

ESG Lab Review

Plexxi HCN for Hyperconverged Infrastructure

Guaranteed Network Performance at Scale

Date: July 2017 Author: Jack Poller, Senior Lab Analyst

Abstract

This ESG Lab Review documents hands-on testing of Plexxi HCN for Hyperconverged Infrastructure (HCI) with a focus on ease of use, performance, resiliency, and best practices for configuration in HCI environments.


The Challenges

With budgets flattening and end-users continuing to demand services that require greater performance and agility from the supporting infrastructure, IT organizations, from the smallest to the largest, are being told that they must do more with less. Traditional IT infrastructure, consisting of silos of dedicated servers, networking, and storage, is extremely complex and expensive to purchase and maintain, and IT must find innovative ways to create scalable solutions.

To limit capital and operational expenses, many organizations have turned to cloud-based solutions, such as infrastructure-as-a-service (IaaS) and hyperconverged infrastructures. Most IT organizations are well on their way to transitioning from traditional data center to cloud-based environments. ESG research indicates that most organizations have, or are on the cusp of having, deployed some form of private cloud, from basic internal cloud (33% of respondents) through advanced internal cloud (46%), all the way to a complete IT-as-a-service operating model (20%) (see Figure 1).¹

Figure 1. Stage of Transition to Private Cloud Computing

Currently, at what point would you say your organization is on its transition to private cloud computing? (Percent of respondents, N=308)



Source: Enterprise Strategy Group, 2017

¹ Source: ESG Research Report, [The Cloud Computing Spectrum, from Private to Hybrid](#), March 2016.


Conventionally, data centers connect application servers to data storage with isolated networks, separating storage traffic from application and client traffic. With the transition to cloud infrastructures, hyperconverged solutions incorporate storage services directly into the application server environment, blurring the lines between client, application, and storage traffic, all of which traverse the same network. Since traditional legacy Ethernet networks were not designed to handle diverse traffic loads (such as compute and storage), this puts additional burdens on the network infrastructure, to accommodate additional types of traffic and additional loads while maintaining the necessary flexibility, scalability, programmability, and resiliency for the modern data center.

The Solution: Plexxi Hyperconverged Networks

Hyperconverged infrastructure is part of the emerging modern software-defined data center (SDDC). HCI provides significant benefits, including the ability to rapidly scale storage and compute as workloads change and new services and applications are introduced into the environment. The SDDC emphasizes delivering agility and cost savings through automation, resource elasticity, data and application mobility, and workflow integration. These benefits are achieved by focusing on tools, integration, and automation, rather than individual components.

Plexxi solutions provide a modern Ethernet network fabric tuned for the needs of HCI and the SDDC. By leveraging intelligent, software-based integration into the HCI management stack and a highly dynamic meshed Ethernet fabric, Plexxi HCN provides agility, ease of use, and security for the SDDC. Plexxi Switches (based on standard merchant silicon white box Ethernet switches) are used to create this highly dynamic meshed fabric, while Plexxi Control delivers fine grained fabric control and a “workload”-based model to align network resources to IT priorities. Plexxi Connect provides processes and tools for managing event-based workflow automation to IT, enabling architects to build truly dynamic and scale-out cloud infrastructures.

Figure 2. Plexxi Solution Overview



Source: Enterprise Strategy Group, 2017

Plexxi and Nutanix

The Nutanix Enterprise Cloud Platform is a hyperconverged infrastructure solution that unites compute, virtualization, and storage in resilient, software-defined scale-out clusters. This platform replaces the silos of servers and storage and eliminates the need for separate management tools and processes. Nutanix is a leading vendor of hyperconverged infrastructure, and like Plexxi, has a strategy of partnering with leading infrastructure vendors with the goal of delivering complete IT solutions to the enterprise.

HCI challenges traditional networking environments with mixed traffic (compute and storage), which have very different network needs and out-of-sync time-based needs. For example, storage traffic is generally high-bandwidth with large packets, and a subset of it is sensitive to latency (e.g., metadata traffic). In addition, storage traffic may peak at different, unrelated times to compute traffic (e.g., storage may peak during backups or large file writes, whereas compute traffic may peak based on user loads, time of day, or other vectors). Compute traffic is generally bursty with smaller packets, and is more generally latency-sensitive. Traditional networks that rely on legacy switching and routing protocols cannot distinguish between these types of workload classes or sub-classes, and therefore the proper resources are not allocated to the proper workloads at the right time. In HCI environments, these two broad workload classes are mixed on the same network. Without a network designed for these types of varied workloads, the full potential of an enterprise's HCI clusters can be impeded.

Plexxi solves this challenge with its dynamic mesh fabric, innovative software-based control mechanisms, and unique integration into leading HCI systems like Nutanix. This intelligent integration allows the Plexxi fabric to allocate separate resources to compute and storage workloads, and to dynamically adjust the fabric allocation as needed. Both Nutanix and Plexxi solutions are architected to start small and scale performance and capacity linearly to meet the needs of growing organizations without introducing management complexity and without overbuilding from the start.

Plexxi Switches – Creating the Dynamic Mesh Fabric

Plexxi uses commodity Ethernet-based switches, such as the Plexxi Switch 3eq, and a novel physical interconnect system called the Plexxi LightRail to create a dynamic mesh fabric. Unlike leaf-spine architectures, Plexxi switches are deployed in a single layer, are physically connected in a daisy chain (i.e., no spine required), and make use of a passive layer 1 optical mesh for direct switch-to-switch connectivity that grows easily as the environment scales—supporting single rack up to cloud-scale deployments with equal ease.

The switch co-resident control software automatically forms a single-tier fabric with no requirement for the user to configure, tweak, and debug messy and complex protocols. This approach dramatically reduces cost and complexity as it creates multiple direct and indirect paths between physical switches. Plexxi Control software provides a global management tool for one or many fabrics, and allows the user to easily override the default fabric allocations with specific workload requirements. This architecture scales linearly, adding fabric capacity, resiliency, and multi-path options with each additional switch.

