EMA
IT & DATA MANAGEMENT RESEARCH, INDUSTRY ANALYSIS & CONSULTING

Prepared for:

Pluribus NETWORKS

The State of Data Center Networking: 2021 Annual Report

May 2021 EMA Research Report

By Shamus McGillicuddy

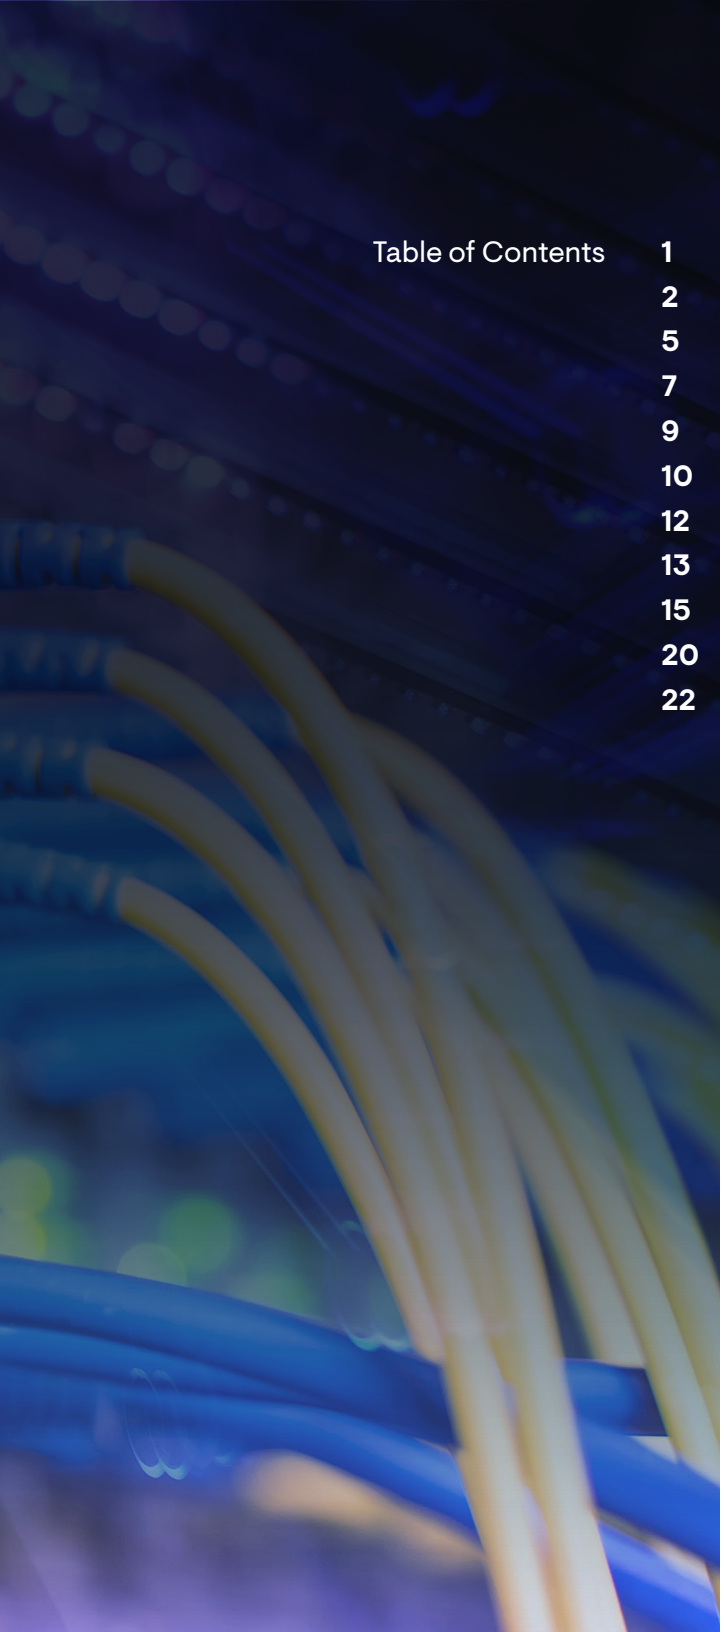

Table of Contents

Executive Summary

In this first annual research report on The State of Data Center Networking, commissioned by Pluribus Networks, Enterprise Management Associates found that IT organizations are aggressively pursuing high-availability application architectures across geographically separated data centers.

These new high-availability data center strategies are severely challenged by network architecture and network operations complexity. To address these issues, IT organizations are transforming their data center networks with network overlay solutions and network automation technology. The findings in this report are based on a spring 2021 survey of 263 North American and European data center infrastructure professionals whose companies operate at least two data centers.

The Cloud is Not Retiring Data Centers

Conventional thinking holds that all applications are moving to the public cloud and that corporate data centers are an endangered species going the way of the dinosaurs. The conventional thinking is wrong.

Yes, it is true that more applications are moving to the public cloud, but the private cloud is still the center of gravity for digital infrastructure. **Figure 1** shows that the overall share of applications living in the public cloud will grow modestly over the next two years, while the share that live in private, on-premises data centers will shrink modestly. Hosted private cloud will also increase somewhat while applications in privately managed colocation data centers will stay rather flat.

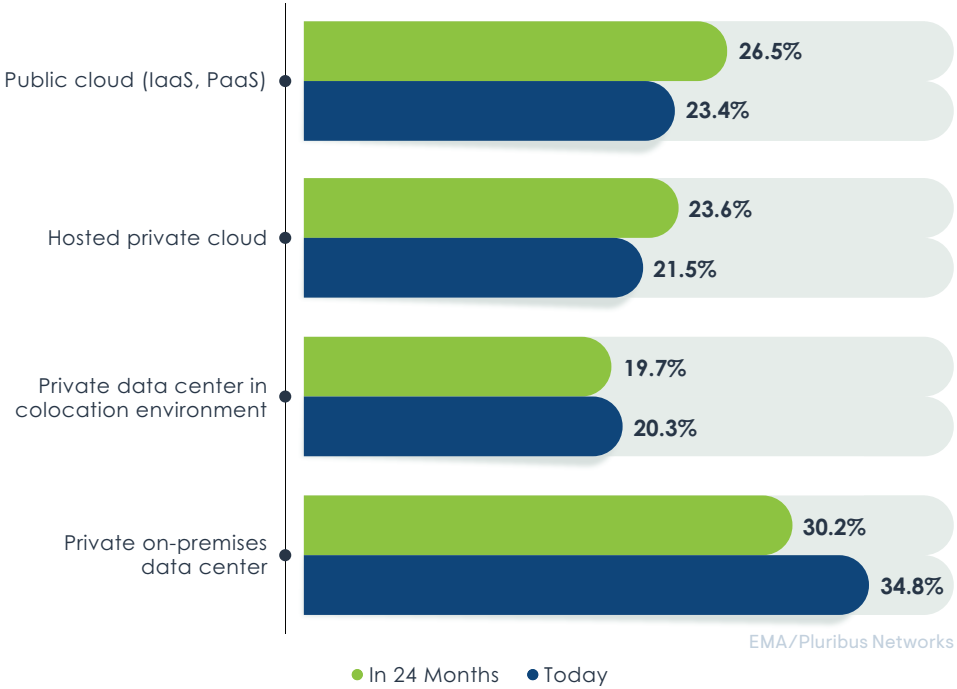

Figure 1. What percentage of your applications reside in the following locations, today and 24 months from now?

Why are many applications remaining in data centers? EMA asked survey participants to identify their top two reasons for keeping some applications in on-premises data centers and private clouds. Their top four responses were:

1. **Security** (48.3%): Many enterprises have application workloads that simply require private security controls and analytics to ensure data is protected. Many survey respondents do not believe that the public cloud providers can meet their security requirements.
2. **Performance** (40.7%): Certain applications have performance requirements that dictate private cloud infrastructure. In some cases, they may want hardware that is optimized for a certain kind of application workload. In other cases, their overall enterprise network is optimized to provide high-performance applications via private infrastructure.
3. **Compliance** (36.9%): Industry and government regulations often restrict certain applications to private cloud and on-premises data center environments. For instance, some financial institutions will not allow customer data to reside in a public cloud. Additionally, many governments have data sovereignty laws that prevent companies from hosting certain classes of data in another country, where the preferred cloud provider is located.
4. **Cost** (29.3%): Many enterprises encounter unexpected cost overruns in the public cloud. Cloud providers often charge data egress fees which make it very expensive to move data in and out of the public cloud. Additionally, over time, the recurring operational costs of running applications in the public cloud can become more expensive than maintaining private infrastructure.

This research surveyed enterprises with a minimum of two data centers, but we found that larger enterprises have far more than that number. As revealed by **Figure 2**, the majority of midmarket enterprises (250 to 2,499 employees) have three to five data centers. And the vast majority of the largest enterprises (10,000 or more employees) have 6 or more data centers, including more than one-third that have 11 or more.

Furthermore, **Figure 3** reveals that more than half of all companies in this research will actually add data center sites to their overall application architecture over the next two years. Nearly 30% are maintaining their current number of data center sites, but they're either modernizing those existing data centers or migrating to new facilities.

