

White Paper

Optimizing Workload Placement in Your Hybrid Cloud

Sponsored by: Dell Technologies, Intel, and VMware

Eric Burgener
July 2020

Sriram Subramanian

IDC OPINION

Many IT organizations are in the process of modernizing their infrastructure, as part of either a digital transformation (DX) strategy or a desire to improve the efficiency, reduce the complexity, and/or lower the cost of IT operations. In conjunction with these efforts, organizations are deploying next-gen apps that are placing more stringent demands than ever on IT leaders to meet agility, performance, availability, and data management requirements. To meet these requirements, IT organizations are looking at cloud, accelerated compute, solid-state storage, distributed data, software-defined infrastructure, scale-out architectures, and artificial intelligence and machine learning (AI/ML) technologies.

Cloud technologies figure prominently in modernized infrastructure. Cloud platforms deliver increased IT agility, provide on-demand access to IT resources, and move organizations to subscription-based pricing models that allow them to align expenditure much more closely with actual usage. Although public cloud may have initially been the way organizations began adopting cloud technology, IDC believes that organizations should start their modernized infrastructure journey by crafting a hybrid cloud strategy first. A hybrid cloud environment delivers traditional on-premises and private and public cloud deployment models under a unified management dashboard, delivering a common operational experience across all platforms and environments. A hybrid cloud strategy provides the most flexibility for optimal workload placement because it offers all three options – each of which offers some unique advantages.

For each workload at any given time in its life cycle, there is generally one option that is better than the others. Organizations should start their modernized infrastructure journey by considering optimal workload placement based on workload usage characteristics and organizational objectives. For organizations crafting a successful hybrid cloud strategy, the ultimate goal is to make strategic workload placement decisions that meet service-level agreements (SLAs), enhance security, and maximize efficiency at the lowest cost.

In early 2020, IDC performed in-depth primary research on how organizations determine workload placement as they evolve through DX and what implications this has for modernized infrastructure objectives and strategies. Key findings include:

- Modernized infrastructure is viewed as a critical determinant of DX success by 91.1% of organizations, with roughly 70% of them involved in server, storage, and/or data protection refresh over the next two years.
- 95.7% of organizations have an executive mandate to leverage cloud technologies, and over half (53.8%) have an executive mandate to support hybrid cloud environments.
- The top considerations in determining workload placement between traditional on-premises and private and public cloud included security, performance, ease of management, availability, cost, compliance, agility, usage patterns, and data gravity.

As organizations optimize their existing IT infrastructures to gain access to the new technologies needed for DX, simplifying a potentially complex environment that includes products and services from multiple server, storage, data protection, and/or public cloud vendors that span on- and off-premises locations is also a consideration. Technology vendors that offer a wide range of well-integrated IT infrastructure solutions and deep partnerships can help cut through much of this complexity.

Our research shows that Dell Technologies' broad infrastructure portfolio offers servers, storage, data protection, and cloud-based services that deliver on modernized infrastructure requirements. The vendor's offerings also include converged and hyperconverged systems solutions, multiple flexible consumption models, and a unified control plane based on VMware technology as well as that of other vendors. Dell Technologies' products support the needed new technologies, and with its range of offerings, the vendor can deliver a consistent operational experience across traditional on-premises and private and public cloud deployment models. The availability of such a comprehensive set of technology solutions from a single vendor offers the opportunity to reduce complexity and simplify vendor interaction. Organizations interested in crafting a well-integrated hybrid cloud environment as part of their modernized infrastructure efforts will want to explore offerings from Dell Technologies and its key partners (which include Amazon, Microsoft Azure, Google, VMware, and Intel).

METHODOLOGY

This white paper uses data from a worldwide survey (IDC's *Modernized Infrastructure Survey*), conducted in January 2020, on modernized infrastructure decisions and workload placement strategies among organizations pursuing DX. The survey spanned multiple enterprise IT infrastructure areas including cloud, servers, storage, and data protection as well as converged and hyperconverged infrastructure. 900 IT directors, solution architects, and IT managers from multiple verticals including financial services, healthcare, discrete manufacturing, IT services, utilities, education, and telecommunications participated in the survey. Respondents were from multiple geographic regions including Australia, Canada, China, Germany, India, New Zealand, the United Kingdom, and the United States and included small and medium-sized business respondents as well as large enterprise respondents.

SITUATION OVERVIEW

IDC defines DX as the continuous process by which organizations leverage digital competencies to innovate with new, more data-centric business models; improve their internal workflows, products, and services; and drive disruptive but positive changes in the external ecosystems in which they compete. DX enables organizations to achieve increased business agility and productivity, and IDC research shows that DX success is strongly correlated with the deployment of modernized infrastructure. Modernized infrastructure refers to the process of adopting new technologies such as cloud-based services, accelerated compute, nonvolatile memory express (NVMe) and the emerging persistent memory products it supports, software-defined infrastructure, artificial intelligence and machine learning, and scale-out architectures as IT infrastructure evolves to meet the needs of DX. More than 91% of survey respondents rated modernized infrastructure as a critical determinant of their DX success.

Cloud technologies figure prominently in today's modernized infrastructure. Organizations are increasingly leveraging cloud technology as part of modernized infrastructure, as evidenced by an increasing percentage of enterprise spend going toward cloud-based infrastructure. 95.7% of survey respondents indicated that they have an executive mandate to move to cloud-based environments – a directive that

includes both private cloud and public cloud. Public cloud platforms enable organizations to outsource much of the infrastructure management responsibility and take advantage of elasticity, easy scalability, and self-service provisioning capabilities. They enable application developers to provision new instances and consume native cloud services on demand, which in turn increases their productivity and ability to innovate. Private cloud platforms deliver the cloud experience while meeting certain other requirements that can be a challenge for public cloud.

