451 Research® | PATHFINDER REPORT

Hyperconverged Infrastructure Requirements to Accelerate Data-Driven Insights

COMMISSIONED BY

CISCO

intel

MARCH 2020

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 THE AUTHOR

HENRY BALTAZAR

RESEARCH VICE PRESIDENT

Henry Baltazar is a Research Vice President for the Storage Channel at 451 Research. Henry returned to 451 Research after spending nearly three years at Forrester Research as a senior analyst serving Infrastructure & Operations Professionals and advising Forrester clients on datacenter infrastructure technologies. Henry has evaluated and tested storage hardware and software offerings for more than 15 years as an industry analyst and as a journalist.

Executive Summary

As organizations look to become more data-driven, their infrastructure resources must evolve to meet the growing burden. The volume of data generated continues to grow at a rapid pace, and what makes matters worse is that this data is being created and processed in many geographically dispersed areas. Time to value is also growing in importance, yet inefficient and outdated infrastructure prevents companies from getting the insights they desperately need in a timely fashion to make strategic decisions – which is the very essence of what it means to be truly data-driven.

Over the last few years, hyperconverged infrastructure (HCI) has emerged as an alternative to traditional infrastructure deployments and has provided a number of benefits including faster resource provisioning and simplified management. HCI deployments allow generalists to take up more of the day-to-day management burden, which previously required expensive and highly trained specialists, such as storage and networking architects.

Data-driven insights are crucial to a company's ability to utilize customer, business and IoT data, so HCI must continue to evolve to deliver accelerated performance to feed critical data-driven decision-making engines. Because data growth is running amok in many organizations, the ability for the platform to scale and accelerate both storage and compute performance is a key capability that organizations consider when selecting an HCI platform.

Organizations need to take advantage of new memory and storage technologies and architectures such as storage class memory (SCM), which can **provide a significant performance increase** as well as resiliency, thanks to its non-volatile nature. Artificial Intelligence (AI) and machine learning (ML) are being weaved into HCI and other infrastructure resources. For example, AI is already starting to provide organizations with proactive guidance to help them optimize utilization and avoid the costly impact of unplanned outages and performance degradation. With new applications, workloads and cloud-native architectures on the rise, the pressure for IT infrastructures to change has never been greater.

In this paper we discuss:

- The current infrastructure challenges limiting organizations from becoming data-driven.
- The evolving role of HCI and its potential to improve infrastructure in multiple dimensions.
- The current state of adoption for SCM and the benefits this medium can provide.
- The operational improvements that can be driven by AI/ML to enhance management tools.

Key Infrastructure Problems

Both enterprises and midsized organizations are facing a range of challenges as workload demands continue to evolve and intensify. The first challenge, which is top of mind for most organizations, is rapid data growth, and that problem only seems to be getting worse every year. In our Voice of the Enterprise: Storage, Budgets and Outlook 2019 survey, respondents said data grew 21% over the previous 12 months, and they expect data to increase 25% in the next year. What makes this issue even more impactful is that these same organizations only expect to see an 11% increase in their storage budgets over the next year. Although organizations can expect to see a cost reduction on a dollar-per-GB basis due to the ongoing improvements made on commodity resources such as flash and hard drives, hardware costs are only part of the equation for many companies since they need to address the operational costs of deploying and maintaining these resources.

As a result, it is not surprising that the lack of self-service infrastructure and application provisioning ranked second in the list of challenges facing on-premises infrastructure, according to respondents of the Voice of The Enterprise Storage: Workloads and Key Projects 2019 survey (Figure 1). Organizations need to ensure that their key stakeholders and applications are able to get the resources they require in a timely fashion. HCI tools allow IT generalists to more easily manage resources, which can help organizations lower their opex since IT specialists are costly to find and retain. This is especially true in data-driven organizations, which must extract insight from data repositories.

