Cisco
Connect

Self-driving Datacenter: Analytics

George Boulescu

Consulting Systems Engineer

19/10/2016

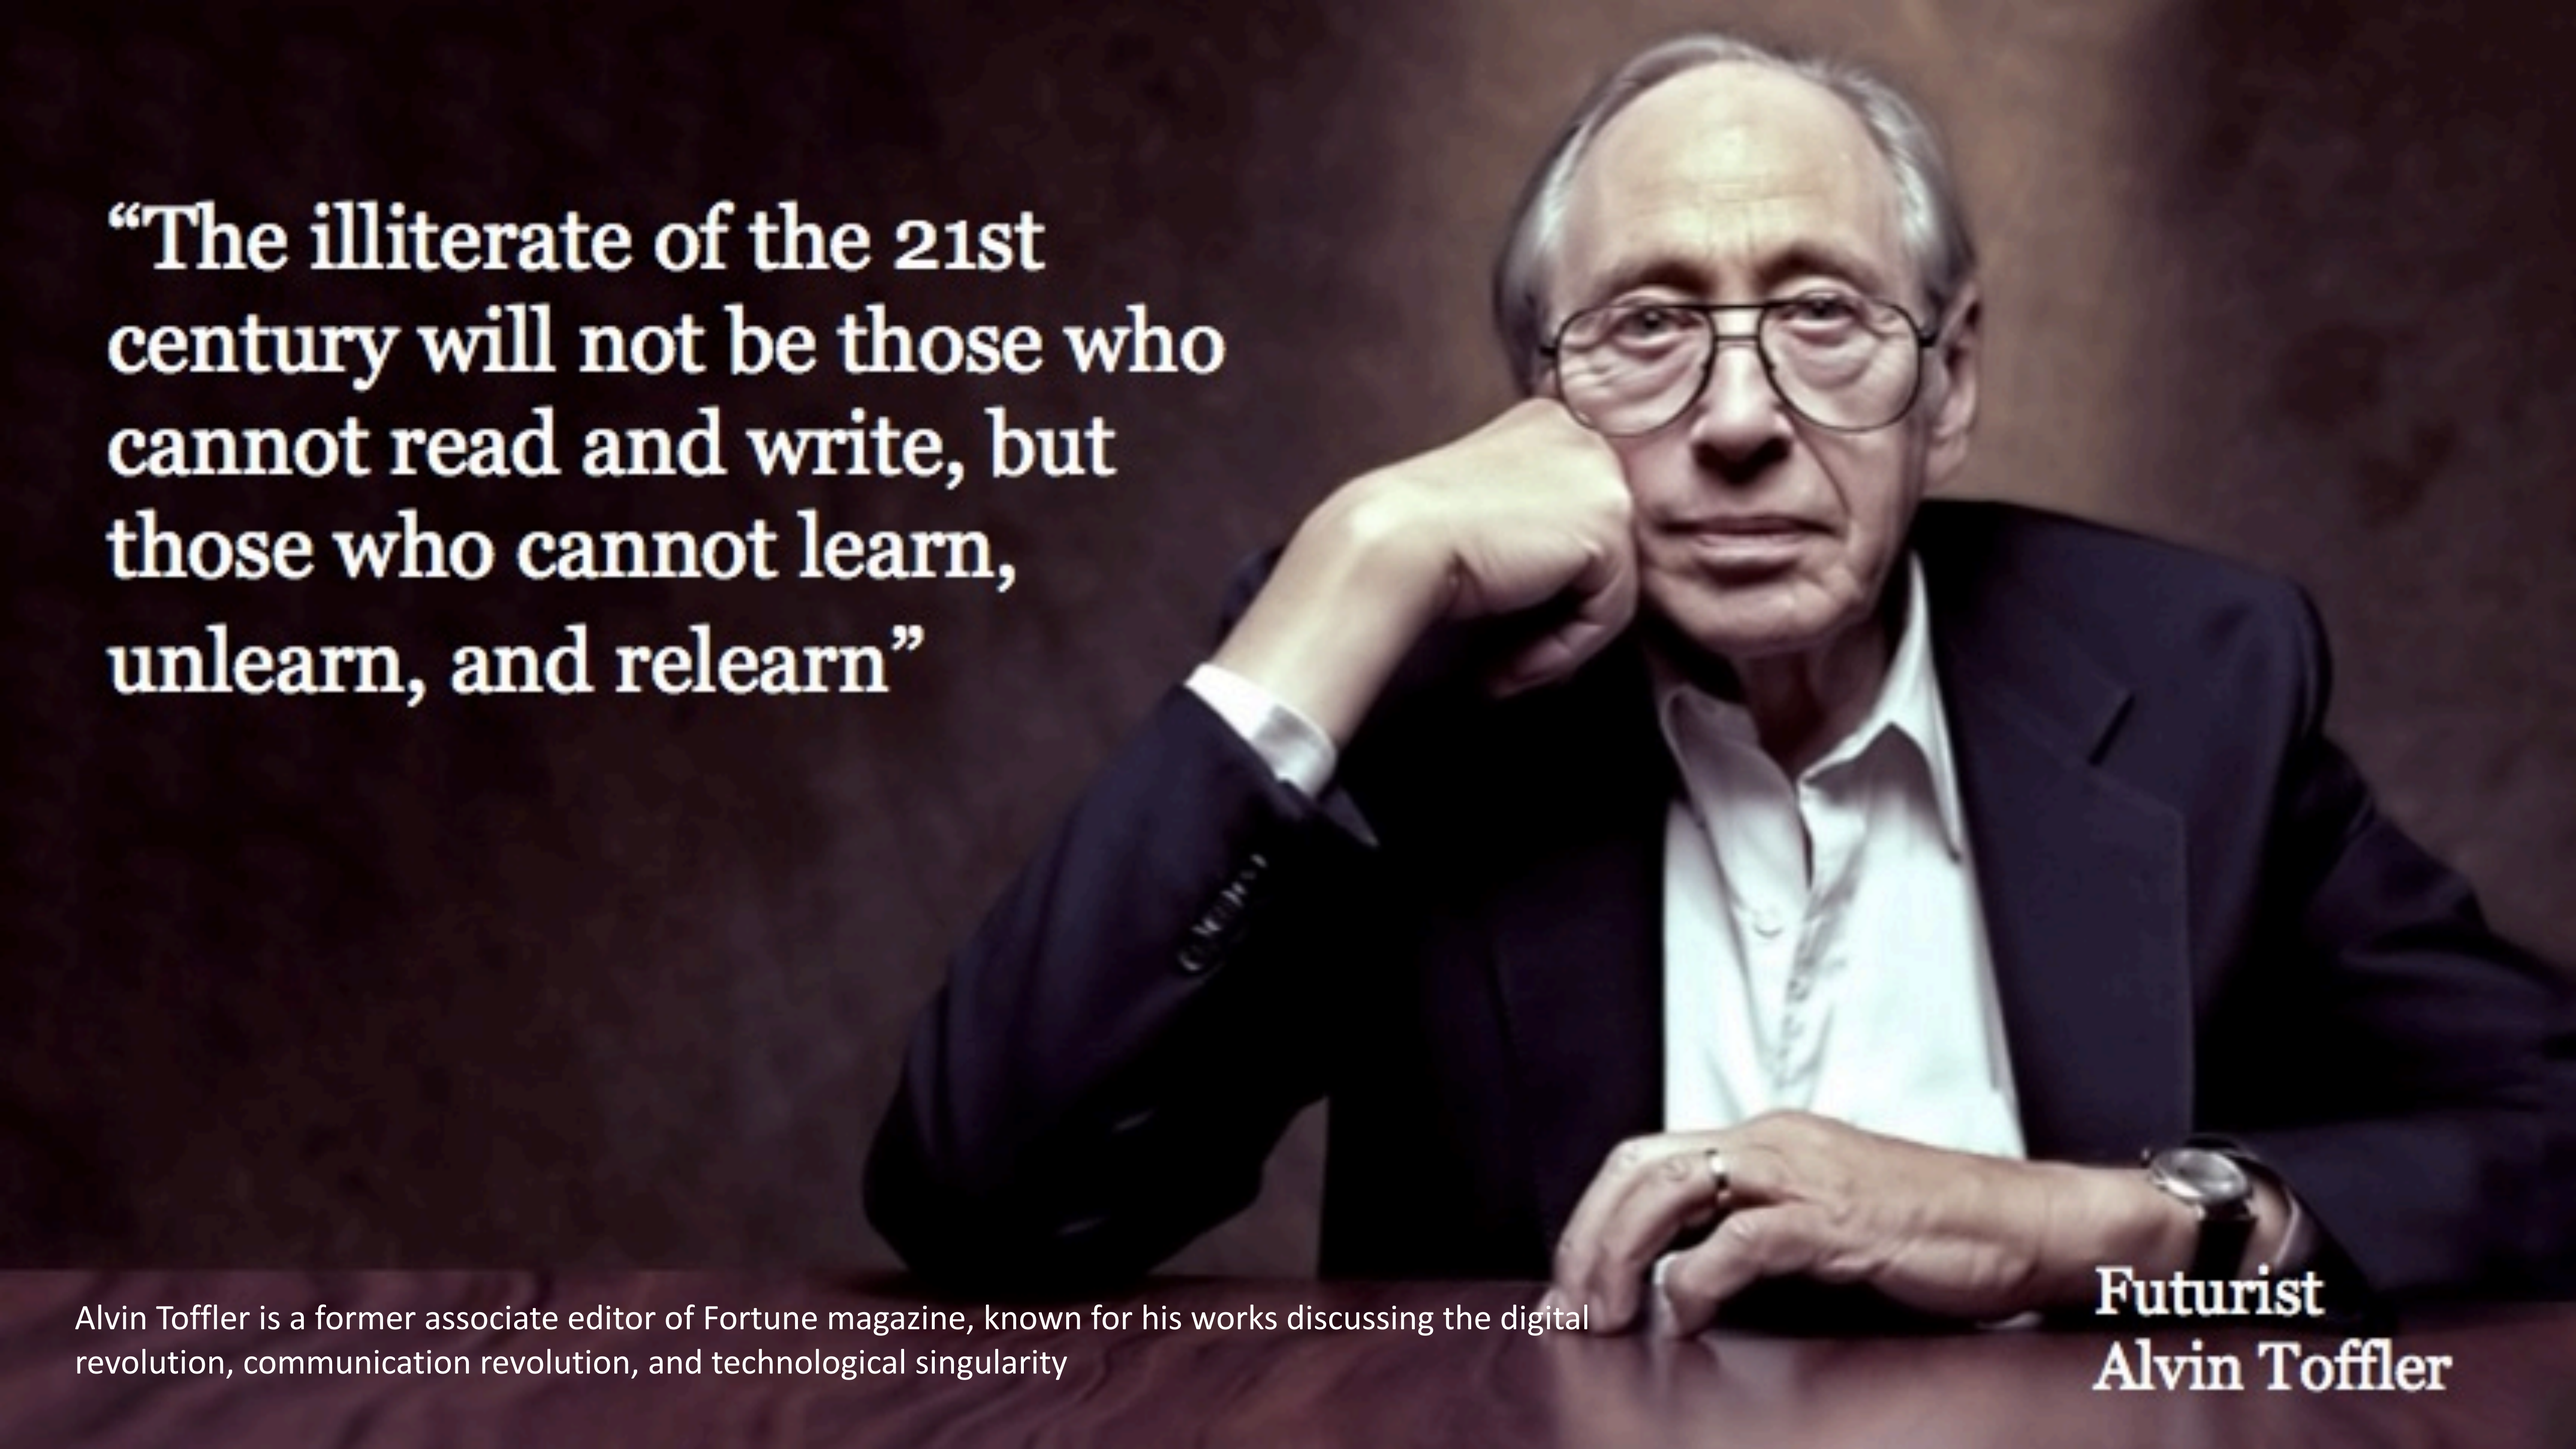

Alvin Toffler is a former associate editor of Fortune magazine, known for his works discussing the digital revolution, communication revolution, and technological singularity

Define Security

The conscious or unconscious acceptance of a risk in relation of the probability of this becoming to be reality in a delta Time …

