The Futurum Group

Lab Insight

A Foundation for Scalable Infrastructure: Dell + Broadcom

AUTHOR **Russ Fellows**
Head of Futurum Labs | The Futurum Group

JULY 2024

Executive Summary

Business leaders are increasingly focused on finding ways to improve existing operational efficiency, and support new cloud and AI applications. These operational imperatives must be met while simultaneously ensuring data privacy and governance is maintained. These challenges have led many organizations to explore ways to build a hybrid cloud that can leverage both public and private cloud resources.

An important component of any IT environment is the networking technology utilized, as the network quite literally serves as the glue that holds these disparate resources together. Increasingly, Ethernet has become the fabric of choice, due to its ubiquity and ability to be utilized for server, storage and now AI environments across the hybrid multi-cloud resources.

The use of proprietary networking, specifically Fibre Channel and InfiniBand is diminishing, due to the added cost, complexity and lack of access to these networks across hybrid cloud environments. In addition to AI, the biggest concern many businesses have is security and cyber resiliency. With cyber-attacks an increasing threat experienced by many organizations, ensuring IT infrastructure is secure is a critical consideration.

The Futurum Group was asked to evaluate Dell's networking options, including the security, management features, and integration with existing and new networking requirements. This included Dell 16th Generation PowerEdge servers with Broadcom 57508 line of Ethernet Network Interface Card (NIC) cards utilizing both Futurum Group and Dell lab facilities.

Summary of Study

The Futurum Group has evaluated the impact of network speed on storage, applications and evaluated the security and cyber resiliency features of Dell Servers with Broadcom NICs.

In Table 1 below we summarize the storage and application networking requirements of a variety of application use cases. The network rates are based on the use of current generation Intel or AMD Dell PowerEdge systems using high CPU core counts, large memory configurations together with virtualization or container native workloads, accessing storage with multiple NVMe drives per node.

HCI Application Requirements / Network Requirements	Storage Networking	Application Networking	Total HCI Networking
Typical VDI Workloads	Moderate: 10 – 25 GbE	Moderate: 10+ GbE	25 GbE minimum
Generalized Virtual Applications	High: 10 – 25 GbE	Moderate: 10+ GbE, low Latency	25 – 50 GbE recommended
Database Clusters – Including Data Protection	High: 25+ GbE minimum	Moderate – 25+ GbE, low latency	50 - 100 GbE recommended
K8s Container Environments	High: 25 GbE minimum	Moderate: Low Latency	25 - 100 GbE or greater
AI/ML Training	Very High: 100 GbE or greater	Very High: 100 GbE or greater	100 GbE + (app dependent)

Table 1: Application Workload Networking Requirements (Source: The Futurum Group)

AI Workloads

As shown in Table 1, AI can be a highly demanding workload for many resources, including network connectivity. Dell is actively working to help firms create and place new AI applications into deployment. Dell together with their partners have developed a set of proof-of-concept AI examples that enable fine-tuning and then inferencing foundational generative AI on Dell PowerEdge servers. The scale-out architecture also leverages Broadcom based 100 and 200 Gb Ethernet NICs along with Dell high-speed ethernet switches and storage to deliver the benefits of AI tools utilizing private clouds.

In particular, fine-tuning large language models (LLMs) can require a distributed infrastructure involving multiple systems with GPUs. A critical aspect of this architecture is the high-speed Ethernet interconnectivity, requiring 100 Gb connectivity in order to eliminate bottlenecks. Other AI workloads that utilize scale out approaches can also be challenging, such as those using Retrieval Augmented Generation (RAG), or distributed inferencing. These architectures require high speed and low-latency connectivity between systems and storage leverage a high-speed ethernet fabric.

Ethernet for Modern IT Infrastructure

An important consideration for IT and other support staff tasked with operating business applications, is how easily the hardware and software interoperate with existing applications and infrastructure. Enterprises have been migrating away from proprietary networks such as Fibre Channel and InfiniBand and instead deploying a common Ethernet fabric. With the advent of high-speed and low-latency Ethernet technology, companies now have the ability to utilize this common technology for traditional server connectivity, along with storage and increasingly for AI workload interconnectivity.