Figure 1: Cost, the perennial pain point for on-premises infrastructure

Source: 451 Research's Voice of the Enterprise: Storage, Workloads and Key Projects 2019
Q: Which of the following are the biggest pain points regarding your current on-premises infrastructure?
Base: All Respondents (n=159)

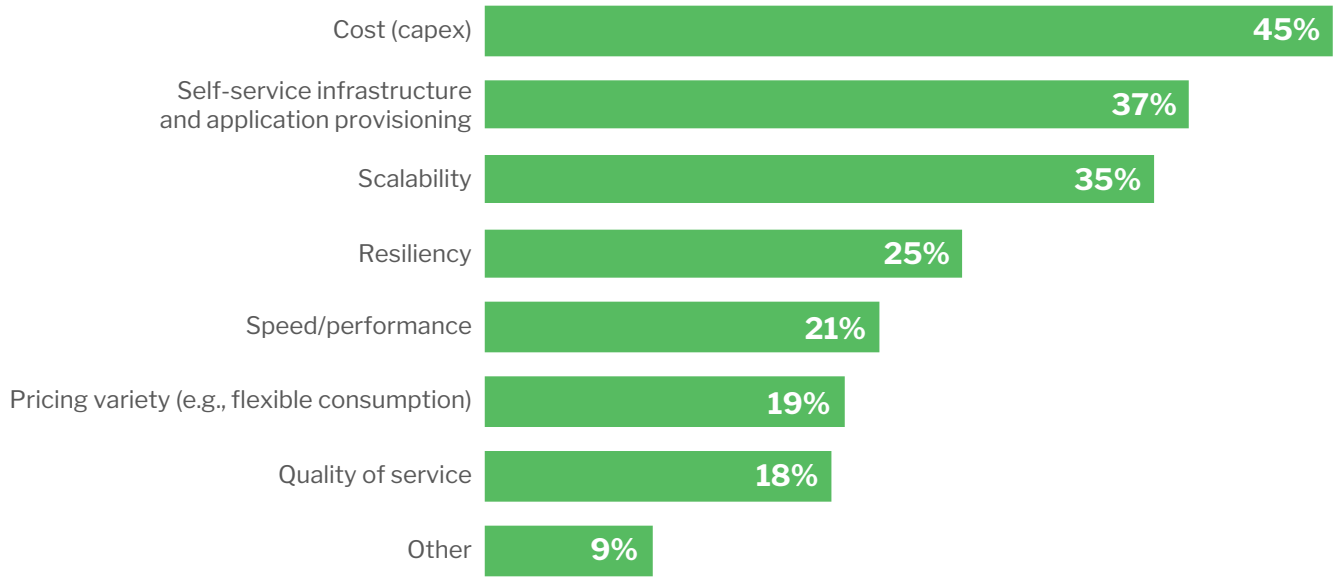

Performance continues to be a major problem for organizations, with 21% choosing it as a top pain point in the study. **As the competitive landscape for organizations becomes more challenging, it's clear that companies should prepare for a world where the ability to make quick data-driven decisions will be a key factor that will separate winners from losers.**

Complicating matters for organizations is that the data they are creating often comes from a variety of geographic areas both inside the organization and outside if it comes from an edge location, third party or a cloud service. One of the most important infrastructure decisions that a data-driven organization has to make is where its data should be processed.

In our Voice of the Enterprise: Storage, Workloads and Key Projects 2019 survey, cost and security were the top factors that organizations are weighing when deciding the execution venue; these were followed closely by application performance (Figure 2). As data-driven insights become increasingly important to deliver enhanced digital services to customers, IT organizations are faced with the need to modernize both compute and storage infrastructure in order to scale both performance and capacity to utilize the growing datasets on the scale of hundreds of terabytes or petabytes. Additionally, moving these large data payloads will make WAN bandwidth a more important issue since that continues to be the preferred method for transport, although the use of physical data shuttles could become a factor. When you factor in egress charges that are common for pulling data out of a public cloud environment, it is clear that organizations need to take data-migration tasks seriously or risk major financial consequences. Organizations also need to factor in the time spent moving the content.

Figure 2: Cost, security and performance impact execution venue decisions

Source: 451 Research's Voice of the Enterprise: Storage, Workloads and Key Projects 2019
Q: In general, which factor is most influential when determining the best execution venue for a workload?
Base: All Respondents (n=503)