Datacenter Evolution

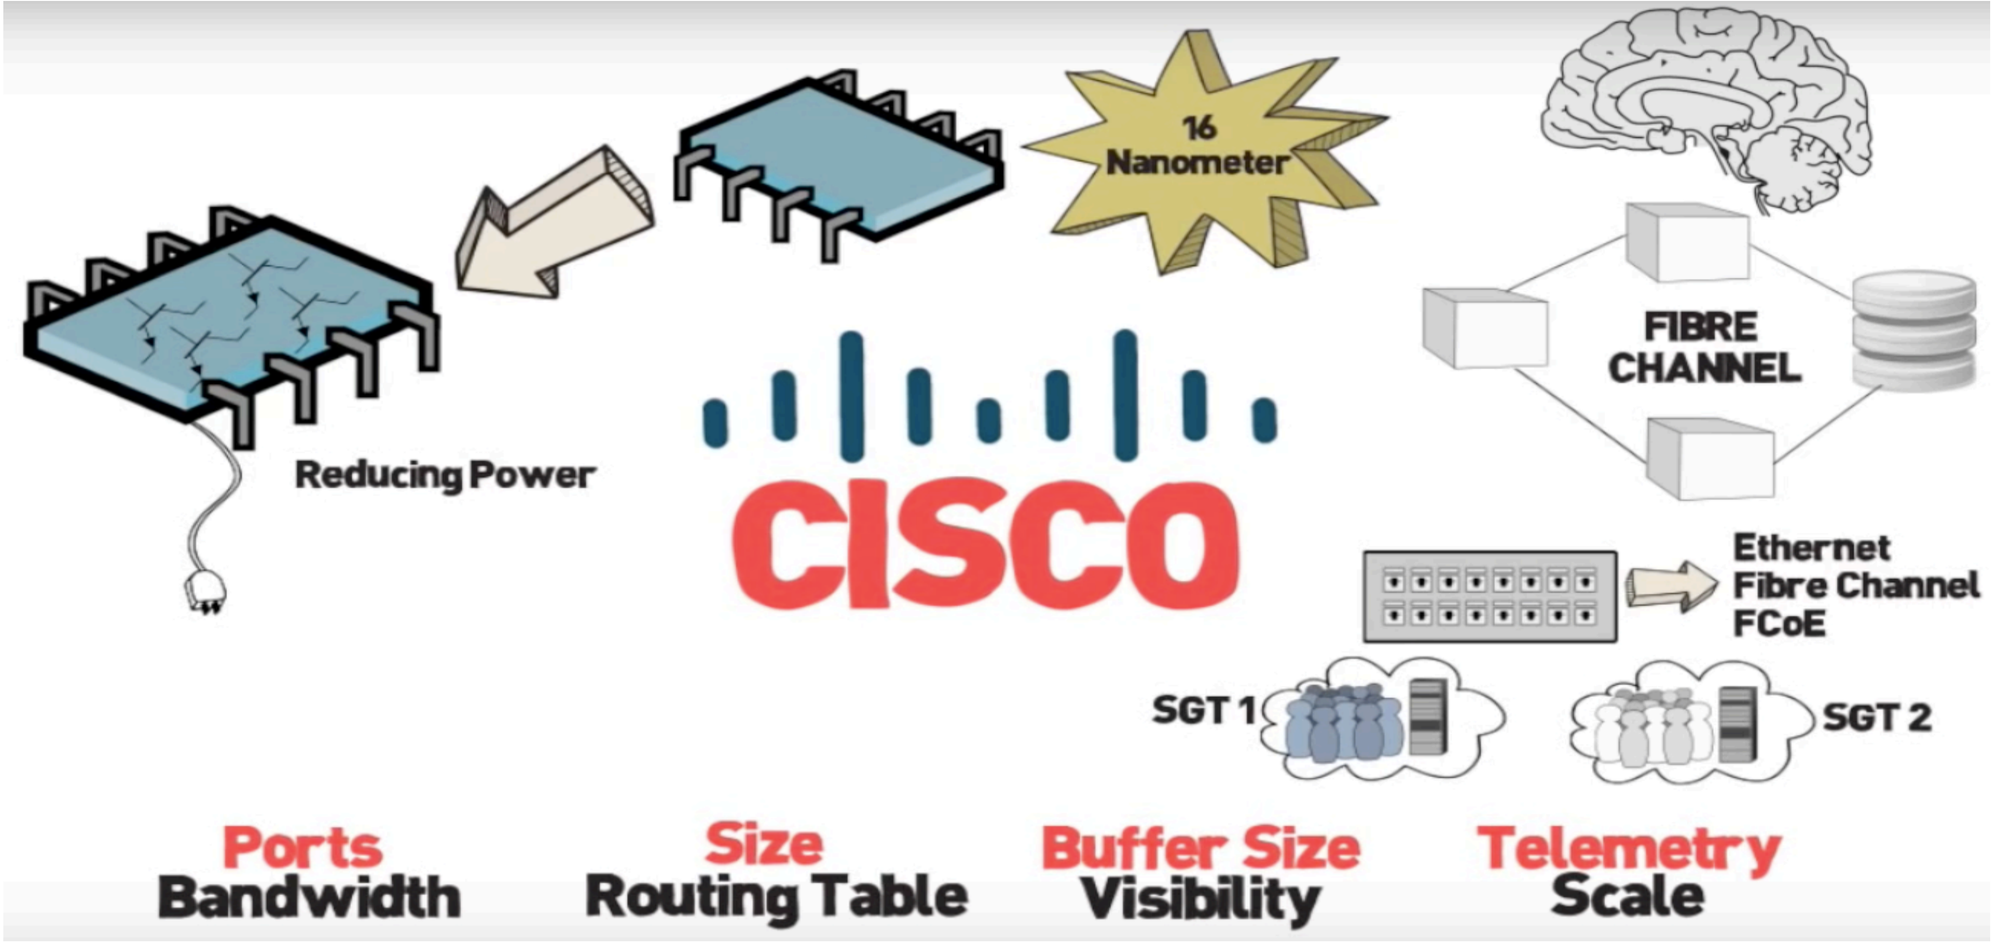

Datacenter Evolution

**"Everyone in the office is getting a netbook.
Little screens make our problems look smaller!"**

We Are at the Cusp of a Major Shift

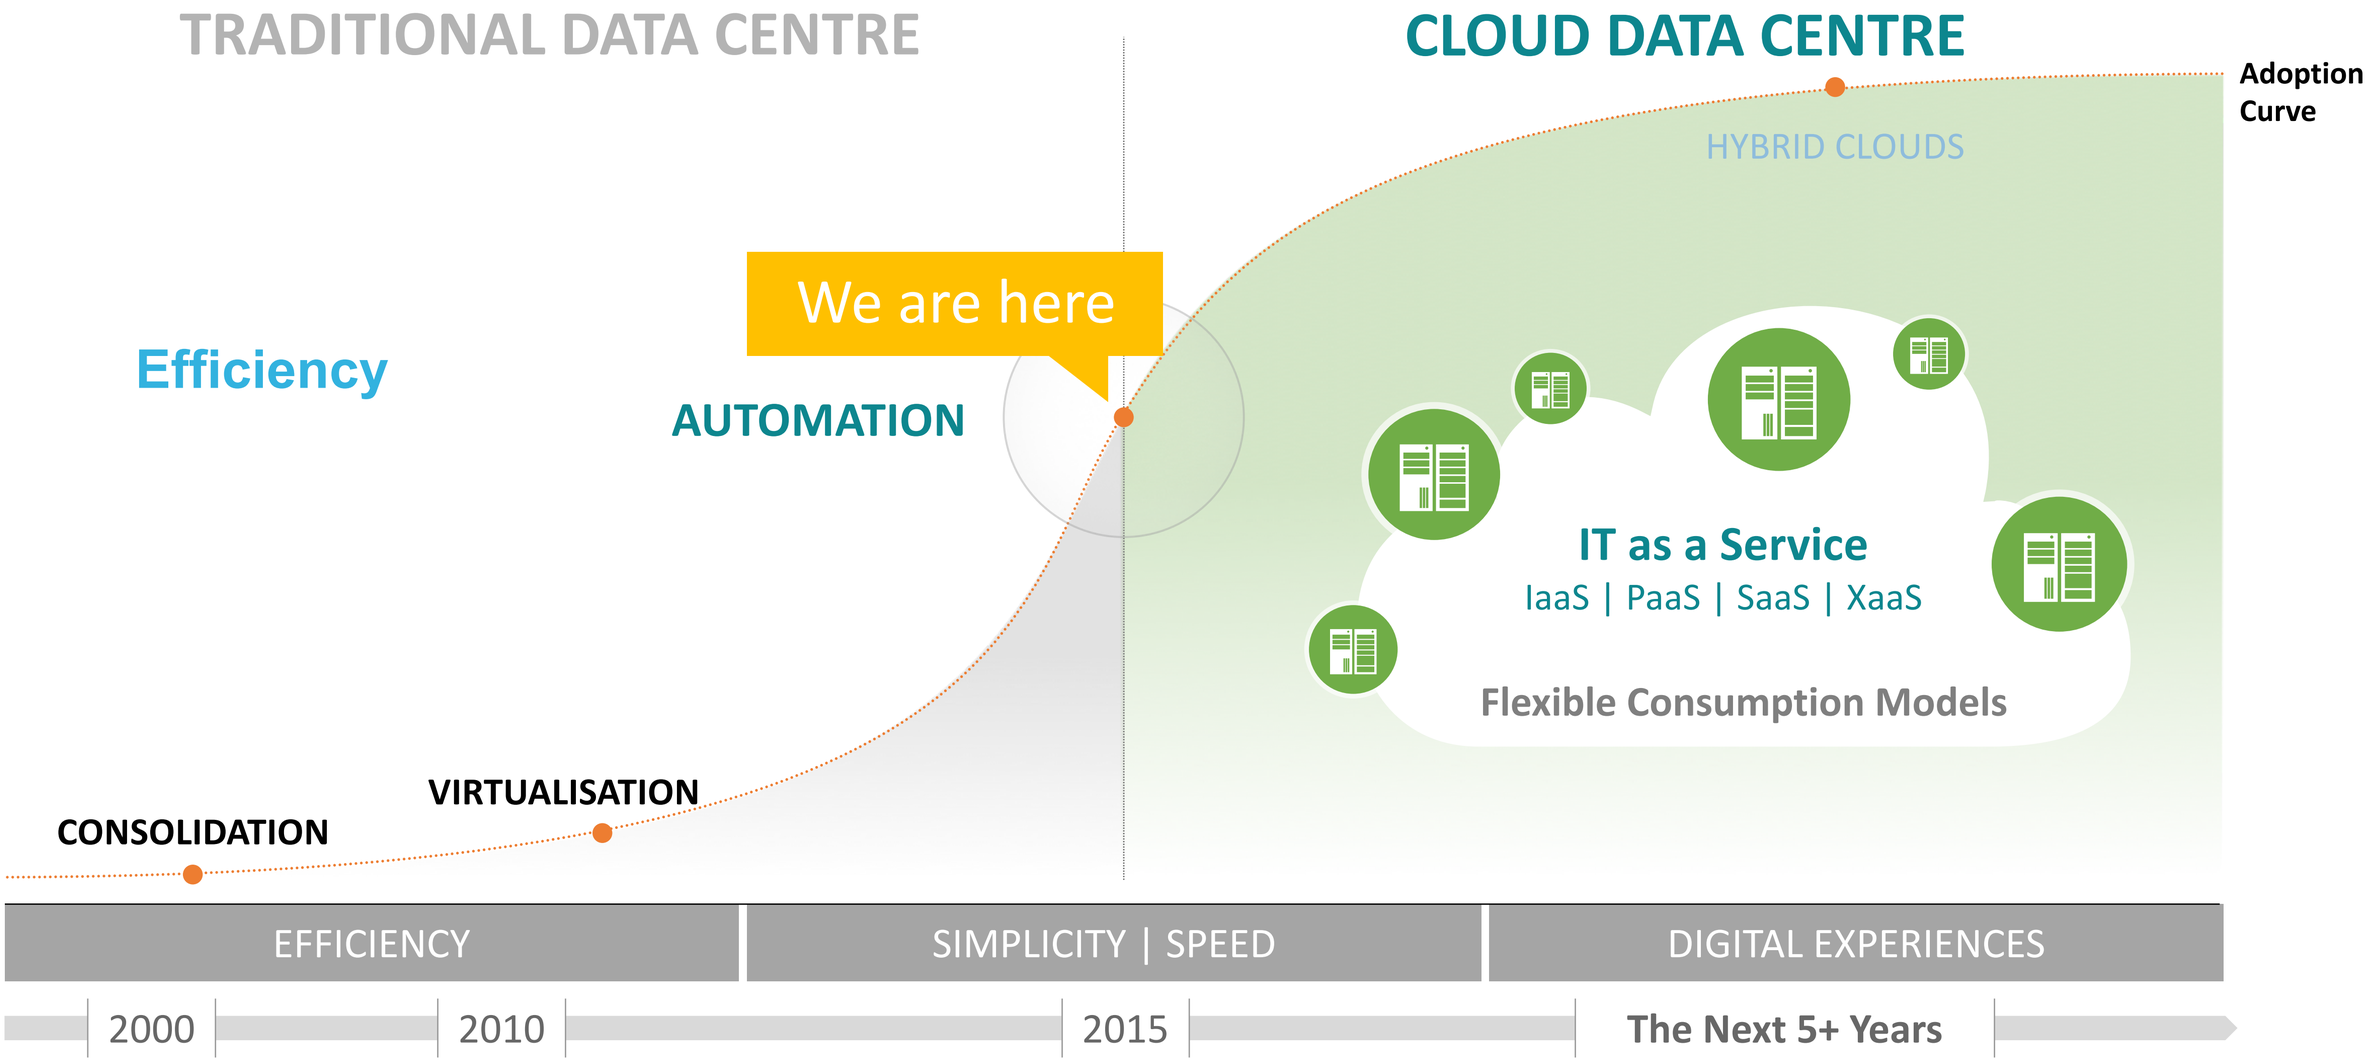

Modern data centers are getting increasingly complex

Big and fast data

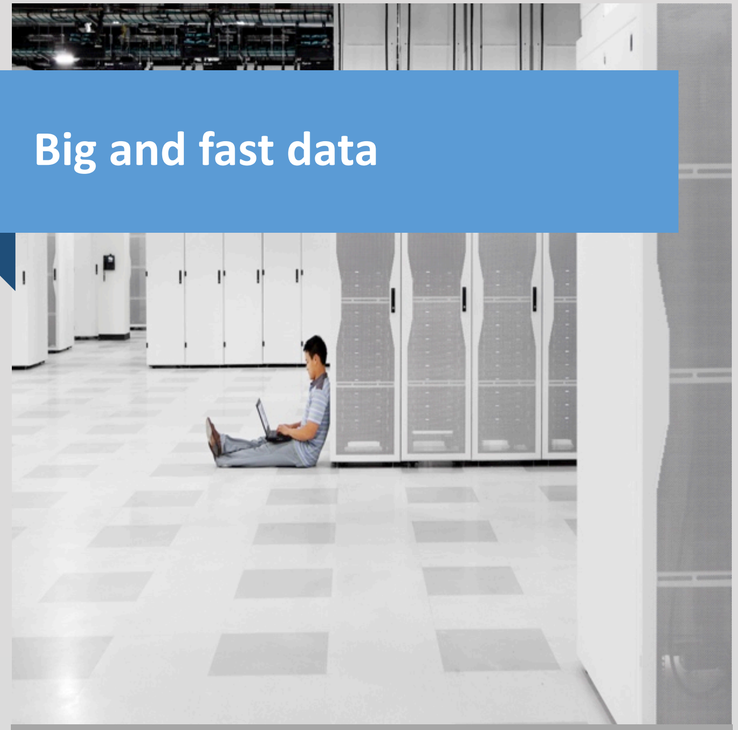

- Increase in east-west traffic
- Expanded attack surface
- Open source

Hybrid cloud

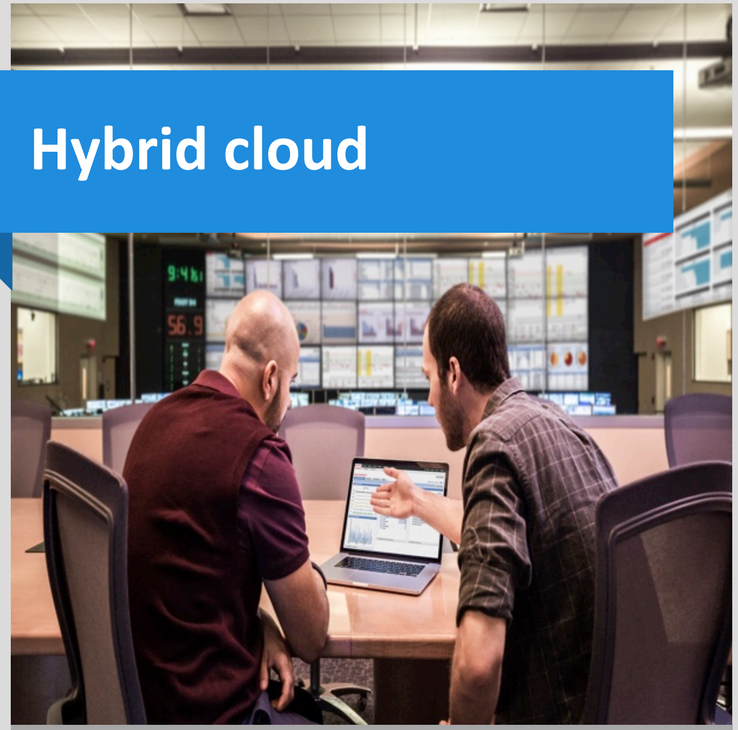

- Zero trust model
- Multi cloud orchestration
- Application portability

Rapid app deployment

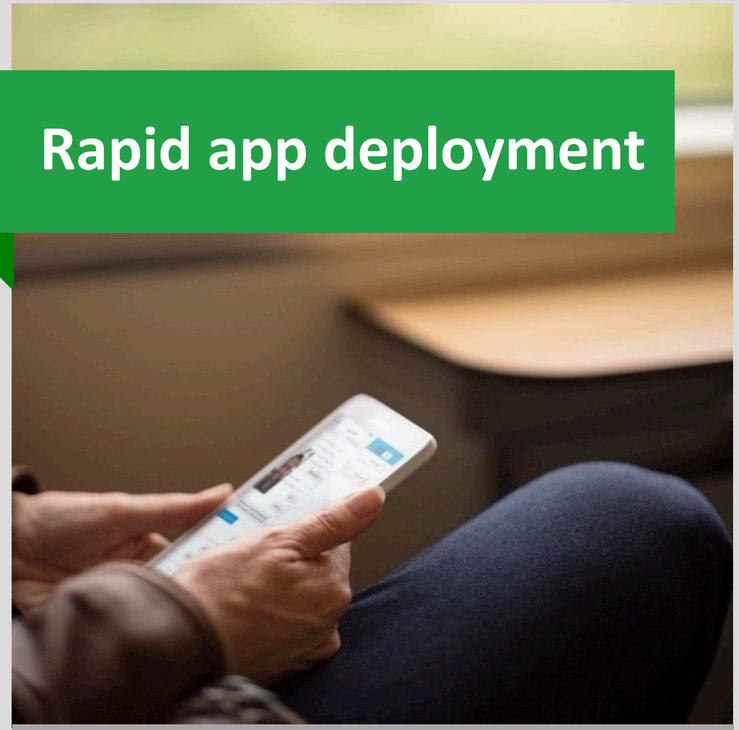

- Continuous development
- Application mobility
- Micro services

What if you could actually **look at every data packet header** that has ever traversed the network without sampling?