Most organizations begin their cloud strategy by investing heavily in public cloud, but 88.9% of survey respondents plan to also deploy private cloud infrastructure in the next 12 months. Public cloud is an important part of an effective IT modernization strategy, but for many workloads, public cloud environments just can't meet important performance, availability, security, compliance, governance, and cost requirements. And even if public cloud may initially be a good match, workload characteristics may change over time. Many organizations surveyed had repatriated workloads from public cloud providers back into their own on-premises infrastructure, citing security, the recent deployment of modernized in-house infrastructure, unpredictable public cloud costs given actual usage patterns, and regulatory mandates as the top reasons for repatriation. Research results like this are why IDC believes that most organizations will ultimately have hybrid cloud environments that will include traditional on-premises infrastructure, on-premises private cloud infrastructure, and public cloud infrastructure. Each approach can offer different advantages, and organizations will seek out the optimal operating model for each workload (a decision that may need to change over time as workloads evolve).

IDC believes that most organizations will ultimately have hybrid cloud environments that will include traditional on-premises infrastructure, private cloud infrastructure, and public cloud infrastructure.


IDC defines hybrid cloud as the usage of IT services (including IaaS, PaaS, and SaaS) across one or more deployment models using a unified management framework. Hybrid cloud platforms include combinations of traditional IT (on-premises), private cloud, and public cloud infrastructure deployments managed by a consistent, unified dashboard. Through hybrid cloud platforms, organizations can place workloads on the most appropriate infrastructure – on-premises and private or public cloud – based on workload requirements and organizational priorities. More than 53% of respondents indicated their intent is to deploy hybrid cloud infrastructure consisting of all three deployment models, and this percentage is expected to increase over time. There is no doubt that, going forward, hybrid cloud will be the optimal IT infrastructure deployment model for most organizations.

Key Objectives in Undertaking Modernized Infrastructure

The demands of DX drive heightened requirements in enterprise infrastructure that can be difficult for legacy environments to meet. Those challenges can include agility, performance, availability, security, and the ability to optimally leverage cloud-based technologies. The objective of modernized infrastructure is to perform a technology refresh to meet these requirements – that is why infrastructure modernization is considered a key determinant of DX success for so many businesses. Organizations have very clear objectives when deploying modernized infrastructure, with the top 4 "most desired technologies" being cybersecurity capabilities, cloud-based data protection, accelerated compute, and solid-state storage (see Figure 1).

FIGURE 1

Most Desired Technologies for Modernized Infrastructure



n = 900

Source: IDC's *Modernized Infrastructure Survey*, 2020

As part of DX, organizations are widely deploying next-gen apps that not only feature improved cybersecurity capabilities, more agile cloud-based data protection, and technologies such as accelerated compute and solid-state storage as well as containers and their associated orchestration platforms (Kubernetes) but also leverage software-based, scale-out design tenets. With privacy a top concern for all IT organizations surveyed, cybersecurity is a critical capability and, in many cases, can be a reason to move to the cloud (when public cloud security capabilities outstrip those of smaller organizations) or, more often, to locate workloads in on-premises infrastructure (when public cloud cannot meet stringent security requirements). Cloud-based data protection offers ease of use, scalability, and efficiencies not available with traditional infrastructure. And with both business and consumer constituencies increasingly expecting real-time response and custom-tailored interaction with suppliers, next-gen apps are often built around performance-oriented technologies such as accelerated compute and solid-state storage media.

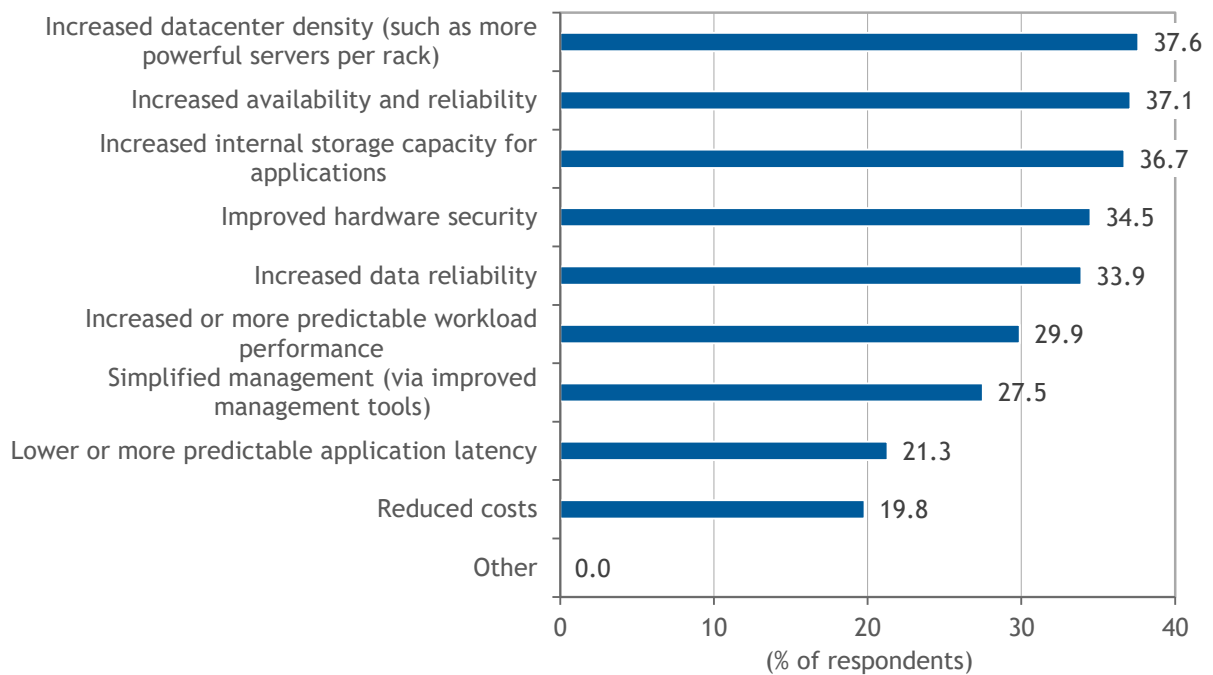
Primary research performed by IDC in early 2020 sheds additional light on objectives and popular technologies deployed as part of modernized infrastructure efforts by organizations undergoing DX. When asked about general strategies for modernized infrastructure, most organizations that deploy a specific strategy are moving most on-premises workloads to software-defined, server-based architectures as they retire legacy infrastructure. IT infrastructure based around software-defined designs offers greater agility and flexibility to meet a variety of different workload requirements. Interestingly, the second most followed general modernized infrastructure strategy focuses more on individual workload requirements, but even these organizations also sought to move to software-defined, server-based infrastructure where they could. 85.1% of organizations deemed software-defined flexibility as either very important or extremely important to their modernized infrastructure efforts.

