PATHFINDER REPORT

451 Research | Advisory

The Future of Infrastructure and the Software-Defined Organization

SEPTEMBER 2018

COMMISSIONED BY

SUSE
We adapt. You succeed.

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events.

NEW YORK
1411 Broadway
New York NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
(Ground floor)
30, Artillery Lane
London, E1 7LS, UK
P +44 (0) 203 929 5700

BOSTON
75-101 Federal Street
5th Floor
Boston, MA 02110
Phone: +1 617.598.7200
Fax: +1 617.357.7495

Executive Summary

Ever since virtualization became popular at the start of 2000, creating a new layer of software that abstracts a given operation from its underlying hardware, the move toward software-defined infrastructure (SDI) has been the direction of travel for enterprise and service provider IT organizations alike. The continued development of abstraction introduces considerable potential for wholesale SDI deployment, allowing a radical rethinking of infrastructure and leading to marked improvements in areas such as flexibility, accessibility and efficiency.

Software-defined infrastructure is all about adaptability and getting rid of the rigid system silos that are now starting to hold companies back from digital transformation. A cornerstone of SDI is its ability to adapt an organization's infrastructure to better support operational agility through automation and flexible resource scaling, thereby making infrastructure simpler to manage and its consumption more cost-effective. We see a core set of essentials as elements in the moves toward the software-defined organization:

The Eight Essentials of Software-Defined Technology

1. Software-defined components, many of them open source, are gaining far greater importance within the broader software infrastructure stack.
2. Compute, storage, networking and security are converging, with management and orchestration providing a single point of access, automation and control.
3. Software-defined projects often start from the top down, requiring customers to assemble all the layers below so that applications match the hardware requirements. A platform can reduce the integration burden and make life easier.
4. Flexibility is a priority. Legacy systems have to run in parallel with 'cloud native' applications and DevOps tools.
5. New infrastructure platforms must be consumable in multiple forms – bare metal, VMs, containers, either on-premises or in the private or public cloud.
6. Hybrid IT must become more elegant. It requires a unified, software-defined control plane, bridging the public and private worlds seamlessly.
7. Developers should be able to build and deploy their applications anywhere. They need a hybrid workspace that supports traditional workloads in VMs and modern apps in containers, and flexibility to deploy across bare metal, private and public cloud.
8. As the data explosion progresses, data efficiency, resiliency, management and mobility are all key requirements that should be abstracted away from the underlying storage and made available across the hybrid platform.

The endgame is a digital infrastructure that offers a better way of tying on-premises resources more closely together, one that incorporates existing legacy applications and infrastructure more inclusively, and one that is prepared for the demands of next-generation workloads, such as artificial intelligence and the Internet of Things (IoT).

Infrastructure is Becoming 'Invisible'

The traditional role of the operating system is changing. Once closely linked to a specific hardware platform, the OS is now part of a broader stack of software components, including the hypervisors and containers, which abstract the applications from the underlying infrastructure. It's the software that 'defines' the capabilities of the underlying infrastructure, which is surely the right way to think about it. After all, the sole purpose of hardware infrastructure is to support the application.

Multiplied even further by the recent explosion of interest in hybrid cloud strategies, this abstraction introduces considerable potential for a radical rethinking of infrastructure in areas such as flexibility, accessibility and efficiency. The concept of SDI extends the abstraction into the previously separate silos of networking, storage and security, providing the basis for new levels of agility through automation and flexible resource scaling. Infrastructure becomes easier to manage, favoring generalist (easy to provision) administrators over specialist (expensive to employ) skills, and its consumption more direct, immediate and cost-effective. Given the pressures and demands of accelerating data growth and the need to accommodate next-generation workloads, even in the face of operational cost cutting, the simplification and automation of infrastructure is becoming a business necessity, and not just a desirable feature.