ACI Architecture

Analytics

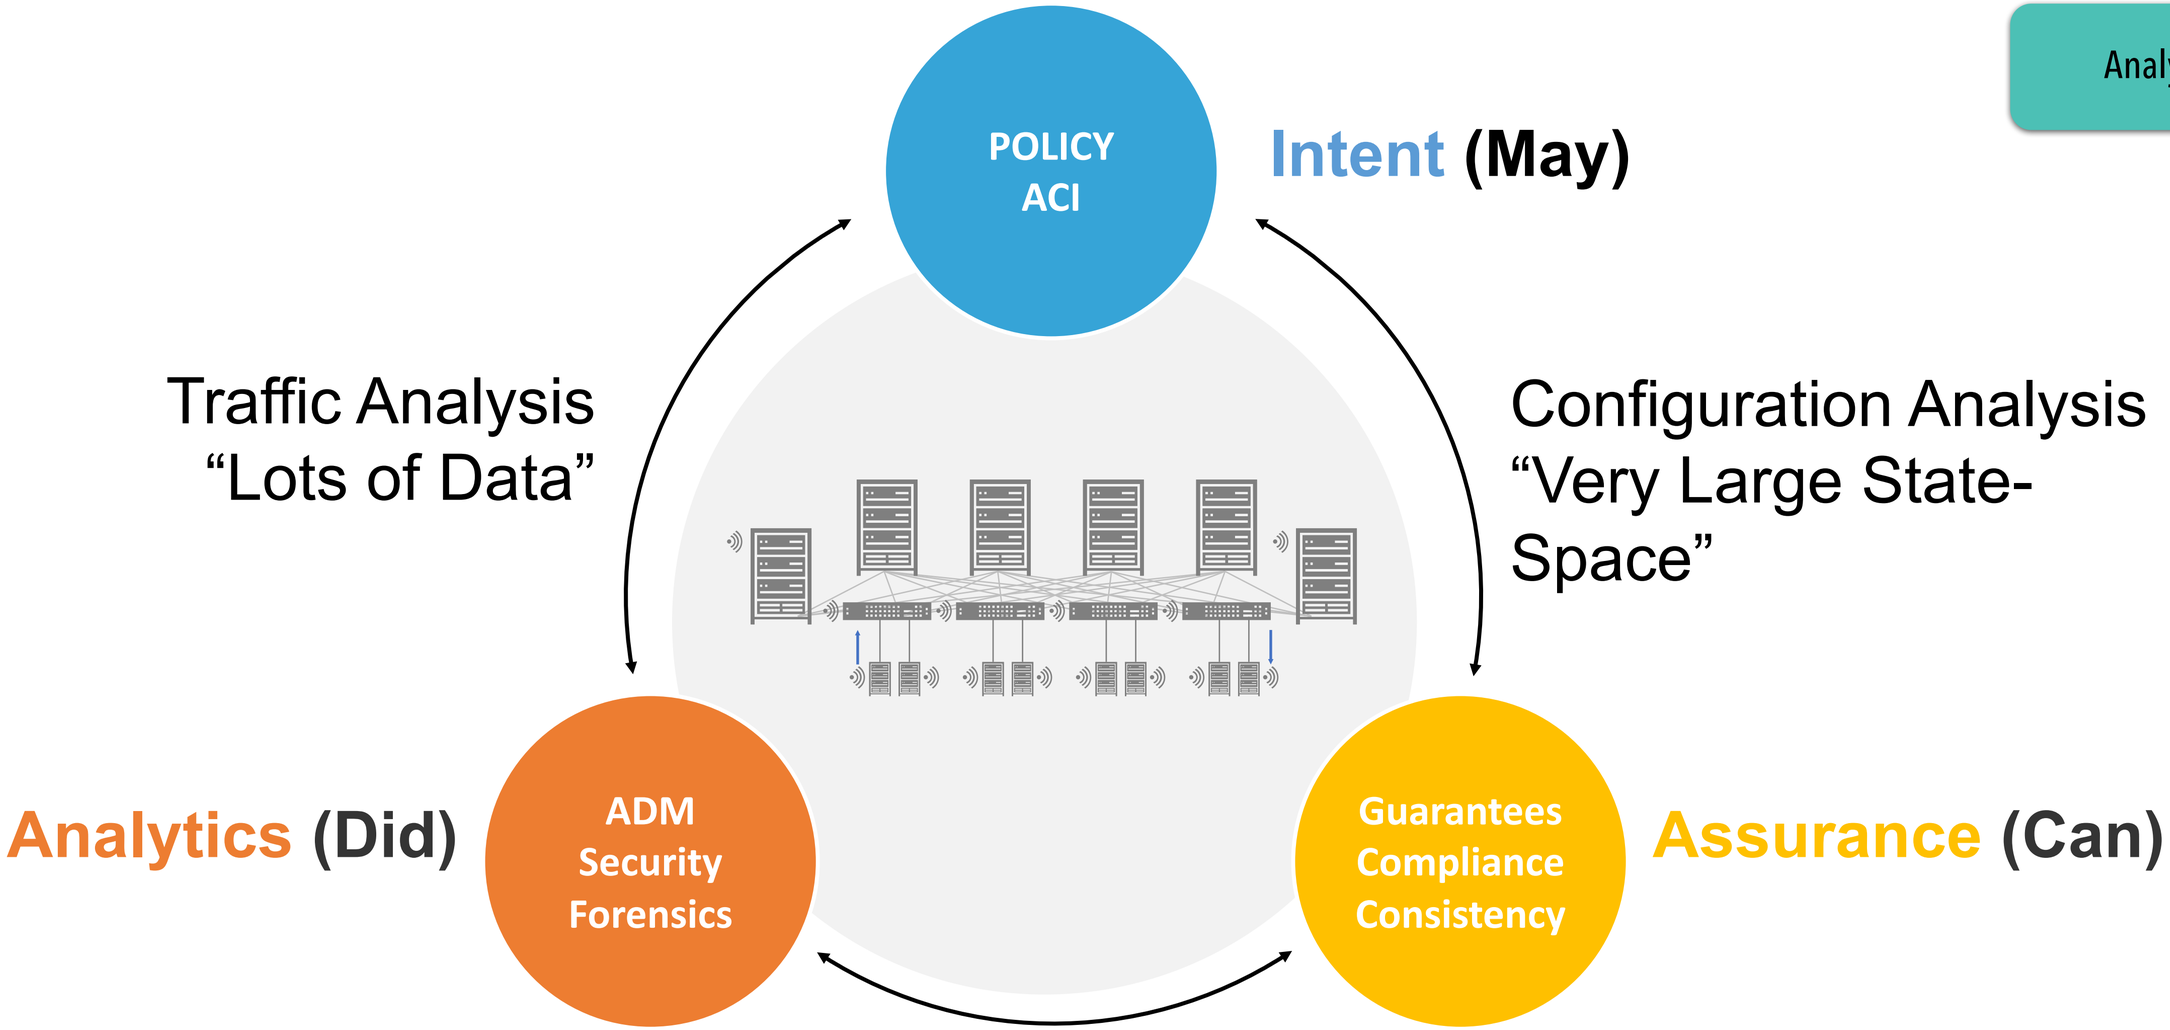

Tetration Analytics Platform

Every Packet, Every Flow, Every Speed

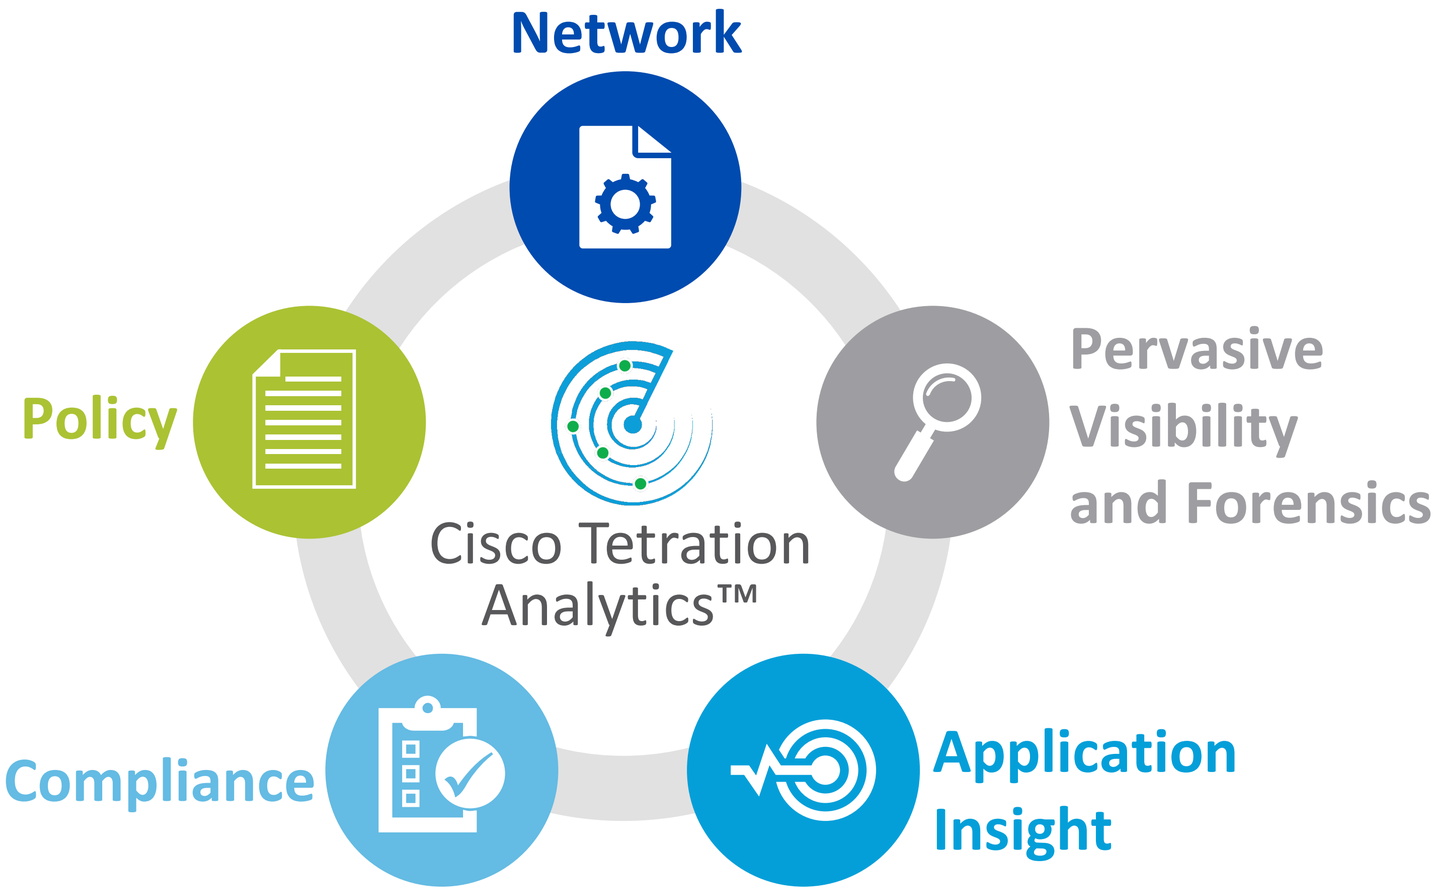

Cisco Connect

Cisco Tetration Analytics

Application Insights

Policy Simulation and Impact Assessment

Automated Whitelist Policy Generation

Forensics: Every Packet, Every Flow, Every Speed

Policy Compliance and Auditability

Cisco Tetration Analytics

Pervasive Sensor Framework

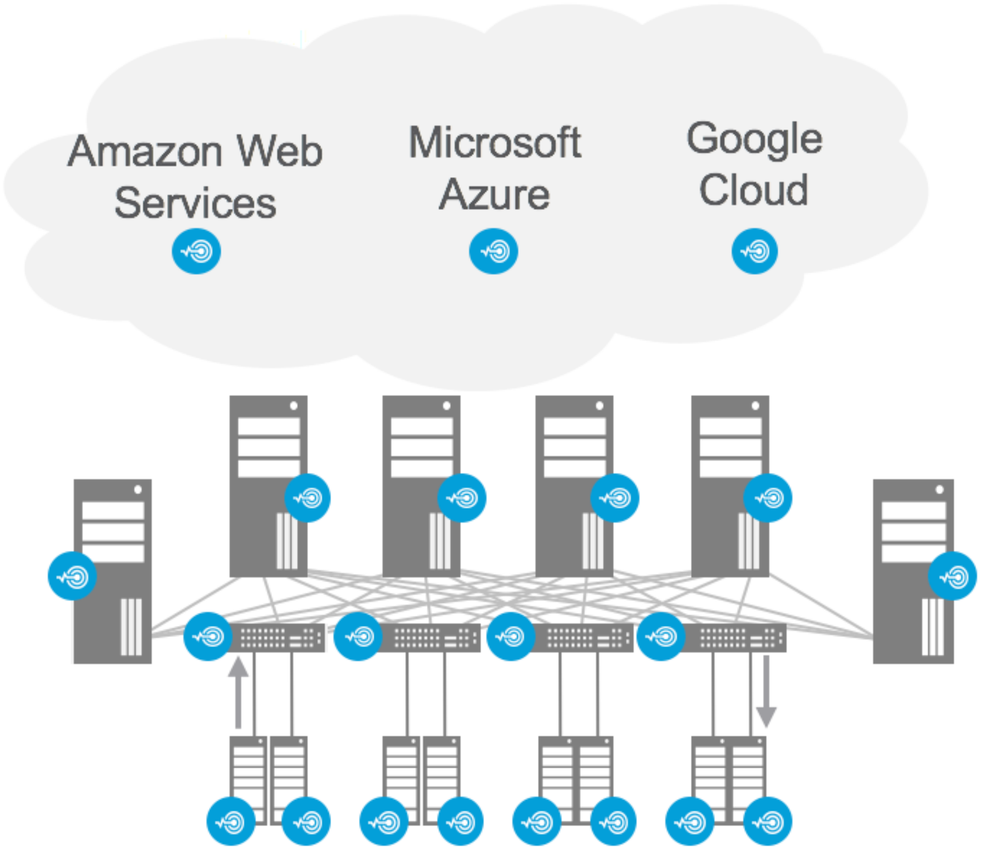

- Provides correlation of data sources across entire application infrastructure
- Enables identification of point events and provides insight into overall systems behavior
- Monitors end-to-end lifecycle of application connectivity

Application Discovery and Endpoint Grouping

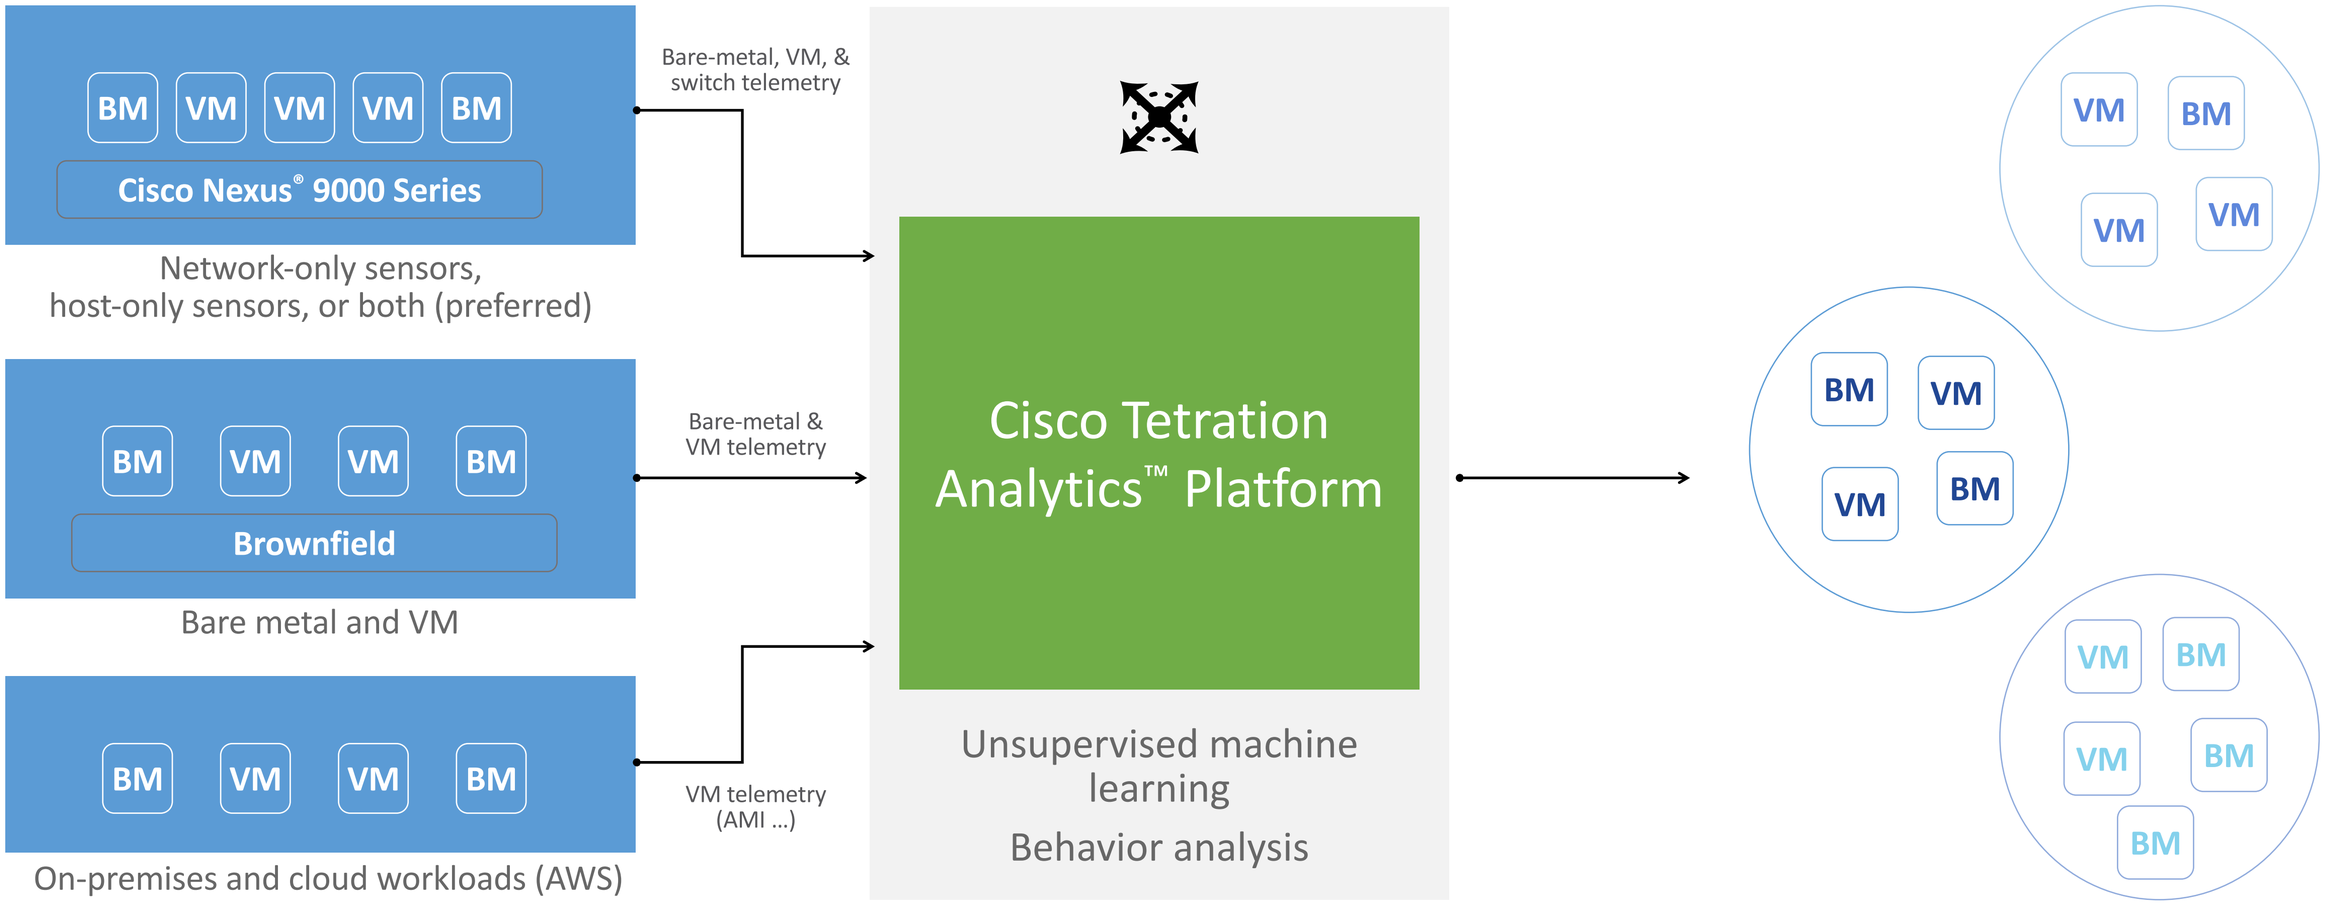

Whitelist Policy Recommendation

Application Discovery

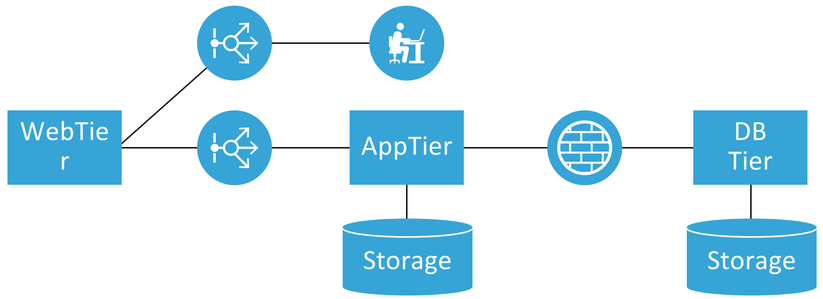

Whitelist Policy Recommendation
(Available in JSON, XML, and YAML)