Server Refreshes Driven by Increased Performance Density Requirements

The top 3 drivers of server infrastructure refreshes are increased infrastructure density, improved availability that would help organizations reduce downtime, and increased server-based storage capacity (see Figure 2). Given these objectives, it's not surprising that 78.3% of organizations preferred to move to one of several newer server architectures rather than just buying more powerful versions of their legacy servers. Converged and hyperconverged infrastructure in particular give organizations looking for improved IT infrastructure density an excellent option for many workload types. Across the board, organizations implementing server refreshes saw strong performance, capacity, manageability, and cost benefits and gained access to a variety of crucial new hardware benefits not available on their legacy servers.

FIGURE 2

Drivers of Server Refreshes



n = 603

Source: IDC's *Modernized Infrastructure Survey*, 2020

Server refreshes can bring a variety of modernized infrastructure benefits:

- Increased performance density was widely achieved, with 78.6% of survey respondents achieving performance improvements of 25% or greater.
- Denser servers enabled denser workload consolidation, allowing organizations to streamline infrastructure going forward. In the survey, 69.6% of organizations achieved datacenter density improvements of 25% or more.

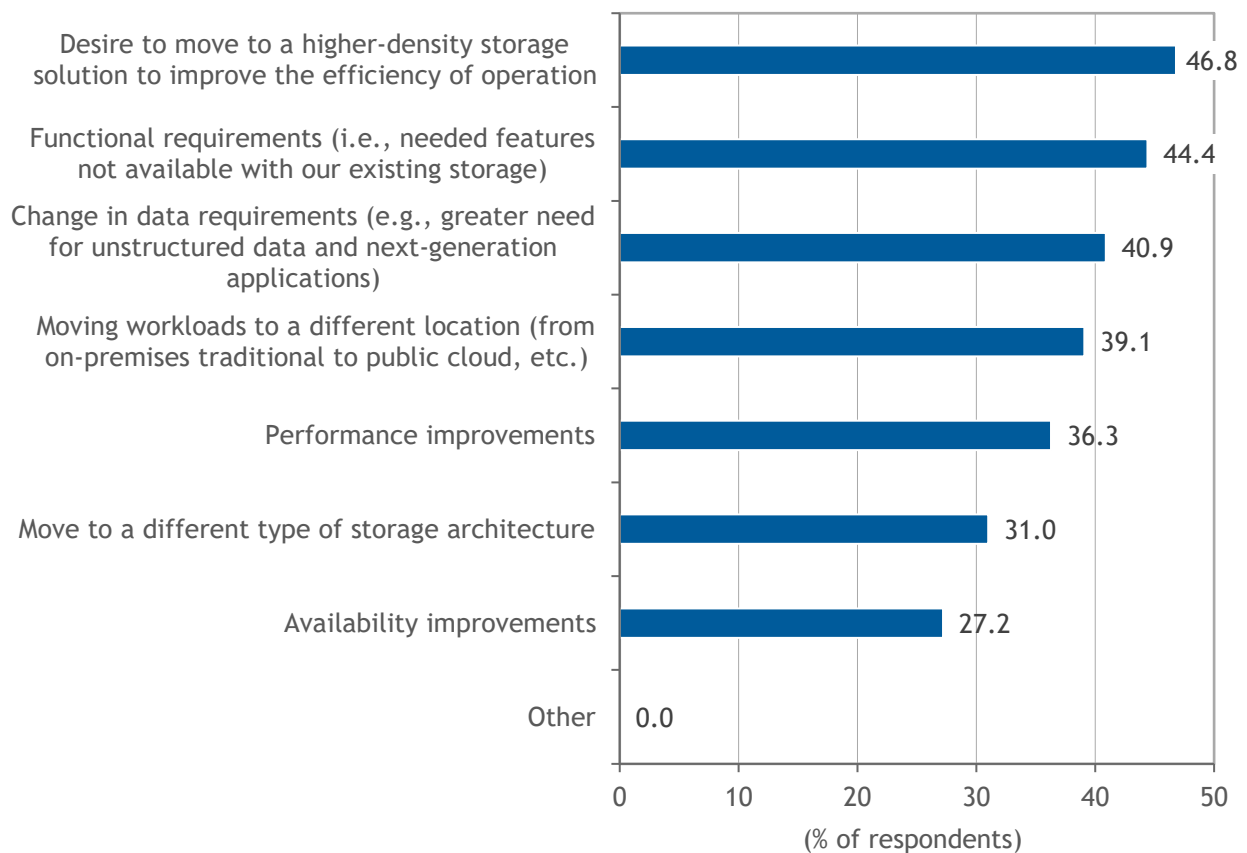
- Respondents also called out hardware benefits they deemed most crucial as part of their server refresh. Better access to compression and encryption technologies such as hardware assist through coprocessor designs, better security features (including certifications, multifactor authentication, secure key management, and Root of Trust capabilities), and more comprehensive health status monitoring capabilities topped the list of their most crucial hardware benefits with a server refresh.

Storage Refreshes Driven by Increased Capacity Density Requirements

With DX, a much greater amount of data must be captured, stored, protected, and analyzed, and this demands a much more real-time and distributed orientation for workloads as well as data. With storage in particular, organizations are interested in increased capacity density and better data management at scale. Among the top drivers of storage refreshes are a desire to move to more efficient, higher-density storage solutions; access to new enterprise-class data services not available with existing storage; and performance improvements (see Figure 3).

