IDC

SPOTLIGHT

Sponsored by: Tetrate

Market demand for cloud-native application networking is growing and will be broad based and ubiquitous in the coming years. The time to modernize network infrastructure and operations for Kubernetes and cloud-native environments is now.

Meeting the Need for Modern Application Networking

June 2022

Written by: Brad Casemore, Research Vice President, Datacenter and Multicloud Networks

Introduction

As organizations advance to the next stage of their digital transformation journeys, they encounter an array of daunting challenges. Many of these challenges involve the need to implement a truly digital infrastructure that can accommodate an increasingly distributed application landscape that extends across heterogeneous application environments, including on-premises datacenters and clouds.

Indeed, hybrid and multicloud challenges are not limited to where applications run. They extend to what applications run on and how well they are supported and delivered by underlying infrastructure. It is here that the rise of cloud-native applications, based on containers and microservices and typically orchestrated by Kubernetes, not only offers tremendous potential but also introduces a range of requirements that cannot be accommodated by traditional infrastructure and operating models.

AT A GLANCE

KEY TAKEAWAYS

Benefits from using a comprehensive and full-featured platform for application layer networking include:

- » Business agility and flexibility
- » Faster time to market for applications and services
- » Greater digital resilience (proactive rather than reactive)
- » Abstracted but comprehensive connectivity and security

Among the infrastructure challenges, those associated with networking are particularly vexing in the context of distributed, heterogeneous, and hybrid environments, inclusive of a growing wave of cloud-native applications. In fact, IDC finds that networking is cited as one of the more problematic areas of Kubernetes implementation and operation. The challenges are compounded as enterprises seek to provide agile, flexible, simple, and secure networking across hybrid/multicloud environments.

To be sure, networking for modern cloud applications is a different proposition from what was required in the client/server era and virtualized monolithic applications. There's no question that networking for Kubernetes can be complex and difficult, replete with notable deviations from classic networking assumptions and constructs.

Containers change not only how applications are developed but also how applications connect across the network. Most focus on containers has been on orchestration and, thus, on Kubernetes, understandably so given its importance to application developers. Nonetheless, networking for modern applications, which can be termed application networking, is critical for production deployments, particularly considering that a Kubernetes network must be automated, elastically scalable, and secure.

Kubernetes abstracts complexity from application developers, providing a clean plug-in interface that eliminates distractions, but complexity remains — displaced now to the operations team, which must ensure that the network continually meets developer expectations. Without the appropriate application networking functionality, Kubernetes deployments lack the connectivity, elastic scalability, and security they require.

The need has never been greater for networking to be closely aligned with developer processes and workflows.

Let's consider the following four distinct facets of Kubernetes application networking:

- **Container-to-container networking.** Containers for networking are in the same pod, analogous to running containers in the host network of a virtual machine (VM) (through the vSwitch).
- **Pod-to-pod networking.** In Kubernetes clusters, pods within a node communicate with each other via IP addresses, with each pod having a dedicated IP address.
- **Pod-to-service networking.** Both containers and pods are ephemeral, meaning IP addresses need to generate and retire continually. An abstraction layer is used to manage and assign IP addressing.
- **Internet-to-service networking.** This type of networking is necessary for access to the internet from a cluster or to make a cluster's services reachable from the internet. For access to the internet from the cluster, egress services are required; for access to the cluster from the internet, ingress services are required. An ingress network in Kubernetes requires an ingress controller, which is a L4–L7 proxy that carries traffic into and out of Kubernetes clusters and services. As such, ingress controllers accommodate ingress-egress, or north-south, traffic flows. Ingress controllers can also be used for visibility and troubleshooting, as well as for security and identity requirements, and they can also address many API gateway use cases. Some organizations choose to combine ingress controller and API gateway functionality (as with the Envoy Gateway project), especially for application networking use cases, while others deploy the functions separately. The benefits of combining ingress controller and API gateway functionality, and of taking a more holistic approach to Kubernetes application networking, can include consistent and simplified operations as well as faster and more effective troubleshooting.

As the preceding discussion illustrates, the need has never been greater for networking to be closely aligned with developer processes and workflows in cloud-native environments. This is because microservices connectivity focuses on application networking for microservices rather than the classic device-based approach to networking, which focused on Layer 2 and Layer 3 of the protocol stack and is associated with traditional application environments and infrastructure.

These modern application environments demand that consideration be focused resolutely on the application and its immediate infrastructure, including not only containers but also VMs and bare metal as well as on-premises deployments and public clouds. It also compels organizations to consider how the application environment will evolve in the foreseeable future, likely encompassing greater adoption of cloud-native applications and clouds.

Security is another key consideration. There are new security implications to address with adoption of Kubernetes. Maintaining seamless connectivity and pervasive security is difficult in across such a distributed, dynamic environment. Security capabilities and protection must be offered against a range of potential exposure, including insider threats and man-in-the-middle attacks. Zero trust policy controls and comprehensive observability are essential.