Policy Enforcement
(Future Roadmap)

Real-Time and Historical Policy Simulation

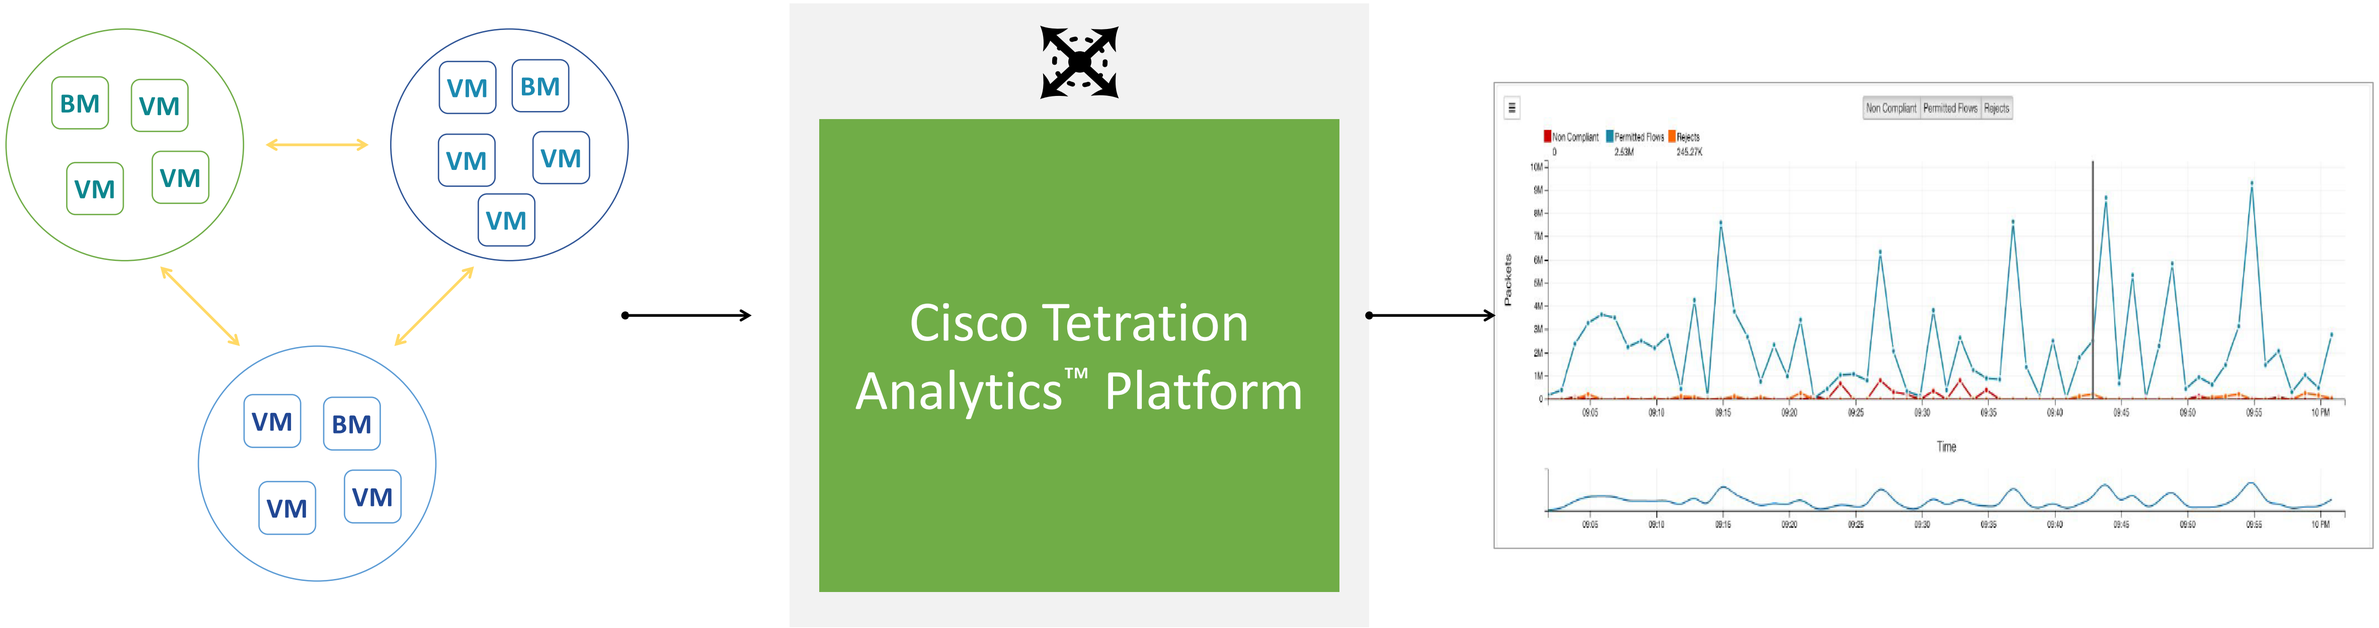

- Validating policy impact assessment in real time
- Simulating policy changes over historic traffic
- View traffic “outliers” for quick intelligence
- Audit becomes a function of continuous machine learning

Policy Compliance

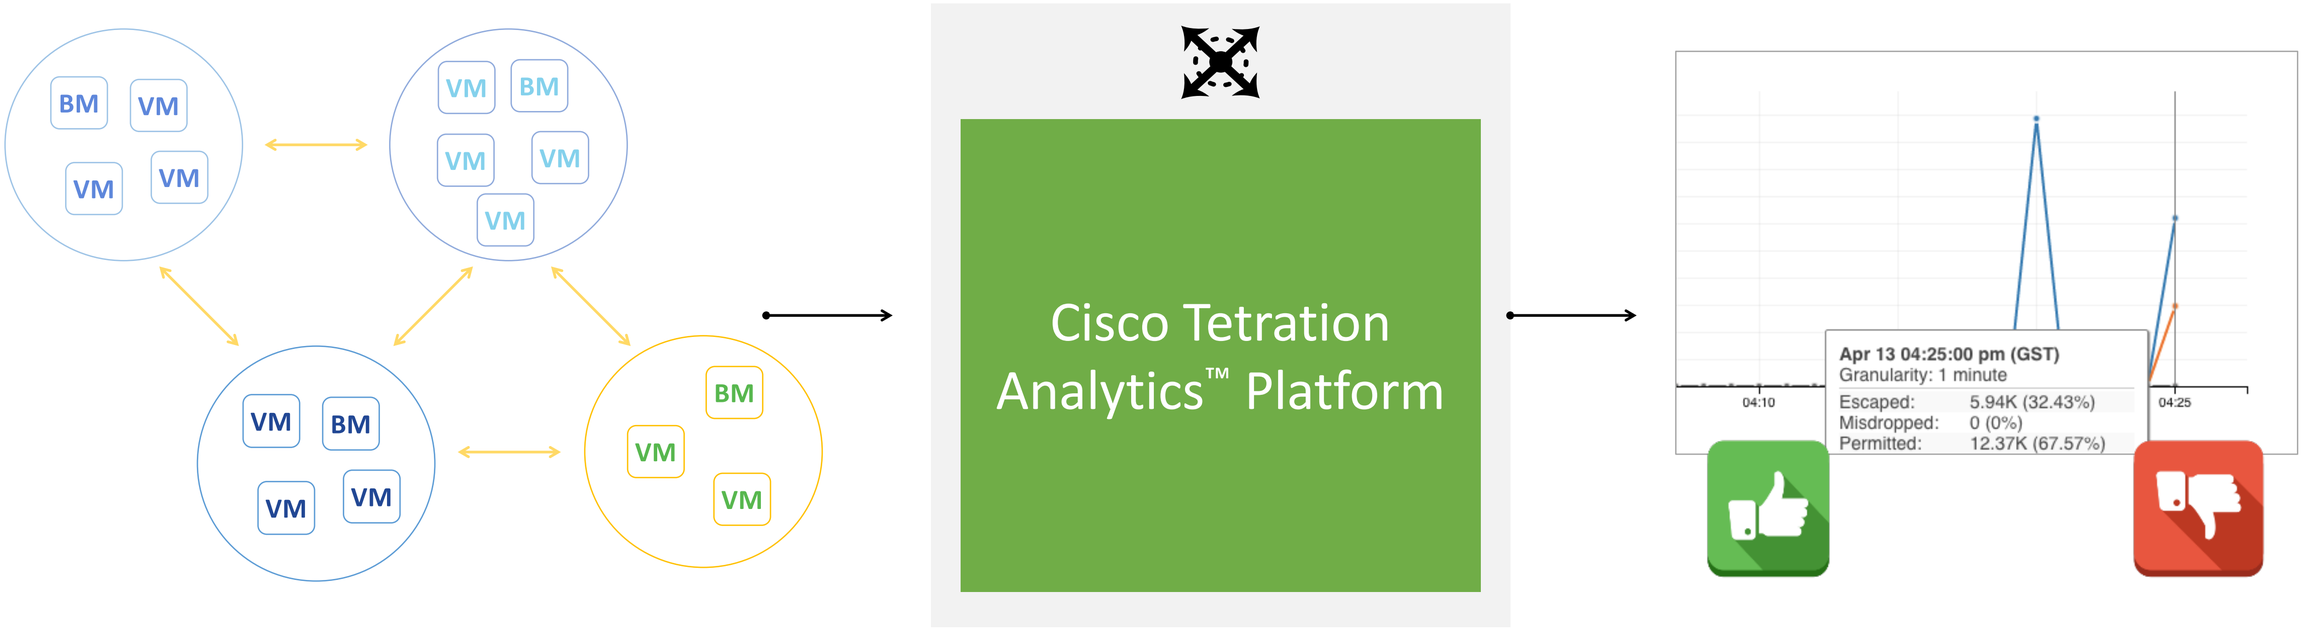

- Identify policy deviations in real-time
- Review and update whitelist policy with one click
- Policy lifecycle management

