Juniper Networks | Engineering Simplicity

White Paper

FROM CLOUD TO MULTICLOUD

A Guide to Battling Complexity in Enterprise Network Design and Operation

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Without specific and explicit design decisions about not only the network architecture but also network operations, enterprises risk drifting toward a future of many clouds, each with their own operational nuances. Those who are not mindful about managing new additions increase the complexity of their network with each variation.

This white paper discusses how to approach the design of the network and operations in the move from cloud to multicloud. The focus is on meeting the business needs of today and tomorrow, battling complexity, preserving the possibilities for options, and supporting IT's role as the steward of technology decisions for the business.

Introduction

Multicloud is about connecting and securing applications from the endpoint to the edge and every cloud in between, as simply as if they were in one cloud. This approach supports the launch of any workload on any cloud across a multivendor environment. Ultimately, multicloud lets IT teams manage resources as a single, cohesive infrastructure with consistent security and operations across all the places in the network. The move to multicloud is about operational transformation, and that requires a different way of thinking.

Enterprises that design multicloud independent of any one provider—either service or technology—maintain control over key decisions that impact their business. This means choosing components that are not only insertable and manageable in the operational model, but also replaceable. Further, enterprises will want to consider their design from end to end—from the data center to the campus, branch, and cloud—as well as from top to bottom—from the connectivity to the orchestration, visibility, and security of the infrastructure.

For most enterprises, the transition to multicloud will be accomplished in steps, with each taken through a natural technology refresh. A framework to stage the path and coordinate decisions across initiatives will assure steady progress specific to the priorities of the business while avoiding decisions that close off future options.

Cloud Becomes Multicloud

As enterprises continue their steady march to the cloud, it's important to note that this evolution is merely an incremental step toward the ultimate destination: multicloud. Whether it's the pursuit of a dual vendor strategy, differentiated capabilities, proximity to users or data, or the realization that old and new must coexist for a time, most enterprises will leverage more than one cloud offering. Therefore, it is critical that enterprises start thinking about their eventual multicloud destination, even as they undertake the smaller cloud movements today.

Agility, Not Cost

Admittedly, the industry dialogue around cloud and multicloud has lacked precision at times. As a result, many businesses are confused about why they should embrace the cloud at all.

There is a prevailing belief that the move to the cloud is primarily about reducing costs, based on the theory that owned infrastructure is more expensive to procure and manage. While it is true that large cloud properties can lower costs based on how they design, procure, and operate their infrastructure, it is somewhat misguided to assume that costs will always decline for enterprises adopting cloud.

In fact, IT leaders who base their decisions on short-term financial gains are unintentionally putting themselves in a precarious position. In the early years of cloud adoption, new infrastructure coexists with the old, meaning teams will need to float their operating expenses as they straddle the present and the future. As a result, at least in the early years, costs nearly always go up. For some workloads, especially persistent ones, costs may even be higher running in the cloud, leading some to launch new initiatives to "uncloud" their workloads as traffic increases.

Does this mean that enterprises should scuttle their cloud endeavors? Of course not. The key to success is having a clear and precise view of why cloud is transformative. Enterprises are not leasing equipment so much as they are leveraging cloud operations. In other words, the move to cloud is about agility, not cost. If teams can deploy software and services more quickly, they put themselves in a better position to deal with the pace of change occurring all around them. It is this macro trend of change that poses the real existential threat to enterprises.

Agility Means Operations

There is a subtle but critical implication of making agility the leading driver for go-forward architectural decisions: it means that the impending transformation will necessarily impose more operational changes than a mere technology evolution would. The future is about moving quickly, which is more of a statement about the workflows that drive the business than about the devices that move the traffic.

So what prevents networking teams from moving fast in their current environments? In a word, complexity. The reliance on pinpoint control over large expanses of distributed devices administered primarily through the command line has led to a networking-wide practice that is rife with complexity and fragility. In an effort to filter out failure, it's no surprise that the dominant means of coping with the frailty of modern networks is to enforce draconian change controls.

If enterprises could afford to evolve at a glacial pace, such an approach would be fine. But change is everywhere, and the pace of change continues to evolve. So as enterprises begin adopting cloud and multicloud architectures, operational agility is the single most important objective.

Invisible Infrastructure

Cloud and multicloud are about delivering infrastructure that is essentially invisible to the user. Ultimately, the IT infrastructure should resemble the power grid: when people flip a light switch, they don't think about the infrastructure that makes it possible—they just expect light. For the cloud, the key to this "invisibleness" will be applications. When users need something, they will not think about the underlying IT infrastructure. Likewise, when developers want to create and launch a new application, they will have tools to set up the network quickly on their own.

To make the promise of an invisible infrastructure possible, there are basic principles that must be followed:

- **Security**: It is almost a cliché to say that security is a key requirement of any enterprise, but security is a fundamental consideration of all aspects of IT. If users, workloads, and data cannot come together in a secure infrastructure, how can that infrastructure fade to the background?

