



PLATFORM SERVICE ASSURANCE FOR NFV

Maryam Tahhan, Eoin Walsh, Tim Verrall, Damien Power, Jim Greene
and John Browne

Agenda

Why do we need SA?

Service Assurance and how it's accomplished today

Platform Service Assurance

Lifetime of an Event

Moving through the phases

Progress to date

Summary

Google spotlights data center inner workings

*“In each cluster's (of 1,800 servers) first year, it's typical that **1,000 individual machine failures** will occur; **thousands of hard drive failures** will occur; **one power distribution unit will fail**, bringing down 500 to 1,000 machines for about 6 hours; **20 racks will fail**, each time causing 40 to 80 machines to vanish from the network; **5 racks will "go wonky," with half their network packets missing in action**; and there's about a **50 percent chance that the cluster will overheat, taking down most of the servers in less than 5 minutes and taking 1 to 2 days to recover.** “ – Jeff Dean 2008*

<https://www.cnet.com/news/google-spotlights-data-center-inner-workings/>

Top Internet Outages of 2016



Overloaded Network



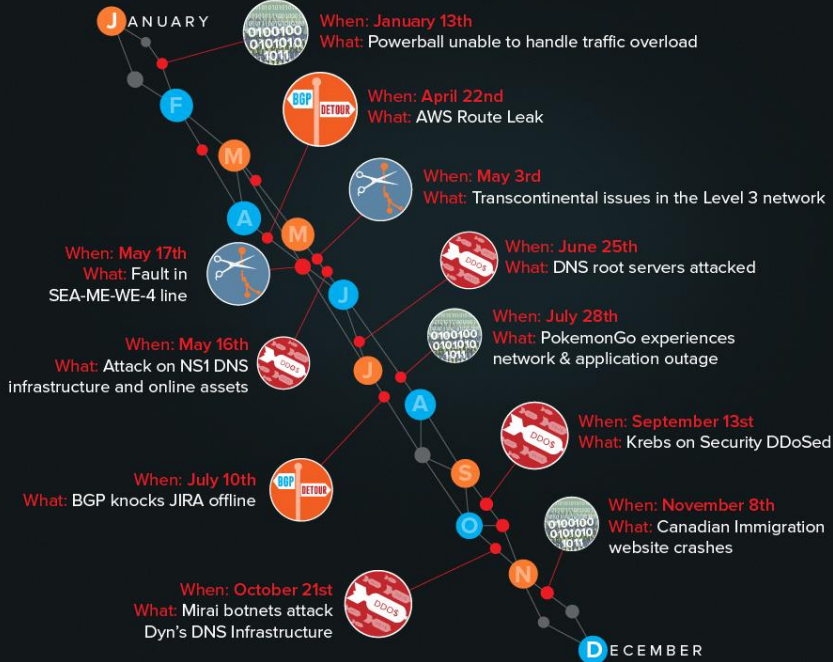
Cable Cut



DDoS Attack



Routing Outages



presented by **ThousandEyes**

“By its very nature, it is expected that networks are bound to have outages and security threats. Smarter networks are not the ones that are built to be fool proof, but the ones that can quickly react to failure and inconsistencies”.

<https://blog.thousandeyes.com/top-internet-outages-2016/>

Why do we need SA?

“Data Centres are powering our everyday lives. Organizations can lose an average of \$400,000++ for one hour of downtime.” [1].

Telco and Enterprise alike are asking how they get and provide Service Assurance, QoS and provide SLA's on the platform and services when deploying NFV.

It is vital to monitor systems for malfunctions or misbehaviours that could lead to service disruption and promptly react to these faults/events to minimize service disruption/downtime.

[1] https://planetaklimata.com.ua/instr/Liebert_Hiross/Cost_of_Data_Center_Outages_2016_Eng.pdf

Service Assurance Today: FCAPS

“Service Assurance” = FCAPS - Fault, Configuration, Accounting, Performance, Security

Fault – Identify component or service failures in the network

Configuration – Setting up or managing the many systems that comprise the network

Accounting – Monitoring data flows for billing, chargeback,

Performance – Measuring data flows and usages for trending, management and analysis

Security – Controlling the access and use of network resources

- Scores of tools are used to collect, report FCAPS data, for graphing and trending
- Thresholds, baselines and watermarks are established and used to decide when to actions are required
- Entire system is used for tracking end-end network service levels and conformance to subscribed service levels

