Application Delivery Infrastructure for Multi-Cloud Enterprises

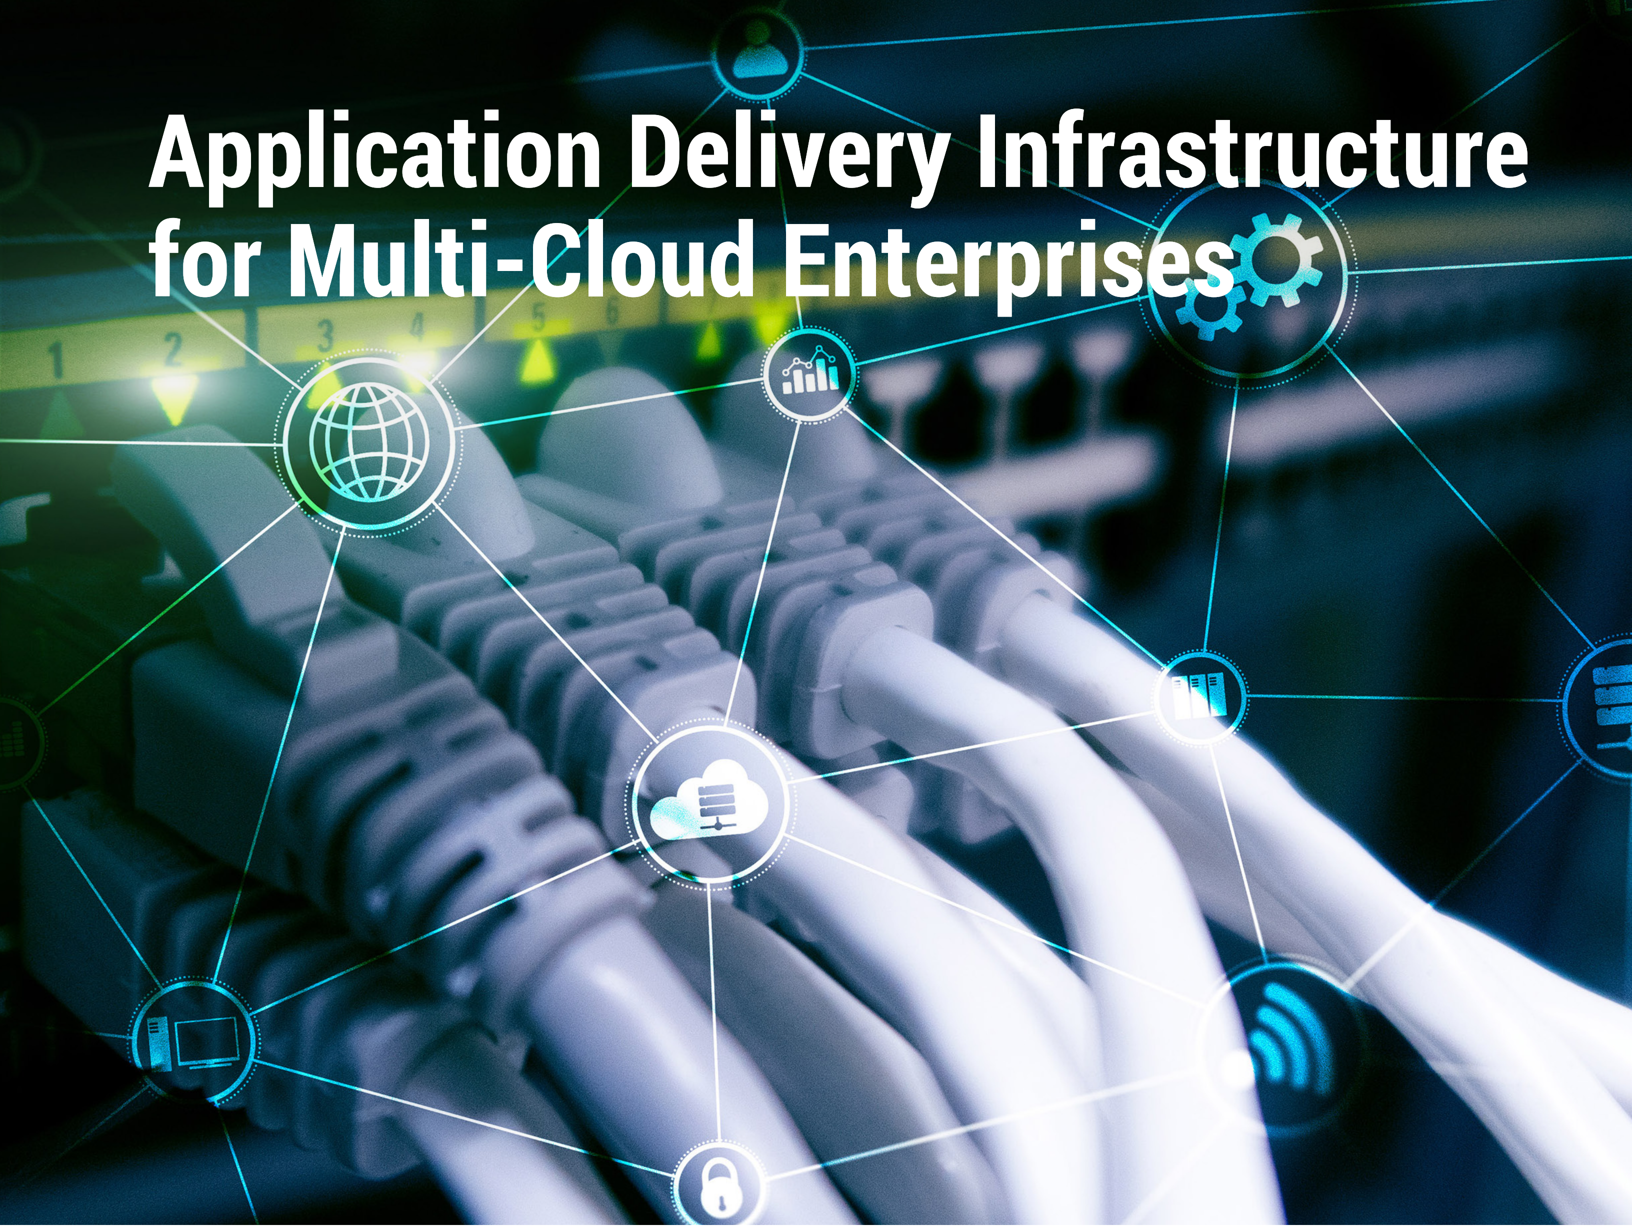

An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) Research Report

By Shamus McGillicuddy

June 2020

Sponsored by:

kemp

NGINX®
Part of F5

Pulse Secure

EMA™
IT & DATA MANAGEMENT RESEARCH,
INDUSTRY ANALYSIS & CONSULTING

Table of Contents

Executive Summary

This Enterprise Management Associates research report explores the state of the art of application delivery infrastructure. Based on a survey of 253 subject matter experts who work for enterprises and cloud providers, the report identifies emerging strategies and technology requirements for load balancers, application delivery controllers, and cloud-native software, such as service mesh.

Introduction

Since the early days of the internet era, application delivery infrastructure has been critical to ensuring application availability, performance, and security. The first generation of devices was comprised of load balancers with a narrow set of functional capabilities focused on distributing user traffic across multiple servers. As applications became more critical to business operations, load balancers evolved into application delivery controllers with an ever-expanding set of functions, from application acceleration and protocol optimization to SSL offload and web application firewalling.

Innovation in application architectures has driven the application delivery infrastructure industry to continuously reinvent itself. For a time, application delivery vendors focused on building massively scalable hardware platforms that could serve the needs of multiple applications from a single appliance.

The cloud computing era shifted application architecture from the monolithic client server model to three-tiered web and mobile applications, and more recently cloud-native, microservices-based applications. These new architectures changed the nature of traffic. So-called “east-west” traffic between application components outstripped the volume of north-south traffic between users and applications. This forced the application delivery infrastructure to shift toward a software-based, dynamic architecture, in which application delivery services are deployed where they are needed, rather than just in a DMZ in front of a monolithic server farm.

Currently, enterprises take a hybrid approach to application delivery infrastructure, with a mix of appliances, software, and services supporting a variety of applications in their data centers and public cloud estates. To some extent the industry has fragmented, as the DevOps teams that own cloud infrastructure work with their preferred application delivery platforms and the data center operations teams that own traditional infrastructure stick with their incumbent platforms. This evolution prompted traditional data center infrastructure vendors, like F5 Networks and VMware, to acquire cloud-native software vendors, like Nginx and Avi Networks. It also sparked a wave of innovation around analytics, automation, and platform design from incumbent vendors, such as the aforementioned F5, Kemp, Pulse Secure, Citrix, A10 Networks, and others.

This Enterprise Management Associates (EMA) research report explores the current state of application delivery infrastructure strategy. Based on a survey of 253 technology professionals, this report identifies the technical requirements, management strategies, and infrastructure roadmaps of cloud-forward enterprises.

Key Findings

- 81% of enterprises have experienced fragmentation in operational and administrative ownership of application delivery infrastructure, leading to increased security risk, compliance challenges, and operational inefficiency.
- Enterprises identified cloud-native software as the application delivery platform of the future, both in private data centers and the public cloud.
- 99% of enterprises consider application delivery infrastructure to be part of their overall security architecture, although the role it plays varies from enterprise to enterprise.
- 82% of enterprises are collecting telemetry from application delivery infrastructure, primarily for security monitoring, capacity planning, and application troubleshooting.
- 95% of enterprises are interested in using AIOps solutions from their application delivery vendors.
- 71% of enterprises say some aspects of their application delivery infrastructure are too difficult and time-consuming to manage.
- The vast majority of enterprises automate application delivery infrastructure management in both the data center and the cloud, but only a minority are satisfied with their ability to automate.
- 85% of enterprises are interested in using service mesh in their cloud-native application environments, and nearly all of them would like some management integration between service mesh and the rest of their application delivery infrastructure.
- 90% of enterprises have made changes to their application delivery infrastructure in response to the COVID-19 pandemic.

Demographics Overview

In April and May 2020, EMA surveyed 253 North American technology professionals who have direct and current experience with their employers' use of application delivery infrastructure. All of them work for companies with application delivery infrastructure deployed in their data centers and the public cloud.

These research participants represent a variety of different perspectives inside technology organizations, as **Figure 1** reveals. Cloud-oriented groups like DevOps and cloud engineering/operations combine to represent more than one-quarter of the survey. Traditional infrastructure teams, such as network operations, infrastructure architecture/engineering, and data center operations, combine to represent nearly 40% of the survey. There is also strong representation from application development and information security teams.

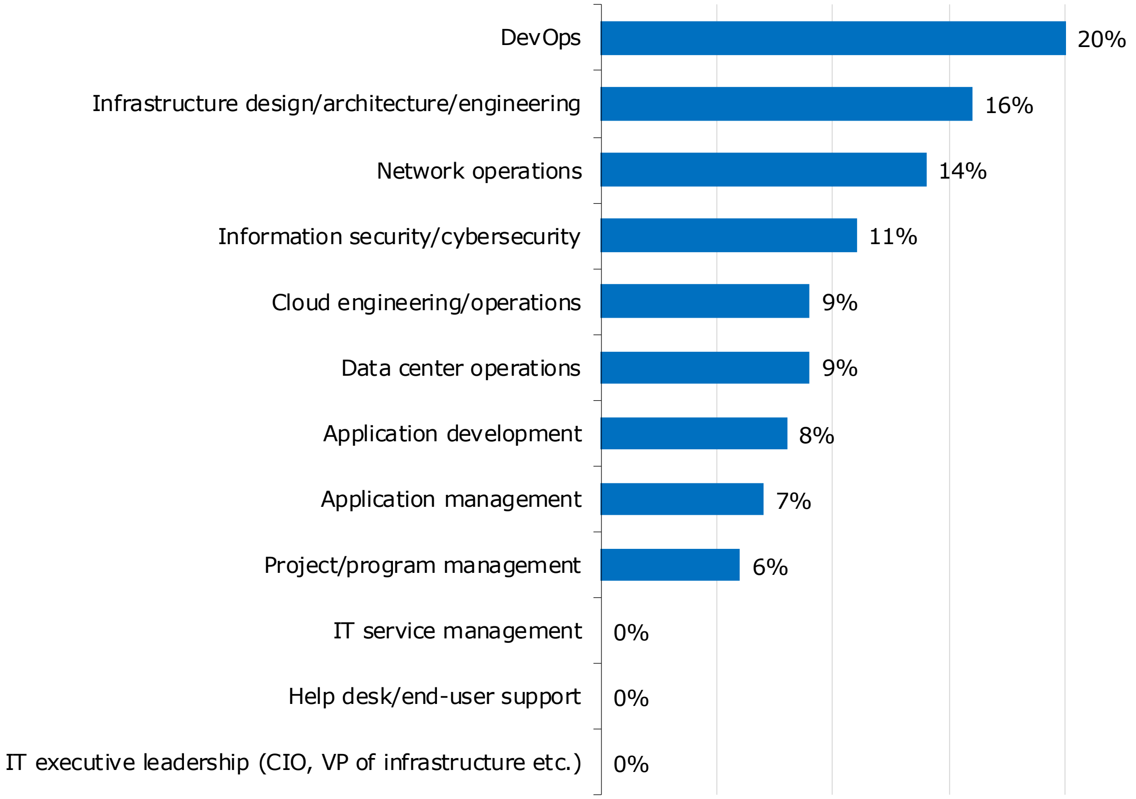

Sample Size = 253

Figure 1. Functional groups represented in the research

Figure 2 reveals that the majority of the companies represented in this survey are large enterprises, with a small contingent of very large enterprises (10,000 or more employees). More than one-third of companies fall into the midsized category (250 to 999 employees). EMA excluded any company with fewer than 250 people.

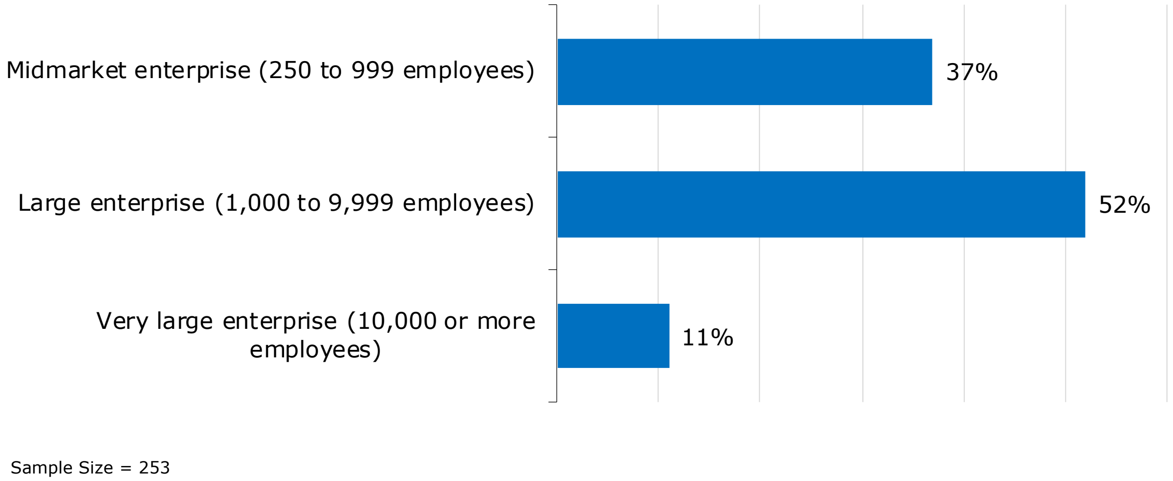

Figure 2. Size of company (employees)

Figure 3 reveals the annual revenue of the enterprises represented in the survey. EMA excluded companies with less than $5 million in revenue. A significant minority of these organizations are billion-dollar companies, but the most common revenue range is $100 million to less than $500 million.

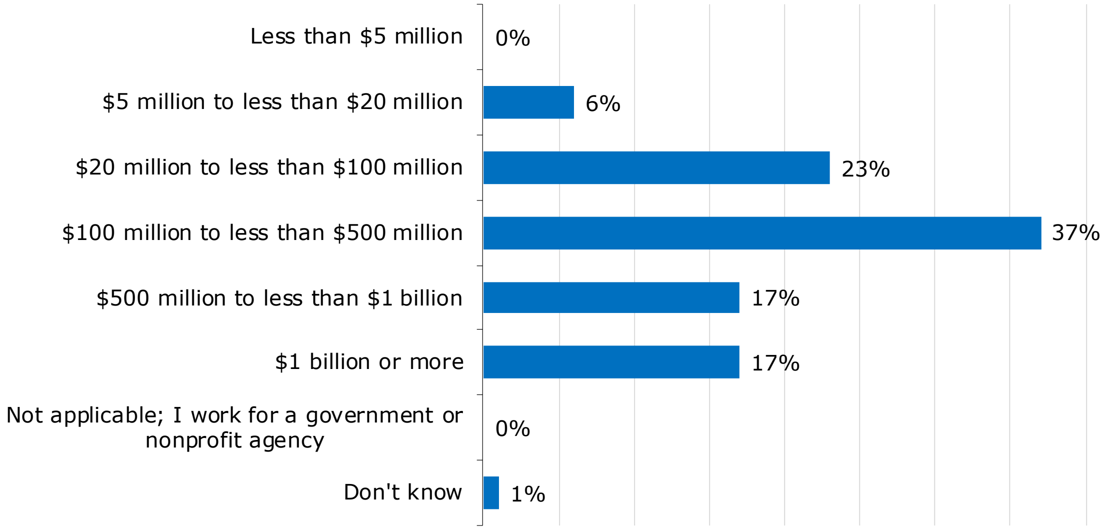

Figure 3. Revenue of represented companies

At least 22 industries are represented, as detailed in **Figure 4.** The most common industries are software vendors and cloud and application service providers. IT-related professional services firms, manufacturers, retail, and financial services are also well-represented. EMA excluded communications service providers and technology sales channel partners from the survey.