Plexxi Switch 3eq features include:

- 1 RU form factor and redundant, hot-swappable power supplies and fans.
- Access Interfaces: Up to 104 x 10GbE, 104 x 25GbE, 26 x 40GbE, 26 x 100GbE.
- LightRail Interfaces: 6 or 12 x QSFP28 (600 Gbps or 1.2 Tbps fabric connectivity).
- Switching Capacity: 3.2 Tpbs.

Plexxi Control

Plexxi Control provides centralized management, automation, orchestration, and visualization of the network deployment, aligning network resources to workload demands. Users gain a real-time view of network capacity and available paths, and tools to control how fabric resources (such as bandwidth, low latency paths, or dedicated access to paths) are allocated to specific mission-critical workloads. Using its innovative Dynamic Fitting Engine, Plexxi Control:

- Builds a model of the physical network, and lets the user define workload endpoints (such as VMs, physical machines and storage devices) and identify relationships and interdependencies between those endpoints.
- Allows the user to define bandwidth, latency, and isolation policies for named workloads that have specific network requirements.

- Uses sophisticated algorithms to create a dynamic topology that intelligently delivers a set of network paths and policies that “fit” the requested workload needs.

Workloads that are not specifically defined in Plexxi Control are automatically load balanced efficiently across the remaining fabric resources, ensuring that priority workloads are serviced first and have deterministic performance, without requiring the user to manage every specific workload or traffic type. The key to this approach is the dynamic mesh fabric created with Plexxi Switches and LightRail, which offer a multitude of direct and indirect paths, giving Plexxi Control many ways to solve a given set of workload needs. With a small amount of information about these workloads, Plexxi Control promotes specific traffic to direct or lightly utilized paths, or can allocate private dedicated paths for highly sensitive or performance-critical workloads (e.g., storage) if needed. It can choose from hundreds of noninterfering paths across the diverse fabric, enabling cost efficiency, higher performance, full utilization of fabric capacity, and greater flexibility. Plexxi Control also enables a workload-centric view of network resources, including a graphical representation of all paths across switches.

Traditional networks that rely on complex but rudimentary protocols that have no workload awareness quickly run into contention issues, forcing the user to continually upgrade her network with more bandwidth even when there is plenty of fabric capacity available. Generally, traditional networks over-provision their fabric links because they can't control the traffic, which is not only wasteful from a capital perspective, but also requires the user to manage, operate, and troubleshoot more network objects for no reason. The ability of Plexxi Control to intelligently fit important workloads to the required network resources and then intelligently load share the remaining fabric results in much higher leverage of the network with less network to manage.

Plexxi Connect

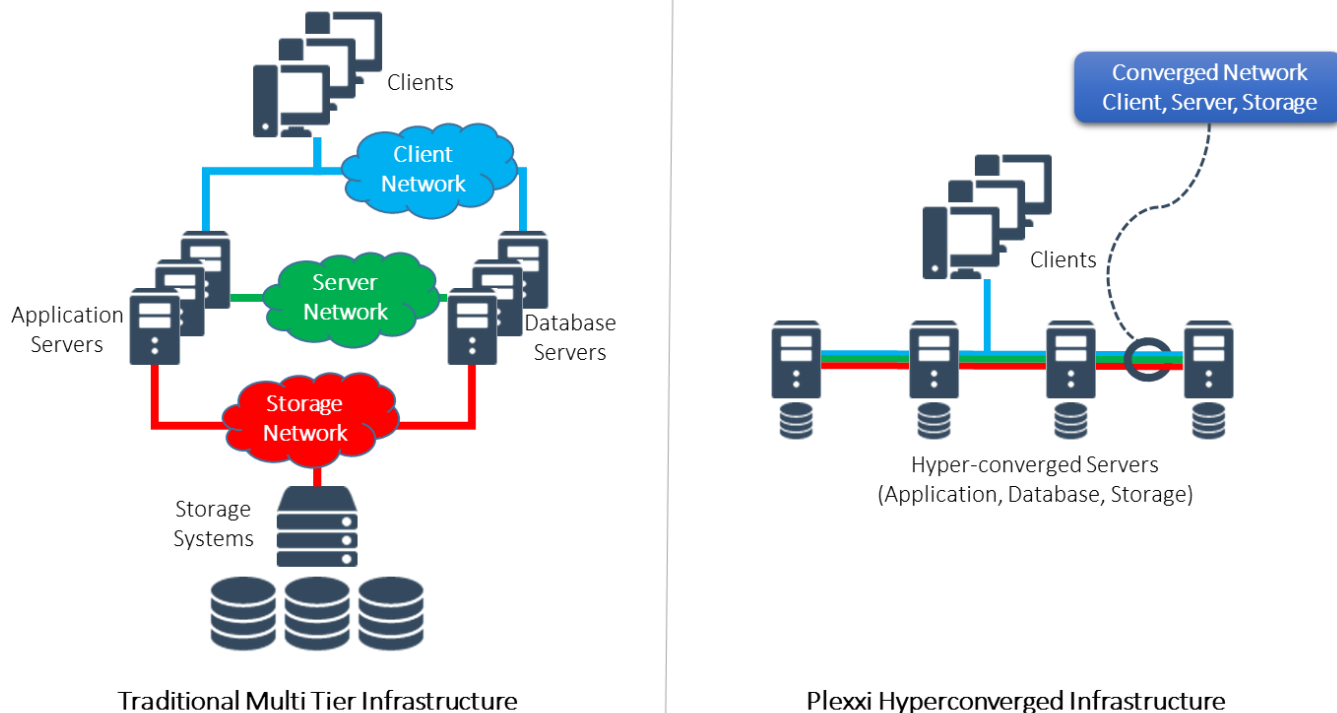
Plexxi Connect provides direct event-based API integration to other systems to provide workload automation for efforts such as configuring and provisioning of ports when a VM is created or moved. It also makes it easy to import information about workload endpoints and groups into Plexxi Control. Plexxi Connect drastically reduces the time and effort needed for organizations to leverage advanced automation capabilities.

