Broadcom Blogs

Tomahawk 3 performance vs. Tolly's Commodity switch

Posted: October 21, 2021

Recently, Tolly Enterprises released a report titled "NVIDIA Spectrum-3 vs. Switch Built with Commodity Silicon." The stated intent of this report is to compare the performance of the Spectrum-3 switch with that of a specific commodity switch. The commodity switch and the settings used by Spectrum-3 are unspecified.

The report claims three main benefits of Spectrum-3 compared to the commodity switch: equitable distribution of bandwidth across multiple servers for equal and different sizes of packets; cloud tenant protection from noisy neighbors; and eight times better burst absorption.

Broadcom repeated the tests in the report with the standard recommended settings to determine how its switch chip stacks against the commodity switch mentioned in the report.

Broadcom sampled Tomahawk 3, the world’s first 12.8T switch with 50G SerDes, in 2018. This switch has been shipped by multiple OEMs and deployed at major cloud vendors since 2019. Broadcom subsequently released three additional 12.8T switches — Trident4-X11 with compiler-based programmability, Tomahawk4-12.8T with 100G SerDes and Tomahawk4-50G-12.8T with 50G SerDes. In addition, Broadcom is shipping in production the industry's first and only 25.6T switches, the Tomahawk 4 family (with 512 x 50G SerDes or 256 x 100G SerDes).

