RESEARCH

A ZK Research Whitepaper

October 2014

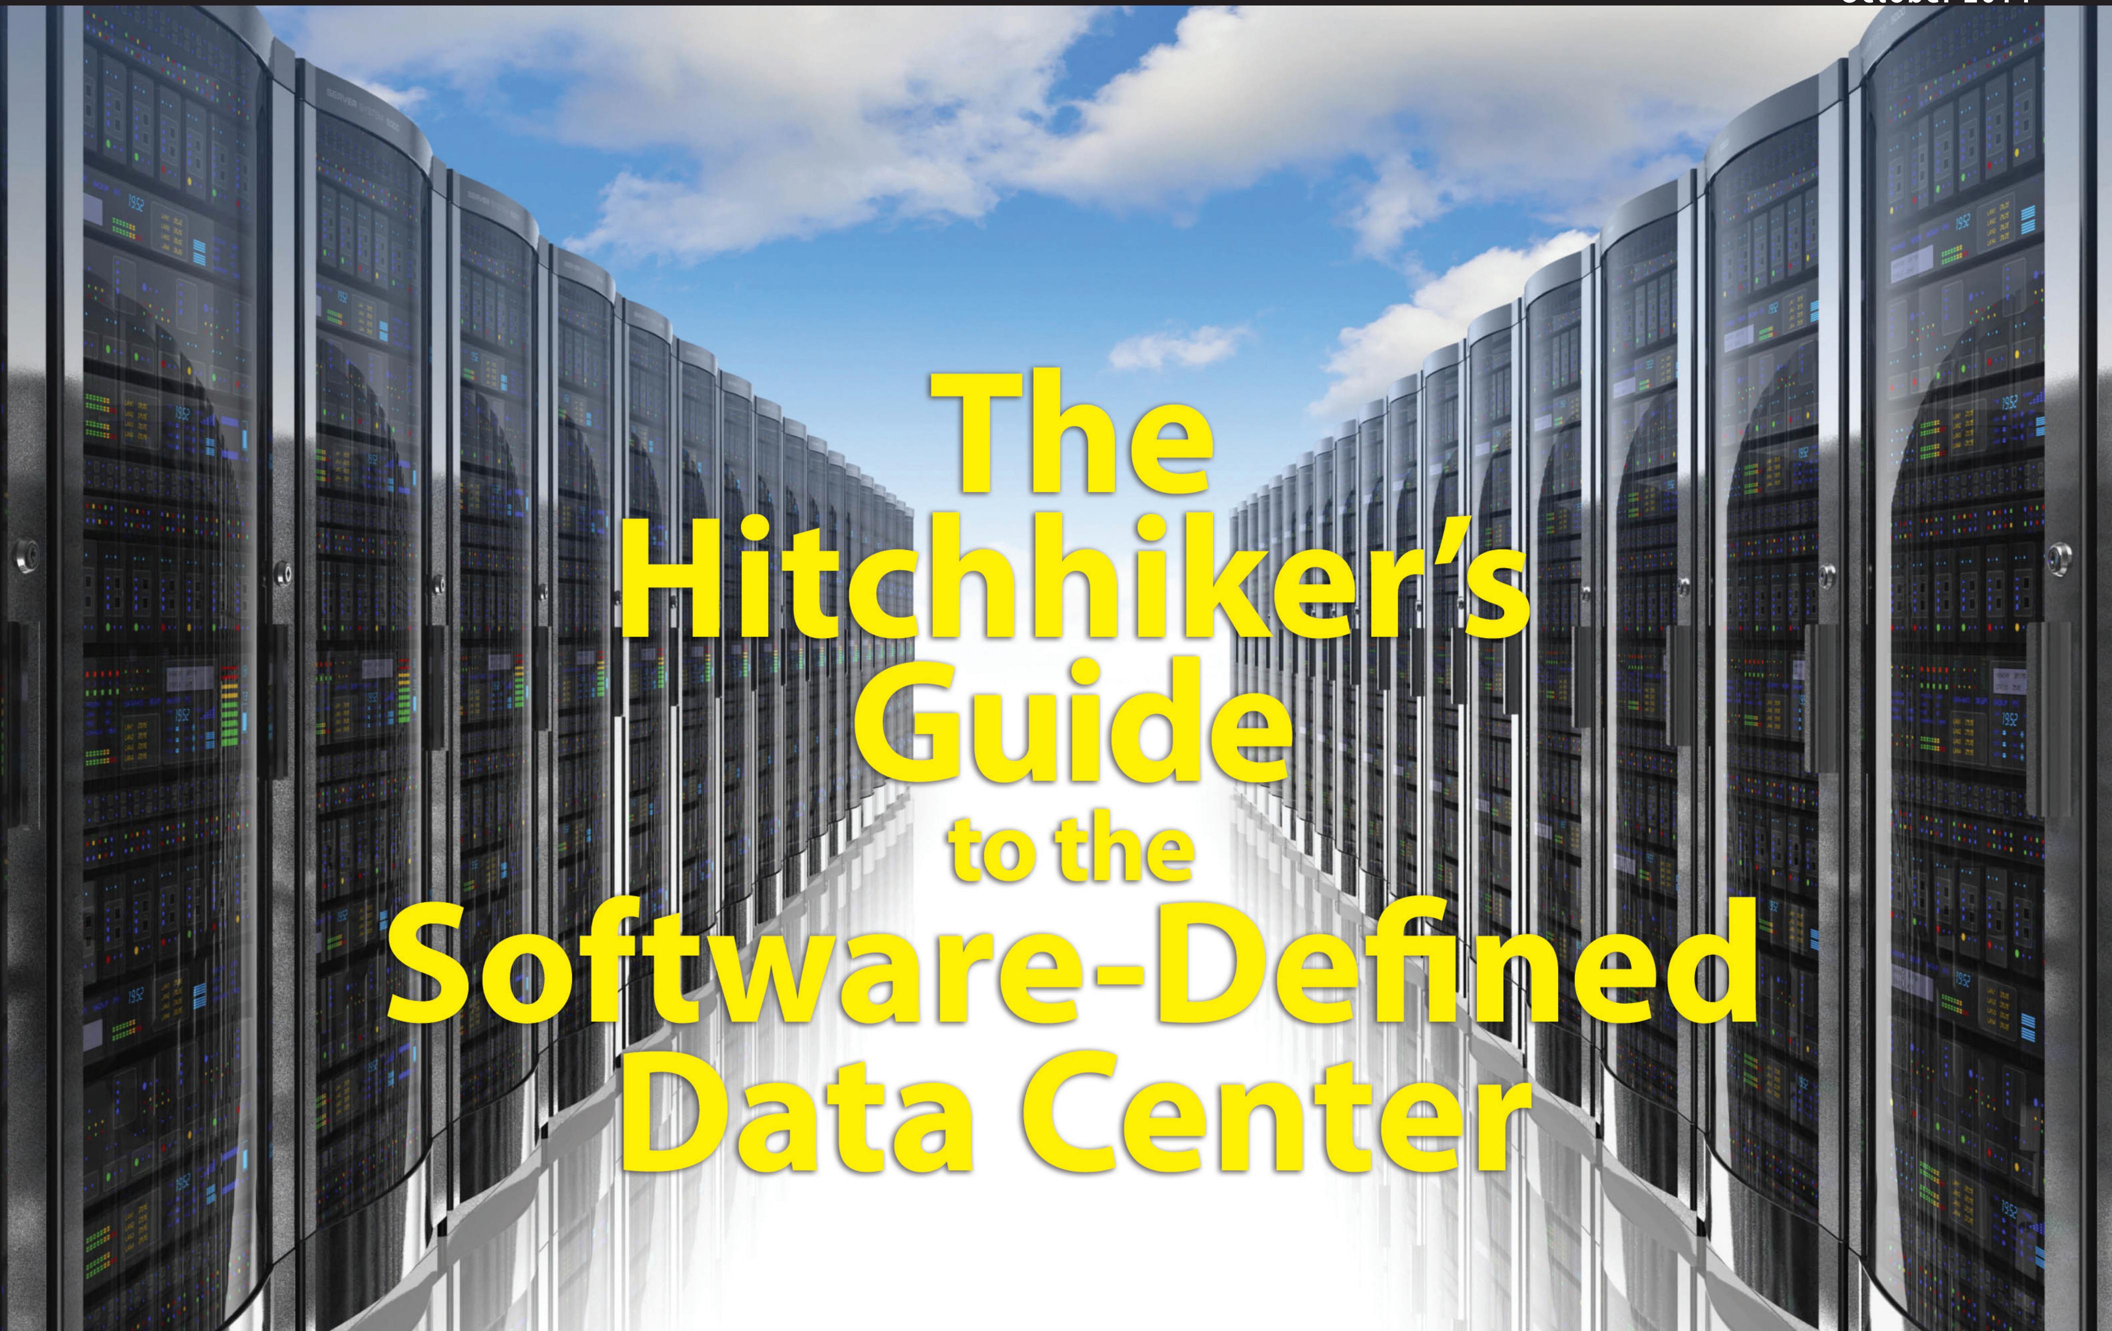

The Hitchhiker's Guide to the Software-Defined Data Center

The Software-Defined Data Center Is the Obvious Entry Point into Software-Defined Networking

The enterprise data center has undergone several major shifts since the introduction of computing as a corporate resource. The computing platform has evolved from mainframe computing to client/server computing to Internet-based computing, and now the industry finds itself in the midst of the next major data center transition—the evolution to the software-defined data center (SDDC). An SDDC can be thought of as an environment where all the data center resources are fluid and can be dynamically moved to the service or application that requires it based on business policy. There has been no single driver for this transition but rather several that have come together simultaneously, including the following:

> **Server virtualization is now the norm.** During the past decade, businesses have been virtualizing more and more workloads. Today there are far more virtual workloads than physical servers (see Exhibit One), and this has driven virtualization in other areas such as storage and other network functions.

> **The use of virtual machine (VM) mobility is increasing.** Server virtualization was initially used to consolidate the number of physical servers by an order of magnitude or more. Today, though, IT leaders are striving to create a more agile data center through the use of virtual machine mobility. VM mobility enables an application to move from one

ABOUT THE AUTHOR

Zeus Kerravala is the founder and principal analyst with ZK Research. Kerravala provides a mix of tactical advice to help his clients in the current business climate and long-term strategic advice. He delivers research and advice to the following constituents: end-user IT and network managers; vendors of IT hardware, software and services; and members of the financial community looking to invest in the companies that he covers.