Futurum Group Comment: *Todays data-center networks must support multiple workloads, including server connectivity, virtualization, along with storage and AI applications. Using the latest generation of high-speed Ethernet from Dell and Broadcom can serve as a common connectivity, helping eliminate proprietary networking to enhance IT and corporate efficiency.*

Broadcom BCM-57508 Ethernet cards are an important aspect of the solution, solving a common bottleneck with distributed systems, the inter-node communications, with both bandwidth and latency as key factors. Broadcom's Peer Direct and GPUDirect RDMA technologies enable data to bypass host CPU and memory, for direct transfer from the network into GPUs and other hardware accelerators. Without these technologies, data is driven by the CPU into local memory and then copied into the accelerator's memory – adding to latency.

Unstructured data access via file and object protocols over Ethernet connectivity is a predominant access method utilized for AI workloads for training and inferencing data sets. Nearly all of Dell's enterprise storage systems support ethernet connectivity, and provide a variety of options that leverage high-speed, low-latency Ethernet fabrics including NFS over RDMA, NVMe over TCP, GPU Direct Storage over RDMA, and RoCEv2 among others.

Dell / Broadcom 57508 NIC

The Dell - Broadcom BCM57508 supports multiple connectivity speeds up to 200 Gb/s and offers multiple standards including RDMA over Ethernet (RoCEv2) with hardware-based congestion control. Broadcom's TruFlow helps to speed and prioritize application specific I/O patterns. This can help to increase application performance, with security based upon Broadcom's hardware secure boot (ROT) and device attestation.

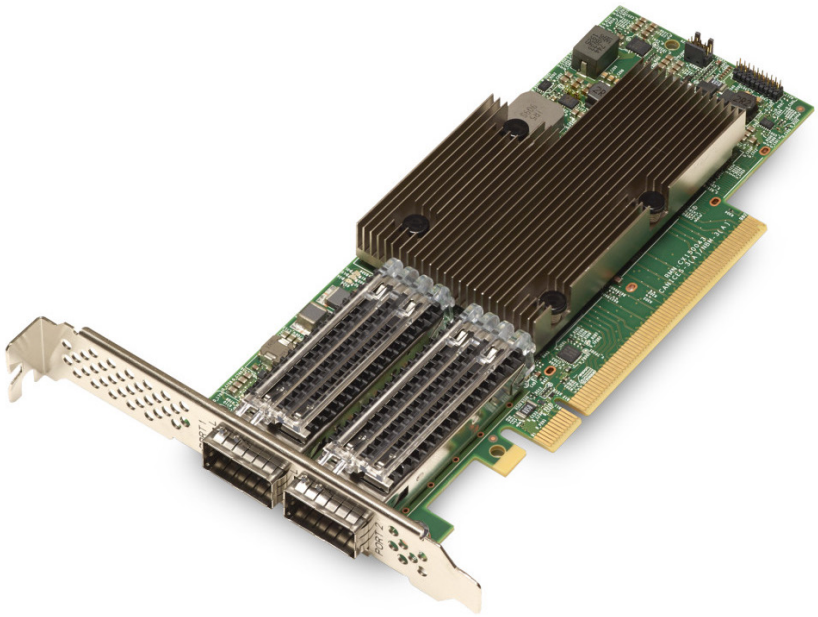

Figure 1: Dell – Broadcom 57508 2x 100 Gb NIC (Source: Broadcom)

Security Features of Dell and Broadcom

System security requires a layered approach, with hardware security features and capabilities utilized by the firmware, that provides a secure foundation for the operating system (OS) and applications running within the OS. The hardware root of trust provides several foundational security capabilities including a secure, tamper-proof key storage location, a unique hardware encryption key, and other facilities required as part of a cryptographic framework.

The latest generation of Broadcom based Dell RAID controllers, NICs and switches now utilize Broadcom's 4th generation hardware secure boot per NIST SP800-193 guidelines. Broadcom also implements the RSA Public Key Cryptography standard (PKCS) digital signature approved by FIPS. This results in the firmware being signed, with the signature itself encrypted, which can then be verified before being used to ensure the integrity of the firmware image. Also, Broadcom uses a secure, tamper-proof FIPS-2 compliant hardware security module to perform the firmware signing. As a result, Dell suplied Broadcom NICs utilize Dell cryptographic keys, thus restricting boot only upon verification of firmware that is digitally signed and authenticated by Dell.

Each aspect of a system must be secured, ideally with each element utilizing industry best practices including Zero Trust principles, enabled via a silicon Root of Trust (RoT) as the foundation of physical security. Utilizing a certified RoT device that provides tamper resistance, the hardware protects the device firmware, which in turn is utilized to ensure critical software security features are verifiably operating securely. Security breach examples could include the exchange of entire PCIe cards, uploading compromised firmware or other attacks that target PCIe cards.