Contrary to commonly held beliefs, it's actually quite rare for enterprises to reduce their overall data center footprint. Only 12.5% are consolidating. This may surprise some industry pundits who have been forecasting a consolidation and reduction of data center sites due to public cloud. Clearly that is not the case. Enterprises may migrate some on-premises data centers to colocation facilities, but it's clear that they want to maintain private infrastructure. Data centers and private clouds will be essential to global digital services moving forward.

More than half of all companies will add data center sites to their overall application architecture over the next two years.

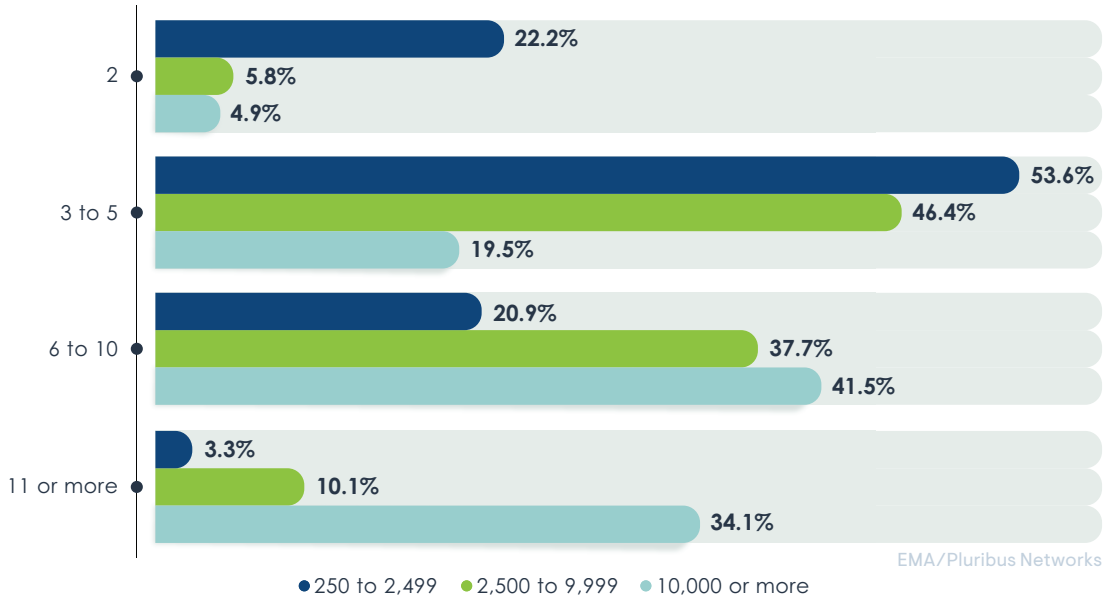

Figure 2: How many data center sites does your IT organization maintain and operate? By number of employees.

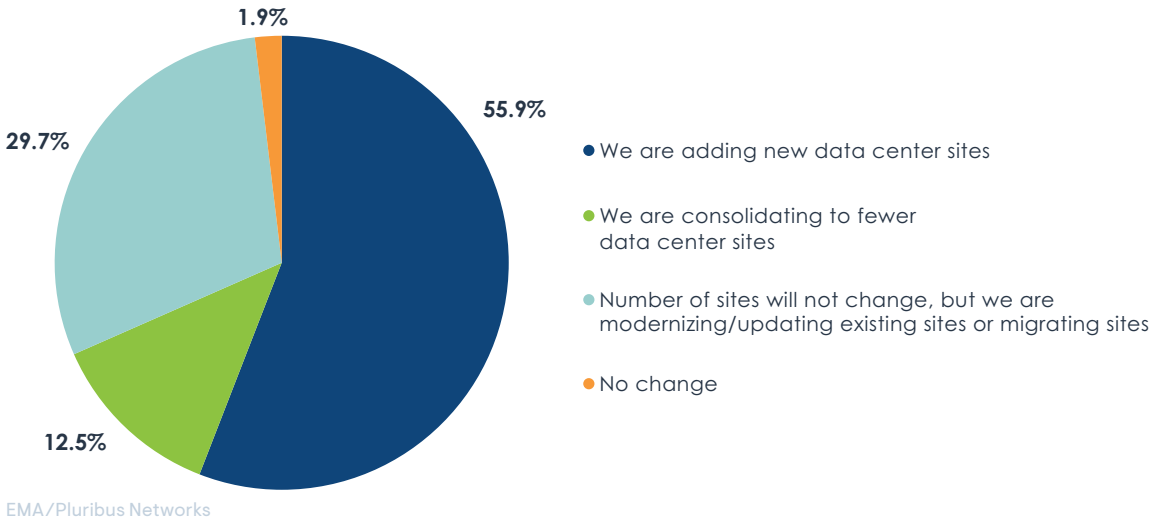

Figure 3. Over the next 24 months, which of the following best describes how your overall data center footprint will change?

