IDC

SPOTLIGHT

Sponsored by: A10 Networks

Adopting a modern and orchestrated application delivery infrastructure empowers organizations to drive innovation, agility, and governance for successful AI, multicloud strategies, and effective cloud operating models.

Modernizing Application Delivery Infrastructure for AI-Powered Applications

January 2026

Written by: Paul Nicholson, Research Vice President, Cloud and Datacenter Networks

Introduction

Applications have been critical to business success for the past few decades, serving the back office, acting as the digital face of organizations, and driving agility, engagement, and innovation for customers, employees, and partners. Digital experiences are core for any organization, and they are evolving faster than ever before with the rise of AI-powered applications or AI enhancements to existing applications.

Modern applications are increasingly distributed, residing not only in on-premises datacenters but also across public clouds, colocation facilities, and edge environments. Many enterprise workloads will remain on premises, while the rest will proliferate across public cloud environments. Organizations are both multicloud, spanning multiple providers, and hybrid cloud, combining private datacenters with other cloud deployments. For example, when organizations were asked what percentage of their AI/generative AI (GenAI) applications they expected to implement within private on-premises systems versus public cloud provider platforms over the next two years, 56.5% stated cloud providers as their preferred strategic partner and 43.5% stated on-premises locations, according to IDC's 2024 *Worldwide AI in Networking Special Report*. Modern applications will require a diverse set of environments to be supported with interconnectivity, networking, and application-related solutions between clouds and on-premises datacenters.

As organizations advance their hybrid cloud and multicloud capabilities, they seek to run applications where they deliver optimal business value, requiring flexibility in workload placement and avoiding vendor lock-in. Strategic workload placement and architectural decisions are essential for optimizing cloud costs and efficiencies, especially as AI-powered

AT A GLANCE

KEY STATS

- As applications continue to be indispensable to the success of digital business, deploying comprehensive and full-featured application delivery infrastructure grows in importance and value.
- Organizations must consider not only how application delivery infrastructure provides essential network and security services for applications spread across increasingly complex hybrid cloud and multicloud environments but also how well this infrastructure enables the agility and operational efficiencies associated with an effective cloud operating model and demanding new AI workloads and APIs.

applications introduce new demands for scalability, performance, and security. The rise of AI-driven workloads is accelerating the need for robust API load balancing because AI models and agents generate massive volumes of API calls and data exchanges across distributed environments that require optimal infrastructure, full uptime, and low latency.

Organizations are also focusing on rationalizing and reevaluating public and private computing resources and optimizing cloud consumption in the AI era. However, many initial projects have not met expectations. The mean share of AI projects in the past two years that have not delivered measurable business outcomes was 55%, according to IDC's September 2025 *Future Enterprise Resiliency and Spending Survey, Wave 7*.

The ongoing evolution in multicloud application development architectures and, increasingly, AI requirements is driving the need for modern application delivery infrastructure.

While considering AI workloads on public cloud environments for agility and innovation, many enterprises are also evaluating data location and costs and deploying AI applications on premises or in other private cloud environments. Cloud spending can exceed initial budgets, prompting organizations to adopt cross-cloud architectures, centralized management and orchestration, and enhanced observability for hybrid cloud and multicloud workloads. Security across clouds remains a top priority, especially as AI applications expand the attack surface and introduce new risks.

The shift to hybrid cloud operating models, which are data driven, automated, and policy driven, supports multicloud deployment, management, orchestration, operational efficiency, and ubiquitous security. Modern application delivery infrastructure is integral to these models, enabling agile service delivery and optimized digital experiences. The evolution from monolithic to cloud-native architectures, with containers and microservices, further increases application distribution and complexity. IDC observes that single applications may span multiple clouds, with back-end databases in one environment and business logic in another, which is a trend that generative AI and advanced analytics amplify.

