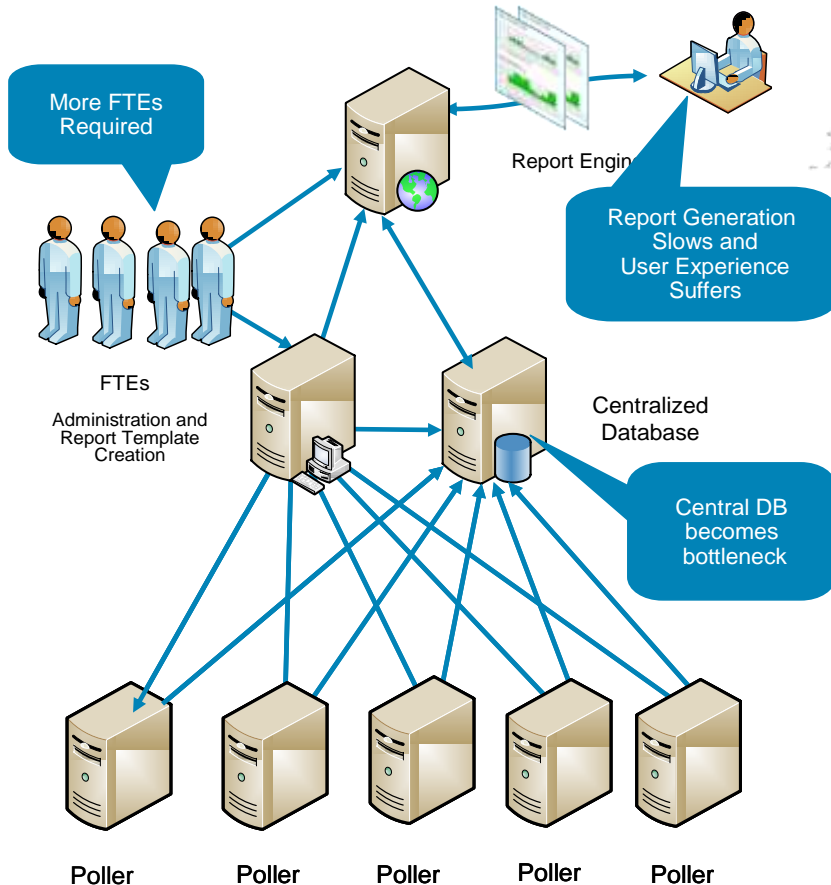


Voice Monitoring

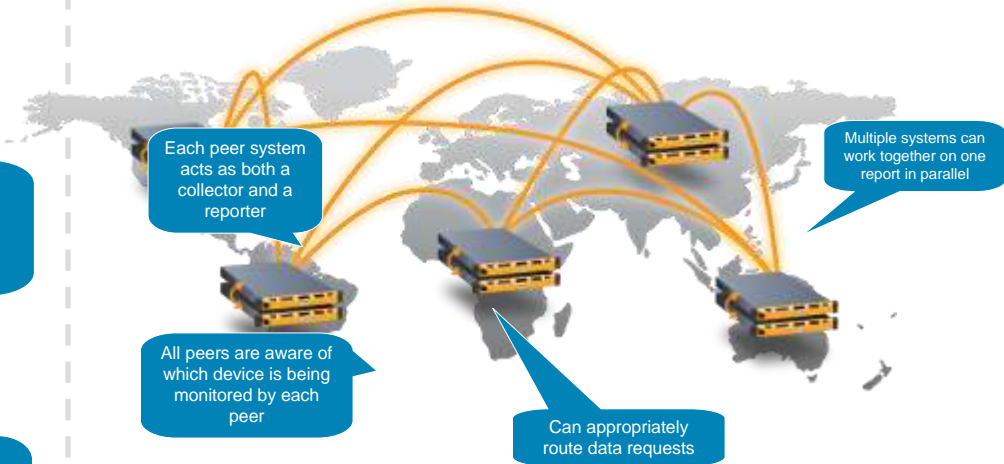


Existing IT Management Solutions Don't Scale, SevOne is Big Data Proven

Traditional Network Management Architecture



The SevOne Cluster: Scale and Speed for BIG Data

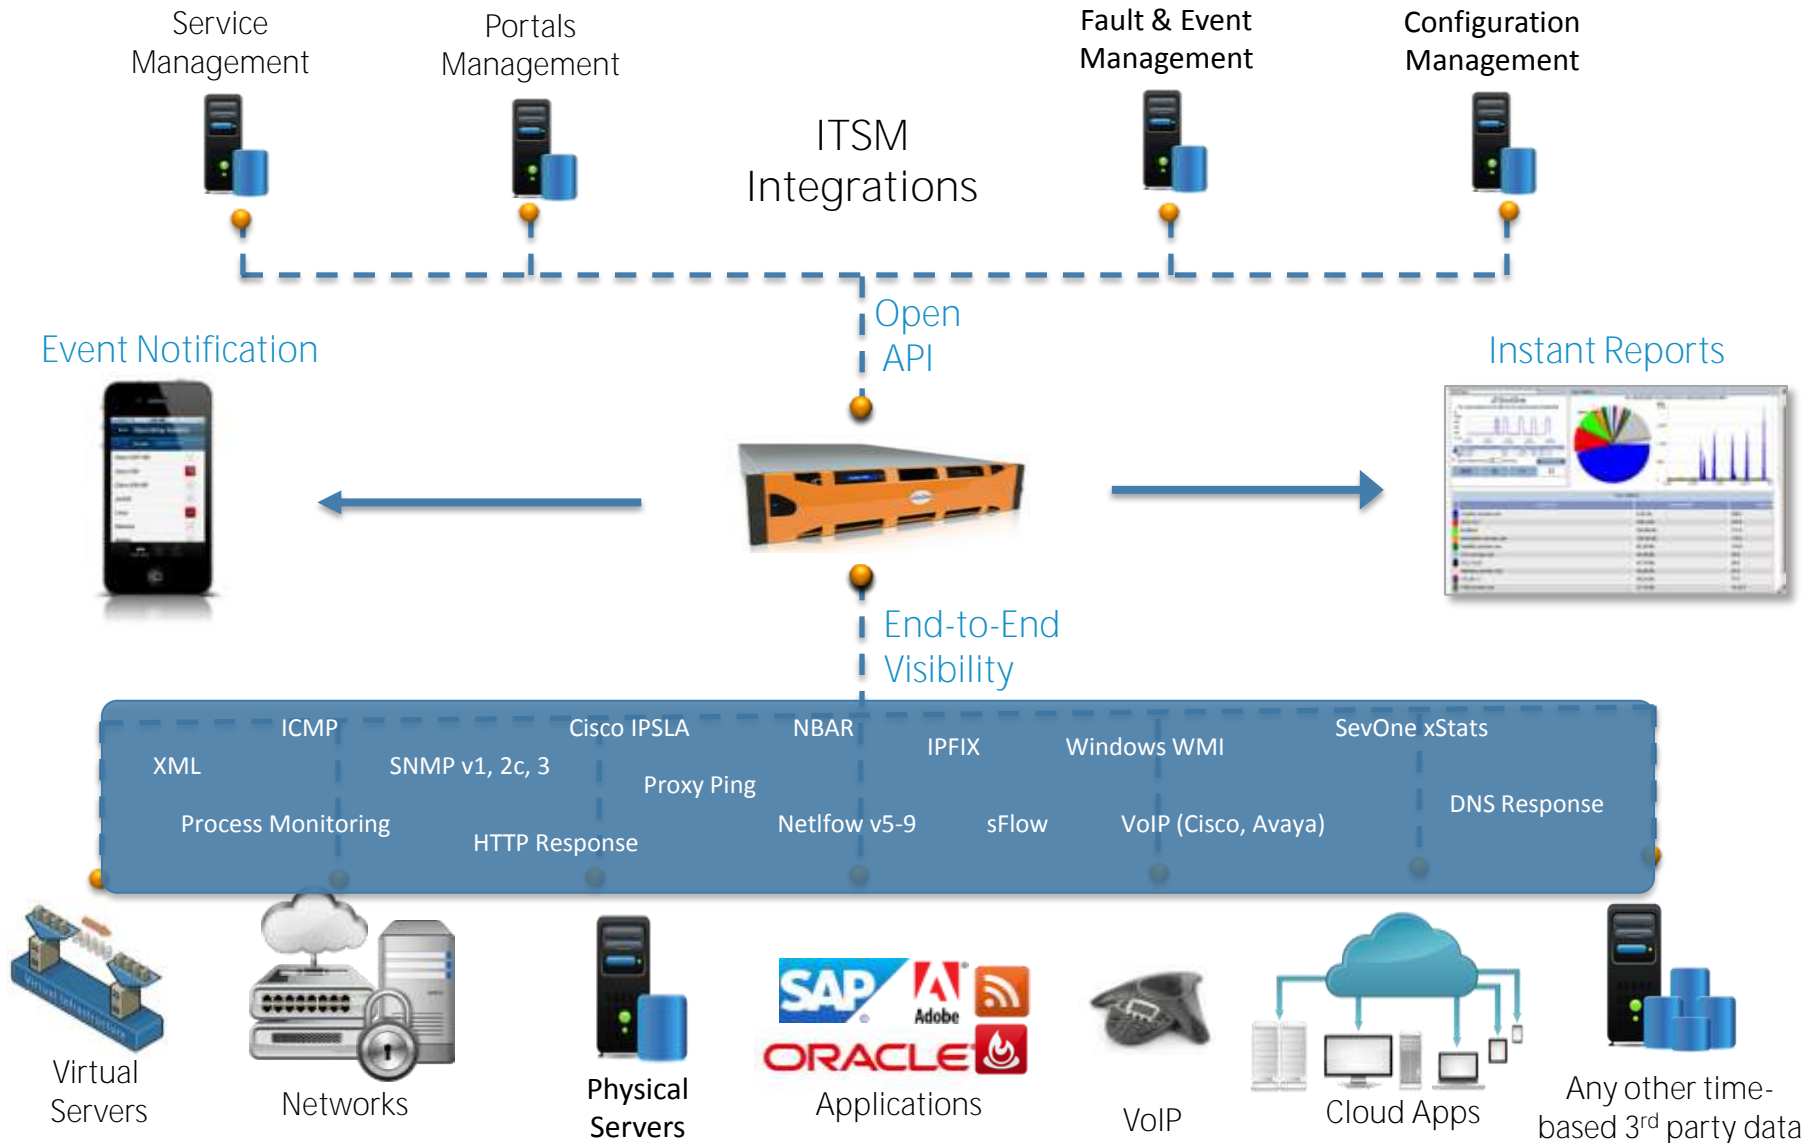


Linear scalability to millions of objects and billions of baselines

Distributed collection and reporting, with no limits

Appliance deployment, footprint 2x-4x smaller than competition

SevOne Open Architecture



- Distribute ALL Processing – Collection and Reporting
- Push Analytics To The Edge
 - Collectors should be able to participate as peers in analysis and reporting
- Push the Edge to Edge
 - Collectors should be close to data sources for timing, scalability and control data reasons.
