

The Future of Data Infrastructure

Composable infrastructure is a new architectural approach that uses NVMe™-over-Fabric to dramatically improve compute and storage utilization, performance, and agility in the data center.

Western Digital OpenFlex™ architecture and products allow storage to be disaggregated from compute, enabling applications to share a common pool of storage capacity. Data can easily be shared between applications or needed capacity can be allocated to an application regardless of location.

With the exponential growth in data, along with the increasing diversity of workflows and demands on IT infrastructure, organizations need to increase speed, agility, and time-to-value for their customers. Emerging as a solution for this, composable disaggregated infrastructure is a new architectural approach, using NVMe™-over-Fabrics, that will vastly improve compute and storage utilization, performance, and agility in the data center.

Western Digital OpenFlex delivers scale-out performance and open composability.

Enabling Fast Data to Live Outside the Server

NVMe-over-Fabrics, or NVMe-oF™, is a networked storage protocol that allows storage to be disaggregated from compute to make that storage widely available to multiple applications and servers. By enabling applications to share a common pool of storage capacity, data can be easily shared between applications or needed capacity can be allocated to an application to respond to application needs. Exploiting NVMe device-level performance, NVMe-oF promises to deliver the lowest end-to-end latency from application to shared storage. NVMe-oF enables composable infrastructures to deliver the data locality benefits of NVMe direct attached storage (low latency, high performance) while providing the agility and flexibility of sharing storage and compute. OpenFlex's scale-out performance is particularly useful for large analytics tasks, large-scale high-performance computing, and other demanding workloads.

REPORT EXCERPT

The following is an excerpt from an independently published 451 Research report, "Composable Infrastructure: A Market Update," released in March 2021.

To purchase the full report or to learn about additional 451 Research services, please visit <https://451research.com/products> or email 451sales451@spglobal.com.

451 Research

S&P Global

Market Intelligence

**THOUGHT
LEADERSHIP**
MAR 2021

Composable Infrastructure: A Market Update

Christian Perry, Senior Research Analyst, Infrastructure

Mike Fratto, Senior Research Analyst, Applied Infrastructure
& DevOps

John Abbott, Principal Research Analyst

Composable infrastructure represents a nascent market with significant opportunity based on current traction with similar software-defined infrastructure technologies. This report examines the current market landscape of composable infrastructure, including ongoing innovation, challenges and organizational dynamics that will impact its future.

About the Authors



Christian Perry

Senior Research Analyst, Infrastructure

Christian Perry is a Senior Research Analyst covering IT infrastructure at 451 Research, a part of S&P Global Market Intelligence. In this role, Christian covers emerging and disruptive infrastructure technologies, including hyperconverged infrastructure and composable infrastructure. He also manages the Voice of the Enterprise products – built on 451 Research's proprietary global network of senior IT decision-makers – covering hyperconverged infrastructure. With more than 20 years of experience tracking and analyzing the IT infrastructure market, Christian brings broad knowledge to research around software-defined and cloud-native technologies that increasingly shape today's market.



Mike Fratto

Senior Research Analyst, Applied
Infrastructure & DevOps



John Abbott

Principal Research Analyst

Executive Summary

Introduction

Once defined largely by a range of theoretical laws, including Moore's Law, the IT infrastructure market is morphing quickly into a dynamic ecosystem that reflects the ongoing deluge of data – and the desire to use it for business advantage. Moore's Law, which posits that transistors on a microchip will double roughly every two years, has been in decline for years due to physical chip limitations and the skyrocketing cost of fabs. The related concept of Dennard scaling, which theorized that power density stays consistent as transistors shrink, is now defunct. Further, today's performance-hungry applications are running into limitations posed by Amdahl's Law, which says that parallelized processing is only as effective as the weakest link in the processing chain.