An added challenge is that neither agility nor security can be compromised or sacrificed. Distributed cloud-native environments must be agile and secure, with inherent security within a cluster and across clusters, clouds, and datacenters. Across this spectrum of workloads and clouds, security must be a primary consideration, going beyond the limitations of conventional perimeter security so that all applications can be protected and secured, constantly and consistently.

Simplicity of operation is also critical because many organizations lack cloud and cloud-native expertise. This cloud-native expertise is often lacking most noticeably in infrastructure and operations teams, which are frequently unfamiliar with technologies such as service mesh. It's here that popular open source projects, including service meshes such as Istio, now under the auspices of the Cloud Native Computing Foundation (CNCF), can deliver not only steadily improving simplicity but also community-based benefits such as knowledge transfer, continuous feature enhancements, and the choice and flexibility deriving from a vibrant ecosystem.

Beyond Service Mesh: Ingress and Egress

While service mesh is an integral component of cloud-native networking, more is needed for enterprises implementing a comprehensive approach to networking for their modernizing application environments. This is because of the requirement to bridge networking across legacy and modern environments and to provide a seamless connectively framework for the former to migrate to the latter over time.

This is where connectivity elements such as API gateways and ingress controllers play an essential role, providing north-south secure networking for inbound traffic. Similarly, egress gateways are required to control outbound traffic. Istio provides both ingress and egress functionality for ingress and egress in Kubernetes environments.

What is needed, then, is an end-to-end application networking platform in which zero trust security is ubiquitous. All of this has to be simple for application developers and operators. The network cannot get in the way of developers and their applications. Instead, it must provide seamless support and continuous availability while remaining unobtrusive and unseen.

Looking further ahead, IDC believes intelligent automation, leveraging real-time telemetry and AI, will make modern application networking more agile, flexible, and secure. Consequently, application networking will become increasingly proactive, providing operators and architects with prescriptive recommendations for traffic management and workload protection and gradually gaining operators' trust to enforce prescriptive policy autonomously.

Benefits

Various benefits accrue from deploying and using a comprehensive and full-featured platform for application layer networking.

For enterprises, the benefits include the following:

- **Business agility and flexibility.** As the modern network becomes as agile and flexible as the developer processes that it supports, the network is no longer a drag on how rapidly applications, the lifeblood of digital business, can advance through the development cycle and move into production.

- **Faster time to market for applications and services.** As the modern network becomes more agile and flexible, it accelerates time to market for new and updated applications and services. It also provides the adaptability and flexibility to support CI/CD processes and ongoing updates.
- **Greater digital resilience (proactive rather than reactive).** With the modern application network in place, organizations achieve much greater digital resilience, especially through observability features that allow a more proactive, predictive, and protective approach to network availability and security.

As the modern network becomes more agile and flexible, it accelerates time to market for new and updated applications and services.

For application developers, the benefits are equally compelling:

- **Abstracted but comprehensive connectivity and security.** As a modernized application layer network provides agile and flexible support for applications without getting in the way, it enables developers to focus on their applications rather than on managing infrastructure, an inconvenience that few developers wish to countenance. In addition, self-service features provide developers with automatically provisioned application infrastructure.

For cloud architects, platform teams, and other modern infrastructure operators, benefits include:

- **Network architecture and capabilities for agile, flexible, reliable, scalable, and secure support for cloud-native and other applications across a heterogeneous infrastructure and clouds (hybrid and multicloud).** A modernization application network provides a simple but robust means of accelerating provisioning of the services that developers and lines of business need. Beyond provisioning, a modern application network enables faster and more effective troubleshooting and remediation of connectivity issues that affect application performance and digital experience.
- **Security and protection for the applications and services at the heart of digital business.** Fully integrated and ubiquitous network security means that architects and DevSecOps teams can achieve the isolation and security required for modern applications.
- **Ability to extend networking wherever it needs to be, across clusters and clouds.** The ubiquity of the network, replete with elastic scale and flexibility, means that it can evolve and adapt to application and business requirements, irrespective of where workloads must run to produce optimal outcomes.

Considering Tetrate

Tetrate Service Bridge (TSB) is an Istio- and Envoy-based application connectivity platform that provides enterprises with a consistent, unified way to connect and secure services across an entire mesh-managed environment, inclusive of the service mesh itself and ingress and egress gateways.

TSB sits at the application edge, at cluster ingress, and between workloads in Kubernetes clusters and traditional compute clusters. Edge and ingress gateways route and load balance application traffic across clusters and clouds, while the mesh controls connectivity between services. Istio is leveraged for the control plane, and Envoy is used at the data plane.

Use cases for Tetrate Service Bridge span cloud migration, zero trust application isolation, security (including user and application authentication, authorization, encryption, and rate limiting), multiplatform support across heterogeneous infrastructure and clouds, and API mediation.

TSB is designed to provide a range of benefits, including business agility, application security and isolation, and business continuity through support for digital resilience.

As a platform, TSB offers a single, centralized management plane for configuration of connectivity, security, and observability across the entire application network. Functionality includes support for multicluster and multicloud environments as well as for multitenancy; connectivity for modern and traditional applications, including those running virtualized or on bare metal; and streamlined acceleration of application-centric operations. The latter capability is designed so that developers can configure APIs declaratively, allowing the platform to apply and enforce intent across the infrastructure.