In addition, Plexxi provides prebuilt Integration Packs that work with common compute and storage infrastructure vendor and open source solutions. These packs provide out-of-the-box automation for VM lifecycle events for VMware, Nutanix, and OpenStack, and leverage innate knowledge of the storage subsystems of these environments to automatically leverage Plexxi's unique fabric isolation capability for that critical workload. For example, Plexxi's Nutanix Integration Pack automatically discovers the cluster VMs that comprise the storage system in Nutanix environments and creates a secure, dedicated set of fabric paths to ensure that cluster operational traffic is never impeded by other user workload traffic. The pack also provides full VM lifecycle automation for both ESXi and AHV hypervisors—automatically configuring port profiles, VLAN membership, and LAG settings—and includes a dedicated GUI that combines the network visualization view from Nutanix Prism with Plexxi's fabric view to provide an end-to-end view of VM connectivity (VM to switch through fabric to switch to VM), which drastically reduces troubleshooting time.

ESG Lab Tested

In the traditional data center, each of the major infrastructure components are connected by separate networks. As shown on the left hand side of Figure 3, application servers communicate with one another and with databases using a separate network. This network may be physically or logically partitioned from the rest of the network. Simultaneously, application servers communicate with storage over the storage network, which can be a SAN (FC, FCoE, or iSCSI), or a private NAS network. In parallel, clients communicate with application servers using a separate network.

Figure 3. Plexxi for Hyperconverged Infrastructure



Source: Enterprise Strategy Group, 2017

With the advent of hyperconverged infrastructures, compute, database, and storage reside on the same hardware. Thus, it no longer becomes necessary, critical, or advantageous to have logical or physical separation between application servers, database servers, and storage. This new hyperconverged infrastructure is shown on the right-hand side of Figure 3.

Plexxi Control for Ease of Use

Administrators manage Plexxi systems using the Plexxi Control GUI, or by integrating Plexxi into other management environments using the APIs in Plexxi Connect. The Plexxi Control dashboard, shown in Figure 4, provides rapid at-a-glance status and health information about the entire Plexxi environment.

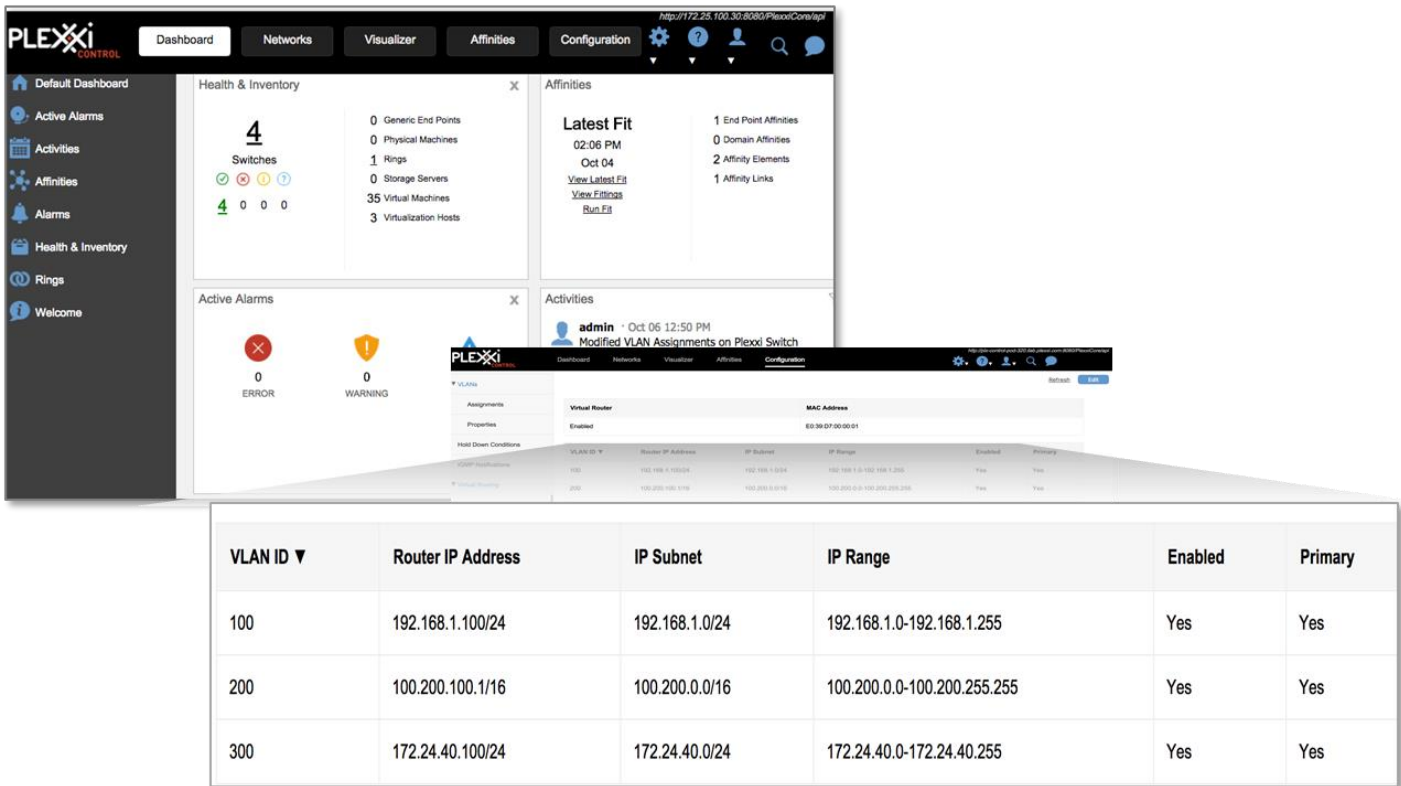
The dashboard is split into a two-by-two grid, with the top left box providing health and inventory information. Switch status is indicated by color and icon, using a green check mark for healthy switches, a red X for sick switches, a yellow exclamation point for switches needing attention, and a blue question mark for switches with unknown status. The box includes an inventory count for switches, physical and virtual machines, storage servers, other generic endpoints, and optical rings.

