White Paper

The Importance of Integrated Networking to HCI

Sponsored by: Dell EMC and Intel

Brad Casemore
June 2017

IDC OPINION

Enterprises have embraced hyperconverged infrastructure (HCI) as a means of building software-defined datacenters (SDDCs). Unfortunately, networking is often overlooked as an integral element in HCI offerings.

This oversight frequently results in suboptimal HCI implementations that not only fall short of meeting tactical enterprise objectives but also create problems that can prevent enterprises from successfully pursuing their hybrid cloud and digital transformation strategies.

In an HCI environment, the network effectively provides the interconnected nervous system for the applications that ride over it. The quality of service, reliability, availability, scalability, and overall performance of the network ultimately determine the success of the HCI system.

Customers can choose to deploy ad hoc networking systems themselves, but that entails all the attendant costs and risks. It also means having to continually reassess, modify, and potentially re-architect the HCI network as the environment scales.

Alternatively, customers can look to an HCI system that features integrated, optimized networking. Such a system should provide network capabilities that ensure the performance of key applications, such as software-defined storage (SDS) workloads, through features such as intelligent queue management.

SITUATION OVERVIEW

Hyperconverged infrastructure has become increasingly popular among enterprise customers as a means of realizing a software-defined datacenter and providing the foundational infrastructure for private cloud.

Increasingly, HCI is adopted within the larger context of digital transformation, which is defined as the process of creating value, growth, and competitive advantage through new digital offerings, new business models, and new business relationships.

That said, enterprises adopt and deploy HCI for several reasons. IDC research has found that hyperconverged systems are embraced by enterprises to address multiple challenges, such as improving IT staff productivity, boosting operational efficiency, reducing capital spending, improving backup/recovery systems, increasing storage/server utilization, consolidating datacenters, reducing

the time needed to provision infrastructure, and improving application performance. It's also adopted as a means of moving away from architectures based on storage area network (SAN) (see Figure 1).

FIGURE 1

Challenges Sought to Address with a Hyperconverged System

Q. What are "all" of the challenges your organization is seeking to address with the use of a hyperconverged system? What is the "primary" challenge your organization is seeking to address with the use of a hyperconverged system?