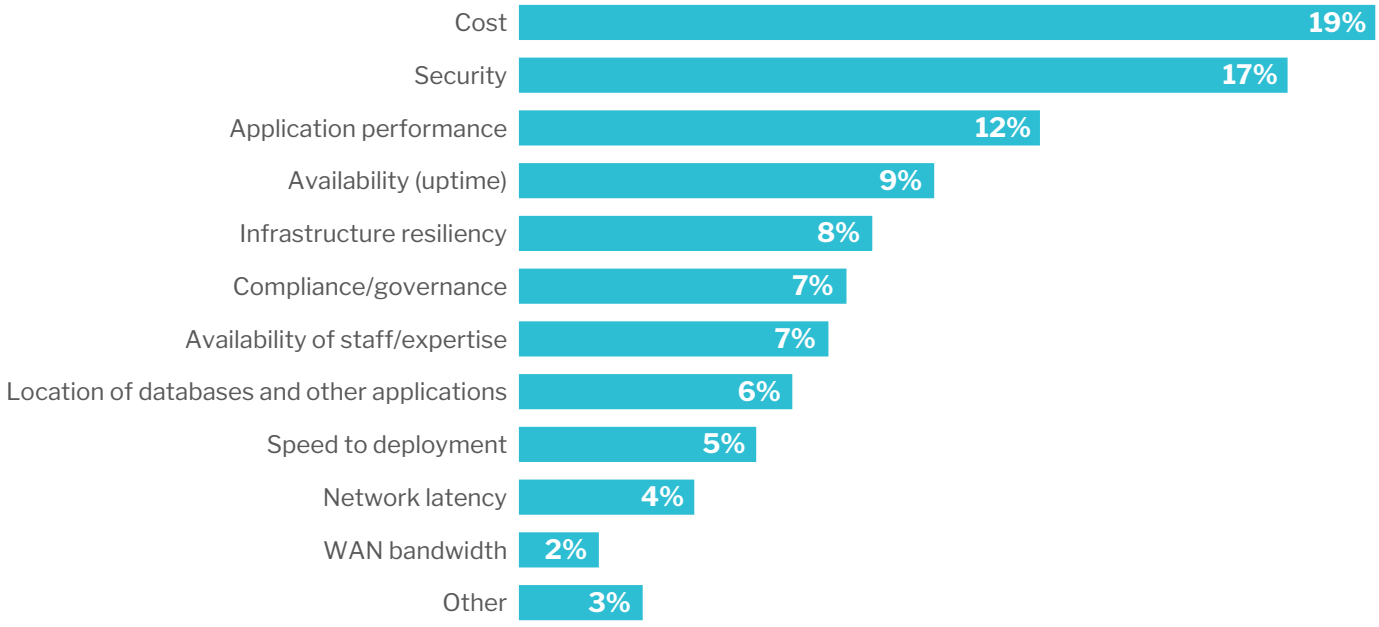

Taking these factors into consideration, it becomes apparent that if large datasets are taking too long to move over the network, there is a need for simplified higher-performance on-premises solutions for edge and remote environments. These offerings should become a potential option for data-driven organizations that want to extract value from their data as soon as possible. To meet these growing needs, we expect to see systems in these deployments that leverage high-performance compute and storage technologies to rapidly process data and quickly send back insights to stakeholders accessing databases in central sites or other environments.

Last but not least, data-driven organizations with highly distributed environments must also face the challenge of managing these resources. Availability of staff/expertise was an influential factor for selecting execution venues for 7% of the respondents, and we believe this will increase as more distributed analytics and IoT use cases become mainstream.

HCI is Expected to Fulfill Key Requirements

HCI adoption continues to ramp up, and customers are using it in a number of ways. In our Voice of the Enterprise: Servers and Converged Infrastructure, Budgets & Outlook survey, 79% of respondents said HCI is a way to simplify infrastructure management/maintenance, while 60% of the respondents viewed HCI as an accelerator for their provisioning and optimization tasks (Figure 3). HCI vendors are investing a large portion of their R&D resources on improving the management capabilities of their platforms, with an emphasis on using AI/ML technologies to proactively locate key issues their systems are reporting from the field and report them to IT managers so that they can avoid being impacted by these issues.

Organizations are also looking to leverage the data replication capabilities of HCI to help them with their data and workload locality challenges, as evidenced by the 51% of organizations that are using HCI to support their hybrid cloud strategies and the 25% that want to use them for remote office/branch office support. The future is also top of mind for HCI customers, with 36% looking to use these products to support emerging technologies such as containers, and 17% enabling greenfield application deployments.

