

Energy Savings & Resiliency with Closed Loop Platform Automation

John Browne
Data Center Group
Intel Corporation
Shannon, Ireland
john.j.browne@intel.com

Krzysztof Kepka
Data Center Group
Intel Corporation
Gdansk, Poland
krzysztof.kepka@intel.com

Patrick Kutch
Data Center Group
Intel Corporation
Hillsboro, USA
patrick.g.kutch@intel.com

Sunku Ranganath
Data Center Group
Intel Corporation
Hillsboro, USA
sunku.ranganath@intel.com

Abstract—This demo showcases the advancements made through utilizing OPNFV Barometer project by introducing local and remote corrective action frameworks with industry standard cloud native orchestrators like Kubernetes, monitoring agents like Collectd, Time Series Database (TSDB) like Prometheus and InfluxDB for widely adopted NFV use cases of Virtual Border Network Gateway (vBNG) and Virtual Cable Modem Termination Systems (vCMTS). Leveraging advanced platform reliability and adaptive power saving telemetry, the demo provides methodologies in to identifying the platform capabilities & provision them for optimal run time power savings and reliability. The two part demo showcases in part one, automated memory fault detection and swap over of vBNG containers to hot stand by containers via local corrective action, while part two showcases adaptive power control of containerized vCMTS workload across multiple compute nodes, essentially providing closed loop platform automation.

Keywords—service assurance, closed loop platform automation, BNG, CMTS, telemetry, DPDK

I. INTRODUCTION

The industry trend to zero-touch Network and Services Management [1] requires operational services to be automated in a similar manner to IT and data center infrastructure. Communication Service Providers often look for improving automation of reliability and power efficiency of their cloud deployments. Many existing methodologies rely on troubleshooting after the fact faults or outages occur. Energy efficiency plays an important role in reducing operational expenditure and the objective is to automate the power control schemes and achieve the best combination of throughput, packet latency, and packet loss while achieving cost savings. The paper selects two use cases, to illustrate closed loop automation, a resiliency use case and a power efficiency use case. The architecture of the use case demonstrations which consist of an open industry standard platform telemetry layer, using collectd, an open monitoring and alerting infrastructure with Prometheus and open orchestration system with kubernetes provides the key building blocks for an open closed loop automation architecture which can expand to many Communication Service Provider use cases.

II. KEY CONCEPTS & DESCRIPTION

The paper outlines two use cases of closed loop automation, a resiliency use case demonstrated using a Virtual Border Network Gateway (vBNG) and a Power Efficiency use case demonstrated using a Virtual Cable Termination Modem System (vCMTS) leveraging Data Plane Development Kit (DPDK) [2].

A. Power Savings with vCMTS

Service providers are looking to reduce power footprint consumed by VNFs and advanced IA platform features like Intel Speed Select help provide required platform support to accomplish this. The first part of this demo showcases ‘intelligent’ scaling of CPU Core Frequencies associated with vCMTS instances running in Kubernetes environment based upon Key Performance Indicators (KPIs) gathered via collectd in order to achieve ‘better’ power consumption as shown in Figure 1. The demo would scale CPU frequencies based on current KPIs. During times when bandwidth requirements are lessened, the CPU core frequencies are scaled down to achieve power savings and vice versa as there is more compute power required for each vCMTS instance.

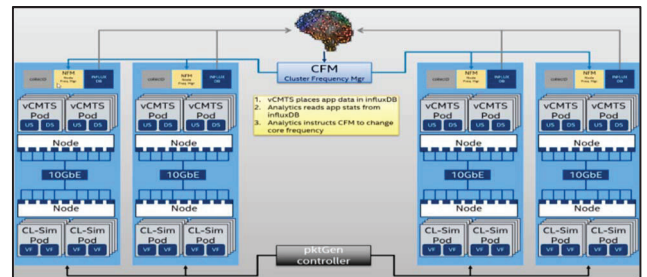


Figure 1: vCMTS deployment topology

This is done in conjunction with reducing the traffic flowing from the traffic generators – the systems are not expected to be able to perform at the same level with a reduction in core frequency.

B. Closed Loop Resiliency with vBNG

Broadband Network Gateway (BNG) is the access point for subscribers, through which they connect to the broadband network. One area of significant interest to Service Providers in recent times is the relationship between automated operations and excellent customer experience [3]. The demo shows that by using Intel Architecture (IA) based x86 platform specific metrics and events, we can monitor the health of the platform & identify issues that may impact the end-user experience of vBNG.