FCAPS is the Operational Scope of Service Assurance, Foundation for SLAs

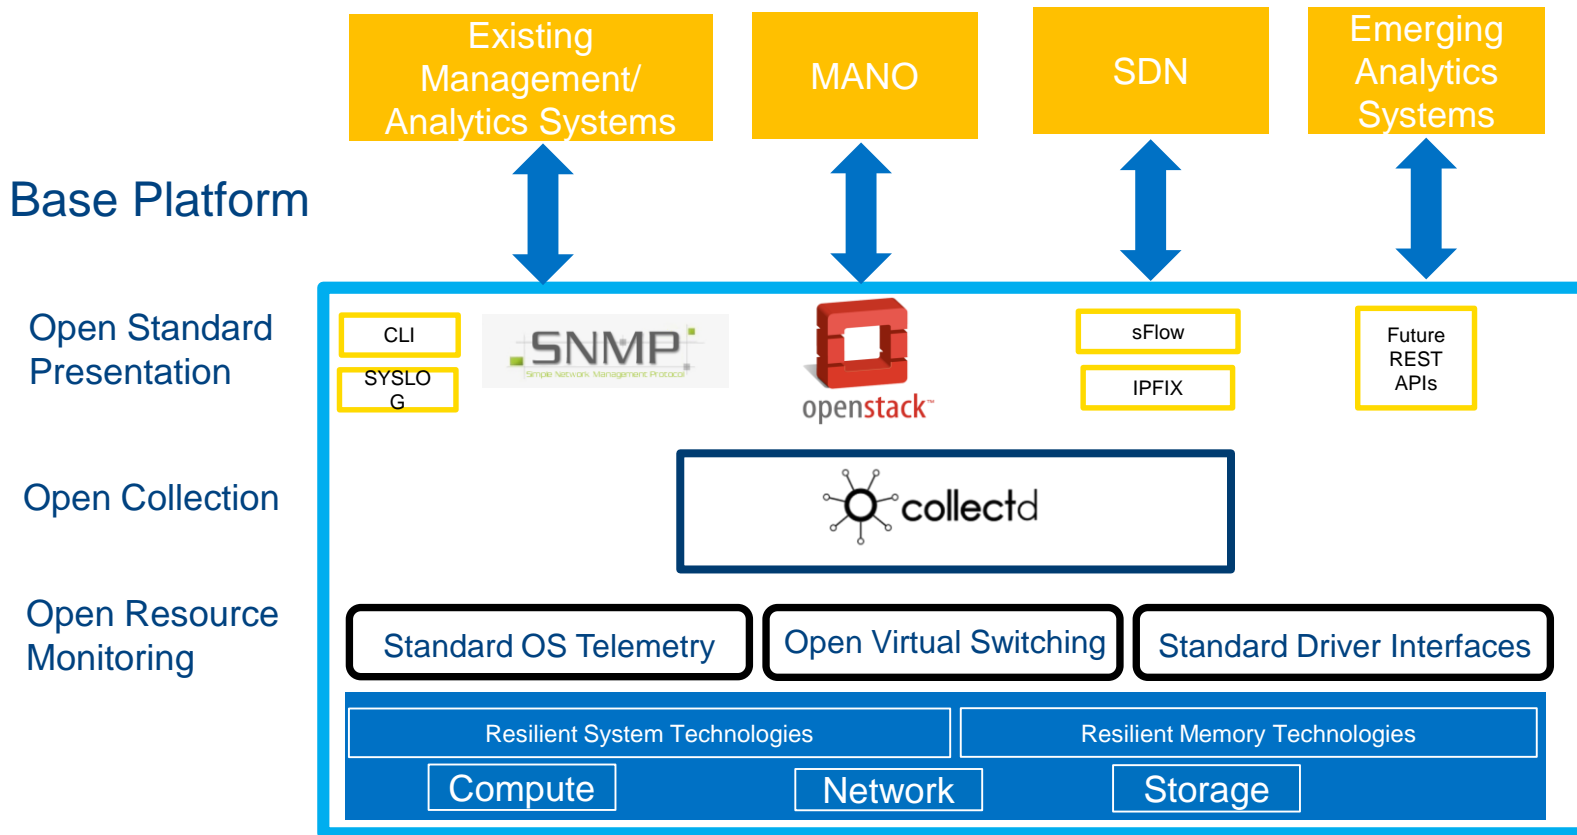
Platform Service Assurance

Cloud Ready NFV Platform with RAS (Reliability, Availability and Serviceability)

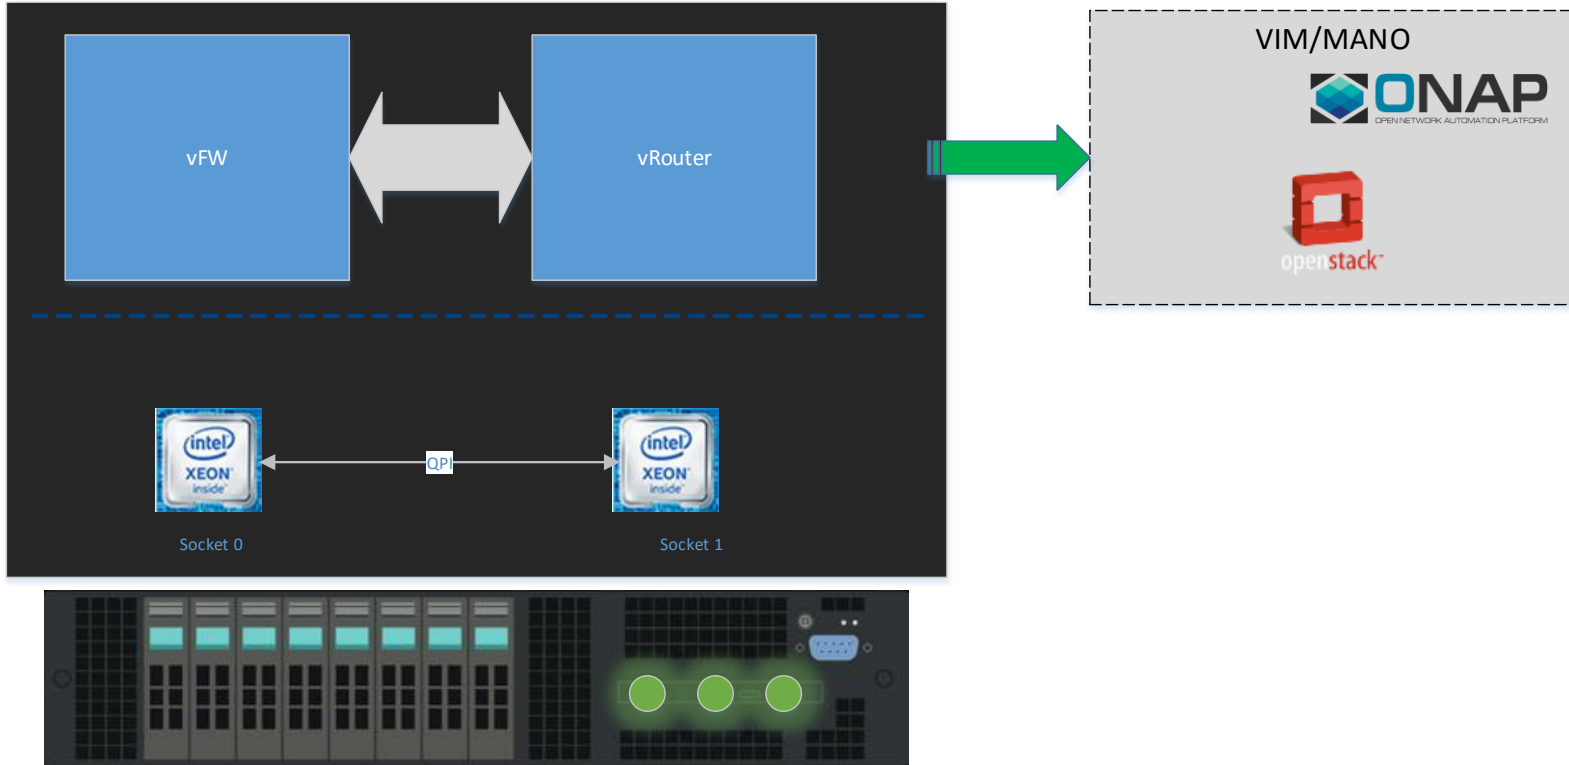
- Provisioning, monitoring and service impacting fault detection for infrastructure platforms
- Enable Reactive and Pro-active fault detection, fault reporting and support corrective actions
- Support current Industry Standard [FCAPS](#), Platform Telemetry and flow Telemetry interfaces
- Support Open Interoperable Interfaces to integrate with High Availability, NFV MANO, and Service Assurance software elements
- Extend Reliability (RAS) to include Software Fault Detection in addition to Hardware Faults

