Network State Collection, Correlation and Analytics Product / RFI Requirements

Version 1.1

A white paper from the
ONUG Network State Collection,
Correlation and Analytics
Working Group

May, 2015

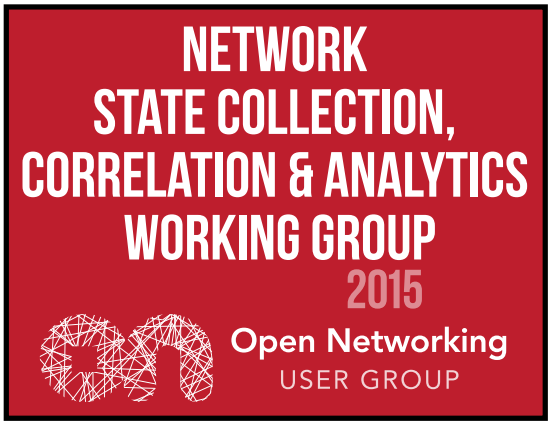

Definition of Open Networking

Open networking is a suite of interoperable software and/or hardware that delivers choice and design options to IT business leaders, service and cloud providers. At its core, open networking is the separation or decoupling of specialized network hardware and software – all in an effort to give IT architects options in the way in which they choose to design, provision, and manage their networks. These technologies must be based on industry standards. The standards can be de-facto as adopted by a large consortium of the vendor community, open in the sense that they are community based, or defined as standards by the prevailing standards bodies. Open networking hopes to deliver on two promises:

1) Decoupling of network hardware and software which mitigates vendor lock-in and shifts network architecture structure options to users
2) Significant reduction of the total cost of ownership model, especially operational expense

Scope

The scope of this document is to provide a set of tactical and strategic requirements that may help guide enterprises in their selection of Software-Defined Networking (SDN) – Network State Collection, Correlation and Analytics (NSCCA) vendor solutions. While the impacts discussed are commensurate with an Information Technology Infrastructure Library (ITIL) service delivery model, enterprises can leverage the information for a Request for Information (RFI) and adapt to scale and suit their current or planned organizational support delivery and maturity capabilities.

Out of Scope

The working group focused its efforts on ***what*** is needed to deliver an open network state collection, correlation and analytics service to the enterprise market. The working group did not focus and does not offer the ***how***; that is, there is no specific protocol(s) or Application Program Interface(s)(API) specified to deliver on the NSCCA service. The working group leaves that work to the vendor and standards communities. Open Networking User Group (ONUG) strongly encourages open interfaces and protocols in the construction of a multivendor interoperable NSCCA service to deliver the greatest value and choice to enterprise IT executives. Further, the working group does not specify ***where*** various NSCCA functions should reside–that is, within physical or vertical products, as modules within existing products or standalone products. That is up to the vendor community.

Executive Summary

This document defines an open architectural framework and common set of functional solution requirements for one of the open networking use cases–Network State Collection, Correlation and Analytics (NSCCA)–identified by the ONUG community. The content of this document is intended as general guidelines for:

- IT enterprise end users to compare vendor solutions and develop formal RFI specifications,
- IT vendors to develop and align product requirements,
- Standards organizations to align their initiatives and efforts to deliver an open approach to network state collection, correlation and analytics.

Defining a common set of solution requirements aligns with ONUG's goal to drive the IT vendor and standards communities to deliver open interoperability networking solutions to provide IT end users maximum choice and flexibility in deploying open networking solutions. The expectation is that this document provides a common baseline, covering a set of enterprise deployment requirements for Network State Collection, Correlation and Analytics solutions. The assumption is being made that this set of requirements will be completed by enterprise-specific requirements to meet specific deployment needs. As the working group developed the NSCCA use case, it found that there are, in fact, multiple use cases based upon the NSCCA open architectural framework, as discussed below. This version of the document provides but a few use cases in what could be the basis for a new computer network industry

NETWORK STATE COLLECTION, CORRELATION & ANALYTICS WORKING GROUP 2015

Open Networking USER GROUP

Open Networking User Group (ONUG)

ONUG is one of the largest industry user groups in the networking and storage sectors. Its board is made up exclusively of IT business leaders, with representation from Fidelity Investments, FedEx, Bank of America, UBS, Cigna, Pfizer, JPMorgan Chase, Citigroup, Credit Suisse, Gap, Inc., and Symantec. The ONUG mission is to guide and accelerate the adoption of open networking solutions that meet user requirements as defined through use cases, proof of concepts, hackathons, and deployment examples to ensure open networking promises are kept.

The ONUG community is led by IT business leaders and aims to drive industry dialogue to set the technology direction and agenda with vendors. To that end, ONUG hosts two major conferences per year where use cases are defined and members vote to establish a prioritized list of early adopter, open networking projects that communicate propensity to buy and budget development. The vendor community stages proof of concepts based upon ONUG Use Cases, while standards and open source organizations prioritize their initiatives and investments based upon them. ONUG also hosts user summits and smaller, regional user-focused Fireside Chat Meet-Ups through the year. ONUG defines six architectural areas that will open the networking industry and deliver choice and design options. To enable an open networking ecosystem, a common multivendor approach is necessary for the following six architecture components::

1) Device discovery, provisioning, and asset registration for physical and virtual devices
2) Automated "no hands on keyboards" configuration and change management tools that align DevOps and NetOps
3) A common controller and control protocol for both physical and virtual devices
4) A baseline policy manager that communicates to the common controller for enforcement
5) A mechanism for sharing (communicating or consuming) network state and a unified network state database that collects, at a minimum, MAC and IP address forwarding tables automatically
6) Integrated monitoring of overlays and underlays

subcategory. Finally, the expectation is that the scope of requirements defined in this document will evolve, and hence, the versioning of this document.

