WHITE PAPER

Enhancing Service Assurance for Virtualized Networks with Intel® Platform Technologies

Authors

Intel Corporation

John Brown
Software Architect

Tim Verrall
Systems Architect

Eoin Walsh
Solutions Architect

Maryam Tahhan
Network Software Architect

James Greene
Technology Marketing Manager

Pat Vaughan
Program Manager

Executive Summary

Providing robust service assurance capabilities is critical in the network transformation to a software-defined and increasingly virtualized network environment. It is vital to monitor systems for utilization and malfunctions that could lead to service disruption in order to facilitate the prompt resumption of normal or improved service. Today, monitoring and management activities throughout the network are supported by discrete systems in fixed service chains with tightly integrated hardware and software products and established management frameworks and assurance tools. In a virtualized environment, one based on Network Functions Virtualization (NFV) and Software Defined Networks (SDN) for example, these activities are more challenging as a result of the disaggregation of hardware and software and the ability to deploy services dynamically.

Service assurance in a virtualized world requires continuous monitoring of the platform hardware, virtualized environment, and services software. The scope of the platform components and tools described in this paper is the provisioning of platform resources, the collection of a growing set of platform performance, fault and other useful data, and the sharing of that data with management, analytics, and orchestration systems such that all the physical, virtual and service resources can all be more thoroughly provisioned, managed, monitored and measured. And of course, to be efficient, all of this should be accomplished using traditional network and systems management tools and in time through heterogeneous supplier environments.

This paper is intended for managers, administrators, and others responsible for planning and implementing the virtualization of their networks. It outlines capabilities that enable the collection and sharing of platform performance, fault and configuration data to conventional service assurance, management and analytics systems, while providing the same rich platform telemetry to NFV orchestration systems and SDN systems. As such, it provides leaders and implementers with an evolutionary path to NFV and SDN as core components in fully automated and predictive management and orchestration systems for modern, efficient networks.

Table of Contents

Executive Summary 1
Introduction 2
The Traditional Approach to Service Assurance 3
Platform Service Assurance in ETSI. . . 4
Service Assurance Requirements of Hardware Resources 5
Intel Infrastructure Management Technologies 5
Service Assurance Requirements of the Virtualization Layer and Virtual Resources 6
From the Platform to Service Assurance 6
Open Source Monitoring 7
Open Collectors 7
Open Standard Presentation 7
Integration into Northbound Management Systems 7
Coming Together: An Example 7
Conclusion 9
References 9

Figures

1. Network Management Today. 3
2. Platform for NFV Service Assurance Aligned with ETSI NFV Architecture 4
3. Intel Infrastructure Management Technologies 5
4. Platform for NFV Service Assurance Layers and Consumers (source adapted from ETSI) 6
5. Data flows for the "Noisy Neighbor" scenario in an enabled environment. 8

Introduction

Communications service providers (CommSP) today have many siloes with service fulfilment stacks providing various end-to-end services, network management, and configuration management capabilities—each based on specific proprietary or fixed-function hardware/software stacks. NFV and SDN technologies promise hardware vendor independence, improved operational efficiency, standardized and open interfaces, and the dynamic chaining of network functions to create services. To begin to gain these NFV/SDN benefits, today's networks are commencing a process of transition moving from or supplementing the traditional physical network functions (PNFs) to corresponding virtual network instantiations (VNFs) of physical functions. For the foreseeable future, networks will continue to be made up of a combination of legacy, virtual (NFV) and SDN technologies, also known as the NFV/SDN "hybrid" network.

Supporting the transition to hybrid networks requires virtualization working side by side with legacy infrastructure, evolving in time towards the real-time, dynamic imperatives of on-demand network adaptation and healing. In both telecommunication and enterprise companies, service assurance is the application of policies and processes to ensure that the services offered over networks meet pre-defined quality levels for optimal subscriber or end user experience. As such, service assurance of hybrid networks becomes critical as legacy and virtual networks co-exist, and it plays a critical role in maintaining the consistency of metrics for the services as the virtual and physical elements of the service interoperate in hybrid networks.