FIGURE 3

Drivers of Storage Refreshes



n = 614

Source: IDC's *Modernized Infrastructure Survey*, 2020

For the survey respondents, storage infrastructure improvements were noticeable:

- 82.2% of respondents improved performance by at least 25%, with 44.1% of respondents improving it at least 50%.
- Newer equipment was much more efficient – 77.8% of respondents reduced their energy and floor space consumption by at least 25%.
- While 69.3% of respondents manage their strategic workloads to at least "four-nines" of availability, 20.9% improved their overall availability to "five-nines" or higher as a result of their storage refresh.

For the last survey result bullet point, the reader should note that many public clouds do not offer an SLA of "four-nines," and among those that do, it often requires an uplift charge. Organizations that require this level of availability tend to make on-premises placement decisions (either traditional IT or private cloud) for those workloads.

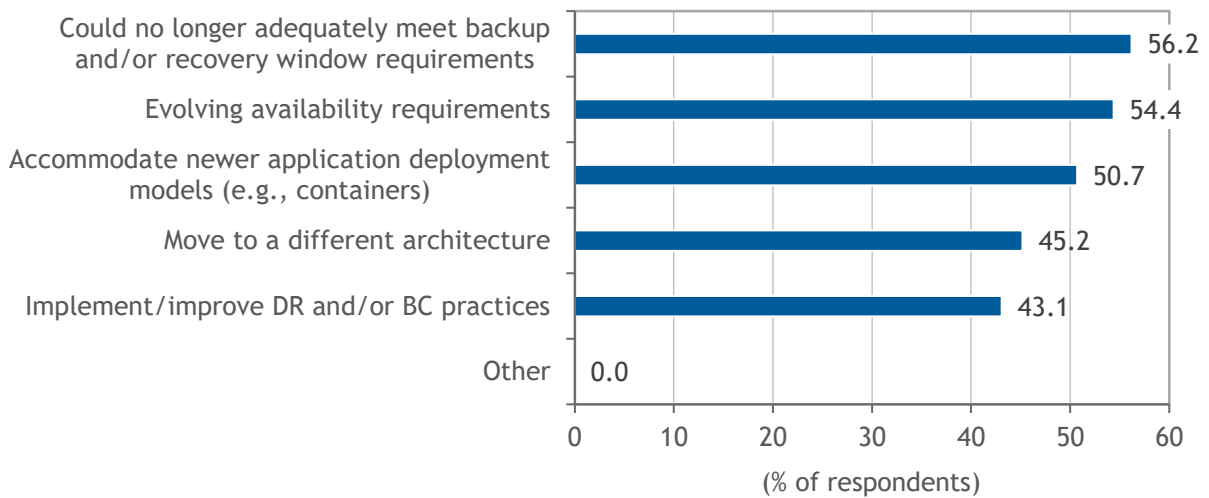
New, multicore CPUs can drive considerably more throughput than older CPUs, making it particularly important for organizations that are refreshing their servers to also do a storage refresh to avoid creating a new bottleneck. Newer storage technologies such as NVMe and solid-state storage are able to use such processing power much more efficiently than older spinning disk technologies, an important factor for latency-sensitive workloads in particular. When combining more powerful compute capabilities with lower-latency storage, server CPU utilization can be significantly increased, potentially driving cost savings not only in new server purchases but also in application licensing costs on those servers (because fewer servers are needed to hit a given performance requirement). Whether or not a system actually needs the performance of NVMe for workloads today, most organizations expect they will need to migrate to it as they continue to implement next-gen apps during their DX journey. The ability to nondisruptively support NVMe and related technologies (such as NVMe over Fabrics) going forward is becoming an increasingly important consideration as organizations refresh both server and storage infrastructure.

Data Protection Refreshes Driven by Modern Applications and Data Growth

For data protection refresh, organizations were focused on handling high data growth across heterogeneous environments while still continuing to meet recovery SLAs and manage costs. The single most important driver for data protection refreshes was that high data growth meant that organizations could no longer adequately meet backup and/or recovery windows with their existing data protection infrastructure and workflows (see Figure 4). Other drivers included the need to meet increasingly stringent recovery requirements, the need to accommodate modern application deployment models such as containers, and a desire to improve disaster recovery and/or business continuity practices. Newer storage architectures, along with access to newer technologies such as cloud and solid-state storage, mean that the data protection landscape has changed significantly for most organizations since their last storage technology refresh. There are significant opportunities not only to lower costs but also to improve data protection performance and efficiency.

FIGURE 4

Drivers of Data Protection Refreshes



n = 631

Source: IDC's *Modernized Infrastructure Survey*, 2020

For our survey respondents, data protection refresh drove noticeable improvements as well as a need for certain key features:

- 85.4% of respondents improved their recovery point objective (RPO) by at least 25%, while 84.8% improved their recovery time objective (RTO) by that same amount.
- 48.5% of respondents cut their backup and recovery job failures by at least half, and 45.9% improved their backup administrative efficiency by at least 50%.
- 67.2% of respondents preferred scale-out data protection infrastructure, and 86.1% considered container support very important or extremely important.

High data growth rates figure prominently in driving the need for data protection refresh for most organizations that undertake it. Modernized storage infrastructure often includes features that can make a significant difference in capacity utilization, such as inline compression and deduplication, and when used in hybrid cloud environments can extend these benefits to data stored in the public cloud. When data is stored in the public cloud, organizations have to implement the right data protection and recovery strategies there as well. While encryption for data at rest is common, the ability to also encrypt data in motion is not, and this capability can help organizations create more secure infrastructure that better protects the privacy of their customers. Protection and recovery may require separate strategies to ensure that data integrity is maintained as it is written and migrated and that recovery is available to meet the range required of RPOs/RTOs. Cyber recovery strategies that can protect against corruption and ransomware attacks are part of that plan, as are snapshots, replication, availability zones, and air-gapped data vaults.

Newer application deployment models such as containers open up the requirement to support faster, more efficient recovery operations when data protection infrastructure supports this level of granularity. In addition, the need to protect and recover the entire container-based environment (including the

application) has become paramount as container-based applications move to production for mission-critical workloads. And with high data growth expected to continue for at least the next five years, the ability to easily scale data protection infrastructure to nondisruptively accommodate business growth is driving interest in the use of scale-out architectures.