Tetration Analytics

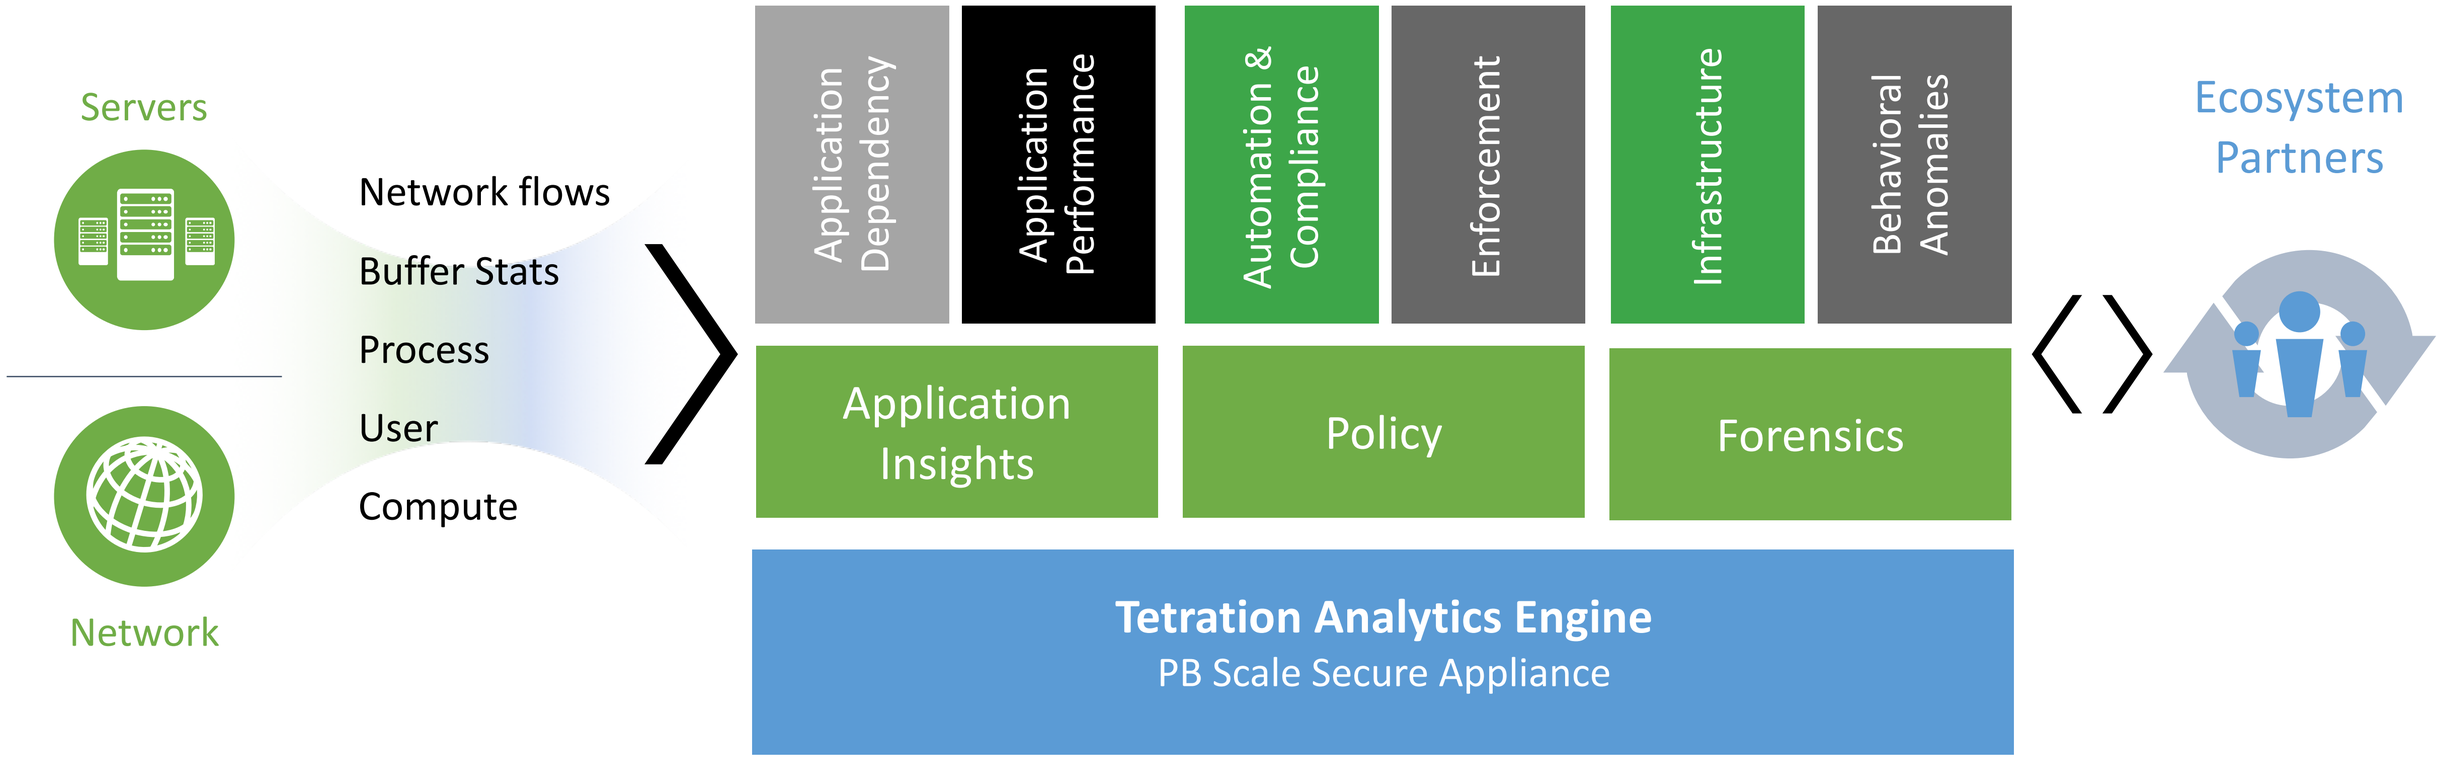

Tetration Analytics Architecture Overview

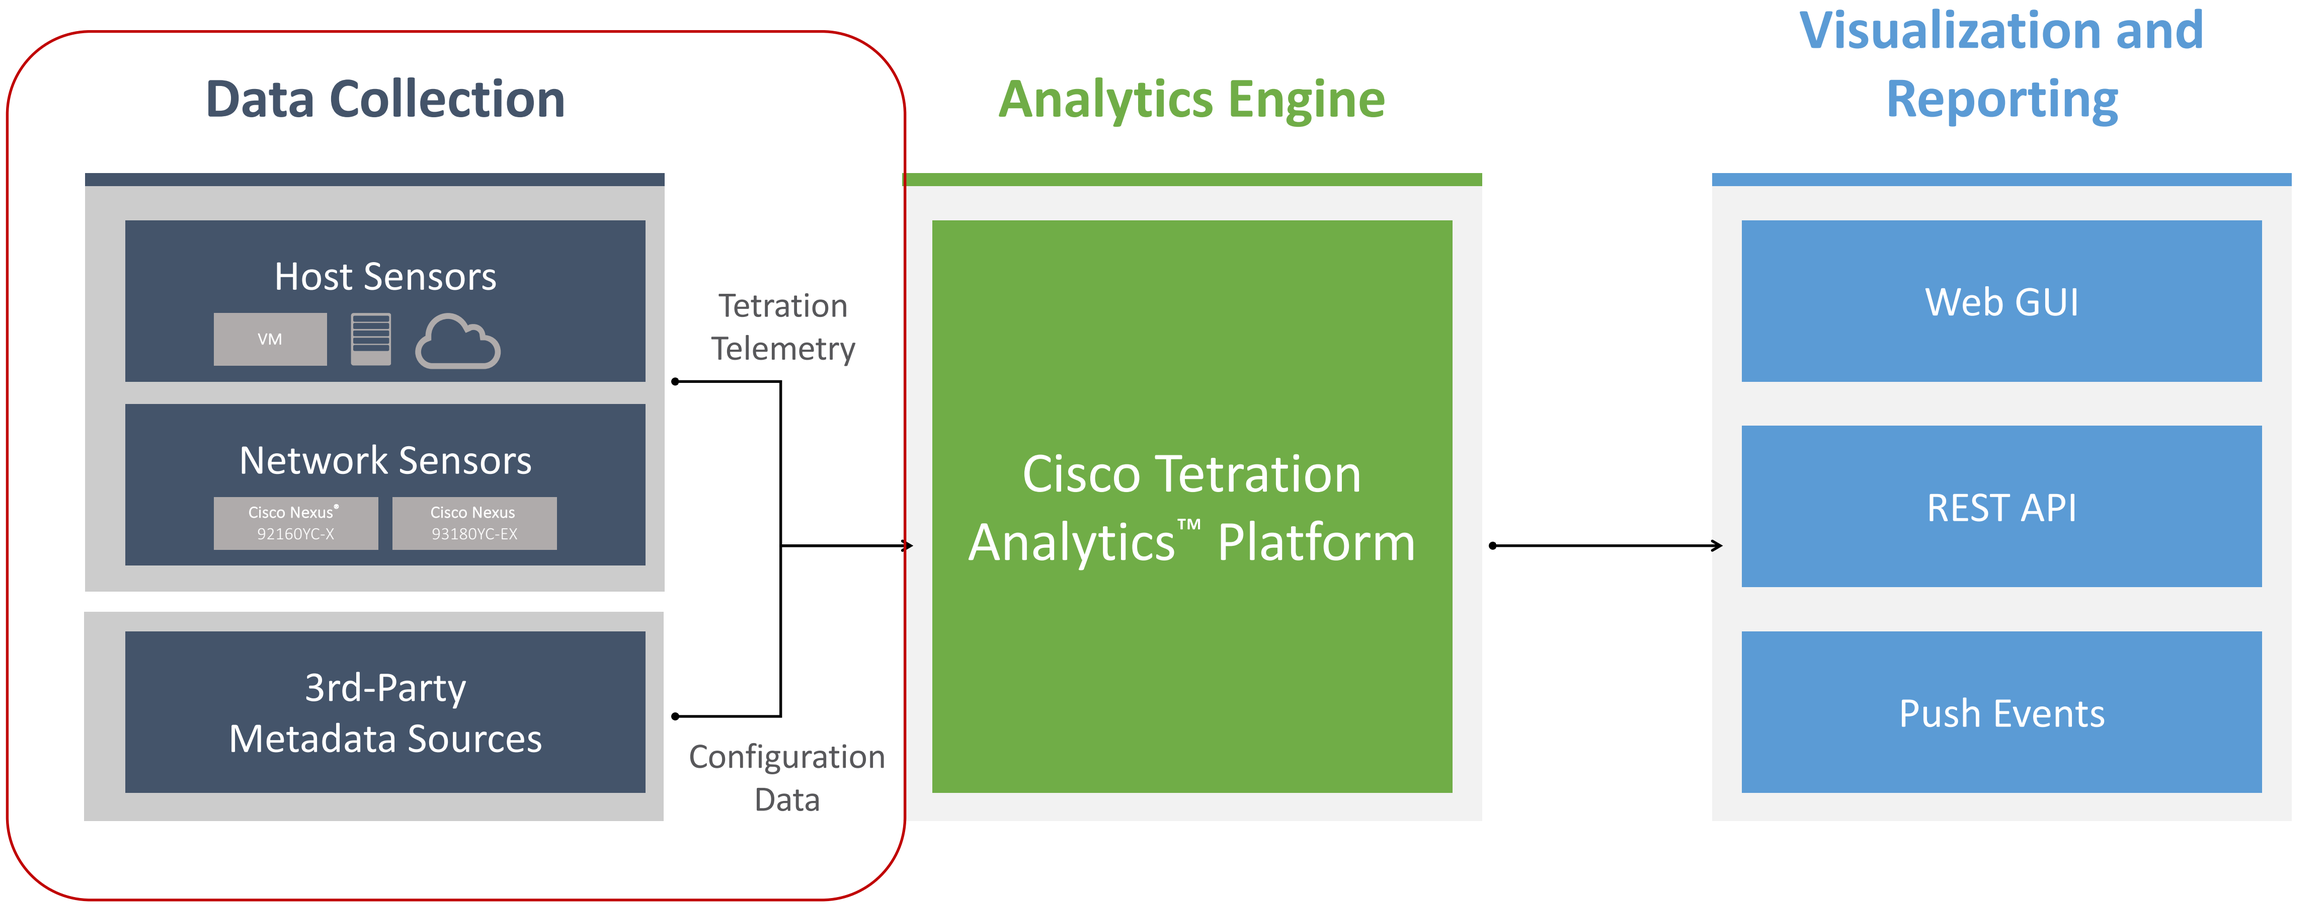

Pervasive Sensors

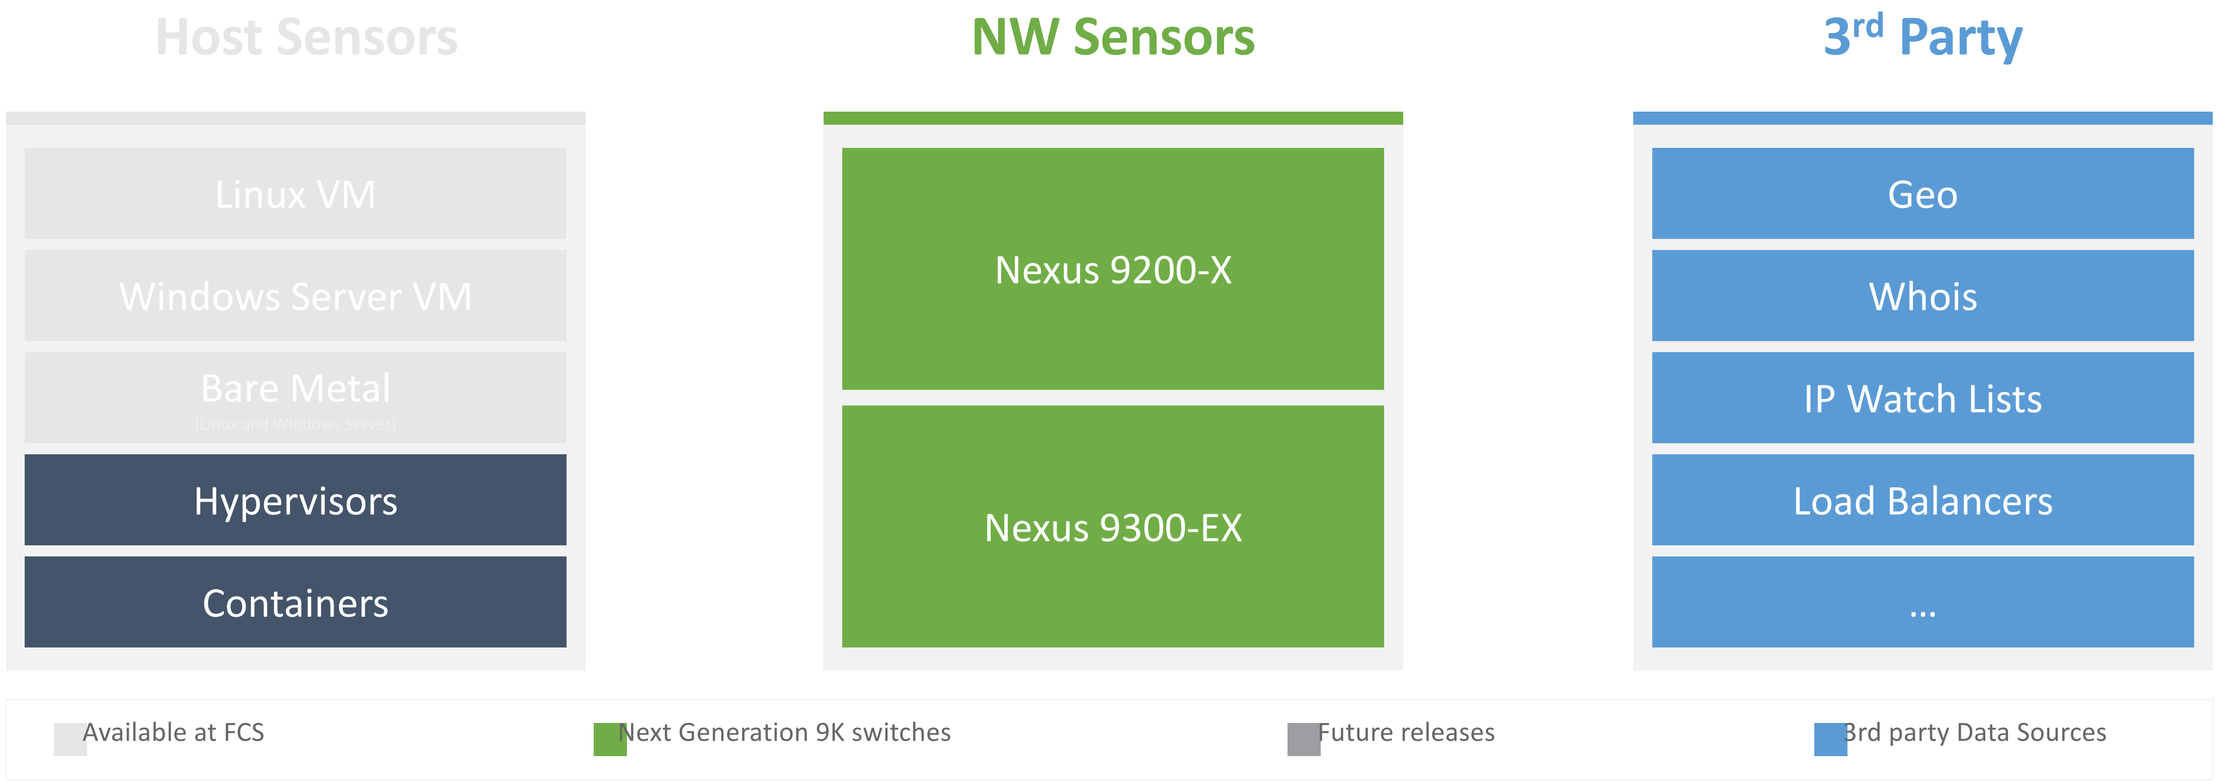

- ✓ Low CPU Overhead (SLA enforced)
- ✓ Low Network Overhead (SLA enforced)
- ✓ Highly Secure (Code Signed, Authenticated)
- ✓ Every flow (No sampling), NO PAYLOAD

Traditional Monitoring Is Showing Its Age

Not suited for Modern Network and Security Operations

Where Data Is Created

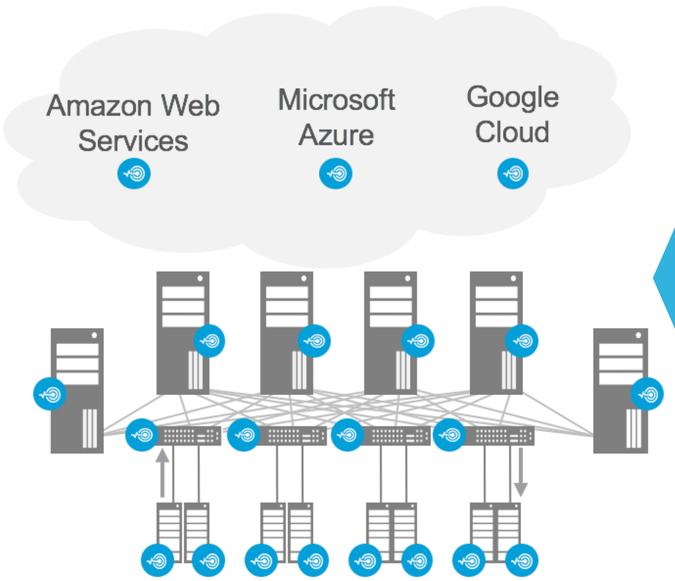

Non Real time

Incomplete

Scale Issues

SNMP

Syslog

CLI

Unstructured

Subject to Change

SNMP Server

Syslog Collector

Scripts

Where Data Is Useful

Storage & Analysis

Strong burden on back-end

Normalize different encodings, transports, data models, timestamps

Streaming Telemetry is a game changer

Monitoring becomes a big data problem

Where Data Is Created

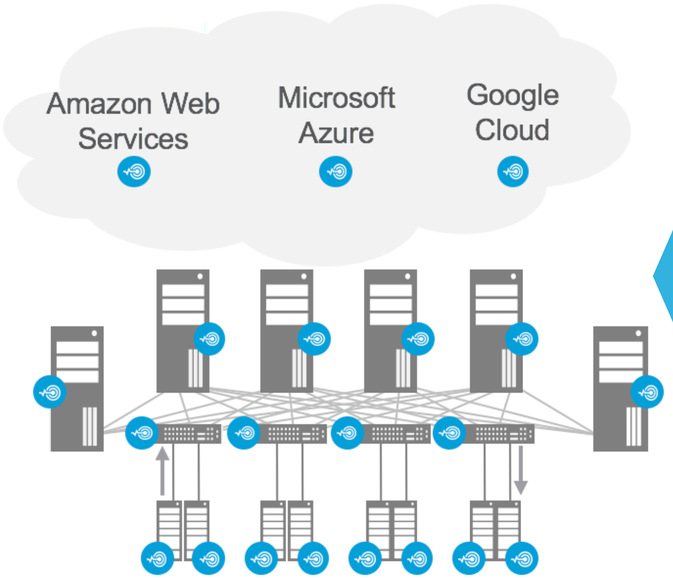

Real time

Removing limitations and complexity

- **Streaming paradigm**
- **Dense Sensor Framework**
- **Increased Data Granularity**
- **Update on every event**
- **Multiple Data Sources**

Where Data Is Useful

Volume – Scale of Data
Velocity – Analysis of Streaming Data
Variety – Different Forms of Data

Big Data and Machine Learning Problem

Why Multiple Sensors?

