Internet Research Task Force (IRTF) Internet Draft Category: Informational R. Krishnan Brocade Dilip Krishnaswamy IBM Research D. R. Lopez Telefonica I+D Asif Qamar Evolv Steven Wright AT&T Norival Figueira Brocade

Expires: April 2015

November 11, 2014

NFV Real-time Analytics and Orchestration: Use Cases and Architectural Framework

draft-krishnan-nfvrg-real-time-analytics-orch-01

Abstract

One of the key goals of NFV is to optimize the infrastructure resource usage while driving operational simplicity. Real-time analytics providing insight into various components such as compute (e.g. dynamic CPU utilization), storage (e.g. dynamic capacity usage), network (e.g. dynamic bandwidth utilization), energy (e.g. dynamic power consumption) is key to not only providing visibility into the NFV infrastructure and thus driving operational simplicity but also optimizing resource usage for the purposes of orchestration. This draft focusses on use cases and architectural framework for real-time analytics and orchestration including Big Data predictive analytics for addressing the aforementioned requirements.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of $\underline{BCP 78}$ and $\underline{BCP 79}$.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/lid-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire in April 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>http://trustee.ietf.org/license-info</u>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Conventions used in this document

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 <u>RFC 2119</u>.

Table of Contents

<u>1</u> .	Introduction <u>3</u>
<u>2</u> .	Real-time Analytics and Orchestration Use Cases $\underline{4}$
	<u>2.1</u> . Enhancements to Real-time Analytics Application 5
	<u>2.1.1</u> . Distributed Predictive Analytics5
	<u>2.1.2</u> . Detecting Noisy Neighbors5
	2.1.3. Addressing security issues due to inconsistent
	configuration
<u>3</u> .	Real-time Analytics and Orchestration Architectural Framework6
<u>4</u> .	Summary
<u>5</u> .	Future Work
<u>6</u> .	IANA Considerations <u>6</u>
<u>7</u> .	Security Considerations <u>6</u>
<u>8</u> .	Contributors <u>6</u>
9.	Acknowledgements

Expires April 2015

[Page 2]

<u>10</u> . References			
<u>10.1</u> .	Normative References $\underline{7}$		
<u>10.2</u> .	Informative References $\underline{7}$		
Authors'	Addresses <u>8</u>		

1. Introduction

Operator Network Function Virtualization Infrastructure Point-of-Presence (NFVI-PoP) locations [ETSI-NFV-TERM] often have capacity, energy and other constraints. Thus, optimizing overall resource usage is an important requirement [ETSI-NFV-REQ]. The general case must consider a distributed (elastic) VNF NFVI platform implementation where VMs running for different VNFs (with different characteristics) can co-exist in the same physical server. This case must address the goal of optimizing overall resource usage through mechanisms like bin-packing [BIN-PACK].

In this context, some of the important challenges faced are:

- . Performance issues due to noisy neighbor effect, where a VM running for a VNF can affect the VM(s) running for another VNF.
- . Security issues, especially due to inconsistent configuration in a dynamic environment where one VNF could affect others.
- . Energy Efficiency given that servers have substantial idle power usage.
- . Resources used (Compute, bandwidth, storage) for the realtime analytics in comparison to the VNF payload resource usage.

The purpose of this document is two-fold. First, it intends to discuss various use cases to describe the above challenges. Second, it will depict an architectural framework for real-time analytics and orchestration, applicable to the above use cases in a multivendor environment.

For the purposes of real-time analytics for orchestration, various metrics need to be collected, stored and analyzed.

Metrics collection: Metric collection may occur at different periods during the lifecycle of the VNF. Metric collection during an onboarding process in a controlled load configuration may provide a baseline for characterization of "normal" operational performance. Such baseline characterizations may be useful for detection "out of normal" performance at a later point in the VNF lifecycle.

Expires April 2015

[Page 3]

Metrics storage: It is recommended to store and analyze metrics locally to minimize the costs of backhaul to remote locations.

Metrics analysis: The assumption here is that the metrics to be collected for analysis would be VNF independent in the sense that they would apply regardless of the type of VNF. Metrics that are specific to particular types of VNF are more appropriate for service specific diagnostics.