  In the context of multicloud, security means a few things. Different infrastructure domains must be connected by secure transport to allow for the geographical separation of users, data, and applications. Fundamentally, this means that security must extend end-to-end, accounting for both the application workloads and the cloud on-ramps over which those workloads are accessed.

  Of course, once end-to-end security is realized, there is the issue of policy management. Policies must be managed centrally and applied uniformly, ensuring that traffic is secure and treated equally regardless of where the user or workload resides. Deep network visibility and multiple points of enforcement throughout an organization's multicloud footprint are required.

- **Ubiquity**: If multicloud brings IT infrastructure closer to utility status, then applications and services must be available everywhere. If users are aware of the underlying infrastructure because their experience differs based on where they are, or the workload is, then the infrastructure is not invisible.

  In a multicloud world, resources and services can be physical or virtual, and they can reside anywhere. The factors that determine where workloads actually get executed are physics and economics, that is to say, performance and cost requirements.

This all needs to be transparent to the user, who does not care where the workload is being run. In fact, if the user can tell the difference, then multicloud has failed to deliver on its promise.

Ubiquity places an operational requirement on multicloud. Workload life cycle management tools like Kubernetes and OpenShift need to integrate with the underlying infrastructure to allow for dynamic workload management. If this integration is intended to unify all layers of the technology—from application to physical devices—it must be considered as part of any architectural decision that impacts the infrastructure.

- **Reliability**: Nothing makes a user more painfully aware of the underlying infrastructure than a service outage. The power grid is hardly noticed until it's unavailable, because that's when entire cities grind to a halt.

  The only way multicloud can satisfy this always-on requirement is if problems can be identified and resolved through the use of automation. This places the burden on monitoring and visibility. The systems that manage workloads need end-to-end visibility extending from the end user to the application, regardless of where either is located.

  Remediation requires operational controls to extend end-to-end as well. Automation enables reliability, reliability enables predictability, and predictability is required for security. Therefore, automation cannot be domain-specific; otherwise, remediation becomes manual, and the end user is once again aware of the underlying infrastructure.

- **Fungibility**: Underpinning all of these requirements is the concept of fungibility. For any of this to be true, resources in a multicloud environment must be fungible, meaning pools of resources need to be essentially interchangeable.

  There will always be constraints—for latency-sensitive workloads, proximity is important, and for cost optimization, some clouds might fare better than others. However, within a class of resources, there can't be any limitations on where workloads can run or what services are available. As soon as limitations are imposed, the user is once again reminded that the infrastructure matters.

  When resources are not fungible, it also means there are constraints on how applications and services are delivered. Pinning a workload to a specific resource pool is decidedly not multicloud.

One Cohesive Set of Resources

The natural conclusion of multicloud is that the enterprise must be managed as a cohesive set of resources. For the infrastructure to be invisible, the domains that enterprises have historically used to break complex networks into smaller, more manageable bits must go away—along with the boundaries between them.

It is impossible to deliver multicloud by relying on traditional design approaches that use separation and containment to manage complexity. For each established boundary, there is overhead—one might call it a crossing tax—for the people, systems, and processes that need to navigate between contexts. Operational drift becomes commonplace and change risk-prone, often taking a lengthy negotiation between the teams on either side of the boundary.

End-to-End

To meet the outlined principles and manage multicloud as a cohesive set of resources, architectures will have to extend beyond just the data center. Yes, the cloud exists in data centers, but it also exists at the edge, especially as multi-access edge computing (MEC) and Internet of Things (IoT) take root. Plus, users need an on-ramp to the cloud, which means that security and policy management must extend end-to-end, from network ingress through the application resources and back. This means that a multicloud architecture will touch campus and branch networks, and the security and automation constructs that make it all work will need to span all places in the network.