Also offered is full life-cycle management for both Istio and Envoy. Configuration safeguards are delivered in the form of validations of Istio configurations and service-level isolation and organizational controls designed to ensure that only correct configurations are invoked at runtime.

At the application edge, TSB offers L7 load balancing across one or more ingress gateways in different clusters over Istio-controlled mTLS. Application ingress is provided through load balancing for service meshes in Kubernetes clusters or traditional workloads. Meanwhile, integrated API gateway functionality is designed to allow TSB to resolve differences between north-south and east-west application traffic flows. As such, traffic is managed at the application edge, at application ingress, and between services traversing the data plane on service mesh sidecar proxies.

TSB's Istio-based service mesh is equipped with security features, traffic management capabilities, and optional FIPS-compatible builds to meet compliance requirements. An array of application layer observability features is also provided.

In addition to a version of the Tetrate Service Bridge that customers deploy and manage themselves, Tetrate Cloud is an option for customers that want a fully managed, cloud-delivered version of TSB.

The Tetrate Istio Subscription, a life-cycle management CLI tool that ensures the use of trusted versions of Istio, is also available. It ensures that upstream Istio distributions work well on the underlying Kubernetes platform, and it has been tested and verified on Kubernetes distributions such as AWS' EKS, Google Cloud's GKE, and Azure's AKS.

Tetrate's founders have been, and continue to be, significant contributors to Envoy and Istio, and they have built considerable experience in working with and understanding Istio and Envoy distributions and use cases. Aside from being a mainstay contributor to Istio, Tetrate recently announced that it is on the steering committee of Envoy Gateway, a new effort within the Envoy proxy open source project to simplify Envoy use in cloud-native application development. Like Istio and Envoy, Envoy Gateway is under the auspices of the CNCF.

Challenges

A number of enterprises and other large organizations are still relatively new to Kubernetes and cloud-native application architecture. Many of those organizations have Kubernetes skills gaps that cannot be addressed easily or cost effectively. Some fail to fully appreciate the infrastructure and connectivity requirements that accompany adoption of Kubernetes. While many have heard of Istio, they might not be all that familiar with it. For Tetrate, the challenge will be simplifying Istio and Envoy so that organizations with these skills gaps can more easily deploy it and use it with assurance throughout the life cycle.

Another challenge is the degree of competitiveness in the service mesh and API gateway/ingress controller realm. Open source service mesh projects Istio and Linkerd are both under the auspices of the CNCF, and a significant number of vendors are active in the market, including major IaaS cloud providers AWS, Microsoft Azure, and Google Cloud.

Conclusion

Greater enterprise emphasis on application networking, adapted to the particular needs of containers and microservices in Kubernetes-orchestrated environments, is an inevitable outcome of modern application architectures and the resultant rise of truly digital infrastructure and operations. The market demand for cloud-native application networking is growing and will be broad based and ubiquitous in ensuing years.

Modern applications require connectivity and networking that are better aligned with developer needs, processes, and CI/CD workflows, capable of rising to the business imperatives of agility and flexibility. At the same time, security and control must not be sacrificed, and enterprises, with varying degrees of cloud-native expertise and skills within their organizations, will also require comprehensive solutions that address the need for simplicity through ease of consumption, deployment, and ongoing operation.

Tetrate Service Bridge has been designed to address the various needs and requirements that confront enterprises as they move to modernize their network infrastructure and operations for Kubernetes and cloud-native environments. If Tetrate can meet and surmount the challenges cited in this paper, IDC believes it will be well placed to provide complete Istio- and Envoy-based ingress, service mesh, and egress offerings to a broad and deep enterprise market. It will also have established a foundation for future enhancements, including increased automated intelligence for even simpler, more agile, and more secure application networking.

About the Analyst

Brad Casemore, *Research Vice President, Datacenter and Multicloud Networks*

Brad Casemore is IDC's Research Vice President, Datacenter and Multicloud Networks. He covers datacenter network hardware, software, IaaS cloud-delivered network services, and related technologies, including hybrid and multicloud networking software, services, and transit networks. Mr. Casemore also works closely with IDC's Enterprise Networking, Server, Storage, Cloud, and Security research analysts to assess the impact of emerging IT and converged and hyperconverged infrastructure.

MESSAGE FROM THE SPONSOR

More About Tetrate

Started by Istio founders to reimagine application networking, Tetrate is an enterprise service mesh company managing the complexity of modern, hybrid cloud application infrastructure. Its flagship product, Tetrate Service Bridge, provides an edge-to-workload application connectivity platform to deliver business continuity, agility, and security for enterprises on the journey from traditional monoliths to the cloud. Customers get consistent, baked-in observability, runtime security, and traffic management in any environment. Tetrate remains a top contributor to the open-source projects Istio and Envoy Proxy. Find out more at www.tetrate.io.

IDC Custom Solutions

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

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com