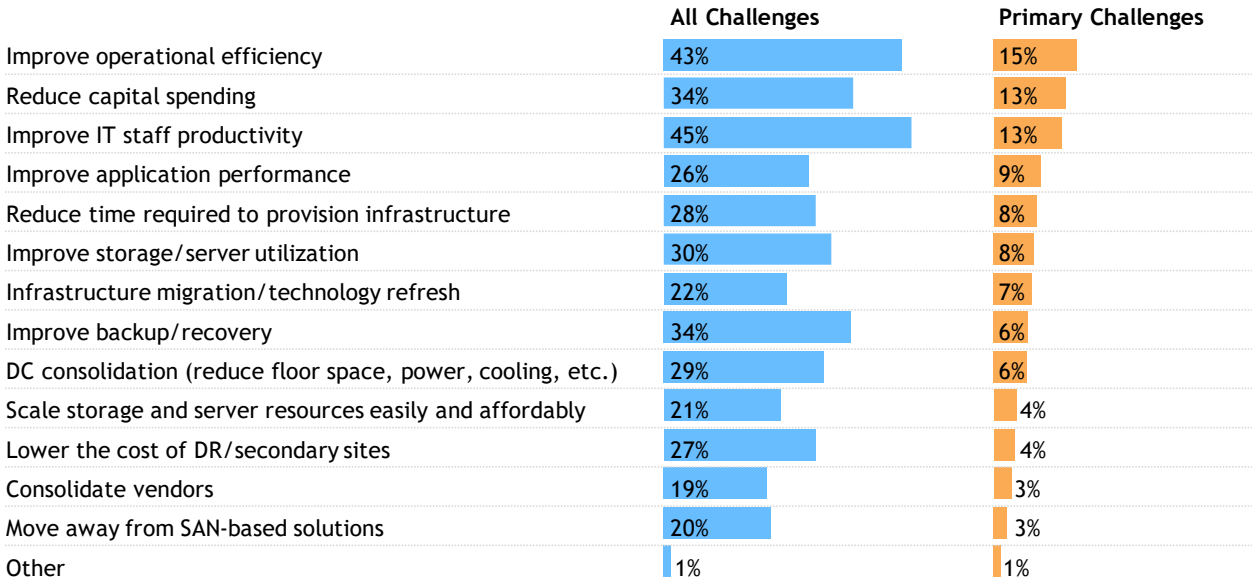

n = 302

Source: IDC's *Hyperconverged Systems Survey,* November 2016

IDC has found that the vast majority of HCI deployments occur at primary datacenters. What's more, IDC has discovered that 75% of HCI deployments replace existing infrastructure, including networking. By its nature, of course, HCI should break down operational and technological silos, allowing for improved IT staff productivity and greater operational efficiency. However, it can only do that if all infrastructure and IT elements, including the network, are properly considered and accounted for.

Customers deploying HCI solutions must focus from the outset on the needs of critical applications. With such an application-centric view, customers can begin to fully appreciate the salience of the network to successful HCI implementations.

Unfortunately, networking is often overlooked as a core element of HCI solutions. In fact, many commercial hyperconverged systems have tended to exclude networking from their solutions. Invariably, considerable attention is accorded to the compute and storage capabilities of HCI systems, but less attention (and sometimes none at all) is given to networking capabilities.

This often results in suboptimal HCI implementations that not only fall short of meeting tactical enterprise objectives but also create problems that can inhibit enterprises from successfully pursuing their hybrid cloud and digital transformation strategies. At a fundamental level, overlooking the network makes it difficult to plan for growth. Consequently, as the environment scales, performance degrades.

For that reason, networking needs to be a critical consideration in any HCI solution, not an afterthought. If networking is treated as a secondary HCI element, significant risks accrue. These risks often manifest in potentially serious issues that affect application performance, availability, reliability, and scalability.

As IDC data attests, HCI systems are generally implemented in demanding datacenter environments, where the network must account for application- and system-design requirements. IDC recommends that customers carefully assess the networking implications of any HCI deployment.

Direct integration of networking into HCI solutions eliminates the risks inherent in treating the network as an afterthought. With integrated and optimized networking, as enterprises scale up HCI compute appliances, networking will not become a problem that can affect the availability, integrity, or performance of business-critical applications. Moreover, from an operational perspective, integrated HCI networking can mitigate or eliminate the problem of dealing with the network as a separate silo.

VxRack FLEX's Integrated Approach to Networking

VxRack FLEX is a rack system architecture designed for flexibility. It is an engineered rack system rather than an appliance, with the compute, storage, and networking components integrated closely together to operate as a single entity.

Dell EMC's VxRack FLEX integrated systems deliver a fully validated and end-to-end network architecture at scale, suitable for an entry-level system comprising a small number of nodes or for a larger system composed of hundreds of nodes. Moreover, VxRack FLEX's Release Certification Matrix (RCM) – the suite of interoperable systems software and firmware releases certified for use on Dell EMC converged and HCI systems – ensures that customers do not have to concern themselves with testing and validating network firmware and software versions. This means that customers need not worry about having to resolve any network issues that might arise from firmware or software upgrades.

VxRack FLEX's integrated networking is designed to support business-critical customer workloads. It constitutes a core element in software-defined infrastructure that supports existing enterprise applications while also being integral to HCI implementations that support infrastructure-as-a-service (IaaS) workloads as well as platform-as-a-service (PaaS) environments and Big Data analytics.

As with all enterprise systems, Dell EMC's integrated networking moves data into and out of the system. In a software-defined datacenter system, however, it also provides the internode fabric over which all the software-defined components communicate to deliver the virtualized compute, storage, and network services. This internode fabric becomes more critical to application performance as the SDDC systems scale. In fact, this internode connectivity, when deployed in conjunction with logical networking best practices, can ensure a highly available software-defined architecture that provides network performance comparable with that of dedicated physical devices such as storage arrays.

In a traditional siloed network environment, enterprise IT has different networks for different purposes. For example, the typical enterprise datacenter might have an Ethernet network providing connectivity

between clients and servers and a Fibre Channel storage area network providing connectivity for storage arrays.

In an HCI environment, Ethernet provides all the networking connectivity, and an Ethernet fabric must be able to prioritize the delivery of certain application traffic flows over others. The network effectively provides the interconnected nervous system for the applications that ride over it. The quality of service, reliability, availability, scalability, and overall performance of that network ultimately define the capabilities of the HCI system.