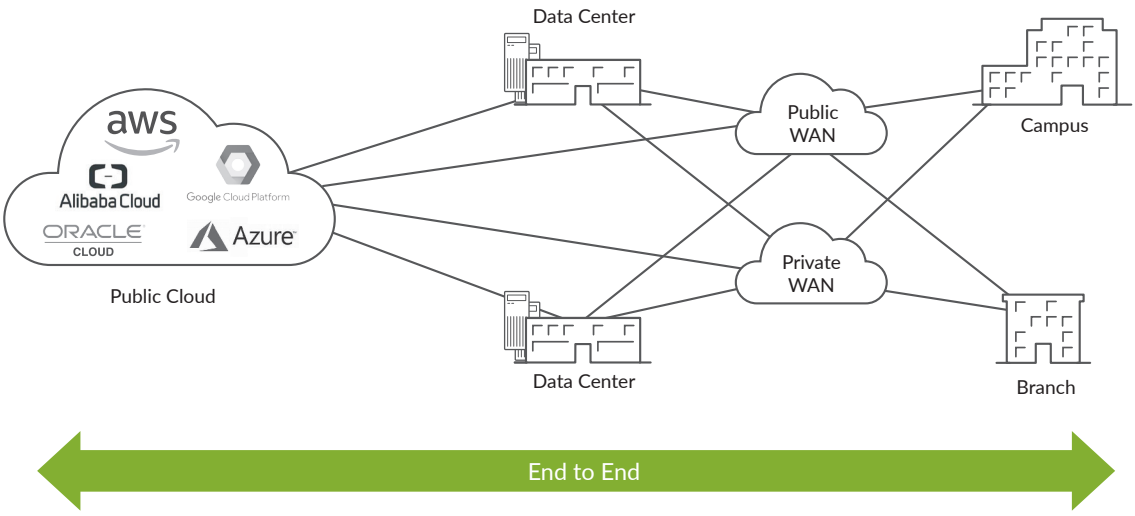

Figure 1: End-to-end multicloud

Top-to-Bottom

Multicloud, however, is more than just connectivity. While the underlying connectivity pieces across all the places in the network are necessary, simply allowing packets to flow is not sufficient.

Managing an expansive infrastructure as a cohesive unit requires end-to-end orchestration as a means of using policy to dictate experience, and network services to enforce the experience. End-to-end visibility is required if operations are to be automated, and end-to-end security is needed to ensure that users, applications, and data are protected.

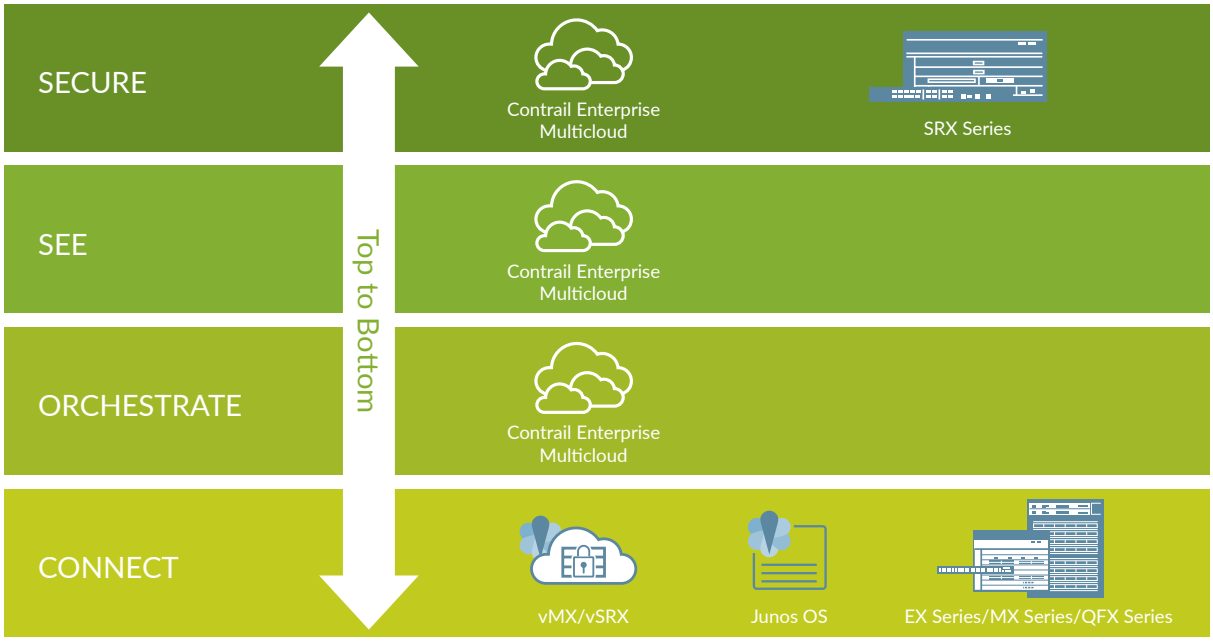

Figure 2: Top-to-bottom functionality in the multicloud

Multicloud vs. Multiple Clouds

It is important to draw a distinction between multicloud and multiple clouds. Many companies begin their journey by adopting Software as a Service (SaaS) offerings like Office365, Salesforce, or Workday as their initial baby steps to the cloud. If this first foray was successful, these companies look to take further advantages of the cloud by turning to more complex applications.

For others, moving to the cloud is an exercise in lift-and-shift. Applications are hosted in AWS, Azure, or GCP rather than within a private data center. As companies explore the operational nuances of various clouds, they might favor one or another for specific applications, typically due to economic or performance requirements.