The top right box provides information on network affinities. Plexxi defines affinities as the higher-order relationships between applications, compute, storage, and the underlying network. Affinities also include the required qualitative behavior for all communication between members, such as performance through proximity, security, and the shared fate of all entities.

The bottom left box provides at-a-glance alarm status, with errors indicated by a red circle with an X, warnings indicated by a yellow shield, and informational alarms indicated by a blue triangle. The bottom right provides a rolling audit trail of administrative activities, ensuring that each administrator is aware of the recent activities of other administrators.

ESG Lab clicked on the **Configuration** button at the top of the dashboard to access the Plexxi system configuration. From the left-hand menu, we selected **VLAN Routing**, and Plexxi Control displayed the inter-VLANs configuration, grouped by the Plexxi Fabric, as shown in Figure 4. Administrators can manage inter-VLAN Routing using this interface, or programmatically through Plexxi Connect APIs. Plexxi Control automatically applies the routing configuration to all switches in the fabric.

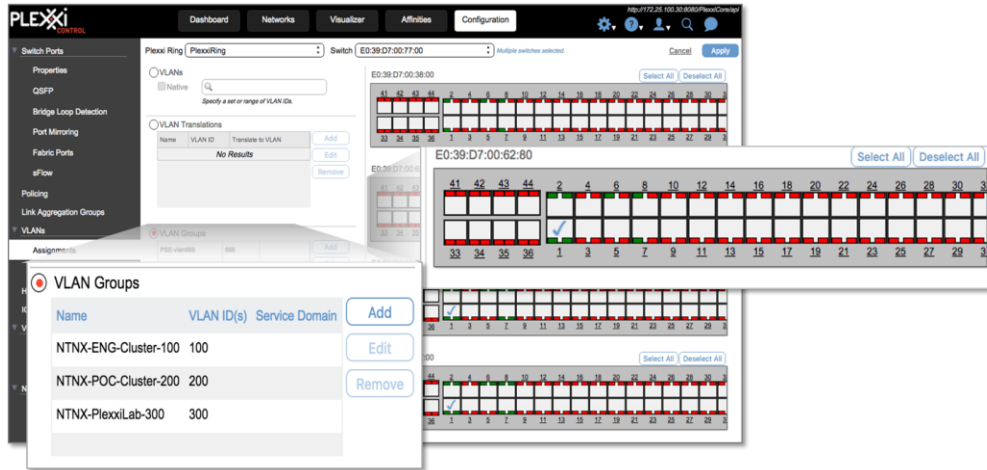
Figure 4. Ease of Management with Plexxi Control



Source: Enterprise Strategy Group, 2017

Plexxi Control can identify the presence of hyperconverged infrastructures. When detected, Plexxi Control can automatically generate the appropriate network configuration for each HCI cluster. As shown in Figure 5, Plexxi Control automatically configured three HCI-associated VLANs, with IDs *100*, *200* and *300*. On the right-hand side of the screen is a graphical representation of the ports for each switch in the Plexxi optical ring, including the connectivity and activity LEDs. A check mark in the port box indicates that port is part of the selected VLAN. Administrators can easily include or exclude ports by clicking the respective box. When satisfied, the administrator can lock down the VLAN configuration to avoid accidental changes in the future.

Figure 5. Dynamic Auto-configuration



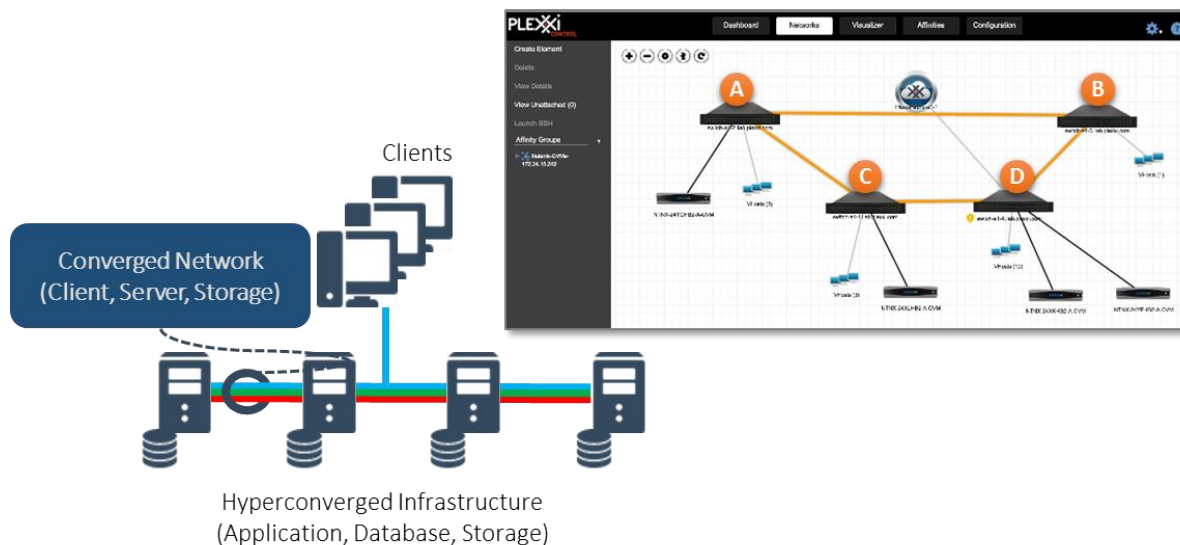
Source: Enterprise Strategy Group, 2017

Hyperconverged Networks – Guaranteed HCI Performance

Next, ESG Lab evaluated Plexxi for hyperconverged infrastructures, looking specifically at workload performance and resiliency. HCI clusters not only converge application and storage servers into the same environment, but also converge network traffic, placing all client, application server, and storage traffic on the same network (unless separate networks are deployed to segregate workload traffic). This puts additional burdens on the network infrastructure in providing sufficient network path bandwidth for all critical HCI components, especially HCI cluster configuration and control services (e.g., Nutanix CVM).