The expected outcome for Network State Collection, Correlation and Analytics enterprise adoption and usage can be summarized, but is not limited to meeting this set of nine architecture requirements and nine use cases:

ONUG Network State Collection, Correlation and Analytics Top Nine Requirements

1. A Network State Collection function to support collecting state data from SDN Controllers and Data Plane Elements with a scalability of at least 1,000 devices and/or endpoints.
2. A Network State Correlation function to support state information aggregation, formatting, storage and correlation for a wide range of SDN Controllers and Data Plane Elements that can scale to tens of thousands of metrics support.
3. A Network State Correlation function to support state information aggregation, formatting, storage and correlation for a wide range of compute and storage devices, services and software modules.
4. A network state database to store data that is collected and accessible across the enterprise via network analytics function.
5. A Network State Analytics function to support big-data analytics and algorithms.
6. A Network State Analytics function to integrate with orchestration systems to take action on the analytics result.
7. A Network State "Pretty Amazing Stuff" (PAS) function to support Network State Collection, Correlation and Analytics extensibility through additional information and open interfaces.
8. A set of Open southbound and northbound interfaces for state information transfer between Network State Collection, Correlation and Analytics functions to support state information transfer.
9. An Open API or protocol between Network State Correlation and "Pretty Amazing Stuff"(PAS) Function to support additional information into Network State Analytic queries.

ONUG Network State Collection, Correlation and Analytics Top Nine Use Cases

1. Query Optimal Service Function or Workload Location,
2. Query Assessment of Service Function or Workload Modification,
3. Query Real-Time Configuration Analytics,
4. Query Real-Time Topology Analytics,
5. Query Real-Time Performance Analytics,
6. Query Predictive Performance Analytics
7. Query Real-Time Fault Analytics,
8. Query Real-Time Security Analytics,
9. Measure SLA Compliance.

Conventions

The management plane consists of all functions needed to configure, monitor and troubleshooting of virtual network overlays, including overlay endpoints and end-to-end connectivity.

The following conventions are used throughout this document. The requirements that apply to the functionality of this document are specified using the following conventions. Items that are REQUIRED (contain the words MUST or MUST NOT) will be labeled as [Rx]. Items that are RECOMMENDED (contain the words SHOULD or SHOULD NOT) will be labeled as [Dy]. Items that are OPTIONAL (contain the words MAY or OPTIONAL) will be labeled as [Oz]. In addition, a priority value of High (H), Medium (M) or Low (L) may be assigned to each item.

The priority will be labeled as [RHx], [DHy] or [OHz] for High priority, [RMx], [DMy] or [OMz] for Medium priority or [RLx], [DLy] or [OLz] for Low priority. The integer values {x, y, z} shall be unique across the document but are not required to be unique across the 3-tuple set {x, y, z}. For example, RM10 and DM10 are allowed whereas RM10 and RL10 are prohibited. Requirements in this document are numbered using increments of 10.

Where needed, related sub-requirements are numbered using increments of 1. The key words "MUST," "MUST NOT," "REQUIRED," "SHALL," "SHALL NOT," "SHOULD," "SHOULD NOT," "RECOMMENDED," "MAY," and "OPTIONAL" in this document are to be interpreted as described in Request for Comments (RFC) 2119. All key words use upper case, bold text to distinguish them from other uses of the words. Any use of these key words (e.g., may and optional) without [Rx], [Dy] or [Oz] is not normative. The priority assignments are defined as follows.

- **High (H):** Functionality that must be supported at day one and is critical for baseline deployment.
- **Medium (M):** Functionality that must be supported, but is not mandatory for initial baseline deployment.
- **Low (L):** Desired functionality, which should be supported, but can be phased in as part of longer-term solution evolution.

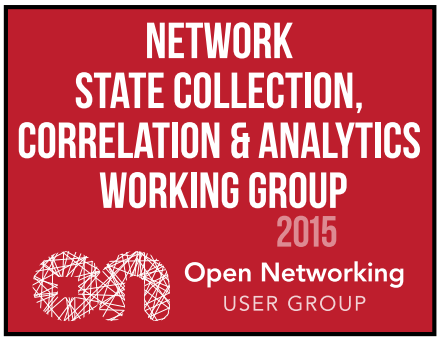

1.0 Problem Statement

Enterprise marketplaces are complex landscapes with a large portfolio of applications, facing potential broad disruption when there are any network changes. Understanding the disruption is complicated, especially with little visibility of how entire networks can be affected by a single device state change. Lack of predictive analytical models that indicate specific network problems that can occur following changes in the state of devices in the network add to the complexity.

2.0 ONUG NSCCA Working Group Open Architecture Framework

The ONUG NSCCA working group has defined "**state**" to be a combination of the following attributes:

Topology: It comprises flow tables, Border Gateway Protocol (BGP) routes, switch Equal-Cost Multipath (ECMP) hash tables, Link Layer Discovery Protocol (LLDP) neighbors, and routing neighbors, and what is behind an IP address where a single IP address hides multiple applications, containers and virtual machines (VMs).

Flows: This includes flow and flowlet tracking (Netflow/IP Flow Information Export (IPFIX)), firewall rule hits (allows/denies), load balancing (number of connections per VIP, per real server, by data type) and Quality of Service (QoS) attributes.