Open Interoperable Interfaces with Open Source Components

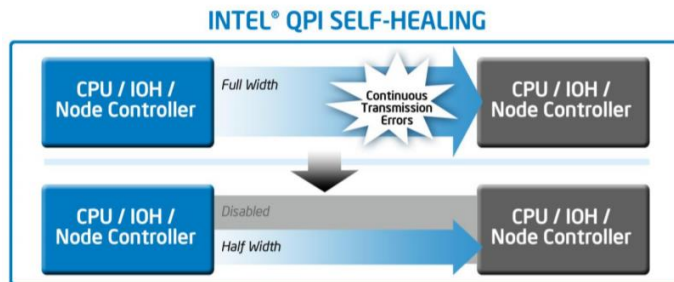


PLATFORM SERVICE ASSURANCE SYSTEM




The lifetime of an event – Before Intel Platform SA



The lifetime of an event – Before Intel Platform SA




The lifetime of an event – Before Intel Platform SA

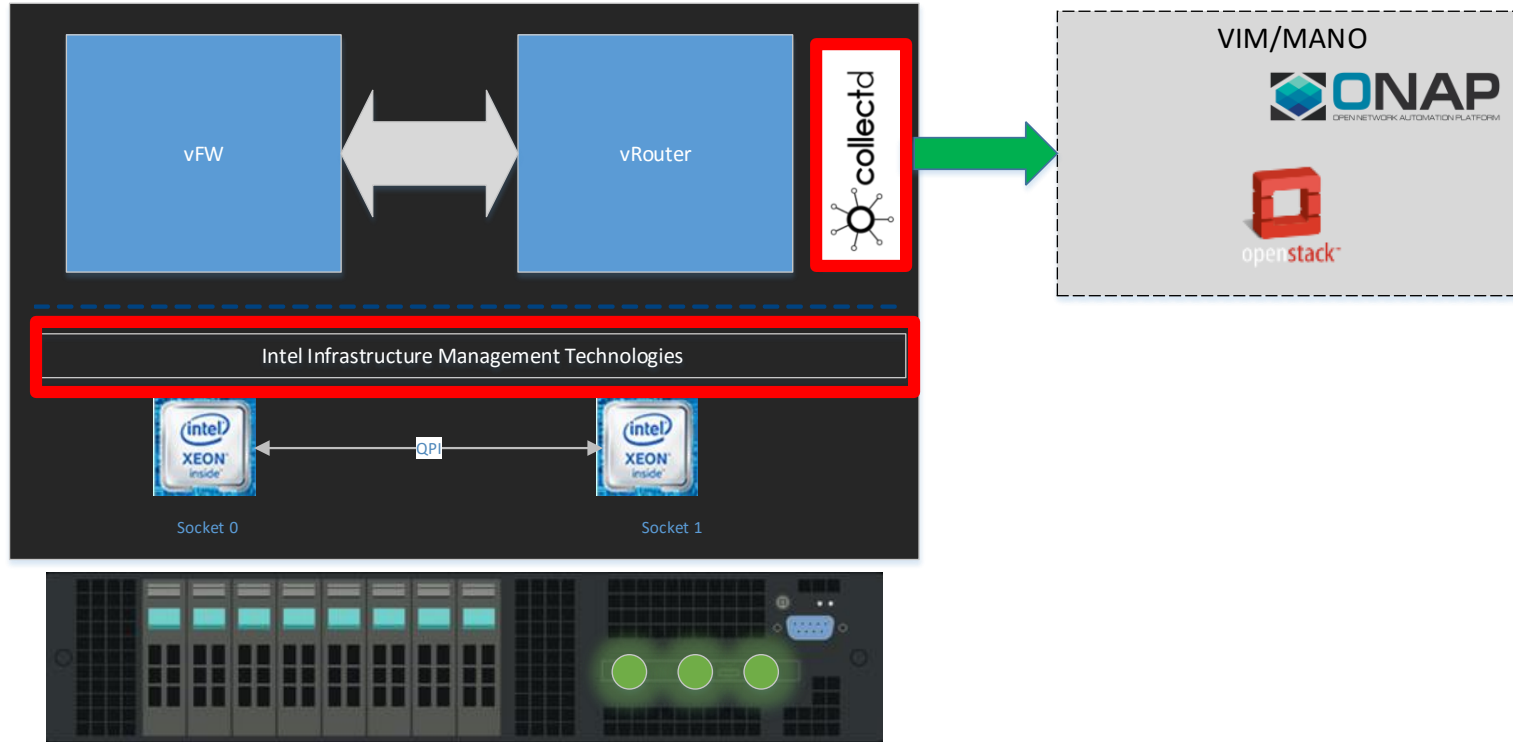


Not aware of platform issues resulting in SLA violation

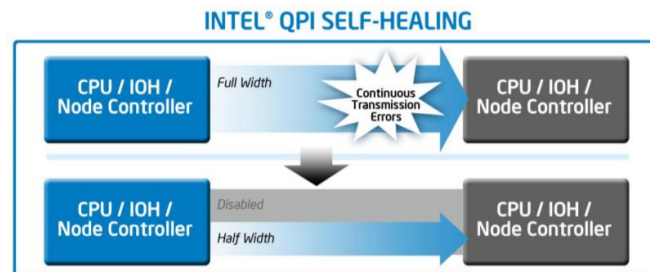
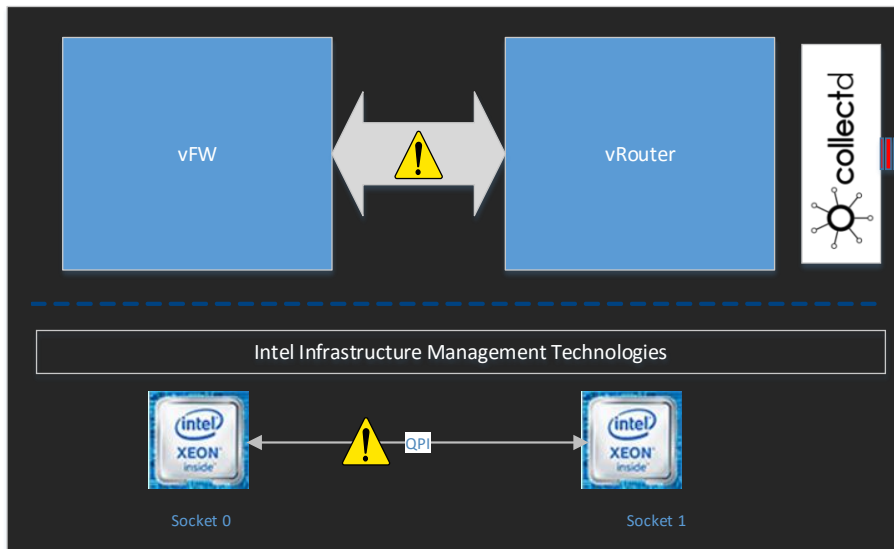
RCA: UNKNOWN
Just sees a degradation in VNF performance



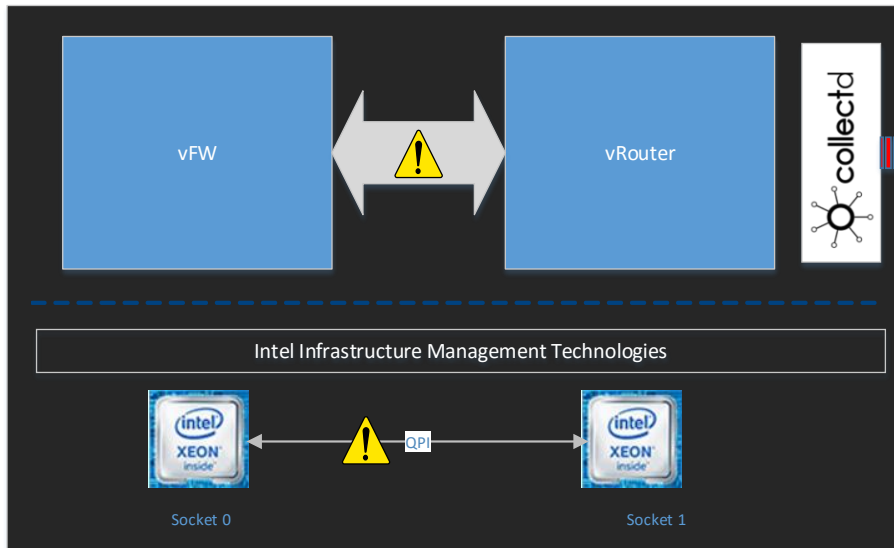
The lifetime of an event – with Intel Platform SA



The lifetime of an event – with Intel Platform SA



The lifetime of an event – with Intel Platform SA




Now aware of platform issues resulting in SLA violations


RCA: NFVI issue resulted in VNF performance degradations

Can make intelligent placement decisions/adjust scheduling policy based on NFVI event

INTEL® QPI SELF-HEALING



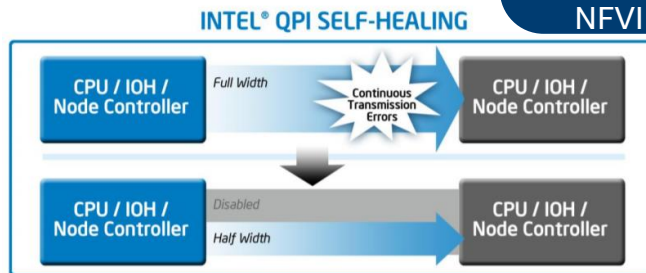
The lifetime of an event – with Intel Platform SA