Converged and Hyperconverged Infrastructure Refreshes Driven by Improving Staff Productivity

With the increasing pace and dynamism of business today, IT agility is top of mind. Converged and hyperconverged infrastructure both enable easier procurement, faster deployment, unified management, and a single point of support contact but differ slightly in their architectural focus. Typically, converged infrastructure packages discrete server, storage, data protection, and networking components into a single SKU, all pre-certified as compatible, whereas hyperconverged infrastructure uses server-based, software-defined building blocks where all resources are pooled. With converged infrastructure, organizations can choose and independently scale configurations based on modular and/or rackmount servers and enterprise-class storage, while with hyperconverged infrastructure, organizations choose from among a variety of server-based building blocks interconnected in a cluster configuration, each providing compute and storage resources, and scale software-defined virtual storage across the nodes in the cluster. Both approaches allow organizations to provision new infrastructure much faster than in the past (see Figure 5).

One of the key value propositions for both converged and hyperconverged infrastructure is that they both simplify and modernize infrastructure. Like with public cloud, ordering what is needed is easier, provisioning infrastructure is faster, organizations do not need to worry about compatibility issues between resource tiers, and the entire "platform" is managed through a single pane of glass and supported by a single vendor. In this way, converged and hyperconverged infrastructures embrace many of the principles of cloud, a factor that not only simplifies on-premises deployments but also makes hybrid cloud management easier.

FIGURE 5

Top Drivers of Converged and Hyperconverged Infrastructure Refresh



n = 389

Source: IDC's *Modernized Infrastructure Survey*, 2020

Survey respondents noted significant benefits to converged infrastructure deployment (which 43.2% of respondents had implemented in their own on-premises datacenters):

- Converged infrastructure deployment reduced the time to provision infrastructure by at least 50% for 67.0% of respondents.
- Converged infrastructure deployment improved application performance by at least 50% for 71.4% of respondents.
- 77.0% of respondents reduced capital costs by at least 30%.

Survey respondents also noted significant benefits to hyperconverged infrastructure deployment (which 31.8% of respondents had implemented in their own on-premises datacenters):

- Hyperconverged infrastructure reduced the time to provision infrastructure by at least 50% for 64.3% of respondents.
- Hyperconverged infrastructure improved application performance by at least 50% for 74.1% of respondents.
- 76.0% of respondents improved productivity by at least 30%.

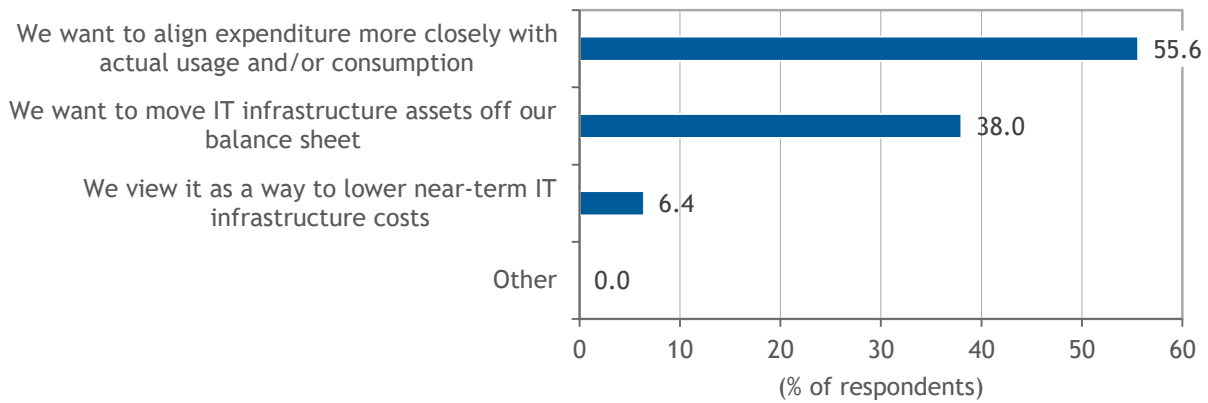
It is interesting to note that improving IT staff productivity was the number 1 driver of converged infrastructure deployment, highlighting that this approach speeds both ordering and installation and simplifies management of compute, storage, and networking resources through its unified control plane. While hyperconverged infrastructure includes that same value proposition, purchasers of that type of infrastructure were more focused on improving application performance. Datacenter consolidation, including not only streamlining infrastructure but also reducing energy and floor space consumption, was a top driver for both constituencies.

Moving to More Flexible Consumption Models

Subscription-based pricing and delivery models provide a number of advantages to organizations. Flexible consumption models deliver IT agility, enable rapid provisioning of new infrastructure, and let IT managers align expenditures more closely with actual usage. They also offer options to move IT infrastructure off the balance sheet when this makes economic sense. All of this makes the business more responsive to changing market conditions and helps make IT operations more efficient. The adoption of subscription-based pricing options for both on-premises and off-premises IT infrastructure (i.e., traditional IT, private cloud, and public cloud) is ramping quickly because it effectively separates the consumption and deployment model choices, providing more freedom to choose the best deployment option for each workload. Top drivers behind this rapid adoption include a need to align expenditure more closely with actual usage and/or consumption and a desire to move IT infrastructure assets off the balance sheet (see Figure 6). While all public cloud services use subscription-based pricing, 60.8% of survey respondents use it for at least some on-premises equipment (including both private cloud and traditional IT infrastructure).

FIGURE 6

Key Drivers of the Desire for Subscription-Based Pricing



n = 900

Source: IDC's *Modernized Infrastructure Survey*, 2020

Many organizations already appreciated the agility of subscription-based pricing for rapid initial provisioning or expansion of existing infrastructure. The "pay as you grow" and "pay per use" consumption models also enable organizations to more flexibly accommodate changing business climates. Although terms will vary by vendor, when the need for IT infrastructure shrinks, whether due to project completion, business downturns, or other economic impacts, these consumption models can provide the ability to quickly scale back both IT resources and costs. This is another side to the "agility" coin where subscription-based pricing models offer advantages over traditional outright purchase models.