As Figure 1 illustrates, the workload environment is in a state of transition. Workloads are migrating to new platforms. Private cloud (on-premises and hosted) expands as the future deployment venue for the majority of IT workloads (28% in 2018; 38% in 2020), but public cloud is expanding faster (17% in 2018; 30% in 2020). Organizations are evenly split between the modernization-in-place (on-premises) approach and strategies involving shifting to off-premises cloud/hosted environments. 'Software defined' is a key element for both.

Figure 1: Workload deployment venues, 2018 and 2020

Source: 451 Research Voice of the Enterprise: Digital Pulse, 2018

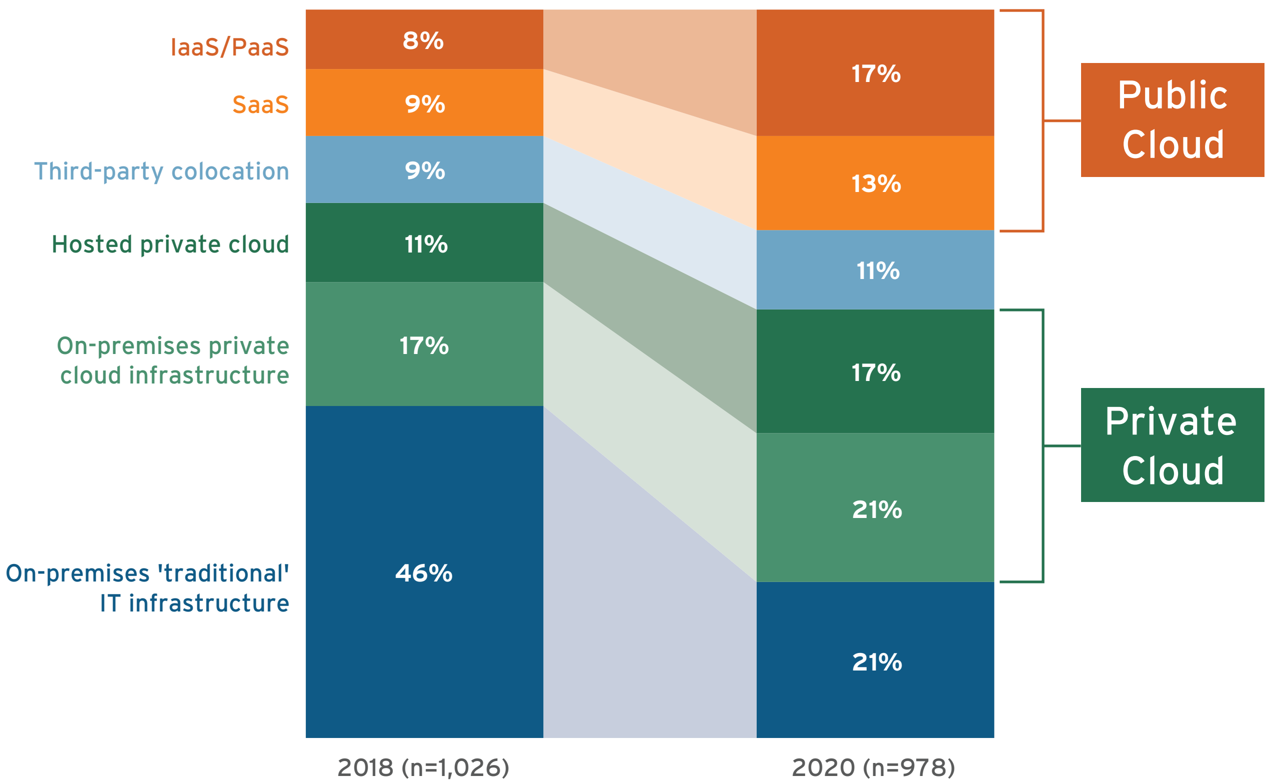

In our view, infrastructure will become invisible over the next few years. Driven by the cloud's consumption-based, service-driven model, the principles of invisible infrastructure are already reshaping the IT and communications landscape. Over the next decade, they will affect the full range of infrastructure elements, including hardware and software, hosting and managed services, networking and telecommunications, and datacenter technologies. Technology consumers are demanding infrastructure that 'just works': instantly available yet always invisible, operating and scaling regardless of specific requirements, and billed and metered in whatever manner the customer prescribes. Successful service providers – internal IT departments as well as external cloud services – will be those that can deliver an 'invisible' experience to clients across a growing spectrum of increasingly sophisticated workloads and applications. Abstracting away this technology filter means transformation projects based upon business outcomes can become a reality.