Now aware of platform issues resulting in SLA violations

RCA: NFVI issue resulted in VNF performance degradations

Can make intelligent placement decisions/adjust scheduling policy based on NFVI event

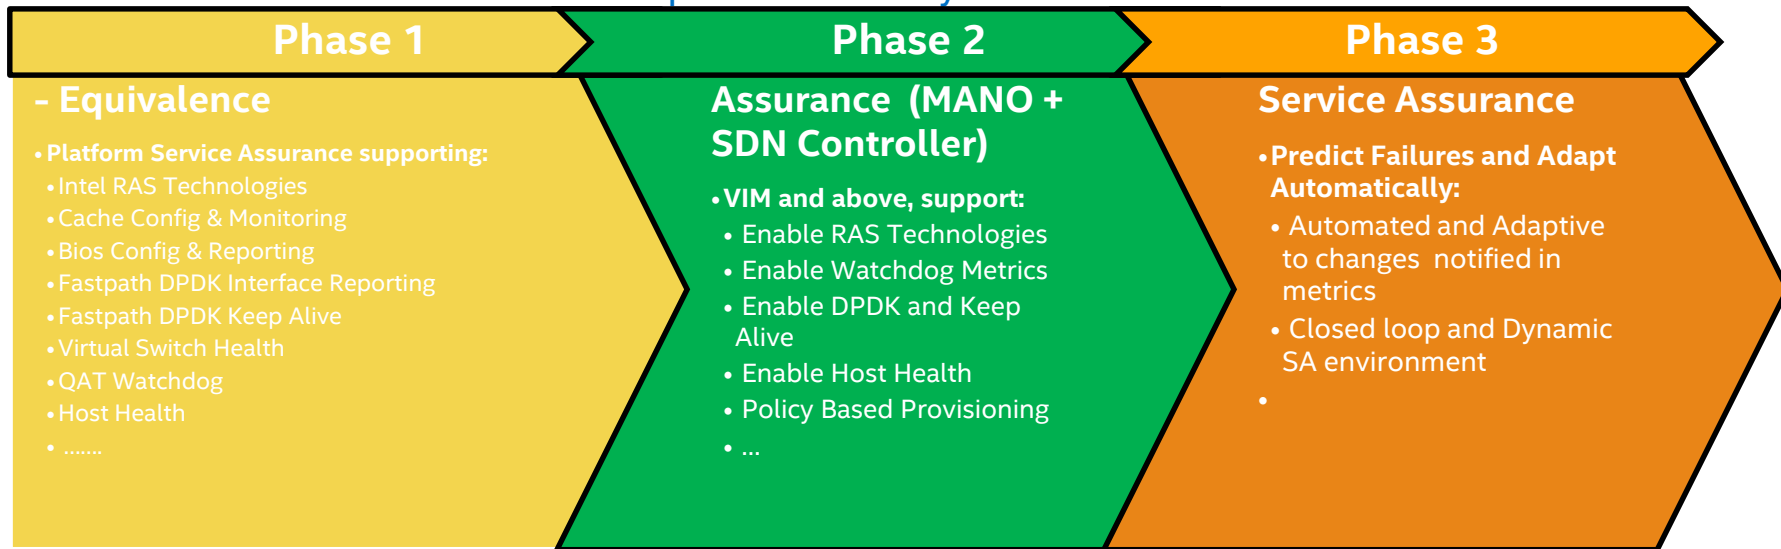


Ensuring the relevant telemetry exposed, understood and actioned End to End to showcase IA centric SA capabilities.



SERVICE ASSURANCE “PHASED” EVOLUTION FOR NFV/SDN

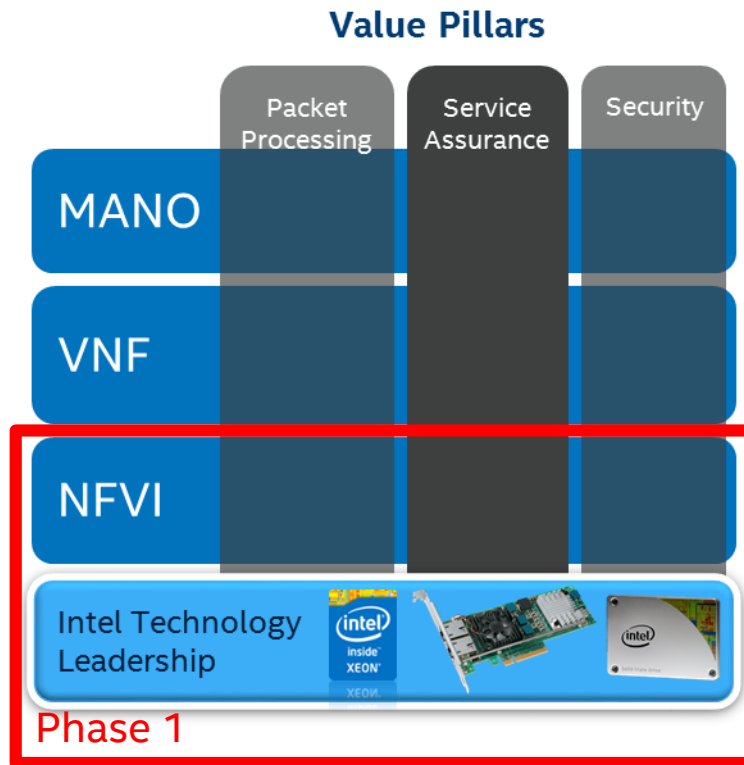
- Strategic Framework for SA “Phase” Evolution
- ✓ Phase 1 - Equivalence (Virtualized + Interworking with existing management systems)
- ✓ Phase 2 - Automated by MANO+SDN Controller
- ✓ Phase 3 - Predict failures and adapt automatically



Evolving from Equivalence towards NFV/SDN Automation

PLATFORM SERVICE ASSURANCE PHASE 1

- Support relevant metrics and events for capacity planning, trending and operational status.
- Ensure IA platform metrics and events are accessible through industry standard interfaces.
- Demonstrate how IA platform technologies can be monitored, consumed and actioned in real time.
- Showcase IA centric SA capabilities.



Ensuring the relevant telemetry is in place and exposed through an industry standard API