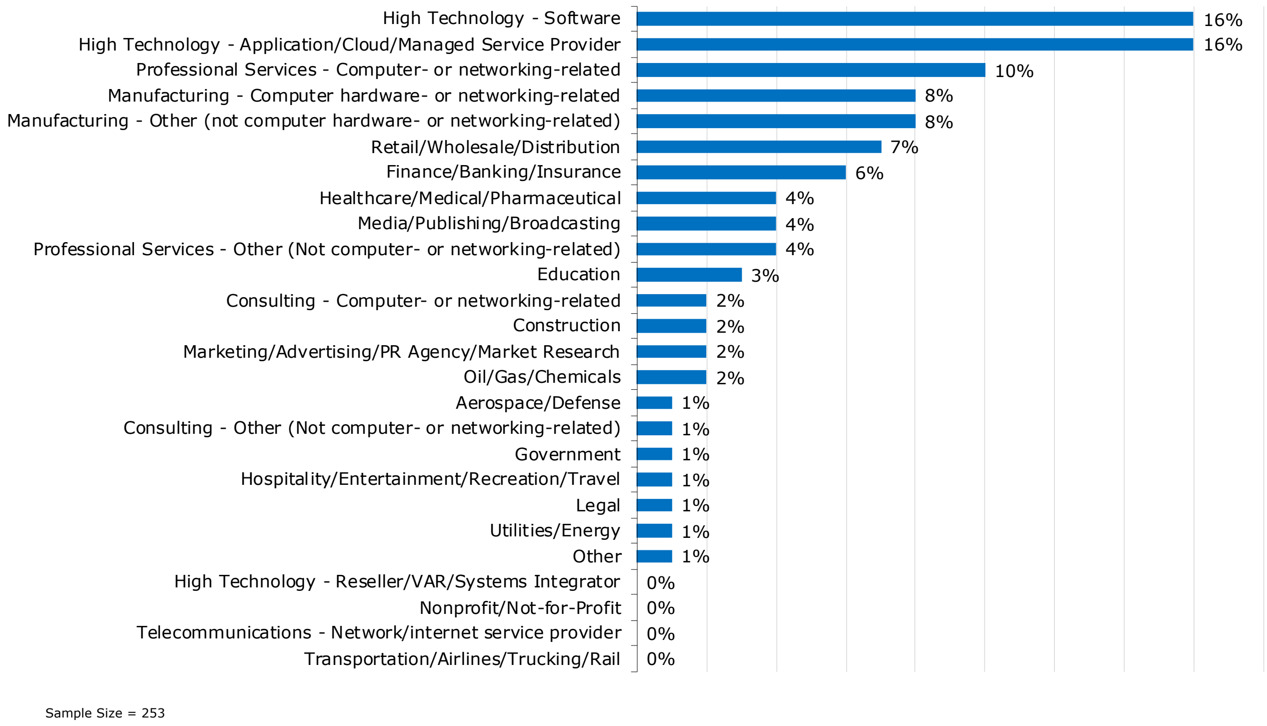

Figure 4. Industries represented in the research

Strategic Drivers

Application Delivery Infrastructure Budgets

EMA polled enterprises about their application delivery infrastructure budgets for context. This survey was conducted well after the onset of the global COVID-19 pandemic, so it reflects a current perspective on budgeting. However, the economic disruption caused by the virus obviously makes this a fluid situation. It's possible that budgets will change through the course of the year.

Figure 5 reveals that the typical annual application delivery infrastructure budget captured in this survey ranges between $200,000 and $1 million, with most falling between $250,000 and $1 million. None of the respondents claimed to have a budget that was less than $50,000, which is reflective of the fact that EMA targeted midsized to very large enterprises.

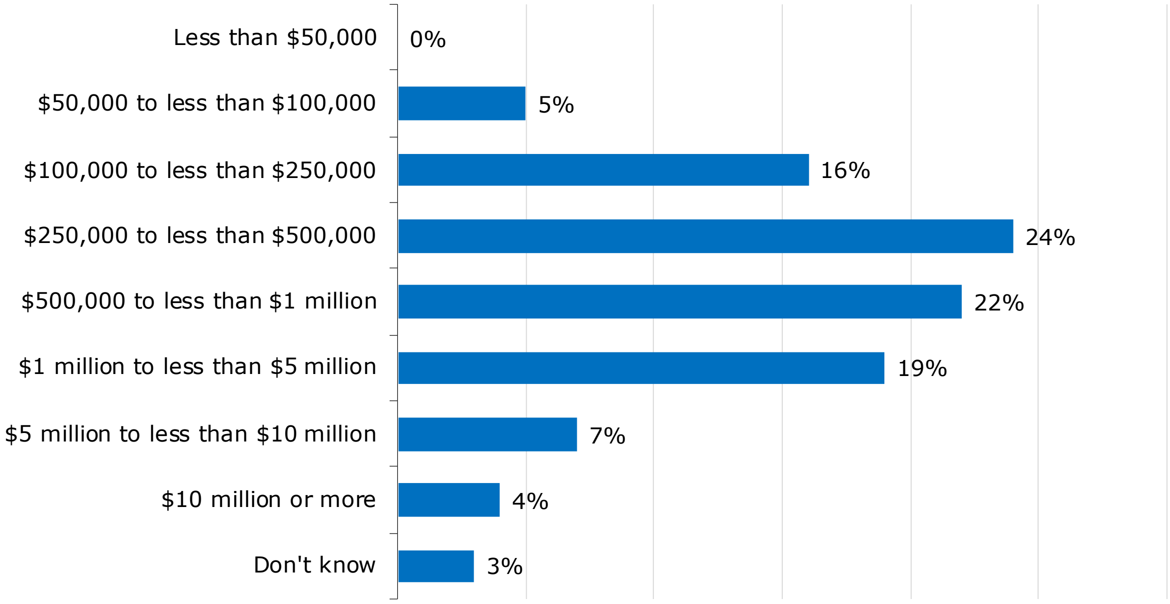

Sample Size = 253

Figure 5. Application delivery infrastructure budgets

Figure 6 tracked budget growth from last year to this year. It shows that 80% of application delivery infrastructure budgets increased this year, typically by 10% or less. Only 3% reported a budget decrease, but a sizable number saw their budget remain flat this year.

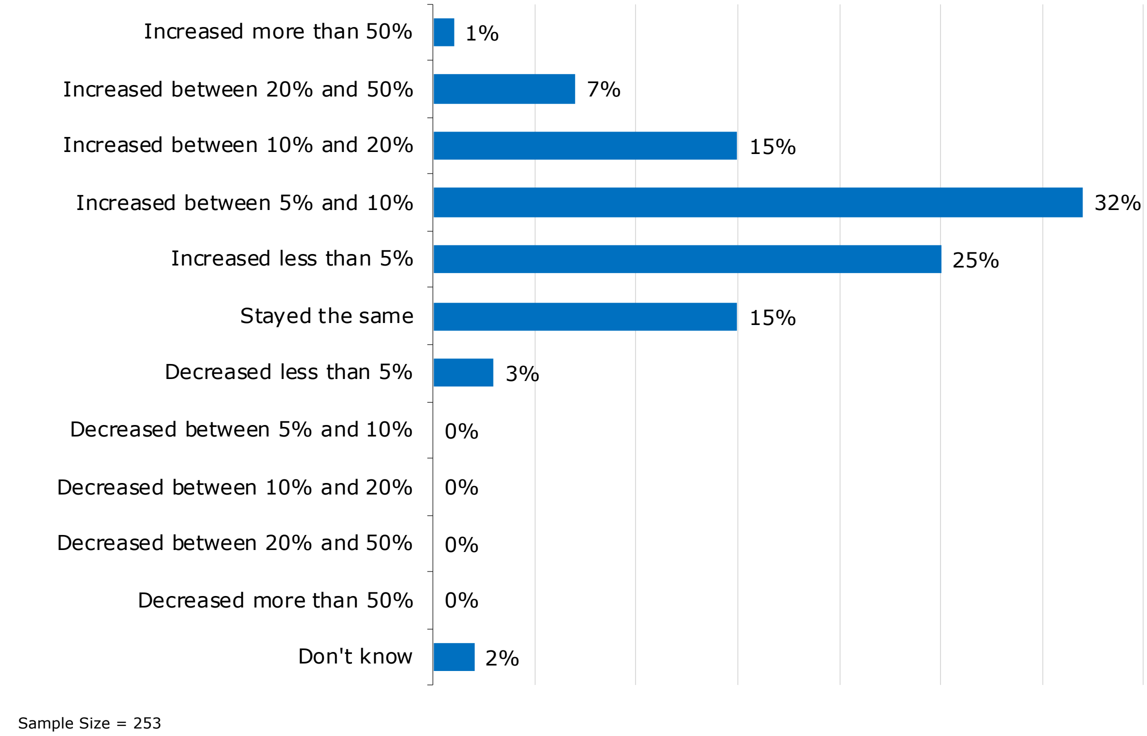

Figure 6. Application delivery infrastructure budget growth

Technical Initiatives

EMA asked respondents to identify up to three technical initiatives that are significantly influencing their application delivery infrastructure strategies. The public cloud is the biggest driver, followed by cloud-native application modernization efforts and hybrid cloud architectures. Data center SDN technology and infrastructure automation are also significant influences.

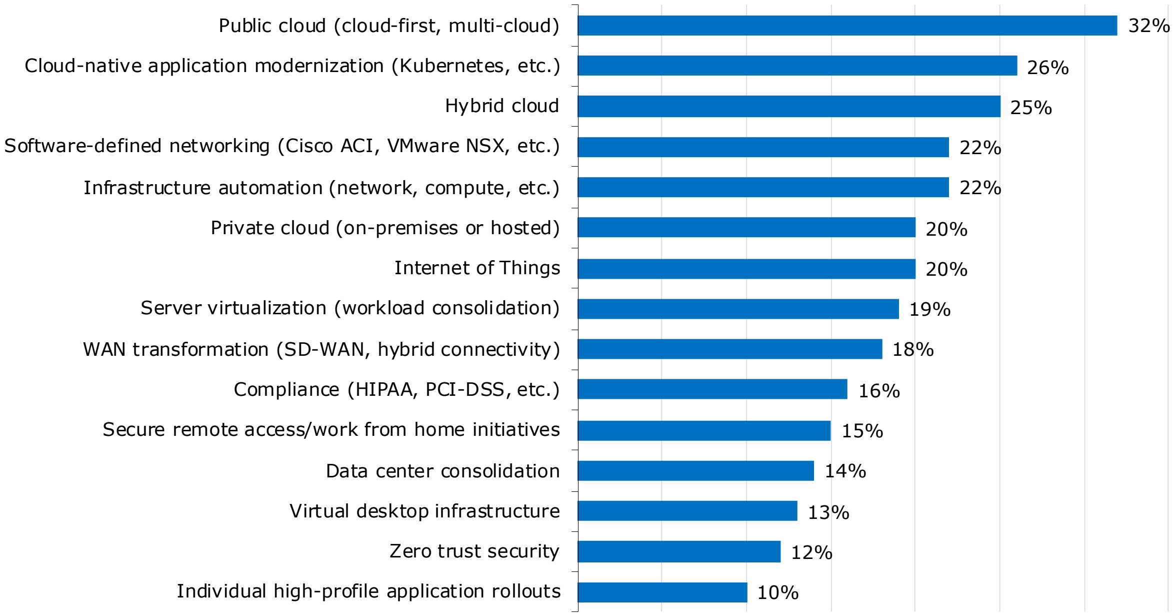

Sample Size = 253, Valid Cases = 253, Total Mentions = 717

Figure 7. Technology initiatives most influential on application delivery infrastructure strategies

Information security professionals were the most likely to cite hybrid cloud as a strategic driver, and network operations professionals were less likely. Cloud engineering/operations are more likely to see private cloud architecture and cloud-native applications as major drivers. Unexpectedly, DevOps professionals were less likely to see cloud-native applications as a driver. Application management teams perceive remote access/work from home initiatives as a major driver.

Enterprises that are successful with their application delivery infrastructure strategy are more likely to cite data center SDN as a driver. Meanwhile, less successful organizations focus more on compliance and zero trust security.

Overall, it's quite clear that cloud architectures, whether public, private, or a mix of the two, are driving application delivery infrastructure strategies.

Cloud architectures, whether public, private, or a mix of the two, are driving application delivery infrastructure strategies.

Application Architecture

Figure 8 examines which application architectures are steering application delivery infrastructure strategy. Clearly, architectures of the past, such as mainframe and client/server, are waning. Instead, the majority of respondents see cloud-native applications and microservices as the guiding architecture. A large minority are more influenced by three-tier web and mobile applications.

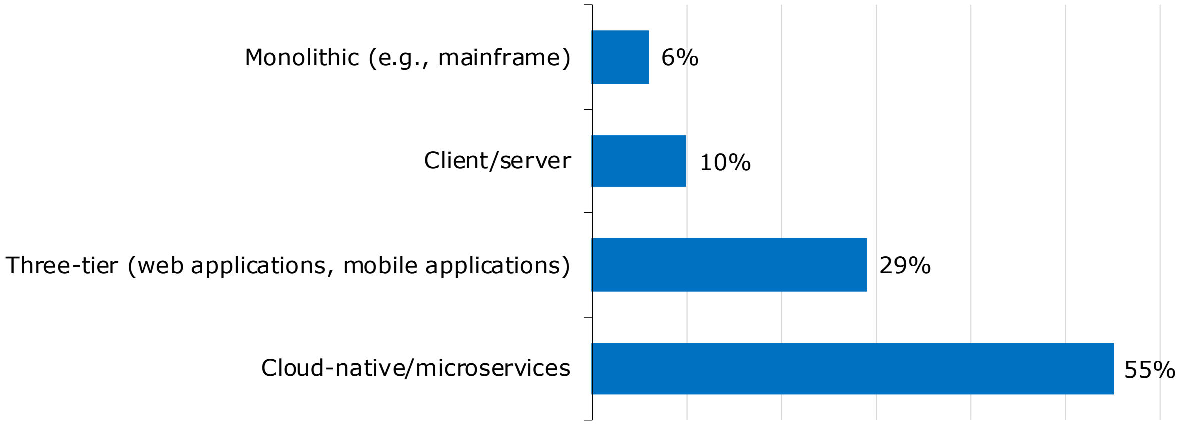

Sample Size = 253

Figure 8. Application architectures most influential on application delivery infrastructure strategy

DevOps and cloud engineering/operations teams are the most likely to see cloud-native as the dominant architecture. Midmarket enterprises also displayed a preference for cloud-native architectures.

These findings confirm that most technology organizations are moving into the future, orienting their application delivery platforms for cloud-native application architectures. Most enterprises will continue to maintain legacy applications, but cloud-native is the future and will drive most infrastructure strategies from now on.

The NetOps/DevOps Divide

IT organizations traditionally have well-defined siloes of responsibility, particularly when it comes to infrastructure. Load balancers and application delivery controllers are network devices. Therefore, IT engineering, network operations, and data center operations are typically responsible for installing and managing them.

81%
of enterprises are dealing with divided ownership of application delivery infrastructure.

As applications have migrated to private and public clouds, enterprises have shifted to using application delivery infrastructure software and cloud-based services. This shift has fragmented ownership of this critical technology across multiple administrative silos. The network infrastructure team still owns a good portion of infrastructure in data centers. However, cloud and DevOps teams have taken direct ownership of infrastructure in cloud environments, leading to divided operations.

Figure 9 shows that 81% of enterprises are dealing with divided ownership of application delivery infrastructure currently. The application development team (60%) is the least likely to perceive this issue, probably because they concern themselves with development and test environments more than production infrastructure. However, DevOps, application management, and information security teams all displayed higher awareness of this issue.