Exhibit One: Virtual Servers Are Now the De Facto Standard

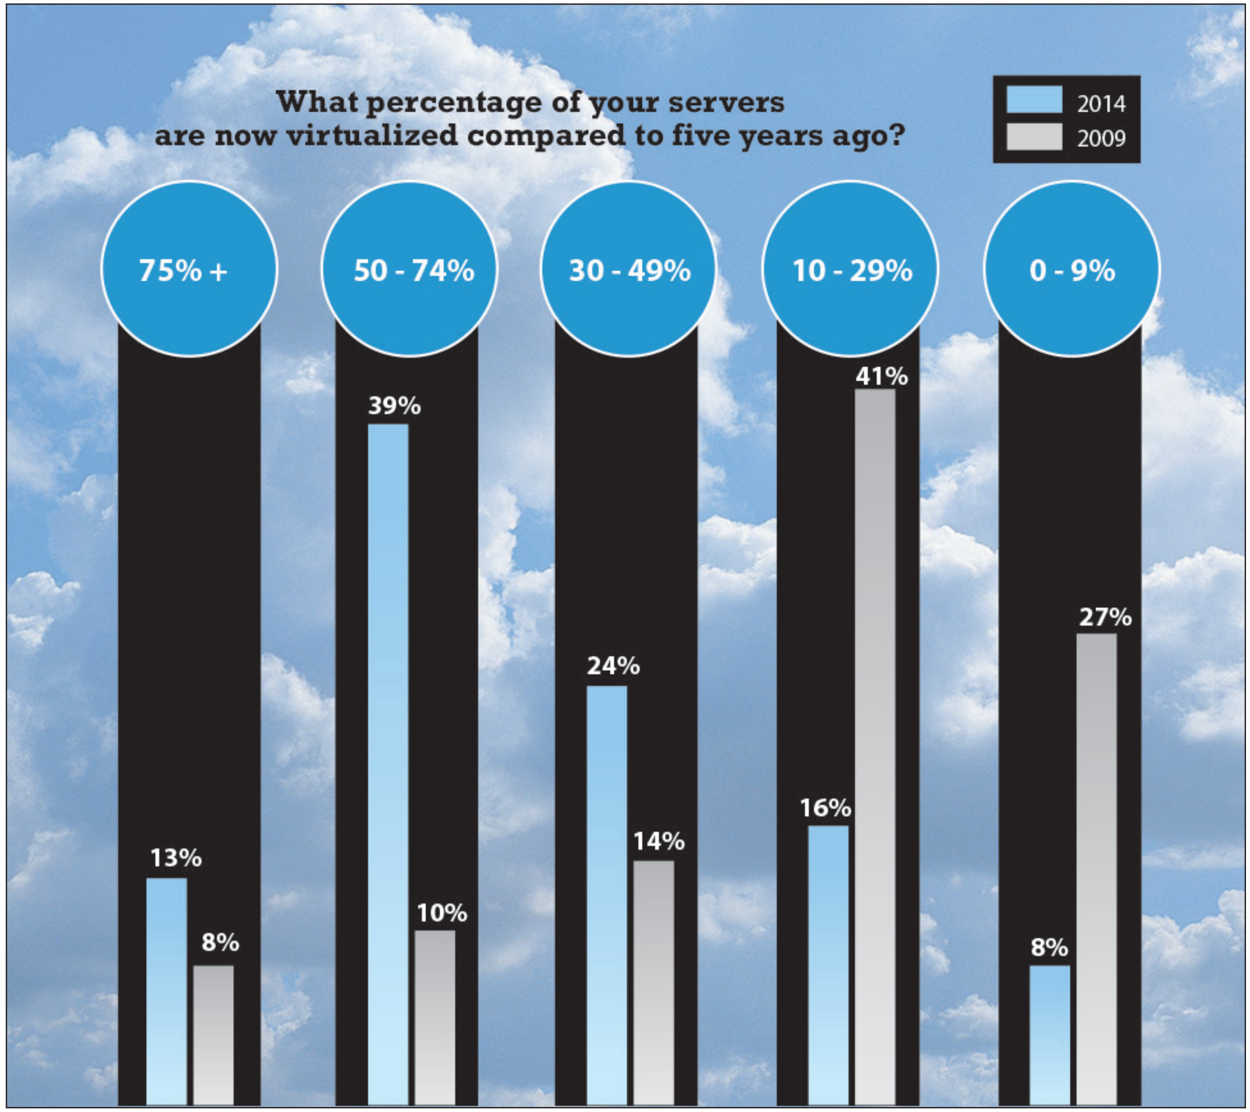

Source: ZK Research Data Center Survey, 2014

Approximately **82%** of an IT budget is now used to maintain the status quo.

physical server to another in real time.

> **Network fabrics are being developed.** Network fabrics have brought the scale and economics of virtualization to the network. The network fabric is a necessary step in the evolution of the virtual data center because it supports all of the compute components in a data center, whether virtualized or not.

> **IT operations and automation have been simplified.** The evolution of the data center has certainly driven costs down and allowed businesses to use technology much more efficiently, driving productivity to new heights. However, the IT environment has become much more complex than in previous years. ZK Research studies indicate that approximately 82% of an IT budget is now used to maintain the status quo, driving organizations to find ways to simplify the current environment. Virtualization allows IT managers to automate many of today's manual operational tasks through the movement of virtual resources. If the movement of VMs is not well coordinated with network operations, the network could be blind to the traffic flows.

> **Cloud computing is a key initiative.** The long-term vision of virtualization is to ultimately move many of the virtual resources to the cloud. Cloud is a key initiative for most CIOs today. In fact, ZK Research forecasts that global cloud revenue will grow from $14 billion in 2013 to more than $26 billion in 2020. However, the success of cloud computing is highly dependent on the network providing the foundation for cloud-based resources to traverse.

> **Data center traffic patterns are changing.** Historically, the majority of traffic in a data

center moved in a north–south direction. Today, virtualization has driven a significant amount of east–west traffic. This change in traffic flow has created new challenges for networks that were optimized for north–south flows.

North–South vs. East–West Traffic

North–south traffic: Traditional data centers are designed with multiple tiers and allow traffic to flow freely, although sequentially, through every tier. Typically, the data center is made up of three tiers: top-of-rack, aggregation and core. Legacy networks are designed to move traffic through each tier, through the core and then back on a similar path, giving it a north–south direction. Given that legacy applications were built in infrastructure silos in which each application had its own servers, storage and network, a network optimized for north–south traffic was sufficient.

East–west traffic: Virtualization, service-oriented architecture (SOA) and cloud computing are based on the principle of shared resources. This means that at any given moment, traffic may need to move from one virtual resource to another in as short a time as possible. The lateral movement of traffic between servers—from left to right and right to left—is known as east–west traffic. A network designed for east–west traffic is much flatter than one with a traditional multi-tier design.

