



## Broadcom Ships Tomahawk 5, Industry's Highest Bandwidth Switch Chip to Accelerate AI/ML Workloads

August 16, 2022

### World's First 51.2 Tbps Ethernet Switch Chip

SAN JOSE, Calif., Aug. 16, 2022 (GLOBE NEWSWIRE) -- [Broadcom Inc.](#) (NASDAQ:AVGO) announced today that it has delivered the StrataXGS® Tomahawk® 5 switch series, providing 51.2 Terabits/sec of Ethernet switching capacity in a single, monolithic device, double the bandwidth of any other switch silicon available on the market today.

"Delivering the world's first 51.2 Tbps switch two years after we released Tomahawk 4, the industry's first 25 Tbps switch, is a testament to the outstanding execution and innovation by the Broadcom team," said Ram Velaga, senior vice president and general manager, Core Switching Group, Broadcom. "Since the introduction of Tomahawk 1 in 2014, Broadcom has consistently executed on doubling the bandwidth approximately every two years. With today's introduction of the fifth generation Tomahawk family, we are proud to say that a single Tomahawk 5 replaces forty-eight Tomahawk 1 switches in the network, resulting in over 95 percent reduction in power requirements. We applaud our customers, partners, and engineers for making this possible."

While data centers continue to experience dramatic growth in network bandwidth requirements, there is also strong motivation to unify the networking infrastructure for general-purpose compute and storage with that of AI/ML compute. AI/ML training clusters are driving the need for fabrics with high-bandwidth connectivity, high radix, and lower job completion time, while operating at high network utilization.

Ethernet offers the best solution for unified network infrastructure, providing the lowest power, highest bandwidth, highest radix, and fastest SerDes speeds, along with a predictable doubling of bandwidth every 18 to 24 months. These benefits combined with its large and vibrant ecosystem, Ethernet provides the highest performance interconnect per Watt and per dollar for AI/ML and cloud scale infrastructure.

To enable the next generation of unified networks, Broadcom is now offering the Tomahawk 5 family. Critical to enabling efficient use of the massively shared infrastructure in large data centers, Tomahawk 5 provides AI/ML workload virtualization with features such as single-pass VxLAN routing and bridging. Critical to minimizing job completion time (JCT) for AI/ML workloads, Tomahawk 5 offers features such as Broadcom Cognitive Routing, advanced shared packet buffering, programmable inband telemetry, and hardware-based link failover.

Tomahawk 5's Cognitive Routing improves network link utilization by automatically and dynamically selecting the most lightly loaded links in the system for every flow that transverses the switch. This is especially important for AI/ML workloads which often have a combination of short-lived mice flows and long-lived, high bandwidth elephant flows with low entropy. Tomahawk 5 includes real-time dynamic load balancing that tracks the utilization of all links, both at the switch and downstream in the network, to determine the optimal path for each flow. It also monitors the health of links in hardware and automatically steers traffic away from failed links. These features provide dramatically improved network utilization and reduced congestion, resulting in shortened JCT.

Also important in improving JCT is minimizing network congestion by controlling the rate of traffic injected into the network by each source. Since network operators employ a variety of different congestion control algorithms at their endpoints (such as merchant or custom NICs), Tomahawk 5 provides extensive programmable inband telemetry on both live traffic and network probes. Real-time metadata can be inserted into traffic at line rate as it traverses the network to collect telemetry on queue size, packet latency, switch utilization, and a variety of other customer-selectable metrics. This metadata can be used for precise end-to-end network congestion control.

To enable the lowest power and lowest cost for physical connectivity, Tomahawk 5 enables a direct 100G PAM4 interface to direct attach copper (DAC), front panel pluggable optics, and co-packaged optics. The flexible, long reach Tomahawk 5 SerDes provides DAC connectivity to all devices within a rack, and even between racks, without the need for retimers or other active components. It can also interface directly to a broad ecosystem of standard front-panel pluggable optical modules.

Additionally, leveraging Broadcom's leading-edge silicon photonics and packaging technologies, Tomahawk 5 will be made available with co-packaged optics using Broadcom's Silicon Photonics Chiplets in Package (SCIP) platform, providing more than 50 percent decrease in the power needed for optical connectivity. Since the same switch silicon provides all these options, customers can choose the optimal I/O for each part of their intra-cluster, inter-cluster, and inter-DC networks with no software porting required.

