



MULTICLOUD TECHNICAL GUIDE FOR NETWORK AND CLOUD ARCHITECTS

*How to connect and secure workloads across many
clouds and operate as if they were in one cloud*

TABLE OF CONTENTS

Introduction: What Is Multicloud?	4
Why Multicloud?	4
The Challenges of Multicloud	5
Multidomain Connectivity.....	5
Multivendor Orchestration	5
End-to-End Visibility.....	6
Pervasive Security.....	6
Complexity	6
End to End and Top to Bottom.....	7
Multicloud Architectural Overview.....	8
Foundational Resources	9
Fabrics	9
IP Fabrics	10
Intent-Based Fabric Operations	12
Unifying Data Center and Campus Fabrics	12
Workload Management.....	13
Instance-Specific Application Environments	13
Workload Abstraction.....	13
Overlay Control.....	15
Unified Underlay and Overlay Management	15
Services Consumption	15
Application Development Life Cycle.....	16
Service-Aware Networking.....	16
Service-Aware Constructs.....	18
Conclusion	19
About Juniper Networks	19

LIST OF FIGURES

Figure 1: End-to-end and top-to-bottom multicloud design.....	8
Figure 2: Multicloud architectural overview	8
Figure 3: Common fabric for multicloud	9
Figure 4: EVPN-VXLAN architecture.....	10
Figure 5: Unified management for multicloud	11
Figure 6: Federated fabric.....	12
Figure 7: Contrail architecture	14
Figure 8: Policy and control	15
Figure 9: Simplified policy enforcement.....	16
Figure 10: Advanced analytics	17
Figure 11: Cross-layer visibility and analytics	18
Figure 12: Service-aware constructs	18

EXECUTIVE SUMMARY

Organizations seeking to connect and secure workloads end to end across many clouds as simply as if they were in one cloud are driving the transition from cloud to multicloud. This move represents a significant operational transformation, and such changes require careful thought and planning. The entire industry benefits from a complete discussion of the key points that comprise that transition.

This white paper discusses how to approach network design, including security and operations, as organizations transition from cloud to multicloud. It begins by clarifying the reasons for change, moves to bounding the problem space, and then discusses key architectural decision points along the way. The emphasis is on meeting the technical needs of today and tomorrow, battling complexity, building operational excellence, preserving options, and supporting the architect's role as the steward of technology for the business.

Introduction: What Is Multicloud?

Multicloud is the management of disparate infrastructure domains as if they were a single cohesive set of resources, regardless of where those resources reside, enabling users to consume those resources from anywhere and in any environment. In this regard, although multicloud is concerned with the underlay transport, it is built primarily around higher layer functionality. Fundamentally, multicloud is more an operational condition than a description of the core infrastructure.

Because multicloud requires a unified management approach, a pointed distinction must be made between “multiple clouds” and “multicloud.” Simply breaking the IT environment into cloud-specific shards separated by bounded domains of infrastructure and operations does not achieve the promise of multicloud. The distinguishing characteristic—multicloud's true benefit—is the ability to operate across both private and public clouds, on and off premises, as a single entity.

Why Multicloud?

In order to understand the desire for multicloud, it's important to first discuss why businesses would adopt cloud technology in the first place. The desire to move from legacy infrastructure to cloud—whether public or private—is commonly driven by either cost or agility.

For companies that find maintaining legacy infrastructure prohibitively expensive, the thought of fulfilling technology needs through a provider designed for that purpose is attractive since it is assumed costs will be lower, allowing the business to focus on core competencies that drive meaningful value.

Yet focusing primarily on cost overlooks the main benefit of cloud. Enterprises are not outsourcing infrastructure so much as they are leveraging the agility of cloud operations. If enterprises can deploy software and services more quickly, they put themselves in a better position to deal with the rate of change all around them. And it is the macro trend of change that poses the real threat to the enterprise business.

While most companies start with modest cloud ambitions, for many—even those of moderate size—this effort will eventually lead to multicloud. Whether it's the pursuit of a dual-vendor strategy, the need to support differentiated capabilities, or the realization that old and new technologies must peacefully coexist (at least temporarily), most enterprises will find themselves supporting a mix of multiple private and public clouds.