Figure 3: HCI is simplifying management while accelerating provisioning

Source: 451 Research's Voice of the Enterprise: Servers and Converged Infrastructure, Budgets & Outlook 2018
Q: What role does hyperconverged infrastructure serve in your organization's hardware infrastructure deployment strategy? Please select all that apply.
Base: All Respondents (n=175)

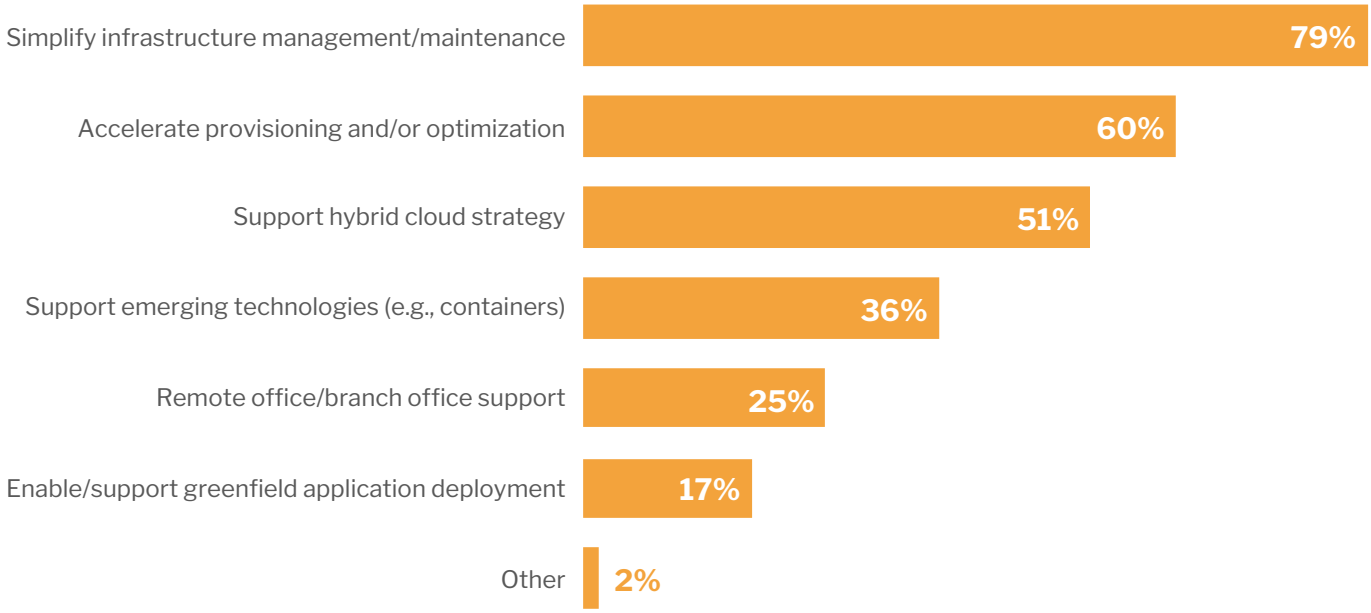

To be successful in progressive IT departments, HCI platforms must deliver storage performance in line with traditional all-flash array technology, and provide scalability in multiple dimensions. Beyond storage capacity scalability, HCI must also add compute in a granular fashion and seamlessly integrate dedicated AL/ML to meet the changing needs of workloads. In the past, HCI was criticized for its inefficiency because both storage and compute came in each node, regardless of whether an organization needed to add only one resource to meet client demands. The more advanced HCI vendors now offer discrete compute and storage nodes that allow customers to add appropriate resources when they are needed.

Figure 4: Scalability and efficiency are HCI selection factors

Source: 451 Research's Voice of the Enterprise: Servers and Converged Infrastructure, Budgets & Outlook 2018
Q: What technology features are most important in selecting a hyperconverged infrastructure platform? Please select up to 3.
Base: All Respondents (n=307)

