

451

Research®

PATHFINDER REPORT

The New Datacenter Network

COMMISSIONED BY




AUGUST 2019

©COPYRIGHT 2019 451 RESEARCH. ALL RIGHTS RESERVED.

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

ABOUT THE AUTHOR



MIKE FRATTO

SENIOR ANALYST, APPLIED
INFRASTRUCTURE AND DEV OPS

Mike Fratto is a senior analyst on 451 Research's Applied Infrastructure and DevOps team covering enterprise networking. He has extensive experience reviewing and writing about enterprise remote access, security and network infrastructure products, as well as consulting with enterprise IT, equipment and software vendors, and service providers.


Introduction

As application architectures change, so does the datacenter, and with it the requirements for datacenter networking. Cloud-native architectures and distributed applications hosted in the datacenter, cloud service, colocation and edge are driving IT to provide fast and reliable networking in a diversity of environments. Demands from users and application architects are pushing for faster provisioning, rapid changes and timely problem resolution that carries significant business impact if not met. New technologies are coming to market to address these demands, including streaming telemetry for performance management; AI and machine learning to make sense of the telemetry data; automation and orchestration and intent-based networking to speed changes reliably; and higher-speed networking to keep pace with demand.

The data from a recent Voice of the Enterprise survey (see Figure 1) not only highlights what is top of mind for enterprise IT, but can almost be used as a progression, starting with containers as another virtualization technology and ending with fully embracing software-defined infrastructure via service meshes.

Figure 1: Most important cloud-native technology or methodology

Source: 451 Research's Voice of the Enterprise: DevOps 2019



Cloud-Native is Complicated

Cloud-native applications are the new normal. 451 Research defines cloud-native as ‘applications designed from the ground up to take advantage of cloud computing architectures and automated environments, leveraging API-driven provisioning, auto-scaling and other operational functions. Cloud-native applications are not limited to cloud applications – we see cloud-native technologies and practices present in on-premises environments in the enterprise.’¹

While cloud-native applications are meant to be movable from one runtime environment to another with minimal changes, or run in multiple environments simultaneously – aka, multi-cloud – the network capabilities have largely been confined to particular environments, forming a patchwork of functions in each location that IT has to understand, configure and optimize individually for each service and environment. The most common denominators are used among all the environments in an effort to streamline operations, which ignores advanced capabilities in each environment that can be beneficial. In the end, the application is on a suboptimal network because supporting the variety of capabilities consistently becomes overly complex, or IT is performing more manual tasks to bridge the capability gaps in each environment, which is inefficient and slows IT’s speed in response to changing demands. A more effective strategy involves unifying the networking at the virtual and management layer, reducing operational overhead and keeping errors from cropping up while taking advantage of advanced features.

If Applications are Distributed, so is the Network

Computing architectures are distributed – whether it is an IoT deployment where some processing occurs on the device or on a local controller at a remote site, a cloud-native application that relies on connections between micro-services on-premises and in cloud services, or a service provider moving workloads out to the network edge to satisfy demands for low-latency processing close to the customer. This often leads to a distributed network that was once contained in a datacenter and now needs to be interconnected into a reliable, manageable network.

Distributed networking isn’t new, but supporting distributed, dynamic applications means capabilities, addressing schemes, and networking services need to be malleable and configurable to where the workloads are to avoid network complexity, lower operational overhead, and enable automated and orchestrated actions across network technologies and locations. This represents one of the transformation hurdles IT has to get over – embracing and managing dynamism in the environment that is simpler and predictable when technical capabilities are consistent everywhere. Another transformation hurdle for IT is how to effectively join the network management domains with a consistent set of capabilities across multiple environments such as on-premises datacenters and one or more cloud services.

1. From ‘The 451 Take on cloud-native: truly transformative for enterprise IT’


Two approaches are adopted. The first is to deploy virtual networking software – virtual routers and switches, for example – where hardware can't be used. Using virtualized products means many of the software capabilities are available in all environments and can be managed seamlessly. But the downside is that more VMs have to be deployed, adding to the licensing cost for the software and VM to run them, increasing management overhead due to complex configurations in cloud services, and adding to the potential bottleneck virtualized networking imposes.

The alternative to virtual network software is abstracting the native cloud services' networking APIs into the existing management infrastructure so that the cloud networks are managed like on-premises switches but are using the cloud services' native networking capabilities. The benefit of using native APIs and capabilities is reliable networking in cloud services with seamless management and operations, but capability gaps will remain, causing inconsistent policy deployments.