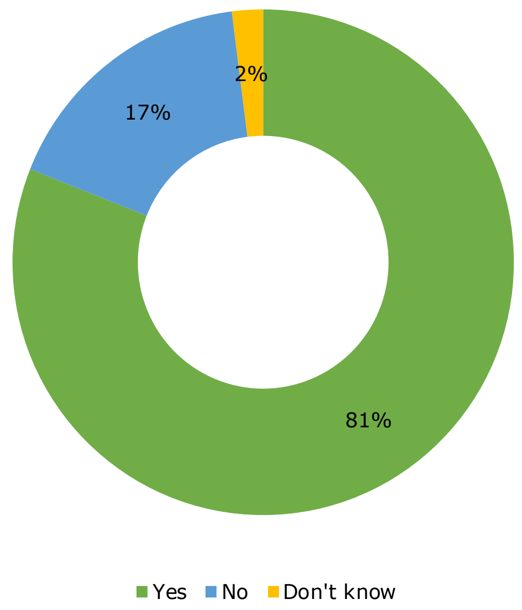

Sample Size = 253

Figure 9. 81% of companies have divided ownership between traditional application delivery infrastructure and cloud-native and public cloud infrastructure

Organizations that are successful with application delivery infrastructure are less likely have a division of ownership, but even they are unable to completely avoid the issue.

Most enterprises see the potential problems of this operational fragmentation. In fact, 85% have taken steps to close the gap and unify management. Midmarket enterprises (92%) are more aggressive with closing this gap, while very large enterprises are lagging (61%).

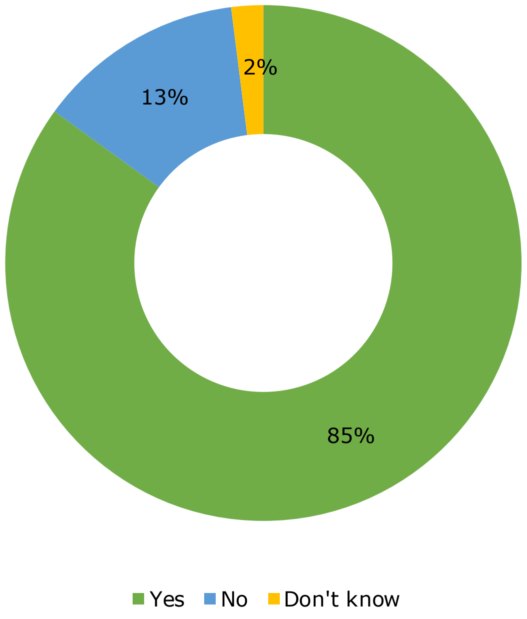

Sample Size = 204

Figure 10. 85% of companies with divided ownership of infrastructure have taken steps to close these divisions

EMA asked research respondents why they felt a need to close this gap in infrastructure ownership. Security risk is the top motivation, as **Figure 11** reveals. Without a unified approach to this infrastructure, operations often have an inconsistent approach to implementing security controls, such as a web application firewall. They take an inconsistent approach to encrypting application data. Cloud engineering and operations teams are especially likely to perceive security risk as a problem. Successful organizations are also more likely to recognize security risk as a problem.

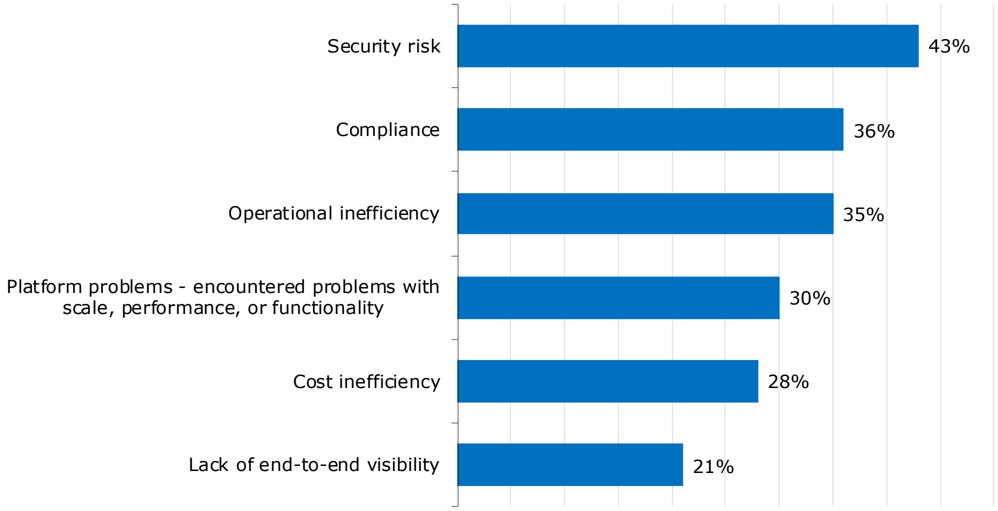

Sample Size = 173, Valid Cases = 173, Total Mentions = 335

Figure 11. Why enterprises want to close gaps in ownership of application delivery infrastructure

Compliance and operational inefficiency are the chief secondary drivers of unifying infrastructure ownership. Some are also perceiving potential platform issues. For instance, the DevOps team might choose a load balancer that can't scale properly. The network operations team might choose a platform that cannot properly serve a microservices architecture.

A lack of end-to-end visibility across different infrastructure environments was only a top driver for one in five respondents, making it the least common challenge.

EMA asked research participants if vendors could help unify this division of application delivery infrastructure operations. As **Figure 12** reveals, 100% of respondents said yes, their vendors could do something to help. Majorities of them said vendors could offer more deployment flexibility and better APIs for platform integration. DevOps and network operations teams have strong interest in better APIs. Successful organizations are also more likely to need improved APIs.

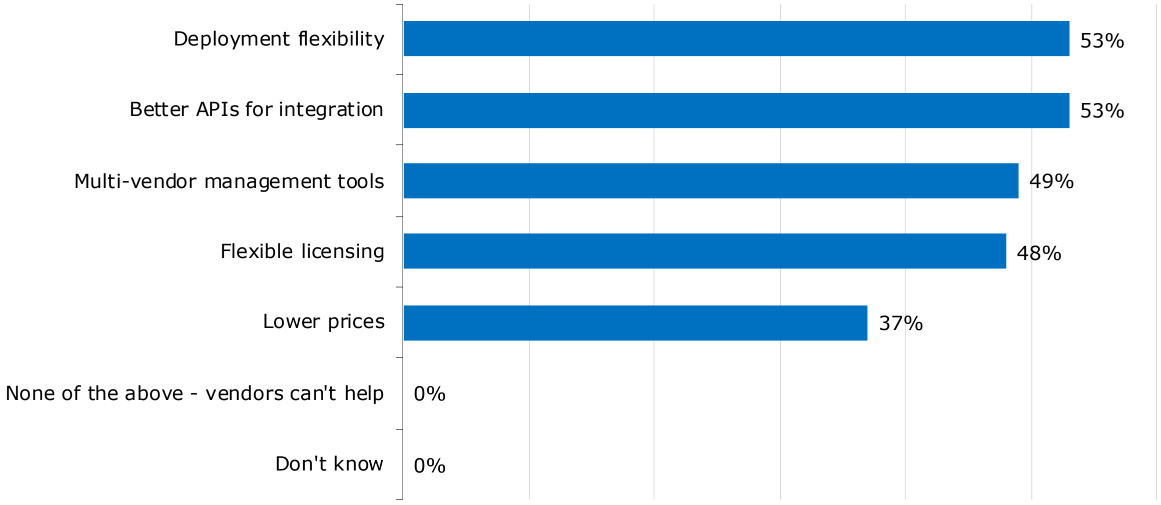

Sample Size = 204, Valid Cases = 204, Total Mentions = 490

Figure 12. Enterprises identify how vendors could help enterprises unify ownership of application delivery infrastructure

Nearly half of enterprises say better multi-vendor management tools and flexible licensing could help them unify infrastructure and operations. DevOps is especially interested in multi-vendor management tools, while infrastructure engineering teams are not. Organizations that are successful with unifying ownership are the most likely to recognize flexible licensing as a way that vendors could help.

EMA found a very strong correlation between enterprises that have moved forward with platform standardization and those that are successful with unifying ownership of application delivery infrastructure.

Standardization on approved application delivery platforms is one way to get everyone on the same page, and 81% of enterprises are trying to close the management gap by standardizing on one or more platforms, as detailed in **Figure 13**. Another 16% say they haven't done this standardization, but they want to do so.

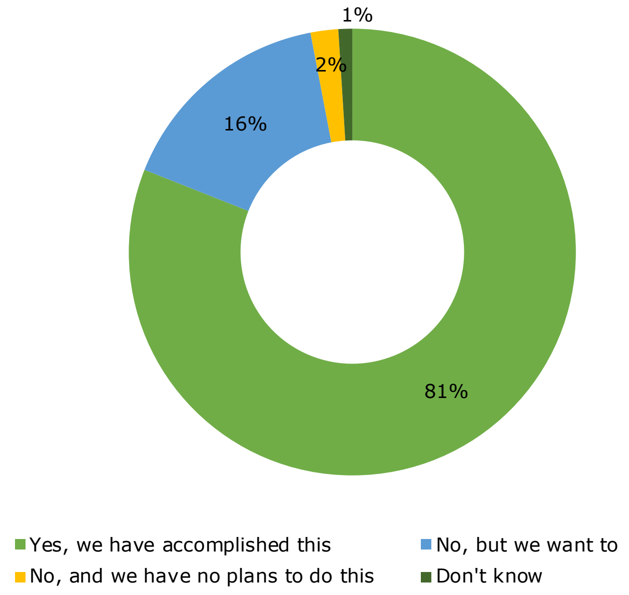

Sample Size = 173

Figure 13. 81% of enterprises that try to close the ownership gap have standardized on approved platforms

EMA found a very strong correlation between enterprises that have moved forward with platform standardization and those that are successful with unifying ownership of application delivery infrastructure.

Application Delivery Infrastructure Platforms: Today and the Future

Data Centers

The majority of enterprises have a wide variety of application delivery platforms installed in their data centers, as revealed in **Figure 14.** Cloud-native software is the most common. This includes emerging technology, such as service mesh software, but it can also include lightweight, open source software solutions favored by DevOps organizations, like HAProxy and NGINX. Application delivery controller software installed on open hardware is also very common. A slight majority of enterprises have virtual application delivery controller appliances running on virtual infrastructure. Finally, slightly less than half are using vertically integrated hardware appliances. A decade ago, this last percentage would have been much higher. It's quite clear that the market has retreated from the high-performance hardware of the client-server era.

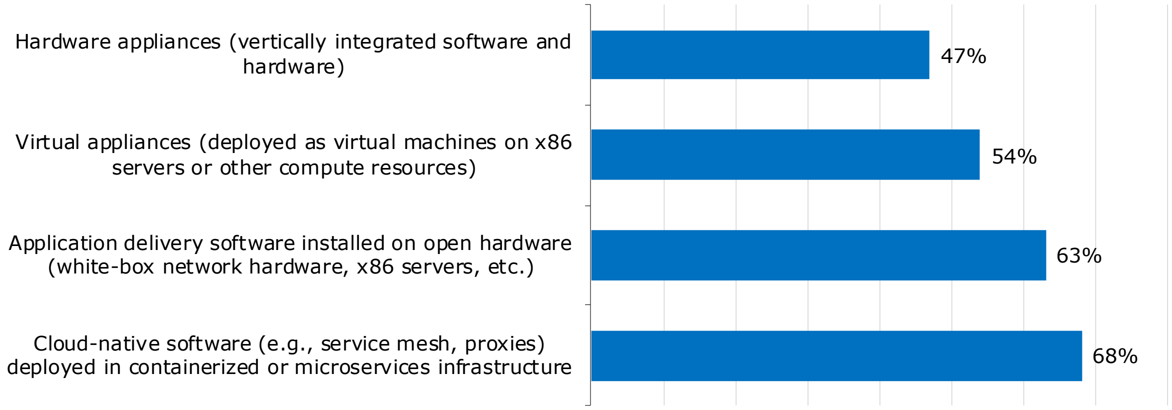

Sample Size = 253, Valid Cases = 253, Total Mentions = 589

Figure 14. Application delivery platforms installed in the data center

EMA found a strong correlation between successful application delivery strategies and installed cloud-native application delivery infrastructure in data centers. EMA suspects that enterprises with a successful approach to infrastructure are more equipped to be aggressive with emerging technology, allowing them to test and adopt cloud-native solutions more quickly.

EMA asked enterprises to project two years out and identify the class of platform that would be most important in their data centers at that time. **Figure 15** reveals that cloud-native platforms are the technology of the future. Only 11% believe hardware appliances will dominate their data centers in two years. It's notable that no enterprises expect to have fully migrated all their applications to the public cloud by 2022. The data center will remain relevant.

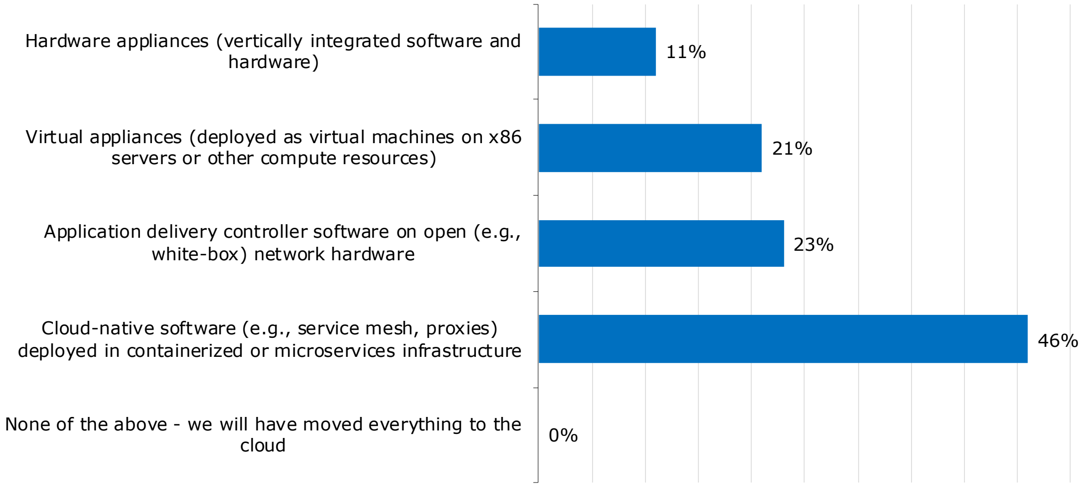

Sample Size = 253

Figure 15. Enterprises identify the platform expected to be most important to their data centers in two years

DevOps and cloud engineering/operations are much more likely to identify cloud-native platforms as their core platform of the future. Infrastructure engineering/architecture and application management are the least enthusiastic about cloud-native platforms. Enterprises that are successful with application delivery infrastructure are more likely to view cloud-native software as the future of application delivery in their data centers.

Research respondents who primarily engage with application delivery infrastructure as managers of budget and product procurement are more likely to believe that application delivery software installed on open hardware is the most important deployment model moving forward. Meanwhile, individuals who deploy and manage this infrastructure or develop and manage the applications that run on that infrastructure tend to view cloud-native software as the future of infrastructure.

Public Clouds

In the public cloud, most enterprises are using application delivery services offered by cloud providers. Smaller majorities are also using cloud-native software and virtual application delivery controllers. Hardware peered to the cloud (e.g., from a colocation provider) is the least common. Successful enterprises are more likely to be using solutions offered by their cloud providers.

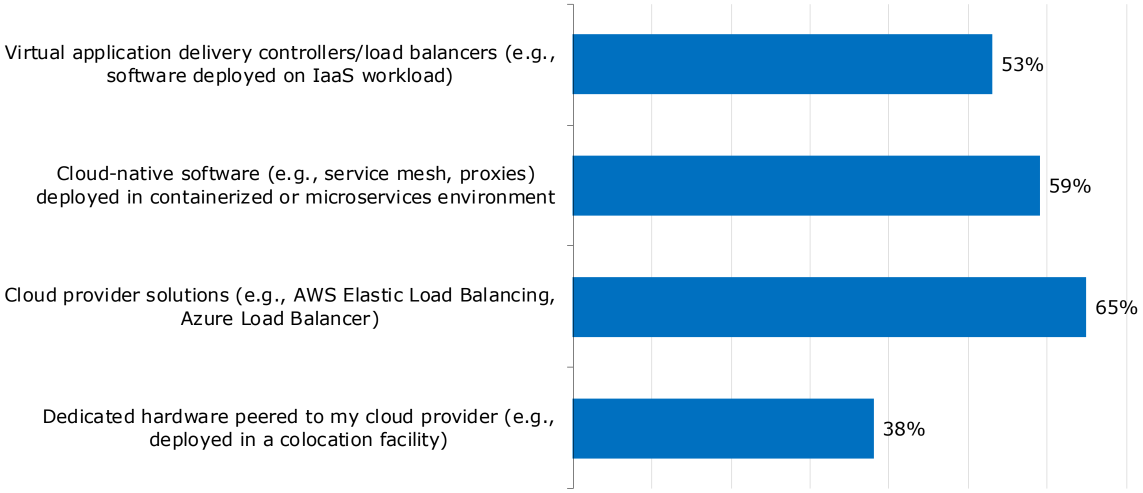

Sample Size = 253, Valid Cases = 253, Total Mentions = 544

Figure 16. Application delivery platforms installed in the public cloud

EMA asked research participants to identify the one type of platform they expect to be most important to their public cloud infrastructure in two years. Application delivery services offered by cloud providers are the most popular. Many are looking at cloud-native software and virtual appliances. Peered hardware is definitely not the future.