In this context, application delivery infrastructure (comprising application delivery controllers [ADCs], load balancers, and integrated security platforms) has become critical for supporting the demands of modern, AI-powered applications. As organizations deploy increasingly distributed workloads across hybrid cloud and multicloud environments, these platforms must deliver not only availability, responsiveness, reliability, and elastic scalability but also advanced API load balancing to handle the surge in AI-driven traffic, directing requests based on tasks, queries, or interactions with large language models. The rise of generative AI and agentic applications is driving exponential growth in API calls and data exchanges, requiring application delivery infrastructure to be centrally orchestrated, software defined, highly automated, and elastically scalable. These solutions must flexibly support all forms of application infrastructure, such as virtual machines (VMs), bare metal, and containers, and operate seamlessly across public and private clouds. Agility features, once exclusive to cloud vendors, are now essential in application delivery infrastructure, including API traffic support, scale-out load balancing, centralized orchestration, and deep analytics and observability for distributed operations.

To provide optimal network and security benefits, application delivery infrastructure should incorporate advances in AIOps and AI-driven observability, which deliver operational efficiency, mitigate human error, and enable rapid troubleshooting and remediation. As organizations manage multiple clouds and disparate application environments, cross-cloud functionality and centralized management become vital for effective operations and resource optimization.

Modern applications, built with cloud-native containers and microservices, further accentuate the need for end-to-end, full-stack, and software-defined delivery and security. Intelligent automation and application-layer services act as connectivity and security front ends, providing encryption offload and dynamic adjustment for Kubernetes clusters as new containers are added or removed.

Controls and functionality can be provided by global server load balancing (GSLB), north-south ingress controllers, and multicluster and intracluster service meshes (for east-west connectivity and application security). The dynamism and diffusion of modern application environments, especially those that AI powers, broaden and deepen the threat landscape. Traditional perimeter security is no longer sufficient; organizations must enhance zero trust initiatives with application-layer security that advanced application delivery infrastructure provides. These platforms can offer a line of defense through mechanisms such as DDoS mitigation, next-generation web application firewall (WAF), and authentication services, ensuring secure, scalable, and resilient operations for distributed, AI-driven digital businesses.

Benefits

The deployment of modern application delivery infrastructure, capable of addressing all the challenges mentioned previously, can deliver the following benefits:

- **Agility with control and without compromise:** Organizations seek to move faster and operate in a more cloudlike manner, but with the discipline and control needed to manage increasingly distributed, AI-powered applications. Modern application delivery infrastructure enables the rapid deployment and scaling of AI workloads and API-driven services while maintaining governance, policy enforcement, and operational discipline across hybrid cloud and multicloud environments.
- **Operational efficiencies, with centralized management and consistent policy:** Centralized management and automation are essential for reducing the complexity of operating across heterogeneous infrastructure and clouds, especially as AI applications drive higher API traffic and distributed workloads. Leveraging AIOps and intelligent automation optimizes staff efficiency, reduces human error, and bridges skills gaps, ensuring a consistent user experience and reliable performance for traditional and AI-driven applications.
- **Better support for business objectives and outcomes:** A platform that supports workload heterogeneity — including AI, cloud-native, and legacy applications — across hybrid cloud and multicloud environments empowers organizations to align IT capabilities with business goals. This flexibility enables optimal workload placement, cost optimization, and the ability to rapidly adapt to changing business needs and AI-driven innovation.
- **Better alignment between developer and operator requirements through role-based access control and application observability:** Role-based access control and advanced observability tools foster collaboration between developers and operators, especially in environments where AI applications and microservices are rapidly deployed and scaled. Enhanced visibility into application performance, API traffic, and security events means that both teams can meet operational and business requirements.
- **Flexibility:** Modern application delivery infrastructure provides flexibility to leverage existing and future workloads, infrastructure, and investments as requirements evolve — whether supporting AI models, API-heavy services, or traditional applications. Flexible licensing, license portability, and consumption models ensure that organizations can deploy application delivery functions wherever needed, across any form factor or cloud.