The Road to The Software-Defined Organization

OPEN SOURCE, VIRTUALIZATION AND THE CLOUD

Open source software has rapidly become a mainstream strategy for reducing cost and development times, maintaining vendor independence and increasing agility. Pre-virtualization, it had been common to dedicate an entire server to one application to avoid conflicts that could severely affect performance. Virtualization has taken the abstraction layer further, lowering dependence on a single operating system, chip architecture or individual server form factor.

The cloud hyperscalers have optimized virtualization for commodity scale-out system architectures. Before Google, mainstream IT ran complex applications (transaction processing, databases) on top of complex infrastructure. Google's software stack was much simpler, primarily geared for running its search engine at scale. It didn't require complex features such as built-in redundancy, so Google was able to simplify its hardware, building it from the ground up using commodity components – to begin with, off-the-shelf motherboards bought from the local Fry's Electronics store. By taking out all the specialist elements, stacking up its bare boards into racks and flattening its infrastructure into a uniform, scale-out distributed architecture, Google achieved the scale it needed at a fraction of the cost. With a consistent and fully scalable hardware platform in place, it set about optimizing its software stack, based on open source operating systems and hypervisors, to run on top of it.

The other giant hyperscale service providers followed a similar course, and most of them built their own systems (or partnered directly with original design manufacturers) rather than buying from traditional vendors whose production systems were mostly geared for enterprise IT environments and still loaded with expensive enterprise feature sets. The systems vendors saw how things were going and started offering their own stripped-down, scale-out systems to win business from the second and third tier of service providers.

OPENSTACK AND THE OPEN COMPUTE PROJECT

The giant hyperscale companies were big and rich enough to build their own software stacks. For smaller players, the alternatives were limited to commercial virtualization layers (expensive) or open source (fewer features and not well supported). There was clearly a gap in the market. Enter OpenStack, which in 2010, emerged from a joint project between Rackspace and NASA. It was intended to provide an open source cloud computing platform for both private and public clouds, although its initial complexity restricted its usage to those with sufficient technical skills to implement it.

As a cloud operating system, OpenStack can control large pools of compute, storage and networking resources throughout a datacenter, centrally managed by administrators but also giving users self-service access to resource provisioning through a web interface. It was built from a series of separate open source projects with common application programming interfaces. The best-supported projects are Nova (compute), Neutron (networking), Swift (object storage), Glance (image service), Keystone (identity service) and Cinder (block storage). There are also data and analytics modules (Trove, Sahara), management tools (Horizon, OpenStack Client), application orchestration (Heat) and a range of monitoring tools.

The emergence of OpenStack led to the establishment of other open source initiatives at the hardware level, most notably the Open Compute Project. OCP, which originated from a Facebook project to handle scale requirements in a cost- and energy-effective manner, was launched in 2011. It established standard designs for compute, storage and networking elements, opening up the initial design specs for industry collaboration.

Some early adopters of OpenStack and OCP found unexpected complexities, as well as integration and support issues, and these undoubtedly slowed down mainstream adoption. However, the central idea of both initiatives – that the value and deployment of software should not be constrained by over-dependence on specific hardware platforms – was firmly established under the mantle of 'software-defined infrastructure.' This significant shift in perception was helped along by several factors. The rapid transition to cloud and hybrid cloud deployment models was the most obvious. But there was also a new excitement over the adoption of containers and microservices, initially for DevOps but soon also for more flexible application deployment. And there was an increasing realization among end-user customers that to effectively support newly emerging cloud-native applications and workloads, an equally new approach to infrastructure would be required.