NFV deployments must be backward compatible with existing service assurance and network management toolsets that cover both the physical and virtualized components. For these components, there is a need to monitor a new set of vectors that span the totality of hardware and software components across the physical and virtualized environments. The following examples show just some of the diversity of NFV deployments and use cases that benefit from a common base platform capable of providing support for service assurance:

- Managed Services - Virtual customer premises equipment (vCPE), virtual home, virtual customer edge (vCE) (WAN edge)
- Enterprise Managed WAN - Virtualizing WAN optimizer, router, firewall, session border controller (SBC), proxy, media gateways, and so on
- Telecommunications Networks – Cloud radio access network (CRAN)/virtual radio access network (VRAN), virtual evolved packet core (vEPC), virtual radio network controller (RNC), virtual IP multimedia subsystem (IMS), virtual broadband remote access server (B-RAS)/broadband network gate (BNG)

Pulling these solutions together will require multiple components to become "service assurance aware." The VNFs must be able to be identified and their resource requirements published such that policies for placement and management can be defined and enforced. The physical and virtual infrastructure must be capable of showing the resources that they are providing or have available, and how they are behaving at an increasingly granular level such that this information can be used in VNF or service placement, monitoring and back-office operations. Intel® has been deeply engaged in the industry to help address the need for service assurance capabilities in virtual network environments. Much of the leadership work has been focused on enabling the reporting, configuration and utilization information of critical platform resources with special emphasis on enabling open source projects. This will be the foundation for instrumenting the service assured NFV infrastructure and provides a resource for the industry to leverage for more rapid adoption.

The Traditional Approach to Service Assurance

The current industry approach to service assurance is based on the Fault, Configuration, Accounting, Performance, and Security (FCAPS) model. In practice, this is roughly:

- Fault – Platform or system faults are critical to the detection, correction, and root cause isolation of events that impact service availability.

  Faults are managed primarily using Simple Network Management Protocol (SNMP) and are typically polled to be used in combination with fault management traps. Syslog parsing is also carried out together with SNMP management.

- Configuration– Configuration is the manageability element of the system and provides provisioning of local resources, services, and interfaces.

  Configuration can be manual or automated using proprietary tools, scripts, also using NETCONF for automating configuration. Representational State Transfer, or REST, interfaces allow for efficient deployment and management of the network resources.

- Accounting – Accounting capabilities provide the usage data for metrics related to billing.

  Accounting leverages flow data from NetFlow*, Sampled Flow (sFlow), or Internet Protocol Flow Information Export (IPFIX) to account for networking metrics, chargebacks, and billing. Call detail records can also be used to diagnose certain types of faults.

- Performance – These include the measurement and reporting of performance indication metrics used to track adherence to service level agreements.

  Performance measurements include NetFlow, sFlow, IPFIX, SNMP, synthetic traffic tools, e.g. Internet protocol service level agreement (IPSLA), to monitor and report on network performance characteristics to compare with service level agreements and other policies.

- Security – Security is essential to monitor, control and record access to platform and network resources.

  To manage access to network resources, various security tools and techniques are used, such as Authentication, Authorization and Accounting (or Triple A) using information from systems including RADIUS and Diameter, Terminal Access Controller Access-Control System Plus (TACACS+), Network Intrusion Protection System (NIPS), NetFlow, sFlow, IPFIX, and more.

In current physical networks, network management is typically performed using SNMP, or other standard collection and configuration tools in conjunction with proprietary tools such as individual element managers that are often required for product-specific management and configuration. The operations support system (OSS)/business support system (BSS) or traditional network management layer common to all the functions and services typically aggregates information from these sources and monitors the state of the entire network and the services being provided. Some typical measurements collected include interface bit-rate, packet rate, packets dropped (per-traffic class queue), queue depths and so on, which are all graphed and trended. Thresholds, baselines, and watermarks are used to decide when network conditions are out of specification, a service is being impacted, or capacity needs investigation. Flow data is also commonly used for tracking network usage, traffic volumes, security, application awareness, quality of service function and troubleshooting. Some primary tools for gathering flow data are NetFlow, sFlow, and IPFIX. In addition, synthetic traffic is used for tracking end-to-end network service level quality today.