The transition to an SDDC will be the most significant IT transformation since the Internet was born because it aims to bring together storage, computing and data networking. Implementing an SDDC has a profound impact on the network. The vision of SDDC is for all data

Exhibit Two: The Network Is the Backplane of the Software-Defined Data Center

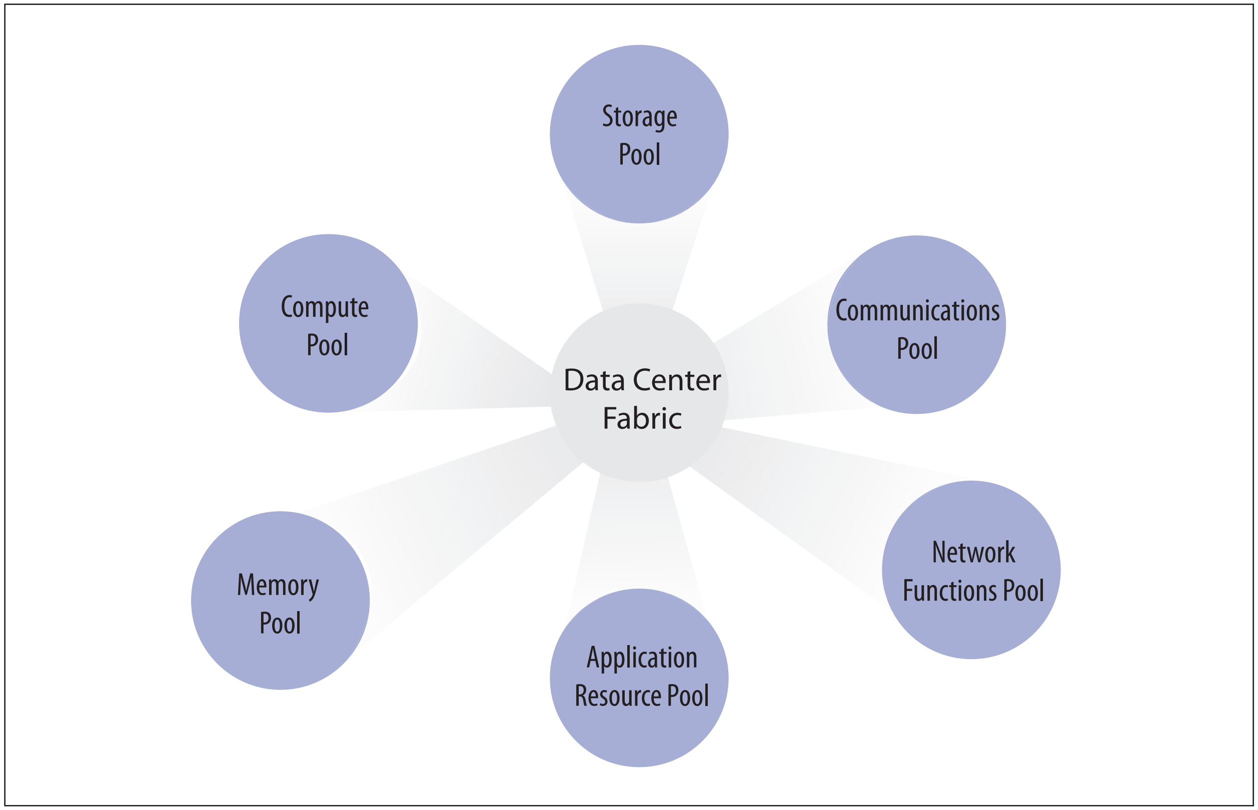

Source: ZK Research, 2014

center resources to exist as pools of virtual resources that can be accessed by whatever service needs them, whether across the data center, across the city or—in theory—across the globe. In this case, the network acts as the backplane for the virtualized data center (see Exhibit Two).

The SDDC will create new demands on the network and gives rise to software-defined networking (SDN). These new demands are driving the shift to an SDN that is better aligned with the other changes in the data center.

Section II: Limitations of the Current Network

The current network architecture used to support data centers is outdated and cannot support the SDDC. The historical value chain in the traditional data center stack assigned the highest value to applications and considered the network to be "plumbing," having the least value relative to the rest of the stack. The network was a necessary part of the overall data center, but because most applications were best effort, it ultimately did not play a critical role in the overall performance of applications.

More specifically, the network of the traditional data center has the following limitations in its ability to support a virtual data center:

> Its three-tier architecture has too much latency to support real-time and large workloads. In a typical three-tier design, all traffic between servers is sent through each tier, to the core and then back through each tier to the other server. All these hops between network

Inefficient use of network ports requires companies to purchase many more ports than necessary.

devices add a tremendous amount of latency that will affect the real-time needs of the virtual data center. The traditional three-tier design will no longer be sufficient.

> **Its use of Spanning Tree Protocol (STP) leads to inefficient network infrastructure usage.** STP prevents routing loops and broadcast radiation by disabling ports that are not part of the spanning tree. These disabled ports are only made active in the event of a failure of one of the active ports. This inefficient use of network ports requires companies to purchase many more ports than necessary.

> **Traditional technology is not designed for virtual or cloud environments.** Traditional data center technology is designed for best-effort traffic such as e-mail and Internet. Cloud computing and the real-time movement of VMs require much lower latency and guaranteed delivery. Best effort is no longer good enough. Additionally, because of the critical nature of cloud computing and virtualization, lossless Ethernet becomes a must have. Legacy network infrastructure was rarely designed with lossless Ethernet as a high priority, or even as a consideration at all.

> **Its poor management tools and high up-front cost drive up TCO.** The portfolios of some of the industry's larger players are made up of a combination of in-house–developed technology and acquired technology. These different systems often have different operating systems (OSs) and management tools. This is often the case in mature markets, where acquisitions are used to increase share as well as fill in technology holes. Trying to manage the end-to-end performance of the disparate components of a data center network is almost impossible

The transformation of the data center is ushering in a new era in networking.

with today's network infrastructure and management platforms.

> The lack of centralized management leads to inefficiencies. Managing traditional networks is a highly distributed process, with numerous touch points and almost no automation capabilities. This means that even the simplest changes to a network can often take weeks or even months to accomplish. Also, the lack of automation means IT is working in a reactive mode that leads to a high percentage of downtime being caused by human error.

> The network is "blind" to the rest of the IT stack. When a change is made in the data center, that change is often invisible to the network. This can mean that when changes are made to the compute or application layers, the network changes are often implemented independently, later or not at all. Alternatively, seemingly unrelated network changes can often impact application performance. Ideally, there would be a high degree of coordinated orchestration across the data center stack to ensure any change made does not adversely affect application or service quality. Similarly, applications are usually blind to the network, and they are typically unable to recognize inefficiencies and request optimization actives.

The transformation of the data center is ushering in a new era in networking. A solution optimized for the SDDC—incorporating high performance, low latency, orchestration, automation and guaranteed delivery—is required.