The solution is application agnostic but we have selected a DPDK based vBNG reference application to showcase the closed loop failover & resiliency capability using a platform generated fault, with orchestration changes resulting in the restoration of vBNG service as indicated in Figure 3. We trigger a platform fault which is reflected via a platform metric and as part of the closed loop error detection and correction. The traffic

will be switched from the active to the standby application to maintain an uninterrupted service. We run two vBNG instances at the start of the demonstration [A] active and [B] standby, on two NUMA nodes on a single server platform. The vBNG instance created by kubernetes will start processing traffic immediately the traffic becomes active.




Figure 3: Closed Loop vBNG demo Topology

III. DEMO PROCESS

A. Closed Loop Automated Power Savings Demonstration with vCMTS

The vCMTS solution utilizes DPDK for packet processing; DPDK usually fully utilizes CPU Cores assigned to the poll mode drivers – they run at full speed even when there is not a lot of traffic. The vCMTS demo will show when the workload (bandwidth) requirements lesson over a period of time, intelligent CPU Core Frequency changes significantly reduce power utilization without sacrificing performance. The demo shows several comparative screens showing power consumption when no scaling is performed, the savings when a cluster centralized agent utilizing machine learning makes frequency changes based upon KPIs gathered from collectd and we hope to show power savings achieved by a node-local agent with an extremely tight reaction time achieved by locally monitoring specific platform KPIs.

B. Closed Loop Automated Resiliency demonstration with vBNG

Stage 1: Server 1, hosts two vBNG instances deployed by Kubernetes. vBNG instance 1 is in active mode receiving traffic on socket 0, the other vBNG instance B is in standby on socket 1 with Collectd running & integrated with Prometheus. Server 2 hosts the traffic generator for the 2 vBNG instances & Prometheus Alert Handler (PAH).




Figure 4: Demo Process Description with vBNG instance

Stage 2: Error scenario is triggered by injecting 5 correctable memory Reliability Availability Serviceability (RAS) fault resulting in DIMM failure on socket 0. Collectd sends these events to PAH. An alarm is raised to alert that a DIMM requires maintenance while triggering the remediation action

Stage 3: In this instance the remediation action is to stop traffic on vBNG instance 1 and start it on vBNG instance 2, resuming the service as normal with minimal disruption to the service for the customer.

IV. HIGHLIGHTS OF INNOVATION

A. Energy Efficiency & Power Savings

The Figure 7 provides a quick visualization of power savings of that can be achieved without sacrificing performance.

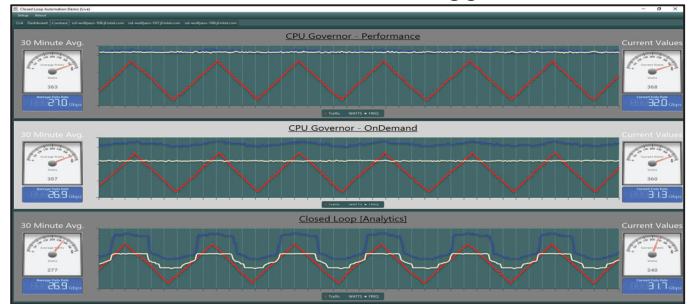


Figure 7: Power Savings of ~80-100 W with Closed Loop Automation

B. Automated Resiliency & Reliability

For Resiliency and Availability scenarios the key system metrics are indicated in Figure 9 [4]. The demo provides insights to reduce down time due to hardware faults, improve resiliency & apply corrective action in an automated way.




Figure 9: Automated Maximization of Service Availability

REFERENCES

- [1] Industry Sepcification Group Zero Touch Network Service Management, ETSI, etsi.org/technologies/zero-touch-network-service-management
- [2] Linux Foundation Projects, Data Plane Development Kit, dpdk.org
- [3] Broadband Forum Technical Report, TR-345 BNG & NFV, broadband-forum.org/technical/download/TR-345.pdf
- [4] OpenSAF and VMware from the Perspective of High Availability - Ali Nikzad, Ferhat KhendekMaria Toeroe Concordia University Ericsson SVM'2013 – Zurich – October 2013