To counter the mounting limitations of silicon-based design, IT teams are increasingly integrating a multitude of modern infrastructure technologies into their datacenters, including accelerators (GPU, ASIC, FPGA, etc.), persistent memory and others. While these work to supplement traditional CPU-based architectures, they also introduce complexity created by the influx of systems to the overall datacenter environment. Although the standard, hardware-centric, three-tier infrastructure approach remains common in today's on-premises IT environments, elements of software-defined resource provisioning and management are materializing in most datacenters, as simply throwing more infrastructure into the mix to scale IT is no longer tenable for organizations that seek efficient, seamless management. The most extreme example of this is composable infrastructure, which disaggregates all hardware and pools compute, storage and networking resources on the fly.

In part because composable infrastructure is a significant departure from traditional infrastructure management, its adoption remains nascent. However, as shown by the meteoric rise of hyperconverged infrastructure (HCI), organizations are willing to invest in modern infrastructure that changes the way they deploy and manage resources. In fact, our research suggests that recent organizational trends, such as the rise of IT generalists, are driven primarily by changes desired by IT teams, rather than as a response to changes in technology. In the meantime, composable infrastructure waits in the wings as organizations continue to modernize their environments to match the innovative pace of public cloud.

Table of Contents

1. Why Composable Infrastructure?	1
<i>Figure 1: Inside the Possibilities of Composable Infrastructure</i>	2
<i>Figure 2: Likelihood To Automate Resource Provisioning</i>	3
<i>Figure 3: Composable Infrastructure Implementation Status</i>	4
<i>Figure 4: Software-Defined Infrastructure Implementation Status</i>	5
Hybrid IT Supports the Case for Composable	5
<i>Figure 5: Overall IT Approach and Strategy</i>	6
Infrastructure Moves Toward Strategic Spending	6
IT Deficiencies Open the Door	7
2. Networking Will Drive Initial Innovation	10
It's All in the Fabric	11
3. Challenges and Market Impact	13
<i>Figure 6: Drawbacks of Automation</i>	14
Composable and the Evolution of IT Team Structure	15
4. Vendor Profiles	17
A3Cube	17
Dell EMC	17
Fungible	18
Hewlett-Packard Enterprise	18
GigaIO	18
Liqid	19
Pensando	19
TidalScale	19
VMware	20
Western Digital	20

5. Conclusions	21
6. Further Reading	22
7. Index of Companies	23

1. Why Composable Infrastructure?

As a ceaseless influx of data continues within datacenters across the globe, enterprises are continually seeking more effective methods of harnessing that data for business advantage. Traditional, three-tier infrastructure composed of servers, storage arrays and networking has steadily evolved to account for the data deluge, offering higher levels of automated provisioning, virtualization and single-pane-of-glass management that eases the overall infrastructure experience. However, the architecture of stand-alone infrastructure limits its flexibility when organizations seek to build private and hybrid clouds that excel with disaggregated resources that can be efficiently pooled and provisioned.

Composable infrastructure represents perhaps the ultimate step in this direction by abstracting compute, storage and networking resources from the underlying infrastructure and presenting them as cloud services. We define composable infrastructure as compute, storage and network modular infrastructure configured into pools of resources that are dynamically allocated for workload requirements. Virtualized platforms, including HCI, use abstraction concepts to optimize utilization of compute and storage, delivering resources from a static set of infrastructure devices. Composable platforms, on the other hand, use modular hardware that can be quickly reconfigured to adapt to changing resource demands. Composed (or recomposed) pools of compute, storage and networking can, in theory, consistently deliver resources for any workload requirement, and they can be delivered in any form of virtualized, containerized or bare metal. This is a major departure from HCI, which can deliver only virtualized or containerized resources, but not bare metal.

The process of modernizing datacenter infrastructure is one fraught with questions and potential complexities. Do we upgrade our existing servers, storage and networking with next-generation versions of the same hardware? Do we methodically introduce wholly different infrastructure, such as HCI, into the environment until it becomes the standard deployment location for all or most on-premises workloads? Do we completely rip and replace our infrastructure with HCI or composable systems? Whichever route is ultimately taken, it appears that most enterprise organizations want a more automated IT environment.