Example monitoring temperature in a room

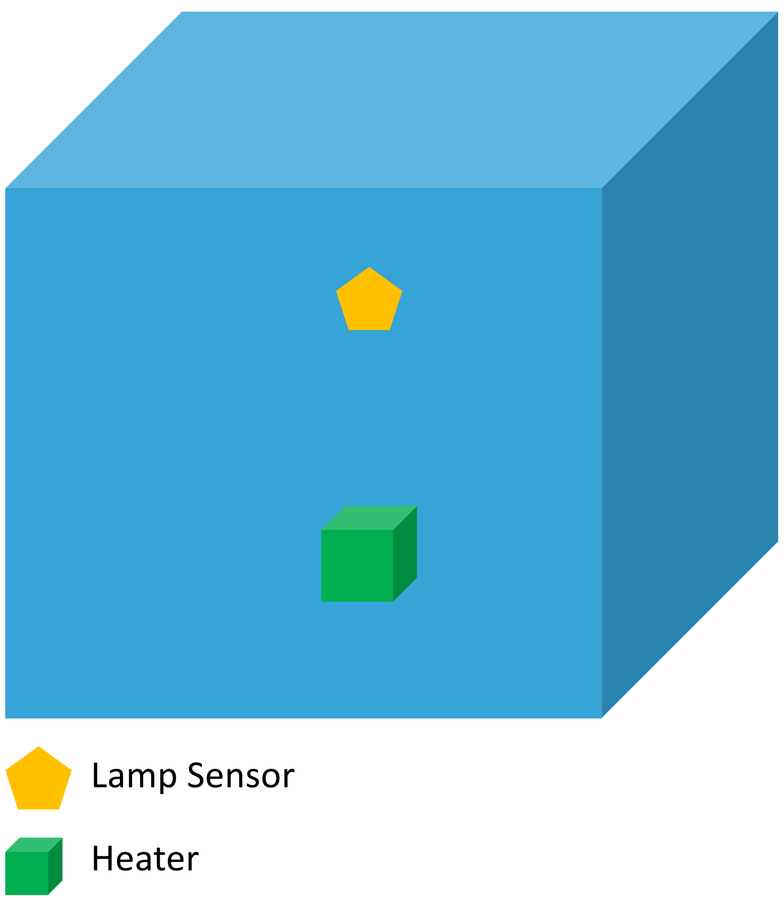

Lamp Sensor

Heater

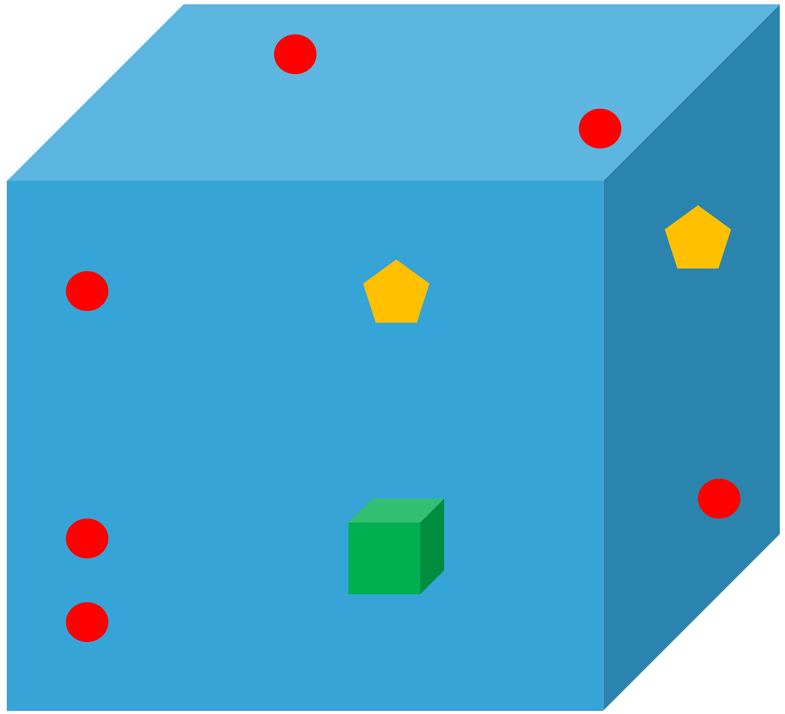

Plug Sensor

Tetration Sensors
Locations

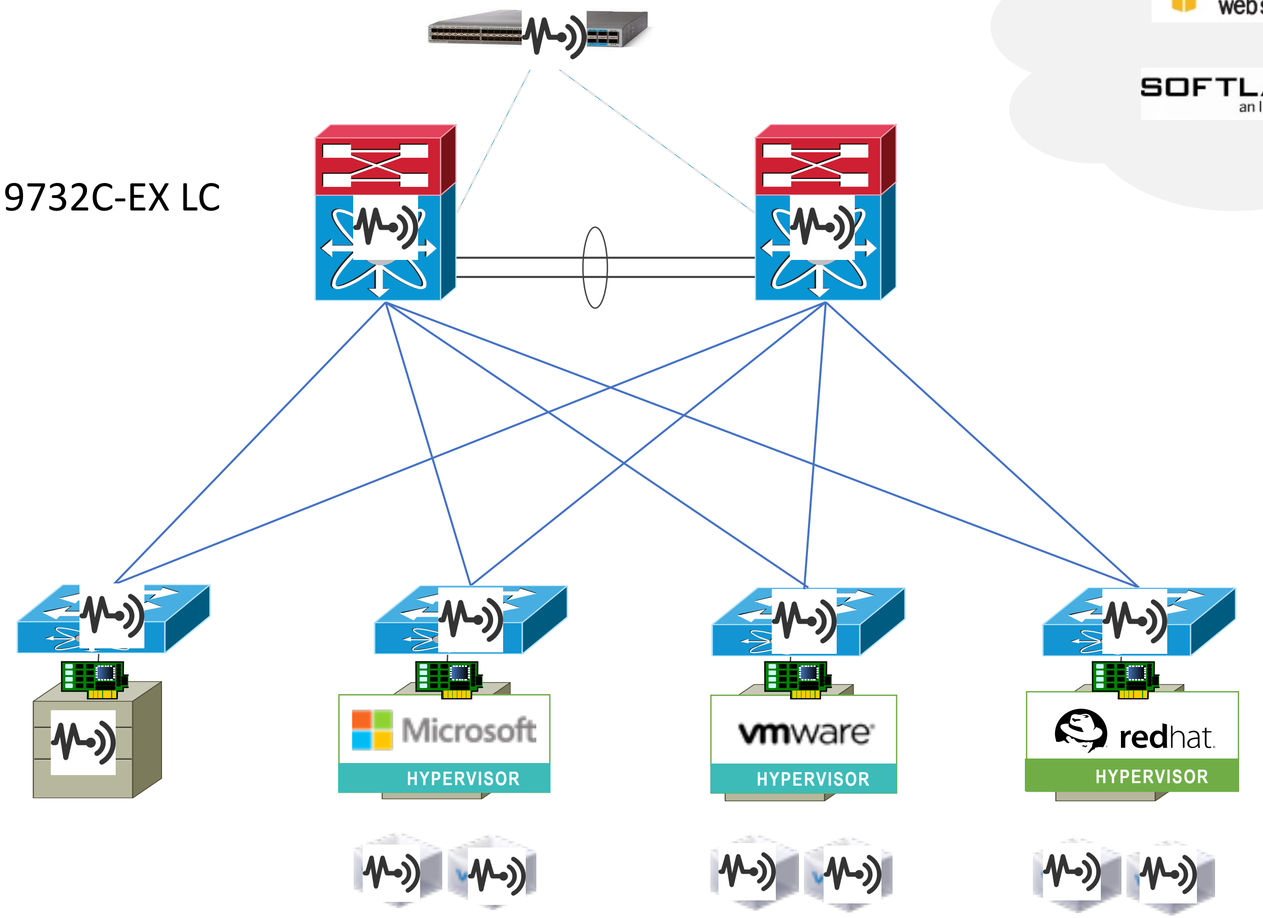

Hardware Sensor
Packet and Flow Events
Buffer and Switch State

Software Sensor
Processes & Socket
Packet and Flow Events

Tetration Cluster

Hardware Sensor

- Embedded Module (Flow Cache)
  - Nexus 92160CY-X
  - Nexus 93180Y-EX & 9732C-EX Line Cards
- Extracts Meta-Data from the forwarding pipeline
  - No latency impact, no performance impact

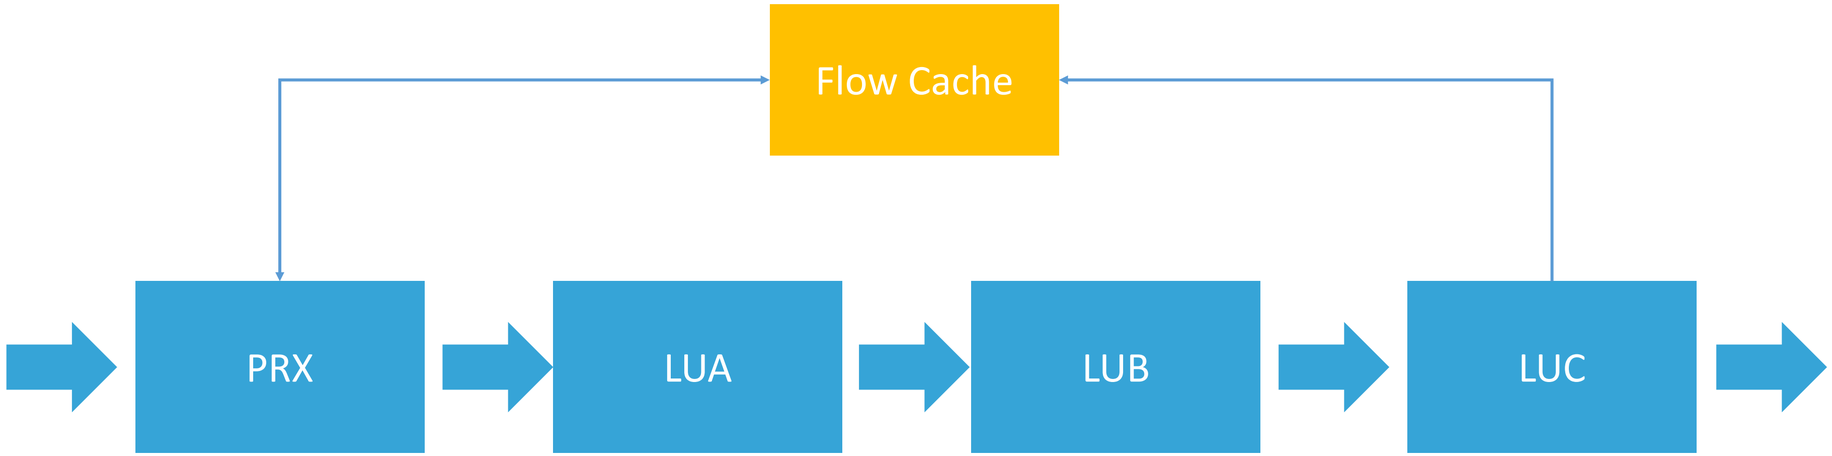

Software Sensor

- Not in the data path
  - Sits in User Space
  - Designed by Kernel Developers
- Secure
  - Code Signed
- SLA Enforcement
  - CPU and BW throttling
- FCS availability
  - Windows
    - 2008 / 2008 R2 / 2012 / 2012 R2
  - Linux
    - RedHat (5.3+, 6.x)
    - CentOS (5.11+, 6.x)
    - Ubuntu (12.04, 14.04, 14.10)