#### StrataXGS Tomahawk 5 Series Key Benefits:

- Enables the next generation of unified data center infrastructure with 64 ports of 800GbE switching and routing.
- Virtualization of general compute and AI/ML workloads with single-pass VxLAN routing and bridging.
- Unparalleled physical I/O options using 512 instances of the industry's highest performance, most flexible, and longest-reach 100G PAM4 SerDes.
- High-precision PTP and SyncE time synchronization.
- Six on-chip ARM processors for high-bandwidth, fully-programmable streaming telemetry, and sophisticated embedded applications such as on-chip statistics summarization.
- Unmatched power efficiency, implemented as a monolithic 5nm die.

"Tail latency is the critical network performance metric for distributed AI/ML training," said Bob Wheeler, principal analyst at Wheeler's Network. "Broadcom recognized the limitations of traditional hash-based load balancing for these workloads and added Cognitive Routing with dynamic flowlet steering to Tomahawk 5. Hyperscale operators can now unify their network fabrics, eliminating specialized interconnects dedicated to only training

clusters.”

Compared to general compute and storage, AI/ML training clusters have unique communication patterns. To minimize job completion time, Tomahawk 5 adds specific features for these workloads and network topologies.

#### **StrataXGS Tomahawk 5 Series Key Features for AI/ML:**

- World's highest radix of 200GbE ports: 256 ports supported on a single chip, enabling flat, low latency AI/ML clusters.
- The industry's most advanced 51.2 Tbps shared-buffer architecture, providing the highest performance and lowest tail latency for RoCEv2 and other new RDMA protocols.
- Advanced Broadcom Cognitive Routing, dynamic load balancing, and support for end-to-end congestion control capabilities specifically designed to handle the large, low entropy flows typical of AI/ML workloads.
- Support for Clos and non-Clos topologies such as torus, Dragonfly, Dragonfly+, and Megafly.
- Hardware-based link failover for improved network resiliency and reduced JCT.

Along with the Trident and Jericho switch families, the Tomahawk series is part of Broadcom's three-pronged strategy of providing optimized switch architectures for different network applications. All these devices share a common programming interface, so customers can easily leverage their software development efforts across different platforms.

Having a strong commitment to open networking, Broadcom has provided both SAI (Switch Abstraction Interface) and Broadcom SDK open APIs for all five generations of the Tomahawk family. Broadcom is one of the industry's largest contributors to SAI and the SONiC network operating system. To accelerate time to deployment, support for SAI and Broadcom SDK are provided on Tomahawk 5 silicon as well as a comprehensive suite of network and device simulation tools.

Broadcom Tomahawk 5, BCM78900, is available today for customers worldwide. For more information on Tomahawk 5 click [here](#).

#### **About Broadcom**

Broadcom Inc. (NASDAQ: AVGO), a Delaware corporation headquartered in San Jose, CA, is a global technology leader that designs, develops and supplies a broad range of semiconductor and infrastructure software solutions. Broadcom's category-leading product portfolio serves critical markets including data center, networking, enterprise software, broadband, wireless, storage and industrial. Our solutions include data center networking and storage, enterprise, mainframe and cyber security software focused on automation, monitoring and security, smartphone components, telecoms and factory automation. For more information, go to <https://www.broadcom.com>.

Broadcom, the pulse logo, Trident, Tomahawk and Connecting everything are among the trademarks of Broadcom. The term "Broadcom" refers to Broadcom Inc., and/or its subsidiaries. Other trademarks are the property of their respective owners.

Press Contact:

Jon Piazza

Corporate Communications

[press.relations@broadcom.com](mailto:press.relations@broadcom.com)

Telephone: +1 310 498 5254

#### **Industry Quotes:**

##### **Dennis Cai, Head of Network Infrastructure, Alibaba Cloud Intelligence**

“Alibaba has deployed multiple generations of Tomahawk-based platforms throughout our network infrastructure. Once again, Broadcom has delivered another robust product in Tomahawk 5, with unmatched bandwidth, radix, power efficiency, and product quality. We are proud to collaborate with Broadcom on Tomahawk 5, leveraging some cutting-edge features such as programmable inband telemetry to build our predictable super performance cluster, and to service the exponential growth of our AI/ML workload and high performance storage requirements.”

