Introduction to Lenovo Network Telemetry

Describes Lenovo switch capabilities for continuous network data gathering

Shows how to use these capabilities to monitor virtual machines

Shows the supported architecture and tools for data analysis

Specifies the available data elements which can be monitored

Chidambaram Bhagavathiperumal

Roberto H Jacob Da Silva

LENOVO PRESS

Abstract

This paper introduces the telemetry capabilities which are included in Lenovo®'s CNOS network switch firmware. These capabilities allow for continuous gathering of a variety of switch-related data elements and also data relating to virtual machines whose hosts are connected to a Lenovo switch. The data are made available through the standard REST API, and can be presented through the use of a variety of available tools.

This paper is for network architects, analysts and planners looking to understand Lenovo's networking switch telemetry capabilities.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction 3
Network Insights. 4
Data center use cases 4
Broadcom BroadView 6
Telemetry for virtual machine visibility 7
Conclusion 8
Authors. 9
Notices 10
Trademarks 11

Introduction

Modern data centers are characterized by a phenomenal growth, with the expansion of cloud, software defined networks, hyper-scale, and big data technologies. As the network size and complexity increases, the need for a telemetry solution becomes evident. Telemetry frameworks aim at providing insights on the network utilization and providing tools for data center wide optimization, which translates into the effective usage of compute and network resources. Lenovo network telemetry alleviates these challenges and provides end-to-end visibility into networks.

With the growth of cloud networks, traditional static networks are moving towards dynamic virtualized networks. The size of such networks calls for an evolution of the telemetry tools. Traditional telemetry techniques based on SNMP polling no longer scale, and APIs providing both fine-grained visibility into application performance and network-wide monitoring have become critical requirements.

Lenovo network telemetry features are built to provide in-depth insights along with need based reporting using REST APIs. While the telemetry functionality is included in CNOS, a tool such as Ganglia or Splunk needs to be used and customized to process and display the data to meet the customer's needs. The architecture, including the tool used to present it, will enable the frequency and completeness of information necessary to provide for reactive and proactive tuning of the network.

The telemetry architecture is shown in Figure 1, below. The figure shows an ecosystem comprised of network switches providing telemetry data, monitoring tools to store, analyze the network telemetry data and orchestration application which collects the feedback from monitoring tools and optimizes network performance.