To realize the full promise of cloud, enterprises must think about their ultimate destination even as they undertake smaller cloud movements today. Table 1 details why today's cloud deployments will ultimately evolve into multicloud.

Table 1: Why Clouds Will Evolve Into Multicloud

Reason	Why Multicloud?
Economics	Choosing among different cloud providers creates economic leverage. Further, cloud tourism is a growing phenomenon in which enterprises experimenting with cloud suffer sticker shock from largely uncoordinated deployments and/or high costs incurred by persistent applications, causing them to take a fresh look at the balance between private and public infrastructure.
Capabilities	As enterprises develop different needs, they are likely to discover differences in cloud offerings. For example, Amazon AWS competes on breadth and flexibility of service offerings, Microsoft Azure on enterprise applications, Google Cloud on machine learning, and Oracle on ERP software.
Availability	Distributing applications offers greater availability and resiliency to cloud service failures. While this diversification happens across clouds, it also inevitably happens within a single provider footprint as enterprises use multiple cloud accounts, regions, and availability zones.
Data Privacy	The need to monitor where data originated, especially for global enterprises, will favor different local cloud providers so that workloads can be collocated with the data they service.
Proximity	Where performance is important, organizations may choose to run a workload in different clouds. For instance, local clouds will offer shorter round-trip times for traffic, while the rise of multi-access edge computing (MEC) might lead to cloud instances running on premises or in a distributed telco cloud.
Transition	The decision to adopt multicloud might be driven by a need to complete the transition from private to public, in which case enterprises will have to bridge both architectures for an undetermined amount of time.

The Challenges of Multicloud

Enterprises moving to multicloud share a common set of challenges that inhibit adoption. This section will focus on these obstacles, helping architects recognize potential problem areas and understand the constraints under which they must operate in order to plot their path forward when designing a multicloud future.

Multidomain Connectivity

While it is true that multicloud is primarily a solution to an operational problem, there is still an underlying need to seamlessly connect and secure islands of resources. Since multicloud is also about controlling the end-to-end infrastructure, this connectivity must also, at the very least, span data center, public cloud, and campus and branch gateways, with full campus and branch integration a longer term goal.

Spanning such a broad set of network locations, both on premises and in the cloud, requires a range of physical and virtual form factors and capabilities. For example, managing traffic over a high-capacity WAN typically requires custom silicon, while data center switches leverage common merchant silicon and branch boxes use x86 CPUs.

A common fabric management platform will also ease network virtualization with seamless connectivity across domains. The virtualized overlay network must span different types of servers as well as virtual/physical networking and security devices so that workloads or users are not exposed to differences in domains. Regardless of the form factor, these devices must continuously communicate with northbound software layers like orchestration, visibility, and security, placing architectural requirements on the underlying connectivity elements, especially around programmability and telemetry.

Multivendor Orchestration

Since a range of form factors and capabilities are required, multicloud deployments will be multivendor, posing a challenge for enterprises looking to unify infrastructure. Providing a common orchestration layer that sits atop a heterogeneous underlay has historically proven difficult; multivendor element management systems (EMS), for instance, have a woeful track record.

Simply put, today's predominant operational model, with its dependency on pinpoint control and manual device-by-device management, will not survive the move to multicloud. In the multicloud, application releases and where application workloads run can change from moment to moment, demanding a network that will adapt just as quickly. As such, multicloud is a significant departure from today's operational model, where the measurement of change is more often in months. Architects must be prepared to design around the new operational practices of applications supporting the digital business.

End-to-End Visibility

If multicloud is predicated on treating a mix of infrastructure as a single entity, then visibility cannot be limited to individual domains. For many existing deployments, unifying visibility poses a significant technological obstacle.

Further complicating the transition is the fact that multicloud will be highly automated. Since the basic premise of automation is to see something, do something, this puts a premium on centralizing monitoring and visibility and expanding the surface area of what can be observed and acted upon in a consistent way. In multicloud, the objective should be to reduce the need for users to discover problems and to relieve operations teams of the need to perform tedious, time-consuming tasks across domain boundaries to determine the root cause when a problem occurs.