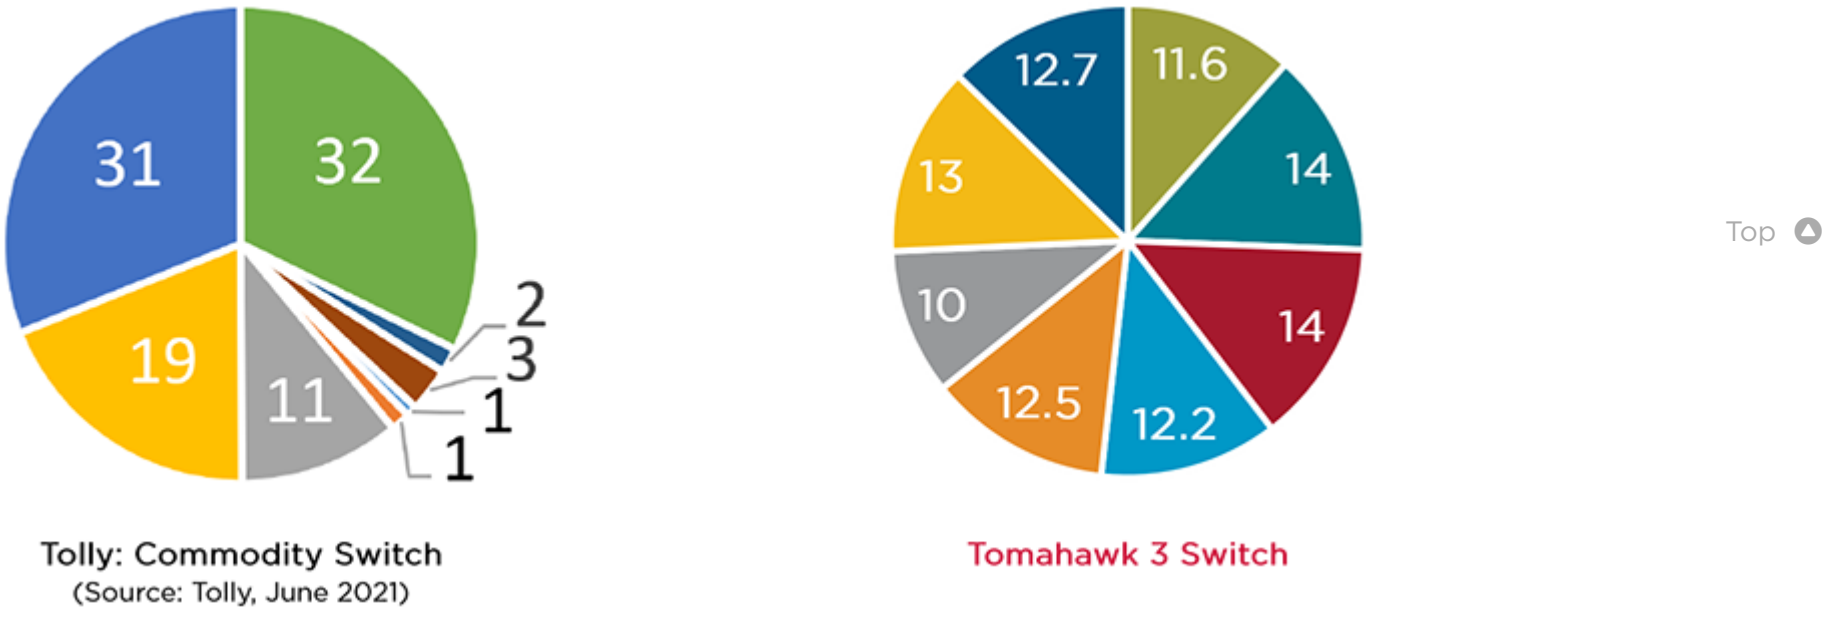

Top

Figure 1: Fairness across ports — 8 x 100G sending to 1 x 100G

We chose Tomahawk 3 for our analysis and conducted some of the tests in the Tolly report to determine Tomahawk 3's performance. Figure 1 shows a test of fairness across multiple ingress ports. In this test, eight ingress ports send large packets to one over-subscribed egress port. The commodity switch exhibits an unequal share of bandwidth for different source ports. Tomahawk 3, on the other hand, shows a relatively equal distribution of bandwidth across all the source ports.

It is important to note, however, that this fairness test is very synthetic in nature since the test equipment sends the same size packets to all ingress ports in a perfectly synchronized manner. This synthetic traffic is not typical of any real-world workload. In a real-world scenario some smaller size packets, such as control traffic, will be interspersed with the larger packets. To reflect this scenario, we also tested Tomahawk 3 with 99.9 percent of the traffic being 9KB packets and 0.1 percent being small-sized packets. Figure 2 shows the result of this test: Tomahawk 3 once again provides an equal distribution of traffic when eight ports send traffic continuously to one egress port.

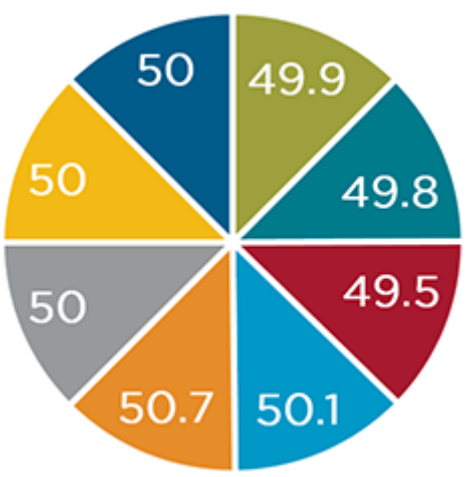

Figure 2: Tomahawk 3 fairness with mix of packet sizes — 8 x 400G sending to 1 x 400G

Another test in the Tolly report involves sending packets with mixed sizes continuously from two ports to a single over-subscribed egress port. It was reported that the commodity switch has unpredictable performance with uneven sharing among the traffic sources. Tomahawk 3, in comparison, exhibits perfectly fair sharing across all ports regardless of the frame sizes or traffic pattern, while Spectrum3 exhibits greater than 20 percent unfairness.