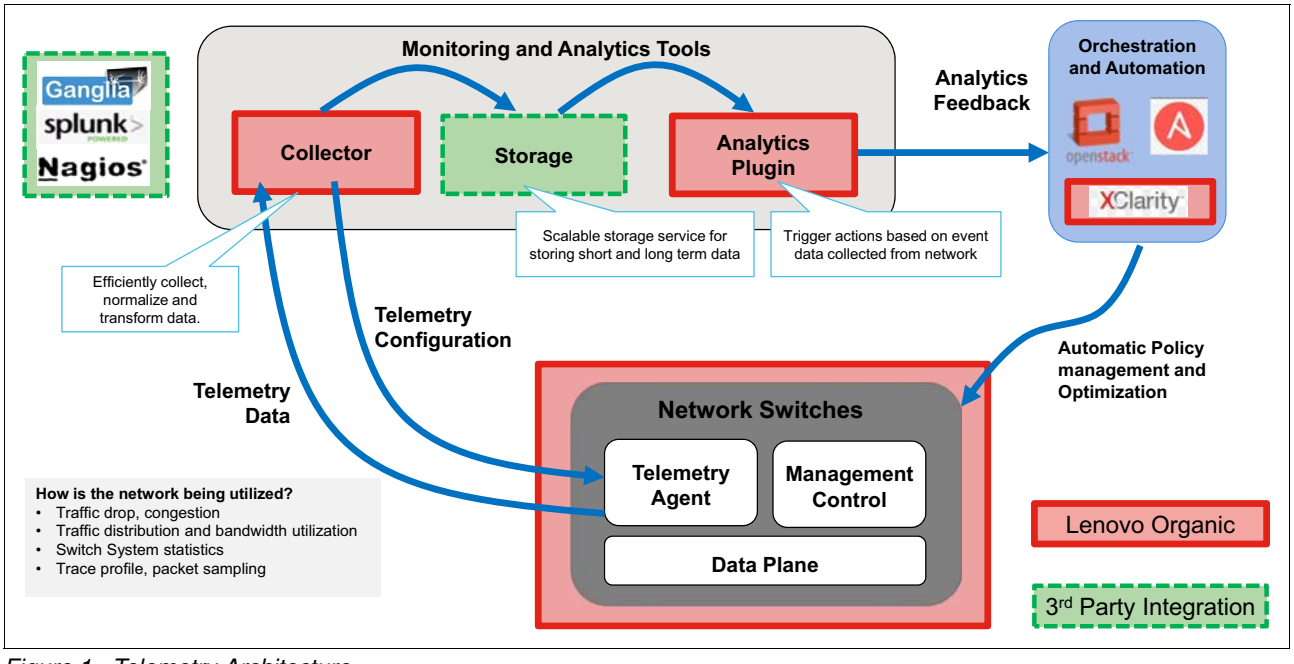

Figure 1 Telemetry Architecture

Network Insights

The Lenovo network telemetry solution reduces operational costs by providing key network insights including:

- Visibility of network utilization allows IT administrators to plan their network in advance and redistribute traffic loads to leverage underutilized infrastructure optimally.
- Identification of workloads per application groups and communication patterns.
- Proactive detection of problems (congestion, hot spots detection) before they occur, through trend monitoring, or after the fact (micro-burst detection).

An example of the type of data that can be collected, and the type of conditions that can be detected is shown in Figure 2.

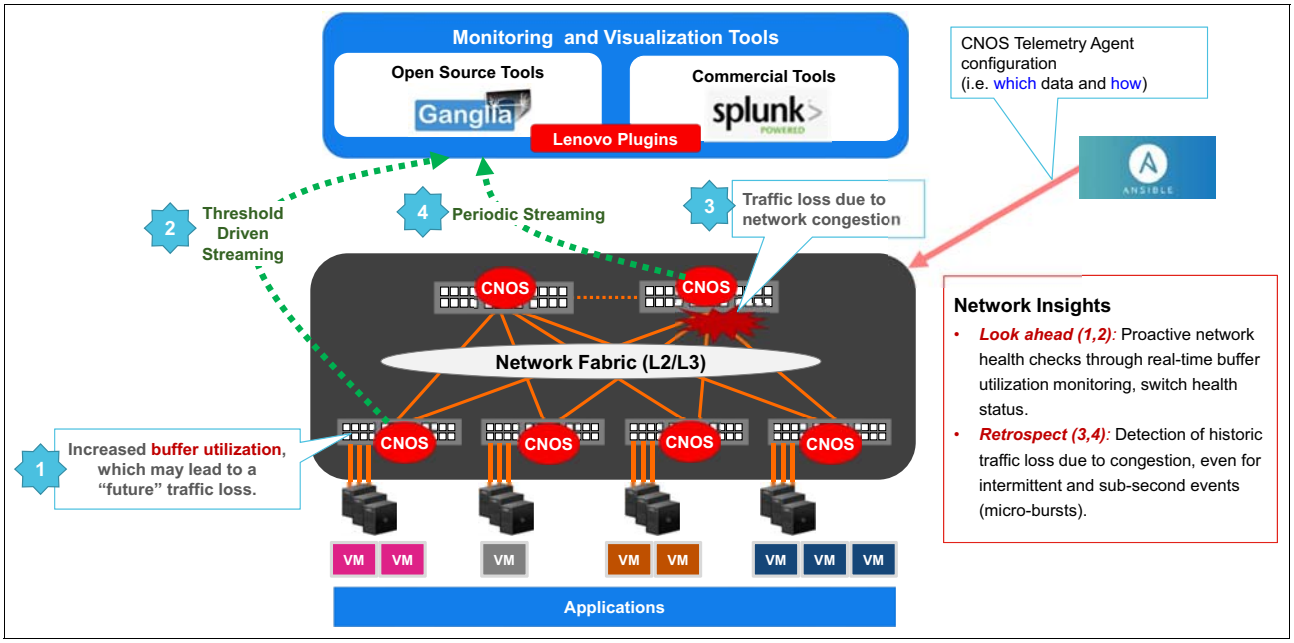

Figure 2 Telemetry ecosystem, with switches and tools, and examples of the types of conditions that can be detected

Data center use cases

The Lenovo network telemetry solution addresses the key challenges on the data center for network monitoring.