Applications are being deployed in a wide variety of locations and environments, which is complicating every facet of IT management (see Figure 2). IT would do well to standardize as much of the application environments as possible to control the chaos.

Figure 2: Cloud or hosted services in use

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2019
Q: Which of the following types of cloud or hosted services, if any, does your organization currently use?



Faster, Always Faster

The integration of developers and operations, DevOps, is a direct result of increasing demands from the business to deploy and change workloads faster than ever before, and for the business to drive those deployments as they see fit. Meeting that demand changes IT's focus from building environments and deploying workloads to automating and orchestrating the building and deployment of those environments and workloads and the networks that support them. This involves making the decisions ahead of time on how workloads are deployed and how networks are built so that they can be deployed by people without requiring in-depth knowledge of those initiating the action. The same holds true for changes taking place on Day 1, Day 2, Day 3 and beyond.

Modern network hardware and software, including management systems and other supporting software such as controllers, performance management and security products, support a variety of integration capabilities like APIs and SDKs and allow IT operations teams to intelligently automate management workflows. The time spent in automating workflows pays off in faster changes with fewer mistakes. Building intelligent workflows capable of gracefully handling errors and exceptions requires the expertise of developers who understand flow control and exception handling, as well as network operations staff who understand how networks can fail and can develop policies defining the proper reactions. This is DevOps in action.

Technology to the Rescue

In support of the diversity of application environments and the business's unrelenting desire for agility, a number of technologies are available to enterprises. Streaming telemetry is a foundational element for data collection and monitoring. Unlike sampled collection methods, streaming telemetry provides data at a steady pace, allowing the identification of transient events such as micro-bursts, which may be missed using other collection methods. Streaming telemetry will be a critical data source used in automation, performance management, and artificial intelligence and machine learning.

AI and ML is already being developed to identify patterns in network traffic and equipment behavior that enables more responsive monitoring and reaction, identifying patterns that would be difficult manually, and enabling experienced professionals to understand dynamic environments and develop widely applicable optimizations. Telemetry, automation and AI/ML lead to closed-loop automation and intent-based networking that adapts and reacts to changes in the environment as defined by IT.

Adapting the Datacenter Network to New Demands


Everything IT does in the datacenter is to support applications whether the applications are on-premises, in a cloud service or hosted at the network edge. As modern application architectures are adopted, identifying an application component and relating it to other components is the first step in applying network and security policies. Where the previous generation of applications could be manually tagged and managed, cloud-native applications are by nature far more dynamic and will require greater use of automation throughout the application's lifecycle. For example, the development and deployment system needs to assign the appropriate tags to applications, which are then used by supporting infrastructure such as an SDN to look up and apply the correct network and security policies. Application identification from traffic analysis is an alternative where application tags aren't defined – and is useful for verifying the application is what it claims – but it's less authoritative than tags applied at deployment.

According to our Voice of the Enterprise: Digital Pulse, Workloads and Key Projects survey, the top three networking priorities are performance-related (see Figure 3). It's the most visible area in which IT can show value for any workload in any environment.

Figure 3: Most important networking-related goals

Source: Voice of the Enterprise: Digital Pulse, Workloads and Key Projects 2019

Q: What are the most important networking-related goals for your organization's IT environment this year?



This example where developers provide vital identifying information about an application, and IT uses it to apply the appropriate set of policies reflects the kind of collaboration taking place in IT departments that have adopted DevOps into their culture. The increased use of automation will also reduce the time to provisioning networking resources. In our 2019 Voice of the Enterprise: Servers and Converged Infrastructure, Budgets and Outlook, only 14% of respondents indicated it took them minutes to provision on-premises networking; slightly more than a quarter of respondents indicated it took hours, days and weeks (26% for each) to provision resources, indicating there is ample opportunity to significantly reduce provisioning time. As your organization embarks or continues its digital transformation, it will also move from mostly manual to mostly automated processes. Enterprises are well underway to automation, with 53% of respondents from the same study indicating they are using automation with manual exception-handling.