Selecting the Optimal Environment for Each Workload

When public cloud services first became available, organizations were running all workloads in on-premises traditional infrastructure. At that time, workload placement strategies typically revolved around determining which workloads could move to public cloud and how (rehost, replace with a SaaS option, re-architect, etc.). This early usage of public cloud exposed organizations to different and better operating models in terms of acquisition, agility, provisioning, infrastructure expansion and contraction, and technology refresh, among others. Much has changed since then, including both deployment and consumption model options. Many of the original advantages of public cloud can now be easily replicated with on-premises IT infrastructure and deployment options that can better meet performance, availability, security, compliance, governance, and cost requirements for many workloads. Today, workload placement strategies need to take these various options into account and focus much more specifically on individual workload requirements to determine where the workload ultimately resides. Hybrid cloud platforms enable enterprises to select the right infrastructure for the right workload, thereby enabling them to make efficient workload placement decisions that help better meet requirements while controlling costs.

Hybrid cloud platforms enable enterprises to select the right infrastructure for each workload, thereby enabling them to make efficient workload placement decisions that help better meet requirements while controlling costs.

Factors Influencing Workload Placement

Among the 10 different enterprise workloads listed in the survey questions, respondents rated collaborative applications, IT infrastructure applications, and business applications as the most critical for their business today. Examples of these application types include email and messaging platforms, databases such as Oracle and SQL Server, virtual desktop infrastructure (VDI), and enterprise resource planning and customer relationship management (SAP, Oracle NetSuite, etc.).

Top initiatives that organizations plan to undertake during DX include IT infrastructure optimization, business applications, and structured data analytics. For example, in the IT infrastructure area, which includes services such as VDI, solid-state storage is critical to smooth out the effects of boot, log-in/log-out, and application storms over the course of the day so that end users see consistent performance. Efficient and performant snapshot and clone technology is critical to maximizing capacity utilization without impacting performance, and for growing businesses, the ability to nondisruptively expand VDI without impacting existing end users is required to meet service-level agreements around availability. Other IT infrastructure components include file and print, directory, and security services, all of which require the peak performance, high availability, and nondisruptive scalability that modernized infrastructure offers. IT infrastructure is considered a "top application" because it is an enabler of other aspects of DX and must be evolved as necessary to support the demanding requirements of DX.

As part of the DX journey, organizations should review their existing workloads to determine whether they need to keep or retire them. If organizations decide to keep them, then they need to determine whether these workloads should reside where they are or be moved to a different deployment model (which could be a new type of modernized on-premises infrastructure or a cloud-based environment). Workload placement decisions should be guided by each workload's usage patterns and requirements. It is IDC's considered opinion that these decisions need to be made assuming a hybrid cloud strategy that offers three deployment models: traditional on-premises, private cloud, and public cloud. This means crafting a hybrid cloud strategy up front to guide technology refresh, modernized infrastructure, and workload placement decisions. The goal of this strategy should be to achieve a consistent cloud experience across all deployment models that support simple and efficient workload placement and data mobility, all managed from a unified dashboard. Workloads can then be evaluated and placed in the optimal deployment model and, if things change over time, easily moved to a different model.

According to the survey, the key considerations that are taken into account when determining workload placement are security, performance, ease of management, availability, and cost (see Figure 7).

FIGURE 7

Factors Influencing Workload Placement



n = 900

Source: IDC's *Modernized Infrastructure Survey*, 2020

Security

Security is the single most important consideration in workload placement, with 76.4% of survey respondents calling it out as the top issue. With the proliferation of workloads across multiple deployment models, some of which have been outsourced to third-party entities, the boundaries of IT environments and who manages them are blurred. Traditional IT environments are protected through firewalls, endpoint security systems, and DMZs. Organizations also have well-established processes to set the right levels of access, enforce governance, and ensure compliance with industry-specific requirements, if any. As workloads and data become more distributed, organizations must evolve security, data protection, and recovery strategies to meet business requirements regardless of which deployment model is chosen. Security in the public cloud is a shared responsibility model that requires actions by both the public cloud provider and the customer. Given that organizations cannot necessarily apply the same tools and practices used for traditional IT infrastructure to cloud-based environments, having the right hybrid cloud strategy in place to protect their most valuable digital assets is very important. As mentioned previously, the purpose of a hybrid cloud is to provide a unified framework for IT services management across a variety of different deployment models, and security is just one of the management areas that benefit from the right hybrid cloud framework.

Performance

70.0% of survey respondents rated performance as the second most important criterion influencing workload placement decisions. Enterprise applications have traditionally depended on the underlying infrastructure for performance and availability. However, performance can mean different things for different workloads – for some, it's purely latency driven; for others, it may mean throughput or depend on data locality. Organizations need to consider what "performance" means for each workload as they decide on the proper deployment model. On-premises deployment options can meet certain performance requirements that are not well met in the public cloud; for example, when the network latency incurred by a hop to a public cloud datacenter may preclude its use for certain very latency-sensitive workloads.

Some workloads that are not designed to take advantage of scale-out cloud architectures will require underlying infrastructure that can meet demanding hardware requirements, while others would benefit from being able to seamlessly add thousands of cores on demand in a public cloud. There are also many accelerated compute and solid-state storage resources available as cloud-native services that can make it quick and easy to access these kinds of resources when organizations either do not already have them or want to deploy them in on-premises infrastructure. In other cases, IT managers may specifically want to have those types of resources deployed in on-premises infrastructure for data gravity and/or other reasons. Data gravity issues must be taken into account as organizations will want to craft workload placement strategies to minimize the movement of large data sets between different locations.

Ease of Management

62.4% of organizations deemed management considerations the third most important driver in workload placement decisions, behind security and performance. Each organization has its own set of management tools and workflows deployed to achieve its business objectives. When deciding deployment models, organizations will need to evaluate which approach best fits their requirements for daily operational tasks such as performance tuning, availability management, provisioning additional resources, executing backup and/or disaster recovery strategies, and supporting new application deployment paradigms such as containers. Deployment models also can offer differentiating advantages for other tasks that are done less often, such as introducing new resources (e.g., storage-class memory or accelerated compute), spinning up additional resources to meet seasonal demands, or managing technology refresh. Part of the reason why organizations seek more IT agility is to better enable their ability to respond to these types of demands.