Historically, companies made network infrastructure decisions based on brand and incumbency instead of technical superiority. But as the era of the virtual data center gains momen-

Over time, more and more data center traffic will flow east–west.

tum, the network will continue to increase in value. Consequently, network managers will find that infrastructure that may have been good enough for Internet computing will not be good enough for virtual computing.

Section III: Network Requirements of the Software-Defined Data Center

The network requirements of an SDDC have created the need for a change in infrastructure and architecture. There can be no compromise between performance, reliability, scalability, security, features and cost. A solid network foundation with infrastructure that has been designed with virtualization in mind is necessary for the long-term evolution and success of the data center. Key attributes for this foundation include:

> A network fabric capable of supporting east–west and north–south traffic flows: Legacy data centers were built with client/server computing in mind, predicated on predominantly north–south traffic flows. Now, factors such as the transition to composite application architectures, virtualization mobility and the interweaving of data center traffic drive flows in a predominantly east–west direction. Over time, more and more data center traffic will flow east–west, causing significant congestion in the aggregation points of a legacy network. A network fabric can be thought of as a single entity where traffic can move east–west and is never more than a single hop away from its destination. This is critical for the movement of virtual workloads, in which any latency will cause a disruption to the business. The network fabric must support both east–west traffic movement and north–south traditional traffic flows.

> **Software-defined networking capabilities:** A software-defined network is a network in which the control functions have been extracted out of the physical infrastructure and abstracted to a centralized control function. The control functionality can be delivered via a centralized controller, or it can be distributed across the network. There can even be a "controller-less" solution in which the functionality is delivered via a cloud service. Any change made to the network can be implemented at this centralized SDN point-of-control and then pushed down to the infrastructure below it (see Exhibit 3). Additionally, the SDN control function has a set of northbound APIs that enable communications with applications. Ultimately, any change made at the application layer could trigger a set of commands in the control layer that could then be pushed down to the network infrastructure layer. Using an SDN can greatly improve the agility of the network and enable automated provisioning and configuration.

> **Automation and orchestration of the network based on open standards:** The ability to automate network operations tasks and orchestrate these changes along with other data center technologies is a key requirement of the SDDC. Some vendors may choose to do this through the use of proprietary protocols. However, in an SDDC, the use of proprietary protocols should be avoided because this limits the ecosystem of other data center solutions that can be used with those vendors' solutions. Open standards are a must when it comes to implementing an SDDC.

Exhibit Three: A Software-Defined Network Architecture

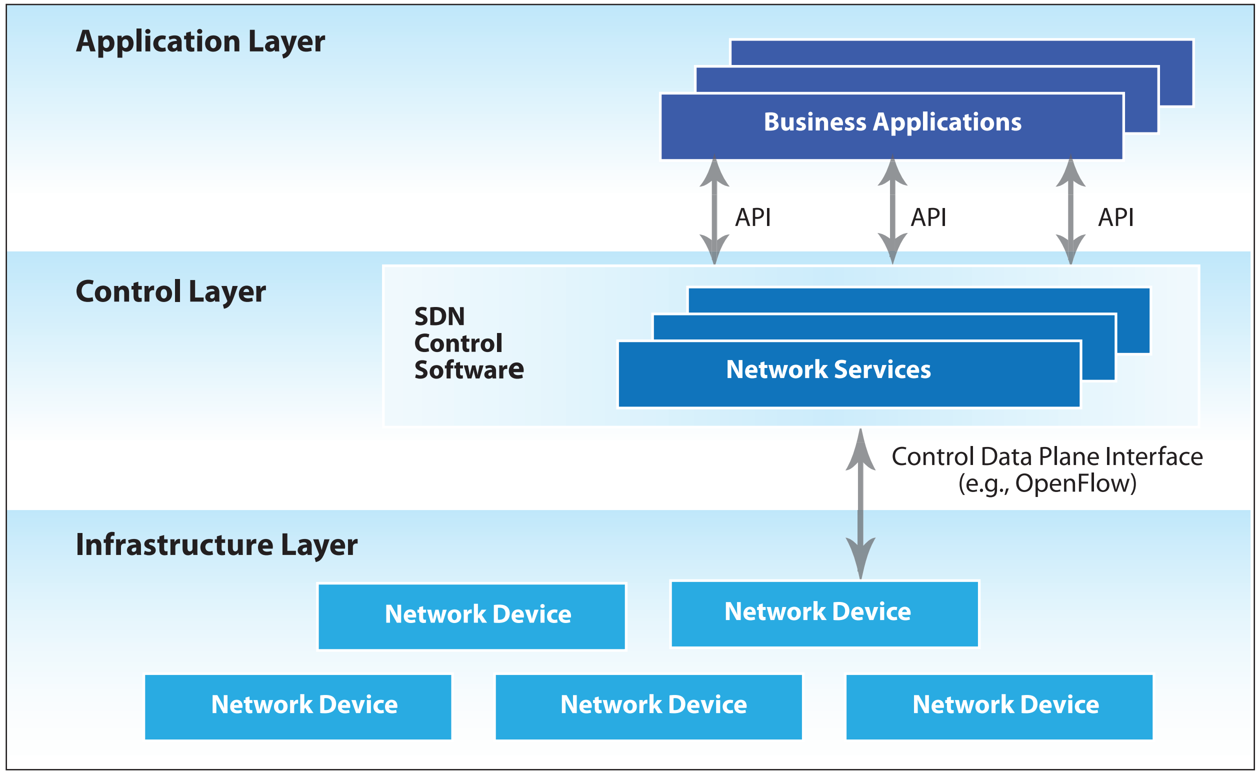

Source: ZK Research, 2014

> **Creation of a data center virtualized backbone:** Bridging data centers has become a primary focus for many companies. Having the ability to move storage traffic between data centers over Ethernet has many benefits. Storage protocols such as Fibre Channel or Fibre Channel-over-Ethernet (FCoE) operate at Layer 2 of the OSI stack. This means that interconnecting storage networks between geographically dispersed data centers is a significant challenge.