Stats: It is made up of interface statistics (e.g., Simple Network Management Protocol (SNMP)), utilization, packet errors, reachability and latency.

Group-Based Policy/Service Function Chaining (GBP/SFC): This encompasses the physical implementation of policy, which defines how an application works, and how endpoint groups and contracts flow through intermediate functions.

Capacity: This defines available or provisioned versus unavailable bandwidth, utilized versus underutilized processors, memory, Transmission Control Protocol (TCP) sockets and physical ports.

Change Verification Functions (CVF): This relates to status of requested changes, if requested changes were applied as directed, and unit sanity checking after a change is introduced.

S**ecurity Posture**: Security posture state may include information (AAA)/802.1x to determine if a device or user is authenticated or infected, if it is a mission-critical device, an operating system version, Media Access Control Security (MACSec) encryption or other security attributes.

Device Identification: State includes device identification information, such as MAC, IP, Domain Name System (DNS) name, user-agent, mobile device unique ID, and authentication token/certificate. Much of the above state information is available in one form or another in multiple network devices and endpoints distributed through an enterprise network. However, it's either not made available by vendors and/or there is no architecture to collect, aggregate, correlate and then analyze this information.

To address this problem, the ONUG NSCCA working group provides an open architecture framework to deliver an open NSCCA service as illustrated in Figure 1. This architecture comprises four **functional** components.

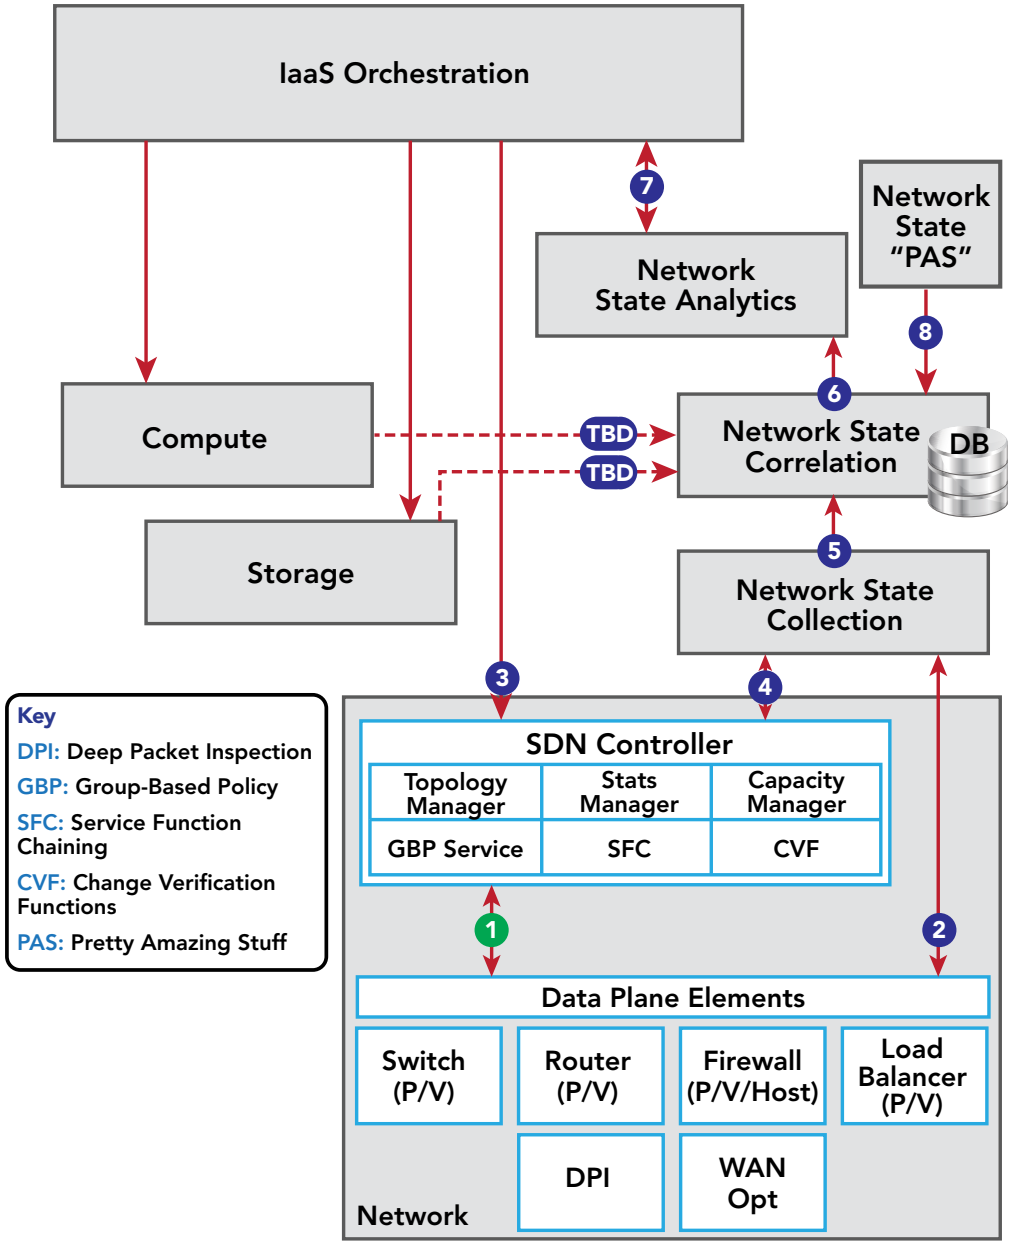

Figure 1.
ONUG Network State Collection, Correlation and Analytics Architecture Framework