- **Ability to meet compliance requirements in hybrid environments:** The flexibility to place and support workloads where they deliver the greatest value is crucial for meeting compliance requirements, especially as AI- and data-driven applications introduce new regulatory considerations. Application delivery infrastructure enables organizations to maintain compliance across distributed, hybrid cloud, and multicloud environments.
- **Reduced operational overhead and costs:** Centralized and automated operations, including full API support for customization, reduce operational overhead and costs. This is especially important as AI workloads and distributed applications increase the complexity and scale of infrastructure management, making automation and API-driven orchestration essential for efficiency.
- **Reliable and consistent security controls:** Unified security policies and advanced controls across multiple cloud providers and on-premises environments are vital for protecting distributed, AI-powered applications. Modern application delivery infrastructure delivers consistent security, compliance, and threat mitigation — including DDoS protection, WAF, and bot defense — across all environments, ensuring robust protection and regulatory adherence.

Considering A10 Networks for modern application delivery infrastructure

The A10 Networks product portfolio is designed to help enterprise organizations streamline their application delivery infrastructure across hybrid cloud and multicloud landscapes. The company's solutions are flexible and scalable, enabling applications to remain continuously available, secure, and optimized for cloud operating models and efficiencies.

The application delivery portfolio includes scalable ADCs available in all common form factors: software, VM, bare metal, public cloud, hardware appliance, and containers for cloud-native environments. This choice enables the right solution for the right requirements. For example, A10 offers lower-latency solutions in hardware appliances for latency-sensitive applications, and software appliances are also available for the deployment of other applications on shared infrastructure. Integrated with the ADC portfolio is A10 Control, which manages and orchestrates A10 ADCs across form factors in heterogeneous application environments. This orchestration extends to other A10 solutions, including CGNAT and IPv6 migration, TLS/SSL Insight, Gi Firewall, and A10 Defend DDoS protection. A10 Control brings automated intelligence to the life cycle of application delivery infrastructure, with centralized management, simplifying application and policy deployment across on-premises datacenters, private clouds, and public clouds. It also provides analytics for actionable insights, helping boost operational efficiency and agility.

Granular, real-time, and per-application analytics provide detailed visibility into application traffic — down to the object level — enabling fast root cause analysis and accelerated troubleshooting for IT and DevOps teams. As DevOps teams push out millions of lines of code each year, real-time visibility across all application environments reveals the state of application performance, errors, issues, and latency between components, ensuring the application delivery environment meets stringent business objectives.

A10 Control is currently available as software for on-premises deployment. Organizations can trial A10 Control and have the flexibility to convert to a production system.

Comprehensive application security across application architectures and clouds

Security exploits and vulnerabilities are constant threats to application availability and integrity. A10 offers a secure service stack that can enhance a zero trust strategy by applying multiple levels of verification and control to each

connection. A10 ADC security controls include TLS/SSL offload, a next-generation web application firewall powered by Fastly, authentication and access management, CAPTCHA user verification, Layer 7 DDoS protection, a datacenter firewall, and other application security capabilities. The A10 Threat Intelligence Service, available by subscription, leverages reputation data to proactively block known bad actors, DDoS attacks, and other threats across multiple configurable categories.

A10 Control provides centralized intelligence for A10's application security. It orchestrates and manages a consistent security policy for multiple security components — covering multiple form factors or clouds — across the application delivery infrastructure. It also provides centralized intelligence for geosensing and geolocation data, helping detect anomalies proactively and minimizing issues that could impair application availability.

Emphasis on simplified operations

Support for an enhanced user experience helps simplify and speed the work of ITOps while improving engagement and satisfaction for application users. To maintain application availability and experience, A10's portfolio includes GSLB, which provides global traffic management to support application resilience and continuity through faster, localized server responses to end-user requests. Cloud redundancy and redirection also support user experience and application continuity, providing proactive resilience in the event of cloud or network outages, prohibitive cloud overage costs, or other unforeseen circumstances. As organizations' hybrid cloud strategies evolve to include multiple public clouds and on-premises environments functioning as a unified system, cloud selection and optimization tools become essential, extending beyond traditional failover and disaster recovery use cases.