When environments are isolated, it is difficult to correlate events across domain boundaries, which dramatically reduces the reach of any operational controls. If telemetry is different across disparate clouds, it must be normalized and translated before it can be used for multicloud purposes. Real-time analytics with intent-based monitoring and alarms for flagging issues, along with data-driven capacity planning to understand where resources need to be spun up or turned down, must be a foundational requirement of any multicloud design.

Pervasive Security

While networking has historically treated security as a perimeter issue, today it has clearly extended beyond the edge. To manage attacks that occur inside the network, security must look at traffic both within and between disparate network subsets. It's no longer enough to put firewall filters and ACLs between network segments; we must be able to isolate traffic streams within those segments, a practice commonly known as microsegmentation.

For multicloud to operate in a cohesive way, consistent protection schemes must be applied across the entire infrastructure, creating a more secure posture for traffic within the campus, across branches, between clouds, and inside the data center.

Policy and control must exist at both the tenant and the application level, and it must be managed in a unified way so that the operational burden of distributing and updating security policy does not render the solution operationally unviable. This requires a single point of management and a single point of end-to-end validation to ensure regulation compliance.

Complexity

Perhaps the biggest challenge to operating a multicloud is complexity. In today's environment, complexity is already so prevalent and debilitating, networks are extremely fragile and intolerant of change; without a good foundational design, networks operating across multiple clouds will add to this complexity, exacerbating the problem.

Unfortunately, complexity is inevitable as a function of the number of devices, users, applications, tools, and so on—a condition of the given ecosystem that cannot be totally eliminated. Good design and automation, however, can offload much of that complexity from everyday operational tasks.

To achieve this, all aspects of an ecosystem must be designed with simplicity in mind. One way to do this is to remove pieces, abstract functionality, and automate workflows (see Table 2). Consistent topologies and common policies, along with automation to ensure proactive management and ease troubleshooting, are essential. Such an approach results in reliable, simplified solutions across the entire multicloud environment.

Try Contrail and Contrail Insights

You can [test Contrail Enterprise Multicloud](#), a cloud automation platform, and Contrail Insights, a cloud operations optimization tool, in a safe, cloud-hosted environment—for free. Visit <http://juniper.net/try>.

Table 2: Taking on Complexity

Action	How to Start
Remove	<p>Cloud companies have innovated on the operations side by reducing the overall number of network elements through judicious design. For example, they have:</p> <ol style="list-style-type: none"> 1. Removed the complexity of diverse architectures by adopting a single architecture across the infrastructure that can scale elastically and accommodate growth. 2. Removed many protocols from their devices, settling on very few (less than 10) which are driven by software automation and control. 3. Standardized on connectivity options and protocols for consistent deployment processes across environments and workloads. 4. Refreshed network devices more frequently, rather than using them for several years, allowing everything to stay current and creating operational consistency. <p>These same design approaches will help enterprises simplify their multicloud design and operation.</p>
Abstract	<p>To make today's networks work, operators must have expertise across multiple devices. If the network includes a mix of vendors, operators must learn the precise language, syntax, and nuances of each.</p> <p>While this approach has been refined over the last few decades, with incremental changes and advancements, it has ultimately resulted in complex, fragile, and unpredictable networks, as evidenced by the draconian controls imposed any time a change is necessary.</p> <p>The idea of abstracted control allows operators to manage the entire network as a single entity, rather than one device or fabric at a time. Abstraction provides one place for operators to execute a workflow, get information, handle integrations, and insert policy.</p> <p>When considering abstraction in a multicloud design, an architect should focus on:</p> <ol style="list-style-type: none"> 1. Identifying the points of abstraction and how these will interpret and convey intent to the network. 2. Enabling operators and developers to interact with the points in language intuitive to their role. 3. Translating the points of control into all represented devices and commands. <p>Given that the scope of multicloud implies heterogeneous environments, how architects plan for abstraction is an essential part of the overall design.</p>
Automate	<p>IT operators should consider how automation, expected to increase agility and improve efficiency, can also enhance reliability. It's not realistic to hope that a network outage won't occur; the best strategy is to actively plan for outages by automating network operations to ensure services are always available.</p> <p>Architects can reframe the automation conversation to capture its full benefits in good design decisions:</p> <ol style="list-style-type: none"> 1. Many teams don't know what they want to automate. The primary goal of automation shouldn't be to eliminate keystrokes; it should be to improve how things work. Delays and faults happen at the boundaries of things: people, systems, organizations. It is within these workflow handoffs, defined by people and processes, where automation will have the greatest impact. 2. The key to automation is getting data from one entity to another. That means the currency of automation is data. Automation designs must consider how to exchange and ensure a common understanding of data. 3. After determining the high-value workflows and methods of data exchange, organizations must determine how to set up and use the automation tools. No matter the provisioning altitude of intent and the presence of autonomous sensors, logic, and actuators in the system, humans will be driving change and providing insight about the system state for data-driven decision making and service-level observability.