Delivering end-to-end security requires the entire solution to utilizing multiple hardware RoT devices, together with software features that leverage the underlying hardware capabilities. Dell's PowerEdge servers together with Broadcom NIC controllers provide this capability by validating each layer, from hardware to firmware and the OS using cryptographic attestation. Dell's iDRAC and other management tools leverage this integrated security environment to deliver a secure, easy to manage solution.

Futurum Group Comment: *A key aspect of the Zero Trust approach is to utilize hardware RoT devices to provide cryptographic verification to build a secure ecosystem. Using Zero Trust, each component first verifies, then trusts their counterpart based upon certified key exchanges and mutual authentication, attestation and authorization. These features are critical for the additional security mechanisms used throughout the Dell servers along with supported devices.*

As part of a standards-based approach to secure system management, the Security Protocol and Data Model (SPDM) standard leverages Zero Trust principles to secure communications over a variety of internal, server transport connections. Dell's 16th Generation servers, along with Broadcom 57508 NICs utilize technologies including public key encryption, and cryptographic signing of certificates to provide device attestation between Dell's servers and Broadcom devices.

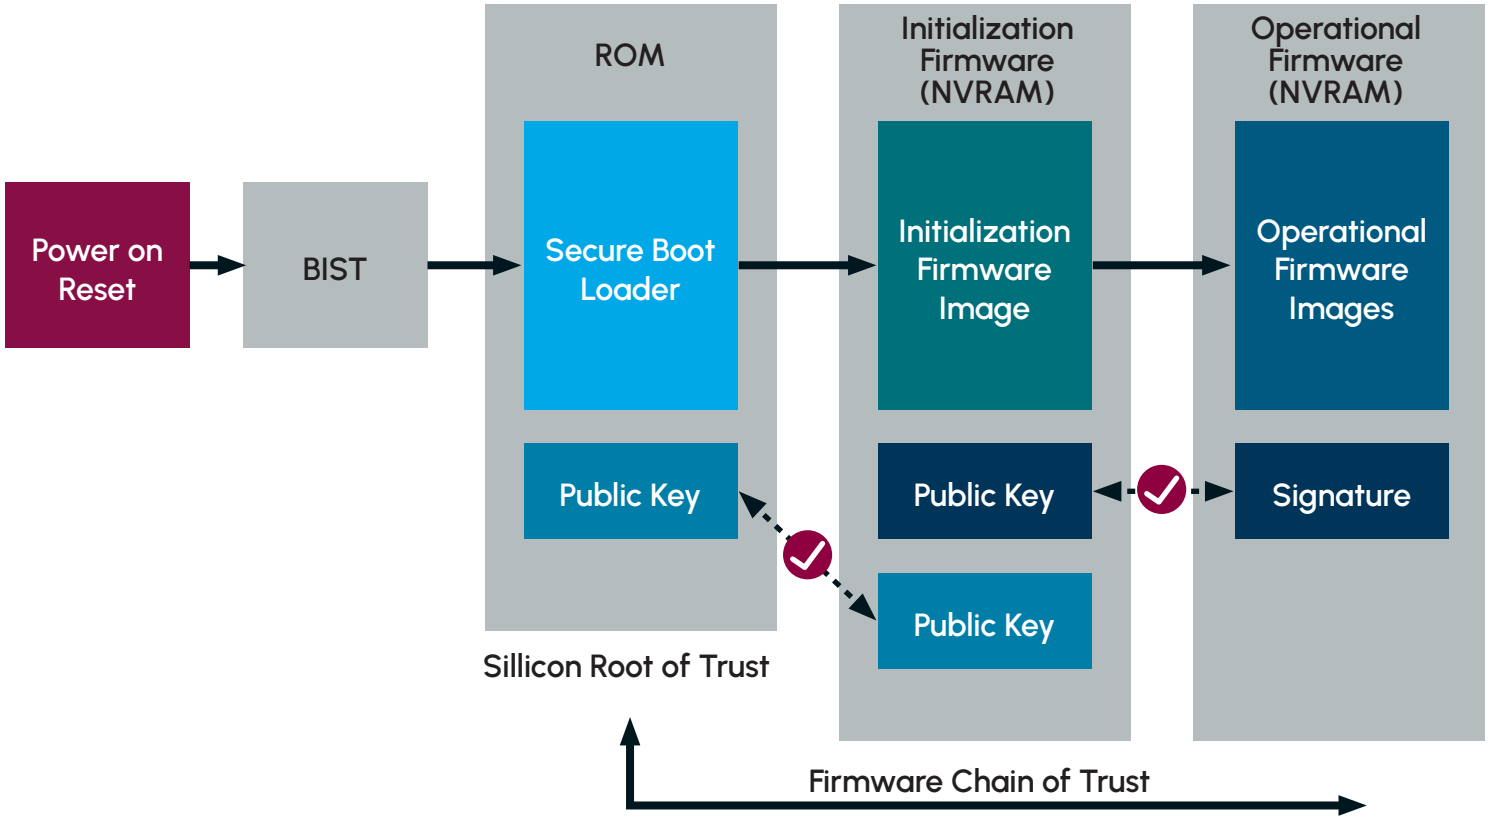

Figure 3: Broadcom's Silicon RoT w/ Secure Firmware Loading (Source: Broadcom)

The Broadcom TruTrustTM technology is capable of secure boot meaning it only executes boot images authenticated by the secure boot loader (SBL). Secure boot functionality is the basis for Zero as it is the root of trust from which all subsequent applications are run. The secure boot capability provides the following functionality:

- Secure Boot: The Secure Boot Loader (SBL) resides in the device read only memory and cannot be modified. This results in a trusted state to then load authenticated images.
- SBL Validation: The SBL cryptographically validates the integrity of the Secure Boot Image before it is executed to ensure that it has not been tampered with maliciously or errantly.
- Boot Image Confidentiality: Dell customizes the image such that only Dell signed images execute on the device.
- Dells Customization provides:
    - Dell security policies for key management to ensure that only signed firmware is authorized for loading onto the Dell / Broadcom NIC.
    - Only code signed by Dell runs on the Dell / Broadcom NICs device. This ensures that device code cannot be tampered with in the field and verifies the authenticity of the image.

Dell PowerEdge 16th Generation Servers

Multiple areas of evaluation and testing were performed with the PowerEdge servers together with PERC RAID cards and Broadcom NICs. One of NIST's specifications (NIST SP 800-193) specifically addresses platform firmware resiliency. It stipulates methods for securing the BIOS, boot ROMS, along with firmware and driver signatures to verify authenticity. This guideline also outlines a method for providing a "secure boot" mechanism, whereby each component verifies subsequent components in the stack from hardware all the way to the operating system.