THE EVOLUTION OF TRADITIONAL HARDWARE

The primary system platforms that software-defined infrastructures expect to run on are scale-out clusters of commodity servers, storage and networking, with most of the specialist features not included, therefore transferring the intelligence to the software layer. These hyperscale-inspired architectures are now moving beyond Google and Facebook into the adjacent sectors of enterprise datacenter services and high-performance computing labs.

However, that's not the whole story. Traditional systems vendors have also been working to reinvent the feature-rich enterprise servers that have been at the center of their product portfolios (and the source of their highest profit margins) for decades. Converged infrastructure systems are in many respects the successors to the scale-up symmetrical multiprocessing enterprise systems of the past. They offer pre-integrated packages of servers, storage and networking, unified software management, and lifecycle management that are certified to work together – and which will continue to work together when individual components are upgraded or patched. This maintains compatibility and compliance, avoids potential downtime and keeps performance optimized so that the original SLAs can be maintained.

Lighter versions of converged infrastructure have emerged in the form of modular building blocks of combined server and storage units, accessed at the hypervisor level. This is hyperconverged infrastructure (HCI), which has found initial traction in small and midrange businesses. HCI is gradually being scaled up and expanded beyond storage to support more demanding enterprise applications and larger installations, but at the expense of some of its original simplicity. The next stage in convergence looks like disaggregation, where pools of resources become available as dynamically 'composable' resources, put together in real time to match the demands of application workloads. From here, specialist resources can rejoin the mix – memory-driven architectures that will run memory- and I/O-hungry applications at maximum performance, or custom accelerators for security, analytics, deep learning and artificial intelligence.

Figure 2: The evolution of invisible infrastructure

Source: 451 Research

COMPOSABLE SYSTEMS	Enterprises are currently voting to bypass the extra internal integration work required by best-of-breed strategies. They favor the pre-integration and simplicity that comes from convergence and its successors. Composable systems take convergence a stage further, providing pre-integrated access to dynamic, automated resources that can react in real time to demand.
MEMORY-DRIVEN ARCHITECTURES	Various forms of non-volatile 'storage class' system memory are under development and on their way to market, enabling apps and analytics to be run entirely in memory. Immediate performance benefits ensue, but over time, the availability of new levels of performance and capacity memory will have a profound effect on the way the software layers are architected.
CUSTOM ACCELERATORS	The recent visible breakthroughs in high-profile speech and image recognition apps, along with renewed interest in AI, are driving a push from the investor and developer communities toward custom silicon accelerators that can run specific workloads with order-of-magnitude performance boosts. This is helping mitigate the slowing down of Moore's Law.
HETEROGENEOUS COMPUTING	All the above trends imply a greater diversity of underlying infrastructure, largely transparent to the software layers above so that 'best execution venue' strategies – the matching of workloads to the most appropriate resources – can be implemented.

These two trends – the first toward less diversified infrastructure based on commodity, bare-bones scale-out systems, the second more of a scale-up, heterogeneous approach implying a greater diversity of underlying infrastructure – may seem to be in opposition. Ultimately, however, they are heading in the same direction, where resources become largely transparent to the intelligent software layers above, directing each workload to the best execution venue. Disaggregated compute, storage, networking and memory resources will be available in fluid, pooled form, rather than rigidly shackled to specific processors or systems as they mostly are today. Specialist accelerators and processors (such as GPUs and FPGAs to support artificial intelligence) will be available for transparent attachment to general pools.

THE SOFTWARE: ELEMENTS OF 'SOFTWARE DEFINED' TECHNOLOGY

The term 'software defined' implies a broadening out of first-wave virtualization adoption beyond compute and into adjacent domains, primarily networking and storage. A software-defined infrastructure framework or cloud platform will offer access to physical servers, storage and networking switches through various software layers: the OS, (Linux or Windows); hypervisors, (KVM, Xen, VMware, Hyper-V and others); storage platforms (Ceph, vSAN and multiple others); and software-defined networking (VMware NSX, Cisco ACI, OpenDaylight, etc.). There will also be management components (operations, DevOps, cluster deployment, orchestration) and perhaps an application delivery layer in the form of PaaS.