Policy Unlocks Potential

Taking a policy-first approach is the first step in streamlining network operations and embarking on automation. While there are many definitions of policy in network IT processes, policy-based management starts with defining the configuration requirements necessary to achieve intended performance and security goals. The policy states IT's intent in an actionable manner. It is then implemented in network hardware and software. In a manual process, the intent is codified in IT's runbooks. In semi- or fully automated processes, the policy is codified in the automation system that forms the backbone of an intent-based network.

That is the benefit of policy-based management. The goals and requirements are predefined, and the implementation details are left up to network IT's capable hands. The policy can be limited to one domain like the datacenter or cloud service, or, as is more likely the case, the policy can be applied across domains, such as spanning datacenters and clouds to branch offices, users, IoT devices or services deployed at the edge.

Policies that are applied end to end are more effective in ensuring the goals are met wherever workloads are placed and consumed – such as improving application performance, an important goal shown in Figure 3 – than policies that are domain- or location-specific. Additionally, applying policies streamlines operations because they can be applied uniformly network-wide. Using a policy to define the development of automated processes simplifies DevOps work by defining the goals that have to be achieved and providing a framework to automate within.

Unified Segmentation

The most common reason to segment networks is for security, such as limiting access to servers to authorized users only. The benefits of segmenting applications extend beyond security, however. A foundational element of containerized applications and microservices is that consumers of the application are only aware of the front-end process they talk to. The application servers and supporting services are tucked away behind the front end to communicate among themselves. Segmenting applications greatly simplifies IT operations, monitoring and troubleshooting because the application components are confined to a well-defined segment. Equally important, networking for the application can be configured as a unit allowing the configuration requirements to move with the application as IT deploys on-premises or in a cloud service.

A robust orchestration platform paired with strong IT operations workflows will be able to adapt to the different network capabilities in each environment. Using virtual network software in the target location means the feature set is more likely to be consistent, but only for the traffic passing through it, and may require additional configuration of the native cloud service network to maintain full isolation. Using the cloud service's native APIs and capabilities ensures network segmentation is maintained in the cloud environment, but there may be gaps in the features and capabilities use that will need to be understood.

AI and ML in Action

AI/ML is poised to become a critical component in ongoing IT operations. AI/ML starts with data – lots of data – then the ML algorithms use the data streamed from network equipment and applications to train the models that will be used to surface insights about the network infrastructure and automate IT operations. We expect AI operations to augment, not replace, IT staff by being able to continually analyze very large datasets quickly and discover anomalies, perform analysis leading to a cause, and generate recommendations for configuration changes or further troubleshooting.

ML, in particular, is already being used in closed-loop automation where application behaviors are assessed and compared to an intended outcome, and automatic actions such as adding new compute and network capacity are added and removed. As applications move to micro-service architectures, the number of compute nodes and the server connections between them will explode in number, making ongoing troubleshooting and root cause analysis extremely complicated without AI-assisted tools to perform data gathering and analysis leading to actionable insights and recommendations. AI-generated analytics can create accurate forecasts and predict when and how to scale (up or out) your IT infrastructure to meet demand before performance suffers. These are practical examples of how AI is helping IT manage its resources more effectively and deliver good applications to the business with high quality of experience.

Figure 4: Critical criteria of network operations features

Source: 451 Research

FEATURE	CRITICAL CRITERIA
REAL-TIME TELEMETRY AND PERFORMANCE DATA	With ever-increasing performance demands, the more rapid the data generation is, the faster automated systems react or alerts can be raised. Generating telemetry data via dedicated ASICs is becoming the norm and has no impact on system performance.
IN-DEPTH TELEMETRY AND PERFORMANCE DATA	Network telemetry indicates streamed, real-time data that describes the state of the underlying hardware and software and can be used to identify anomalies such as micro-bursts and resource exhaustion that is missed by other data-collection routines.
MACHINE LEARNING	Algorithms that mine data looking for related patterns, ML can be performed on the device or against a data repository. Supervised ML is when a person influences the machine learning, whereas unsupervised ML is emergent.
CLOSED-LOOP AUTOMATION	This occurs when the automation system takes data from the network, selects an appropriate configuration change, implements the change, verifies that the change was properly made, and measures the result.