2. Real-time Analytics and Orchestration Use Cases

A real-time analytics application periodically collects metrics (also called information in this document) from individual VMs, VNFs, physical servers, network elements etc. regarding various subsystems such as compute (e.g. dynamic CPU utilization), storage (e.g. dynamic capacity usage), network (e.g. dynamic bandwidth utilization), energy (e.g. dynamic power consumption) through polling. The real-time analytics application computes the average utilization for VMs, VNFs, physical servers, networks etc. regarding the various sub-systems such as compute (e.g. average CPU utilization), storage (e.g. average capacity usage), network (e.g. average bandwidth utilization), energy (e.g. average power consumption).

Using the average utilization information, the real-time analytics application provides real-time visibility into the operating point of the VNF in the NFV Node thus driving operational efficiency.

The NFV orchestrator uses the average utilization information from the real-time analytics application to determine the appropriate time to scale up/down the running software instances. Typically the thresholds for scale up/down are manually programmed into the system - this may not be performance optimal since the workloads and deployment scenarios can substantially vary.

In addition, predictive analytics based on machine learning techniques [MACHINE-LEARNING-BOOK] can be used by the real-time analytics application to automatically determine the appropriate thresholds for scale up/down the running software instances for differing workloads including events related to social behavior (think of a YouTube video going viral) and deployment scenarios. This information can be used by the orchestrator for optimizing overall performance and maximizing energy efficiency. Maximizing energy efficiency comes from the fact that by determining the appropriate thresholds for scale up/down the workloads can be consolidated into a minimum set of physical resources so the rest of the unused physical resources can be completely powered off to avoid

Expires April 2015

[Page 4]

any idle power consumption. [SPEC-BENCHMARK] analyzes the power profile of physical servers from various vendors; the active idle power consumption of physical servers could be as much as 30%.

<u>2.1</u>. Enhancements to Real-time Analytics Application

<u>2.1.1</u>. Distributed Predictive Analytics

A real-time analytics application could be notified of significant events by individual running software instances of VMs, VNFs etc. or by infrastructure elements such as physical servers, hypervisors etc. This helps reduce the rate of polling by the real-time analytics application and also helps in reacting to significant events such as overload much faster. The challenge in this case is to determine the appropriate thresholds (e.g. average power consumption has been higher than x Watts for t seconds) for event notification.

Predictive analytics engines which use machine learning techniques [MACHINE-LEARNING-BOOK] can be used to determine the appropriate thresholds per running software instance and infrastructure element for different workloads and deployment scenarios. These predictive analytics engines can run in various nodes in the infrastructure in a distributed predictive analytics architectural framework.

2.1.2. Detecting Noisy Neighbors

In the context of multiple VNFs, "Noisy Neighbor Effect" could be defined as follows: the VM running for one VNF can affect the performance of a VM running for another VNF in the case where they are using the same physical resources (physical servers, physical network elements). A real-time analytics application could help in detecting and mitigating the noisy neighbor effect. A good example is the case where the VMs running for two VNFs share the same physical server, are memory access intensive (load balancers, firewalls etc.) and have correlated memory access patterns for the given workload and deployment scenario.

Real-time big data analytics techniques [RT-ANALYTICS-BOOK] can be used by the analytics application to determine such correlation patterns which can affect performance in real-time. Additionally, predictive analytics based on machine learning techniques [MACHINE-LEARNING-BOOK] can be used to predict the frequency and duration of such correlation patterns. This information can be used to create dynamic anti-affinity rules for VM placement and migration including redundancy considerations - e.g. VMs of VNF "A" cannot co-exist with VMs of VNF "B".

Expires April 2015

[Page 5]

2.1.3. Addressing security issues due to inconsistent configuration

NFV configuration is expected to be dynamic, especially in the edge NFV PoPs where capacity is limited; a very good example is handling a viral event such as mobile gaming application. While autonomic networking techniques could be used to automate the configuration process including modular updates, it is important to take into account that incomplete and/or inconsistent configuration may lead to security issues. Distributed VNF implementations (e.g. VMs of single VNF which span different physical servers) typically use an eventually consistent configuration model [CAP-THEOREM] for scalability reasons -- this poses additional security challenges.