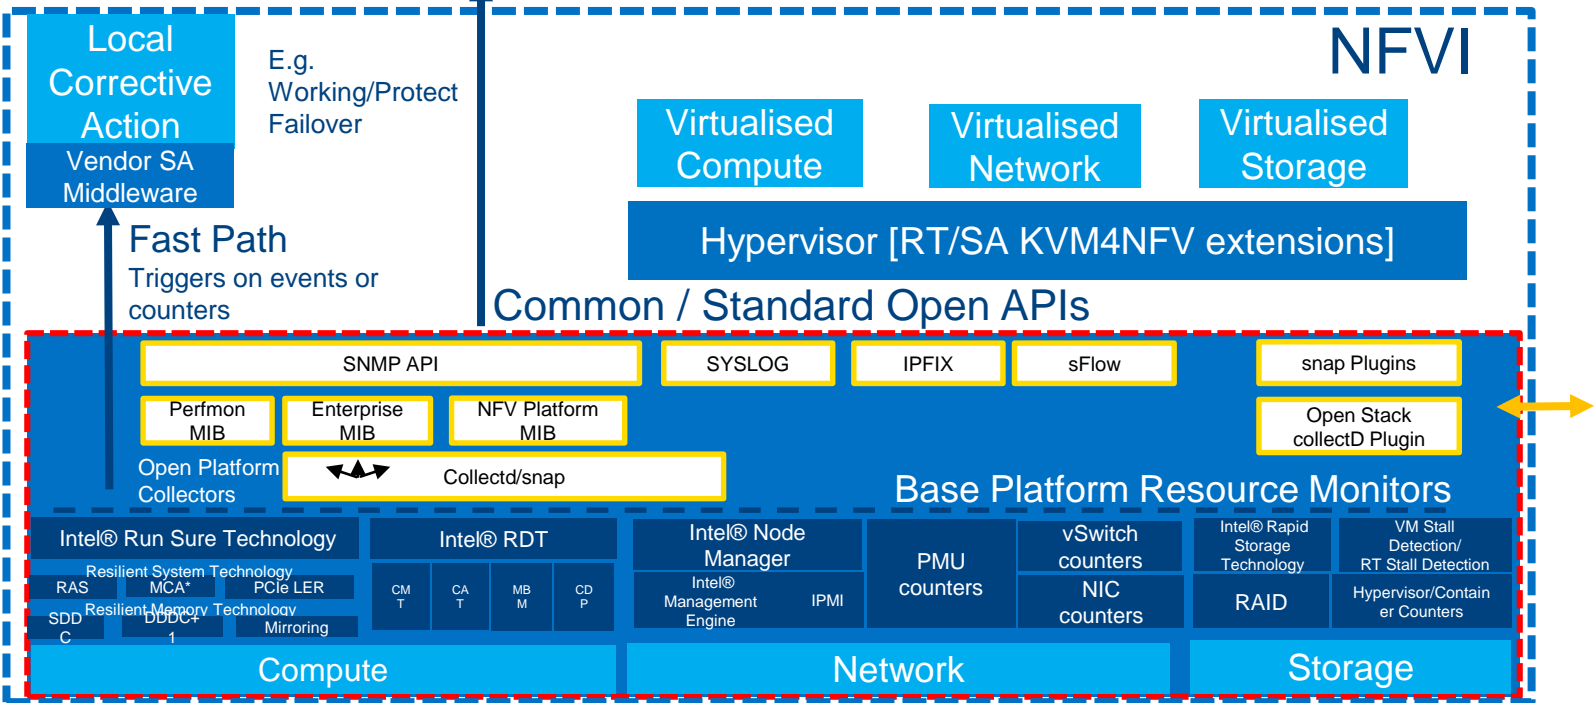
PHASE 1: CARRIER GRADE BASE PLATFORM SA (MONITORING)

Includes
NetFlow
Collectors

Monitoring/
Analytics
Systems

Slow Path
Periodic Pull 1/15mins

Standard Open
APIs
Intel Components



Orchestrator

VES

VIM

Vitrage

Aodh

Ceilometer

Congress

Gnocchi


Neutron

Cinder

Watcher



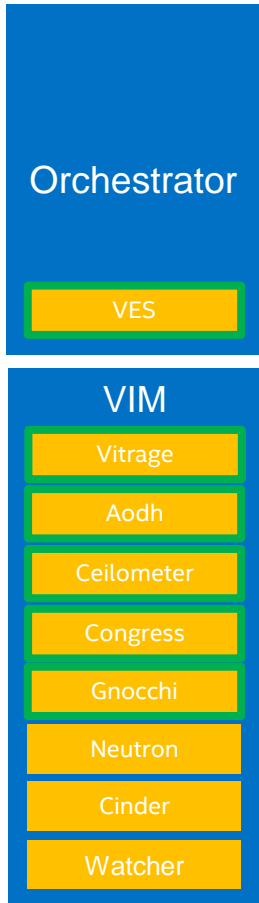
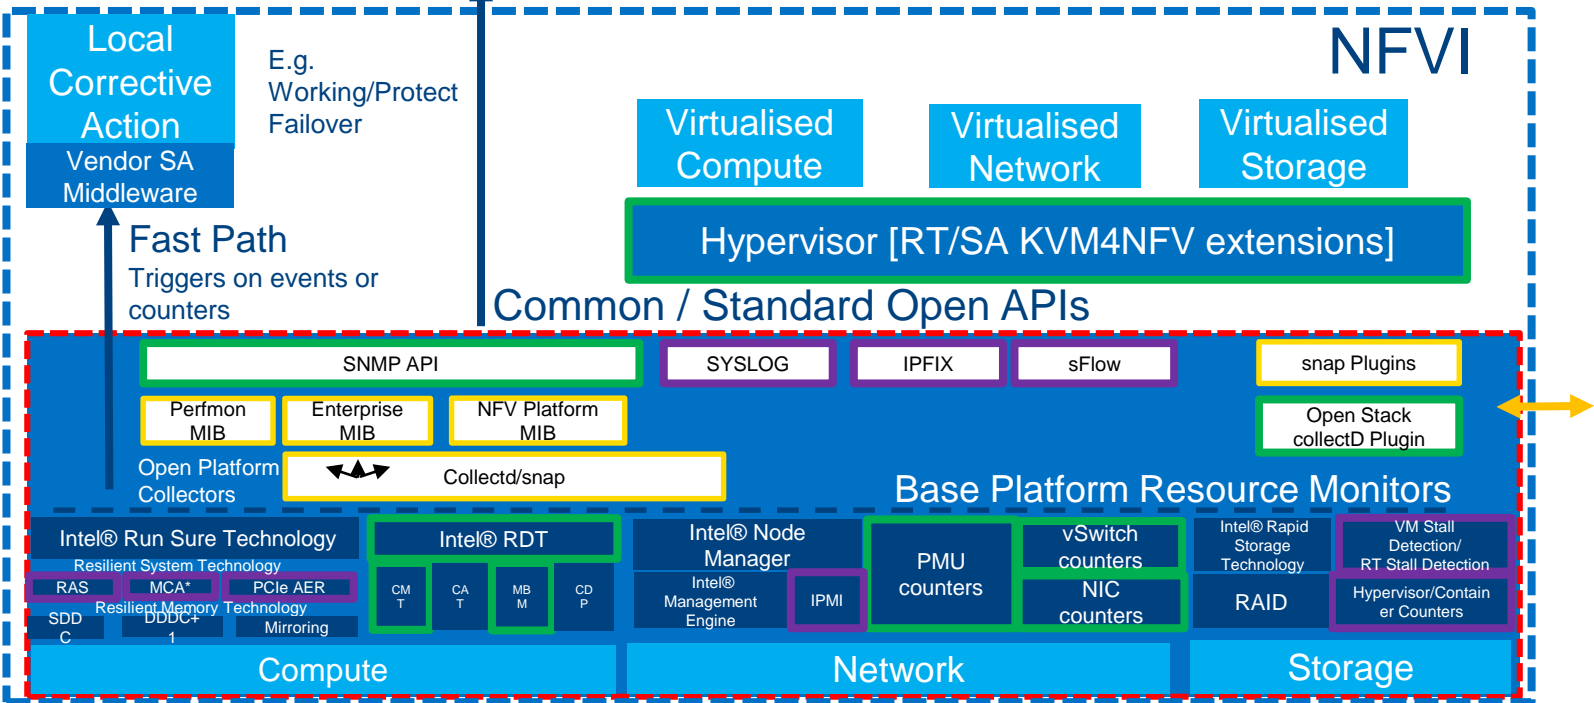
VISUAL: PHASE 1 PROGRESS