End to End and Top to Bottom

When discussing multicloud, architects must be careful not to narrow the design scope and risk exacerbating the challenges detailed above.

Since both users and applications can be anywhere, a true multicloud design must include a complete end-to-end perspective to accommodate full communication. Users, for instance, can access applications from a campus, remote branch, home office, or public space, while applications may be deployed in on-premises data centers, in various public clouds, or dynamically moved to edge compute as well as serverless environments where they may only live for a short period of time. Because applications must be available to all users, regardless of location, the multicloud architecture should touch all places in the network.

Architects must also consider more than just connectivity between users and applications. Simply allowing packets to flow is not enough; the multicloud also requires networking functions to extend from top to bottom in order to secure everything, monitor performance, and orchestrate policy. Managing a multicloud environment in a centralized way requires end-to-end orchestration to implement policies and enable automation, as well as end-to-end visibility to understand where resources, users, and applications reside. End-to-end security is also required to protect users, applications, and data. Figure 1 depicts a full end-to-end, top-to-bottom multicloud design.

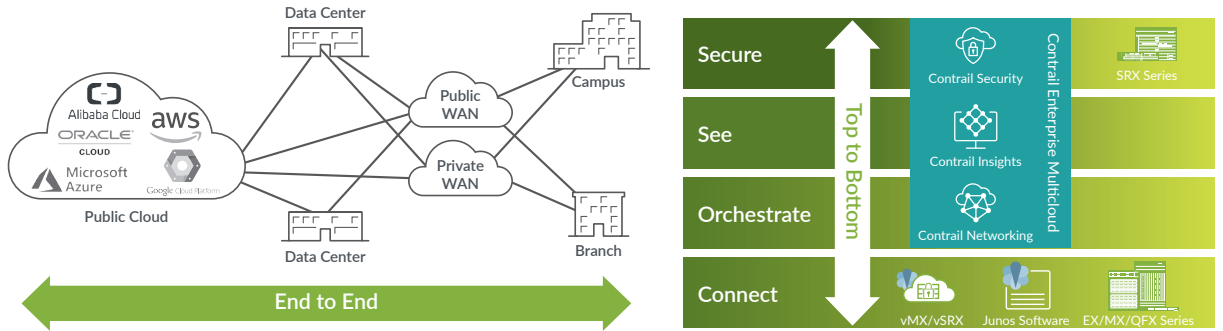


Figure 1: End-to-end and top-to-bottom multicloud design

Multicloud Architectural Overview

Taking the scope and potential challenges of multicloud design into consideration, the primary building blocks for multicloud can be broken into three layers (see Figure 2):

- **Foundational Resources:** The underlying compute, storage, network, and security elements that form the foundation for any workload infrastructure.
- **Workload Management:** The workload constructs such as virtual machines (VMs), containers, and serverless architectures, as well as broader workload life-cycle management frameworks like OpenStack, Kubernetes, OpenShift, and various public clouds.
- **Service Consumption:** Ultimately, users consume services, typically through applications. This layer decouples infrastructure and services by transparently abstracting the foundational and workload management layers into a set of services for each application.