NETWORK STATE COLLECTION, CORRELATION & ANALYTICS WORKING GROUP
2015
Open Networking USER GROUP

1) **Network State Collection:** A function to collect network state information from various sources, including network devices such as controllers, switches, routers, etc., and network services (service functions) such as firewalls, load balancers, deep packet inspectors and workload applications, etc.
2) **Network State Correlation:** Network state information collected by the Network State Collection function is sent to a Network State Correlation function for aggregation, formatting and storage. To assure a wide view of how network state changes impact applications and IT service delivery, compute, storage and PAS inputs are defined.
3) **Network State Analytics:** While the Network State Correlation function aggregates, formats and stores state information, the Network State Analytics function provides computational processing, transforming state data into usable use case applications, such as new service functional placement, service function moves, network design "what if" analytics and much more.
4) **Network State "Pretty Amazing Stuff" (PAS):** PAS is any external source that would provide valuable information to be correlated with state. PAS may be an authorization service that provides device security posture information or an application server that is providing feedback to the network about application performance from a user perspective. PAS provides a way to inject innovation into the NSCCA ecosystem.

The above ONUG NSCCA open architecture provides a framework to collect, format, aggregate, correlate and store network, compute, storage and other state information for processing via an analytics engine to deliver network-wide predictive network design assist service and trouble identification.

3.0 Use Cases

Figure 2 illustrates a few examples of use cases or applications that the analytics engine may provide, equipped with a wide view of correlated and analyzed network state information.

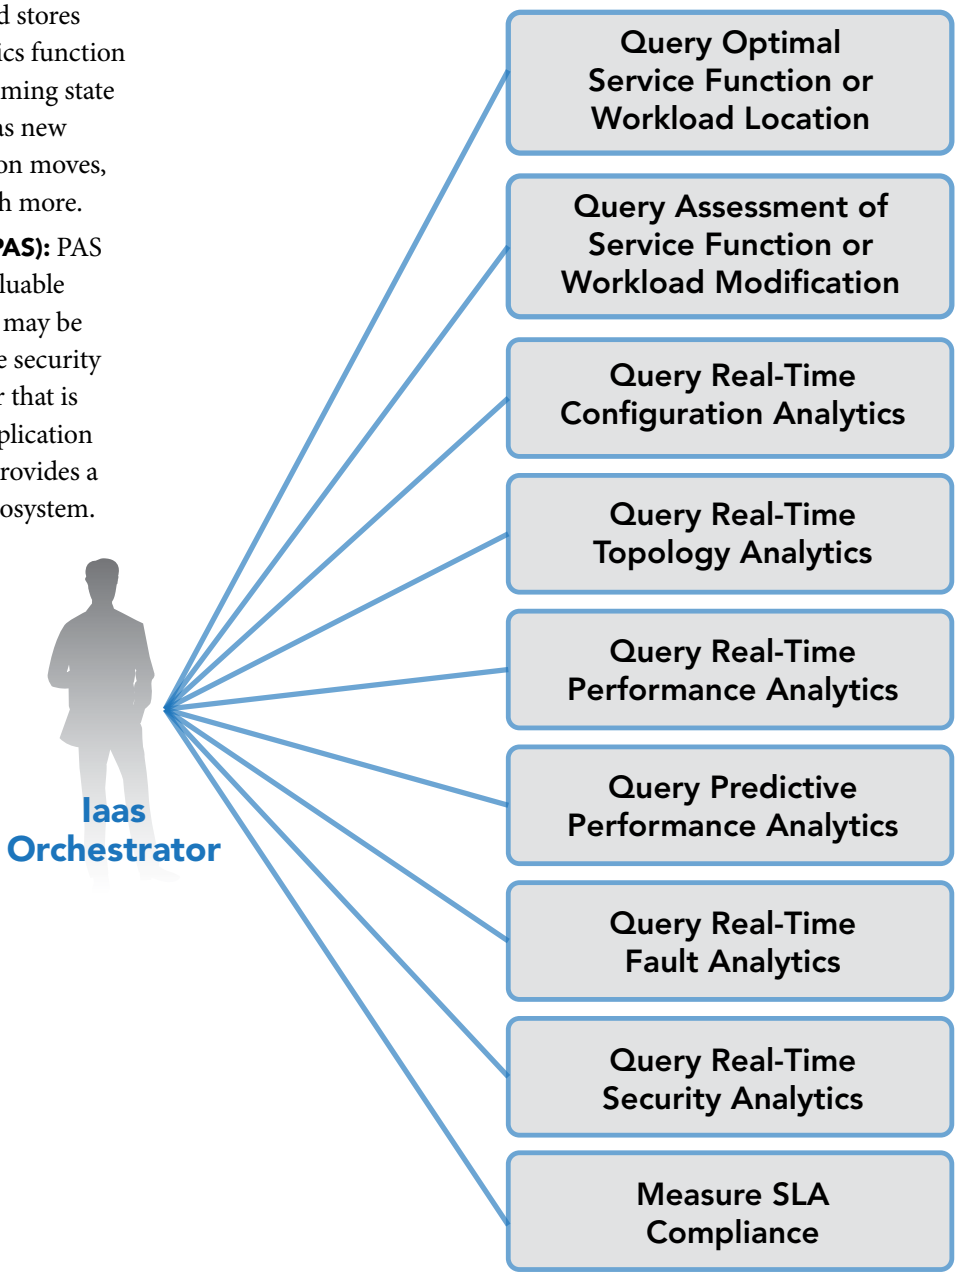

Figure 2.
ONUG NSCCA Use Case Diagram

NETWORK STATE COLLECTION, CORRELATION & ANALYTICS WORKING GROUP
2015
Open Networking USER GROUP

"What if" Analysis: Another use case is to provide a series of "what if" predictive analytics, such as what may be the impact of high load or loss on a network link or network node failure.