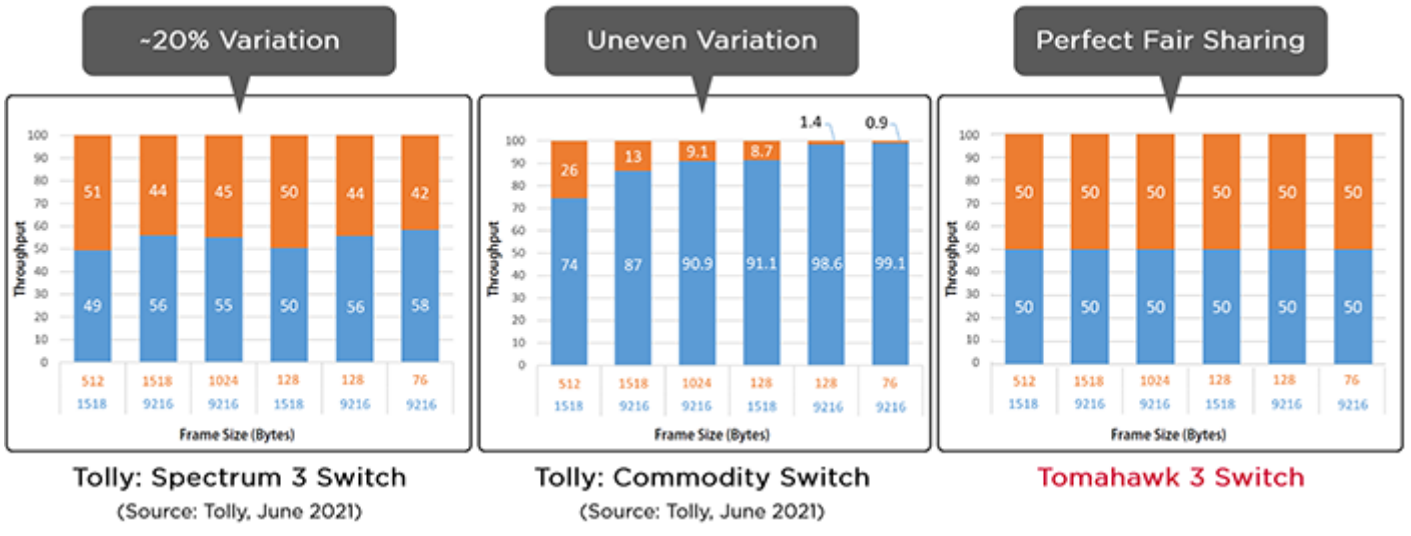

Figure 3: Fairness across packet sizes

Click to expand image.

Finally, we ran a microburst absorption test to determine how many packets are stored before they are dropped when an egress port is congested. In a typical incast scenario, traffic arrives from multiple input ports. It was reported by Tolly that the commodity switch showed a minimal buffer for burst absorption. Broadcom's Tomahawk 3 switch, in comparison, implements a shared-buffer architecture with dynamic sharing and provides 62MB of burst absorption. It allows for egress port burst absorption, while maintaining fairness across ports when there is loading on all the ports. Note also that Tomahawk 3 implements a rich set of admission control and scheduler policies that can be tailored to specific customer deployments to provide optimal performance.

When considering switches from a particular chip vendor, it is important to consider the breadth of the vendor’s roadmap as well as the speed and predictability of their execution. Broadcom provides the industry’s leading switch portfolio with chips up to 25Tbps and with products optimized for multiple networking use cases. For example, the Tomahawk family is optimized and deployed at scale in hyperscale datacenters for compute and ML/AI applications.