For example, public cloud provides a subscription-based pricing model, offers better IT agility than traditional IT infrastructure, may allow administrators to rapidly access new technologies through native cloud services, and enables organizations to offload management responsibilities to the cloud provider. Private cloud can offer many of those same benefits with on-premises infrastructure that can meet certain performance and availability requirements that public cloud providers won't contractually guarantee, but it requires customers to manage that infrastructure themselves. Other workloads may be best kept on traditional IT infrastructure for any number of workload requirement or business objective reasons. To optimize workload placement, organizations need to have all three deployment models available to them.

Different public clouds (e.g., Amazon and Google) offer cloud-native services that may differentiate them from other cloud providers, and many organizations use multiple public clouds. Different public clouds have different operating models and management interfaces, and those differ from what organizations use for on-premises infrastructure. Managing each environment using the cloud

provider's native approaches necessarily means that IT managers need to learn and use different management dashboards and workflows, introducing significant additional complexity. The objective of a well-crafted hybrid cloud strategy is to move to more of a common operating model across not only different public clouds but also all deployment models that may be under management.

Availability

A very high availability requirement makes a workload strategic for more survey respondents than any other criteria – 33.3% of respondents chose it as the most important characteristic. Other characteristics that make a workload strategic are that it be responsible for driving competitive differentiation and that it exhibit a very high performance requirement. Availability was also a key criterion determining workload placement, with 58.7% of organizations considering it during the placement decision.

Most public cloud providers will not write a contract that guarantees more than "four-nines" availability, even if their actual performance over a given time period happens to exceed that. Many public cloud-based storage services carry an even lower availability guarantee than that. With on-premises infrastructure, IT managers have more control to build specific configurations that they know can provide higher availability than what public cloud providers can contractually guarantee. For workloads that need to be managed to "four-nines" availability, public cloud can be appropriate (although IT managers will need to evaluate the additional costs that may be associated with a public cloud provider guaranteeing that level of availability). If workloads require higher availability than "four-nines," then public cloud may not be the best option. While public cloud data protection and disaster recovery can be a valuable component of modernized infrastructure, some workloads may require on-premises deployment to ensure compliance with strict RPOs/RTOs or country-specific regulations for data locality.

Cost

With its massive scalability, flexible agility, and low price-per-gigabyte costs, public cloud can be a cost-effective platform for workloads that need to retain data over long periods of time, require unpredictable and relatively rapid access to new IT resources that may be seasonal or project driven in nature, and have usage patterns that do not require a lot of ongoing data access or data movement (data gravity considerations). On-premises deployment options can be more cost effective when workload growth and resource utilization are more predictable; when performance, availability, and/or security requirements are not well met with public cloud capabilities; when regulations require that administrators physically know where data resides at all times within certain geographical boundaries; and/or when data access and usage patterns are more frequent. Most enterprises have both types of workloads in their application portfolio, and a hybrid cloud strategy ensures that optimal workload placement decisions can be made based on individual application requirements while still providing a consistent operational model across all environments.

Those considering public cloud for the first time should realize that the use of the public cloud often requires administrators to familiarize themselves with new sets of workflows. There will be a learning curve as organizations handle day-to-day cloud-based operations. Often, there is also considerable cost required for replatforming applications to optimize for public cloud. For organizations looking to enable self-service with public cloud-based offerings, there may be additional training and even governance concerns that don't exist with on-premises infrastructure. It is important to note that our survey results indicate that when workloads are repatriated from the public cloud back to on-premises locations, 82.1% of respondents identified cost as a driving factor.

While the agility of the public cloud has been rightly billed as an attractive financial advantage, the availability of different consumption model options for on-premises IT infrastructure can provide many of the same benefits in this area. Subscription-based pay-as-you-grow and pay-per-use options for on-premises infrastructure allow IT to much more closely tie expenditures to actual resource consumption and can also allow organizations to move IT assets off the balance sheet. Many vendors offer both customer-managed as-a-service and fully managed services options using their IT infrastructure products and provide organizations the ability to physically locate equipment in either on-premises or off-premises locations. The subscription-based pricing model is no longer a unique advantage of the public cloud – organizations can now choose from different deployment and consumption model options for products and services that have nothing to do with the public cloud.

DELL TECHNOLOGIES: A COMPREHENSIVE PORTFOLIO FOR MODERNIZED INFRASTRUCTURE

Dell Technologies is a large provider of both consumer and enterprise IT products and services. The vendor's extensive enterprise technology portfolio includes products from Dell EMC, Secureworks, Virtustream, and VMware, and its offerings cover cloud software, servers, storage, data protection, networking, converged and hyperconverged infrastructure, software-defined datacenter, and cloud platforms as well as enterprise infrastructure software in the virtualization, storage, security, data protection, and hybrid cloud control plane markets. In the enterprise storage market, the vendor's portfolio comprises primary storage platforms, unstructured data storage, and software-defined storage solutions and features many of the key technologies – solid-state storage, AI/ML-driven management, cloud technologies, software-defined infrastructure, and scale-out architectures – that organizations undertaking DX are most interested in. Dell EMC servers and Dell EMC storage and data protection brands hold industry-leading market share positions by revenue in their respective industries and can be used across virtual, bare metal, and cloud-based environments. Those brands include PowerEdge servers, PowerMax/PowerStore/PowerVault/PowerScale storage, and PowerProtect data protection.

Dell Technologies' products feature Intel processor technology and leverage unique capabilities such as Intel QuickAssist Technology (QAT). Dell PowerEdge servers as well as Dell EMC PowerMax, PowerStore, PowerVault, PowerScale, VxRail, and PowerProtect systems are built with second-generation Intel Xeon Scalable processors. Intel QAT provides a software-enabled foundation for security, authentication, and compression to increase the performance and efficiency of platforms that use this technology. The Dell EMC PowerStore unified storage array is the vendor's first storage platform to leverage this technology to drive hardware-assisted compression, and organizations can expect that they will see Intel QAT leveraged in other Dell EMC server and storage systems over time.