3.1 Query Optimal Service Function or Workload Location

Field	Description
Use Case Number	1
Use Case Name	Query Optimal Service Function or Workload Location
Description	In this Use Case, the Infrastructure as a Service (IaaS) Orchestrator queries the Network State Analytics function for the physical and virtual proximity of all elements in a service chain in order to effectively reduce east-west traffic flows before a new service function is deployed within the network. The Network State Ana-lytics function will provide the optimal placement of a new service function or workload. Current and projected network load information will be pro-vided to predict the impact of a new service function or workload being implemented. In addition to new ser-vice function placement and changes to a workload placement, optimiza-tion applications provides assistance to IT architecture/design teams and/or automated IT service delivery services such as cloud or VM or-chestration tools.
Actors	IaaS Orchestrator
Pre-Conditions	The service must be identified and provisioned. The network resources participating in the service chain must be deployed, provisioned and activated. The requested service function must be identified and ready for deployment.
Post-Conditions	The IaaS Orchestrator has all the information necessary for Service Function placement.
Alternative Paths	TBD
Assumptions	The service and resources participating in the service and requested service function must be known and identifiable. Network connectivity to the resources must be established.
Reference	TBD

3.2 Query Assessment of Service Function or Workload Modification

Field	Description
Use Case Number	2
Use Case Name	Query Assessment of Service Function or Workload Modification
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function to predict the impact of a service function or workload move or a change in demand condition. This use case would provide information so that operations groups may be proactive to changes or moves of a service function in an effort to optimize the virtual topology. Predictive information due to changes in the service chair or chaining load conditions of the network upon the service function move or change in demand conditions is highly desired.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.3 Query Real-Time Configuration Analytics

Field	Description
Use Case Number	3
Use Case Name	Query Real-Time Configuration Analytics
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function for real-time network configuration state information. Configuration state information applies to network devices, service functions, workloads, etc.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.4 Query Real-Time Topology Analytics

Field	Description
Use Case Number	4
Use Case Name	Query Real-Time Topology Analytics
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function for real-time network topology information. Network topology information includes information, such as an aggregate number of network devices, service functions, workloads and links. This also might include capacity, oversubscription, and hardware/software lifecycle stage, etc. Another example might be an aggregated number or score that reflects the complexity of a network in terms of number of devices, links, capacity, oversubscription, security posture, hardware/software lifecycle stage, etc.

Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.5 Query Real-Time Performance Analytics

Field	Description
Use Case Number	5
Use Case Name	Query Real-Time Performance Analytics
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function for real-time performance information. Performance information includes quality of service and quality of experience metrics for applications, network devices, service functions and workloads. The Network State Analytics function provides real-time troubleshooting information so as to correlate a change in application performance due to a change in network state.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.6 Query Predictive Performance Analytics

Field	Description
Use Case Number	6
Use Case Name	Query Predictive Performance Analytics
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function for predictive performance information. This use case provides IT teams with a simulation of application performance laid upon a grid of network state to predict the likely performance of an application before it is deployed.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.7 Query Real-Time Fault Analytics

Field	Description
Use Case Number	7
Use Case Name	Query Real-Time Fault Analytics
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function for real-time fault and diagnostic information. Fault and diagnostic information includes alarm, error and diagnostic test results for applications, network devices, service functions and workloads. An example might include a map and graphic representation of the overall health of a network including dynamic correlation of traffic bursts, security events, etc., which may feed into an automatic action to be taken defined by policy.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.8 Query Real-Time Security Analytics

Field	Description
Use Case Number	8
Use Case Name	Query Real-Time Security Analytics
Description	In this Use Case, the IaaS Orchestrator queries the Network State Analytics function for real-time security information. Security information includes security posture, security events for applications, network devices, service functions and workloads.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

3.9 Measure SLA Compliance

Field	Description
Use Case Number	9
Use Case Name	Measure SLA Compliance
Description	In this Use Case, the IaaS Orchestrator measures the Service Level Agreement (SLA) conformance/compliance via the Network State Analytics function. SLA measurements may be performed on the basis of a per customer/business unit, service function, workload, bandwidth, latency, etc.
Actors	IaaS Orchestrator
Pre-Conditions	
Post-Conditions	
Alternative Paths	TBD
Assumptions	
Reference	TBD

4.0 Requirements

4.1 Network State Collection

The Network State Collection function is expected to provide state collection across a heterogeneous set of virtual and physical network devices and endpoints, which will evolve over time. Network State Collection is to exploit ratified standards for "situational awareness" protocols or APIs to capture contextual awareness of network devices. Open interfaces between network state collection and device scenarios include:

SDN Controller(s): An open interface between Network State Collection and SDN controller(s) is required to extract topology, stats manager, GBP service, service function chaining, capacity management and change verification function information. This is labeled in Figure 1 as interface 4.

Data Plane Elements: An open interface is required for Network State Collection such that data plane elements can be accessed directly or indirectly via SDN or Network Service Virtualization (NSV) and/or network function controllers. These elements include both physical and virtual switches, routers, firewalls, load balancers, deep packet inspection and wide area optimization devices and network services. This is labeled in Figure 1 as interface 2.