With Plexxi, administrators can isolate critical HCI control traffic to ensure that other workloads, such as client and application workloads, do not interfere with critical HCI control traffic. To test the impact of workload isolation for critical HCI control traffic, ESG set up an HCI cluster of four nodes connected by four Plexxi switches as shown in Figure 6. Two hyperconverged servers were connected to Plexxi switch **D**, and the remaining two servers were each connected to switches **A** and **C**. Clients were connected to each of the four switches.

Figure 6. Network Topology

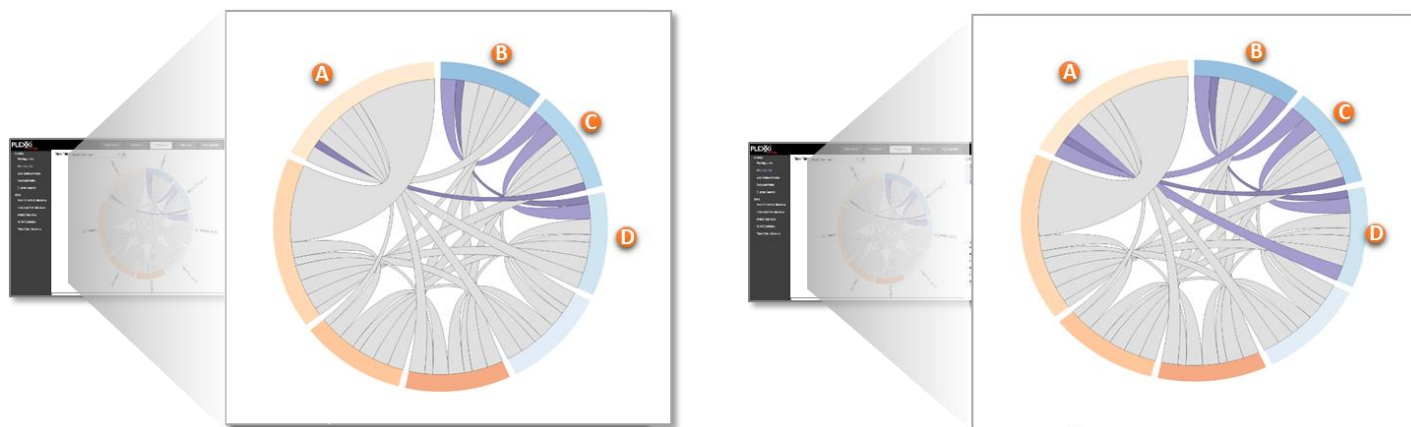


Source: Enterprise Strategy Group, 2017

In this configuration, the hyperconverged cluster sends storage traffic between switches **A**, **C**, and **D**. Using the **Peering Links / Affinities** option, Plexxi Control displayed the visualization of the network flows between the switches, as shown in Figure 7. The network flow diagram shows in a mesh, bands of links between peer switches (the width of the band indicating the quantity of links), the bands in purple indicate where Nutanix CVM traffic flows, and the darkness of the purple indicates the quantity of links used by the CVM. The grey bands simply designate “non-CVM” workloads.

The initial configuration, on the left-hand side of the visualization, shows traffic flows from switch **A** to switch **B** indirectly through switch **C** (as you can see no purple between **A** and **C**). ESG Lab modified the affinity, reserving additional bandwidth on the optical ring and providing a direct path from switch **A** to switch **B** to provide even better performance for the hyperconverged infrastructure. The result is shown in the visualization on the right-hand side.

Figure 7. Network Flow Visualization



Source: Enterprise Strategy Group, 2017

To maintain storage consistency and resiliency across the cluster, the HCI cluster generates multiple network operations/workloads. In addition to moving data from client to disk, the cluster must create and move copies of the data to other cluster member disks for data protection. The HCI cluster must also update the storage metadata in all storage

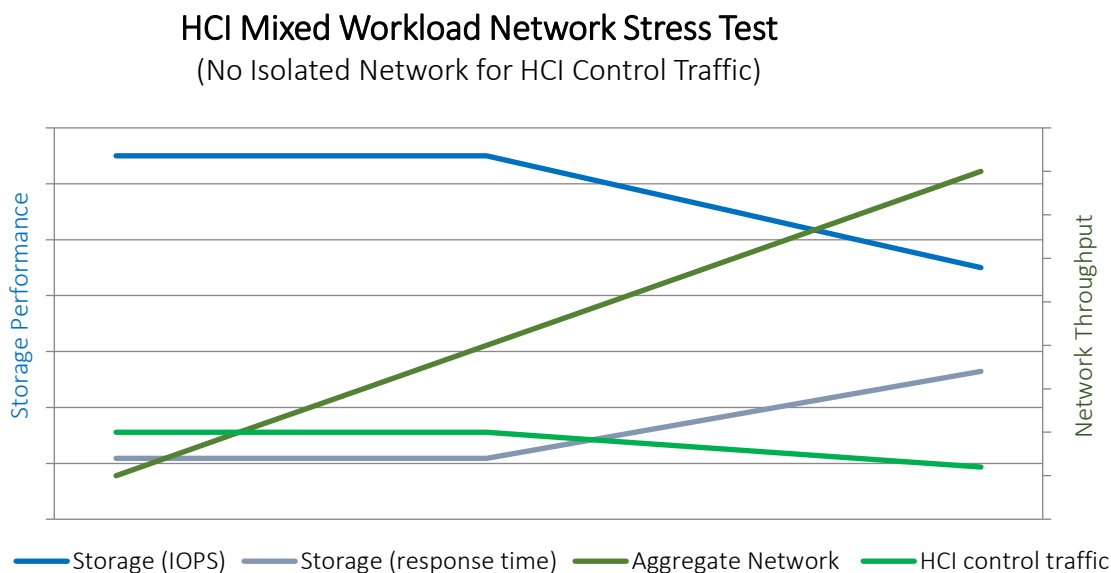
controllers in the cluster. The HCI cluster also manages the migration of virtual machines between physical servers (a.k.a., vmotion). As a result, storage operations place heavy demands on the network supporting the hyperconverged cluster.