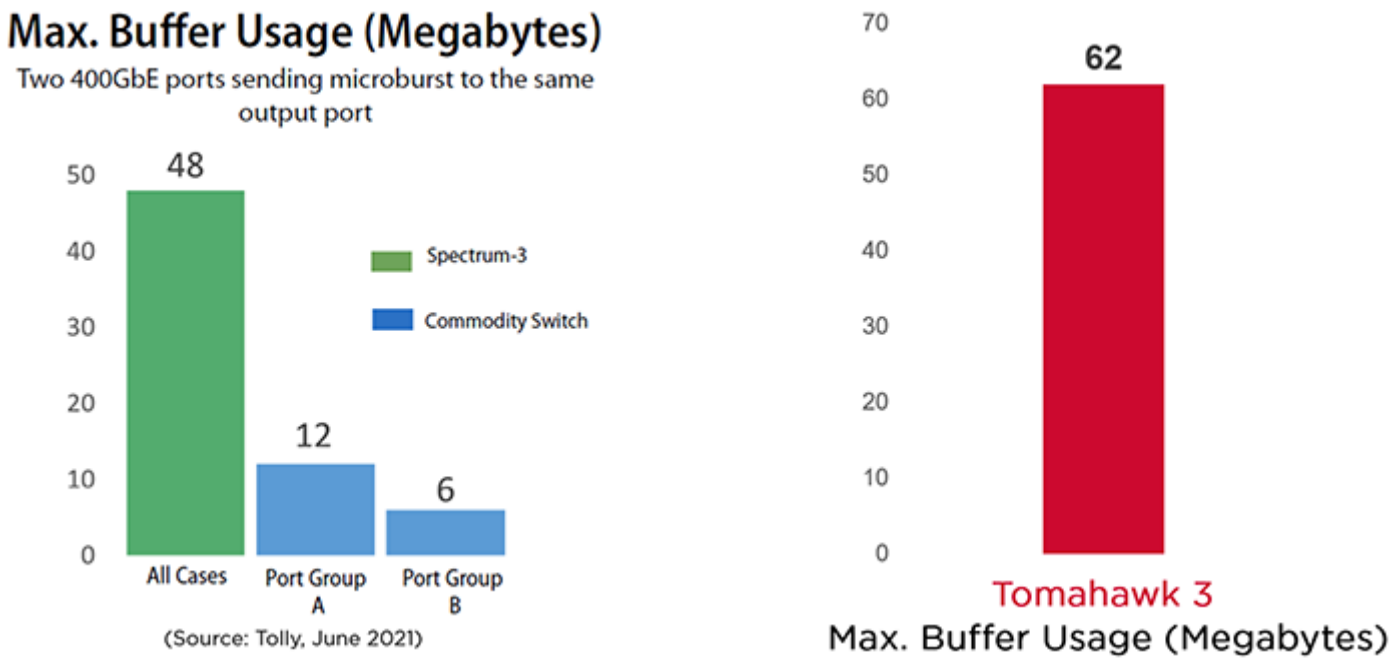

Figure 4: Microburst absorption test

Broadcom is a firm believer and a leader in deploying ML/AI networking using Ethernet/IP fabrics instead of proprietary protocols such as Infiniband. Such proprietary technologies are not scalable (can only be deployed in small islands) and, lacking a healthy ecosystem, are very expensive. In comparison, Broadcom’s Ethernet/IP switches offer the industry’s highest performance, configurability and programmability while leveraging the vibrant Ethernet ecosystem.

Finally, we invite customers to discuss their unique network needs with Broadcom, so we can provide the optimal solution (product and configuration) for your application.

Learn more

Access the Buyer's Guide and explore the entire merchant silicon line: Merchant Silicon Networking Chips

Connect

Follow us on Twitter and LinkedIn for real-time information on Broadcom's industry-leading silicon switch and router solutions.

Author

Recent Posts

Share

Related Posts

Packet fronthaul: Linchpin of the ORAN 5G network

The introduction of 5G has transformed the fronthaul segment of the Radio Access Network (RAN) from a short, simple link connecting a radio and a base station into a disaggregated, heterogeneous netwo...

Ethernet switching solutions at Broadcom

An industry leader, Broadcom is a premier fabless semiconductor company. We focus on the communications industry, developing silicon for a wide-range of products that keep the globe connected, includi...

The Butterfly Effect: Reducing the carbon footprint in the world

Broadcom develops three distinct switch product lines to address distinctly different customer needs -- the enterprise, service provider and data center markets. Standard features are leveraged across...

See more posts related to Broadcom Blogs