Real-time analytics techniques [<u>RT-ANALYTICS-BOOK</u>] can be used by the analytics application to determine communication pattern anomalies due to incomplete and/or inconsistent configuration in real-time by analyzing event logs. Additionally, predictive analytics based on machine learning techniques [MACHINE-LEARNING-BOOK] can be used to predict the frequency and duration of such communication pattern anomalies. A simple example is a flow-specific firewall rule which never got installed due to reasons such as control plane messaging issues, data plane table full condition etc.

3. Real-time Analytics and Orchestration Architectural Framework

TBD

4. Summary

TBD

5. Future Work

TBD

6. IANA Considerations

This draft does not have any IANA considerations.

7. Security Considerations

- 8. Contributors
- 9. Acknowledgements

None.

Expires April 2015

[Page 6]

10. References

<u>**10.1</u>**. Normative References</u>

<u>**10.2</u>**. Informative References</u>

[ETSI-NFV-WHITE] "ETSI NFV White Paper," http://portal.etsi.org/NFV/NFV_White_Paper.pdf

[ETSI-NFV-USE-CASES] "ETSI NFV Use Cases," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_N FV001v010101p.pdf

[ETSI-NFV-REQ] "ETSI NFV Virtualization Requirements," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/004/01.01.01_60/gs_N FV004v010101p.pdf

[ETSI-NFV-ARCH] "ETSI NFV Architectural Framework," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.01.01_60/gs_N FV002v010101p.pdf

[ETSI-NFV-TERM] "Terminology for Main Concepts in NFV," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/003/01.01.01_60/gs_n fv003v010101p.pdf

[OPENSTACK] "OpenStack Open Source Software," https://www.openstack.org/

[OPENSTACK-CONGRESS-POLICY-ENGINE] "A policy as a service open source project in OpenStack," <u>https://wiki.openstack.org/wiki/Congress</u>

[OPENSTACK-CELIOMETER-MEASUREMENT] "OpenStack Celiometer," http://docs.openstack.org/developer/ceilometer/measurements.html

[OPENSTACK-NOVA-COMPUTE] "OpenStack Nova," https://wiki.openstack.org/wiki/Nova

[NFV-MANO-SPEC] "NFV Management and Orchestration Framework Specification," <u>http://docbox.etsi.org/ISG/NFV/Open/Latest_Drafts/NFV-MAN001v061-</u> %20management%20and%20orchestration.pdf

[BIN-PACK] Coffman, Jr., E., M. Garey, and D. Johnson. Approximation Algorithms for Bin-Packing -- An Updated Survey. In Algorithm Design for Computer System Design, ed. by Ausiello, Lucertini, and Serafini. Springer-Verlag, 1984.

Expires April 2015

[Page 7]

[SPEC-BENCHMARK] "SPEC Benchmark Results: HP Proliant DL380p Rack Server," <u>http://i.dell.com/sites/doccontent/shared-content/data-</u> <u>sheets/en/Documents/Comparing-Dell-R720-and-HP-Proliant-DL380p-Gen8-</u> Servers.pdf

[CAP-THEOREM] Eric Brewer, "CAP twelve years later: How the "rules" have changed", IEEE Explore, Volume 45, Issue 2 (2012), pg. 23-29.

[MACHINE-LEARNING-BOOK] Ian H. Witten et al., "Practical Machine Learning Tools and Techniques, Third Edition," Morgan Kaufmann, 2011

[RT-ANALYTICS-BOOK] Byron Ellis, "Real-Time Analytics: Techniques to Analyze and Visualize Streaming Data," Wiley, 2014

Authors' Addresses

Ram (Ramki) Krishnan Brocade Communications ramk@brocade.com

Dilip Krishnaswamy IBM Research dilikris@in.ibm.com

Diego Lopez Telefonica I+D Don Ramon de la Cruz, 82 Madrid, 28006, Spain +34 913 129 041 diego.r.lopez@telefonica.com

Asif Qamar Evolv asif@asifqamar.com