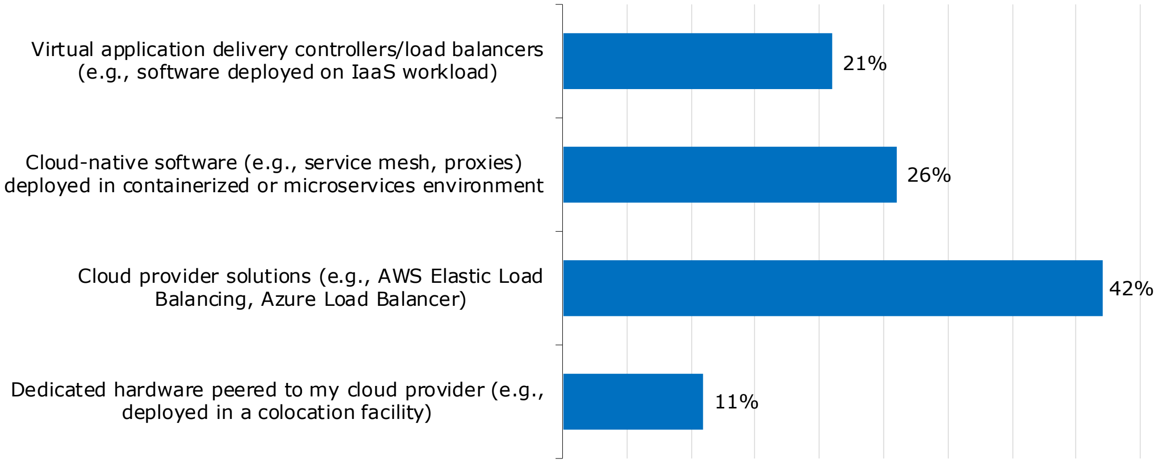

Sample Size = 253

Figure 17. Enterprises identify the platform expected to be most important to their public cloud infrastructure in two years

DevOps professionals are more likely to consider cloud provider solutions as essential to the future. Application developers are more likely to think cloud-native software is the future.

Application Delivery Infrastructure Requirements

Core Application Delivery Functions

This section examines the application delivery functions that are most important to enterprises at the time of this report. As **Figure 18** shows, in the data center global load balancing, secure application gateways, data compression, and API gateways are the most valuable capabilities.

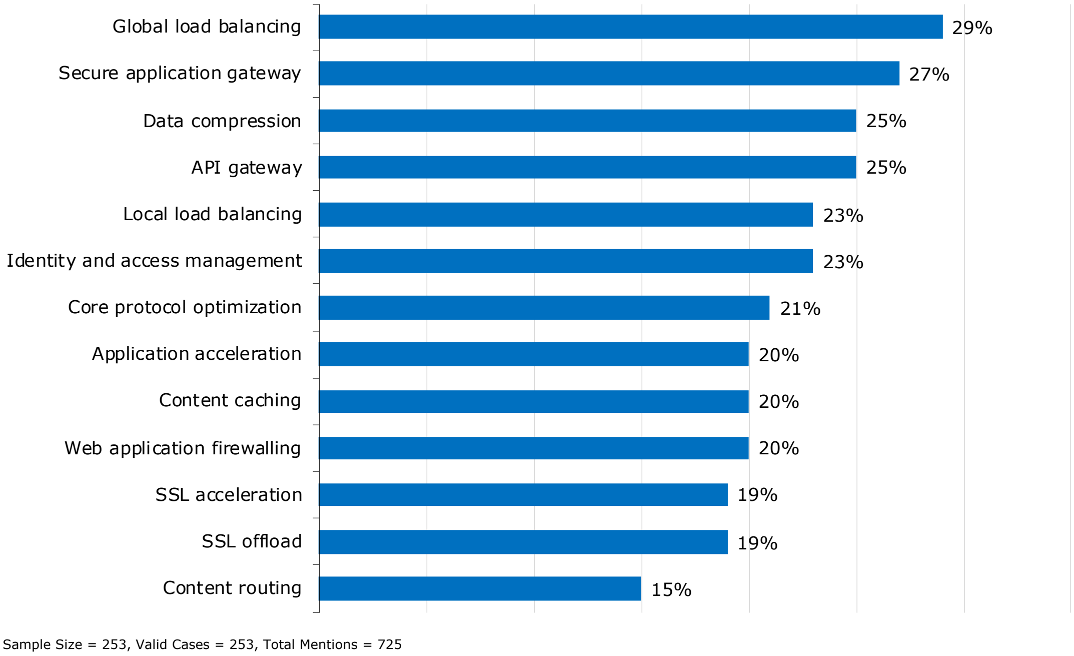

Figure 18. Most important application delivery functions in data center deployments

Local load balancing and identity and access management are also valuable. Local load balancing is particularly important among cloud engineering and operations teams, but not as important within infrastructure architecture and engineering teams. Very large enterprises are also more likely to prize local load balancing. DevOps teams are more interested in API gateways, but infrastructure architecture and engineering teams are not.

Other functions like core protocol optimization, application acceleration, and content caching are lower priorities, probably depending on what kinds of applications enterprises consider strategic to their business. DevOps professionals are less likely to value application acceleration, but application management and project management do prize it. Application developers want content caching, but DevOps are unlikely to.

SSL acceleration, SSL offload, and content routing are the least popular functions. Successful application delivery teams are more likely to recognize the value of application acceleration. Less successful organizations place more value on SSL offload. DevOps is very unlikely to value SSL acceleration, but cloud engineering and operations tend to see more value in it. Information security teams are the most likely to value SSL offload, while DevOps is least likely.

> *Successful application delivery teams are more likely to recognize the value of application acceleration.*

Enterprises have some different priorities in their public cloud infrastructure. Global load balancing remains paramount, and data compression also remains a top-three priority, as **Figure 19** reveals. However, local load balancing has much more value in the public cloud. Global load balancing is less popular among DevOps, but more so among data center operations.

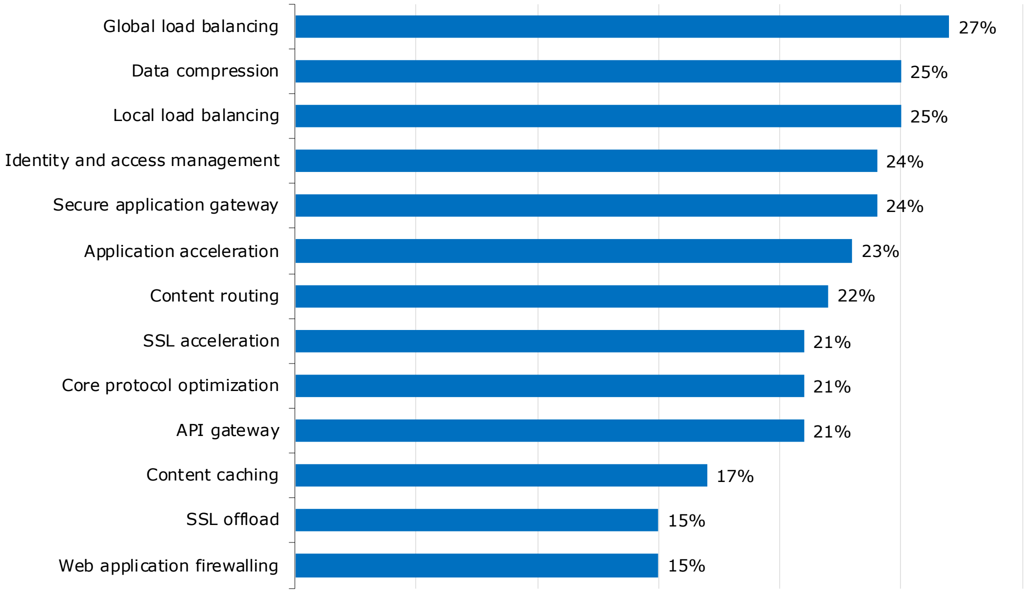

Sample Size = 253, Valid Cases = 253, Total Mentions = 708

Figure 19. Most important application delivery functions in public cloud

Identity and access management and secure application gateways round out the top five, the latter being less important in the public cloud than it is in the data center. Content routing is more important in the public cloud, but web application firewalls are a lower priority. Identity and access management is of particular importance to very large enterprises.

As they do in data centers, successful application delivery teams are more likely to value application acceleration in the cloud. However, less successful organizations place more emphasis on local load balancing and SSL offload.

Platform Requirements

EMA explored platform requirements for both data centers and public cloud environments. In the data center, enterprises have seven priorities, starting with multi-cloud support. In other words, enterprises want the application delivery infrastructure that they deploy in their data center to be deployable in their various cloud environments, too. This points to the need for a standard platform across public and private infrastructure. Very large enterprises are the most likely to value multi-cloud support.

Figure 20 shows that enterprises identified several secondary platform requirements, including multi-node clustering, throughput scalability, resource efficiency, resiliency, observability, and SSL/TLS transaction capacity. The lowest priorities are programmability (both data plane and control plane) and single-node scalability.

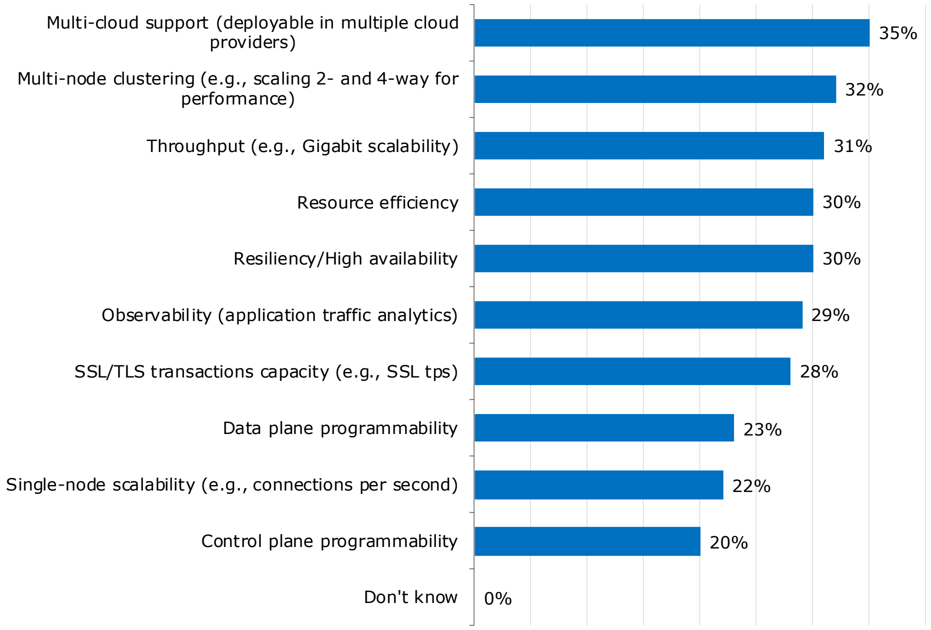

Sample Size = 253, Valid Cases = 253, Total Mentions = 710

Figure 20. Platform requirements in the data center

Resiliency is a high priority for network operations, cloud engineering/operations, and project management, but a lower priority for infrastructure engineering and DevOps.

In the public cloud, multi-cloud support is the most important platform requirement. This points to the fact that many enterprises want to standardize their application delivery infrastructure across multiple cloud providers.

Figure 21 shows that resiliency stands out as the chief secondary platform requirement. Resource efficiency, observability, and throughput round out the top five. Midmarket enterprises are the most likely to prize resource efficiency. Data plane programmability and single-node scalability remain low priorities, as they did in the data center. However, control plane programmability is a bit more important in the public cloud.

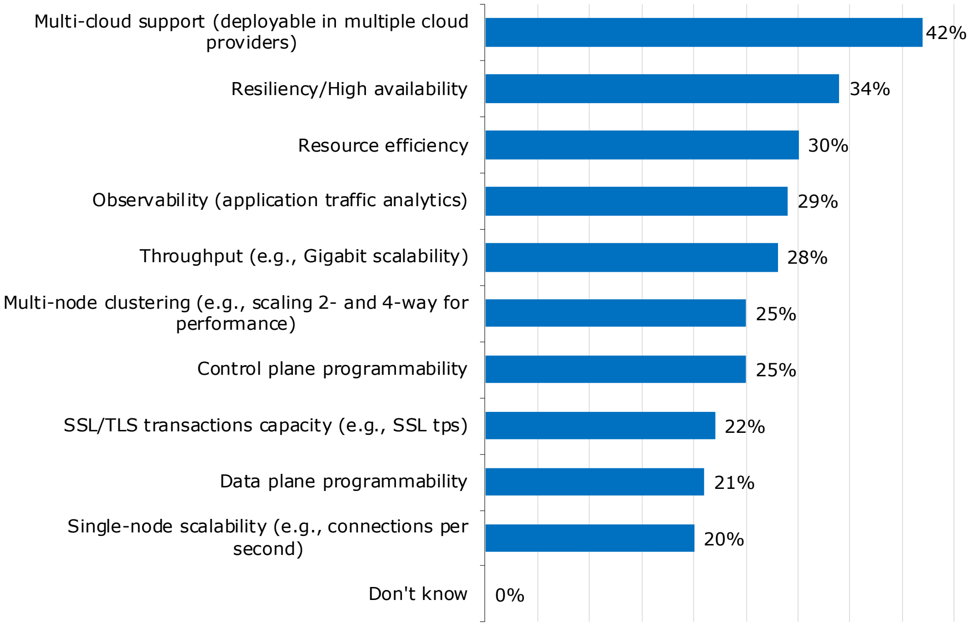

Sample Size = 253, Valid Cases = 253, Total Mentions = 702

Figure 21. Platform requirements in the public cloud

Network operations and DevOps have different points of view about platform scalability. Network operations teams are more likely to focus on single-node scalability, while DevOps is more focused on multi-node clustering. Multi-cloud support is especially important to cloud engineering and operations, but a lower priority to application development and infrastructure architecture and engineering.

Support of Security Architectures

While application delivery infrastructure is essential to application performance and availability, it is also a critical component of a security architecture. Enterprises should keep this in mind as they define their technical requirements. In fact, **Figure 22** reveals that 99% of enterprises believe that application delivery controllers and load balancers are part of their security architecture, including 59% who say this infrastructure a critical part of their overall security stack.

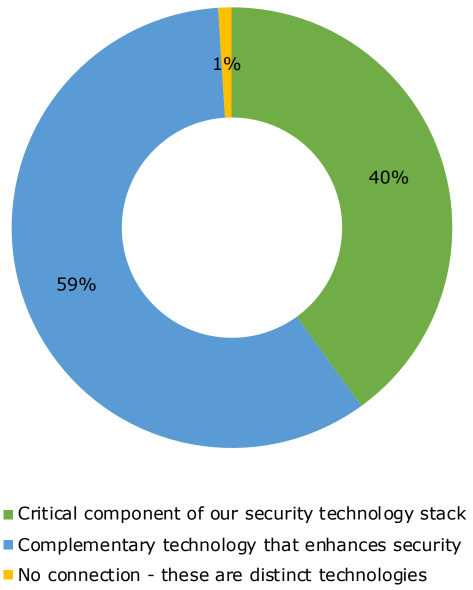

Sample Size = 253

Figure 22. The role that application delivery infrastructure plays in security architecture

> *99% of enterprises believe that application delivery controllers and load balancers are part of their security architecture, including 59% who say this infrastructure is a critical part of their overall security stack.*

The most common role these solutions play is as an enhancement to policy enforcement, as **Figure 23** indicates. Enterprises are enforcing some policies at the application delivery infrastructure layer rather than at the traditional security stack. This can close gaps in policy enforcement and save resources in the security stack.

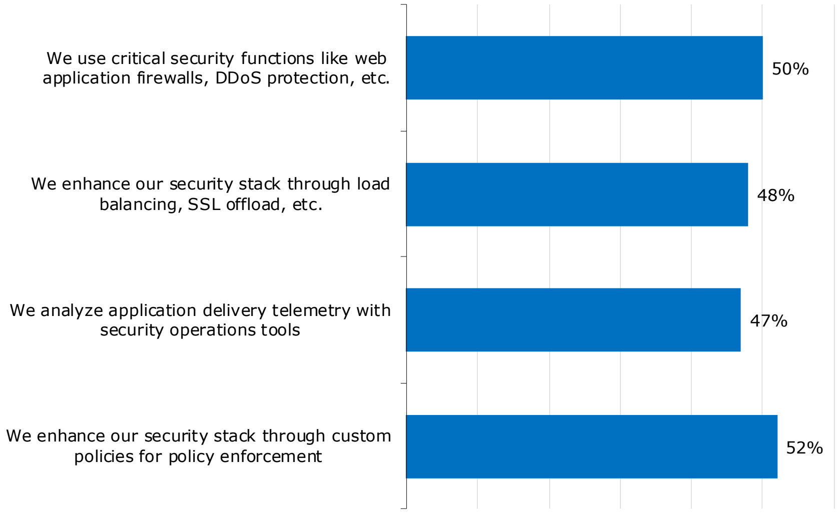

Figure 23. How application delivery infrastructure supports a security architecture

Half of enterprises enhance their security architecture by using critical security functions in their application delivery infrastructure, such as web application firewalls and DDoS protection. Slightly less than half rely on key features, such as load balancing and SSL offload, which can protect their security solutions from oversubscription. Forty-seven percent analyze telemetry from this infrastructure with their security operations tools.