Network State Collection to Network State Correlation Interface: An open interface between Network State Collection and Network State Correlation is required to facilitate data transfer. This is labeled in Figure 1 as interface 5.

4.1.1 Functional Requirements

The following general requirements are defined for the functionality of the Network State Collection:

RH-10 The Network State Collection function **MUST** provide compute processing and storage available to collect state data from 1,000 network devices and/or endpoints.

RM-20 The Network State Collection function **MUST** provide the ability to record information at different precisions for different use cases.

For example, during real-time debugging, there should be the ability to have fine-grained metrics collected every 5 seconds to 1 min., whereas during normal operation recording may be collected only at 5 min. intervals.

RM-21 The Network State Collection function **MUST** provide the ability for a user configurable collection time interval.

RM-30 The Network State Collection function **MUST** provide a mechanism to notify subscribers of network state changes on selectable criteria, such as per flow or application basis.

For example, a Network Analytics Application should be able to register for an Application's Network State change for events like configuration update or availability change. These can be used for root cause and correlation along with performance and other metrics.

4.1.2 Interface Requirements

The following general requirements are defined for communication flow of the Network State Collection (interface 2, 4 and 5):

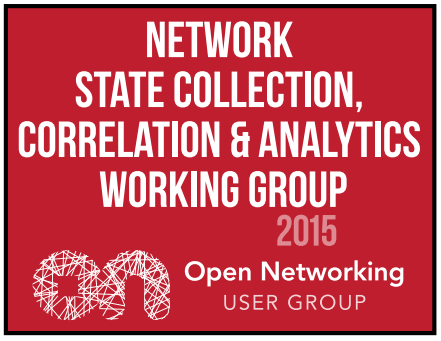

RH-40 The Network State Collection function **MUST** support an open interface and protocol to the SDN Controller function, labeled as interface 4.

RH-50 Interface 4 **MUST** be robust enough to support the extraction of topology, stats manager, GBP service, service function chaining, capacity management and virtualized functions information and metrics listed in Appendix A.

RH-60 Interface 4 **MUST** support both push and pull capabilities for data exchange with the SDN Controller function.

RH-70 Interface 4 **MUST** support a response time to either a push or pull command of less than 50 ms for a network of 1,000 network devices.

The following provides a prioritized list of state data to be extracted from various software managers/modules within an SDN controller(s):

RH-80 Interface 4 **MUST** support collection of data associated with the SDN Controller Topology Manager function.

RL-90 Interface 4 **MUST** support collection of data associated with the SDN Controller GBP Service function.

RH-100 Interface 4 **MUST** support collection of data associated with the SDN Controller Stats Manager function.

RM-110 Interface 4 **MUST** support collection of data associated with the SDN Controller SFC function.

RM-120 Interface 4 **MUST** support collection of data associated with the SDN Controller Capacity Manager function.

RM-130 Interface 4 **MUST** support collection of data associated with the SDN Controller CVF function.

RM-140 Interface 4 **MUST** support collection of data associated with the SDN Controller NSV and Network Services Function (NFV) function.

The following provides a prioritized list of state data to be extracted from various software managers/modules within an SDN controller(s):

RH-150 The Network State Collection function **MUST** support an open interface and protocol to the Data Plane Elements, labeled as interface 2.

Data Plane elements provide an open interface either directly as (interface 2) or via SDN controller (interface 4).

RH-160 Interface 2 **MUST** support a push or pull command option to push and pull read-only network state data from a single Data Plane Element to Network State Collector function.

RH-170 Interface 2 **MUST** support a response time to either a push or pull command of less than 50 ms for a network of 1,000 network devices.

The following provides a prioritized list of state data to be extracted from various data plane physical and virtual devices:

RH-180 Interface 2 **MUST** support collection of read-only network state data from a Switch (p/v) Data Plane Element.

RH-190 Interface 2 **MUST** support collection of read-only network state data from a Router (p/v) Data Plane Element.

RM-200 Interface 2 **MUST** support collection of read-only network state data from a Firewall (p/v/host) Data Plane Element.

RH-210 Interface 2 **MUST** support collection of read-only network state data from a Load Balancer (p/v) Data Plane Element.

RL-220 Interface 2 **MUST** support collection of read-only network state data from a Deep Packet Inspection (DPI) Data Plane Element.

RM-230 Interface 2 **MUST** support collection of read-only network state data from a Wide Area Network (WAN) Optimization (p/v) Data Plane Element.

The following general requirements are defined for communication flow between Network State Collection and State Correlation (interface 5):

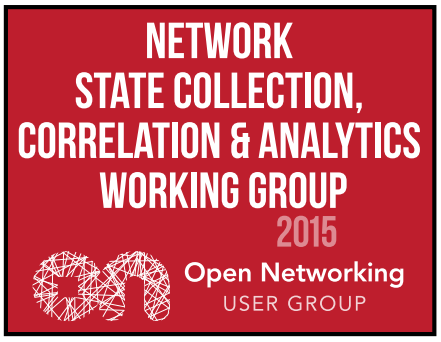

RH-240 The Network State Collection function **MUST** support an open, multi-vendor interface and protocol for information exchange with the Network State Correlation function, labeled as interface 5.

The protocol/interface needs to be robust enough to support the extraction of collected state information from SDN controllers and data plane elements and in the future, compute plus storage elements.

RH-250 Interface 5 MUST support a push or pull command option to push and pull network state data from the Network State Collector function to the Network State Correlation function.

