

TECHNICAL BRIEF

BroadView

Building an Open Source Data Center Monitoring Tool Using Broadcom BroadView™ Instrumentation Software

What is BroadView?

Broadcom's BroadView™ software suite showcases the innovative instrumentation and telemetry features of Broadcom's industry-leading silicon. It includes agents, applications, and software that are open-sourced, so customers can easily adapt them into their solutions. The instrumentation features include Buffer Statistics Tracking (BST), packet tracing, and ECMP resolution, with others to follow. These features enable advanced network optimization tools for end-to-end network monitoring, leading to improved network performance and optimal network usage.

BST provides the ability to monitor the MMU buffer utilization, enabling operators to proactively detect microbursts and take actions to improve network performance when congestion occurs.

Features:

- A programmatic interface via REST APIs.
- Runs on OpenNSL, the Broadcom SDK, or a Network OS.
- SDN-based reference implementations for easy adoption.
- Deeper network telemetry that helps operators diagnose network problems, optimize network resources, and meet strict service-level agreements.

OVERVIEW

With network performance becoming an increasingly critical factor in the success of any business, network operators are under constant pressure to ensure that their networks are responsive, 'always-on', and capable of meeting the ever-growing demand from applications they run. Providing operators with deeper instrumentation and telemetry data about the network helps operators diagnose network issues, plan and fine-tune the network to provide improved performance, and make optimal use of network resources.

Traditionally "congestion" has been associated with switch ports being utilized at close to line rate. In a congestion scenario, packets can be dropped by the switch or flows may be backpressured due to lack of buffer space. However, more recent analysis has uncovered the existence of microbursts, where a rapid burst of packets at line rate can temporarily overflow the switch buffers and cause packet loss or backpressure. Typically, these microbursts do not last long enough to be detected by traditional switch counters such as SNMP or port statistics.

BroadView instrumentation software leverages unique telemetry features offered by Broadcom silicon, such as Buffer Statistics Tracking (BST), to detect these microbursts and enable software to take corrective action. BroadView offers a rich set of instrumentation parameters that provide network administrators and applications detailed knowledge about the state of the network. The BroadView Suite includes the BroadView agent and applications that interface with the agent, presenting and analyzing data they receive from the agent and tuning the network based on the insights gained.

This technical brief focuses on the BST feature of the BCM56850 and BCM56960 switches from Broadcom's XGS[®] family.

The BroadView agent resides on the switch and interfaces with the switch silicon via OpenNSL or the Broadcom SDK. The agent can be modified to use any NOS-provided API to access the silicon. The agent collects data related to the instrumentation features (e.g., BST statistics) and makes it available to the application, which can reside on the network device or on an external controller/collector. In this demo, Ganglia with the BroadView Plugin is the external controller/collector.

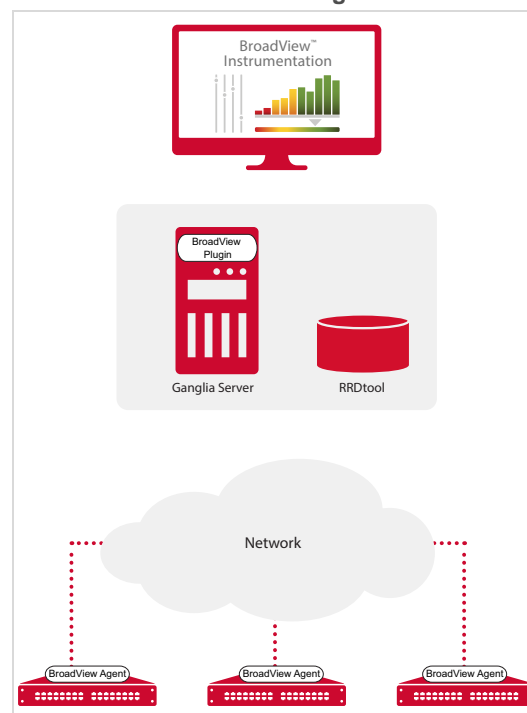
The BroadView agent communicates with the application using REST APIs, with data exchanged in the JSON-RPC (2.0) format. The agent supports both the 'pull' model of operation, where the collector requests and obtains data, and the 'push' model, where the agent sends periodic reports, asynchronously.

The agent must be configured with parameters such as the IP address and port number for communicating with the Collector and how frequently reports are to be sent. This configuration is made via a text file that is applied during agent startup. If a startup configuration file is not available, the default configuration is used.