Application Delivery Infrastructure Management

Management System Integration

EMA asked enterprises to identify which IT systems they integrate with application delivery infrastructure. IT service management and public cloud management and orchestration are the top integration priorities, as **Figure 24** details. Everything else is secondary. Data center SDN is the lowest priority. However, data center operations teams and DevOps teams indicated much more interest in SDN integration. Also, successful application delivery teams identified data center SDN integration as one of their highest priorities.

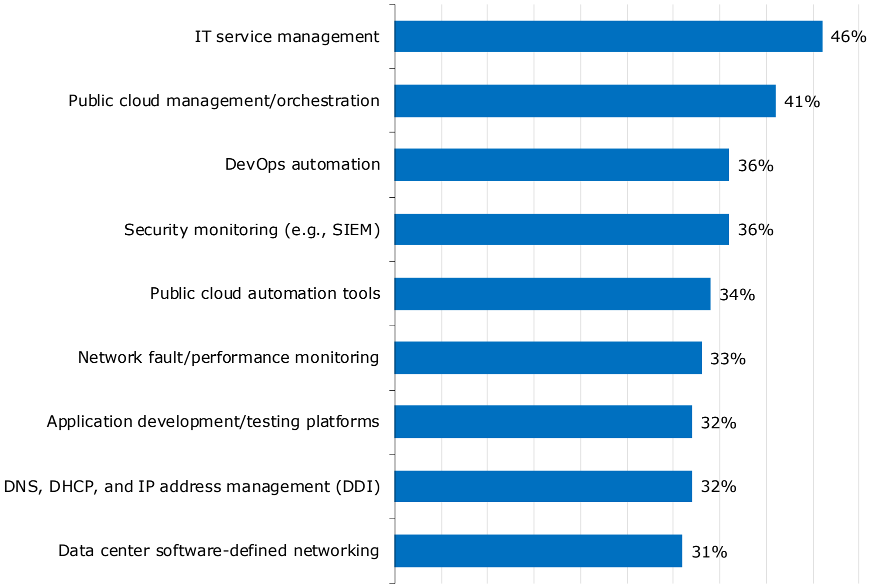

Sample Size = 253, Valid Cases = 253, Total Mentions = 816

Figure 24. Integration priorities for application delivery infrastructure

DevOps teams expressed more interest in integration with DevOps automation. Application development teams, application management teams, and project management teams all showed stronger interest in integration with application development and testing platforms. Cloud engineering, application management, data center operations, and information security all expressed stronger interest in integration with security monitoring. Very large enterprises are the least likely to prioritize integration with security monitoring tools and cloud management/orchestration tools.

Telemetry and Observability

Application delivery controllers and load balancers are well-situated for providing visibility into application traffic. Thus, enterprises can collect metrics and telemetry from this infrastructure to understand security, application performance, and network performance. **Figure 25** reveals that 82% of enterprises are collecting telemetry and statistics from their application delivery infrastructure.

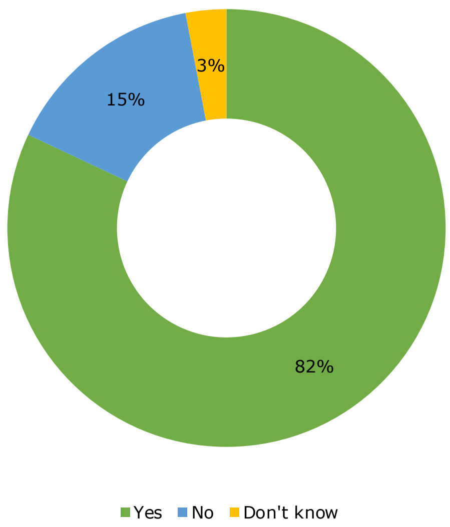

Sample Size = 253

Figure 25. "Does your organization collect and analyze telemetry and statistics from your application delivery infrastructure?"

Organizations that are successful or somewhat successful with application delivery infrastructure are more likely than unsuccessful organizations to collect this telemetry.

82%

of enterprises are collecting telemetry and statistics from their application delivery infrastructure.

EMA asked research participants to identify the tools they use to analyze this data. According to **Figure 26,** the average enterprise is using two different tools to analyze this data, and none of them stand out as the de facto choice. However, third-party security monitoring tools are the most popular. The second most popular approach is to export the data to a data repository for analysis by a standalone analytics solution, such as the popular open-source ELK stack (Elasticsearch-Logstash-Kibana).

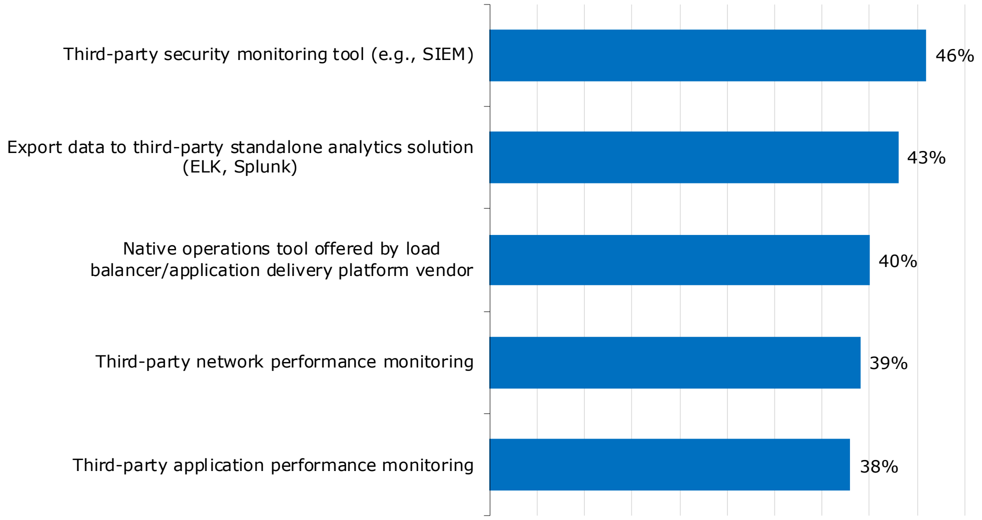

Sample Size = 208, Valid Cases = 208, Total Mentions = 429

Figure 26. Tools used to analyze application delivery infrastructure telemetry

Successful application delivery teams demonstrated a preference for native tools offered by their load balancer and ADC vendors. Data center operation teams are more likely to send this data to an application performance management tool. Infrastructure engineering teams and cloud engineering/operations teams are more likely to send data to security monitoring tools.

Given that security monitoring is the primary destination for this telemetry, it's not surprising that the most valued benefit of this data analysis is the ability to detect security incidents and attacks, as revealed in **Figure 27.** The other top benefits are predictive capacity management and accelerated triage and troubleshooting. Cloud engineering and operations teams are particularly likely to recognize the benefit of accelerated triage and troubleshooting. Network operations, however, is less likely to perceive this benefit.

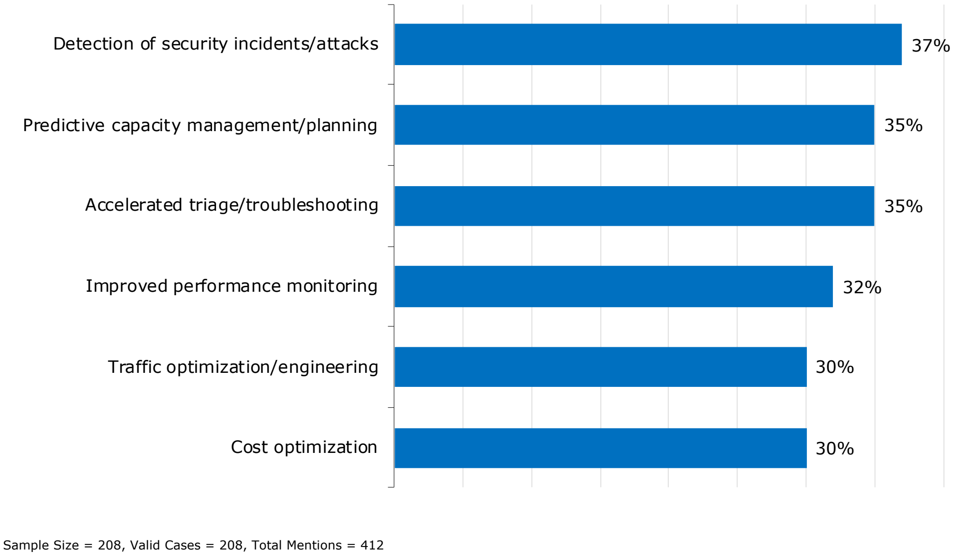

Figure 27. Most important benefits of telemetry analysis

Although improved network performance is only a middling benefit overall, successful application delivery teams identified it as one of their highest priorities. Cost optimization and traffic optimization are the least important benefits of this analysis, but still prioritized by three in ten enterprises.

AIOps and Application Delivery Infrastructure

Some vendors of load balancers and application delivery controllers have started developing AIOps capabilities, applying machine learning and statistical analysis to the telemetry collected from their products. EMA asked research participants if they would use such offerings to enhance management and automation of their infrastructure. **Figure 28** reveals that 85% said yes, and nearly half of them said this capability would be essential to their operations. Successful application delivery teams are more likely to identify these AIOps capabilities as potentially essential.

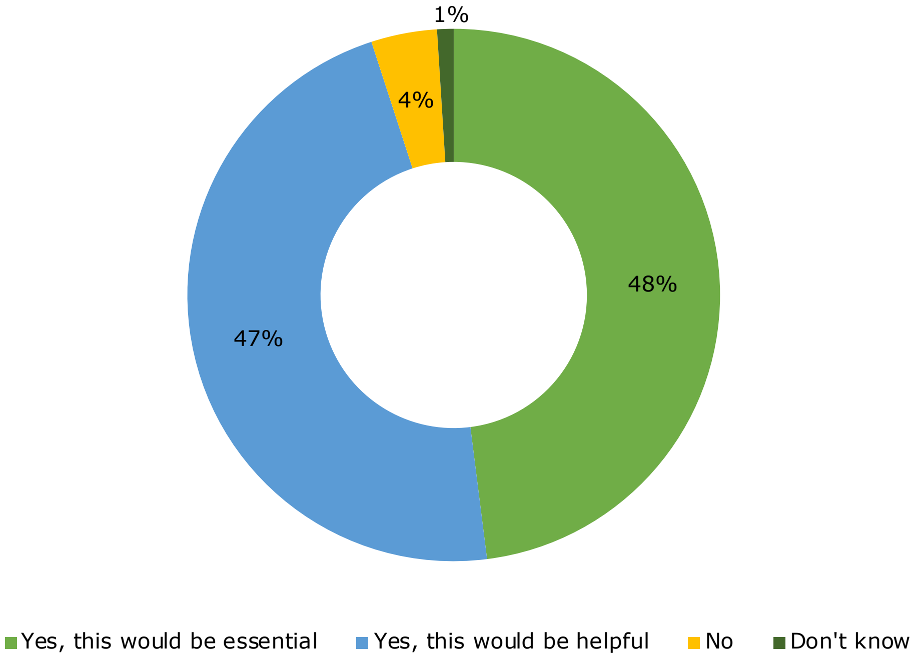

Sample Size = 253

Figure 28. "If your application delivery infrastructure vendor offered AIOps capabilities to enhance management and automation, would you use them?"

Midmarket enterprises are the least enthusiastic about AIOps capabilities. They are less likely to see such offerings as critical.

Enterprises have four top use cases for these AIOps capabilities, according to **Figure 29.** First, they are interested in automated root cause analysis of service problems. Second, they are interested in user behavior analytics. Automated remediation of service trouble and predictive alerting round out the top four. Automated remediation is especially appealing to DevOps, application management, and project management, but not of much interest to the infrastructure engineering team.

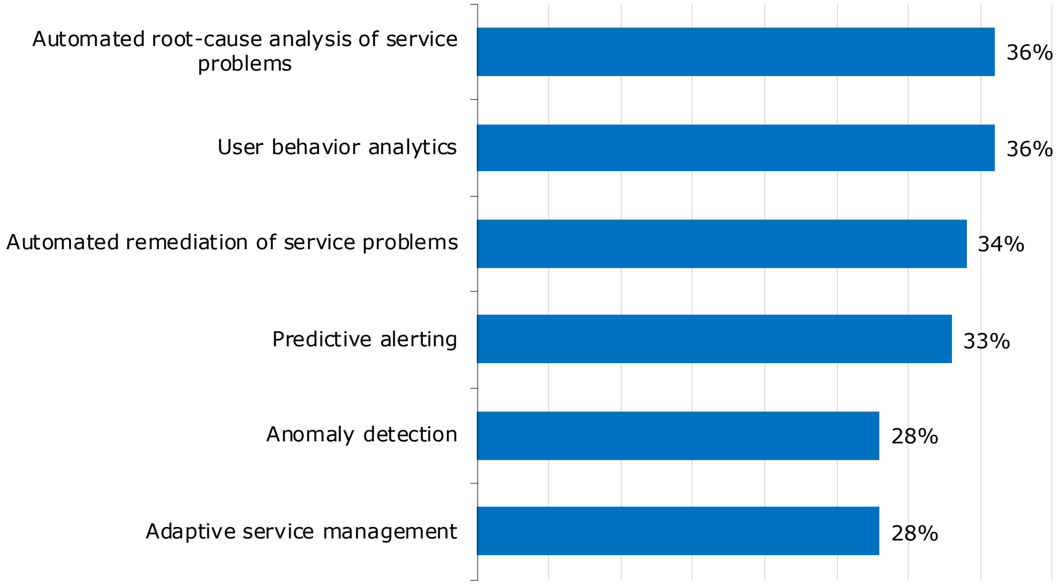

Sample Size = 239, Valid Cases = 239, Total Mentions = 468

Figure 29. Top use cases for AIOps solutions offered by application delivery vendors

Anomaly detection and adaptive service management are the lowest priorities. However, application management professionals consider anomaly detection a top use case, while network operations and infrastructure engineering teams are both very uninterested in this opportunity.

EMA observed some variation of priorities based on how an individual interacts with this infrastructure. Anomaly detection is more appealing to people who primarily research and evaluate application delivery infrastructure, but it is less interesting to people who own and operate this infrastructure and people who develop and deploy applications on the infrastructure. On the other hand, these latter two classes of people (infrastructure owners and application owners) are very interested in automated root cause analysis. People who primarily own infrastructure budgets and make procurement decisions are the least likely to see the value of automated root-cause analysis.

88%

of enterprises automate management of their application delivery infrastructure in their data centers, although only 31% have extensive automation.

Infrastructure Automation

Infrastructure Automation in the Data Center

EMA asked enterprises to characterize their automation of application delivery infrastructure in their data centers. **Figure 30** revealed that 88% of enterprises automate management of their application delivery infrastructure in their data centers, although only 31% have extensive automation. The majority automate only some management tasks.

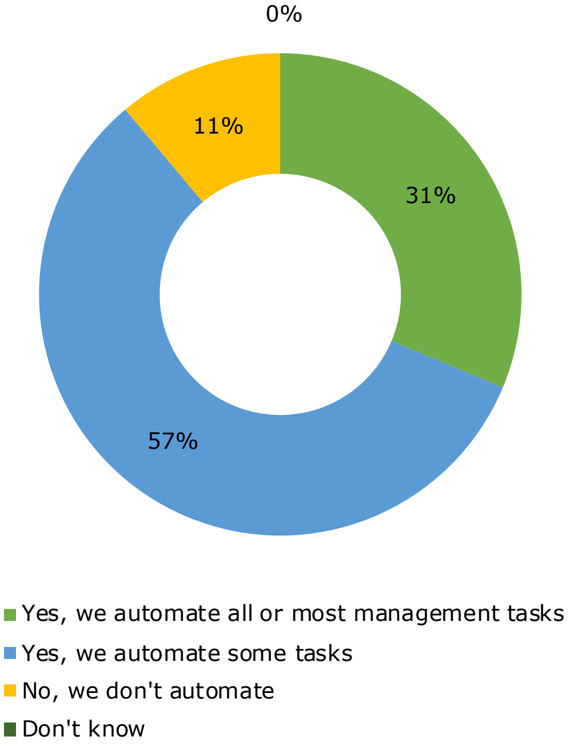

Sample Size = 253

Figure 30. "Does your organization automate management of load balancing and application delivery controllers in your data center(s)?"