A driving force behind the push toward composable infrastructure is the disaggregation of resources, which theoretically unlocks a range of benefits across IT environments. Figure 1 looks at those possibilities.

Figure 1: Inside the Possibilities of Composable Infrastructure

Source: 451 Research, 2021

POSSIBILITY	BENEFIT
Purchase infrastructure resources for individual domains when required.	Cost savings Efficient resource utilization (e.g., less overprovisioning)
Create VMs on the fly for specific workloads.	Provisioning efficiency
Use any server, PCI endpoint device, storage, persistent memory (e.g., 3D XPoint), accelerator (e.g., GPU, FPGA, ASIC) and/or network adapter.	Infrastructure flexibility and freedom Precisely tuned resources for specific workloads
Choose operating system (e.g., Windows, Linux), operating environment (e.g., bare metal, VMs, containers), orchestration tools, storage OS, AI framework, big data frameworks and resource/job managers.	Software flexibility and freedom Match software to in-house expertise
Use ready-to-run third-party orchestration software.	Deep flexibility and customization via communication through command line, RESTful APIs and Redfish APIs

According to our *VotE: Compute Infrastructure, Budgets & Outlook 2020* study, 43% of organizations say their preferred approach to IT infrastructure management is automated with manual exception handling, while 40% say they prefer manual infrastructure management with limited automated tools. A closer inspection of the data reveals that senior IT managers are more likely than IT/engineering managers and staff to prefer automated IT management, indicating the possibility that the staff physically configuring infrastructure and provisioning resources prefer greater control over IT – and may represent speed bumps on the road to wide-scale automation. There remains a perception across the IT community that automation will lead to some loss of control over IT resources, but our discussions with organizations indicate that the introduction and expansion of automated tools and policies does not appear to diminish control; rather, these processes simply change the plane of control from manual to automated.


When examining the likelihood to automate certain IT processes (rather than all IT management), the picture becomes more telling. For example, 77% of organizations are either already automating resourcing provisioning or likely to automate it. This shift toward automated provisioning is larger in organizations with at least 10,000 employees, where 89% are already automating provisioning or are at least somewhat likely to do so, while only 3% are not at all likely (see Figure 2). These organizations are also far more likely to be executing digital transformation strategies; 70% of them already have digital transformation strategies in motion, compared with just 44% of organizations with fewer than 10,000 employees. They are also more likely to be more bullish on new technology adoption, with 75% classified as Technology Leaders in our study, compared with 67% among organizations with fewer than 10,000 employees.

Figure 2: Likelihood To Automate Resource Provisioning

Source: 451 Research's Voice of the Enterprise: Compute Infrastructure, Budgets & Outlook 2020

Q. What is your organization's likelihood to automate each of the following processes? – Resource provisioning

Base: All respondents




This data indicates that large organizations are significantly separated from the customer pack in terms of infrastructure modernization. In general, larger environments represent a more promising opportunity for composable infrastructure. According to our *VotE: Hyperconverged Infrastructure, Strategy & Workloads 2020* study, 19% of organizations with at least 1,000 employees have composable infrastructure in use, compared with just 9% for organizations with fewer than 1,000 employees (see Figure 3). Similarly, 33% of larger organizations are considering composable but have no plan to implement, compared with 23% of smaller organizations. Considering the provisioning power of composable platforms and the desire to automate that process among large organizations, the higher adoption and interest among these organizations makes sense.

Figure 3: Composable Infrastructure Implementation Status

Source: 451 Research's Voice of the Enterprise: Hyperconverged Infrastructure, Strategy & Workloads 2020

Q. What is your organization's current implementation status for composable infrastructure? Note: By 'composable infrastructure,' we mean 'compute, storage and network modular infrastructure configured into pools of resources that are dynamically allocated for workload requirements.'