Includes
NetFlow
Collectors

Monitoring/
Analytics
Systems

Slow Path
Periodic Pull 1/15mins



Current Status I

Plugin	Description	Collect D	Collectd Version	OpenStack	RHOSP	Additional Info
DPDK Stats Plugin	A read plugin that retrieve stats from the DPDK extended NIC stats API.	Upstreamed	5.7.2	Pike	RHOSP 12	<p>Metrics and Events reported by these plugins can be found here: https://wiki.opnfv.org/display/fastpath/Collectd+Metrics+and+Events</p> <p>User guide: http://artifacts.opnfv.org/barometer/docs/index.html#document-release/userguide/index</p>
RAS Memory Plugin	A read plugin that uses mcelog to check for memory Machine Check Exceptions and sends the stats for reported exceptions	Upstreamed	5.7.2	Pike	RHOSP 12	
RDT Cache Monitoring	A read plugin that provides the last level cache utilization and memory bandwidth utilization	Upstreamed	5.7.2	Pike	RHOSP 12	
Huge Pages	A read plugin that retrieves the number of available and free hugepages on a platform as well as what is available in terms of hugepages per socket.	Upstreamed	5.7.2	Pike	RHOSP 12	
BIOS Plugin		Was not accepted upstream	N/A	N/A	N/A	
vSwitch Stats	A read plugin that retrieves interface stats from OVS.	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	
vSwitch Events	A read plugin that retrieves events (like link status changes) from OVS.	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	
DPDK Events Plugin	A read plugin that retrieves DPDK link status and DPDK forwarding cores liveliness status (DPDK Keep Alive).	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	
DPDK Events Plugin	A read plugin that retrieves DPDK link status and DPDK forwarding cores liveliness status (DPDK Keep Alive).	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	
Libvirt	A read plugin that uses virtualization API libvirt to gather statistics about virtualized guests on a system directly from the hypervisor, without a need to install collectd instance on the guest.	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	

Current Status II

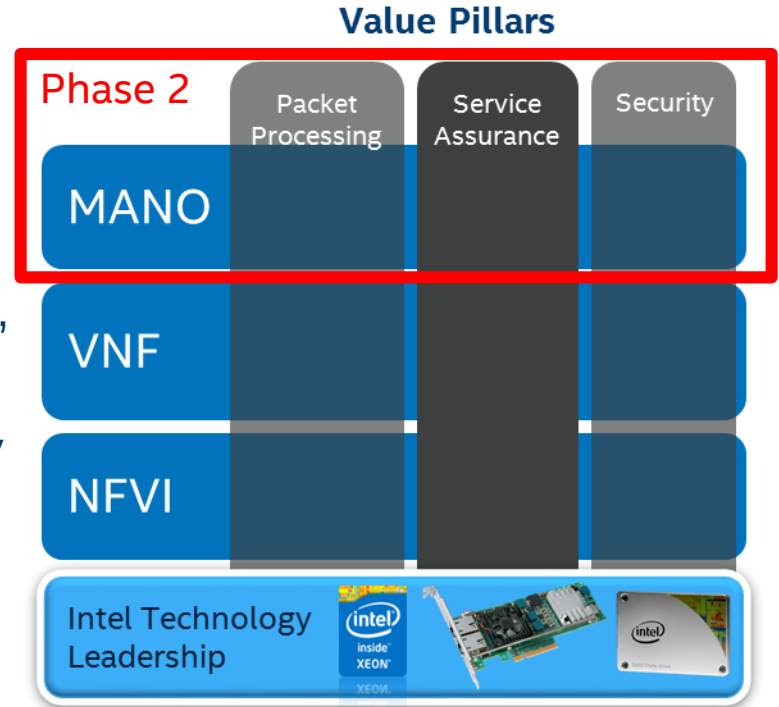
Plugin	Description	Collect D	Collectd Version	OpenStack	RHOS P	Additional Info
PMU Plugin	A read plugin that collects performance monitoring events supported by Intel Performance Monitoring Units (PMUs). The PMU is hardware built inside a processor to measure its performance parameters such as instruction cycles, cache hits, cache misses, branch misses and many others. Performance monitoring events provide facilities to characterize the interaction between programmed sequences of instructions and microarchitectural sub-systems.	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	<p>Metrics and Events reported by these plugins can be found here: https://wiki.opnfv.org/display/fastpath/Collectd+Metrics+and+Events</p> <p>User guide: http://artifacts.opnfv.org/barometer/docs/index.html#document-release/userguide/index</p>
SNMP Write Plugin	A write plugin that will act as a AgentX subagent that receives and handles queries from SNMP master agent and returns the data collected by read plugins. The SNMP Agent plugin handles requests only for OIDs specified in configuration file. To handle SNMP queries the plugin gets data from collectd and translates requested values from collectd's internal format to SNMP format. Supports SNMP: get, getnext and walk requests.	Upstreamed	5.8.0 (not released yet)	Queens (if released in time)	RHOSP 13	
Gnocchi	A write plugin that pushes the retrieved stats to Gnocchi. It's capable of pushing any stats read through collectd to Gnocchi, not just the DPDK stats.	Upstreamed		Available since Stable/Mitaka in the collectd-ceilometer-plugin repo		https://github.com/opnstack/collectd-ceilometer-plugin
Aodh	A notification plugin that pushes events to Aodh, and creates/updates alarms appropriately.	Upstreamed		Available since Stable/Mitaka in the collectd-ceilometer-plugin repo		https://github.com/opnstack/collectd-ceilometer-plugin

Current Status III

Plugin	Description	Collect D	Collectd Version	OpenStack	RHOSP	Additional Info
Legacy/IPMI Feature	A read plugin that reports platform thermals, voltages, fan speed, current, flow, power etc. Also, the plugin monitors Intelligent Platform Management Interface (IPMI) System Event Log (SEL) and sends appropriate notifications based on monitored SEL events.	Submitted Under Review	TBD	TBD	TBD	<p>Metrics and Events reported by these plugins can be found here: https://wiki.opnfv.org/display/fastpath/Collectd+Metrics+and+Events User guide: http://artifacts.opnfv.org/barometer/docs/index.html#document-release/userguide/index</p>
RAS Other Plugins	A read plugin that uses mclog to check for cpu, IO, QPI or system Machine Check Exceptions and sends the stats for reported exceptions	Submitted Under Review	TBD	TBD	TBD	
QAT Plugin		Ready for upstreaming	TBD	TBD	TBD	
PCIe AER	A read plugin that monitors PCIe standard and advanced errors and sends notifications about those errors	Ready for upstreaming	TBD	TBD	TBD	
Storage	Not started					

PLATFORM SERVICE ASSURANCE PHASE 2

- From Faults towards Full Automation
- Today: Reactive Fault detection...
- Tomorrow: Automated “Watch”, Decide, Act, Learn
 - Different layers of the stack automatically act on critical alerts.
 - Apply analytics for greater efficiency, agility and proactivity.