Congestion detection

Scalable cloud networks have expanded, and currently are deployed over multi-tier topologies supporting east-west traffic with high data rates ranging from 10Gb to 100Gb. Therefore, it has become very difficult to predict and analyze the network performance using standard network interface statistics.

As a result, a richer set of parameters providing in-depth insights into the state, errors, and performance of the network devices has become paramount. Moreover, the ability to quickly detect changes in traffic patterns, based on application needs, provides a key requirement for monitoring applications.

Lenovo network telemetry delivers a framework capable of detecting network congestion regardless of the duration of the event. The detection of network congestion, and hence increased latency, can be performed at various levels including device, interface, traffic direction (i.e. ingress versus egress traffic), traffic type (i.e. unicast versus multicast), and with class-of-service.

This wealth of information is crucial for monitoring applications because it not only identifies congestion/hot-spots, but also provides the traffic pattern causing the bottleneck. This depth of visibility helps network administrators to change policies, which can guarantee lossless behavior for time sensitive traffic, and optimize the network.

Data center planning

Virtualization technologies support dynamic deployment and expansion of applications running in the data center. This flexibility creates a challenge for the network administrators related to how to proactively keep the network infrastructure up-to-date, so it can handle the applications' requirements as far as connectivity, bandwidth, and performance.

The Lenovo network telemetry solution provides network administrators with information about current and historic network resources (switch buffers, interface buffers) utilization. This knowledge is an important asset to help network administrators to proactively understand how the network infrastructure is being utilized, and plan for its expansion based on relevant parameters such as bandwidth, interface utilization.

Open APIs

SNMP has been the default option for monitoring network devices. Unfortunately, SNMP has several disadvantages including complex OID management, limited scalability, and no real-time monitoring data support.

In order to address these challenges, the Lenovo network telemetry provides REST APIs for real time monitoring of data with asynchronous reporting, which equates to streaming of data driven by events. This approach makes the Lenovo Network Telemetry solution suitable for cloud analytics ecosystems. Nevertheless, the framework is flexible and does support multiple data collection techniques:

- **On Demand**: Periodic "pull" of telemetry data from an analytics application, which gives the user the ability to retrieve any data at anytime.
- **Time Series**: Periodic "push" transfer of telemetry data, which is useful for information that changes at low frequencies, or for the cases where the user wants to build a database of historic state.
- **Triggered**: Asynchronous "push" of reports on pre-configured threshold breach.

Lenovo has developed a collector plugin for Ganglia, show casing a few use cases including congestion detection and capacity planning. The plugin is available on GitHub at:

https://github.com/lenovo/networking-telemetry

Figure 3 on page 6 shows the available data elements that can be captured and displayed.