A flexible, inclusive licensing model

FlexPool offers capacity pooling that allows organizations to purchase licensed capacity on demand and allocate it as needed across hybrid cloud and multicloud locations. This is especially useful when deployment requirements are not fully known or are subject to change. Benefits include:

- A portable, consumption-based software model
- Quick customer adoption and license mobility
- An inclusive feature set — no additional licenses needed

Other A10 solution features to address ADC trends include:

- A holistic, secure application services approach that treats on-premises datacenters and clouds as one system, enabling greater simplicity, agility, consistent manageability, and security and supporting an effective and ubiquitous cloud operating model, including high-performance TLS/SSL (with support for cipher suites such as ECC)
- Multicloud and hybrid cloud optimization, redundancy, and selection
- Advanced customization with a TCL-based scripting language (A10 calls this aFleX) for application transformation, allowing for the removal of barriers so custom applications can move to the cloud when basic load-balancing capabilities do not provide the scripting that the application requires to modify application traffic behavior
- Flexibility in control, centralized management, and GUI, CLI, or API with full CLI parity

- Advanced DNS load balancing to improve security and optimize performance, reduce back-end DNS server requirements, and improve overall utilization; includes recursive DNS offload, DNS application firewall capabilities, and DNS over HTTPS (DoH) or DNS over TLS (DoT)

A10 acquired ThreatX in early 2025 and plans to further integrate its web application and API protection (WAAP) functionality with the A10 ADC portfolio, which is a logically adjacent enhancement, especially with the focus on the increased API demands that AI deployments are driving.

Challenges

A major enterprise challenge also impacting suppliers and partners is that many organizations have yet to articulate or implement an effective cloud operating model that includes a consistent and unified approach to critical elements such as application delivery infrastructure, especially for managing AI traffic and applications. In practice, this often results in the use of disparate application delivery solutions from multiple vendors and cloud providers, without a consolidated strategy for hybrid cloud and multicloud environments. This fragmented approach leads to operational and cost inefficiencies, increased complexity in handling high-volume AI-driven API traffic, and heightened security vulnerabilities. The inertia of siloed environments and traditional IT models continues to hinder many organizations, making it difficult for them to adapt to the demands of distributed, AI-powered workloads.

Another challenge is the emergence of new buying centers and influencers in cloud-native and AI-driven environments. Cloud architects, platform teams, and AI specialists are increasingly familiar with cloud-native load balancing, service meshes, and open source solutions. Suppliers of application delivery infrastructure must engage these stakeholders meaningfully, addressing their specific requirements for managing distributed AI applications, API load balancing, and seamless integration across hybrid cloud and multicloud environments. In addition, A10 must keep pace with rapidly evolving AI and encryption standards and requirements — such as Model Context Protocol (MCP), post-quantum cryptography, and other emerging protocols and traffic pattern shifts — to ensure secure, compliant, and future-proof application delivery infrastructure.

Conclusion

As organizations face growing complexity and unprecedented distribution in their application environments — including the surge in AI traffic and applications — application delivery infrastructure with advanced local and global load balancing and robust security functionality becomes a critical enabler of digital business success and resiliency. Managing AI-driven workloads alongside traditional applications demands infrastructure that can seamlessly orchestrate and elastically scale across hybrid cloud and multicloud environments, ensuring consistency and predictability for VMs, bare metal, containers, and public cloud deployments.

Modern application delivery infrastructure is essential for supporting AIOps capabilities and enabling an effective cloud operating model, especially as AI applications drive higher API volumes and dynamic traffic patterns. If A10 Networks succeeds in addressing these evolving challenges, it will be well positioned to deliver enterprise customers a comprehensive, modern, and ubiquitous application delivery infrastructure capable of managing AI and traditional workloads alike and delivering the full spectrum of benefits outlined in this paper.

About the analyst

Paul Nicholson, *Research Vice President, Cloud and Datacenter Networks*

Paul Nicholson is IDC's research vice president for Cloud and Datacenter Networks. He provides thought leadership and actionable insights on cloud and datacenter networking markets and technologies. Paul has a deep understanding of the networking market along with its business and application requirements, technologies, product road maps, competitive differentiation, and go-to-market strategies, enabling him to provide informed guidance for vendors, cloud providers, enterprise IT buyers, and practitioners.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
One Beacon Street
Suite 33100
Boston, MA 02108, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com