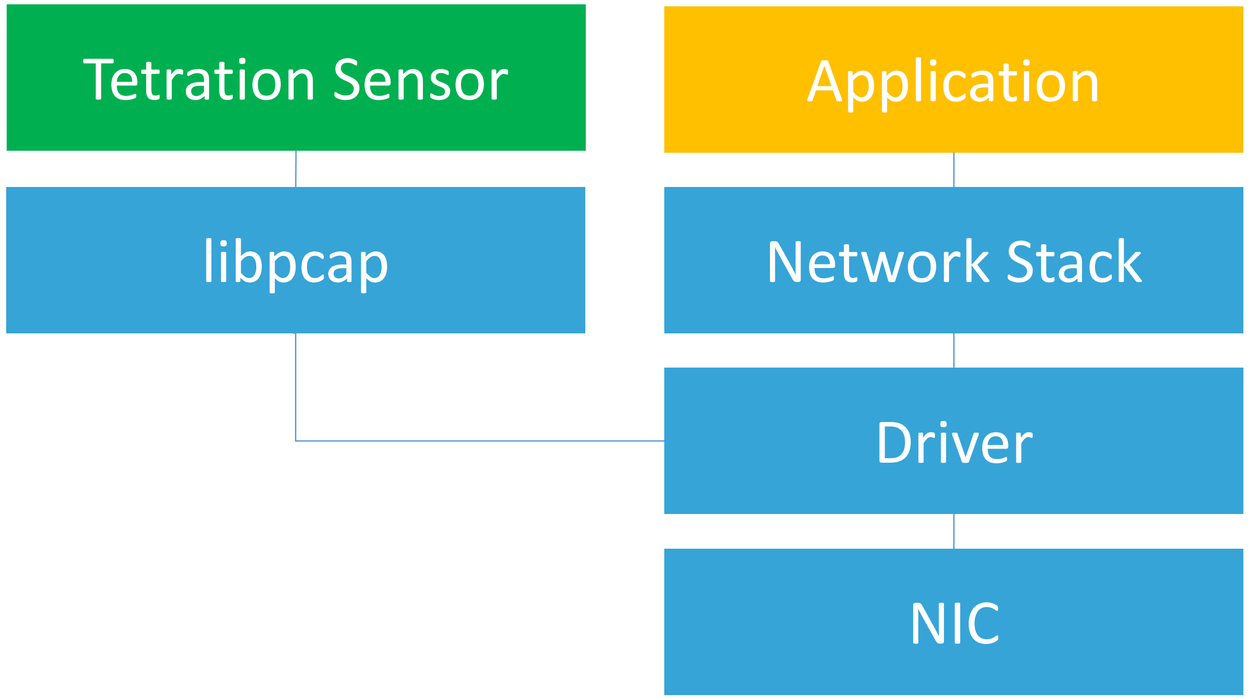

Methods to deploy the sensor

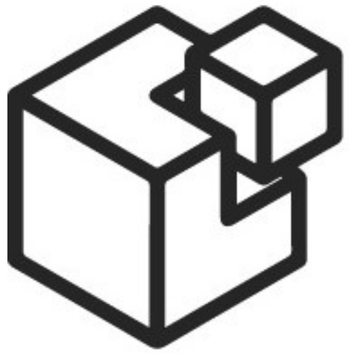

SALTSTACK

Coming soon to a GitHub near you

github.com/datacenter

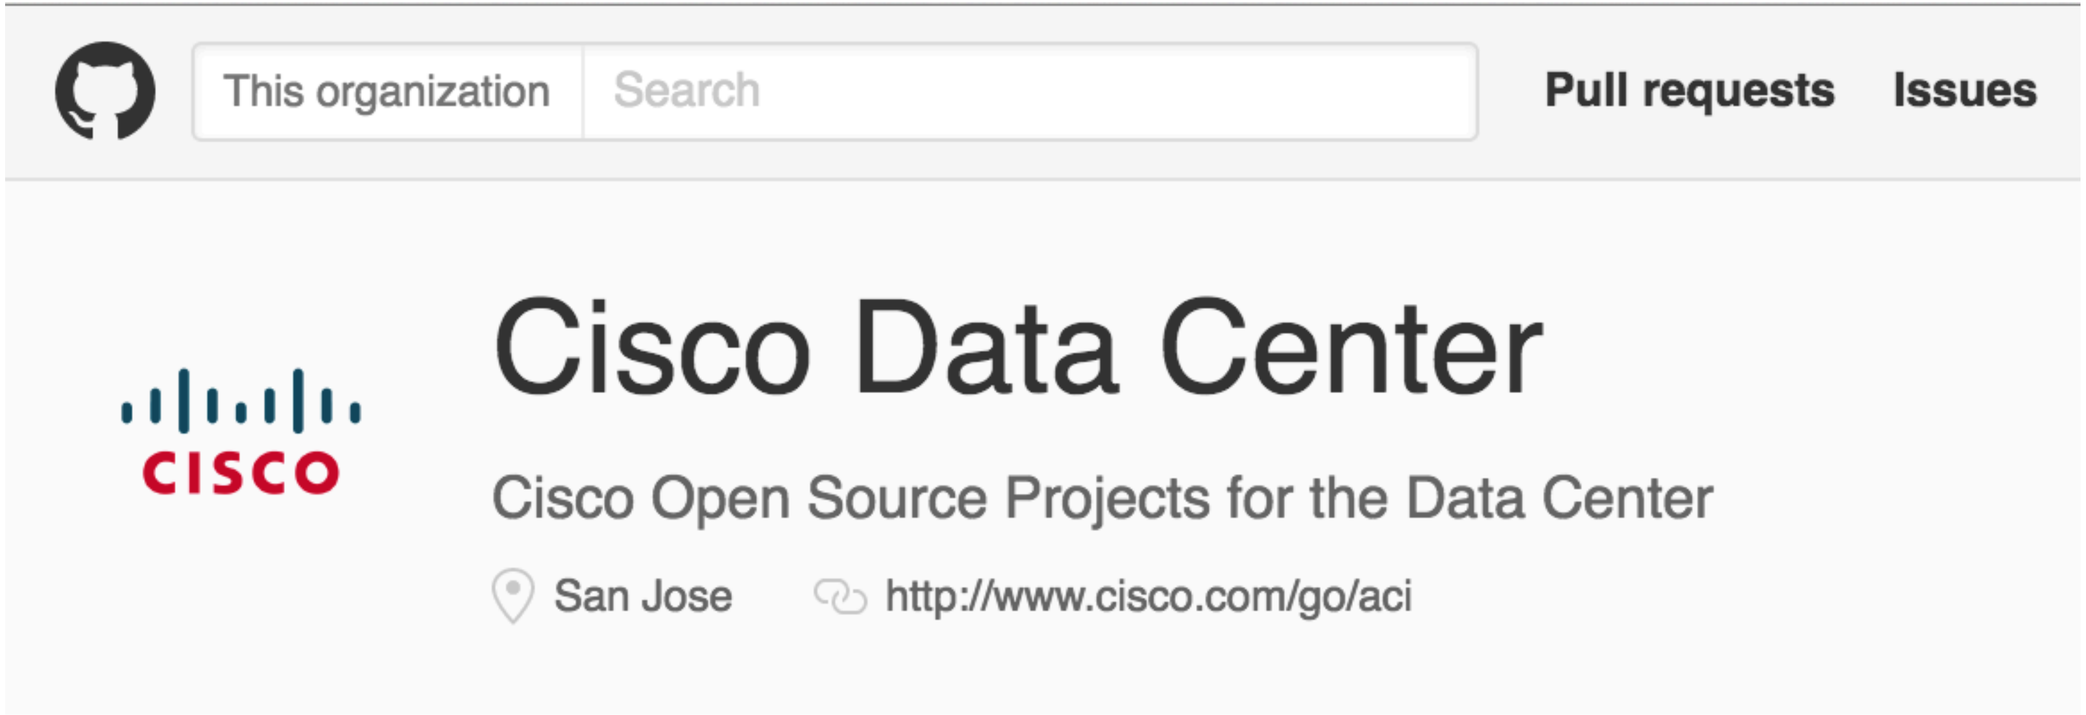

Tetration Analytics Architecture Overview

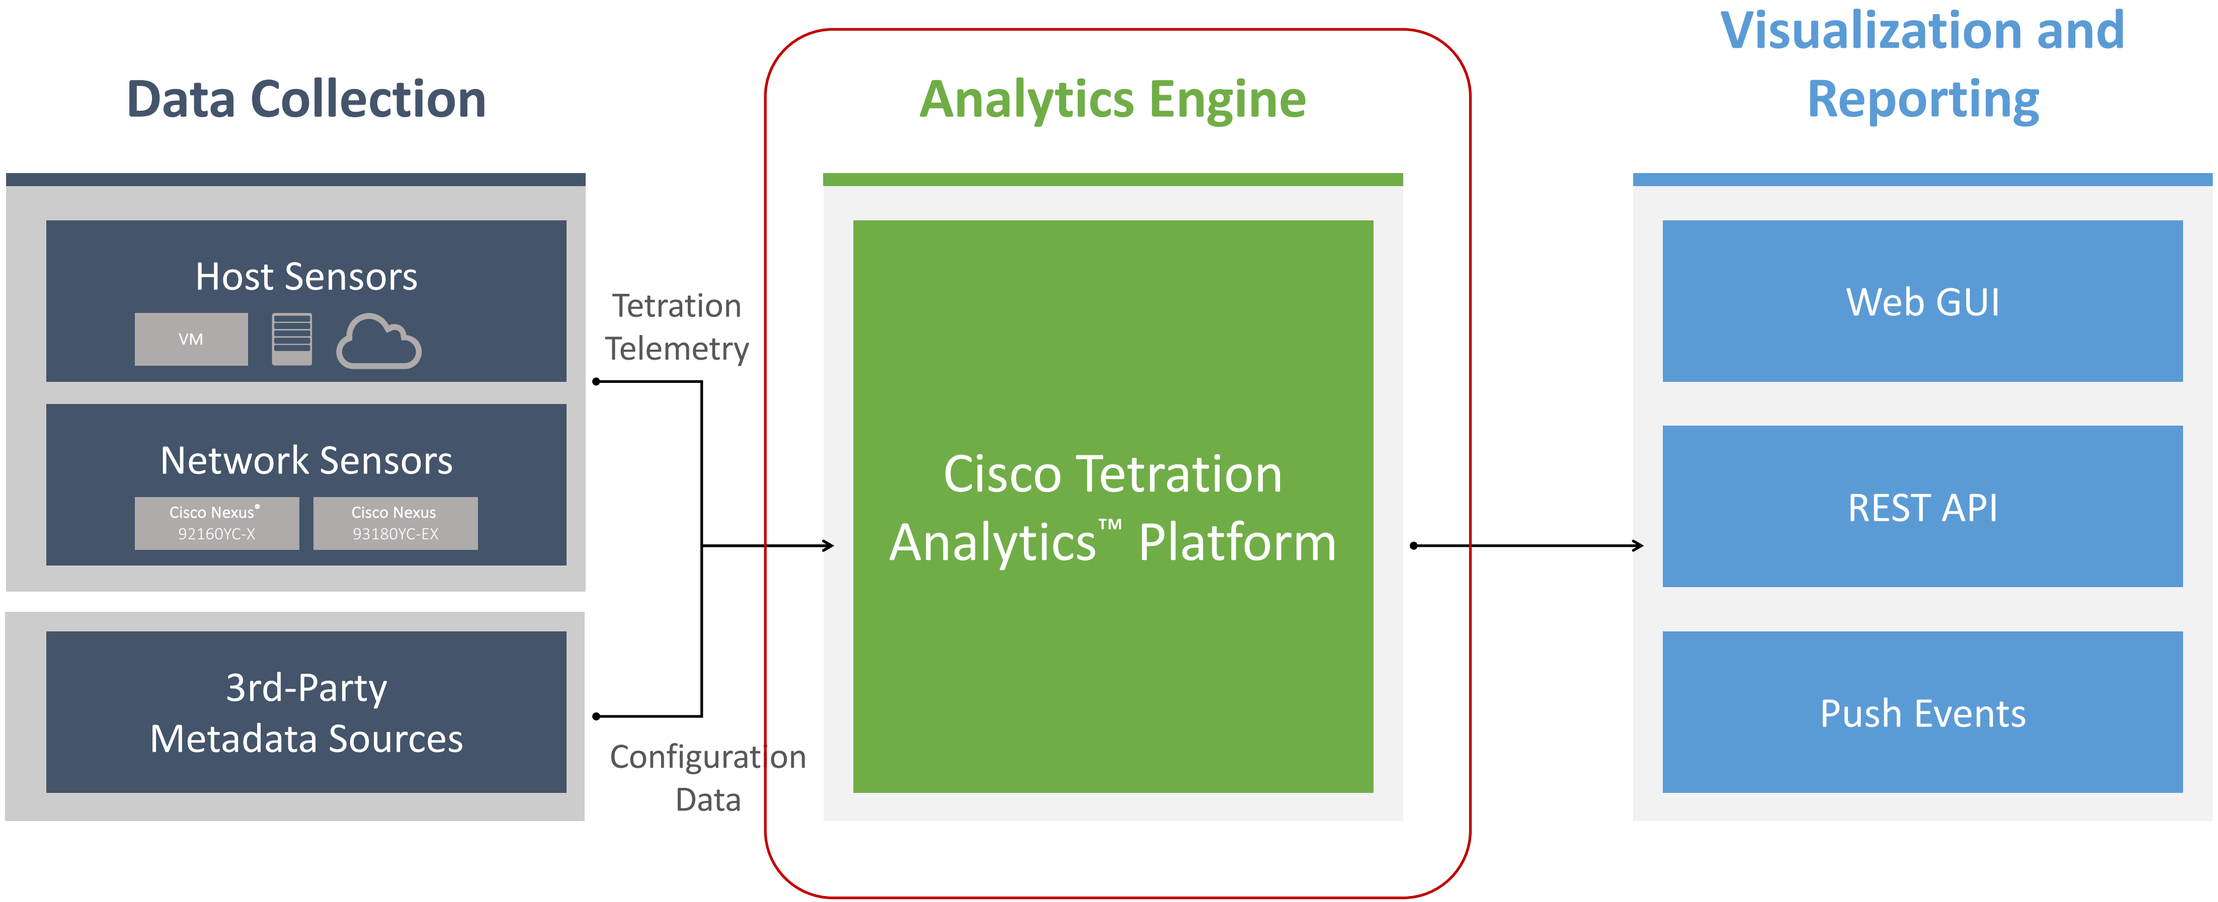

The Analytics Cluster
Components

- Hadoop Based Platform
  - Self managed
  - One touch deployment
- Tiered System
  - Heavy Compute for Machine Learning
  - Caching for light speed queries
- Extensibility (future)
  - Messaging Bus
  - API Access

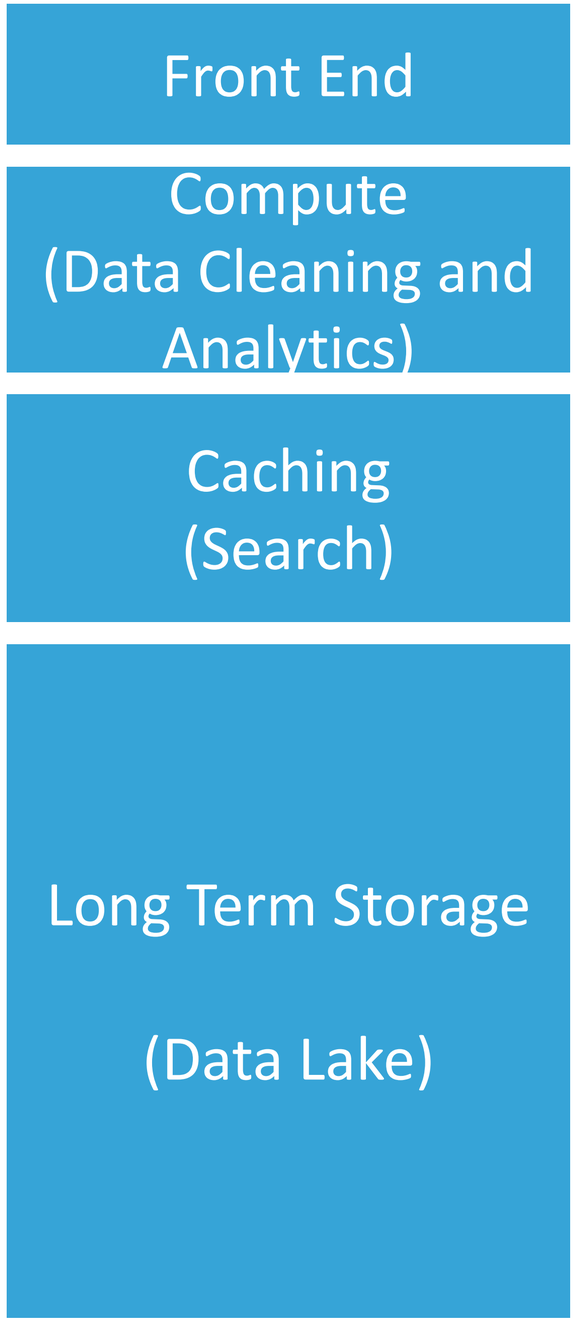

The Analytics Cluster Appliance

- The Analytics Cluster operates as an appliance
  - Avoids the need for in house Big Data, Analytics expertise
  - Supported by Cisco TAC
- Self Monitoring
  - The cluster leverages a sensor architecture to track it’s state and provides event based notifications for
- Software upgrades and full install are all automated

FCS Analytics Cluster Configurations

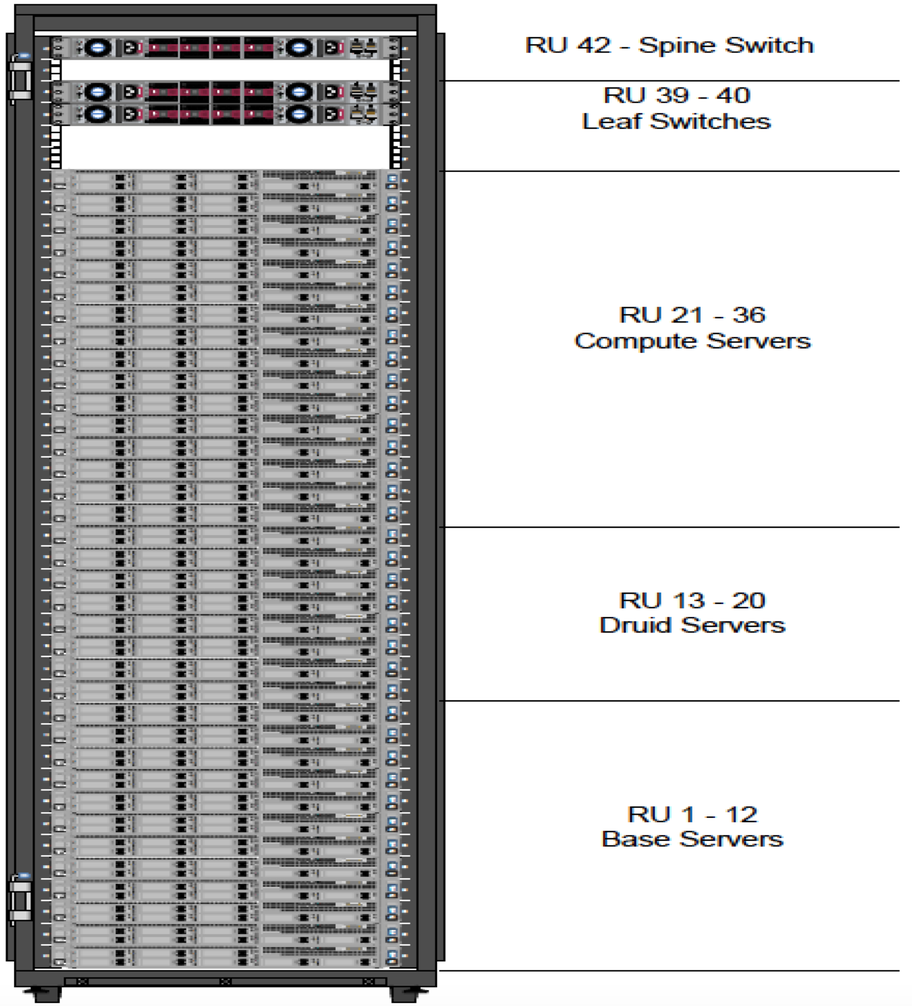

4 x 3-Phase PDU
22.5 KW Peak Power

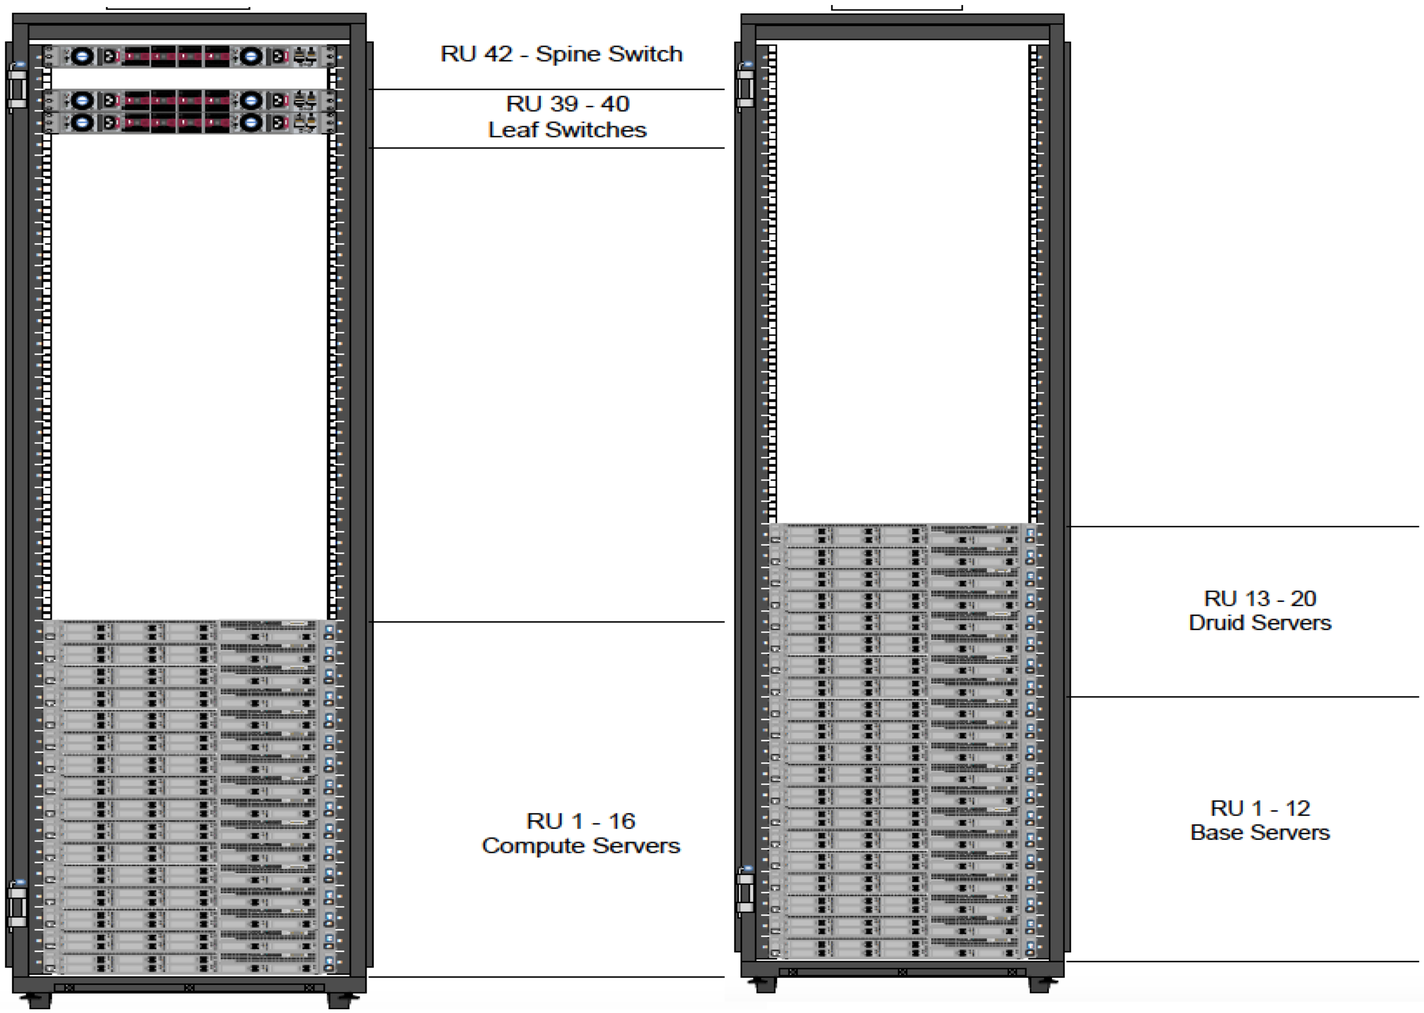

4 x 1-Phase PDU
11.5 KW Peak Power

Options for Future Cluster Models

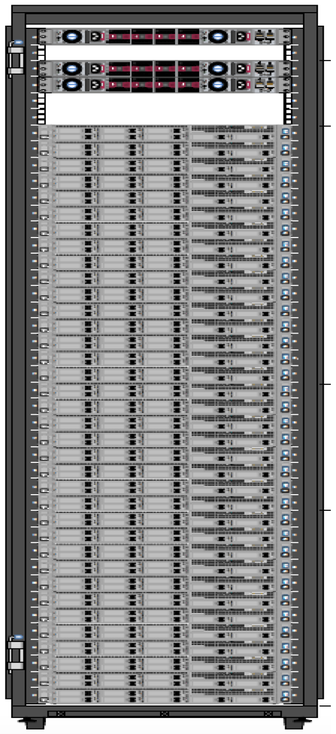

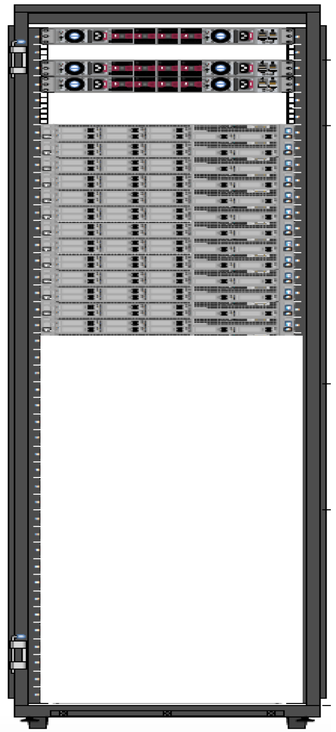

Analytics Engine
The Platform

- Hadoop Based Platform
  - Self managed
  - One touch deployment
- Tiered System
  - Heavy Compute for Machine Learning
  - Caching for light speed queries
- Extensibility (future)
  - Messaging Bus
  - API Access