Dell Technologies has focused on creating a comprehensive hybrid cloud strategy that centers around VMware virtualization technology. VMware Cloud Foundation (VCF) provides the unified control plane for a hybrid cloud strategy that lets organizations modernize their datacenters, migrate to cloud-based technology, scale on demand, and simplify hybrid cloud operations that span on- and off-premises deployment models and multiple public cloud platforms. Dell Technologies is using VMware Cloud Foundation to deliver consistent operations across its entire portfolio, supporting and integrating with multiple server, storage, and data protection solutions. Dell Technologies is also developing API technology to integrate and enhance the overall VMware experience (e.g., VxRail API integration into the VCF SDDC stack, Virtual Volumes, VMware APIs for Array Integration, VMware APIs for Data Protection, vRealize APIs, VMware Cloud on AWS API, Site Recovery Manager, and vSphere APIs) and improve and automate operations as well as offer unique functionality that provides better performance,

availability, monitoring, data protection, and recovery. Embedded Intel VT in Dell servers enables live data migration using VMware vMotion. Dell Technologies Cloud lets organizations use Dell technology across all major public clouds, completing the hybrid cloud strategy by enabling consistent management (through VMware technology) and a single-vendor experience for workload portfolios that span traditional on-premises and private and public cloud deployment models. Dell EMC VxRail is the first hyperconverged system fully integrated with VMware Cloud Foundation SDDC Manager. The Dell Technologies Cloud Platform, VCF on VxRail, delivers a simple and direct path to the hybrid cloud with one complete automated platform and consistent operations across VMware clouds.

With its broad portfolio of consumption-based and as-a-service solutions, Dell Technologies on Demand (DTOD) provides maximum flexibility to consume products and services in the manner that best fits an organization's business objectives. This allows organizations to choose deployment and consumption models separately, and this flexibility helps drive better business outcomes. In primary research conducted by IDC in early 2020, the business value of DTOD was strongly validated (see *Business Value of Dell Technologies On Demand: A Study of Usage-Based Consumption Models for Storage*, IDC #US46191920, April 2020). Among other positive impacts, DTOD enables Dell Technologies' customers to enjoy on average 25% lower acquisition costs, achieve on average 92% faster provisioning for new capacity, experience on average 54% fewer incidents of unplanned downtime, and lower their storage cost of operations by an average of 23% per year.

CHALLENGES/OPPORTUNITIES

As organizations consider the best way to optimize their infrastructure, they will likely be starting with server, storage, data protection, networking, infrastructure software, and other offerings from a variety of different vendors. This heterogeneous, on-premises IT infrastructure reality adds to management complexity. Migrating workloads to public cloud can help alleviate some of this complexity, but most organizations will operate hybrid cloud environments with workloads using all three deployment models: traditional on-premises, private cloud, and public cloud. This highlights the critical importance of crafting the right hybrid cloud strategy. Managing complexity will continue to be a challenge going forward for IT managers who lack a viable hybrid cloud strategy.

Vendors such as Dell Technologies that can offer well-integrated enterprise IT infrastructure solutions, available from a single source supplier, that span all three deployment models can alleviate much of this complexity. Dell Technologies is a market share leader in servers, primary and data protection storage infrastructure, and software-defined datacenter platforms, so it has a proven track record that spans all IT infrastructure areas. Over time, Dell Technologies and VMware have made significant investments in technology that mitigates multicloud complexity, and today, they offer a comprehensive platform for hybrid cloud that simplifies the management of modernized infrastructure. DX and modernized infrastructure are both opportunities for simplification through better hybrid cloud strategies.

CONCLUSION

When considering an optimal workload placement strategy across a portfolio of enterprise applications, today's organizations have more deployment model options than ever before. While early public cloud experiences validated many of the advantages of cloud-based deployment models, they have also pointed out some of the disadvantages of public cloud. Today's organizations need more than just traditional on-premises and public cloud-based options – they also need a private cloud-based option – and this means they need a comprehensive hybrid cloud strategy.

Historical outcomes have indicated that the right way to craft optimal workload placement strategies is to start with a hybrid cloud strategy. This approach gives organizations three options (traditional on-premises, private cloud, and public cloud) instead of only two options (traditional on-premises and public cloud) and gets businesses closer to optimal workload placement faster.

At the very core of the definition of hybrid cloud is the idea of a unified control plane, or management framework, that provides a cloud experience across an organization's entire IT infrastructure – traditional on-premises, private cloud, and public cloud. This framework needs to support simple and efficient workload and data mobility as well as comprehensive, cross-stack monitoring to ensure organizations have the visibility to meet their performance, availability, recoverability, and cost goals. Public cloud can be a good fit for certain workloads, but it may not be the best option for other workloads. On-premises private cloud infrastructure that replicates the agility and flexibility of the public cloud, in terms of both deployment and consumption models (i.e., subscription-based pricing), provides a way to meet workload-specific performance, availability, security, compliance, governance, and cost requirements for which the public cloud is not a good fit.

The ability to place workloads optimally on modernized infrastructure and migrate those workloads to different deployment models when things change is critical to achieving efficiency in a digitally transformed business.

The ability to place workloads optimally on modernized infrastructure and migrate those workloads to different deployment models when things change is critical to achieving efficiency in a digitally transformed business.

In building these technologies and capabilities into the IT infrastructure, organizations will potentially have to deal with many different products and vendors with many different management interfaces. The complexity of this fragmented, multivendor ecosystem will be a real challenge for many IT organizations. Dell Technologies offers a broad enterprise infrastructure portfolio that includes servers, storage, data protection, networking, and converged and hyperconverged infrastructure, as well as cloud offerings, all manageable through a single unified control plane that is jointly engineered with VMware. The vendor's products support needed new technologies, and with its range of offerings, Dell Technologies can deliver a consistent operational experience across traditional on-premises and private and public cloud deployment models that reduce complexity. Organizations interested in crafting a well-integrated hybrid cloud environment as part of their modernized infrastructure efforts will want to explore Dell Technologies' offerings.

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

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2020 IDC. Reproduction without written permission is completely forbidden.