RH-260 Interface 5 MUST support a user configurable response time to either a push or pull command, based on a time interval.

4.2 Network State Correlation

The Network State Correlation function and database are expected to aggregate and correlate state information from a wide range of systems, including compute, storage, PAS and network state collection. For this version of the ONUG NSCCA working group, the focus is on network state collection plus the introduction of PAS. From a networking point of view, think of Network State Correlation and database as NetFlow, sFlow and SNMP rolled into one. The Network State Correlation and database function provides compute processing and storage to format streams of state information for entry into a network state database. It is both a state depository as well as state correlation between layers. The state correlation function correlates state information between network layers–that is, changes in one layer and the impact on other layers is a key functional requirement. The Network State Correlation and database should capture and model network topology, interdependencies, service chaining, etc., providing state data to support "what if" analytics and more.

4.2.1 Functional Requirements

The following general requirements are defined for the functionality of the Network State Correlation function:

RH-270 The Network State Correlation function **MUST** support an openinterface using standard index and metadata format to normalize and aggregate multiple heterogeneous streams of state information.

RH-280 The Network State Correlation function **MUST** support a single flow to generate multiple records from different Network State Collection functions.

DH-10 The Network State Correlation function **SHOULD** "link" these records and keep this relationship such that it can be queried for further deep analysis.

DH-290 The Network State Correlation function **MUST** support the ability to aggregate and archive metadata for turnkey "big data" analysis

RH-300 The Network State Correlation function **MUST** support the ability to index large number of metrics (tens of thousands) across different network components.

RH-310 The Network State Correlation function **MUST** support a data-aging function to provide summarization functions such that the overall state of the network state database does not crash under the weight of endless amounts of data.

DH-320 The Network State Correlation function **SHOULD** support historical snapshots of the state database minimized to one week with a user selectable frequency.

4.3 Network State Analytics

The Network State Analytics function is a big data algorithm and structure capable of processing millions of rows of network state data. The Network State Analytics function turns a vast array of state data into application use cases as described above. Output of the Network State Analytics function, such as "new service function placement," will be fed to an Infrastructure as a Service (IaaS) orchestrator for new service function configuration. That is, a set of use case applications which will directly interface to an IaaS orchestrator while others will be standalone, such as "Network Health" or "What If Analysis" use case applications.

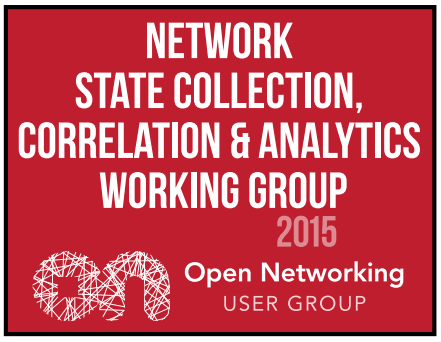

4.3.1 Functional Requirements

The following general requirements are defined for the functionality of the Network State Analytics function:

RH-370 The Network State Analytics function **MUST** support the ability for subscribers to request filtered events for a flow, application, VM or Host to allow focused analytics rather than requiring all applications to listen on all network state events.

RH-380 The Network State Analytics function **MUST** provide both a batch analytics capability that does off-line/non-real-time analysis of large datasets as well as a streaming analytics capability that can perform near-real-time analysis on a data stream.

DH-381 The Network State Analytics function in need of historic correlation function data **SHOULD** do its own archival backed by sufficient storage.

4.3.2 Interface Requirements

The following general requirements are defined for communication flow of theNetwork State Analytics function (interface 6 and 7):

RH-390 The Network State Analytics function **MUST** support an open, multi-vendor interface and protocol for information (e.g., network state data, event notifications, etc.) exchange with the Network State Correlation function, labeled as interface 6.

RH-400 The Network State Analytics function MUST support an open, multi-vendor interface and protocol for information (e.g., network state data, event notifications, etc.) exchange with the IaaS Orchestrator, labeled as interface 7.

4.4 Network State PAS

The Network State PAS is an external source or input into the Network State Correlation function that provides valuable information. A PAS may be an authorization service, such as an active directory, to determine if a device or user has been authenticated or an application server providing feedback to the network about application performance from a user perspective. PAS is an innovation injection point and network state ecosystem enabler.

4.4.1 Functional Requirements

The following general requirement is defined for the functionality of the PAS:

Functional requirements for PAS will be addressed in a future version of this document.

4.4.2 Interface Requirements

The following general requirements are defined for communication flow of the PAS (interface 8):

RH-360 The PAS function MUST support an open, multi-vendor interface and protocol for information exchange with southbound applications, such as the Network State Correlation function, labeled as interface 8.

5.0 Recommendations for Standards

For completeness, the specific areas in Network State Collection, Correlation and Analytics solutions, where there is a need for open well-defined and vendor-agreed-on technology standards, can be summarized as follows:

Network State Collection to SDN Controller and Data Plane Elements Interface: open southbound interface for state collection from SDN Controllers plus network devices and services.

Network State Collection to Network State Correlation Interface: open northbound interface for the aggregation, indexing and storage of state information into a state correlation database.

Network State Correlation to Network State PAS Interface: An open east-west interface for the exchange of PAS information to enable authentication, application performance, etc., to be added to the state correlation database.

Network State Correlation to Network State Analytics Interface: open northbound interface to support state database queries from the Network State Analytics function.

Network State Analytics to IaaS Orchestrator Interface: open northbound interface to support queries based on the identified use cases from the Iaas Orchestrator.