- Support Many Data Sources - Normalize Data Early
- Parallelize Everything
- Minimize footprint
 - Bandwidth, Data Center
- Keep it Flat
- Keep it Open

SevOne – Defining BIG Data



PayPal

Total Data Store: 100 TB

Twitter

New Data: 4962 Tweets per second
Peak: 25,000 Tweets per second

Facebook

Queries/sec Peak: 13,000,000
Changed Rows/sec Peak: 3,500,000

SevOne

Existing SevOne Production clusters

New Data per Node (base): 21,000/s
New Data per Node (peak): 200,000/s

New Data per Cluster (base): 2,100,000/s
New Data per Cluster (peak): 20,000,000/s

New Data Volume: 1TB+ / day
Total Data Store: 100TB+

Number of Nodes: about 100

Combined Xerox and Comcast Clusters are processing more new data rows at base load than Facebook

DNC 1000HF (single appliance)

- 15,000,000 Flows / minute
- 4,500,000 rows of data / minute

Source: ITNews – <http://www.itnews.com.au/News/317811,twitter-paypal-reveal-database-performance.aspx>

Who Uses This?

Telecom & MSP



Banking & Finance



THOMSON
REUTERS

CREDIT SUISSE

Morgan Stanley



News & Media



Many Others...





CASE STUDY

- Initial Capacity 16,000,000 KPIs / min
- Current Capacity 160,000,000 KPIs / min
- Projected Capacity for Full Deployment 2,500,000,000 KPIs / min
- Cut performance operations cost by over 75%
- Inefficient service based model w/ 8 FTEs replaced by a self-service model w/1/2 FTE
- now thousands of self-service users
- SevOne is now Comcast's performance "source of truth" & common communication platform

"SevOne has never failed to identify a potential degradation or outage"

"With SevOne, we now can access real-time reports in seconds that use to take hours"



Thank You