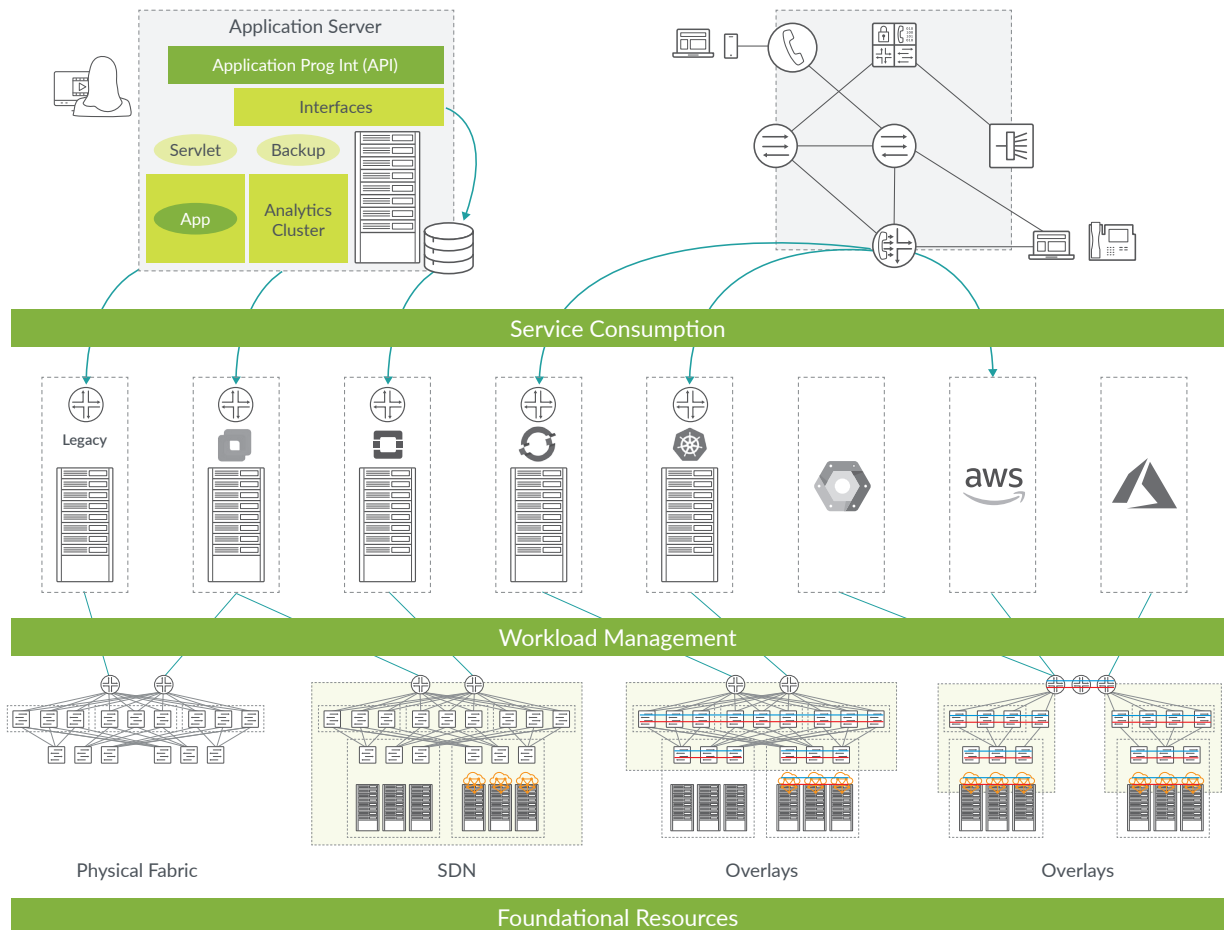


Figure 2: Multicloud architectural overview

Foundational Resources

From a networking perspective, the foundational resources are primarily responsible for transport. Tenants, users, and applications get logical overlay connectivity across a shared underlay infrastructure. This layer also includes major management constructs such as SDN controllers (including overlay managers) to manage it all.