##### **Igal Elbaz, SVP and Network CTO, AT&T**

“AT&T's leadership position in the rollout of the 5G infrastructure requires strategic collaborators like Broadcom to provide high-quality merchant silicon products to support our ongoing efforts in developing a robust open and disaggregated white box ecosystem. These time-to-market innovative solutions can also be deployed in any node throughout our network. AT&T applauds Broadcom for delivering yet another industry first in the single-chip Tomahawk 5 silicon.”

##### **Gang Cheng, Distinguished Engineer and Head of Network Architecture, Baidu**

“Broadcom's latest Tomahawk 5 product is an impressive technical achievement. We are excited to see the continued growth of the merchant silicon ecosystem, so that Baidu can continue to build a world-class network infrastructure.”

##### **Lark Wang, General Manager of System Department, ByteDance**

“The growth of our short video platforms over the past few years demonstrates the importance of having a scalable, high bandwidth, high performance and efficient network in data centers. Tomahawk 5 will continue to fulfill these metrics, and ByteDance is thrilled to see the readiness of Tomahawk 5 in the merchant silicon ecosystem, which will enable ByteDance to develop its new generation switches and network infrastructure.”

##### **Dave Maltz, Technical Fellow and Corporate Vice President, Microsoft Azure Networking**

“At Microsoft, we believe that merchant silicon, hardware independence, and open-source protocol and management stack is the optimum solution for running converged hyperscale data center workloads including compute, storage, and AI/ML. We congratulate Broadcom's continued contribution to this vision with the milestone achievement to sample Tomahawk 5 at 51.2 Tbps with robust support for SAI on both silicon and simulation models. Broadcom is one of the top contributors to SONiC and SAI open-source initiatives, with Tomahawk 5 representing the fifth generation of the Tomahawk family to support SAI and SONiC.”

##### **Steve Scott, Technical Fellow and Corporate Vice President, Microsoft Azure Hardware Architecture**

“Microsoft Azure offers best-in-class infrastructure for HPC and AI/ML workloads of all sizes. Broadcom's Tomahawk 5 provides important benefits for

next generation HPC and AI/ML requirements. Tomahawk 5 advances the scale, telemetry, and advanced features needed to support HPC and AI/ML network requirements, within the framework of the open and innovative Ethernet ecosystem.”

**Sage Zou, Vice President, Tencent Cloud**

“As Tencent continues to scale our global infrastructure, it is critical for us to partner with Broadcom for ground-breaking products such as Tomahawk 5. We are glad to see that Broadcom will offer both the DAC and the CPO connectivity. Undoubtedly, Tomahawk 5 will be the cornerstone of our next generation of low-cost and high-performance cloud network.”

**Jim Keller, President and CTO, Tenstorrent**

“I’m really excited about next generation high bandwidth switches. AI workloads move more data than any other workload I’ve designed for. Great to see progress on performance, QOS and power. We’re going to use all of it.”

**Anshul Sadana, Chief Operating Officer & Senior Vice President, Arista Networks**

“The Arista 7060X and modular 7388X Series running EOS provide a compelling range of high radix solutions for Hyper-Scale Cloud data centers, powered by multiple generations of the Broadcom Tomahawk architecture. Tomahawk 5 innovations enable 800G ready, power efficient, and high radix Ethernet solutions for Cloud scale compute and storage networks and converged AI/ML.”

**Steven Yoe, Co-President, Chief Technology Officer, H3C**

“As a long-term strategic partner, H3C is very excited about Broadcom’s launch of the market-leading Tomahawk 5 switch series. This chip family will change hyperscale data center network infrastructure and implementation by enabling the use of 5nm technology, 400GbE+ ports, and more importantly, introducing the world’s highest performance Ethernet switch ASIC at 51.2 Tbps. We are looking forward to using Tomahawk 5, with its significant innovations, to bring next-generation breakthroughs to H3C’s DC switch portfolio. This will help us rapidly deploy leading network solutions to meet the demanding growth of hyperscale cloud data centers.”