Sample Size = 263

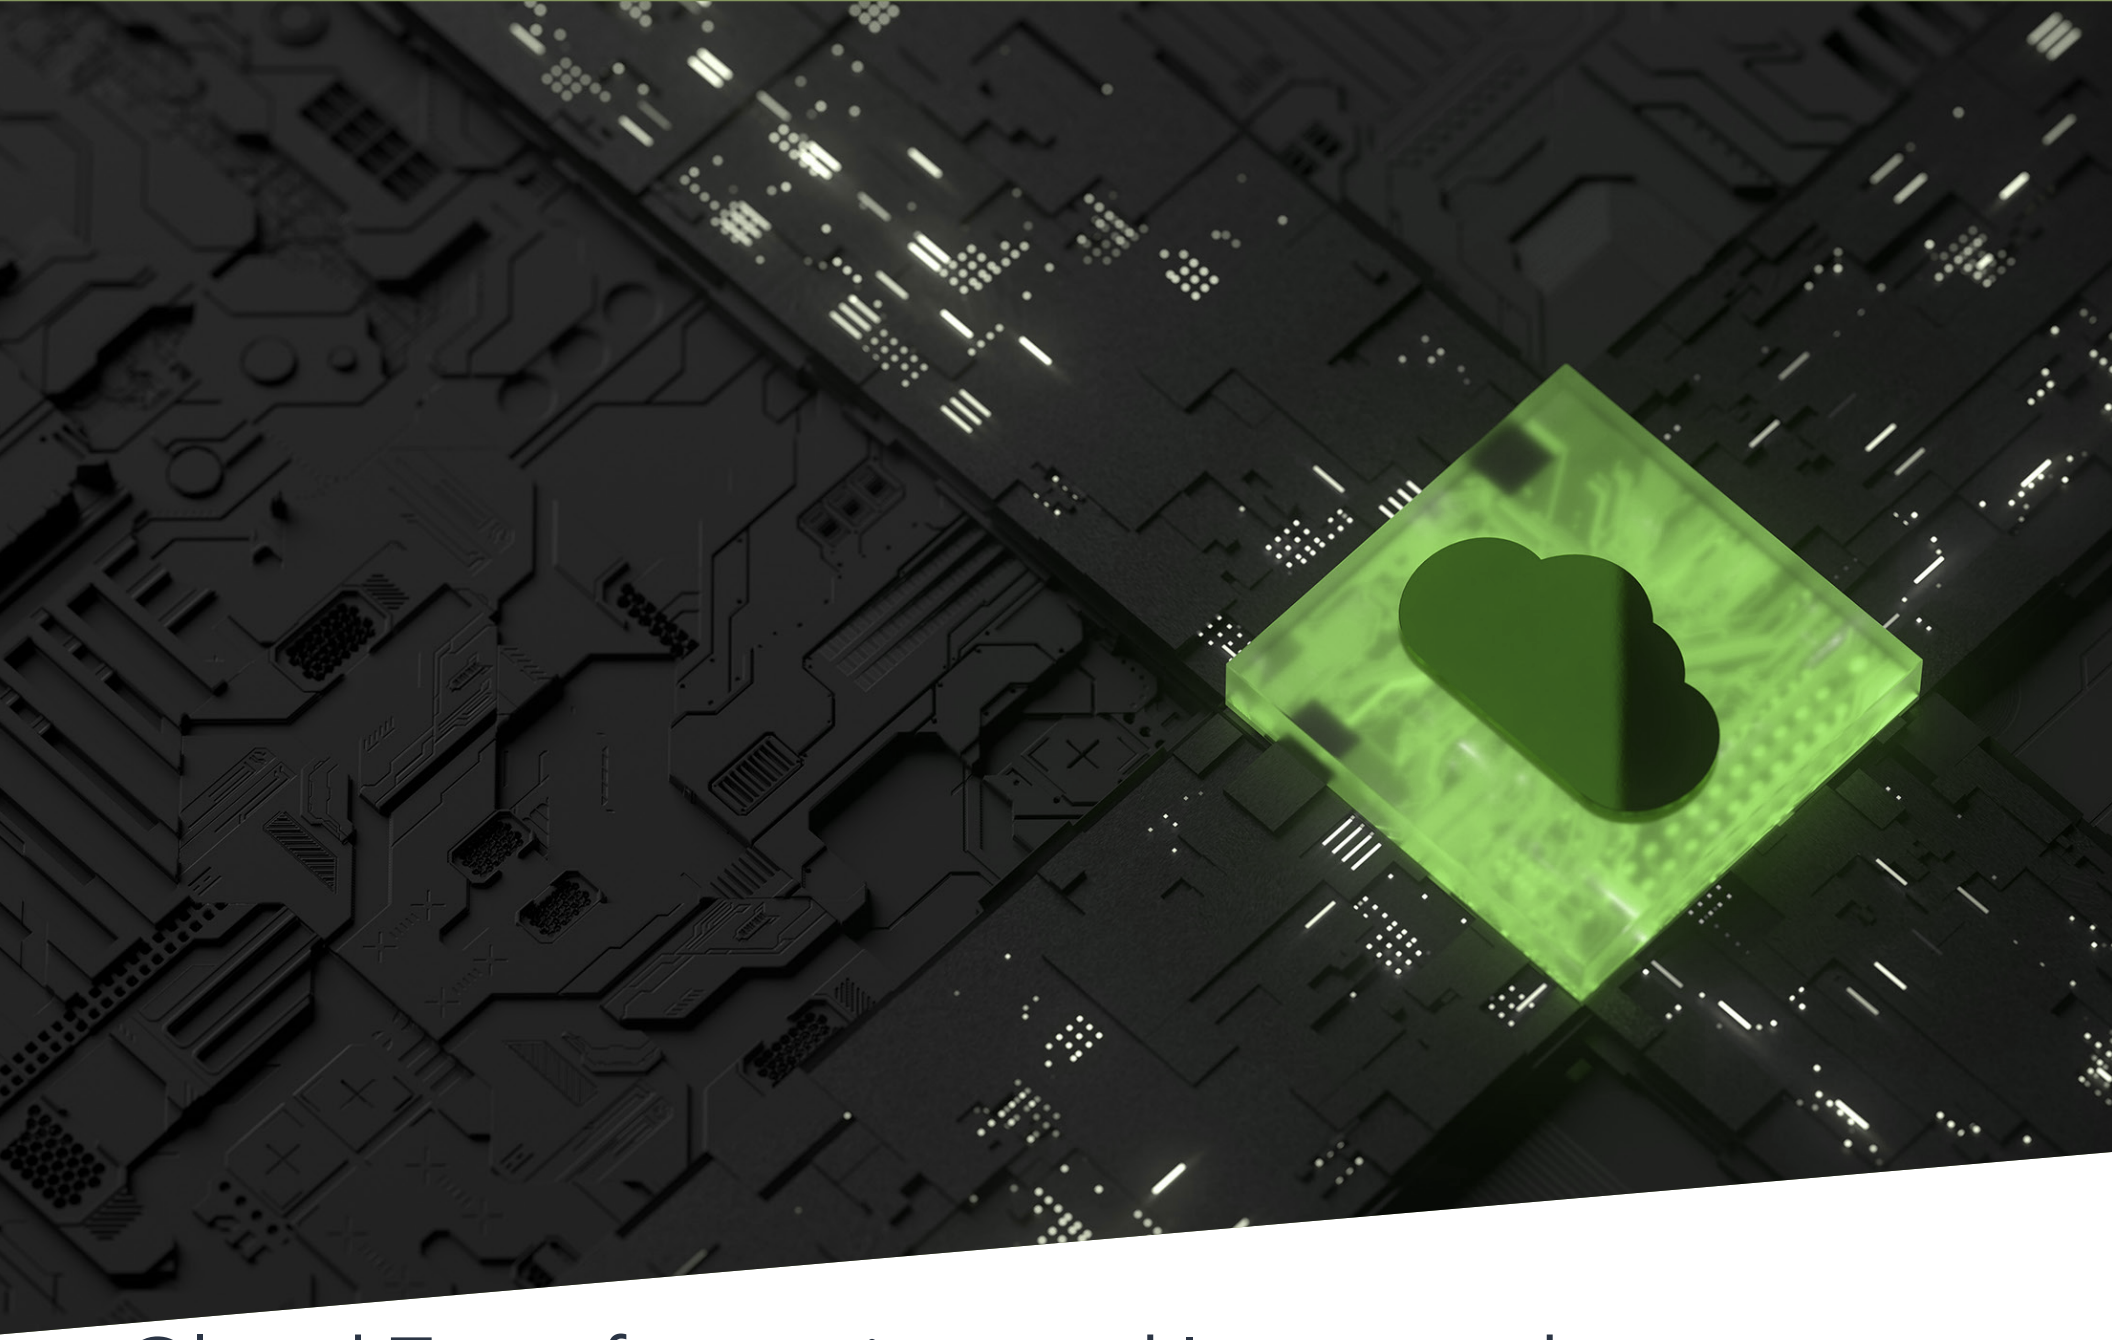

Cloud Transformation and Improved Application Experience are Driving an Increase in Data Center Locations

What's driving the continued importance and increased number of data centers? EMA asked IT professionals to identify the top two drivers of their multi-data center strategies. The top two answers were:

1. **Cloud transformation** (54.8%): Enterprises are establishing private cloud infrastructure in their data centers so that they can play a role in hybrid cloud and multi-cloud architectures. They are also adding data center sites to provide more on-ramp locations to the public cloud.
2. **Application experience** (39.5%): Enterprises are leveraging multi-data center footprints to develop global availability zones for applications.

Less frequent, but still prominent drivers (26% to 31%) include security initiatives, disaster recovery, increased numbers of remote workers, compliance, and edge computing.

High-availability data center architectures, such as active-active and active-hot standby configurations, can ensure that private cloud infrastructure are able to support global availability zones and participate in a multi-cloud architecture. **Figure 4** shows that most companies are pushing hard toward these high-availability architectures.

The majority of companies have only limited deployments of high-availability data centers today, but within two years more than 80% will have broad deployments of these architectures. This research reveals that IT organizations have ambitious plans for their data center footprints, and they are moving to active-active architectures.

The majority of companies have limited deployments of high-availability data centers today, but within two years more than 80% will have broad deployments.

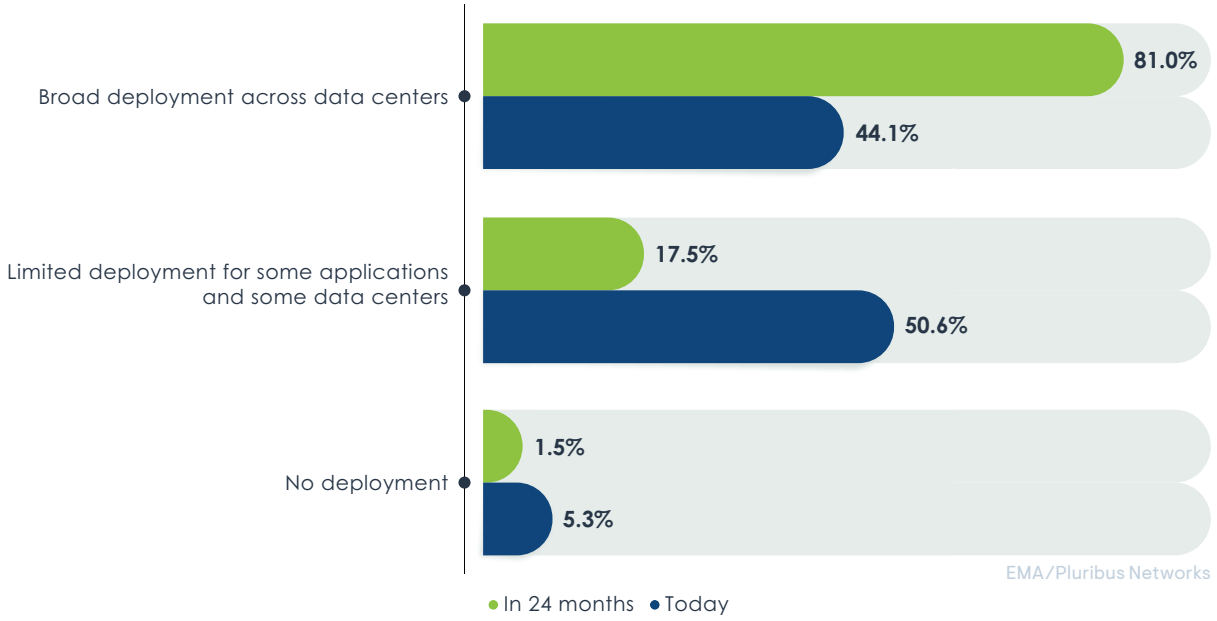

Figure 4. Deployment of active-active or active-hot-standby data center architectures to enable high application availability with near-immediate failover

Networking is the Biggest Obstacle to Deploying Multi-Site Data Centers

Figure 5 identifies the primary barriers to establishing high-availability, multi-data center architectures. Complexity of network architecture and complexity of network operations are the biggest issues. Application architecture complexity and skills gaps are just secondary issues. And less than 15% of enterprises are held back by a lack of a business case for high-availability data centers.

Enterprises that have a larger number of data centers and those who are adding more data centers to their overall infrastructure footprint are the most likely to struggle with network architecture and network operations complexity. As they add more sites to a multi-cloud architecture, they see this complexity compound.

Clearly, companies need to transform and modernize networks to eliminate the complexity that undermines this vision of global application availability and hybrid multi-cloud. They will need to establish data center networks that can span multiple sites and are highly automated, with comprehensive visibility and a consistent policy driven approach.