References

[ONUG web site] - ONUG white paper, "Open Networking Challenges and Opportunities."

[Open Virtual Switch] – Apache project Open vSwitch.

[draft-mahalingam-dutt-dcops-vxlan] - IETF Informational RFC, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks."

[draft-sridharan-virtualization-nvgre] - IETF Informational RFC, "NVGRE: Network Virtualization using Generic Routing Encapsulation."

[OVSDB] - IETF Informational RFC, "The Open vSwitch Database Management Protocol."

[draft-smith-opflex] - IETF Informational RFC, "OpFlex Control Protocol."

[draft-ietf-l2vpn-evpn-07] - IETF Draft RFC, "BGP MPLS Based Ethernet VPN."

[Networking API v2.0 (CURRENT)] - OpenStack Networking API specification.

[RFC2233] - IETF Draft Standard RFC, "The Interfaces Group MIB."

ONUG Network State Collection, Correlation and Analytics Working Group

Nick Lippis	The Lippis Report	Co-Chairman	Neal Secher	BNY Mellon	Co-Chairman

IT Executives

Vijay Balasubramanian	Bank of America	Member	Kevin Irwin	Fidelity Investments	Member
Brian Hedstrom	datavision	Member	Bryan Larish		Member
Nick Feamster	Princeton University	Author, Member	Jem Pagan	JNK	Member
Michael Githens	ixia	Testing Partner	Sean Wang	UBC The University of British Columbia	Member
Piyush Gupta	JPMorgan Chase & Co.	Member			

ONUG Demonstrators

Satish Grandhi	avni	Member	Gaurav Rastogi	AVI Networks	Member
Ananda Rajagopal	Gigamon	Member	Pascale Vicat-Blanc	f5	Member

NETWORK STATE COLLECTION, CORRELATION & ANALYTICS WORKING GROUP 2015

Open Networking USER GROUP

Appendix A Performance, Error and Resource Metrics

A.1 Network Metrics

L2 / L3 Metrics Link Utilization (Aggregated for Host and VM): Bandwidth utilization measured against maximum bandwidth allowed on the link.

Link Traffic (Bandwidth for Host and VM): Network bandwidth measured as bits per second in both receive and transmit direction

Link Packet Loss on per Link Basis (Host and VM): Network packets dropped in both receive and transmit direction

Link Mean Packet Arrival (Host and VM): Averaged time duration between packets arriving on a link over a period of time.

Link Packets / sec (Host and VM): Averaged rate of packets in receive and transmit direction over a period of time.

Per Flow Available Capacity between Two Endpoints: Network bandwidth available between any two network nodes.

Per Flow Connectivity between Two Endpoints: Available network paths between any two network nodes.

Per Flow Availability (up time of the network connectivity): Network per flow link utilization (compared to against VMs link utilization)

Per Flow Link Latency: Network latency measuring effective bandwidth between two nodes.

Per Flow Bandwidth: Network bandwidth measured as bits per second in both receive and transmit direction.

L4/L5 (Network Counters) Per Flow RTT: Network Round-Trip Time (RTT) experienced by the logical traffic represented as a network flow over a period of time.

Per Flow Duplicate ACK Retransmits: Number of duplicate Acknowledgement (ACK) retransmits indicating network loss for every TCP sessions.

Per Flow SACK Retransmits: Number of Selective Acknowledgement (SACK) retransmits over a period for TCP sessions.

Per Flow Timeouts: Number of TCP timeouts over a period of time.

Per Flow Out of Order Packets: Number of out of order packets in a TCP session. This can be helpful in understanding stability of underlying networking conditions.

Per Flow SYNs: Number of the new TCP connection attempts. It is useful for measuring load and potential DOS attacks.

Per Flow Connection Drops: Number of the bad TCP connections over a period of time.

Per Flow Connections: Number of successful TCP connections over a period of time.

Per Flow Zero Window Size Events: Number of zero window size events that can help in understanding network performance issues in both client as well as backend.

Per Flow Server Flow Control Events: Number of TCP control events on per flow.

A.2 Application Metrics

Per App Errors: Application errors. In case of L7 HTTP application it is number of 4xx and 5xx responses.

Per App Errors Percentage: Measures the error percentage against the total number of requests.

Per App Latency: Averaged client transaction latency computed by adding response latencies across all requests.

Per App Requests: Averaged number of requests received by the application.

Per App SLA: Measure reflecting application's performance meeting SLA requirements like percentage of responses within threshold specified by admin.

Per App Responses and Their Breakdown: Breakdown of responses by the Hypertext Transfer Protocol (HTTP) response codes.

Per HTTP App Navigation Timings: Measure of application's latencies as perW3C navigation timing standard

Per App SSL Performance: Metrics representing number of Secure Socket Layer (SSL) handshakes.

Per App SSL Error Counters: Metrics representing percentage of SSL Errors and its breakdown.

Per App SSL Protocol Breakdown: Breakdown of the SSL sessions by the SSL/TLS (Transport Layer Security) protocol use. This can be used for analyzing security characteristics of the application.

Per App SSL Session Key Exchange and Certificate: Breakdown of the SSL sessions by key exchange algorithm for an application.

A.3 Infrastructure Metrics

Per VM/Host Resource Utilization and Saturation Metrics: CPU, Memory, Disk, I/O, etc. Any metric that can represent the utilization of the resource and saturation of a system are very helpful in estimating capacity and understanding data center operations.

Per VM/Host Resource Availability Metrics: Metrics that measure unavailability of resources due to sharing can be critical in understanding and triaging operational issues.