Very large enterprises are the most likely to say they don't automate data center infrastructure. Meanwhile, nearly all midmarket firms are doing some kind of automation. Successful application delivery teams are more likely to have extensive automation, while somewhat successful teams are more likely to have only limited automation.

Budget is an issue with data center automation. Enterprises with modest budget growth (less than 10%) and enterprises with flat or shrinking budgets are ten times more likely than enterprises with higher budget growth to report no automation of data center infrastructure.

Enterprises revealed two clear preferences for automation tools in their data centers, as shown in **Figure 31**. First, they use native automation tools provided by their application delivery infrastructure vendors. Nearly as many prefer using a third-party commercial automation solution. Homegrown software and one-off scripts are the least common approaches to automation currently.

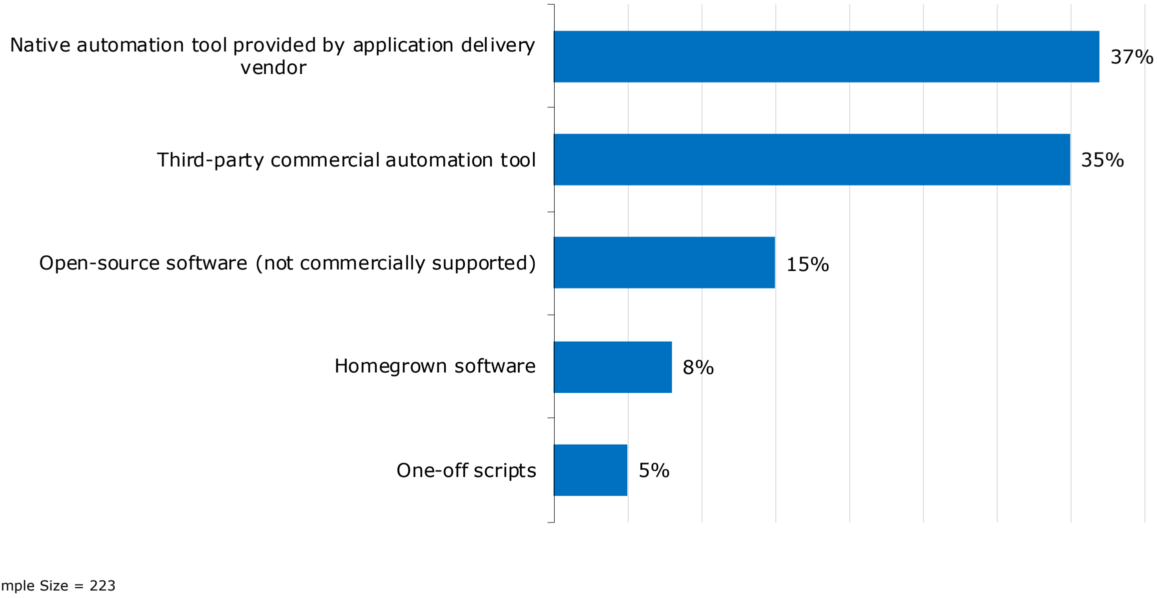

Figure 31. Primary approach to infrastructure automation in data centers

Very large enterprises are the most likely to rely on scripts for data center infrastructure automation.

Unsupported open-source software is a distant secondary preferred automation approach; however, data center operations, application development, and application management teams are more likely to prefer open-source, while infrastructure engineering teams are much less likely to prefer it.

Only 41% of enterprises that automate application delivery infrastructure in their data centers are satisfied with that automation. **Figure 32** shows that more than half are only somewhat satisfied, suggesting that they see room for improvement. Cloud engineering teams are more likely to be happy with data center automation, but several groups are less likely to be happy, including infrastructure engineering, DevOps, data center operations, and project management.

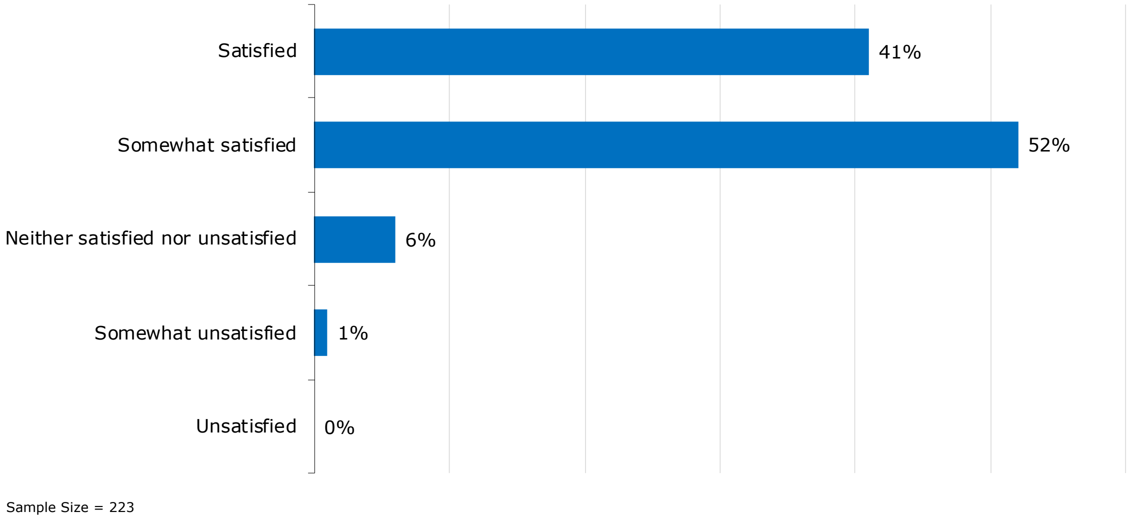

Figure 32. Satisfaction with data center infrastructure automation capabilities

Successful application delivery teams tend to be satisfied with data center infrastructure automation. Meanwhile, somewhat successful teams are more likely to be only somewhat satisfied with this automation.

Only 41% of enterprises that automate application delivery infrastructure in their data centers are satisfied with that automation.

Infrastructure Automation in Public Clouds and Cloud-Native Application Environments

Extensive application delivery infrastructure automation is slightly more common in the public cloud and in cloud-native environments, as detailed in **Figure 33**. However, the number of enterprises that claim to do no automation at all is also slightly higher than in the data center.

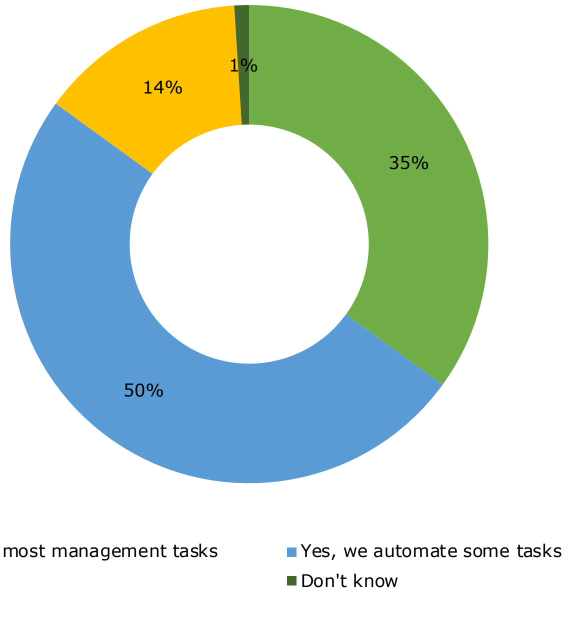

Figure 33. “Does your organization automate management of load balancing and application delivery controllers in public cloud and/or cloud-native platform environments?”

As with the data center, budget growth correlates with automation in the cloud. Organizations with flat or shrinking application delivery budgets are 12 times more likely than those with high rates of budget growth to have no automation in the cloud.

Automation tool preferences are very different in the cloud, as described by **Figure 34.** Third-party commercial automation tools are the clear favorites. Whereas native automation tools offered by infrastructure vendors are the most popular class of solution for data center infrastructure, this capability is a distant second choice in the cloud. Half of all large enterprises in this survey rely on third-party commercial software. Midmarket firms were less likely to do so.

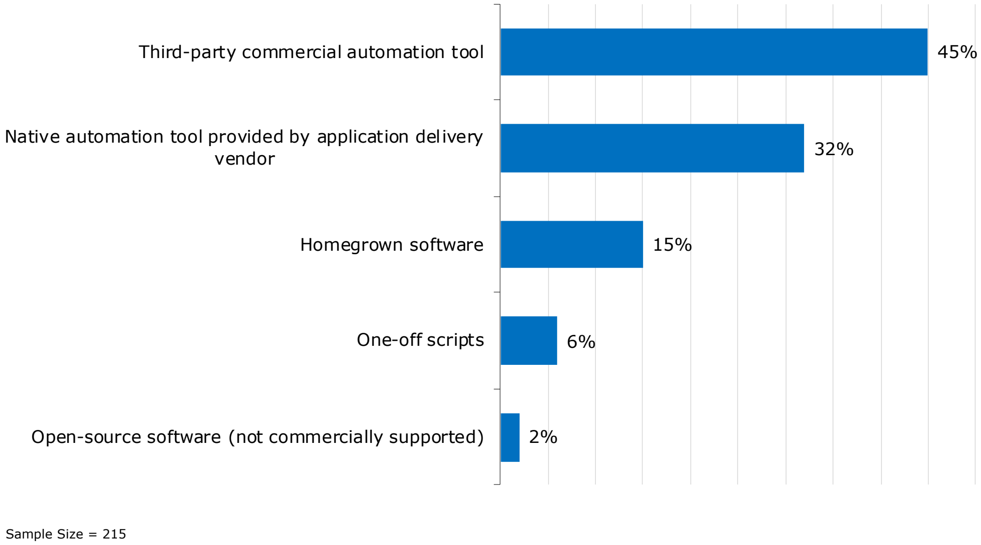

Figure 34. Primary approach to infrastructure automation in public cloud and cloud-native infrastructure

Homegrown software is more popular in the cloud than it is in the data center, while open-source software is only preferred by a handful of companies in the cloud. Homegrown software is a popular approach for midmarket companies. Very large enterprises are more likely than others to rely on scripts.

Satisfaction with cloud infrastructure automation is almost identical to satisfaction with data center infrastructure automation, as **Figure 35** verifies. However, EMA observed some variation by IT group. NetOps teams and cloud engineering/operations teams are more likely to be satisfied, and infrastructure engineering and data center operations are less likely to be happy.

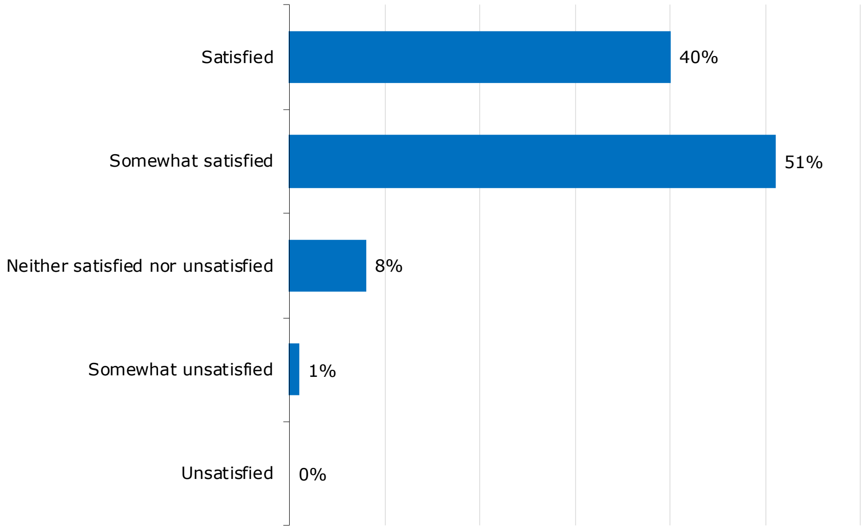

Sample Size = 215

Figure 35. Satisfaction with public cloud and cloud-native infrastructure automation capabilities

Successful organizations are more likely to be fully satisfied with their application delivery infrastructure automation in the cloud. Less successful teams tend to be only partially satisfied.

License Model Preferences

As the application delivery infrastructure industry has moved away from vertically integrated hardware appliances, the perpetual software licenses that were typically sold as part of those appliances have fallen out of favor, both in the data center and the public cloud. As **Figure 36** reveals, the most popular license model for application delivery infrastructure is now a subscription model. Perpetual licenses and pay-as-you-go licenses are distant secondary preferences. Pay-as-you-go is slightly more popular in the public cloud, while perpetual licenses are a little more popular in the data center. Open-source and freemium licenses are the least popular.

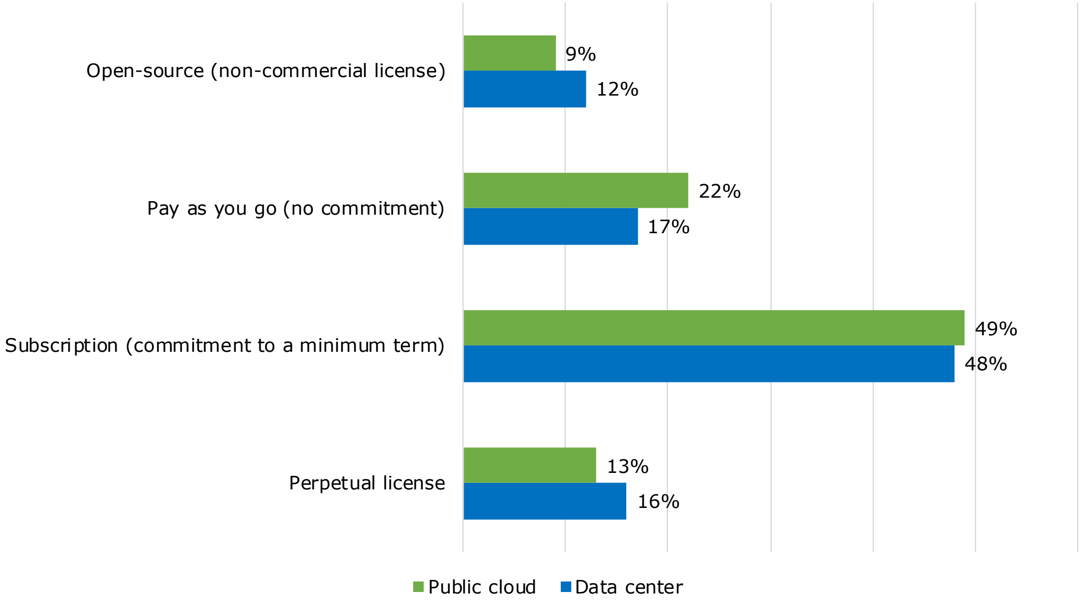

Figure 36. Preferred license models for application delivery infrastructure

Successful application delivery teams are a little more likely than others to prefer a perpetual license in data center and cloud environments, suggesting that this might be the ideal licensing approach in many circumstances. Still, subscription licenses are the most popular license model for enterprises, regardless of how successful they are.

Cloud engineering and information security teams are more likely to push for subscription licenses in the data center, while network operations, infrastructure engineering, and application development are less interested in subscriptions. Application development teams are very likely to prefer open-source licenses in the data center.

In the public cloud, infrastructure engineering is more likely to push for a subscription license. Network operations is more likely to prefer a pay-as-you-go license in the public cloud. Application development prefers open-source licenses in the cloud.

Very large enterprises are more likely than others to seek perpetual licenses in the data center and the cloud. Midmarket companies are most likely to seek subscription licenses in the data center.

Service Mesh Perspectives

Service mesh is an emerging technology associated with cloud-native, containerized application platforms. It is essentially an infrastructure layer built into a microservices application architecture. In the classic service mesh deployment, service mesh proxies are deployed as “sidecars” side by side with the containers that comprise the application. In these environments, there isn’t a standalone infrastructure layer of application delivery controllers and load balancers. The leading service mesh offerings today are open-source software, such as isitio and linkerd.

85%
of the enterprises in this survey have interest in using service mesh.

EMA found that 85% of the enterprises in this survey have interest in using service mesh, as **Figure 37** indicates. Infrastructure engineering teams showed the strongest interest, while network operations teams showed less interest. Individuals who primarily engage with application delivery solutions as owners and operators of infrastructure are the least likely to be interested in service mesh, which is unsurprising given that many of them may end up being bypassed by the cloud-native application platform owners who deploy them.

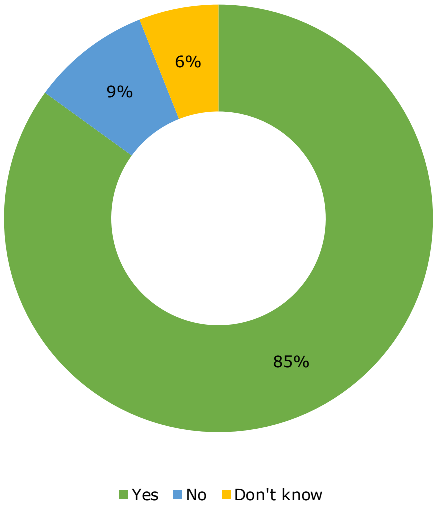

Sample Size = 253

Figure 37. “Is your organization interested in using service mesh in its cloud-native application environments?”