For **compute,** we already have server virtualization, which aims to simplify the task of running multiple workloads on fewer physical boxes. Hypervisors and virtual machines are the most common way to consume compute resources through cloud services or software-defined frameworks. However, bare-metal (i.e., single-tenant) servers have made something of a comeback where high performance and predictability are factors. Over the past few years, containers have also seen rapid adoption as a lightweight alternative (or sometimes a supplement) to virtualization (see Figure 3). Containers package up just the elements an application needs to run (runtime, dependencies, libraries, binaries, configuration, etc.) rather than the full OS, as in a virtual machine, making them more portable and modular and less resource-hungry, and they can be used as the basis for delivering applications as a set of microservices.

Container adoption is accelerating at a fast pace, with nearly half of 451 Research's Voice of the Enterprise respondents either currently using containers or planning to use them in the near future. A significant component of software-defined infrastructure, containers enable the support of next-generation workloads through increased automation, orchestration and seamless data mobility.

Figure 3: Adoption status for containers

Source: 451 Research Voice of the Enterprise: Servers & Converged Infrastructure, Workloads & Key Projects 2018

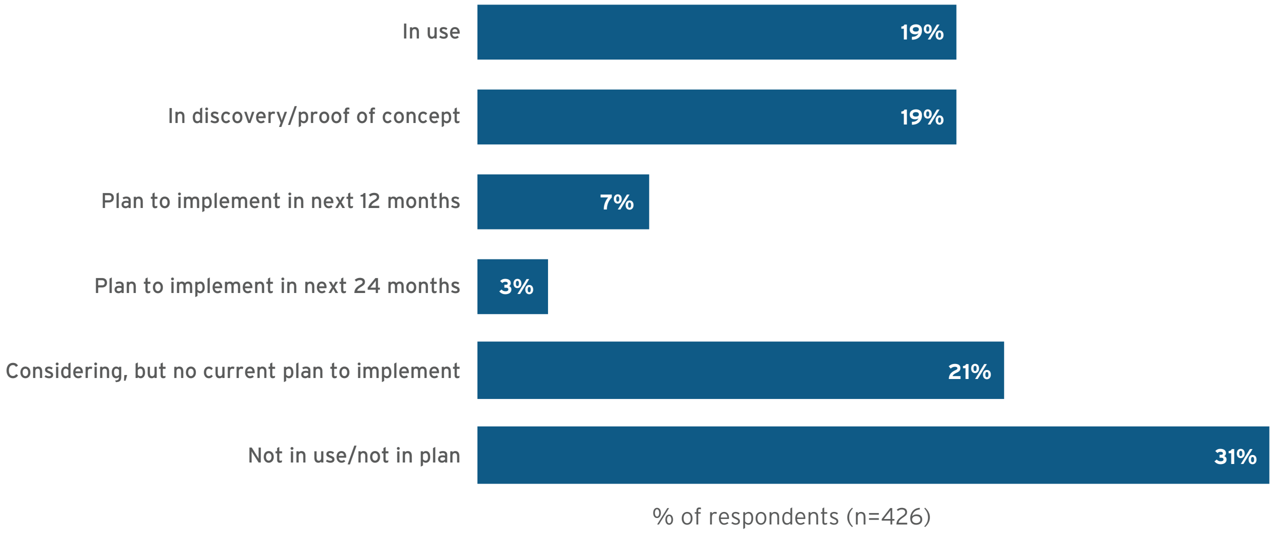

In the **storage** world, the classic monolithic array is now more than three decades old and often expensive, complex and stove-piped. Since running storage traffic over Ethernet became possible, however, advancements such as network-attached storage (combining storage and file systems), unified storage (i.e., file- and block-level convergence), storage virtualization, and scale-out storage software (pooling the internal storage resources of multiple x86 servers) have begun to change the situation. Solid-state storage and tiered caching schemes have also helped the process of placing storage resources closer to the compute node, as has a new generation of integrated storage appliances, designed specifically for virtualized environments and increasingly using flash storage in place of spinning disks.