The promise of multicloud, however, is not about simply fracturing the IT environment into cloud-specific shards that encompass infrastructure and operations for bounded domains. Multicloud is about securely and consistently managing resources—regardless of whether they reside in a private or public cloud—as a single, cohesive infrastructure.

Multicloud Must Be Multivendor

While the ultimate promise of cloud requires infrastructure to be invisible so that where a workload resides or where users log in doesn't matter, the realities of legacy means that enterprises run the most fragmented and complex networks.

Moving to multicloud cannot be about discarding what already exists. Therefore, multicloud must be multivendor—not only because most businesses will ultimately want to use multiple public clouds, but also to unify operations with what exists today. Multicloud must, in its design, adhere to the concept of "open," taking both existing and future infrastructure into account to ensure interoperability and programmability while avoiding unnecessary lock-in.

A design principle made famous by Amazon is helpful when assessing each decision:

"Some decisions are consequential and irreversible or nearly irreversible—one-way doors—and these decisions must be made methodically, carefully, slowly, with great deliberation and consultation. If you walk through and don't like what you see on the other side, you can't get back to where you were before."

With each change, whatever the immediate needs, choosing multicloud-ready solutions will help businesses move forward without sending them through one of these one-way doors.

Migration to Secure and Automated Multicloud

Ultimately, multicloud will represent a journey for enterprises. It is more than a shrink-wrapped product or even a single project; it requires the thoughtful coordination of many decisions across multiple years. Evolution will naturally encompass changes to overarching architecture, to the products and tools that underpin that architecture, and to the people and processes that manage it all. Something like the 5-step general model shown in Figure 3 is helpful when considering the strategy for change.

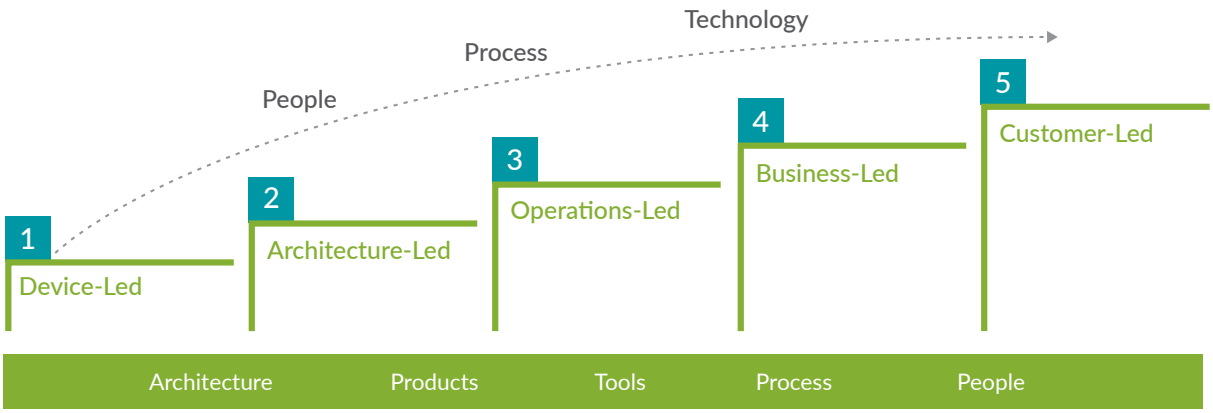

Figure 3: The 5-Step general model

At the highest levels, the model serves as a general compass for companies evolving from predominantly device-led to architecture- and operations-led IT, to ultimately business- and customer-led. At the company level, this is a useful tool for framing the discussion with individual teams so that they eventually converge around a multicloud infrastructure.

Still, migrations are not about general movements. They require specificity. They must be grounded in tangible actions built on capable technology. There will be specific paths for each place in the network: data center, campus, branch, and public cloud. Enterprises might evolve different parts of the infrastructure at different times, but the thoughtful culmination of those efforts will eventually lead to a full multicloud future.

Conclusion

As the industry moves towards its multicloud future, complexity has become the number one enemy for all enterprises. While complexity has been an anchor until now, as enterprises cope with managing the operational changes associated with multicloud, complexity represents an existential threat. Those that cannot tame the beast will find their futures extremely challenging.

This means that enterprises have to use every refresh or expansion opportunity to continue their journey to multicloud and simplify their overall operating environment. The journey will be more than just technological; it requires enterprises to evolve their architectures, processes, and people; and it demands a different approach to planning, procuring, deploying, and managing infrastructure.

Changes of this magnitude are extraordinarily rare—they might happen once or twice in a generation. While such changes pose a threat to companies ill-prepared for the transition, they also represent tremendous opportunity for those who capitalize. The future favors the agile, and multicloud.

About Juniper Networks

Juniper Networks brings simplicity to networking with products, solutions and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily. At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable and secure networks to move at the speed of business.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.0.207.125.700
Fax: +31.0.207.125.701

JUNIPER NETWORKS | Engineering Simplicity