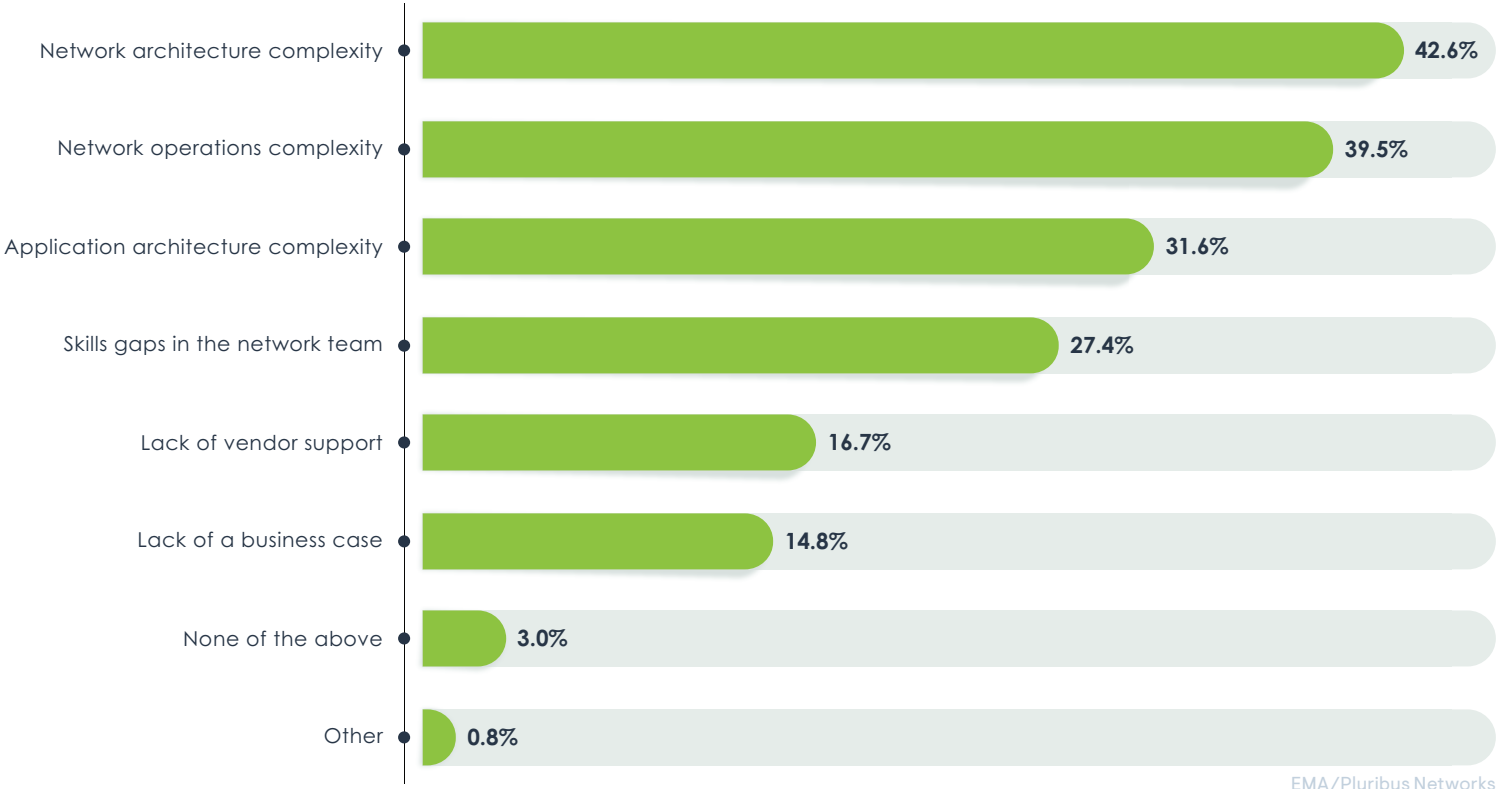

Figure 5. Most challenging aspects of adopting active-active or active hot-standby multi-data center architectures

How Enterprises are Transforming Networking for Data Centers of the Future

EMA suspects that traditional approaches to network architecture and operations are holding back data center strategies today. To reduce the complexity that is challenging multi-site data center architecture, infrastructure teams must build multi-site networks that are flexible, abstracted, and automated.

Network Virtualization via Overlays

Modern data center networks typically feature leaf/spine "underlay" architectures with Clos fabric designs. These networks provide consistent bandwidth and single-hop reachability for all application components, which provides predictable performance for east-west traffic. However, the physical underlay can still be complex and difficult to manage, limiting agility, speed to service, and availability. Infrastructure teams are attempting to reduce this underlay complexity with network virtualization via overlay technologies.

Overlay data planes are typically based on tunneling protocols such as VXLAN and GENEVE. There are two approaches to data center overlays today. The first is a host-based or compute-based solution, with overlay network endpoints (e.g., VTEPs) deployed on a hypervisor. The second is a switch-based overlay where the overlay tunnels terminate on the switch network processor. Switched-based overlays can either use a software-defined networking (SDN) control plane solution or a BGP EVPN-based protocol control plane.

All of these overlay technologies add a virtual abstraction layer on top of the network underlay that reduces architectural and operational complexity. The abstraction layer allows enterprises to treat the underlay as a simple and scalable IP fabric that is easier to maintain. The overlay becomes a service layer for the network, where ongoing data center network engineering and operations are performed. Because the overlay is software-defined, it typically delivers a higher level of agility and services that can be deployed faster versus delivering services from the underlay.

Overlay solutions usually offer hooks for automation and orchestration pipelines, making network administration in the data center easier. For instance, the infrastructure team can maintain the IP fabric underlay using their well-established tools and processes. The DevOps team can use their automation tools to deploy Layer 2 and Layer 3 services in the overlay, including policy frameworks, quality of service tiers, and general moves, adds, and changes. The overlay can also provide a more scalable approach to network segmentation than traditional VLANs, which can prevent or slow down lateral movement inside the data center if the perimeter is penetrated.

Figure 6 shows robust adoption plans for overlays. While only 33.5% of companies have overlays deployed today, more than 80% will have them within the next two years. Overall, 92.6% of companies plan to leverage overlay technology in their data centers at some point in the future.

While only 33.5% of companies have overlays deployed today, more than 80% will have them within the next two years.

This research found that two-thirds of enterprises are using or planning to use a switch-based overlay (e.g., SDN or BGP EVPN).

EMA asked research participants to rank the importance of network overlay use cases in their data centers. Network virtualization for agility and automation emerged as the most important, followed closely by network segmentation for security. The use of overlays to stretch network fabrics across data center pods or sites ranked a very close third. Stretching overlays across data center sites is the foundation of the active-active architecture many enterprises are pursuing with their multi-data center strategies today.

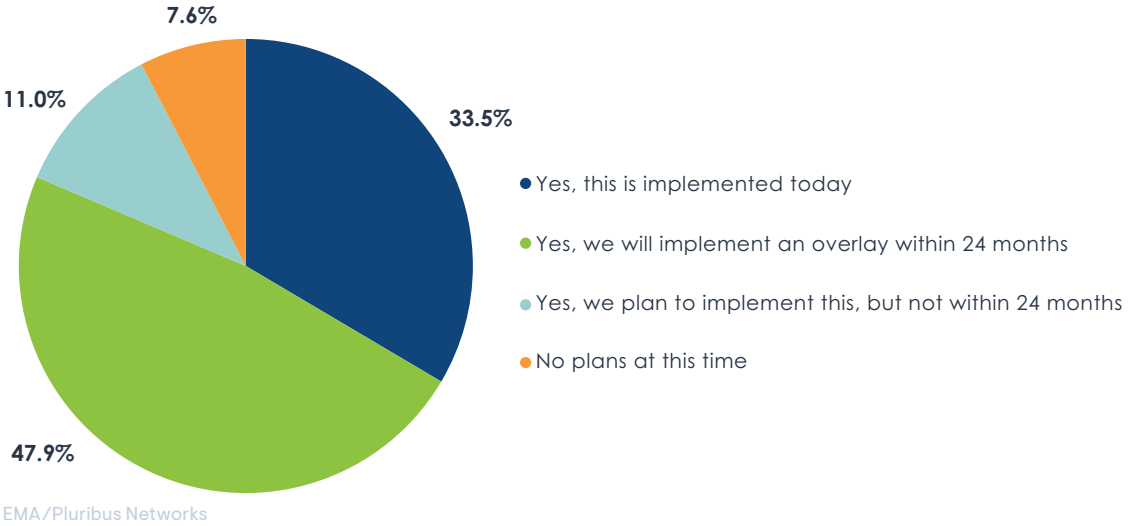

Figure 6. Does your organization have any plans to deploy a network overlay in one or more of your data centers?

Stretching Network Overlays for Active-Active Application Architecture

The high-availability architecture of active-active and active-hot standby applications requires a common network environment across data centers. Infrastructure teams can achieve this by stretching a network overlay across data center borders. This provides stretched Layer 2 and Layer 3 services. Stretching a Layer 2 overlay on top of a Layer 3 underlay is a reliable architecture that can enable live migration of workloads and support protocols that prefer Layer 2 (e.g., vSAN) for data replication. When integrated into the global network architecture, stretched overlays allow companies to direct application requests to the best location, improving overall experience.