Storage virtualization began the process of breaking hardware lock-in by relocating core storage services away from embedded array controllers, allowing all capacity (from multiple systems vendors) to be consolidated and managed consistently through a central software layer. Various block, file system and object storage software layers have also been used as the basis for scale-out storage systems and cloud storage, and today, several open source components have momentum, including HDFS, Gluster and Ceph. Software-defined storage can provide numerous benefits over traditional arrays, but the primary areas of focus have been on scalability, consistent performance and agility. A software-only approach can leverage the rapid innovation of commodity hardware while cutting out the dangers of vendor lock-in through proprietary hardware. Cloud-native offerings also allow vendors to extend their data protection and management capabilities to data residing in a public cloud.

As for **networking**, physical on-premises switches have remained pretty much the same ever since Cisco introduced its first network switches in 1993, except that they now have more and faster ports. They are typically run at only a fraction of their potential capacity and weren't designed to cope with virtualized resources and dynamic, moving workloads. Also, the old-fashioned tiered networking approach is inefficient when it's applied to virtual and cloud infrastructure: there are just too many hops between one VM and another. To meet these challenges, vendor-neutral open source initiatives (such as OpenFlow and Open vSwitch) led the evolution to software-defined networking (SDN). Various flavors of I/O virtualization can potentially cut out three-quarters of network connections, divide bandwidth to suit demand and workloads, and provide an on-ramp to converged fabrics. Bringing more efficiency to the network makes the most sense for service providers selling VMs on to their clients – if they can increase VM density from, say, 20 VMs per server to 150 VMs per server, that can make a huge difference to profitability.

Network Function Virtualization

Network function virtualization and orchestration is a distinct market opportunity for SDN, and it is almost exclusively carrier-driven. The first installations are now reaching production, although most are still modest at this stage. While SDN is about the abstraction of the control and data planes from networking equipment, NFV focuses on the separation of network functions from hardware so that the functions themselves can be run as software. These functions can span the entire ecosystem of network applications, including load balancers, DPI, firewalls, policy servers, DNS, network address translation, IMS, etc. However, two applications in particular are key to the transition to NFV: virtualized evolved packet core (for core mobile network providers converging voice and data traffic on LTE networks) and virtual customer premises equipment (deployed at the edge of the fixed network to deliver network services to enterprises and consumers).

One of the main values of SDN is a simplification of the way in which networks are interconnected. In physical networks, lashing systems together means chasing physical cabling between chassis and boxes. While this offers a flexibility limited only by the creativity of cable routing, keeping track of what went where creates a significant burden of documentation, and once connections are in place, changing them requires overcoming formidable inertia. Network virtualization offers dynamic creation and the ability to change interconnections without physical access to systems, and enables VMs to remain independent of physical location and free from dependencies on the physical network infrastructure. To achieve this, extensions to the networking stack must provide VM-level visibility and granular network configuration and control. VxLAN, a means of supporting enough virtual LANs through the use of overlay networks underlying large multi-tenant clouds, is a primary example.

Finally, the **management** of software-defined infrastructure requires an independent management platform to bring together the separate domains, directly analogous with cloud management platforms. The aim is to provide a common framework for enterprises to manage many distributed and disparate hybrid cloud and IT infrastructures, without asking administrators to constantly jump between screens or tools. There's an added advantage when the management platform shares a common set of APIs with the OpenStack platform, enabling tighter integration. Users need to be able to make data-driven decisions when selecting cloud applications and services, including (but not limited to) cost, compliance, utility, governance and auditability. Management platforms able to support a mix of old- and new-style IT services can also act as a bridge to help address the challenges of moving to cloud.