Since the underlying network must span the data center, cloud, and anywhere applications and users reside, it includes a mix of physical and virtual devices deployed both on premises and in the public cloud. The need for a common orchestration layer demands that the underlay devices have standards-based, programmable interfaces; a mix of closed, proprietary systems would simply lead to more complexity.

Fabrics

Multicloud strives to unite underlay, overlay, and public clouds as a single, common fabric that spans the entire infrastructure to leverage a unified operational framework (see Figure 3). This federated fabric is created by using a common set of protocols in order to combine disparate infrastructures and enable both hardware and software—including virtual devices and cloud-native infrastructure—to seamlessly coexist.

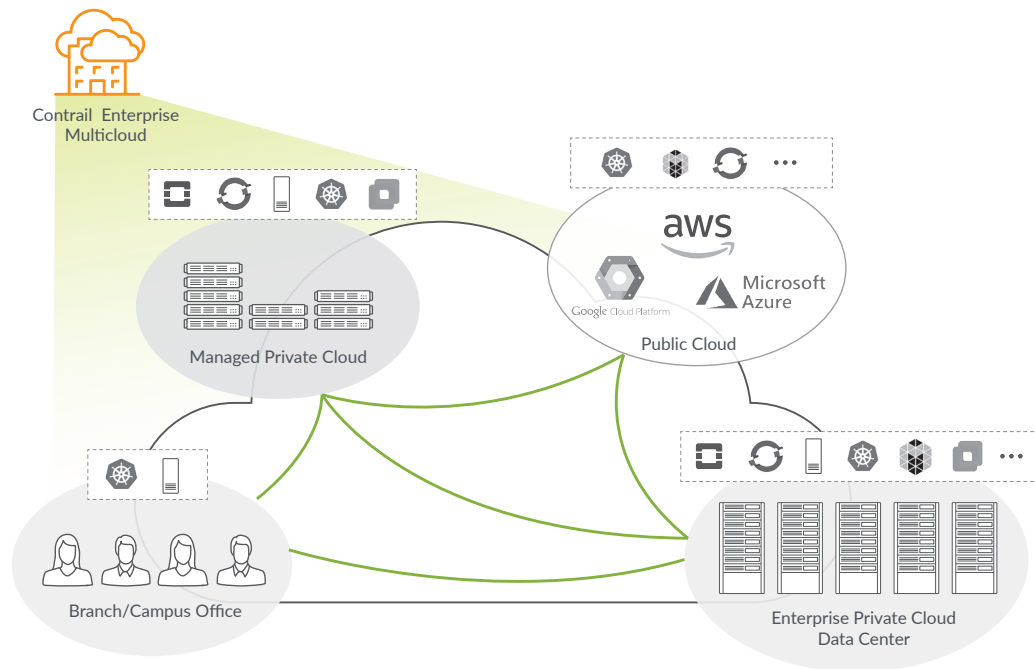


Figure 3: Common fabric for multicloud

Over the last few years, data center networks have standardized on the leaf-and-spine architecture. This Clos network type has layers of switching infrastructure connected in a nonblocking mode to simplify the network and reduce latency for east-west traffic.

Network teams use this architecture to build two types of fabrics: Ethernet fabrics that operate at Layer 2, and IP fabrics that operate at Layer 3. As more and more applications are designed to operate at L3, many organizations are choosing to operate their underlay network as an IP fabric. Designing the network as a simple IP fabric underlay with an orchestrated overlay supports the growing need for a more dynamic network in order to keep up with the frequent changes triggered by new application requirements.

To realize these benefits, the hyperscalers (AWS, Azure, Google Cloud, and so on) build their networks as an IP fabric underlay with a logical overlay, and many enterprises are also now adopting this design. Further, the approach enables network construction with standard protocols, thus preserving the possibility for multiple vendors in the network, either as preferred vendor practice or allowing for times of vendor transition, to reduce vendor lock-in.