Figure 7 reveals that more than 95% of enterprises with overlay interest want to stretch those overlays across geographically separated data centers. Companies that have already broadly implemented active-active or active-hot-standby architectures across their data centers are twice as likely as others to be stretching overlays across those data centers today. While most infrastructure teams have implemented or plan to implement stretched overlays, 14.4% of companies want to stretch overlays but haven't found a workable solution. This points to some issues with various vendor overlay architectures that make it more difficult to enable stretched overlays.

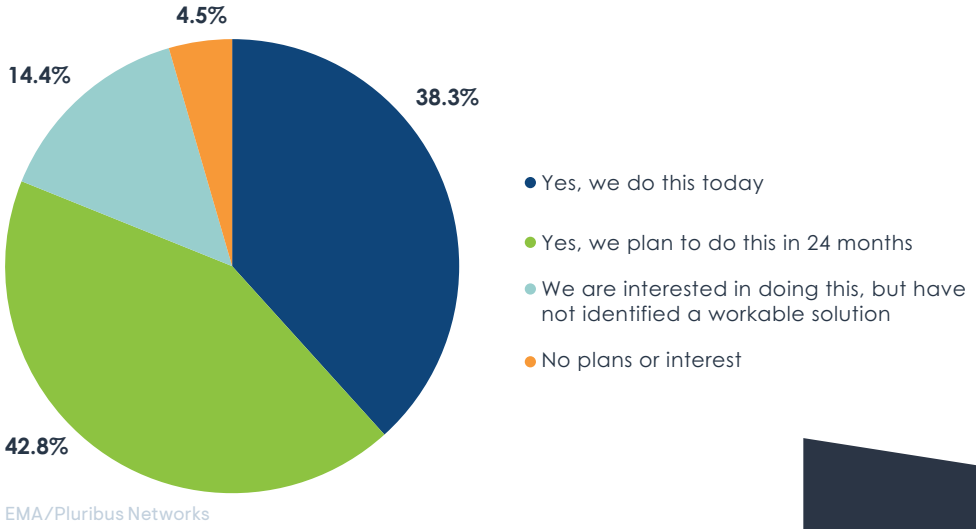

Figure 7. Do you have any interest in stretching your data center network overlay across multiple geographically separated data centers?

95% of enterprises with overlay interest want to stretch those overlays across geographically separated data centers.

Apply Network Automation to the Underlay and Overlay

While network overlays provide abstraction, not all of them provide sufficient automation. EMA analysts have spoken to multiple infrastructure teams who have applied third-party automation tools to overlay management. While reducing complexity with overlays, infrastructure teams should also think about how to apply network automation to the transformed data center network. For instance, infrastructure teams that have broadly deployed active-active data center architectures are nearly three times more likely to have broad deployments of network automation in place today than infrastructure teams with only partial active-active architectures.

70% plan to have a high level of automation implemented by 2023.

Enterprises need to reduce the number of manual processes that persist in Day 1 and Day 2 network management. Automation makes networks more agile and responsive to change, and it also makes them more reliable. EMA research has found that 25% of all network trouble is caused by manual errors,[1] something that automation can help eliminate. It also addresses skills gaps in the network team, where high level network engineers are often stretched too thin to address business priorities.

Figure 8 shows that over the next two years infrastructure teams are poised to dramatically expand the amount of network automation they have in their data centers. While only 36.9% claim to have network automation broadly deployed in their data centers today, 70% plan to have a high level of automation implemented by 2023.

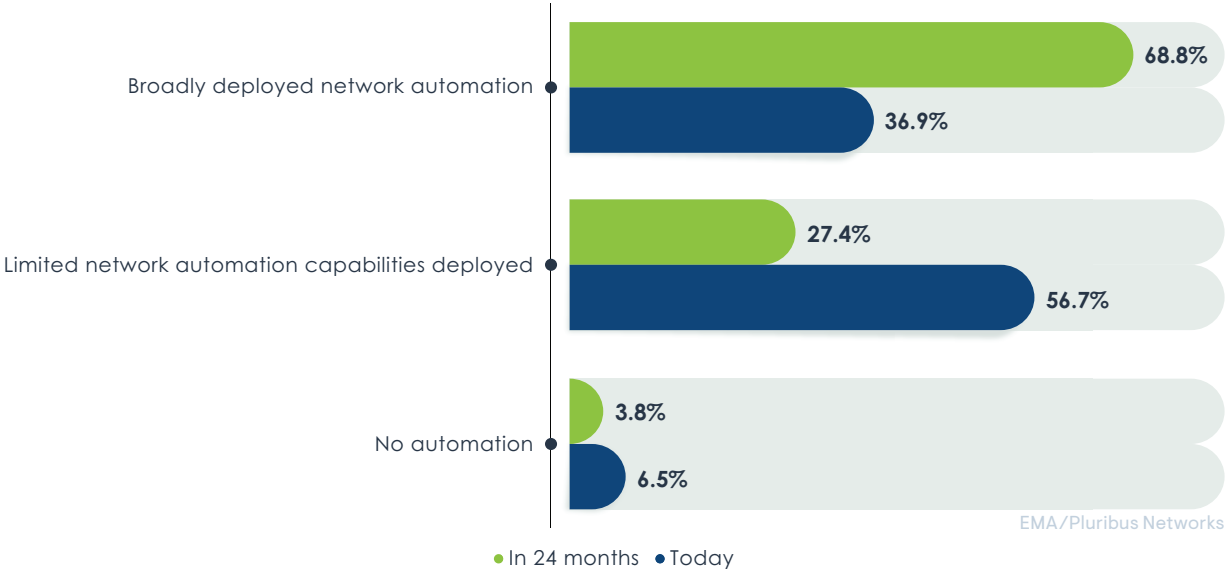

Figure 8. Level of network automation in your data center networks today and in 24 months

[1] EMA, "Network Management Megatrends 2020," April 2020.

Figure 9 reveals how infrastructure teams are planning to automate data center networks. It shows what they are using today and what they expect to be using in 24 months. There are two mainstream solutions. Today, nearly half are using SDN solutions that program the underlay and/or overlay. Slightly more than half are using server-based overlays with SDN automated control planes. Both of these classes of network automation will grow modestly in use over the next two years.

Internally developed automation (e.g., Python scripts and modified open-source tools) is also quite popular today, but these solutions can add a great deal of technical debt, especially libraries of one-off scripts that can break every time a switch vendor introduces a new version of its network operating system, or that can drift due to lack of script governance. Notably, these internally developed tools will decline in favor of commercial solutions over the next two years.

Third-party network configuration solutions are used by more than a third of companies today, but they will experience some contraction by 2023. Script-based configuration solutions, which are often favored by DevOps teams, are the least popular solutions for network automation today but will experience moderate growth in use over the next two years.

More than half of the respondents will use software defined networking (SDN) automation for overlay and/or underlay networks. EMA suspects this will be one of the most popular solutions moving forward.

What's most clear from these responses is that infrastructure teams are typically using more than one automation solution, both now and in the future. EMA research has repeatedly found that IT organizations use multiple tools for network automation.