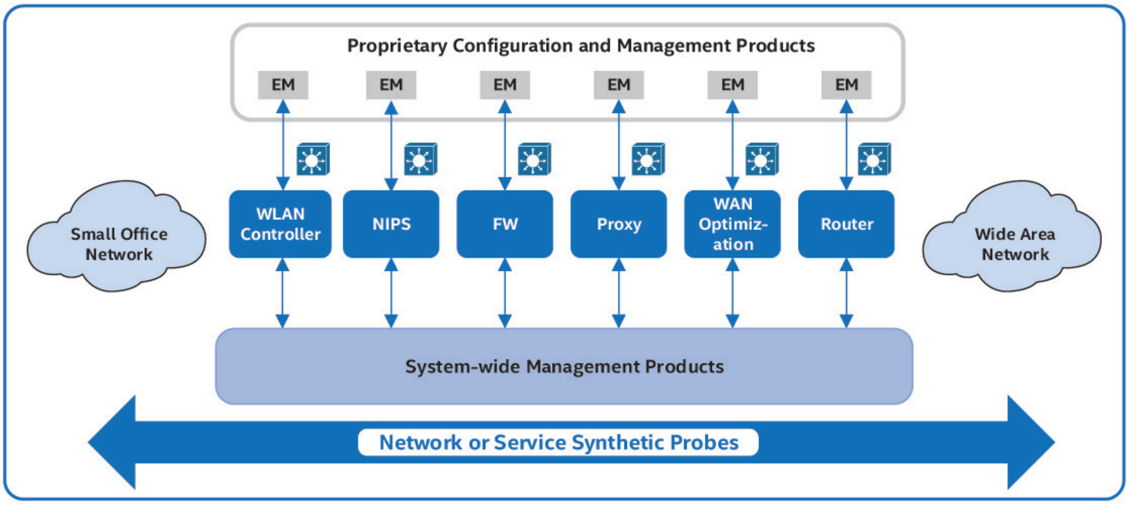
Proprietary Configuration and Management Products
EM
EM
EM
EM
EM
EM
Small Office Network
WLAN Controller
NIPS
FW
Proxy
WAN Optimiz-ation
Router
Wide Area Network
System-wide Management Products
Network or Service Synthetic Probes

Figure 1. Network Management Today

While companies see the potential cost-reduction, deployment-simplification and agility benefits of NFV deployments, they are also asking how to provide service assurance and quality of service (QoS) at a solution level for NFV deployments and how to deliver compliance with traditional "five nines" (99.999% up time) availability and reliability requirements. Simply, if virtualized appliances are to displace or supplement physical appliances, the service, reliability, and manageability must be equivalent or better than that available in networks today. The absence of deterministic service assurance levels is increasingly being identified as a barrier to broader NFV adoption and deployment. The assured reliability of services is not possible if you cannot assure or measure the reliability of the underlying platform. However, innovations on the platform and the enablement of service assurance solutions promise to break this barrier. Thus, platform service assurance provides interesting and essential capabilities to enhance network service assurance in virtualized, software defined networks.

The challenge is to enable the same level of functionality and integration in an open, standards-based NFV/SDN environment, while also dealing with some of the unique challenges that an NFV model might introduce (such as dealing with "noisy neighbors" or prioritizing shared resources for more critical applications). These challenges have been the area of deep focus by standards groups such as the European Telecommunications Standards Institute (ETSI), Intel, and the broader ecosystem.

Platform Service Assurance in ETSI

Platform service assurance is enabled on a NFV base platform that includes physical compute, storage, networking, virtual switch, host OS and hypervisor. The platform provides functionality to configure physical resources and monitor both physical and virtual resources. Platform telemetry provides capabilities that give greater insight and control to virtual and physical attributes and is aligned with the *ETSI NFV architecture as defined in the ETSI Network Functions Virtualisation (NFV) Architectural Framework* (ETSI GS NFV 002).

A service-assured platform allows the enforcement of policies and provides telemetry and fault information while minimizing downtime using industry-standard open interfaces. Integration with service assurance practices enable the identification and resolution of faults or degradations quickly with minimal service interruption. This includes the ability to proactively locate, diagnose, and repair degradations or malfunctions in service quality, improving availability of the service and minimizing the impact to users and subscribers. The same metrics used to measure and record service quality may also be used for capacity planning and issue avoidance. In this context, platform service assurance can be characterized as supporting the provisioning, monitoring, and service impacting fault detection for the NFV infrastructure that enables reactive and proactive fault detection, fault reporting, and supports corrective actions.