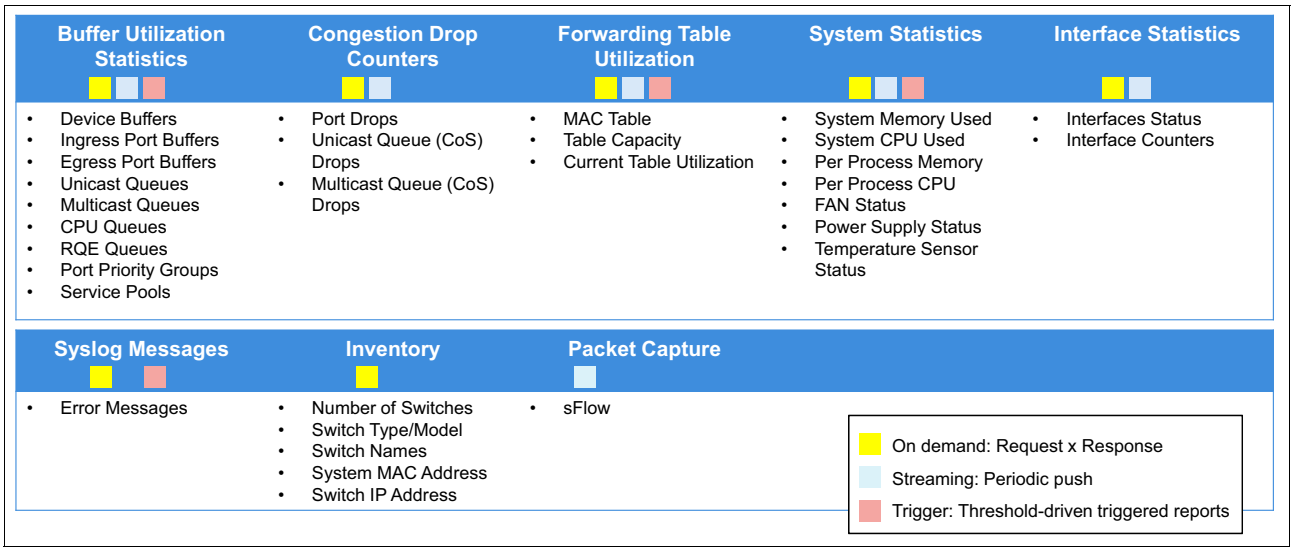

Figure 3 Available telemetry data elements

Broadcom BroadView

Lenovo has integrated Broadcom BroadView libraries to its next generation cloud network operating system (CNOS) on Broadcom based TOR switches. Lenovo CNOS with BroadView integration provides in-depth insights of network congestion and hardware resource utilization, helping network administrators proactively detect failure and do long term capacity planning.

The BroadView telemetry solution involves two software components:

- **BroadView Instrumentation Agent**: A software component, which runs on the switches, and it is responsible for obtaining data from the switch and exposing this information using a variety of transfer techniques.
- **BroadView Analytics Application**: An analytics application, which runs outside the switches, typically on an external server or controller, and is responsible for collecting telemetry data from various switches. The analytics application requests and/or receives data from the agent via REST API.

CNOS integration with BroadView

The BroadView agent is integrated into Lenovo CNOS. The basic agent has been enhanced to customize the interfaces exposed via REST API to make them more generic and user friendly. For example, human readable interface names are used (e.g. Ethernet1/1) rather than Broadcom specific internal numbers.

Lenovo has also integrated with Ganglia, the open source distributed monitoring system for high-performance computing clusters, but is equally usable on Lenovo switches in non-HPC environments.

Lenovo is also looking to extend the scope of the analytics applications that they can interact with our telemetry agent by developing plugins for Splunk (a commercial monitoring and logging software product), and integration with OpenContrail (a Software Defined Network management product.

Telemetry for virtual machine visibility

In today's virtualized data centers, virtual machines connect to each other through virtual switches which, in turn, attach to the physical infrastructure. This deployment adds an abstraction layer in which the network administrator typically has no way to correlate the virtual machine connectivity information with the underlying network infrastructure.

Lenovo CNOS includes Network Policy Agent, which has been integrated with the Nutanix Acropolis and VMware ESXi hypervisors to provide virtual machine management as well as virtual machine visibility.

The functions provided by the Network Policy Agent are:

- Virtual Machine Visibility
  - Workload visibility including virtual machine information as well as data about the overlay network
  - Auto-discovery of vNetwork topology
  - Virtual network statistics
- Auto-Provisioning of Network
  - Automatic configuration of physical networks based on the discovered topology (e.g. port and VLAN configuration are automated)
- Automatic VM-aware Network updates
  - Ongoing, dynamic updates to the physical network configuration in response to new virtual machines, updated virtual machines, and deleted virtual machines
  - Increased efficiency (since multiple switches are dynamically updated; with manual updates the administrator performs one at a time or has to enable other automation tools)
  - No dependency on human intervention
  - Eliminates errors from manual configuration

Figure 4 on page 8 illustrates how Lenovo ThinkAgile™ Network Orchestrator provides a network administrator with the visibility of virtual network infrastructure along with details of workloads or VMs.