Core to the Nutanix HCI cluster is the CVM, which orchestrates these operations across the cluster. To guarantee performance of this core HCI component, Plexxi enables the administrator to identify the CVM workload and designate a specific traffic path that is isolated from all other network traffic. This can be done automatically via Plexxi Connect Integration Packs.

ESG Lab used the [Vdbench](#) open source benchmark tool to simulate a real-world mix of virtual machines running a mix of application workloads including e-mail, OLTP database, web server, file server, and backup jobs. After the virtual machines hit a steady state level of performance, a network emulator was used to increase the aggregate network throughput over time. Two test runs were performed back-to-back: one with shared network fabric links and no HCI control isolation, and the other with HCI control traffic isolated from all other network traffic with Plexxi HCN.

Figure 8 summarizes how performance was impacted during the first test with a shared network fabric. Note how the increase in network traffic over time resulted in a decrease in application performance (decreased storage throughput with increased storage latency) and a drop in HCI control throughput. In a typical HCI environment, this increase in network throughput could be due to a variety of reasons including an increase in the number of applications and users, a burst of Internet upload/download activity, virtual machine migrations, HCI cluster expansion and recovery.

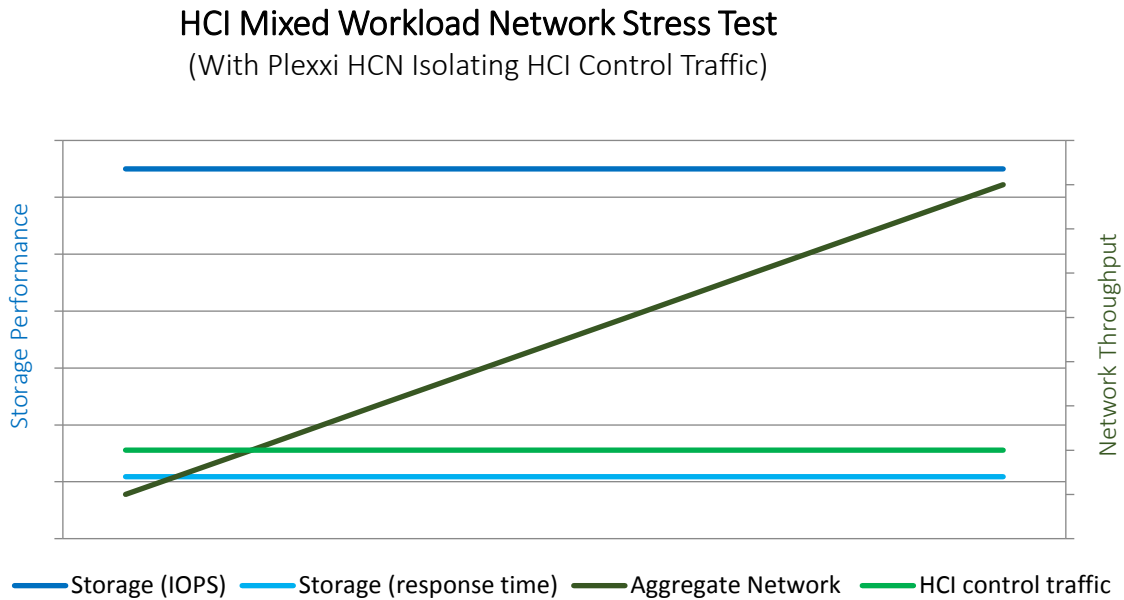
Figure 8. Application Workload Results: No Isolated Network Path for HCI Control Traffic



Source: Enterprise Strategy Group, 2017

As you can see in Figure 9, performance remained unaffected when the HCI control traffic was given exclusive ownership of a specific network path with Plexxi HCN. The storage performance of the mixed application workloads (storage IOPS and latency) and the HCI cluster throughput were not impacted as aggregate network throughput increased over time.