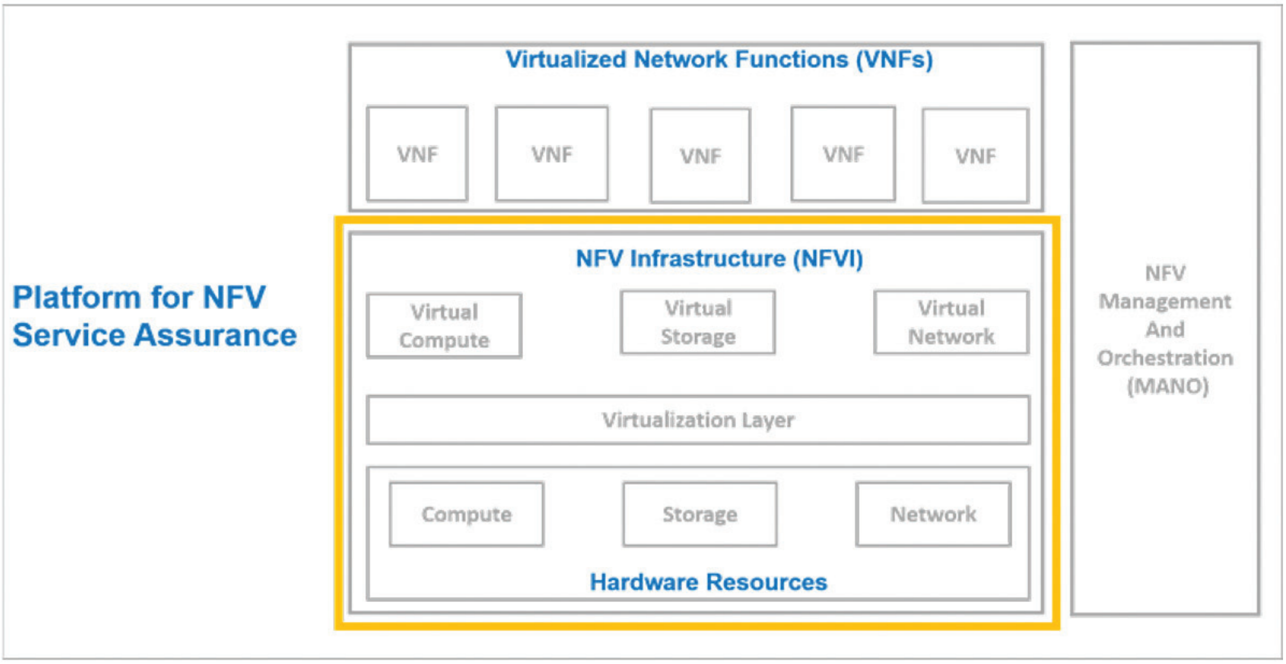
Platform for NFV Service Assurance
Virtualized Network Functions (VNFs)
VNF
VNF
VNF
VNF
VNF
NFV Infrastructure (NFVI)
Virtual Compute
Virtual Storage
Virtual Network
Virtualization Layer
Compute
Storage
Network
Hardware Resources
NFV Management And Orchestration (MANO)

Figure 2. Platform for NFV Service Assurance Aligned with ETSI NFV Architecture

Service Assurance Requirements of Hardware Resources

Hardware resource attributes that are useful for providing service assurance capabilities for NFV architectures focus on three key capabilities: hardware partitioning, hardware resiliency, and monitoring. Hardware resource partitioning enables allocation of key hardware resources to meet service resource requirements. Hardware resiliency is required to enable hardware components to auto-detect and correct transient hardware errors where possible (such as error correction in memory or PCI transfers). Monitoring capability is required to detect and report hardware resource metrics and faults to enable corrective action. As a whole, these capabilities unlock opportunities for real innovation to provide support for service assurance within the physical and virtual components of the NFV infrastructure.

Intel Infrastructure Management Technologies

Service assurance can only be supported when all the relevant parts of the system can be configured, managed, and recorded. From a platform perspective, this includes compute, storage, and networking components, such as, but not limited to, CPUs, cores, ports, links, accelerators, hypervisor, virtual switches, disks, and real-time delivery traffic QoS indicators. The platform can provide three critical sets of functionality that meet needs for service assurance uses. These include:

- Provisioning: Enabling configuration of specific service levels based on workload or service priority for:
  - Platform and workload: Includes allocating or partitioning platform resources such as CPU, memory, cache, and network bandwidth
  - Platform network interfaces: Includes setting bandwidth, QoS, rate limits per workload or VNF, and protecting bandwidth for each VNF
- Monitoring: Enabling deeper management and tracking of specific service levels:
  - Platform counters to track usage and performance to configured parameters
  - Network counters to track usage and performance to configured parameters
  - Service monitoring probes to record service levels
- Presentation: Reporting to enable reaction to service level changes:
  - Human intervention for threshold violations or failures
  - Dynamic intervention for threshold violations or failures
  - Support for the detection of trending against configured parameters and the enabling of capacity plan changes based on those trends

The Intel hardware platform provides the ability to provision, monitor and report on the environment via a rich and growing set of features covering power, fault, security, utilization, I/O, thermal, performance, capacity and more. With this information, it is possible to enable enhanced service assurance at the platform level and from there do much more. Due to this utility, the platform features (many of which are established and available for several CPU generations) are referred to as Intel Infrastructure Management Technologies. These are a suite of platform technologies that help to monitor, manage, and control the resources that support data and services management use models such as service assurance or software defined infrastructure.

As shown in Figure 3, there are numerous technologies that fall under the Intel Infrastructure Management Technologies label. Intel has published feature briefs that will describe how many of these can be enabled to enhance service assurance use models for NFV. Note that not all technologies are available on all platforms. Please check with your system manufacturer for supported features.