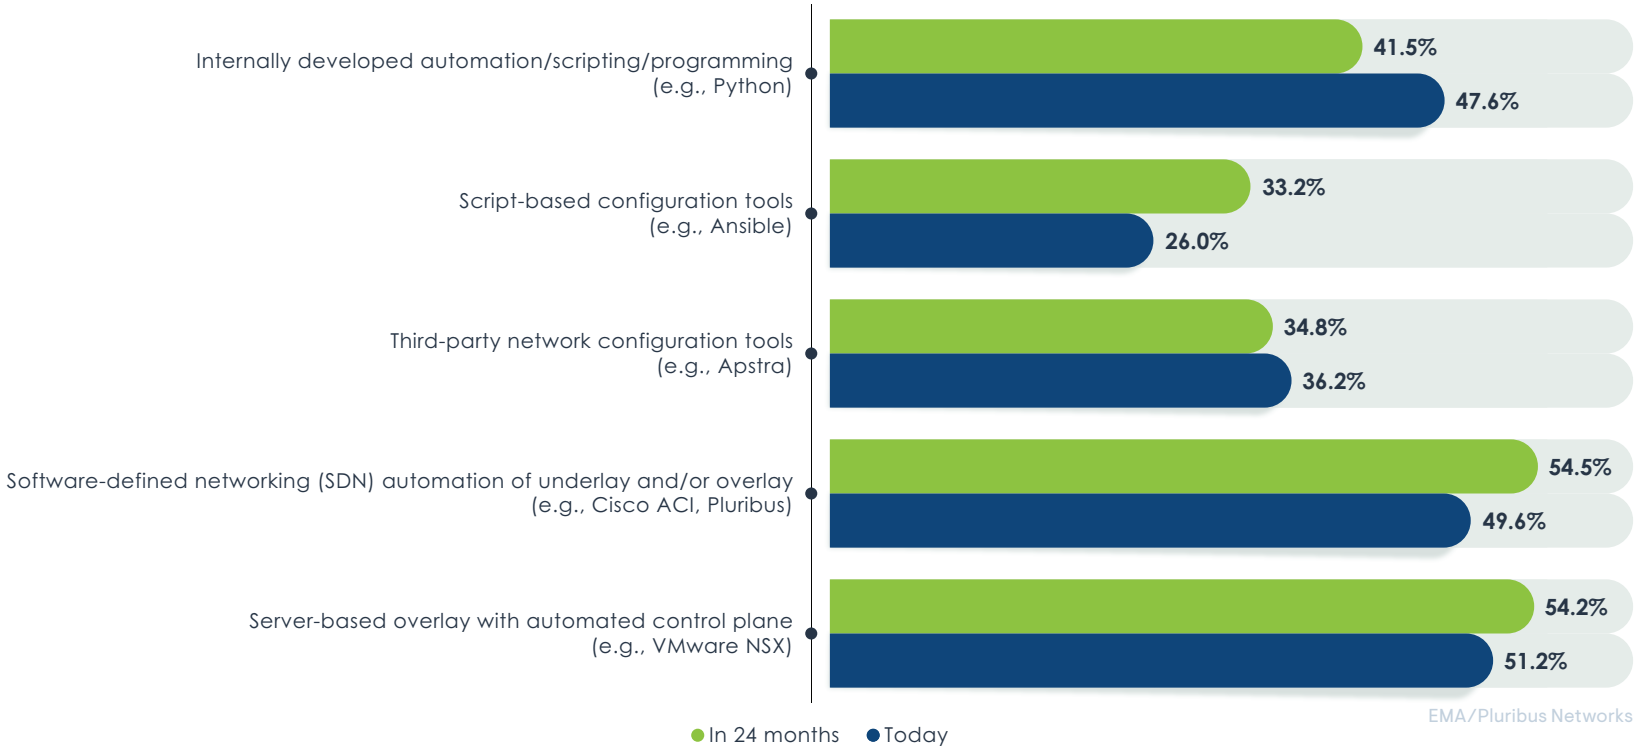

Figure 9. Types of data center network automation in use today and in 24 months

Pitfalls to Avoid with Next-Generation Network Architecture

If it were easy to transform network architecture and network operations, everyone would have done it already. Every network would be fully automated and network engineers would be much less stressed out.

The reality is different. Networking technology is fundamentally more complex than other IT disciplines because of what the technology world demands of it (any-to-any connectivity, high-availability, low latency, complex policy controls). It takes hard work to transform a network. Each step on the path forward will present technical and business challenges for most infrastructure teams. It's important to be aware of these challenges before embarking on a data center network transformation project, because forethought can help you avoid them.

Figure 10 reveals the top challenges that companies typically encounter when using overlay technology to transform their networks. The number-one issue is cost. Overlays can add a new line item to a budget. Host-based overlays are usually highly automated but can be very expensive with per host-CPU license fees, and in some cases SDN controller and gateway license fees and hardware costs. They can introduce a layer of infrastructure that didn't exist eight years ago.

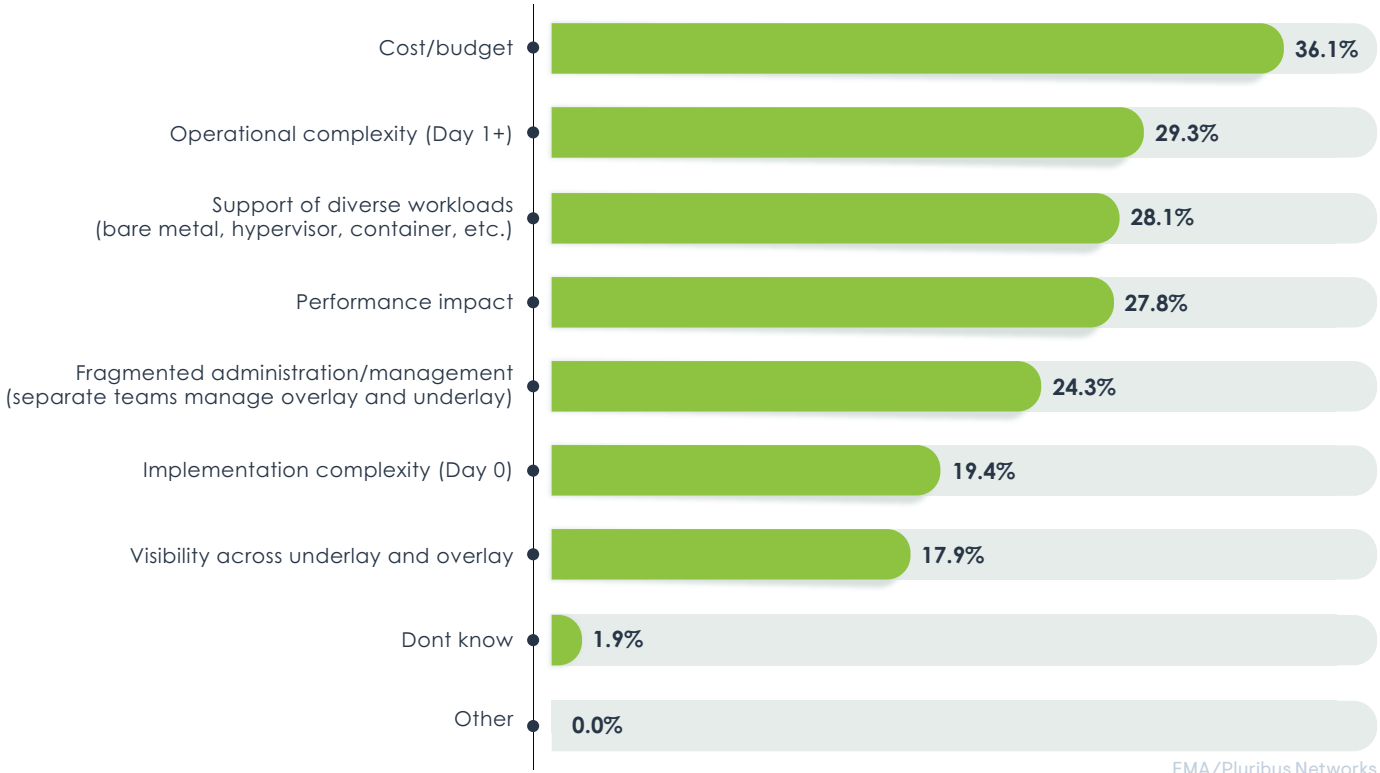

Figure 10. Top challenges associated with data center network overlay technology

Switch-based overlays have the benefit of only incurring license costs at the top-of-rack (ToR) switch layer versus compute-based overlays that charge a network license for every host. On average there are 20 hosts or servers per ToR switch pair—BGP EVPN is a IETF standard control plane implementation—but even this can add costs over and above the standard switch and switch OS. Some vendors might charge an additional license to support an BGP-EVPN feature. More importantly, BGP EVPN is not inherently automated, which could result in the need to buy additional external third-party automation solutions.

Controller-based SDN solutions are highly automated but require additional expenses for external controller hardware and software at every data center site, and multi-site orchestrators between sites. Industry best practices dictate that three controllers be deployed at every site to ensure high availability of the control plane. Controllerless SDN solutions are highly automated and they can incur additional license costs for the SDN functionality. However, they do not require external controllers and therefore may be lower cost overall.

The second most challenging issue with overlays, according to the survey, is operational complexity, including monitoring, troubleshooting, and moves, adds, and changes. An overlay can simply add more knobs and buttons to keep track of. Protocol-based solutions (e.g., BGP EVPN) can require significant box-by-box configuration efforts for every service deployed. This issue is no doubt a driver for the increase in automation shown in Figure 7.

Workload diversity is the number-three challenge. This is particularly an issue for host-based solutions, which can struggle to integrate bare-metal workloads and other devices like access switches and IoT gateways. To understand just how challenging this issue can be, note that the leading compute-based overlay vendor is encouraging its customers to migrate to a new product in order to provide more support of diverse workloads. Switch-based solutions have an easier time aggregating these different traffic sources.

The fourth leading challenge is the performance impact of overlays. When routing and switching is performed on general-purpose compute for a host-based overlay, some infrastructure teams may run into these issues.

The last major challenge is fragmented management of overlays and underlays. In many cases, separate teams are responsible for the overlay and underlay. Coordinating operations across the two groups can be a challenge. For instance, the network operations team might own the underlay, but the server team might own the overlay. These teams often perceive each other as adversaries or sources of annoyance, rather than partners.

High and complex operational overhead is the leading problem with stretching an overlay.

Implementing a network overlay presents one set of issues. Stretching that overlay across multiple data centers introduces another set of challenges. **Figure 11** reveals that infrastructure teams primarily struggle with four issues when they implement these architectures.

High and complex operational overhead is the leading problem. Monitoring and troubleshooting of an overlay that stretches across the WAN may require new tools and processes. Ideally, solutions should provide per-flow visibility across multiple sites without the need for external devices like TAPS and packet brokers.

Implementation complexity is the second-biggest issue. Many overlay vendors originally designed their solutions for a single data center and retrofitted them to support multi-site use cases. For example, these solutions often require a "controller of controllers" to stitch together multiple sites. Infrastructure teams should be aware that some overlay solutions have been better optimized for multi-site use cases. There are countless other issues that can make implementation complex. Infrastructure teams may need to coordinate management of the underlay and overlay at each data center. They might need to reconfigure data center interconnects. They may need to integrate and configure third-party automation systems. And they'll need to coordinate the IP address space across sites. The research found that enterprises with a higher number of data center sites struggle with implementation complexity more often. Companies with a large number of sites should design their multi-site networks for elegant n-site scaling.