Enabling Operations Transformation via NFV and Service Assurance

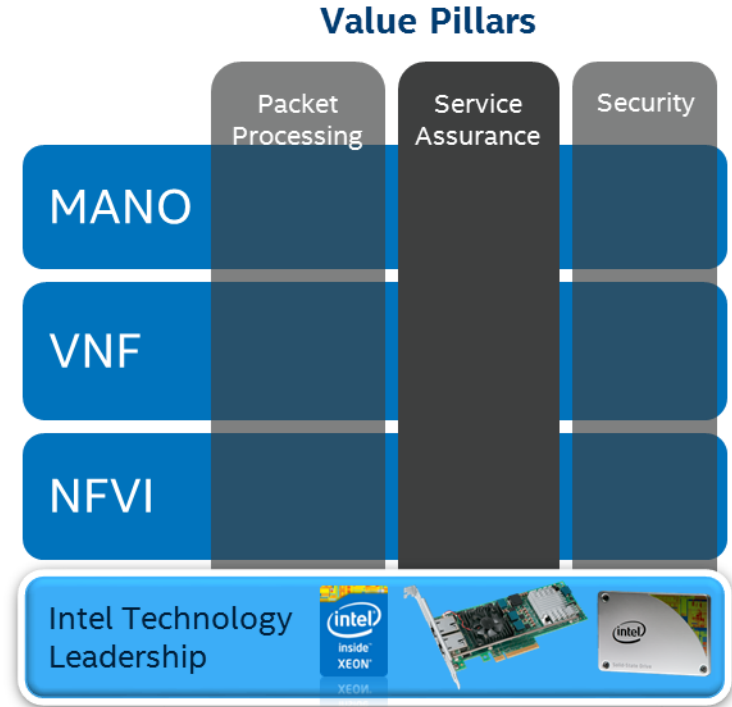
Next Steps

Priority	Use Case	Realization
1	VIM and MANO: Ensuring the telemetry is understood and acted on	<ul style="list-style-type: none">• NFVI Stats and events are available to the VIM and MANO through internal and customer collaboration (Gnocchi, Aodh, Ceilometer, Vitrage, congress, and VES (ONAP DCAE)).• Integrate with key monitoring and policy projects in ONAP.
1	Machine Learning and Analytics Reduce the dataset/deduce the KPIs/key events that are needed for failure prediction, auto-healing...	<ul style="list-style-type: none">• Collaboration with a number of data scientists to leverage their ML framework for feature selection and orchestration.• Integration with OPNFV projects to provide NFVI data associated with VNF characterization/fingerprinting.• Integration with PNDA through OPNFV Bamboo
2	Micro-services and Containers Integration: Platform telemetry, health and capabilities reporting	<ul style="list-style-type: none">• Container integration.• Devise a platform health Micro-Service to evaluate the platform condition on demand, or periodically – integrate with MANO (ONAP).• Devise a platform Capabilities Micro-Service and integrate with MANO inventory management (ONAP).

SUMMARY

PLATFORM SERVICE ASSURANCE SUMMARY

- Identified as major blocker across the ecosystem to NFV deployment
- Standard Open Interfaces to help all our customers go faster by giving them the best reference platform to build on
- Requirement: Expose underlying platform features and controls



Intel can help lead acceleration of Service Assurance for SDN/NFV but it needs ecosystem partners

Questions to the ONAP Community

- Does NFVI telemetry fit into DCAE?
- What is going to provision collectors like VES in ONAP today?
- What will manage the physical infrastructure in ONAP? Is that a better fit for collectd?
- Will a collector control plane API be defined for DCAE?





OPENSTACK BOSTON SUMMIT: NOISY NEIGHBOUR CORRECTION

NOISY NEIGHBOUR DETECTION AND CORRECTION DEMO

NOKIA



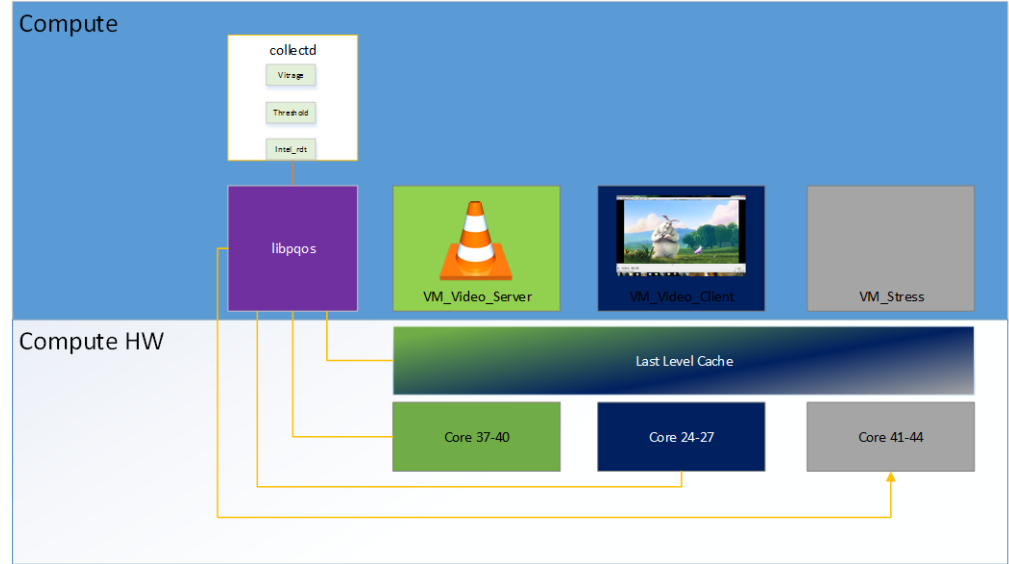
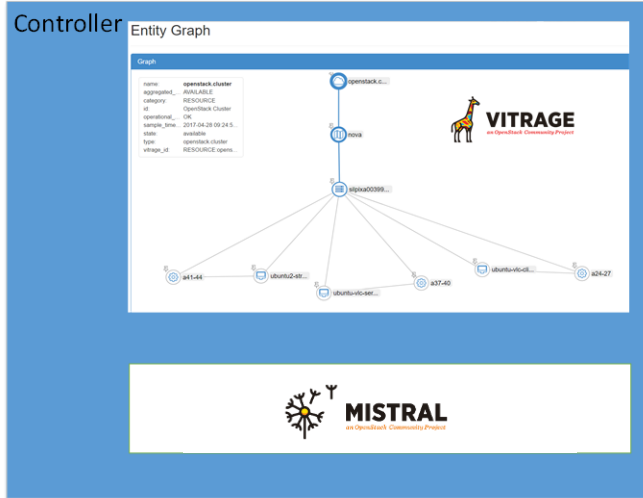
Question:

What technologies are available to detect and correct a Noisy Neighbour in an OpenStack environment without migrating workloads?