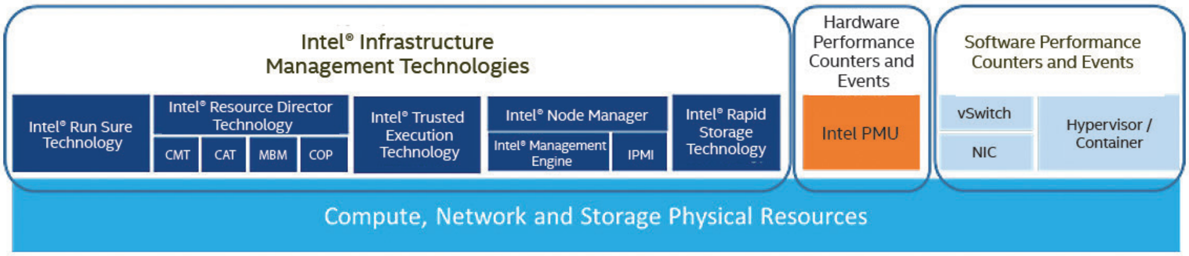
Intel® Infrastructure Management Technologies
Intel® Run Sure Technology
Intel® Resource Director Technology
CMT
CAT
MBM
COP
Intel® Trusted Execution Technology
Intel® Node Manager
Intel® Management Engine
IPMI
Intel® Rapid Storage Technology
Hardware Performance Counters and Events
Intel PMU
Software Performance Counters and Events
vSwitch
NIC
Hypervisor / Container
Compute, Network and Storage Physical Resources

Figure 3. Intel Infrastructure Management Technologies and other useful telemetry sources

Resource Partitioning and Monitoring Using Intel Resource Director Technology (Intel RDT)

Intel Resource Director Technology (Intel RDT) provides mechanisms to partition key platform resources, such as CPU cache. In addition, Intel RDT provides features to track cache utilization and memory bandwidth on Intel Xeon E5v4 and Intel Xeon E7v4 platforms.

Cache and memory platform features can be configured and monitored on a virtual platform deployed in a traditional managed infrastructure with the same capabilities available in an NFV/SDN deployment. The network operator can use NFV alongside the existing physical network functions, achieving a similar service level from the virtualized network equipment as that obtained from the fixed function equipment. This helps to provide a seamless path to NFV without degrading services or decreasing supported service levels.

Resiliency and Monitoring Using Intel Run Sure® Technology

Intel Run Sure® Technology is, itself, a suite of features that provide the capability to auto-detect and correct transient hardware errors. They provide advanced CPU, memory disk monitoring and recovery features, which includes the detection and reporting of potentially service impacting faults. Intel Run Sure Technology includes two groups of technology functions: Resilient system technologies and resilient memory technologies, providing the platform with built-in reliability, availability, and serviceability (RAS) features. Resilient system technologies integrate processor, firmware, and software layers that allow the system to diagnose and/or recover from previously fatal errors. These include processor RAS features, such as error correcting code (ECC) and parity check, clone detection and cataloging method (CDCM), Intel QuickPath Interconnect (Intel QPI) healing, corrected machine check interrupt (CMCI), machine check architecture (MCA), and CPU hot-add.

Resilient memory technologies ensure data integrity and enable systems to keep running reliably. These include RAS features such as memory demand, single device DRAM correction (SDDC), memory mirroring, Intel Scalable Memory Interconnect (Intel SMI) reliability, and failed dual in-line memory module (DIMM).

Service Assurance Requirements of the Virtualization Layer and Virtual Resources

Of course, one still needs to know what is happening in the virtual environment to manage the entire solution. The virtualization layer requires an open interface to manage virtual machines and monitor virtualized resources. The libvirt toolkit provides tools and an open-standard interface to manage virtual machines and other virtualization functionality, such as storage and network interface management. Other host and hypervisor management tools (including elements of the virtual infrastructure management (VIM) are gaining similar instrumentation for reporting on virtual infrastructure health, capacity and utilization.

Virtual network functions (VNFs) that require low-interrupt latency and timing correctness place extra requirements on the virtualization layer. To meet the low-latency requirements, the open hypervisor KVM can be extended with functionality provided for example by the OPNFV* project, KVM4NFV. In another step of progress, a growing set of VNFs are becoming more “environment aware,” providing a manifest of resources they require that allow these VNFs to be more effectively managed in environments where the underlying resources are granularly identifiable and the VIM and MANO can use this information for workload placement and service management.

From the Platform to Service Assurance

Service assurance is enabled through alignment of communications through three layers: the presentation layer, the collection layer, and resource telemetry interfaces.