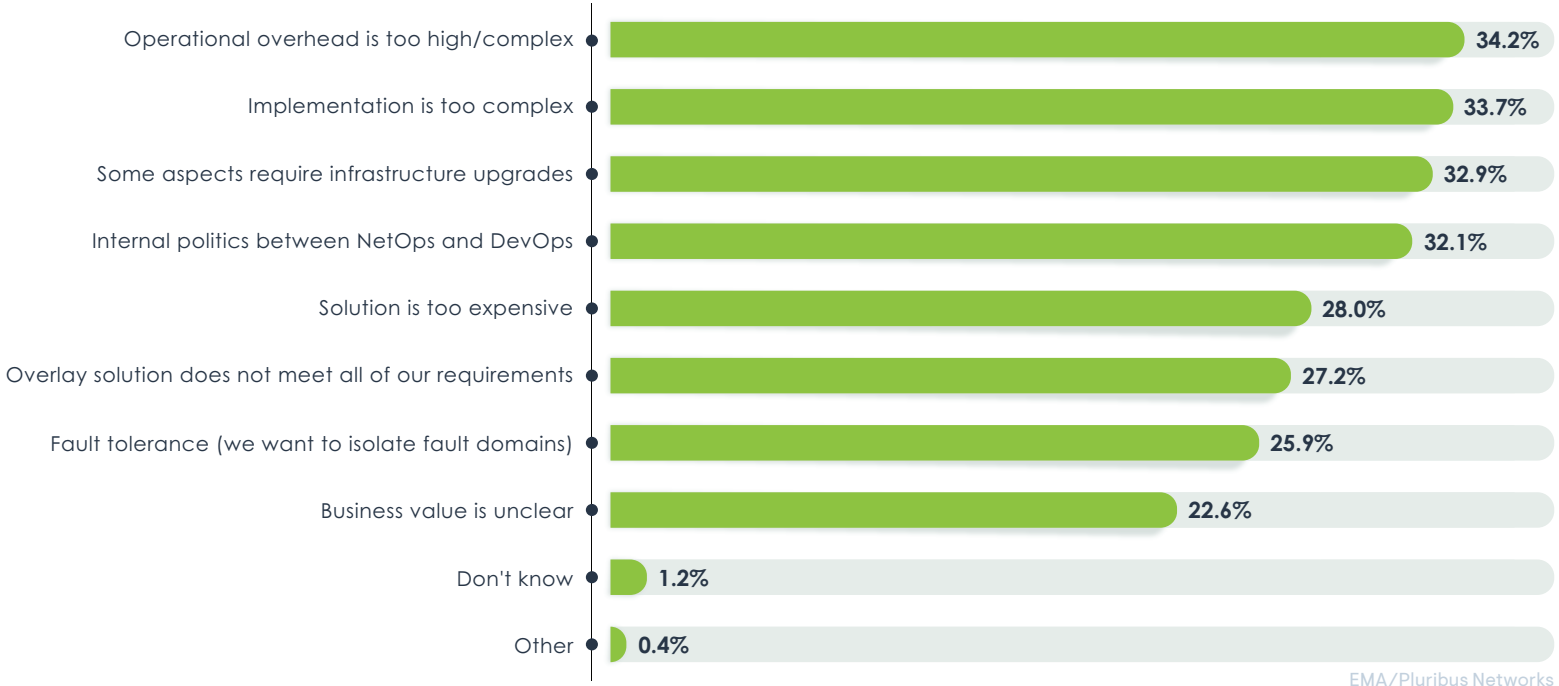

Figure 11. Top challenges associated with stretching a network overlay across multiple geographically separated data centers

Sample Size = 246

Thirdly, some aspects of stretching overlays might require infrastructure upgrades. The overlay might lack interoperability with third-party network components, such as spine switches and the routers used for data center interconnect. One or more data centers may also require a network underlay upgrade, especially if the overlay solution requires a specific hardware vendor. This will add complexity and cost to the project.

Internal politics between NetOps and DevOps is the last of the major challenges with stretched overlays. This can involve divided ownership of different components involved in implementation and operations. For instance, the network team might own the underlay networks in each data center and the interconnect between the data centers, while the DevOps team owns the overlay. The network team might think it's crazy to stretch the overlay, or the DevOps team might try to implement it without the network team's help. This is another issue that is more prominent with companies that have larger numbers of data centers. IT leaders need to make sure that all stakeholders are on the same page when stretching overlays across data centers.

Conclusion

Data centers remain a strategic asset for enterprises that are digitally transforming. While approximately 25% of workloads are migrating to the public cloud, the majority of applications will remain in the private cloud for the foreseeable future. Data centers will play a critical role in the private cloud footprint of most enterprises. These new multi-cloud architectures, where private cloud is the anchor point, are driving enterprises to add new data center locations and to implement active-active architectures to provide improved application availability and performance.

Unfortunately, legacy approaches to data center networking will hold many companies back. Enterprises need to virtualize and automate their networks using overlay networking and automation approaches such as SDN. And they need to stretch those virtualized networks across geographically separated data centers to support high-availability data centers and private clouds.

This research has proven that infrastructure teams need to implement networking solutions that are designed to span across multiple data centers and that can reduce complexity with integrated automation capabilities. Finally, all of this must be achieved affordably and with minimal complexity.

As you embark on your digital transformation journey, you should evaluate the state of your data center networks. Are they capable of supporting a high-availability, multi-cloud architecture? If not, what steps do you need to take to modernize and transform your networks? This report offers an early roadmap on how to get to where you need to be to support your digital future.

Demographics

This research is based on a survey of 263 technology professionals who are responsible for data center network infrastructure and operations. The survey was conducted during the spring of 2021. The following charts provide a demographic overview of the survey respondents.