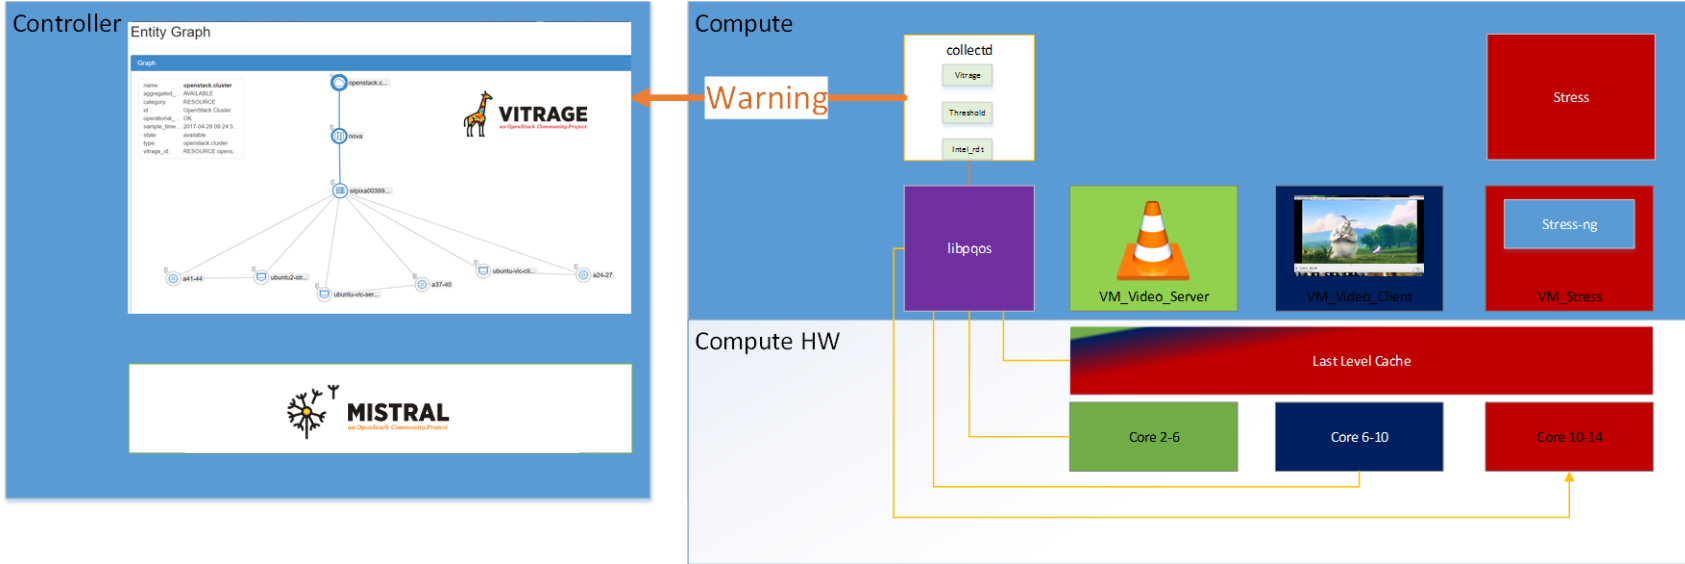
Answer:

- Noisy neighbours in an NFV/OpenStack environment can be detected with **collectd** using **INTEL® RESOURCE DIRECTOR TECHNOLOGY (RDT)**.
- **Notifications** of this event are propagated to **Vitrage**, where the problem and its impact can be **root caused, visualized** and **exported**.
- Finally, in response to the event and the analysis **corrective actions** will be taken to fix violations on the platform using **Mistral + INTEL® RDT**.

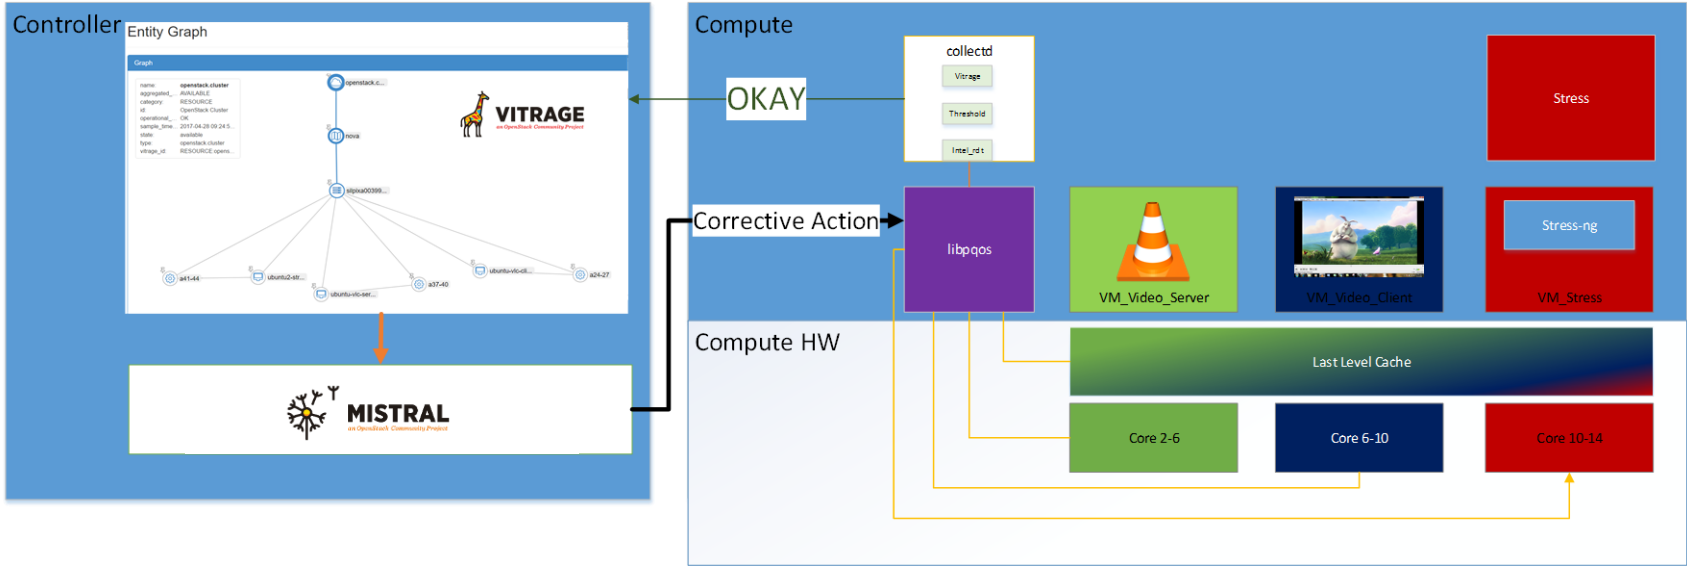
NOISY NEIGHBOUR DETECTION AND CORRECTION DEMO



NOISY NEIGHBOUR DETECTION AND CORRECTION DEMO



NOISY NEIGHBOUR DETECTION AND CORRECTION DEMO



NOISY NEIGHBOUR DETECTION AND CORRECTION DEMO

NOKIA



The screenshot displays a network management interface with three main sections:

- Entity Graph:** A hierarchical tree diagram showing network entities. At the top is 'openstack.dns...', followed by 'nova'. Below 'nova' is 'openstack...', which branches into several 'ubuntu-' nodes (e.g., 'ubuntu-44', 'ubuntu-45', 'ubuntu-46', 'ubuntu-47', 'ubuntu-48', 'ubuntu-49', 'ubuntu-4a', 'ubuntu-4b', 'ubuntu-4c', 'ubuntu-4d', 'ubuntu-4e', 'ubuntu-4f', 'ubuntu-4g', 'ubuntu-4h', 'ubuntu-4i', 'ubuntu-4j', 'ubuntu-4k', 'ubuntu-4l', 'ubuntu-4m', 'ubuntu-4n', 'ubuntu-4o', 'ubuntu-4p', 'ubuntu-4q', 'ubuntu-4r', 'ubuntu-4s', 'ubuntu-4t', 'ubuntu-4u', 'ubuntu-4v', 'ubuntu-4w', 'ubuntu-4x', 'ubuntu-4y', 'ubuntu-4z').
- Video Player:** A VLC media player window showing a video of a white rabbit holding a carrot. The URL is 'http://172.24.4.7:8080/'.
- Alarms Analysis:** A table with columns: Time Stamp, Name, Resource Type, Resource ID, Severity, Type, and RCA. The table is currently empty.

At the bottom left, there is a graph titled 'RDY Value' with a y-axis from -1.0 to 1.0 and an x-axis from 08:39:30 to 08:44:00. The graph shows 'No data points'.

https://www.youtube.com/watch?v=DmSiXS_tUEY

Questions/backup