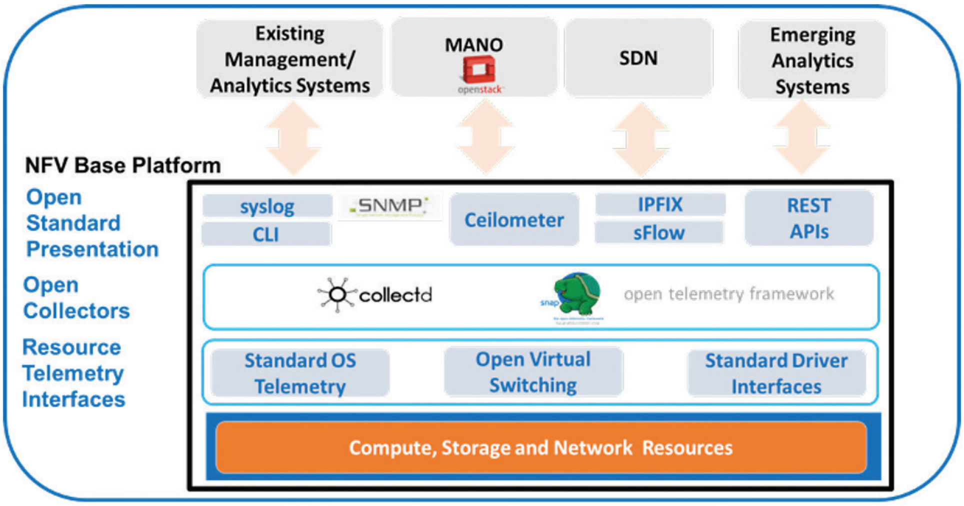
Existing Management/ Analytics Systems
MANO
openstack
SDN
Emerging Analytics Systems
NFV Base Platform
Open Standard Presentation
Open Collectors
Resource Telemetry Interfaces
syslog
CLI
SNMP
Ceilometer
IPFIX
sFlow
REST APIs
collectd
snap
open telemetry framework
Standard OS Telemetry
Open Virtual Switching
Standard Driver Interfaces
Compute, Storage and Network Resources

Figure 4. Platform for NFV Service Assurance Layers and Consumers

The presentation layer provides open industry standard interfaces to report metrics and telemetry as well as providing open interfaces to provision the platform resources. The collection layer uses open collection agents, such as collectd and Snap, to aggregate metrics and provide those metrics to the presentation layer. A collection agent provides a single aggregation point for all metrics on the platform and simplifies the translation effort when rendering to multiple reporting interfaces in the presentation layer. The resource telemetry interfaces provide all the platform metrics gathered from the hardware and software. These metrics are provided to the collection layer, which, in turn, aggregates, thresholds and translates the metrics for the presentation layer.

Open Source Monitoring

The resource telemetry interfaces provide open standard interfaces to both hardware and software resources for the reporting of telemetry and fault information. Some examples include:

- Hypervisor – libvirt provides an open interface to retrieve statistical metrics for the utilization rates of domains, vCPUs, memory, block devices, and network interfaces.
- Virtual Switching – sFlow provides an open interface to counters and flow telemetry.
- Hardware Capacity Monitoring – Intel RDT provides metrics on cache utilization and memory bandwidth utilization. Intel RDT can be used to detect "noisy neighbor" VNFs allowing corrective actions to be initiated.
- Hardware Resiliency Event Monitoring – Hardware events are monitored using standard Linux kernel mechanisms and reporting methods, which includes reporting to syslog. The collector agents (collectd and Snap) can detect, count, and report the hardware errors (such as those detected/corrected by Intel Run Sure technologies) reported to syslog.

Open Collectors

The collection layer provides a common focal point for reporting information provided by the resource telemetry interfaces. The most commonly used open industry standard collectors are collectd and Snap. Both collectd and Snap provide resource telemetry collection and fault collection.

A collection layer simplifies the interfacing of resource telemetry by providing a common open local interface, effectively providing a common bus for all faults and telemetry.

Platform collectors such as collectd and Snap, use a plugin-style architecture to gather telemetry and faults from platform resources. Intel has contributed new plugins to collect faults and metrics for collectd and Snap, which, together with active community contributions, provide a rich set of telemetry made available through OPNFV and specific related projects. This telemetry and fault data includes hardware, software, hypervisor metrics, container metrics, and virtual switching performance. Collectd and/or Snap, in turn provides the counters and event data to a local SNMP agent and any other open APIs that must report or display the information. The counter sets provided include the following:

- Data Plane Development Kit (DPDK) metrics and events
- Virtual switching metrics and events
- Accelerator metrics and events
- Cache utilization metrics and events
- Hypervisor metrics and events
- Container metrics and events
- RAS metrics and events

Open Standard Presentation

The presentation layer provides the various open-standard interfaces required to expose the platform performance or fault metrics to other layers for use. The interfaces, including provisioning, telemetry and fault interfaces, are available for consumption by management, analytics, and SDN systems. The presentation layer is multi-generational and extensible, in the sense that interfaces can be exposed to interoperate with existing management systems and NFV Management and Orchestration (MANO) at the same time. Through the open standard presentation layer, platform service assurance functions and data may interoperate with existing enterprise and telecommunications FCAPS systems, using industry standard open APIs, such as SNMP.