Root bridge STP-based VLANs are not robust enough to transport storage. Additionally, these networks do not support optimized path selection, which is necessary to minimize the latency of the storage traffic. Advanced network standards such as Shortest Path Bridging (SPB) and Transparent Interconnection of Lots of Links (TRILL) are being positioned to replace the outdated STP and allow Layer 2 domains to be extended between data centers, creating a single, seamless backbone for storage traffic. To ensure greater interoperability in the future, it's critical that the bridging technology be based on a standard. However, the Layer 2 issue is only one area that must be addressed. The ability to support dual-homing of servers and Layer 3 routing must also be part of the design criteria for the next-generation data center. Solving only one piece of the equation is insufficient and, realistically, unacceptable in the enterprise.

> **Architectural choice in next-generation data center design:** The most obvious and most commonly discussed migration strategy from the current three-tier data center architecture is to move to a two-tier architecture (see Exhibit Four), in which all the devices in a single rack connect to a top-of-rack switch. These top-of-rack switches are horizontally interconnected. Instead of running to an end-of-rack switch, they are connected directly to the core switch, thereby eliminating the tier of switches at the end of each row. This leads to TCO savings that are both capital (removal of hardware) and operational (simpler design means simplified management) in nature. There is no one right solution here. Some companies will opt to stay with a three-tier architecture, while some will choose a two-tier environment. The network infrastruc-

Exhibit Four: Two- and Three-Tier Architectures

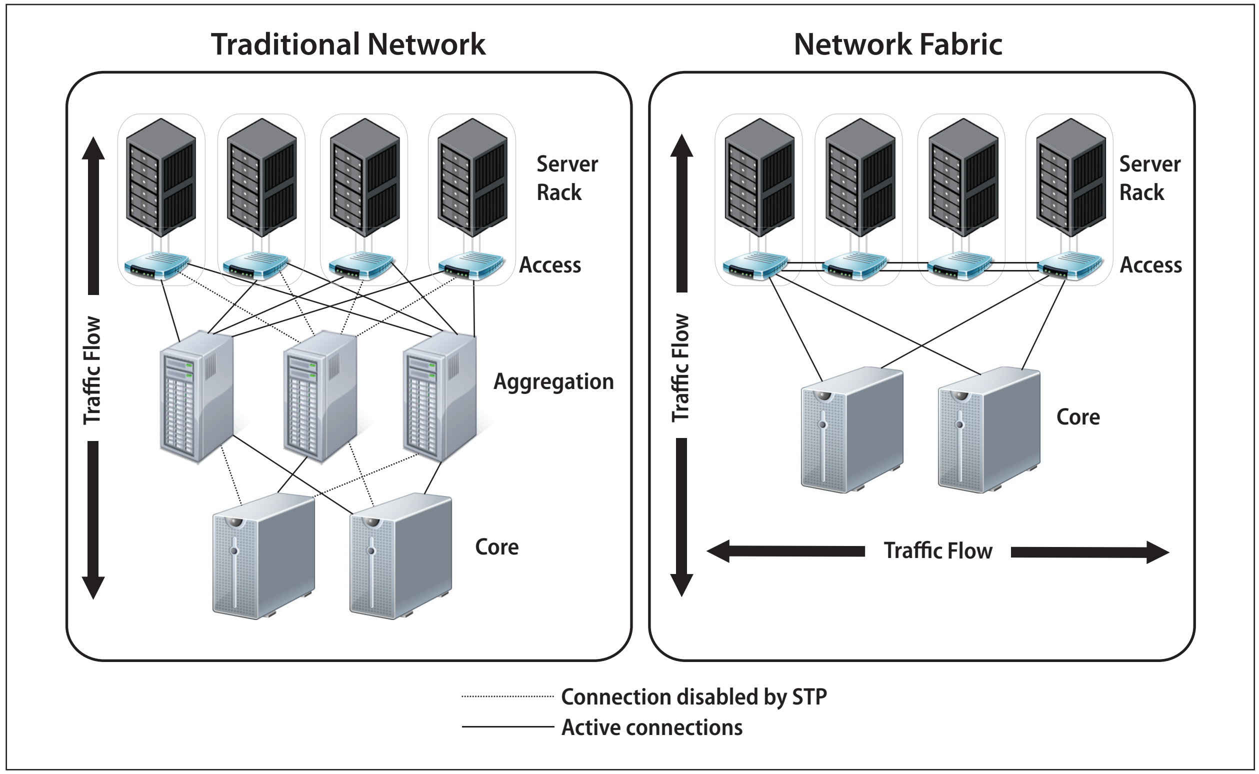

Source: ZK Research, 2014

ture must be capable of supporting both options.

> **Optimization for blade switching:** If the data center is using a blade server chassis with active blade switches, a top-of-rack switch may still be required based on the amount of east–west traffic. This preserves the three-tier data center architecture. But if the blade switch is strictly doing "pass through," then the top-of-rack switch is the first network device the com-

pute traffic encounters, driving more east–west traffic in a two-tier architecture. And if there is a minimal amount of east–west traffic, the optimal solution is to route traffic from the blade switch tier directly to the core and eliminate the top-of-rack switch. These are all viable options based on traffic flows, and the network needs to support all of them.

> **Active/active switching:** STP is widely used in data center networks today to prevent network traffic loops—that is, traffic that takes a circular path and does not reach its destination. This can happen when multiple links are used between network switches for redundancy purposes. STP operates by disabling one of the two active connections and putting the second link in a passive standby mode until the primary link fails. Although STP does achieve the goal of avoiding routing loops, it is a highly inefficient protocol because up to half of the ports could be disabled. Ideally, the network would be active/active, meaning all ports are active and able to pass traffic simultaneously. This leads to a lower cost, simpler network that performs better. The active/active fabric needs to be a topology-independent solution, meaning it can be a fully meshed fabric or partially meshed network. The network also must be interoperable with any device supporting link aggregation (LAG) and must allow for sub-second failover and recovery for always-on access.

Ideally, the network would be active/active, meaning all ports are active and able to pass traffic simultaneously.