Base: All respondents, abbreviated fielding



Although composable adoption remains nascent, there appears to be more opportunity moving forward, with 33% of large organizations considering it. We also see pockets of opportunity on a broader level, particularly when we examine the uptake of software-defined infrastructure (SDI), which can include just one element of the infrastructure stack, such as networking. According to *VotE: Digital Pulse, Budgets & Outlooks 2021*, 29% of organizations have SDI in use, 15% have it in discovery/POC, and another 28% plan to implement it or at least are considering it.


A closer examination of those using SDI reveals that larger organizations and digital transformation leaders are far more likely to be using it (see Figure 4). For example, 41% of organizations with 1,000+ employees are using SDI compared with just 20% of organizations with fewer than 1,000 employees. Similarly, 38% of digital transformation leaders are using SDI, compared with just 20% of digital transformation learners and laggards. Perhaps most telling about the future of SDI is the difference between these organizations in terms of not using it or having no plans to use it – large organizations and active digital transformers are far less likely to have no plans for it.

Figure 4: Software-Defined Infrastructure Implementation Status

Source: 451 Research's Voice of the Enterprise: Digital Pulse, Budgets & Outlook 2021

Q. Please describe the level of usage within your organization for each of the following technologies. – Software-defined infrastructure (SDI)

Base: All respondents, abbreviated fielding



The bricks that support full-scale composable infrastructure are currently being laid by SDI, whether through technologies that enable deployment and provisioning automation, software-defined networking (SDN) or other innovations. The continued penetration of SDI into enterprise IT environments could ultimately set the stage for complete infrastructure composability, at least in terms of IT team comfort and experience. These steps are critical to further adoption of composable infrastructure because frustration with SDI – for example, with complex configuration or perceived skill requirements – could push some customers to abandon further on-premises infrastructure modernization and accelerate migration to public cloud services.

IT Deficiencies Open the Door

A misconception in the market is that composable infrastructure requires a complete rip-and-replace of existing infrastructure.

The shift to a fully software-defined environment generally will not occur overnight or across an otherwise short period of time. Depending on the approach, an organization may choose to deploy purpose-built composable appliances, or software designed for use with existing infrastructure. The key behind either deployment choice is a unified API that provides one interface for the full management of the composable infrastructure, including discovery, inventory, configuration, provisioning, updates and diagnosis. This unified API also enables staff to program a wide range of infrastructure elements with a single line of code.

For some organizations well along the digital transformation path, pieces of an existing infrastructure environment are modernized to allow for composability, similar to how an organization migrates portions of its environment to public cloud over time. In practice, it is likely only the most well-equipped, well-funded and well-skilled teams that encounter broad success with these rollouts, but it is nonetheless notable that composability is occurring without the deployment of purpose-built composable infrastructure.

“We used some outside tools and put together internal stuff to do our own software-defined datacenter applications. So clients will be able to go into our service management package and say, ‘I want to have this.’ The package will figure out, ‘You can have these in the cloud. You can have these on premise, these at the primary datacenter, these at the backup datacenter,’ and it’ll provision them for them, just like you can in the cloud, but it’ll be integrated together.”

- Mid-level manager, 50,000-99,000 employees, \$10bn+ revenue, Healthcare

Source: 451 Research’s Voice of the Enterprise: Compute Infrastructure, Budgets & Outlook 2020

On the storage side, the roots of composability already reside in existing technology, including storage-area networks (SANs), where pooled storage resources provide high levels of flexibility for large application estates that present a wide, consistently shifting range of storage requirements. Network-attached storage (NAS) deployments also deliver pooled storage resources, albeit without the dedicated network provided by SANs. Despite the flexible advantage of delivering shared storage, a common denominator across SAN and NAS is their dependence on networks to deliver resources, which generates complexity (particularly with SAN) and performance issues when network traffic is heavy. Direct-attached storage (DAS) eliminates the network reliance, but these deployments cannot pool storage and cannot be scaled to anywhere near the extent of SAN or NAS.