Interfaces in the presentation layer include current industry-standard command line interfaces (CLIs), syslog, local debug port, Secure Shell (SSH), SNMP, syslog, and MANO supported by OpenStack* interfaces including ceilometer and Enhanced Platform Awareness (EPA). RESTful APIs can get information from the common open collection layer tools, such as collectd and Snap. Additional networking telemetry is provided by sFlow and IPFIX, which provide interface metrics and flow telemetry.

Integration into Northbound Management Systems

By providing a presentation layer supporting open industry standard interfaces, the platform provides a consistent set of rich platform telemetry to a wide range of management and analytics systems. The platform can interwork with current-generation management systems based on SNMP and analytics systems based on SNMP, syslog, IPFIX, and sFlow. The platform interworks with MANO layers by providing OpenStack* interfaces made through ceilometer, gnocchi or Aodh. In addition to the current list of open interfaces, future RESTful interfaces can be built on top of the open collector layer and provide a consistent set of telemetry to future interfaces under development.

Coming Together: An Example

There is a lot to take in here: lots of functionality, lots of roles that need to be filled, and lots of layers between physical systems, virtual services and the management infrastructure that runs the network and the business. It often helps to examine an example to see how the pieces come together. Let's consider the "noisy neighbor" example mentioned previously. In this scenario, multiple VMs are residing on a common system. In this example, assume that one VM is a critical video streaming application, the second VM hosts a less important workload and has a spiky traffic profile. In an undifferentiated architecture, these two neighbors would be sharing resources with no prioritization and limited ability to detect contention when resources became scarce due to peaks in demand. In an environment configured to enable enhanced service assurance for the virtual network, much more is possible.

In this environment, Intel RDT is able to detect contention when the two workloads are utilizing the last level cache, with the critical video streaming application suffering as a result of the spiky neighbor utilization pattern—causing jitter and degraded performance. Intel RDT monitors this utilization contention and a plug-in allows the information to flow into the collection layer via collectd. From there, thresholds allow the triggering of SNMP calls into the network management system and the administrator can decide to take appropriate action—moving one of the workloads, granularly assigning resources based on priority, or some other response.

Note that this is just a single example. Similar flows can be drawn using any number of different technologies. For example, recurring memory errors detected by Intel Run Sure technologies could trigger an event through its plug-in into collectd. Collectd could similarly trigger an SNMP message into the network management infrastructure to drive either a manual response or perhaps some pre-programmed response (e.g. migrate a critical app or service onto a host that is not experiencing memory issues).

There are two important points to extend here. The first key point is that this capability is not restricted to legacy CLI, SNMP, network management and collectd environments. For a more greenfield or cloud-centric environment, a very similar flow from the platform (Intel RDT detecting and reporting on cache utilization contention) through Barometer* plug-ins to OpenStack via Gnocchi or Aodh tools. Within OpenStack, the administrator can be notified and take the appropriate action. And in fact, it is entirely possible that these two models can co-exist in a hybrid model.

This leads to the final key point: that this basic plumbing can quickly evolve into more robust automation. In the flow cited above, the administrator is left making some choice and likely taking an action. But when the resources are identifiable and configurable, it is a small step to envision policy-driven actions as a response to the warnings and events detected within the network. Perhaps the commands to move a competing workload from a host or to prioritize resources for a critical workload could be entirely automated. Intel and partners have been showcasing just such a scenario using the Intel RDT/noisy neighbor example above in OpenStack environments—using Intel RDT and collectd plug-ins for threshold violation monitoring to pass data to the OpenStack Vitrage. Vitrage is triggered and an alarm is passed to the Mistral workflow service to trigger a response—to set Intel RDT controls to allocate cache priority to the more critical application. In this demonstration, the result is a predictable video stream.