The Software-Defined Organization

Why should enterprises and service providers take notice of the industry's growing emphasis on software-defined technologies? A positive impact on business value must be the end goal. The cornerstone of software-defined technology is its ability to adapt an enterprise's infrastructure to better support operational agility through automation and flexible resource scaling, thereby making infrastructure simpler to manage and its consumption more cost-effective. Adoption rates are closely related to other trends, such as the popularity of converged infrastructure appliances and the migration toward public cloud services. For instance, avoiding vendor lock-in (either on-premises or in the cloud) should be a business imperative, not just an option. If only for fiduciary responsibility, organizations are adopting multiple cloud services. This is especially true for large enterprises, which spend millions of dollars a year on cloud services, where cloud spending is a visible line item on the balance sheet, and where executives want to be certain of getting the best option available at the lowest cost.

Organizations are looking for better ways of tying their on-premises resources more closely together and incorporating their legacy applications and infrastructure more inclusively. Workloads, data and processes are becoming increasingly fluid across multiple on-premises, hosted, private and public cloud services, requiring a new approach to hybrid cloud management – one with a uniform means for access control, billing and provisioning, capacity management, cost control and performance analysis (among others). Enterprises are demanding that IT vendors craft a holistic platform to allocate workloads strategically to the best execution venue and do so while managing business continuity across the entire hybrid system. The demands of accelerating data growth and requirements to get ready for next-generation workloads, such as artificial intelligence and IoT, add another level of urgency, compounded by a continuing increase in the population of generalists over specialists in the workforce.

A software-defined infrastructure platform should reduce the integration burden and, in theory, make life much easier for the customer. Those increasingly hard-to-find specialist storage and networking experts will become less essential over time, giving way to more generalist administrators working with higher-level tools. However, many software-defined projects start from the top down, requiring customers to find all the software drivers required for applications that match the platform's hardware requirements. And the skills gap won't be solved overnight. That's where trusted business partners come in; they can help clients achieve their goals without having to invest too heavily in increasing headcount. And a platform must be future-proof, with no barriers to the adoption of emerging and future technologies. Platforms continue to transform at a rapid pace. Software-defined infrastructures must leave room for the incorporation of new APIs, microservices, hybrid computing, memory-based computing and new silicon workload accelerators, as well as the extension of intelligence to the edge though the Internet of Things.

Next Steps and Conclusions

Platforms continue to transform at a rapid pace. APIs have made it much simpler for one application to talk to another without heavyweight hardwiring. Now smarter, unified APIs are codifying increasingly sophisticated aspects of data integration across both the development and operations domains, encouraging the development of open partner ecosystems. Microservices allow a similarly agile approach to integration: A developer can make changes to one of 20 separate microservices in an application without having to worry about breaking the whole set. Built on top of containers, clearly defined microservices, along with API-based integration, will enable more advanced technologies – such as machine learning – to be dropped into applications and optimize specific functions more easily, letting businesses react faster to changes in the market. Meanwhile, cutting-edge developments such as memory-driven computing and edge servers to support IoT are on the near horizon.

Changing developer and IT admin best practices show the broader impact of these technology advances. 'No-ops' hybrid IT management is becoming a destination point: A converged management console, service portal and app store can enable customers to compose, deploy and scale hybrid clouds that support all applications, and to manage production compliance and lifecycle governance. Developers can access components pre-validated by IT operations for scaling and regional compliance. On the IT operations side, the curated tools reduce risk, and the hybrid workspace furthers the links between the development phase and subsequent production deployment, governance/compliance and lifecycle.

Strong leadership with a disruptive attitude is necessary to guide an organization through rapid changes in technology, customer behaviors and business models. Trusted business partners with a proven history in helping businesses transform, with the relevant experience in software-defined platforms and tools, will likely prove invaluable. It is a journey that forces companies to move beyond business as usual to embrace organizational change and agile business processes and systems to support transformation. Digitally empowered customers are demanding new forms of engagement with organizations, driving digital transformation projects in the process. No longer willing to be 'sold to,' they are forcing companies to adopt the latest applications, analytics and infrastructure to deliver differentiated experiences. For these companies themselves, this is not a luxury, but a necessity for survival. Businesses need to embrace digital transformation if they are to improve, connect and radically change business processes, enhance customer experiences and harness innovation.