Enterprises that have seen strong growth in their application delivery budgets are much more likely to have interest in service mesh than companies with flat or shrinking budgets.

Service mesh is not application delivery controller software. It is not comparable to the virtual application delivery controllers that enterprises might deploy as a virtual machine on a hypervisor host.

No vendors have necessarily claimed this space. Thus, EMA asked research participants who indicated interest in service mesh to reveal how they might prefer to consume service mesh solutions. The two most popular options are solutions offered by cloud-native infrastructure vendors (e.g., providers of Kubernetes solutions) and solutions from vendors of load balancer or application delivery software. The former preference indicates an affinity to aligning service mesh with application development platforms and processes. Infrastructure engineering teams have the strongest interest in this option.

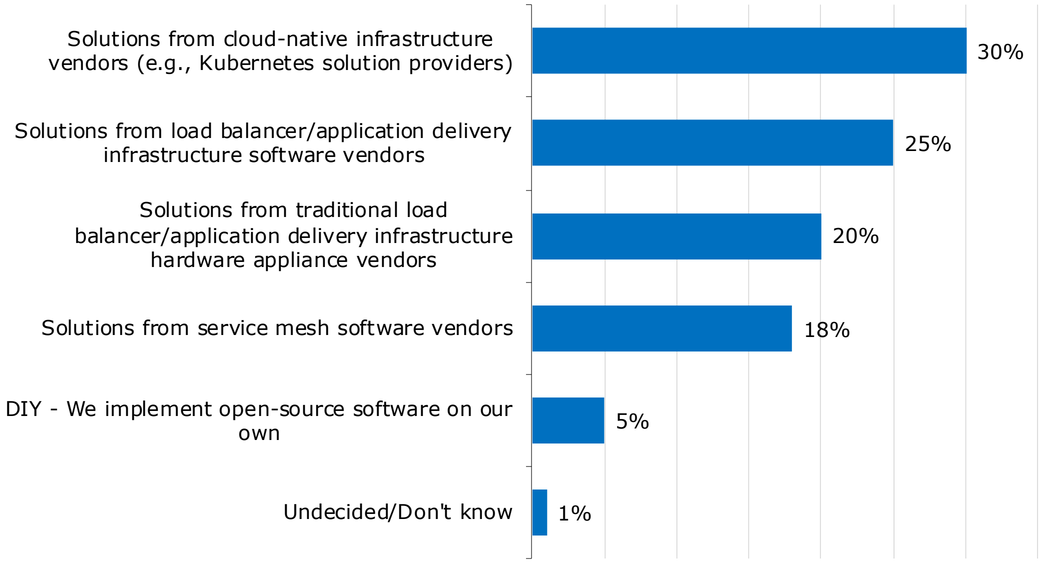

Sample Size = 214

Figure 38. How enterprises want to consume service mesh solutions

The latter preference for load balancer and application delivery controller software vendors indicates an interest in aligning service mesh with network software, but not necessarily traditional hardware. Information security professionals have a strong preference for this consumption model. Shockingly, application development teams prefer solutions from these infrastructure vendors, while infrastructure engineering teams do not. These latter findings are the reverse of what EMA anticipated.

Solutions from traditional hardware vendors are the third-most preferred consumption option, followed by solutions from service mesh software specialists. Hardware vendors are more likely to be the preference of application management, data center operations, and project management teams. Individuals who primarily engage with application delivery as owners and operators of infrastructure are the most likely to prefer solutions from application delivery controller hardware vendors, revealing that they are turning to their legacy vendors for guidance with this emerging technology. Unsurprisingly, individuals who primarily engage with application delivery as a developer and deployer of applications on infrastructure are the least likely to turn to a hardware vendor.

Service mesh specialists are more likely to be the preference of cloud engineering teams. Very few enterprises want to take a DIY approach to service mesh by implementing open-source software on their own without commercial support.

Management Integration Between Service Mesh and Application Delivery Infrastructure

EMA asked research participants whether they need to integrate management of service mesh with management of their traditional load balancer and application delivery controller infrastructure. Only 2% said no. Nearly half recognize this integration as critical, while another half see it as only helpful. Network operations teams, infrastructure engineering teams, and application management teams are more likely to say this integration is only helpful.

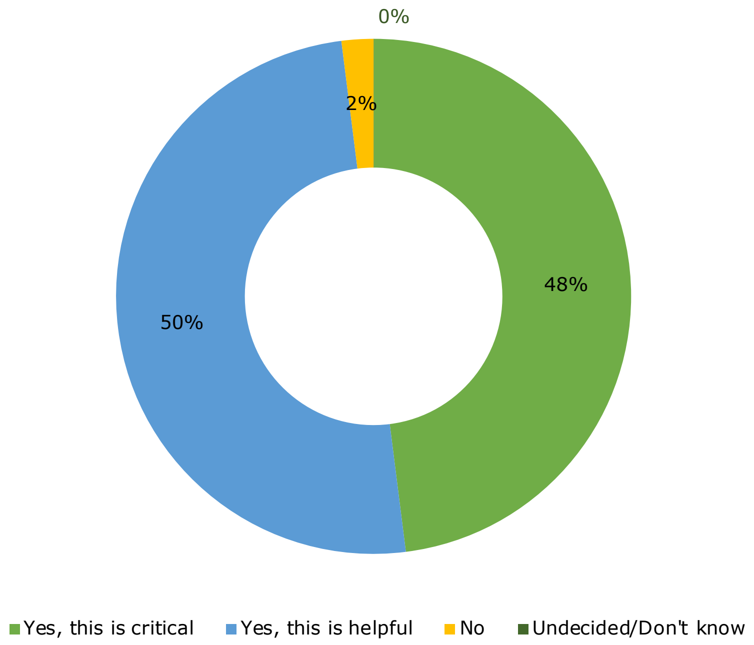

Sample Size = 214

Figure 39. “Does your organization need to integrate management and orchestration of service mesh technology with management of traditional load-balancer/application delivery infrastructure?”

Critical Service Mesh Capabilities

Research participants revealed that service authentication is the most beneficial capability of a service mesh, as detailed in **Figure 40**. Secondary benefits are observability/visibility and encryption. Traffic control and policy enforcement are the lowest priorities.

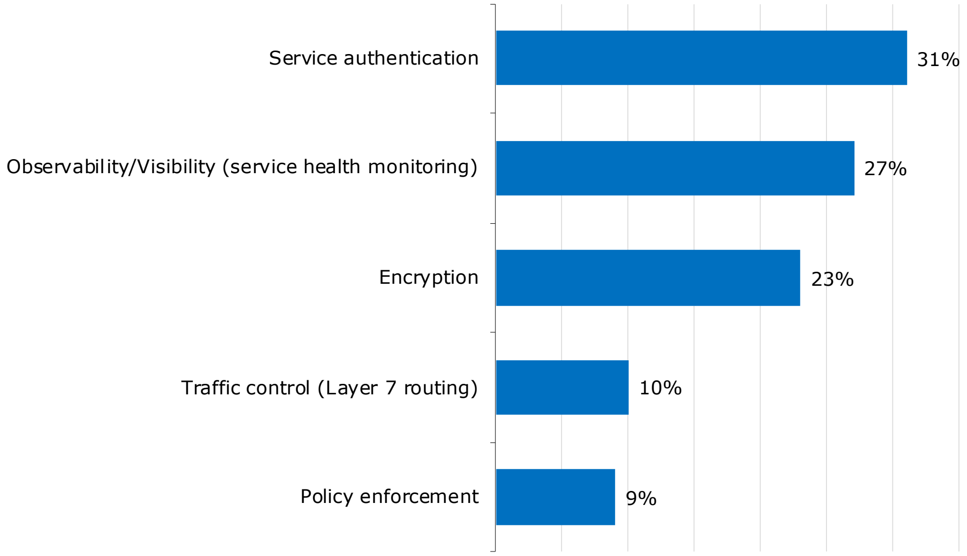

Sample Size = 214

Figure 40. Most beneficial capabilities of service mesh

Infrastructure engineering, cloud engineering, application development, and data center operations teams were all more likely to single out observability. Very large enterprises expressed stronger interest in policy enforcement, while very few midmarket companies were particularly interested in this capability.

The Impacts of the Coronavirus Pandemic

EMA started work on this research just as the World Health Organization declared a pandemic of the COVID-19 virus. Thus, EMA added some questions to the survey about how the public health response has impacted application delivery infrastructure, particularly given the massive surge in people forced to work from home.

First, EMA asked respondents to indicate what changes they made to this infrastructure in response to the business conditions dictated by the pandemic. **Figure 41** shows that only 10% claimed to have made no changes in response to the situation. Overall, the average enterprise is taking two actions. The most common responses are the addition of new security functionality, increased infrastructure throughput, scaled up session capacity, and increased use of automation. Many are also adding infrastructure to a new cloud provider, suggesting that they are scaling out overall services in the cloud.

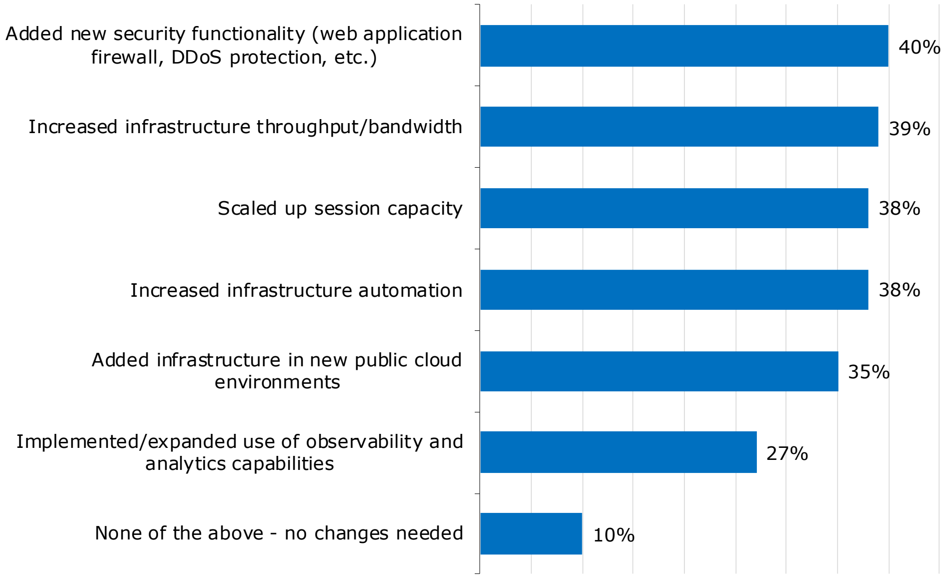

Sample Size = 253, Valid Cases = 253, Total Mentions = 575

Figure 41. Changes made to application delivery infrastructure in response to the COVID-19 pandemic

Data center operations teams and security teams are more likely to say they had scaled up session capacity, while DevOps are less likely. Application management teams are more likely to increase their use of automation, while data center operations teams are less likely.

The least common response to the pandemic is the implementation or expansion of observability and analytics capabilities.

Very large enterprises are twice as likely as other companies to report having made no changes to infrastructure in response to the pandemic. Twenty-one percent of enterprises with flat or shrinking budgets said they had made no changes in response to the pandemic, making it obvious that companies with limited resources are in a tight spot with this crisis. Companies with modest budget growth are more likely than those with strong budget growth to report the addition of new security functionality. Meanwhile, companies with strong budget growth are more likely to increase their use of automation.

EMA also asked respondents to describe how they are using this infrastructure to ensure business continuity and the productivity of users working from home during the pandemic. First, enterprises are relying more on the scalability and performance that this infrastructure delivers to applications, as **Figure 42** reveals. Second, they are using this infrastructure to manage the bandwidth consumption of individual remote user sessions. Network operations teams are especially likely to manage bandwidth consumption of remote users, while DevOps and information security are not. Very large enterprises are also more likely to perform this bandwidth management of user sessions.

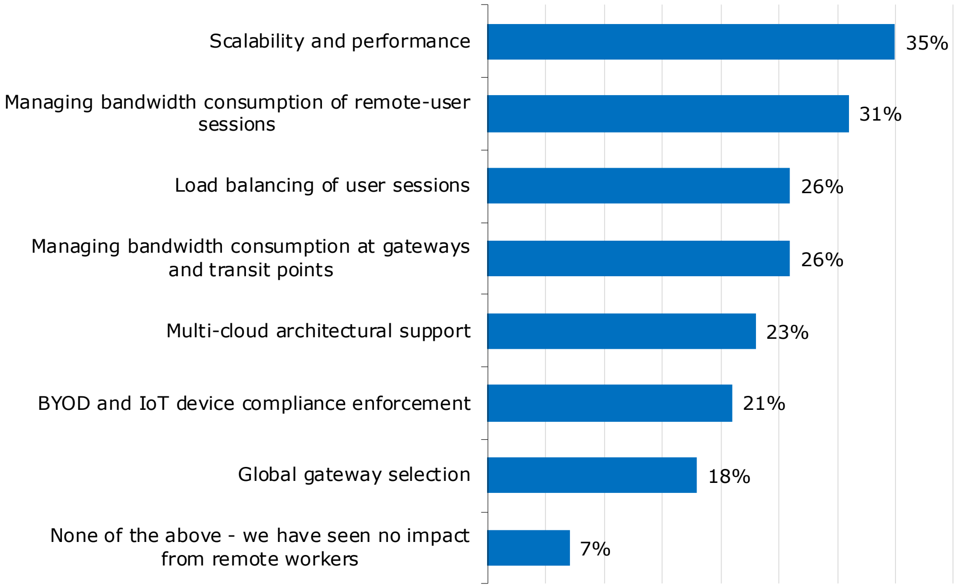

Sample Size = 253, Valid Cases = 253, Total Mentions = 473

Figure 42. Aspects of application delivery infrastructure most important to ensuring business continuity and productivity of users working from home during the COVID-19 pandemic

Load balancing of user sessions and managing bandwidth consumption at gateways and transit points are the secondary capabilities. Managing transit point and gateway consumption is especially popular for infrastructure engineering and application development teams. Cloud engineering teams are more likely to focus on load balancing of user sessions.

Global gateway selection and BYOD/IoT device compliance enforcement are the least important capabilities. BYOD and IoT device compliance is a lower priority for successful application delivery teams.

Succeeding with Application Delivery Infrastructure

Only 38% of enterprises are fully successful with their use of load balancers and application delivery infrastructure. More than half are somewhat successful, meaning they see room for improvement. Application development teams are the most likely to rate their efforts as successful. Their focus is on services deployed in development and test environments, rather than in production. Infrastructure engineering, application management, project management, and information security teams are all more pessimistic than the average respondent about success.

Only 38% of enterprises are fully successful with their use of load balancers and application delivery infrastructure.

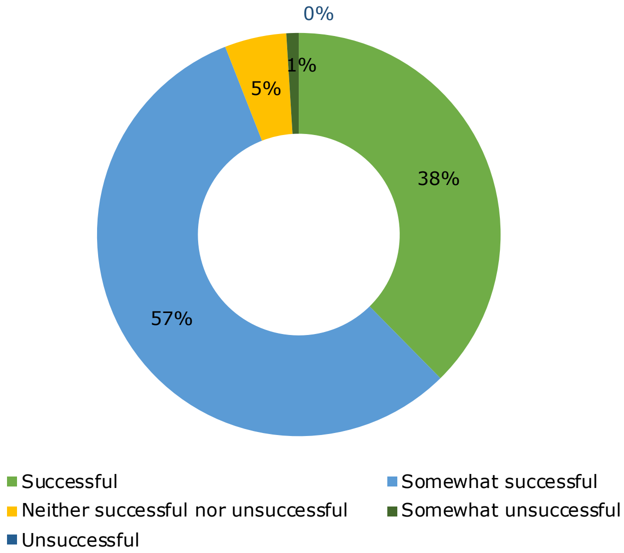

Sample Size = 253

Figure 43. Overall success with load balancer and application delivery infrastructure

Enterprises must be willing to spend on this technology. Success with application delivery infrastructure correlates strongly with budget growth. Those with the highest rates of budget growth are the most likely to be successful with this infrastructure. Companies with flat or shrinking budgets are much less likely to succeed.

Business and Technical Challenges

Figure 44 reviews the top business challenges that enterprises encounter with their application delivery infrastructure. A skills gap is the biggest challenge with application delivery infrastructure. Some of the various teams that work with this technology lack the expertise to manage it properly. Change control/role-based access, budget shortfalls, and conflicts between teams are the chief secondary challenges. Data center operations teams are more likely to complain about budgets.