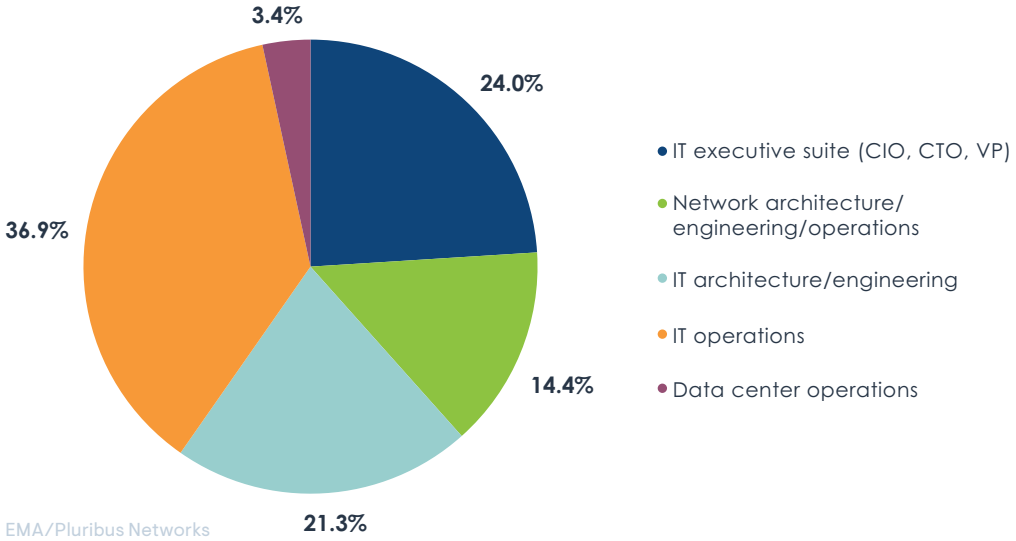

Figure 12. Survey respondents by IT group

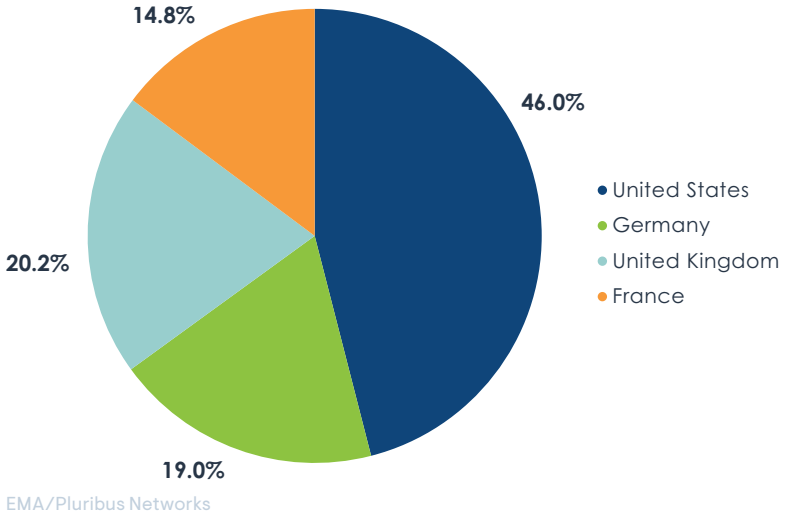

Figure 13. Survey respondents by location

Sample Size = 412

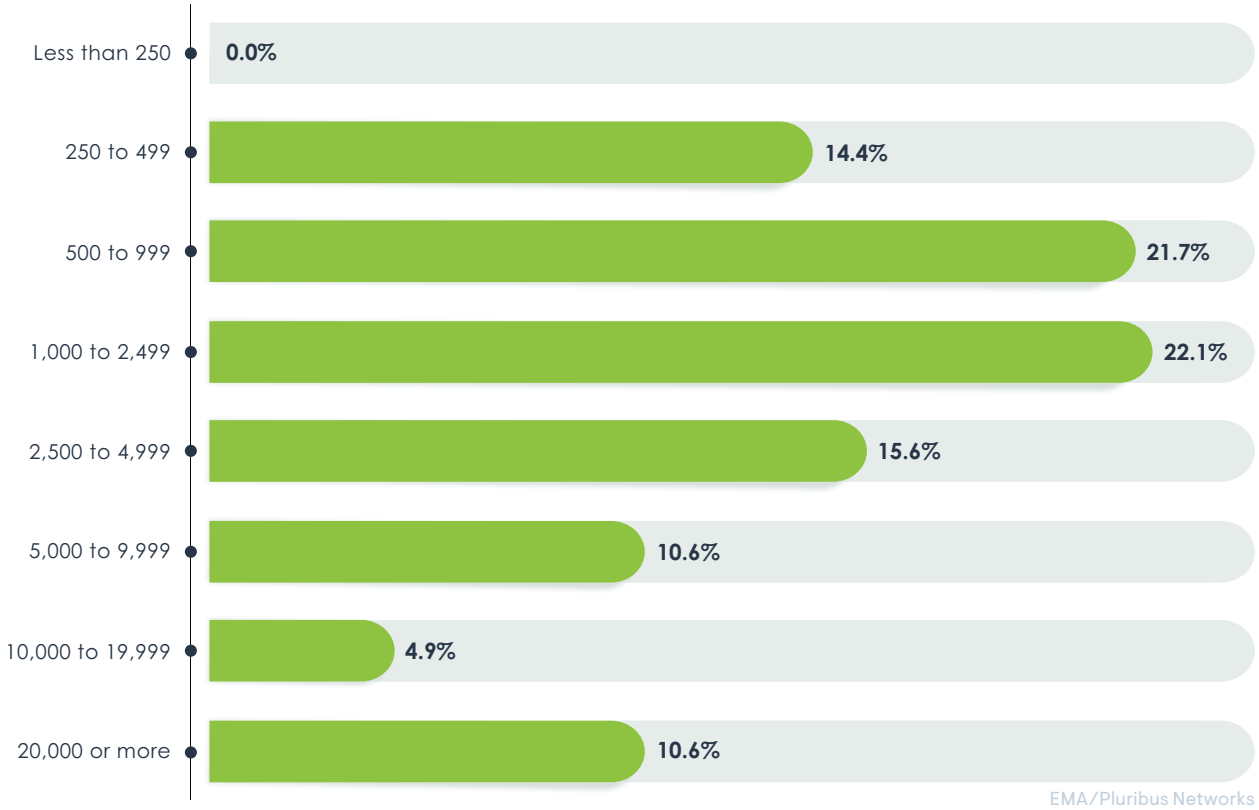

Figure 14. Company size by employees

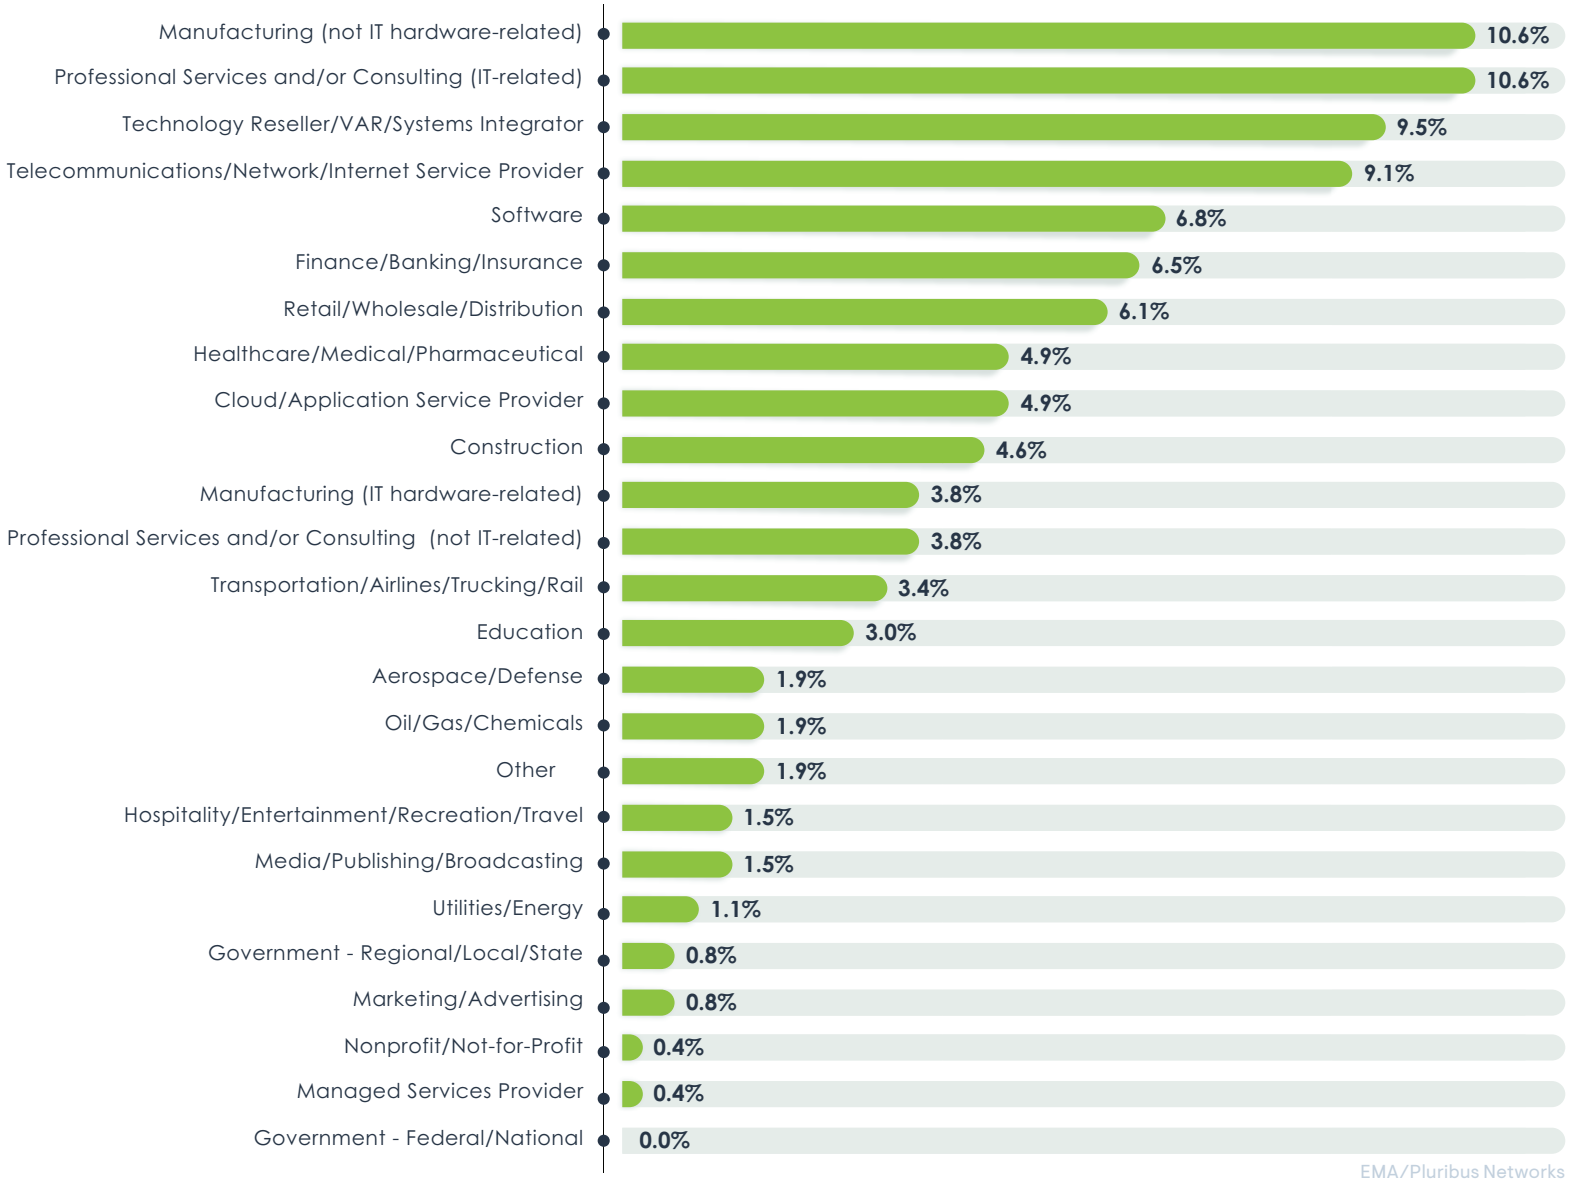

Figure 15. Industries

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals, and IT vendors at **www.enterprisemanagement.com**. You can also follow EMA on **Twitter** or **LinkedIn**.

1995 North 57th Court, Suite 120, Boulder, CO 80301 +1 303.543.9500 **www.enterprisemanagement.com**

4096.051921