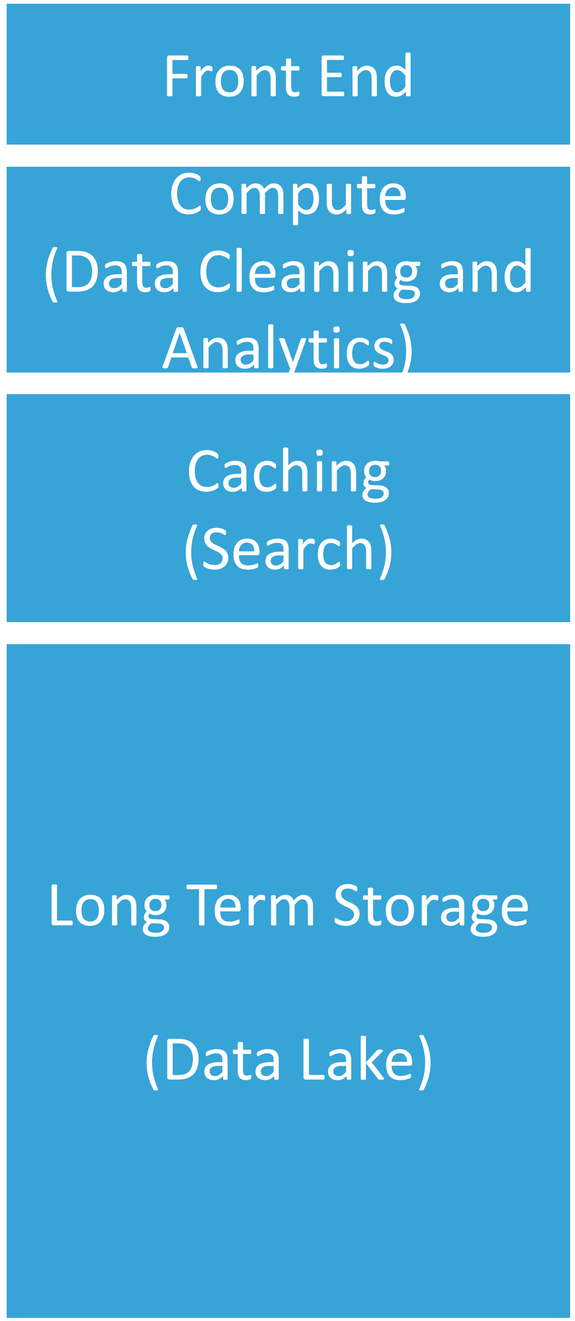

Front End

GUI, RESTful API, Messaging BUS

- Servers hosting front end processes
- GUI and Operational Interfaces
- RESTful API (post FCS)
- Messaging BUS (post FCS)

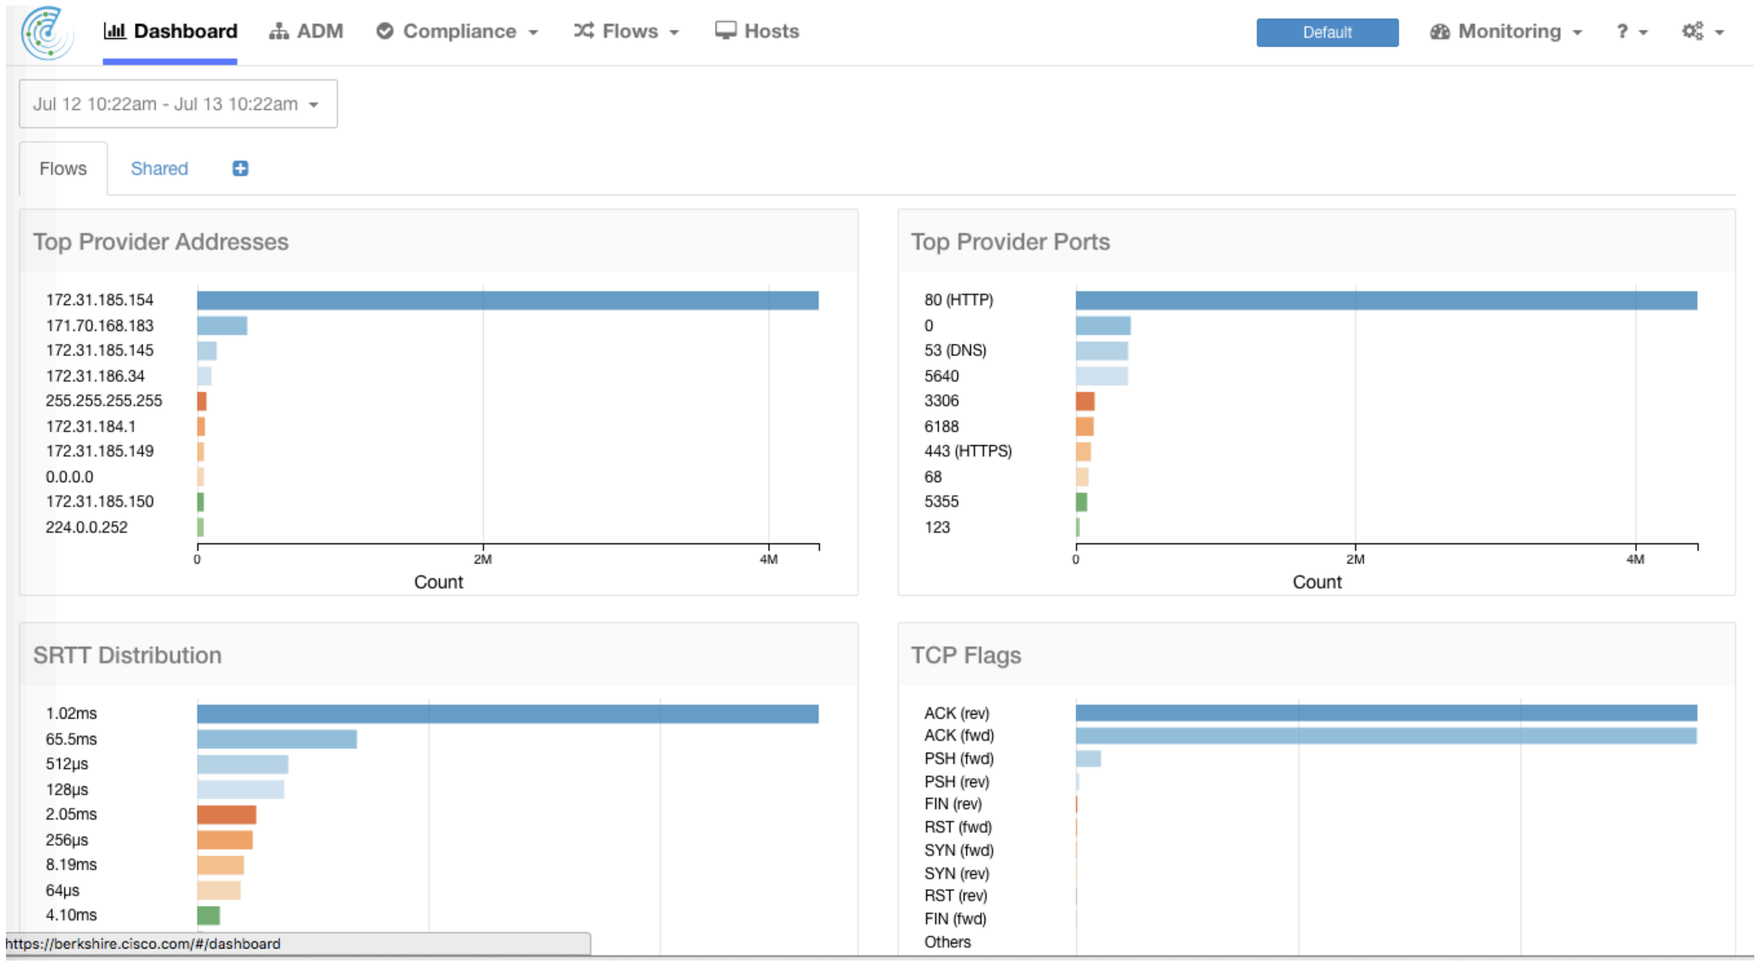

ACI Architecture

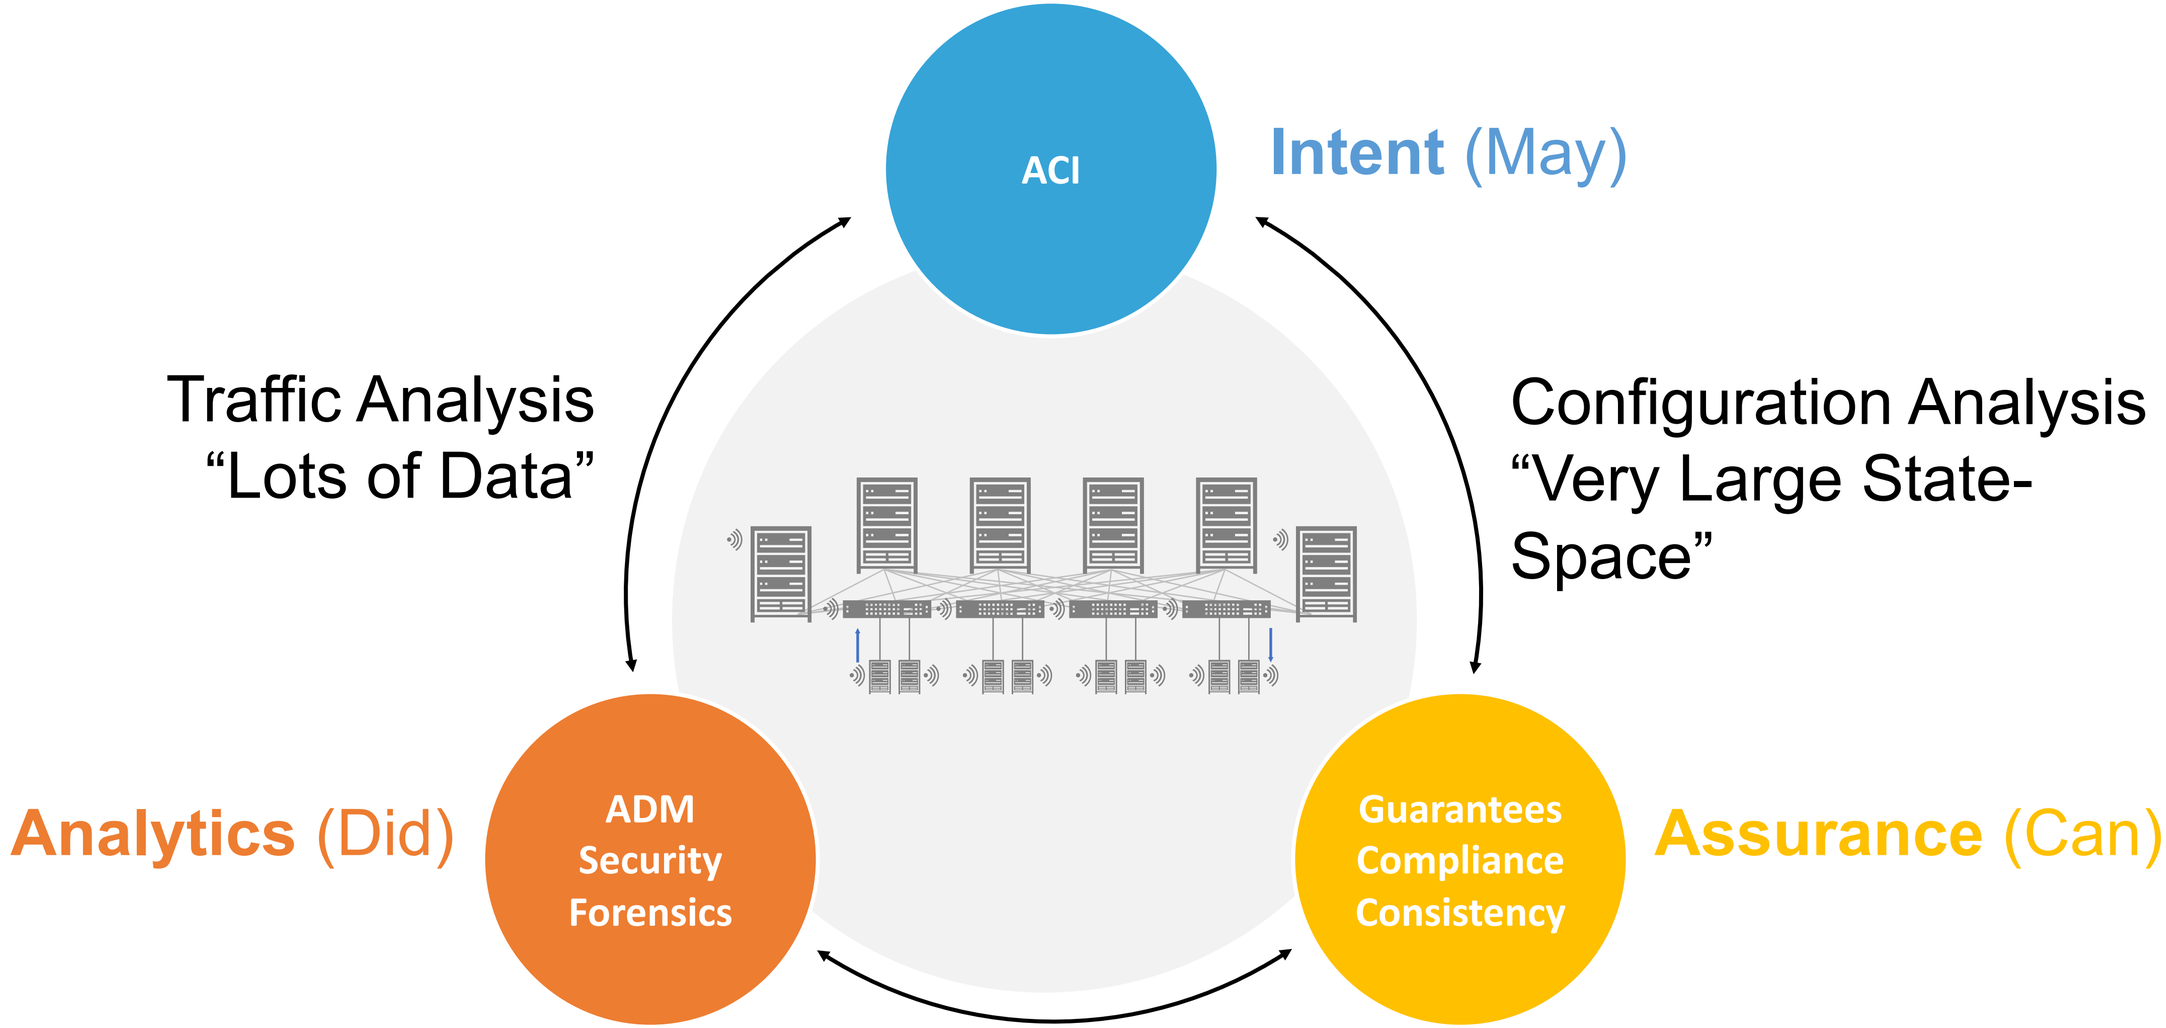

Summary

Pervasive flow telemetry that supports infrastructure for multiple data centers at scale

Ready-to-use solution to address critical data center operational use cases

Self-monitoring and eliminate the need for in-house big data expertise

Open platform and northbound APIs enable transparent integration

Accelerated adoption and comprehensive Solution support with Services


CISCO