> **High-density 10 Gbps Ethernet (GbE) with seamless upgradability to 25, 40 and 100 GbE:** Speed matters in the SDDC. The bigger the pipe, the more information can be passed across it. However, companies often must choose between line cards with high density and

lower speeds, and high-speed cards with fewer ports. Ultimately, the network switch should support high-density 10 GbE and then provide an upgrade path to 25, 40 and 100 GbE when those standards are ratified and products are commoditized. Ideally, the network would be upgradable from 10 to 25, 40 and then 100 GbE without a forklift upgrade, allowing for minimal disruption and the best TCO scenario. It's worth noting that Ethernet fabrics move network traffic much more efficiently, making wire-speed performance less of a showstopper in data center networks.

> **Road map to converged storage:** Most data centers today have two networks: an Ethernet network for data traffic and a parallel storage network that uses Fibre Channel as its transport protocol. FCoE was developed as a way to bring these technologies together as a single combined network, providing not only infrastructure consolidation benefits but also operational efficiencies. However, based on a series of storage manager interviews, ZK Research estimates that the penetration rate of FCoE is still less than 5% of the total storage market today. The main reason FCoE adoption has been slow is that it requires a forklift upgrade of the storage network infrastructure, which can be very disruptive to a company. Because the incumbent technology, Fibre Channel, offers 16 Gigabit speeds and FCoE only offers 10 Gigabit speeds, the disruption to the business from the drop in speed and technology transition outweighs any other potential benefits of making the move to FCoE. Organizations that want to start enjoying the benefits of converged storage now should consider iSCSI as a practical solution today,

Exhibit Five: Multi-Tenant Network

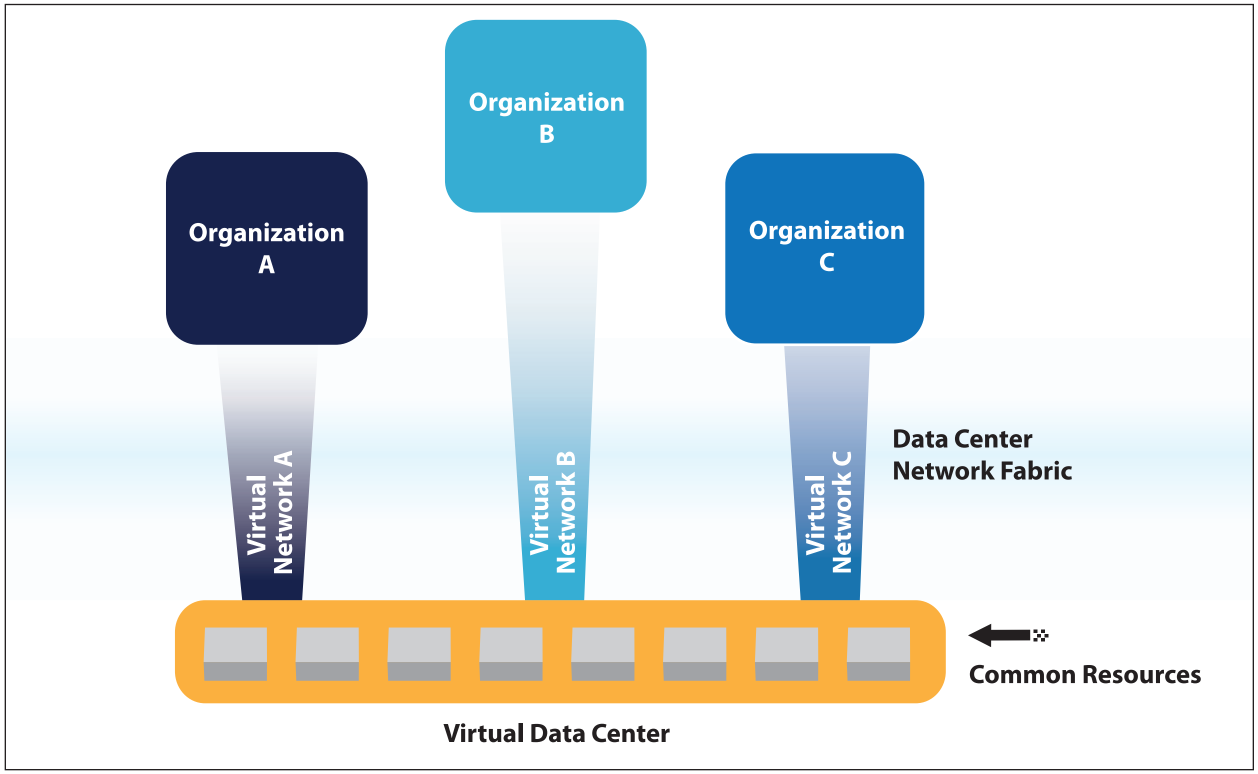

Source: ZK Research, 2014

and then look at FCoE as a potential long-term solution when 25, 40 and 100 Gigabit Ethernet speeds become mainstream. Customers that need storage convergence today can look at options such as ATA over Ethernet (AoE), iSCSI and network-attached storage (NAS).

> **Multi-tenant capabilities:** Many organizations require a network that can separate traffic. This could be the case when the IT organization supports many subsidiary companies or the busi-

ness units require it. This requirement is very common in verticals such as state and local government, education and financial services. In these situations, the network must be a single physical network that looks and acts like several logical networks, each with its own unique needs and requirements. To operate efficiently, the network must automatically separate traffic into distinct virtual networks but have access to a common set of data center services (see Exhibit Five).

To operate efficiently, the network must automatically separate traffic into distinct virtual networks but have access to a common set of data center services.

Section IV: What to Look for in an SDDC Solution Provider

The demands of the SDDC are driving network evolution faster than ever before. The network will play a key role in determining the ultimate success or failure of fulfilling the vision of an SDDC. This is a new role for the network; consequently, network decision-makers must shed their old-school buying strategies that are based primarily on brand or vendor incumbency and instead evaluate network infrastructure based on its ability to support the needs of the evolving data center.

However, enterprises might not know what to look for from a network solution provider, especially with the transition to virtual computing currently under way. Enterprises evaluating solution providers must consider the following factors:

Standards-Based Solutions

There are several ways for solution providers to meet the challenges of network evolution.