In the solution described here, Ganglia is used as an application to provide visualization of the BST statistics provided by the agent. Ganglia is a scalable distributed monitoring system for high-performance computing systems such as clusters and Grids. It is based on a hierarchical design targeted at federations of clusters. It leverages widely used technologies such as XML for data representation; XDR for compact, portable data transport; and the RRDtool for data storage and visualization (see www.ganglia.info).

The BroadView agent running on each switch sends its statistics report using a REST API to the Ganglia server, both periodically and when a threshold for a counter (if set) is reached. The Ganglia daemon gathers the data and displays it in a graphical format. The graph can be shown as line graph or a bar graph. BroadView enables the user to monitor the traffic and get notifications upon congestion. The BST statistics include thousands of counters related to buffers and can show both ingress as well as egress values for unicast and multicast traffic.

BroadView and Ganglia Demo



BENEFITS OF BROADVIEW IN THIS DEMO

In this demo, Ganglia provides a single point of access to view packet buffer utilization. The BroadView solution is extensible in that it allows the Ganglia interface to combine the counters across all priority groups in a port to a single counter for that port. The solution also supports network fabric view by visualizing metrics from multiple switches at the same time through Ganglia.

The benefits of using BroadView include:

- The ability to visualize advanced telemetry to monitor and detect microburst-like congestion scenarios.
- Detailed statistics displayed in an easy-to-use graphical format.
- User control of which statistics are reported.

- A web interface that allows access from anywhere.
- Easy integration of the Instrumentation agent with the network OS of your choice.

Having a proactive statistics tool that can give details per queue can be a great help in monitoring and maintaining network stability and encouraging optimum network usage, particularly in datacenter environments.

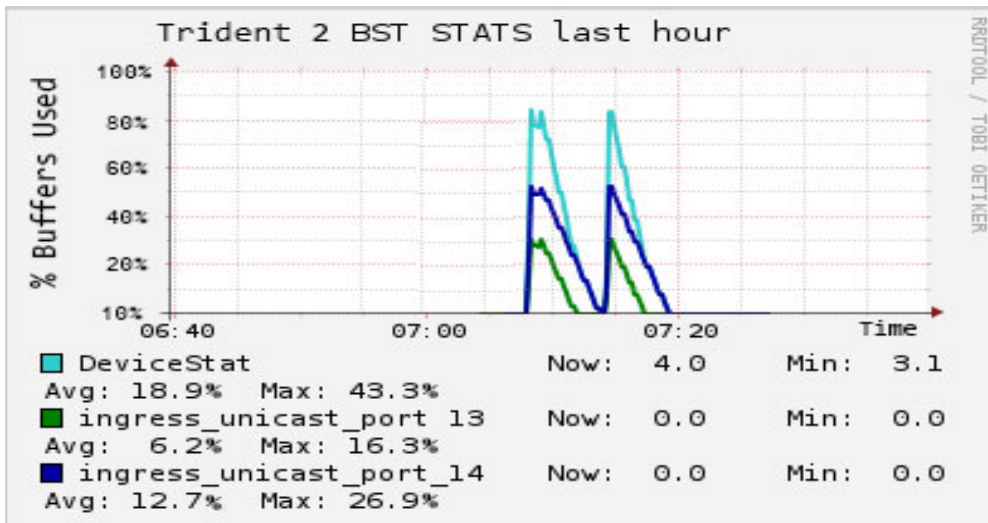
DEMO COMPONENTS

- A server node on an x86 host running Ganglia software
- A Broadcom BCM56850-based switch node running the BroadView BST application over the SDK
- A traffic generator/sink such as Ixia or Spirent (three 10G/40G fiber ports)

DEMO FLOW

- Traffic is set up and the switch is powered up running Linux. Relevant BroadView code is downloaded on the switch and initialization steps are taken. The two links connected to switch B are configured to run at 40% traffic.
- A third link is configured to send a 30% traffic in a burst, causing congestion.
- The BroadView agent is started on the switch.
- Ganglia is started on the x86 server.
- The buffers on the ingress port on the switch capture the details using BroadView (for a port for a priority group).
- BroadView can be extended to display combo counters; for example, the counters for a port for all priority groups.
- For more information, refer to the BroadView Ganglia solution application note.

Buffer Statistics Tracking Demo Application



ABOUT BROADCOM

Broadcom (NASDAQ: AVGO) is a diversified global semiconductor leader built on 50 years of innovation, collaboration and engineering excellence. Broadcom's extensive product portfolio serves multiple applications within four primary end markets: wired infrastructure, wireless communications, enterprise storage and industrial & others. Broadcom is changing the world by Connecting everything®. For more information, go to www.broadcom.com.

BroadView-TB201-R • March 22, 2016