**Michael Bushong, Group Vice President, Juniper Networks**

“Broadcom’s impressive achievement of doubling the bandwidth for data center platforms on a single chip with Tomahawk 5 allows Juniper Networks to keep providing an extensive routing and switching portfolio with unmatched power efficiency. Combining best-of-breed silicon technology with Junos OS, the industry’s most advanced network operating system, addresses the leading-edge requirements of our customers.”

**Ken Kutzler, Vice President Hardware Engineering, Nokia**

“Broadcom’s Tomahawk 5 64x800G switch becomes a valuable tool to help migrate our data center customers into the 800GE era. 800G optics are generally available this year and their use makes sense right away from both a commercial and power perspective compared to 400GE.”

**Edgar Masri, CEO, Accton Technology Corporation**

“Accton is excited to see Broadcom continue to lead the industry and deliver the first 51.2 Tbps switching ASIC. Accton, and its subsidiary Edgecore Networks, see first-hand the strong interest from hyperscale operators and we are proud to lead the charge in developing world-class 400GbE+ platforms based on Tomahawk 5 for our valued hyperscale customers.”

**Daniel Huang, General Manager, Alpha Networks Inc**

“Broadcom has once again established market leadership by delivering the world’s highest performing single-chip Ethernet switch in Tomahawk 5, optimized specifically for Cloud AI/ML workloads. Alpha is excited to work with Broadcom and our customers to deliver state-of-the-art 400GbE and 800GbE platforms for our web-scale and OEM customers worldwide with best in class economics and unparalleled product quality and reliability.”

**Steven Dorwart, Vice President, Service Provider Solutions, Celestica Inc.**

“Celestica is proud to continue our close collaboration with Broadcom in developing a rich portfolio of open networking platforms based on the Tomahawk series. Tomahawk 5 introduces not just exceptional speeds and feeds, it also incorporates new features that greatly help with AI/ML workloads running in the largest data centers throughout the world. We look forward to offering a new wave of Tomahawk 5-based platforms that will enable our customers to grow their business.”

**Victor Cheng, Sr. VP & GM of Delta ICTBG**

“Delta congratulates Broadcom in delivering yet another game-changing product. The power and cost-per-gigabit savings Tomahawk 5 enables will usher in a new era of 400GbE+ networking systems that allows network operators unparalleled capability to process AI/ML workloads with innovative features like adaptive routing. Delta looks forward to leveraging Tomahawk 5 to bring market-leading 400GbE+ networking solutions to market.”

**Jack Tsai, SVP and General Manager, Inventec Corporation**

“With the launch of the 51.2 Tbps Tomahawk 5 silicon, Broadcom continues to lead the industry with flawless execution, technology innovation, and merchant silicon ecosystem enablement. We look forward to offering a new portfolio of platforms based on Tomahawk 5.”

**Liu Zhongdong, President and CEO, Ruijie Networks**

“It is stimulating that Tomahawk 5 doubles the switching capacity again in a monolithic 51.2 Tbps chip! By integrating the industry leading 100G SerDes and advanced switching/telemetry features, Tomahawk 5 is fully ready for the 400G/800G cloud-scale networks. As the leading and dedicated switch solution vendor, Ruijie has enriched our product portfolio with Tomahawk 5. We can now support China and global customers to build the next generation compute/storage/AI&ML converged networks.”

**Mike Yang, Senior Vice President of Quanta Computer Inc. and President of Quanta Cloud Technology**

“As a leading data center infrastructure solution provider, QCT is honored to collaborate with Broadcom to showcase the innovations of the Tomahawk 5. QCT is thrilled to develop both the DAC and CPO version of Tomahawk 5 to service our customer requirements.”

**William Lin, CEO, Wistron**

“Wistron is excited to work with Broadcom to deliver Tomahawk 5 based switching products to market. Tomahawk 5 will provide unprecedented density, performance, and efficiency for 400GbE+ network platforms. We look forward to delivering new Tomahawk 5 platforms to enable the ever increasing need for more bandwidth from complex AI/ML workloads.”

**Vincent Ho, CEO, UfiSpace**

“As Broadcom’s premier networking ODM provider, UfiSpace is focused on delivering the best-of-breed products in Ethernet switches for fixed and

modular based systems. We are thrilled to see the availability of Tomahawk 5, as we expect its new features will enable our customers to take their next generation of services to unprecedented heights.”



Source: Broadcom Inc.