The integrated networking in VxRack FLEX systems is optimized for HCI workloads, including software-defined storage and hypervisor flows. These application flows need to be classified and prioritized, with storage replication, for example, receiving different treatment than some other applications, such as web traffic. With VxRack FLEX, hypervisor traffic and storage traffic can be properly prioritized and treated differently from other types of traffic, avoiding the "noisy neighbor" problem that can affect the performance of virtualized applications running over traditional Ethernet networks. This is where capabilities such as active and intelligent queue management provide such value.

Indeed, in an HCI datacenter network, application flows are critically important. How those flows are forwarded determines the efficacy and value of the underlying network.

Generally, there are two broad classifications of flows: short flows, frequently called "mice flows," and long flows, often called "elephant flows." The vast majority of application flows in the datacenter are represented by mice flows, while the minority are elephant flows. That said, the overwhelming majority of bandwidth consumed in the datacenter is attributable to elephant flows. While mice flows are far more prevalent in number, elephant flows overwhelmingly predominate in bandwidth consumption.

Since the applications associated with elephant and mice flows are inherently different, they have similarly differing requirements for bandwidth, traffic loss, and latency. Typically, elephant flows are associated with applications that involve large data transfers, such as backup and virtual machine migration. By contrast, mice flows usually consist of short queries, responses, and control messages exchanged between application hosts. With mice flows, which are typically sensitive to latency and do not have high tolerance for packet loss, TCP retransmission time-outs can significantly degrade application performance.

In distributed application architectures, mice and elephant flows are carried together across the same queues on network links. This is where active and intelligent queue management becomes so important. For mice flows, packet loss and queuing latency need to be mitigated during periods of link congestion, whereas elephant flows must be accorded high throughput and requisite bandwidth without causing a degradation in the application performance of mice flows.

Switches on a network fabric can balance these objectives by providing sufficient buffer space for elephant flows while using TCP's congestion control mechanisms to prevent the larger application flows from monopolizing the queue and compromising the performance of mice flow applications. The objective, which can be achieved with a switch fabric that implements intelligent queue management, is to enhance all flow completion times to achieve greater overall application performance.

As such, datacenter switches in HCI environments must be capable of providing differentiated queue management for different application flows. Datacenter switches with simple buffering capabilities and

similarly basic approaches to queue management cannot provide this functionality. As such, they are unable to prevent mice flows from suffering the adverse latency effects during periods of congestion.

Because they are incapable of providing selective buffering for mice flows, conventional datacenter switches are compelled to buffer all packets at a congested link. This is why some switches have deep buffers. But switches with increased buffer space carry higher costs and also incur longer queuing latency, which increases flow completion times and affects application performance. Moreover, even when the switch attempts to buffer all packets during congestion, elephant flows could still consume all available buffer space, leaving none for mice flows.

Integrated, optimized networking in Dell EMC's VxRack FLEX provides enterprise customers with intelligent queue management that can prioritize and ensure the performance of key applications, such as SDS workloads, by applying differentiated queue management schemes that accommodate application requirements even during periods of link congestion (see Figure 2).