HCI addresses the complexity and scalability issues inherent in these legacy storage mediums through the implementation of fully virtualized storage and compute resources. However, a standard HCI platform is designed only for virtualized applications or containerized applications

running within VMs, which presents limitations on the compute side of the resource equation. The typical application estate in large enterprises contains a large number of virtualized applications, but bare metal resources remain in the mix for certain requirements. Some of these are performance-related, as bare-metal servers do not face the resource overhead that VMs do, and some applications benefit from access to features present in the underlying compute hardware. Other factors also influence the continued use of bare-metal servers. For example, some applications may have license restrictions that could lead to escalating costs if virtualized, or may have been developed for use with bare metal and would need expensive reengineering to virtualize. Further, some organizations continue to have IT staff that is well-trained around bare-metal server deployment and provisioning and can operate at efficient levels on par with staff overseeing virtualized resources.

Public cloud service providers (such as AWS and Google Cloud) offer a wide range of CPU and GPU bare-metal instances, and some organizations have become adept at using these resources in a cost-effective manner. While highly effective for technical bare-metal requirements, controlling costs and the complexity involved with running applications in multiple locations and potentially outside of a purposely designed hybrid cloud management system can prove challenging for some teams, which might need to assign different staff members or subteams to manage on- and off-premises infrastructure. This challenge, in theory, can be avoided if all infrastructure services – whether on- or off-premises – are delivered as a service.

“I don’t think we’ll ever move away from a bare metal completely, obviously, I think, like most people. So we’ll probably have our traditional storage architecture running on the proprietary hardware. But being able to expand it to the cloud, and run that same software, I think it’s the benefit [of SDI]. Just be an expansion of our datacenter, whether that’s a DR perspective or whether that’s a cloud-only offering, that we could replicate back and forth.”

- IT/engineering manager/staff, 10,000-49,000 employees, \$10bn+ revenue, Energy

Source: 451 Research’s Voice of the Enterprise: Hyperconverged Infrastructure, Technology & Platform Innovation 2020

The agile design of composable infrastructure allows resources to be composed and recomposed, depending on ongoing application needs. When a job is completed based on an application request, compute, storage and/or networking resources are freed and placed back into the resource pool. Perhaps most importantly, these actions occur without the need for manual oversight, which is a critical step toward building a cloudlike environment that supports all applications (not just those natively designed for cloud). Because the average large IT ecosystem contains multiple infrastructure variants, including bare metal and virtualized servers, GPUs and other accelerators, VMs and containers, and a mix of SAN, NAS and DAS storage, composability can remove the burden of manually assigning the right resources to applications.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its Web sites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.

Western Digital OpenFlex

The architecture allows storage to be disaggregated from compute, enabling applications to share a common pool of storage capacity. Data can easily be shared between applications or needed capacity can be allocated to an application regardless of location. By decoupling software from the underlying platform, organizations can build solutions with greater flexibility spanning the portfolio of OpenFlex infrastructure offerings. This provides a decisive step forward in reducing the cost of ownership for data center deployments to come, and it helps address all the problems caused by the volume, variety, velocity, access speed and consistency of data.

In addition to enabling NAND flash media access using NVMe-oF, the OpenFlex architecture ultimately enables disk and other IT components such as GPUs, FPGAs and even tape to be accessed via NVMe-of so that all data center storage can be addressed in the same way. The Western Digital NVMe-oF architecture is a huge step towards the software-defined data center—allowing storage to be assigned to applications without regard for where it is physically located. This is the essence of “composable infrastructure” where physical resources (compute, networking, storage) can be logically and dynamically configured and treated as a resource for a specific application without the need for physical configuration. Western Digital has established the [Open Composable Compatibility Lab](#) (OCCL). The OCCL is where different vendors can come together to test and verify their Open architecture. The OCCL underscores Western Digital’s commitment to having a dynamic community of technologies supporting composable disaggregated infrastructure (CDI), which is an important step in providing the flexibility needed to keep up with the rapid rate of change in today’s information environment.

Learn more about the Western Digital OpenFlex [architecture](#) and [product line](#).

Western Digital®