Management, Alerting and Reporting

A critical part of Dell's cybersecurity capabilities are the management tools available. The Futurum Group Labs validated numerous functions and capabilities of Dell's enterprise management solutions, iDRAC, Open Manage Enterprise and CloudIQ which together provides management, alerting and reporting capabilities.

- Dell's iDRAC: system level management software pre-installed on all Dell PowerEdge servers, providing secure out-of-the box management capabilities for individual systems.
- OpenManage Enterprise (OME): Designed for IT staff to monitor and manage Dell servers within a datacenter. OME provides roll-up features to aggregate information while still providing the ability to manage individual systems, either from within OME or via linking to iDRAC.
- CloudIQ: Enables multi-site and enterprise-wide monitoring and management, with high-level management and reporting while also providing drill-down system management of some features and function.

Final Thoughts

Businesses are searching for ways to improve their efficiency, by standardizing on a common set of technologies that can be leveraged across multiple applications. With companies implementing hybrid cloud as a standard operating model, it is important to maintain consistency across both public and private cloud deployments.

A new set of workloads are adding to IT operational challenges, including support for cloud native, containerized applications, along with more traditional virtualized applications and new AI based applications. The challenges include ensuring cyber resiliency through security mechanisms while implementing infrastructure that can operate both in public and private cloud settings.

Additionally, the ability to build and manage both the hardware and software infrastructure helps companies compete effectively, while balancing their corporate security concerns and ensuring their data is not compromised or released externally.

Utilizing Dell's new 16th Generation PowerEdge servers together with Broadcom NICs enables companies to implement a common fabric for connecting servers running hypervisors, containers or AI workloads. Using Ethernet as a common high-speed fabric can support AI GPU communications along with server and storage connectivity.

Dell's server, storage and switch lines can all be leveraged to provide a flexible, scalable infrastructure to support whatever application demands arise. Importantly, this can all be done while maintaining a high cyber resiliency posture, through the NIST cyber security framework.

Important Information About this Report

CONTRIBUTORS

Russ Fellows
Head of Futurum Labs | The Futurum Group

PUBLISHER

Daniel Newman
CEO | The Futurum Group

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

CONTACT INFORMATION

The Futurum Group LLC | futurumgroup.com | (833) 722-5337