FIGURE 2

VxRack FLEX Provides Intelligent Queue Management

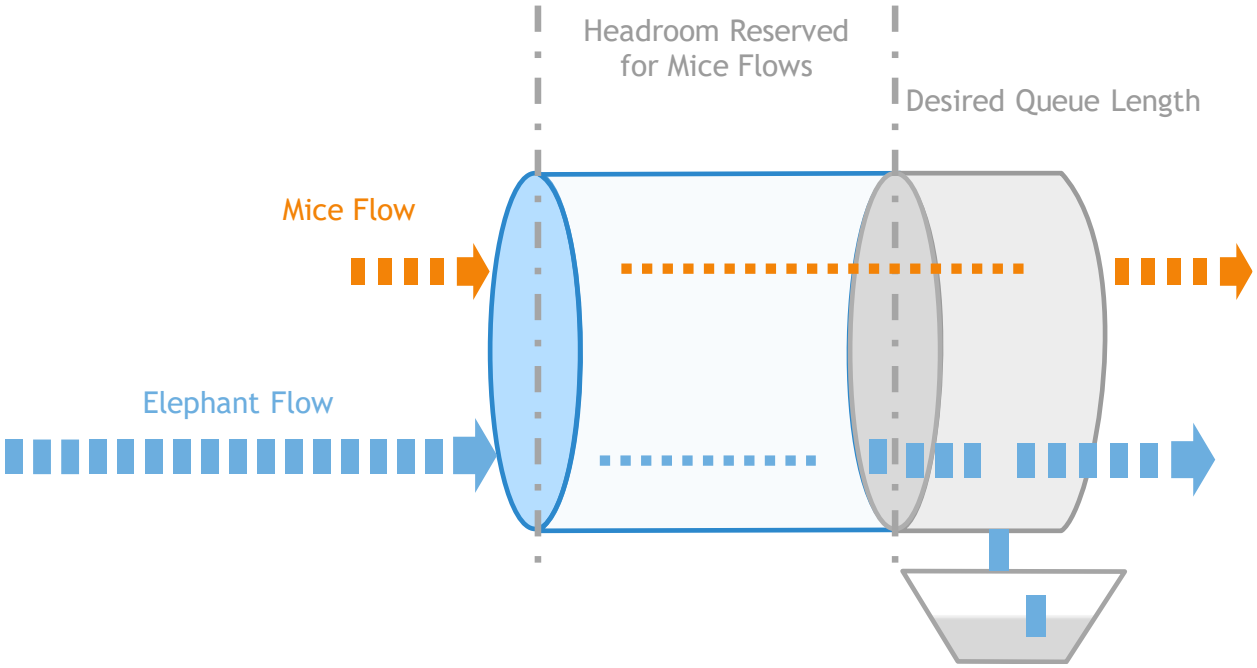

Source: IDC, 2017

Conversely, customers that choose to deploy ad hoc networking capabilities while building their own architectures will have to account for these requirements. This could result in the need to continually reassess, modify, and potentially re-architect their network as environments scale.

With respect to scale, VxRack FLEX's pre-designed and pre-engineered fabric facilitates datacenter-scale growth. The integrated networking makes it easy for customers to start with small environments and scale to encompass hundreds of nodes. Dell EMC also provides onsite resources, including a

network expert to design and plan for future growth. This ensures that all networking considerations are addressed for present and future needs, from bandwidth requirements to port densities, oversubscription rates, and east-west traffic.

Ultimately, customers must ask whether they want to design and build HCI networking systems themselves, assuming the costs and risks associated with such an undertaking, or look to a vendor such as Dell EMC to provide the networking for them in systems that are integrated and optimized for HCI application performance.

An additional consideration is that VxRack FLEX's integrated networking provides an SDN-ready architecture. Customers can deploy either VMware's NSX or Cisco's ACI, deriving benefits that include policy-based automation, network programmability, and orchestration, plus the security benefits of microsegmentation. Either SDN option — NSX or ACI — can deliver simplified operations, making network management easier and more agile. Moreover, SDN at scale can provide additional benefits, such as the ability for service providers to support multitenancy and isolation.

Even customers that don't perceive a need for SDN today will want the option of being able to implement it later. In that respect, VxRack FLEX integrated networking facilitates easy addition of nodes to support a future SDN deployment as the need arises.

CONCLUSION

This paper has examined the need for integrated networking in HCI systems. It has also considered Dell EMC's VxRack FLEX integrated networking, which is designed to support existing enterprise applications as well as infrastructure-as-a-service workloads, platform-as-a-service environments, and Big Data analytics.

Many commercial hyperconverged systems on the market have tended to exclude networking from their solutions. This often results in suboptimal HCI implementations that fail to deliver on their promise and compromise SDDC and digital transformation initiatives.

In an HCI environment, the network embodies the interconnected nervous system for application delivery. The reliability, availability, scalability, and overall performance of the network ultimately define the capabilities of the HCI system.

Through capabilities such as active and intelligent queue management, the integrated networking in VxRack FLEX systems is optimized for HCI workloads, including software-defined storage and hypervisor flows. This integrated, optimized approach to HCI networking eliminates the risks associated with treating the network as an afterthought, and it helps ensure that SDDC deployments fulfill their promise and meet their objectives.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com