Many vendors choose to use proprietary protocols and solutions to develop new products. This can sometimes short-cut development time and enable vendors to get products to market early. However, in the long term, it causes vendor lock-in and impairs a customer's ability to choose best-of-breed products in the future. A standards-based solution guarantees interoperability with other best-of-breed products and ensures a wide variety of choices. For example, there are a number of possible replacement protocols for STP, but so far, only SPB has emerged as a viable standards-based option. Despite being itself standardized, TRILL has failed to engender any standards-based implementations, which may lead to interoperability issues in the future.

A Migration Path Away from Current Network Topology

The data center is the most important IT asset for the majority of organizations today. It contains all of a company's critical data, and any disruption could cripple an organization. Because of this, any architectural change must be based on a simple, reliable migration path and not a "rip and replace" strategy. The solution provider must be able to support the existing topology while providing immediate benefits where the new technology is deployed. For example, in a three-tier data center, a migration path might consist of first replacing only the core switching infrastructure. This would provide the core with the benefits of a network fabric where required without disrupting the compute access layer. Eventually, the rest of the network can be migrated at the organization's own pace. This allows IT departments to get comfortable with the new technology without disrupting business operations.

A unified management platform is just as important to the ongoing reliability and efficiency of the network as network hardware is.

A Fabric Solution That Extends Past the Data Center

Although the current trends in the data center create a compelling need for a network fabric that is initially data center–centric, organizations must consider a long-term strategy that extends the fabric outside the data center. Virtual machine mobility, IP-based storage and other trends require low-latency, seamless connectivity within the data center. Applications such as desktop virtualization, video conferencing and unified communications require the same levels of network performance—but all the way out to the access edge. Any network fabric solution deployed today must have the ability to extend into the campus and ultimately out to the access edge.

Unified Management

Network hardware is obviously a very important piece of the overall solution, but it's not the only one. A unified management platform is just as important to the ongoing reliability and efficiency of the network. A unified management tool needs to provide the following:

> **Application awareness:** The management tool must be able to tie network virtualization to application virtualization through the automated provisioning of the network to facilitate application-driven VM mobility.

> **End-to-end network management:** This includes end-to-end provisioning, monitoring and troubleshooting of virtualized and physical elements.

> **Performance and fault management:** Key functions here are topology discovery, fault correlation and network discovery. Ideally, the management tool provides these functions in a multivendor environment.

>**IP flow management:** This management tool must be able to monitor network flows for traffic analysis, capacity planning and usage analysis.

> **Wizard-based configuration management:** This provides network managers with a step-by-step method of streamlining the configuration and provisioning of the network, significantly reducing the number of outages due to human error—currently the largest cause of downtime in networks today.

> **Management and Orchestration (MANO):** A key requirement of this management tool is to orchestrate the configuration of virtual services across the network fabric. This simplifies many of the complex tasks needed to create the fabric, such as configuration of an SPB infrastructure.

Network Infrastructure Optimized for the SDDC

Many network products in data centers today were designed in the early 2000s or even the late 1990s. The network infrastructure for the virtual computing era must be designed and

The shift to SDDCs is transforming the data center faster than at any other time in the history of IT.

optimized for the unique demands of the evolving data center. Key criteria for the network infrastructure include the following:

> **Support for software-defined networking:** Although the majority of organizations have yet to begin an SDN deployment, SDNs will play a significant role in an SDDC over time. Any network infrastructure being deployed in a data center today should enable a path to SDN including support for both programmability (e.g., OpenFlow) and orchestration (e.g., OpenStack, OpenDaylight).

> **Fully redundant hardware devices with no single point-of-failure for continuous operations:** This should not be an option but a standard part of the solution.

> **An OS designed with the data center in mind:** A few network vendors have been building solutions that use the same OS across all network devices. However, the needs of the data center are significantly different from those of wiring closets and branches, and the OSs should be optimized accordingly.

> **Future readiness:** The solution must allow for seamless upgrades to higher speeds and greater throughput when the needs of the organization dictate.

Section V: Conclusions and Recommendations

The shift to SDDCs is transforming the data center faster than at any other time in the history of IT. Much of that shift is being driven by the increased use of virtualization. SDDCs have already had a significant impact on the software and server industry, and they will have a similar impact on networking because the network plays a critical role in mobilizing virtual workloads. An effective SDDC will enable companies to allocate virtual IT resources to the applications and services that require them on demand, creating a truly agile IT infrastructure.

Such a data center positions the network to become the key point of competitive differentiation for organizations as they look to capitalize on the flexibility and cost efficiencies of virtualization. This introduces new requirements for network design and the choice of solution provider.

To realize its full potential, the network needs to undergo a major transformation. Customers must move their decision criteria away from elements such as market share and vendor incumbency. Instead, they should focus on factors such as how resilient the network is, how the network simplifies architecture and how it enables virtualization to be used more broadly. With this in mind, ZK Research recommends the following steps to help companies evolve to an SDDC:

1. **Remove the complexity from the network.** Legacy networks are far too complex to support the vision of an SDDC. Shift to a simple, fabric deployment in which many of the points of complexity have been removed. A software-defined network can centralize the control and management of the network, simplifying network operations.
2. **Automate data center operations.** After the network has been simplified and control

has been centralized, network professionals should automate as many standard operational tasks as possible. This will speed up configuration changes, enhance time-to-service, reduce the amount of downtime and enable the network operations team to focus on more strategic initiatives.

3. **Evolve applications to be network aware.** The next step in the shift to an SDDC is to ensure that the applications are aware of the network and the data that's available.

4. **Enable real-time application feedback into the controller.** Once the applications have become network aware, the application should send feedback into the controller through the northbound APIs. This can help optimize application performance with no human intervention required. At this point, APIs need to transition from being simply "northbound" (i.e., by implication, only unidirectional, from network to application) to being "northside" (i.e., supporting communications bidirectionally between network and applications).

Ultimately, the majority of operational tasks in a software-defined data center should be automated. This can significantly reduce the operational costs of running a data center, which accounts for 40% of the overall TCO.