Figure 9. Application Workload Results: Isolated Path for HCI Control Traffic



Source: Enterprise Strategy Group, 2017

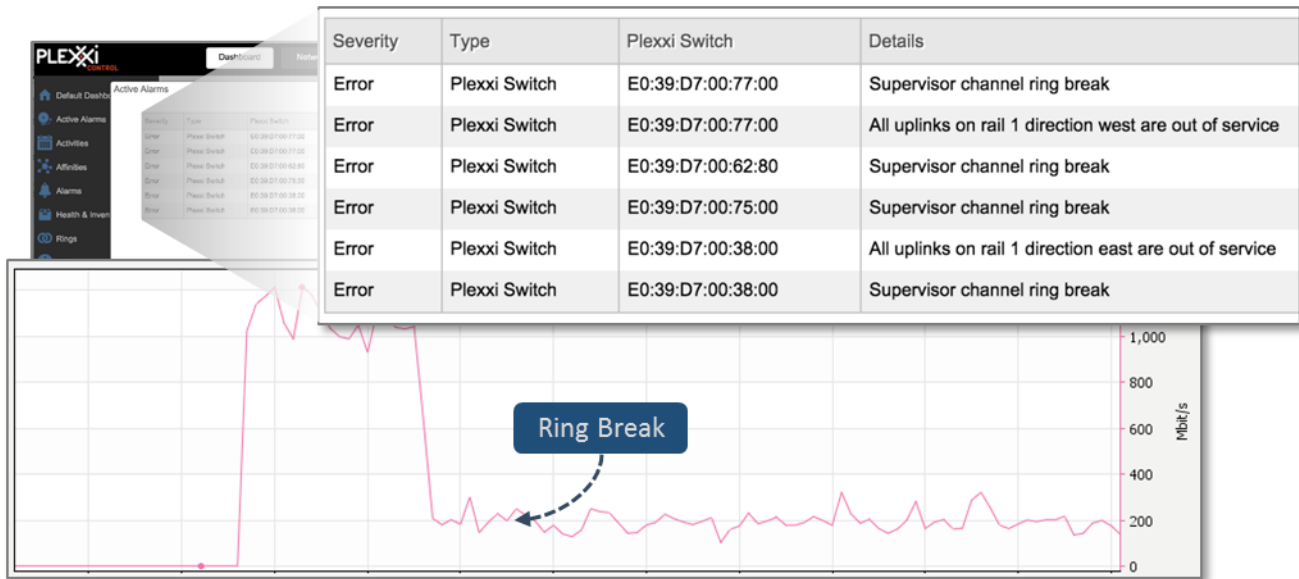
Plexxi Network Resiliency

Hyperconverged infrastructures include many performance optimization features. One feature that affects the network is automatic tiering, where the HCI storage controllers store most-frequently-used “hot” data in high-performance SSD. As the data “cools,” or becomes less frequently accessed, the HCI storage controllers move the data to lower-performance disk drives. ESG Lab observed that after the ten-minute vdbench tests, the HCI solution continued to move data across the network.

ESG Lab disconnected a fabric interconnect from one of the Plexxi switches to simulate a catastrophic cable failure. As shown in Figure 10, both switches on either side of the optical interconnect immediately generated alerts visible in Plexxi Control.

Each affinity that used that fabric interconnect was automatically updated by the switches to use an alternate path (the Controller is not required in this operation, so it worked instantly). Network traffic was automatically rerouted, to connect the two switches. Thus, there was no noticeable drop in network traffic or network bandwidth, and the HCI cluster continued normal operations.

Figure 10. Plexxi Network Resiliency



Source: Enterprise Strategy Group, 2017

i Why This Matters

Cloud computing comes with a set of challenges that are fundamentally different from the traditional IT environment. Traditional network systems were not designed with these challenges in mind. As organizations scale into the cloud, they need an approach that is built specifically for the cloud and that provides them with the context they need to connect and manage operations across their entire environment.

The Plexxi network system was designed and built with the purpose of dynamically connecting cloud and hyperconverged infrastructures. ESG Lab was able to confirm that organizations can use Plexxi to employ workflow-centric network connectivity and management. The speed and ease of configuration and management was clear, and we could ensure optimal network connectivity for multiple HCI clusters.

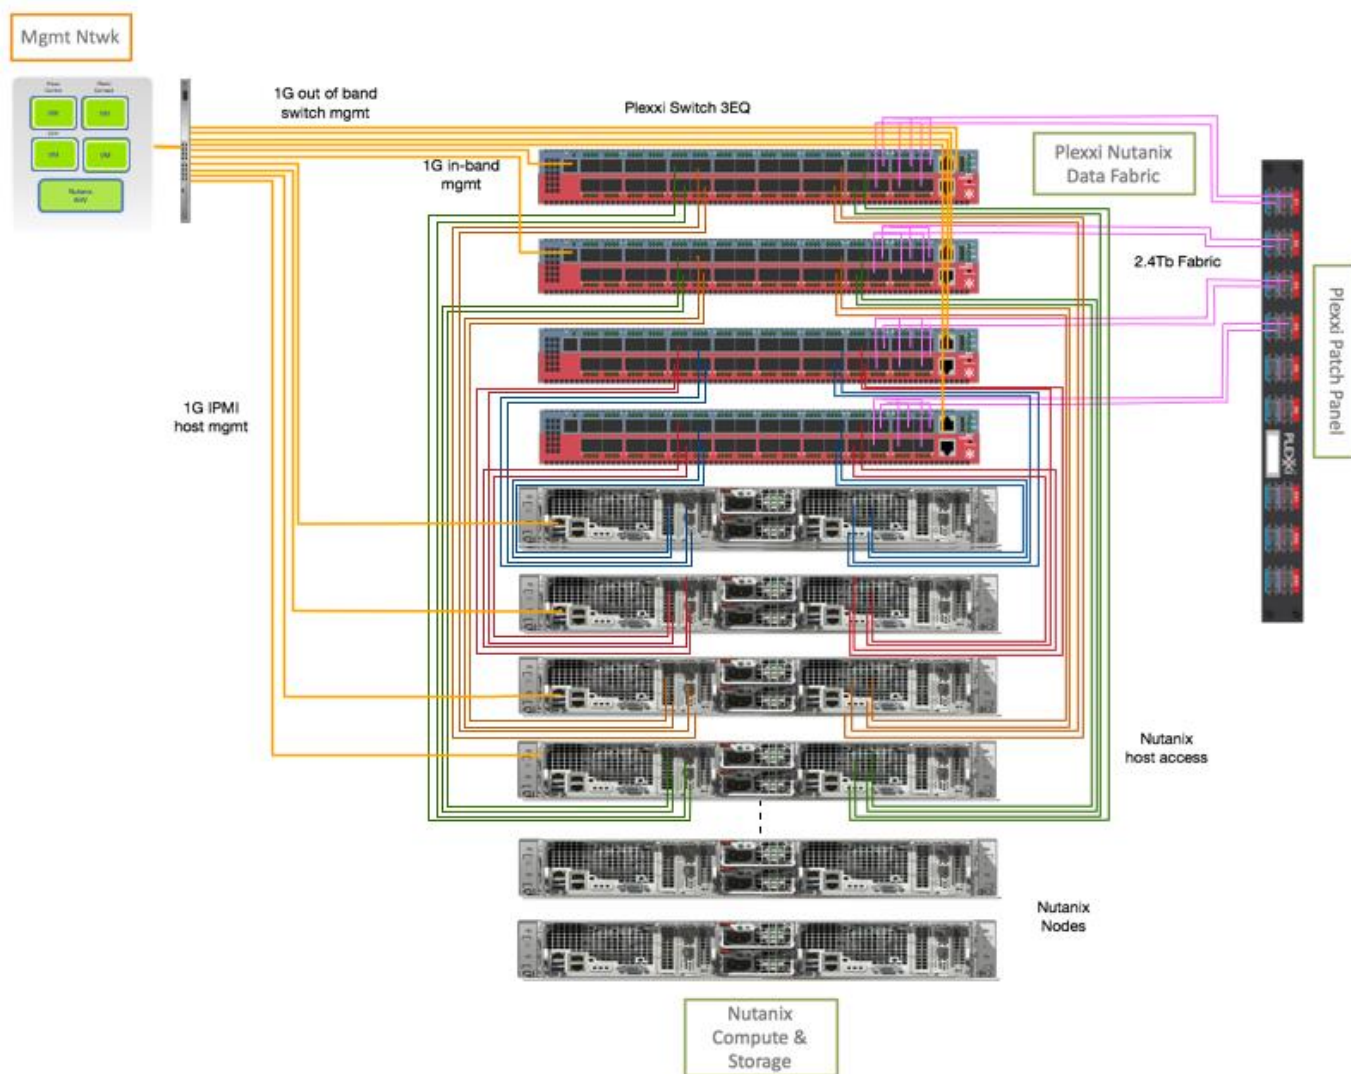
Plexxi networks proved to be resilient in the face of catastrophic failures, and to have the performance to meet the demanding needs of HCI clusters.