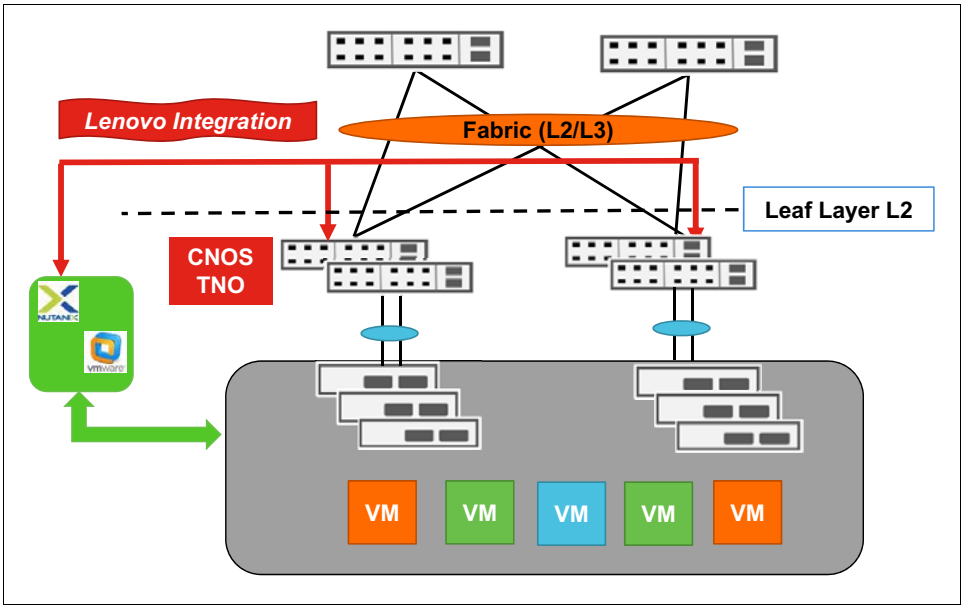

Figure 4 ThinkAgile Network Orchestrator

Conclusion

The Lenovo telemetry solution is a powerful tool which allows the administrators to reduce their network infrastructure operational cost by providing in-depth insights into the network. The rich data set, event driven architecture, and open APIs are the foundation for achieving workload visibility, real-time streaming of analytics, proactive planning, and high-level centralized problem resolution. Customers can choose their preferred tool to store, analyze, and present this data.

The telemetry functionality is included in Lenovo CNOS firmware as a standard feature, and is available on Lenovo switches which use CNOS, including (as of this writing) G8272, G8296, G8332, NE1032 and 1032T, NE1072T, NE2572, and NE10032.

Authors

This paper was produced by the following team of specialists:

Chidambaram Bhagavathiperumal is a Software Architect for Lenovo. He performs network architecture design and implementation for hyper-converged infrastructure, network telemetry and SDN solutions that involve Lenovo products. He joined Lenovo as part of the System x® acquisition from IBM and has 14 years of experience in networking products. He is also the author of 17 issued U.S patents on networking solutions and has published multiple white papers. Chidambaram holds a holds a BS degree in Computer Science and Engineering from MS University in India.

Roberto H Jacob Da Silva is a Lead Software Architect for Lenovo Networking. He joined Lenovo from IBM and has over 20 years of networking experience working in Belgium, Brazil, Spain and the United States. Roberto is the author of 15 issued U.S. patents on LAN switching and is a former Distinguished Member of Technical Staff at Alcatel-Lucent, the title given to the company's top 60 R&D engineers worldwide annually. Roberto graduated from the Technological Institute of Aeronautics (Sao Paulo, Brazil) with a BS in Electronics Engineering.

Thanks to the following people for their contributions to this project:

- Scott Lorditch
- David Watts, Lenovo Press

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on May 18, 2018.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp0894

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®	System x®
Lenovo(logo)®	ThinkAgile™

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.