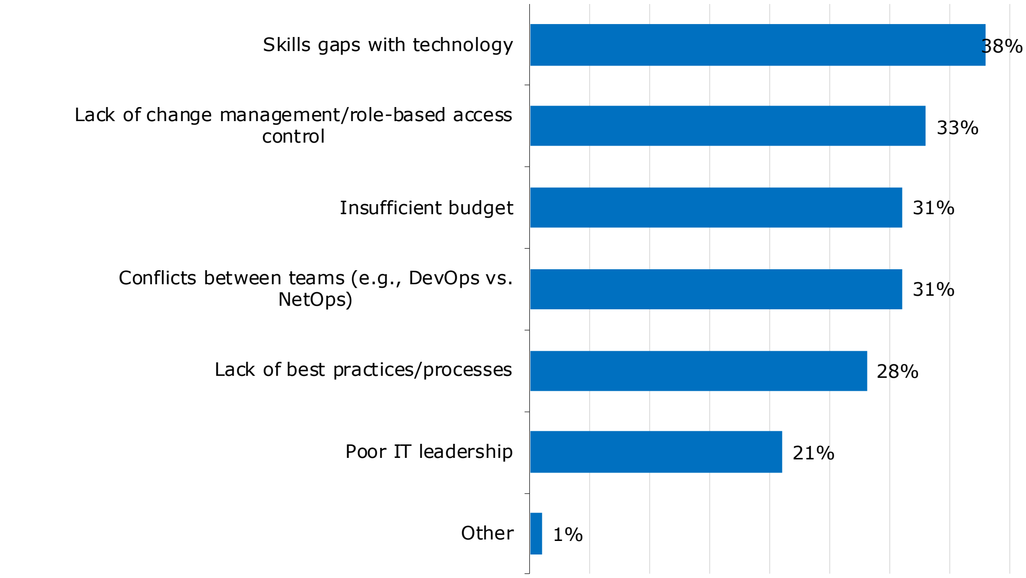

Sample Size = 253, Valid Cases = 253, Total Mentions = 462

Figure 44. Business challenges associated with load balancers and application delivery infrastructure

Poor IT leadership is the least prominent business challenge. Data center operations are more likely to perceive such a problem, while infrastructure engineering and DevOps teams are less likely to see this issue. More than half of very large enterprises revealed that skills gaps are one of their biggest business challenges. Poor IT leadership is the one business challenge that distinguishes successful and somewhat successful teams. Somewhat successful teams are more than twice as likely to struggle with this issue.

> *A skills gap is the biggest challenge with application delivery infrastructure. Some of the various teams that work with this technology lack the expertise to manage it properly.*

Figure 45 reviews the technical challenges that companies grapple with in application delivery infrastructure. First, enterprises are struggling most often with the performance limitations of their application delivery infrastructure, which suggests that they need to upgrade the throughput of these solutions. Next, they are struggling with infrastructure scalability.

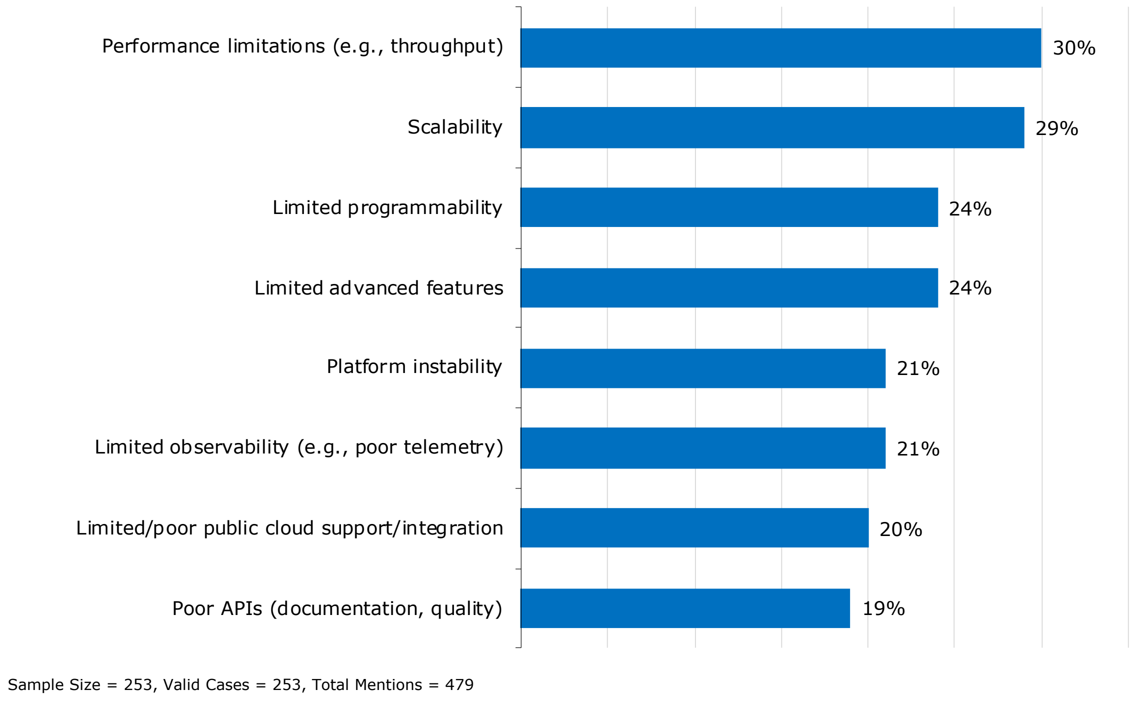

Figure 45. Technical problems associated with load balancers and application delivery infrastructure

The two leading secondary challenges are limitations with programmability and advanced features. The lack of programmability partially explains the significant interest enterprises have in automation of this infrastructure.

The least common challenge stems from poor APIs offered by platform providers. However, application management and data center operations teams identified APIs as one of their most problematic issues. Infrastructure engineering teams were the least likely to complain about them.

Figure 46 reveals that 71% of enterprises believe that some of the application delivery infrastructure functions they use are too difficult or time-consuming to configure and manage, which again explains the high interest this research found in automating this infrastructure. EMA observed that some groups have different perspectives on this issue. The application development team is very unlikely to say yes to this question, but infrastructure engineering, application management, data center operations, and information security are more likely to say this is a problem.

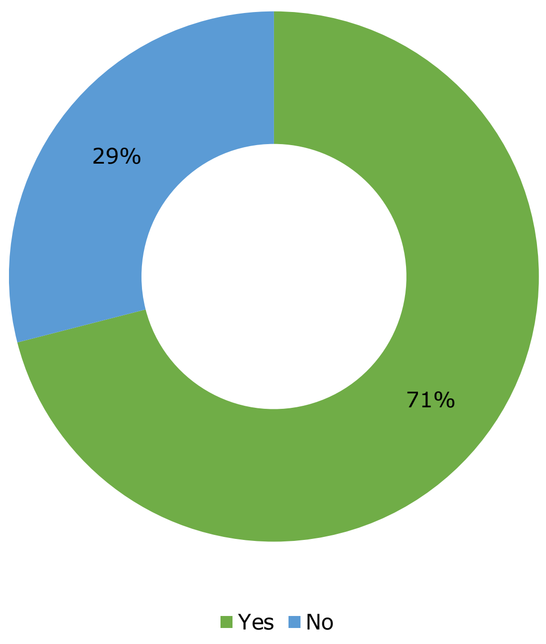

Sample Size = 253

Figure 46. "Do you believe that some functions of your application delivery infrastructure are too difficult and time-consuming to configure and manage?"

Very large enterprises are the least likely to say this is a problem for them. Successful organizations are also less likely to say managing application infrastructure is too difficult and time-consuming.

Figure 47 identifies the functions that present the most management problems. Global load balancing, content routing, and SSL acceleration are the three functions that are most difficult and time-consuming to manage. Application acceleration and secure application gateway are also relatively more difficult. Data compression and identity and access management present the fewest management challenges.

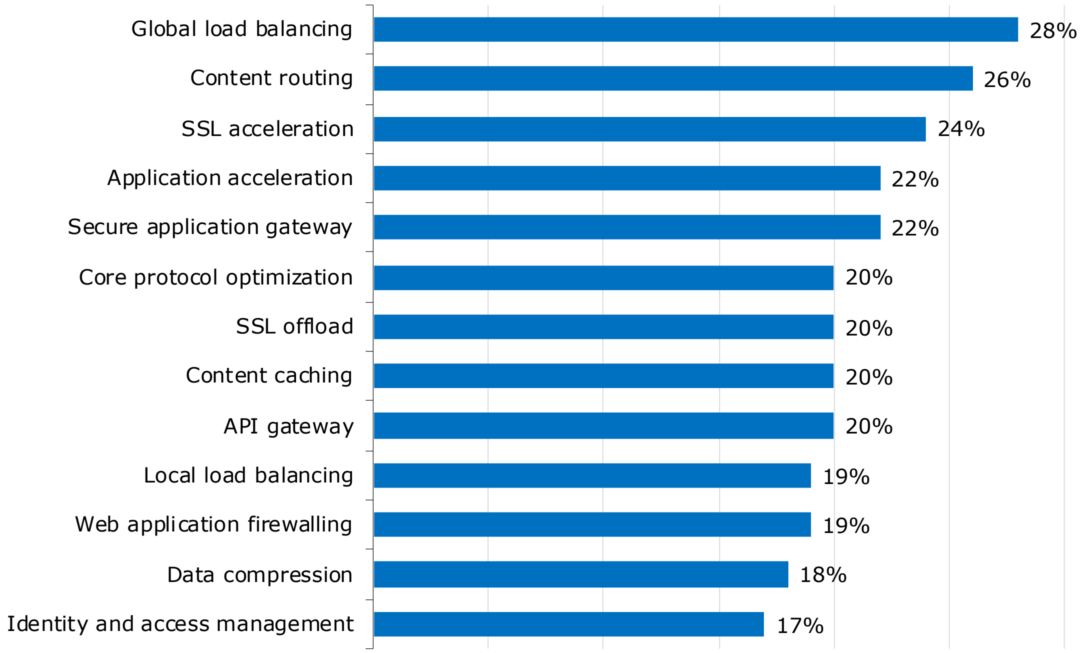

Sample Size = 179, Valid Cases = 179, Total Mentions = 493

Figure 47. Application delivery functions that are the most difficult and time-consuming to manage

Although successful organizations are less likely to struggle with application delivery management, those that do struggle pointed to web application firewalls as one of their biggest management headaches. Very large enterprises are twice as likely as other companies to have difficulty with identity and access management.

Conclusion

This research shows that the days of appliance-based application delivery infrastructure are over. Enterprises are embracing cloud-native software, both in the data center and the cloud. Furthermore, service mesh is top of mind for most enterprises.

This transition to software and the cloud has fragmented infrastructure engineering and operations, but enterprises are working to close this gap, particularly because it presents a security risk. Infrastructure vendors can help with this issue by adding multi-vendor support to their own management tools and increasing deployment flexibility. The former will be particularly helpful as enterprises adopt service mesh.

Enterprises are particularly interested in application delivery infrastructure automation. They also want to leverage the telemetry that is available within this infrastructure. Beyond simple telemetry generation, enterprises would welcome AIOps capabilities from their infrastructure vendors to improve operations and security.

EMA was disappointed to find that only a minority of companies are fully successful with this infrastructure, although automation and AIOps appear to be paths to improving this issue. Enterprises must press forward with finding ways to advance their success with application delivery infrastructure because it remains critical, even as application architectures evolve and migrate to the cloud. In fact, this research found that this infrastructure has proven itself indispensable to business continuity during the current COVID-19 pandemic. It’s good that most enterprises are expanding their budgets for this technology.

EMA will continue to follow this market and adjacent markets closely with its ongoing research into infrastructure and operations.

Case Study: Atlassian Builds SaaS Offering in AWS with Pulse Secure Virtual Application Delivery Controller

This case study was provided by Pulse Secure.

When enterprise software company Atlassian moved its SaaS offering from its private data centers to Amazon Web Services, it partnered with Pulse Secure for application delivery infrastructure.

For several years prior to its move to AWS, Atlassian had used Pulse Secure Virtual Traffic Manager (vTM) as a software-based Layer 7 application delivery controller (ADC) within its own managed data centers. In addition to removing the burden of managing a network of data centers, the move to AWS provided Atlassian with several additional advantages, including more deployment flexibility and enhanced resiliency through its global presence.

However, the built-in AWS application traffic management capabilities did not have the advanced features offered by the Pulse Secure vTM, such as TrafficScript and customization that Atlassian had benefited from within its data center SaaS solution.

Compared to open-source alternatives that would require integration between multiple elements to gain a comparable level of functionality, Pulse Secure vTM offers a single platform designed to integrate seamlessly with any application deployed on Amazon Web Services. This includes load balancing, application scalability, and fine-grained application control.

Atlassian also deployed the Pulse Secure Services Director to allow flexible sharing of Pulse Secure vTM capacity. This enabled cost-efficient resource allocation based on demand, further helping the expansion of its operations and move to the cloud. As a result, the company can deploy resources wherever and whenever they are needed, which improves business agility, operational flexibility, and scalability.

Benefits

Alongside ease of migration, another key capability was the vTM's built-in support for Terraform, an open-source tool for describing and automating the provisioning and configuration of application services. Atlassian also benefited from the enhanced bandwidth management capabilities within Pulse Secure vTM, including Request Rate Shaping. This capability allowed Atlassian to spread available resources more evenly across connections to ensure that all SaaS users were gaining an equal level of performance with the lowest levels of latency.

Pulse Secure vTM also helped the Atlassian IT team track and troubleshoot performance by inserting custom tracing tags into live traffic for collection and analysis. "Using Pulse Secure's vTM TrafficScript capability, we were able to insert Zipkin tags into transactions that help gather timing data needed to troubleshoot latency problems across our microservice architectures," said Nicolas Meessen, senior principal network engineer at Atlassian. This JSON data is imported into a centralized trace service and provides the Atlassian IT team with a rich set of dashboards to quickly identify and fix any service issues.

The implementation of Pulse Secure vTM on AWS continues to be a success and Atlassian is handling more than 60,000 dynamic content requests per second through multiple clusters of virtualized appliances. “Pulse Secure vTMs are a critical part of how we deliver our SaaS service,” says Meessen. “The technology helps us to provide the best possible customer experience while offering us the flexibility to grow and adapt as our products and services evolve.”

With the combination of its progressive use of cutting-edge technology and processes that include Pulse Secure vTM, Terraform, and its own software, such as Jira and Confluence, Atlassian is moving toward a more DevOps-focused culture. “Elements such as Pulse Secure vTM that support open standards and transparency really match the philosophy of Atlassian, and we look forward to working closely in the future,” Meessen adds.

ASOS uses Kemp in Microsoft Azure to Handle 167 Million Website Visits on Black Friday

This case study was provided by Kemp.

As a fully online business, European online fashion giant ASOS requires continuous uptime all year round. Any downtime or slowness from the store can lead to lost sales and an impact on customer loyalty. As part of an improvement project to ensure availability of service, ASOS reengineered their online store to convert it from a monolithic platform running in legacy data centers into a suite of interacting microservices deployed on the Microsoft Azure Cloud. This allowed them to run the core store functions in the European and North American Azure instances, where most of the customer base resides, and also to deploy time-critical components to Asian Azure instances as required. The move to the global Azure Cloud infrastructure also provided ASOS with the ability to deliver the best performance possible to customers by using local Azure resources, and also allowed them to maximize availability by load balancing globally across the Azure instances.

However, a crucial part of making this high-workload cloud infrastructure work would be a load balancing infrastructure that could support it on a global scale at even the most stressful periods.

ASOS required an application delivery solution that could load balance across multiple geographic regions, be fully supported in the Azure Marketplace, provide a scriptable interface for deployment and management, and provide flexibility to scale on demand and support continuous delivery

Benefits

ASOS ultimately found that Kemp LoadMaster met these requirements with integrated Global Server Load Balancing (GSLB), full support across all Azure marketplaces, and a RESTful API to enable rapid provisioning and configuration of load balancer instances.

Working closely with Kemp Professional Services, ASOS designed and deployed a global load balancing solution that has delivered premium performance and availability during peak shopping periods. Online retail has periods when site visitors and transactions are much higher than normal; for example, in the run-up to the December holiday season or events such as Black Friday and Cyber Monday. During one such event, the ASOS store had 167 million site visits and handled up to 33 orders per second via the Kemp and Azure solution.

In addition, Kemp's Metered Licensing addressed the requirement for scalability and flexibility. Metered Licensing allows ASOS to use and pay for the traffic throughput per month irrespective of the number of load balancer instances that are deployed. This allows the fabric of the network to flex based on traffic needs and not be hampered by the number of licenses available. Metered licensing is ideal for businesses like ASOS that deal with seasonal events and meet the test and development requirements of continuous delivery.

ASOS eliminated the challenges associated with downtime and traffic spikes due to busy shopping periods and retail events. The adoption of the Metered Licensing model also allowed ASOS to only pay for the capacity they needed without having to estimate in advance the number of load balancers and licenses needed.

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 or blog.enterprisemanagement.com. You can also follow EMA on Twitter, Facebook, or LinkedIn.

Corporate Headquarters:
1995 North 57th Court, Suite 120
Boulder, CO 80301
Phone: +1 303.543.9500
www.enterprisemanagement.com
3990.06012020-rev062920KempCS