Plexxi HCN for Nutanix – Best Practices for a Resilient Modern Data Center

Plexxi Management Connectivity

Plexxi switches have a 1Gbps out-of-band port, which is the recommended way to manage the switches (you can optionally configure in-band, but that is not best practice). Plexxi Control communicates through this out-of-band port with each switch. Plexxi Connect needs to communicate with both Plexxi Control and Nutanix Prism. Typically Control and Connect are put on the same management network, and access for Connect to Prism is required through redundant in-band connectivity to the fabric. Figure 11 shows how this can be configured.

Figure 11. Plexxi Network Configuration – Redundant In-band Connectivity



Source: Enterprise Strategy Group, 2017

Host Connectivity

Modern hosts (and this includes HCI hosts) have at least two 1Gbps ports, and at least two 10Gbps (or higher speed) ports. Plexxi recommends that you provide at least quad 10Gbps ports for use by storage and workload traffic. This allows both in-band and out-of-band redundant connectivity and isolated storage and workload fabric paths to be built by Plexxi.

For each host, at least one, 1Gbps port should be dedicated to IPMI (host console access), and either the same port or another 1Gbps should be earmarked for out-of-band host access (that is, access that does not traverse the network fabric). Its uncommon for the 1Gbps out-of-band connections to be redundant, as they are not active fabric paths.

For each host, at least two 10Gbps ports dedicated to CVM traffic, and at least two 10Gbps ports dedicated to user workload traffic are required. It is not expected that 20Gbps is required for CVM, but 10Gbps is the minimum port speed and physical redundancy is a must.

In any open vswitch-based system (like Nutanix AHV), a second bridge/bond must be created for the out-of-band traffic—in a VMware setup, this is typically the default setup.

The Bigger Truth

ESG's annual IT spending intentions survey in 2017 revealed that data center modernization is tied for third in the list of CIO whiteboard initiatives, with 15% of respondents citing it as their top initiative.² IT organizations are taking to heart the lessons of the cloud- and hyper-scale data centers in driving efficiencies. However, the modern data center with a fully virtualized and hyperconverged infrastructure presents multiple networking challenges to IT organizations, including management, performance, agility, and resiliency.

The Plexxi network solution for the modern cloud and hyperconverged data center includes Plexxi switches with an internal or external integrated optical mesh fabric interconnect for dynamic management of application flows. Plexxi Control software provides intelligent, automated, application-driven management, letting cloud builders design their networks based on workflows. Plexxi Connect integrates the Plexxi solution with cloud automation and DevOps environments.

In hands-on testing, ESG Lab validated that managing the Plexxi solution was simple. We could quickly configure workflow-based affinities to optimize the network for Nutanix Enterprise Cloud Platform. Reconfiguring for other HCI solutions was equally quick and easy.

Plexxi demonstrated that performance of the HCI solution was guaranteed even under heavy workload usage. The solution was extremely resilient. Even when we pulled a cable to simulate a catastrophic failure in the optical mesh fabric interconnect, network operations continued unabated.

Enterprises looking to take advantage of the latest innovations in cloud infrastructures to build highly reliable, resilient, maintainable, and programmable network infrastructures to support hyperconverged infrastructures and the software-defined data center would be smart to take a close look at Plexxi's network infrastructure solutions.

² Source: ESG Research Report, [2017 IT Spending Intentions Survey](#), March 2017.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.



© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.



www.esg-global.com



contact@esg-global.com



P.508.482.0188