FEATURE	CRITICAL CRITERIA
BUFFER AND QUEUE MANAGEMENT	Buffers and queues store network traffic that needs to be sent but can't. Buffer and queue management tends to be first in, first out, but modern networking gear supporting demanding applications at 100Gbps and 400Gbps utilizes smarter traffic prioritization.
APPLICATION PROGRAMMING INTERFACES	APIs are a way for developers to call functions on another system to send a command or retrieve information. Language-agnostic, the API can be called using any supported protocol, typically HTTP/HTTPS, and can be easily integrated into the network application. Vendors should not make frequent changes to how APIs are called nor deprecate APIs without ample time for developers to adapt.
SOFTWARE DEVELOPMENT KITS	SDKs are software libraries that are imported into applications and provide a way for developers to integrate with systems. Unlike APIs, SDKs are language-specific, and updates on the target systems may require updating the SDKs, which can be disruptive to IT operations. APIs should be the preferred integration method over SDKs.

Reaching Across Vendors

One of the biggest challenges enterprises face is automating workflows in multi-vendor environments. It's common for enterprises to deploy products from multiple vendors, such as using one vendor for routing and switching, a vendor for application delivery controllers, a vendor for network security, a vendor for IP address management and so forth. Configuring the network for an application or during an application move or a scaling event means automating the changes across all of these systems.

APIs and SDKs from networking vendors along with other development tools are critical for integrators and DevOps teams to automate processes across the entire network seamlessly and reliably. Support for standardized protocols and models such as NETCONF and YANG eases integration due to broad support and reduces the amount of effort developers spend on integration. Having existing support for the products an enterprise uses is important, but equally important is the speed and ease in which new products can be added and updated. Popular products are likely to be supported first, but products with lower market share or custom software will have to be added as needed. Enterprises will also want to speed up modernization steps to enable better automation and integration capabilities in the network.


Consistent Networking Everywhere

Datacenter networking has gotten very complex as IT has become tasked with supporting physical, virtual and container networks on-premises, in colocation datacenters, and one or more cloud services, each with its own management interfaces, features and capabilities. Complicating matters are demands from the business to deploy faster, respond quickly to change requests, and to proactively maintain the network, reducing downtime and the impact of outages.

Workload lifecycle management is getting more complex with management systems popping up like mushrooms, none of which integrate, and is unsustainable for IT. Unifying management is very important for all facets for IT, including networking on-premises, in colocation datacenters, in cloud services, and with physical, virtual and container environments (see Figure 5).

Figure 5: Importance of unified IT management system

Source: 451 Research's Voice of the Enterprise: Digital Pulse, Workloads and Key Projects 2019



Container environments are particularly difficult because the container management systems want to take over all aspects of provisioning, including the networking within the container environment, or hand those duties off to a service mesh that is responsible for interconnection, service identification, authorization and load balancing. Yet, service meshes are well suited for that role. The lines of responsibility are blurring between network IT and others as the ways applications are deployed and managed diversify.

One way to gain control of network environment diversity is to abstract the differences via a unified management and intelligent orchestration system that can identify and respond to capability gaps where they reside. Chasing consistency across the entirety of network functions with the same product set is impossible because new environments are introduced by stakeholders outside of network IT and need the unique capabilities in those systems. Unifying at the management and orchestration layer allows IT to use unique capabilities in the physical, virtual, cloud and container networks and manage them as a consistent entity.

Conclusion

While SDN continues to make strides in simplifying network operations and improving network performance in the datacenter, the diversity of application environments is splintering the network at a time when the business and other parts of IT want faster changes. Technologies such as streaming telemetry, AI/ML, automation and orchestration, and better traffic management schemes paired with faster networking are coming to market to help reduce operational overhead, foster integration, promote workflow automation and improve application performance.



Liberate your business by taking your network to where the data is.

Critical Components:

- Automation for every stage of the IT lifecycle
- Consistent policy model for scaling multicloud networks
- Pervasive security across the entire network

www.cisco.com/go/dcnetworking

PATHFINDER | THE NEW DATACENTER NETWORK

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

© 2019 451 Research, LLC and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with 451 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. 451 Research disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although 451 Research may discuss legal issues related to the information technology business, 451 Research does not provide legal advice or services and their research should not be construed or used as such.

451 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.



NEW YORK

Chrysler Building
405 Lexington Avenue,
9th Floor
New York, NY 10174
+1 212 505 3030



SAN FRANCISCO

505 Montgomery Street,
Suite 1052
San Francisco, CA 94111
+1 212 505 3030



LONDON

Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 203 929 5700



BOSTON

75-101 Federal Street
Boston, MA 02110
+1 617 598 7200