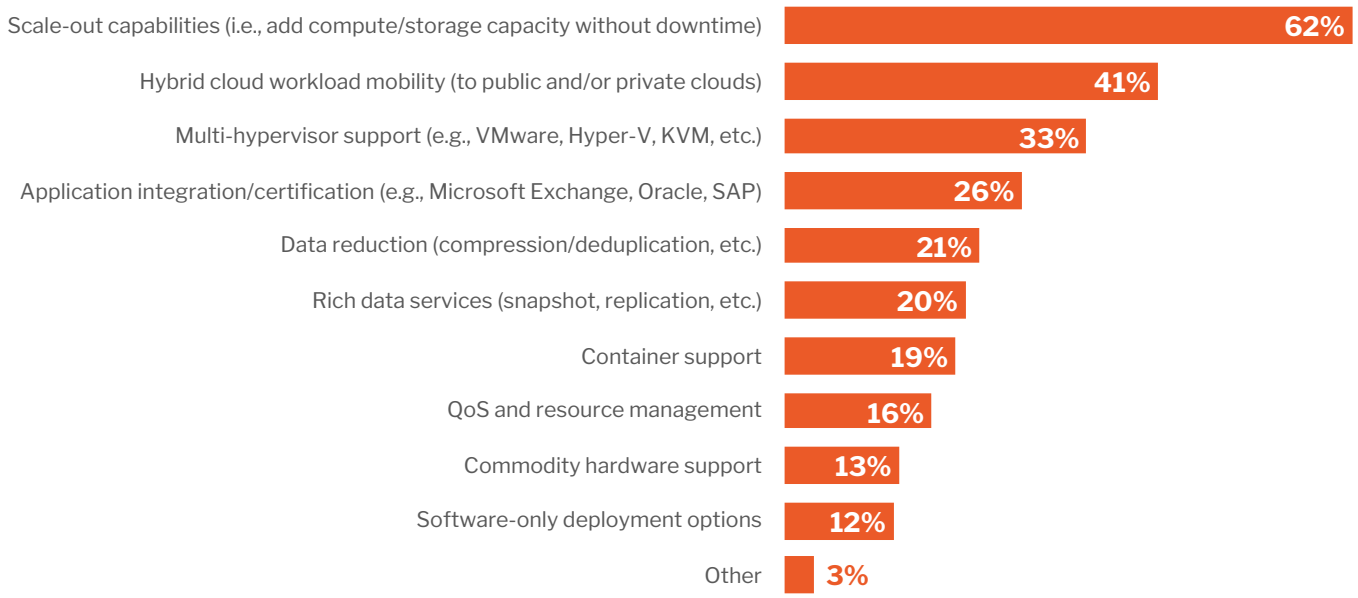

Another factor in the performance scalability challenge is the ability of HCI platforms to continue to provide high performance and low latency as the amount of capacity being managed scales up. The use of HCI to support applications like databases continues to grow ~10% YoY to 26% (Figure 4); thus, performance scalability becomes crucial for handling the increasingly large datasets and volume to deliver data-driven insights. SAS and SATA drive interfaces create IO bottlenecks since these drive interfaces were not designed to utilize the full capabilities of flash media. Trying to meet IOPS requirements leveraging this technology quickly drives up cost due to the increased number of HCI nodes required to support the disks needed.

Another technology that is gaining traction is NVMe flash drives, which can deliver lower-latency performance compared with SATA and SAS SSD drives that populate most all-flash arrays today. While NVMe drives can deliver increased storage performance, the lack of standardized protocols for reliability, availability and serviceability is hindering mainstream adoption. This is a key consideration for organizations when exploring shared-architecture technologies like HCI. As the prices for NVMe media and SCM continue to drop, we expect these technologies will become mainstream for both enterprises and midsized organizations.

This phenomenon is creating a market opportunity for SCM, which can provide higher performance than flash while costing less than DRAM. SCM is still in the early adoption phase. In 451's recent Voice of the Enterprise: Storage, Workloads & Key Projects 2019 survey, 22% of respondents said they have SCM deployed in servers, storage arrays and HCI platforms today. The adoption of SCM was higher for large companies (more than 1,000 employees), with 30% in use compared with 18% for organizations with fewer than 1,000 employees.

Conclusion

Recommendations:

- Reduce time to value for data by accelerating storage and compute performance with high-performance HCI offerings.
- Reduce operational costs with HCI (management tools and benefit of CPU/storage consolidation).
- Increase performance and reduce overall storage costs – leverage new higher-performance storage (SCM, NVMe) in both caching and capacity tiers to reduce the number of drives needed to meet performance demands.

Learn more how Cisco HyperFlex All NVMe HCI solutions accelerate storage performance to deliver data-driven insights at: *http://cs.co/hyperflexallnvme*

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. Founded in 2000, 451 Research is a part of S&P Global Market Intelligence.

NEW YORK
55 Water Street
New York, NY 10041
+1 212 505 3030

SAN FRANCISCO
One California Street,
31st Floor
San Francisco, CA 94111
+1 212 505 3030

LONDON
20 Canada Square
Canary Wharf
London E14 5LH, UK
+44 (0) 203 929 5700

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

451 Research®