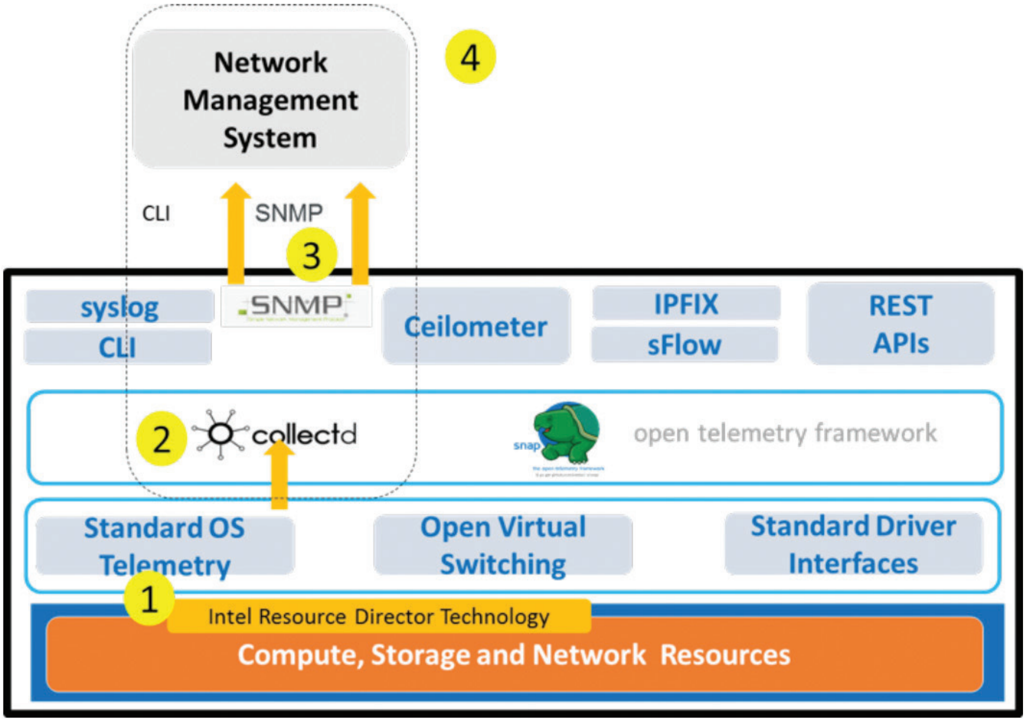

Figure 5. Example data flows for the "Noisy Neighbor" scenario in an enabled environment

Conclusion

Intel is uniquely positioned to bring the worlds of connectivity, computing, and cloud together to provide the infrastructure for driving enhanced IT efficiency and to unlock business agility with the adoption of NFV and SDN. Intel's newest technology platforms are delivering the processing, storage, virtualization, security, I/O, acceleration, management, and analytics capabilities that will power these high performance, efficient, scalable, agile, and increasingly automated networks. These capabilities are largely provided by Intel Infrastructure Management technologies and enabled for service assurance usages through plug-in capabilities available for popular collector tools.

Networks evolve through a combination of legacy, virtual (NFV) and SDN technologies, the NFV/SDN "hybrid" network of traditional fixed-function network devices and new NFV solutions. It is an imperative that these technologies are integrated and managed with tools, processes, and techniques that operators use for service assurance today, interoperating towards existing FCAPS management systems while enabling the path to the truly automated networks of the future.

Rich telemetry exposed by the platform through standard open interfaces can feed management and analytics systems including machine learning systems. Machine learning can provide deeper insights into telemetry datasets and potentially provide greater system reliability and efficiencies, as the learning systems model and adapt based on telemetry.

Intel is leading the development of key standards and partnering with industry-leading service and equipment providers through the Network, Storage, and Cloud "Builders" programs to foster the development of more efficient, powerful, flexible, interoperable, and automated IT infrastructure. These collaborations with a dedicated focus on enabling integration with traditional service assurance tools and techniques--promise to deliver solutions that provide IT with the tools to adopt the open platforms that support business innovation at lower costs, faster time to deployment, and greater ease and confidence. For more information on Intel Builders program participation, go to: https://builders.intel.com.

References

#	TITLE	LINK
1	ETSI Network Function Virtualization (NFV) Architectural Framework (ETSI GS NFV 002)	http://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/002/02.01.01_60/gs_NFV-IFA002v020101p.pdf

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: http://www.intel.com/design/literature.htm

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at http://www.intel.com/ or from the OEM or retailer.

Intel, Intel QuickPath, Intel Scalable Memory Interconnection, Intel SMI, Intel Resource Director Technology, Intel RDT, Intel Run Sure Technology, Intel Infrastructure Management Technologies, Intel Trusted Execution Technology, Intel Rapid Storage Technology, Intel Node Manager, Intel Management Engine and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others. Copyright© 2017, Intel Corporation. All rights reserved.

© 2017 Intel Corporation Printed in USA 0717/HM/DJA/PDF ♲ Please Recycle 336200-001US