IP Fabrics

Operating the network as an IP fabric offers direct L3 connectivity to applications along with the benefits of greater scale and resiliency. When L2 is required, Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) transforms the IP fabric to support end-to-end Ethernet connectivity. The standard for Ethernet connectivity over IP fabrics is EVPN-VXLAN (see Figure 4), which allows networking teams to build fabrics without relying on proprietary protocols. This also makes the deployment of multivendor networks possible, which in turn enables enterprises to evolve without unnecessarily stranding existing assets.

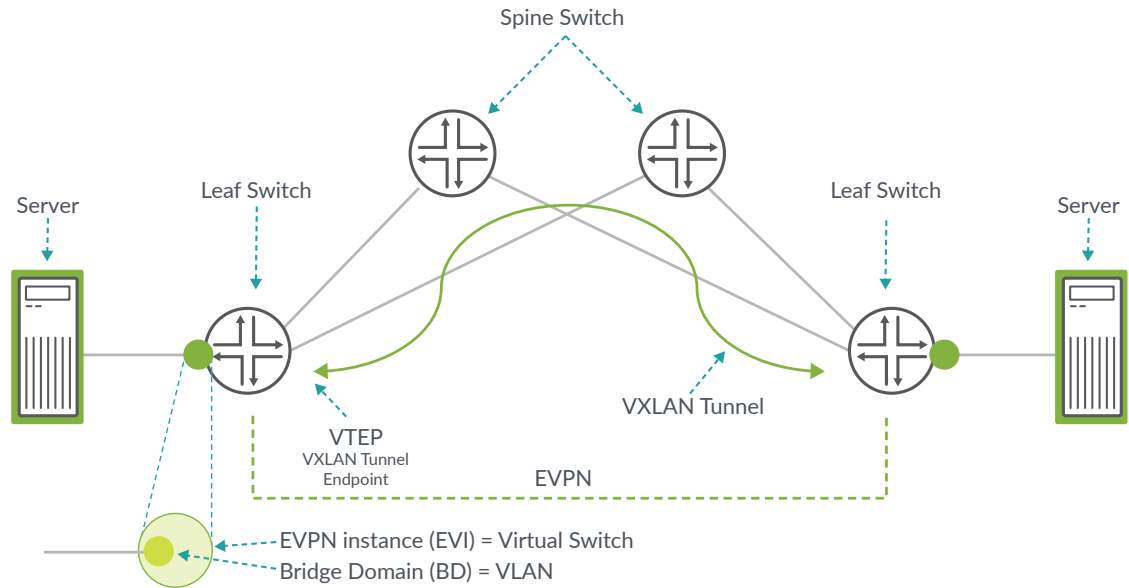


Figure 4: EVPN-VXLAN architecture

Since EVPN runs over an L3 IP fabric, it simplifies networks by reducing reliance on fragile L2 constructs like Spanning Tree. While most architects starting fresh would choose L3 over L2, the reality is that legacy applications and equipment exist in their data centers. For enterprises with applications that depend on L2 connectivity, EVPN-VXLAN stretches L2 domains across routed interfaces by encapsulating (tunneling) Ethernet frames. The coexistence of L2 and L3 makes EVPN-VXLAN the ideal technology to facilitate the transition from legacy to modern multicloud environments. Additionally, the technology can originate/terminate EVPN-VXLAN overlays on server nodes as well as on hardware network devices, allowing the same overlay to span physical and virtual elements and create a single, seamless fabric.

EVPN-VXLAN can be configured from the CLI, as well as with automation tools such as Ansible. It is also possible to orchestrate underlay and overlay setup through multicloud orchestrators like Juniper® Contrail® Enterprise Multicloud, with its unified managed interface, Contrail Command (Figure 5).

Test Drive the Juniper vQFX and EVPN-VXLAN

Download free Juniper Networks [vQFX virtualized data center switch](https://www.juniper.net/vqfx) software with a prebuilt EVPN-VXLAN network and build a virtual network on your laptop or desktop. Visit <https://www.juniper.net/vqfx> to test drive the software functionality to learn more about the QFX Series data center switches and how to build EVPN-VXLAN fabrics.

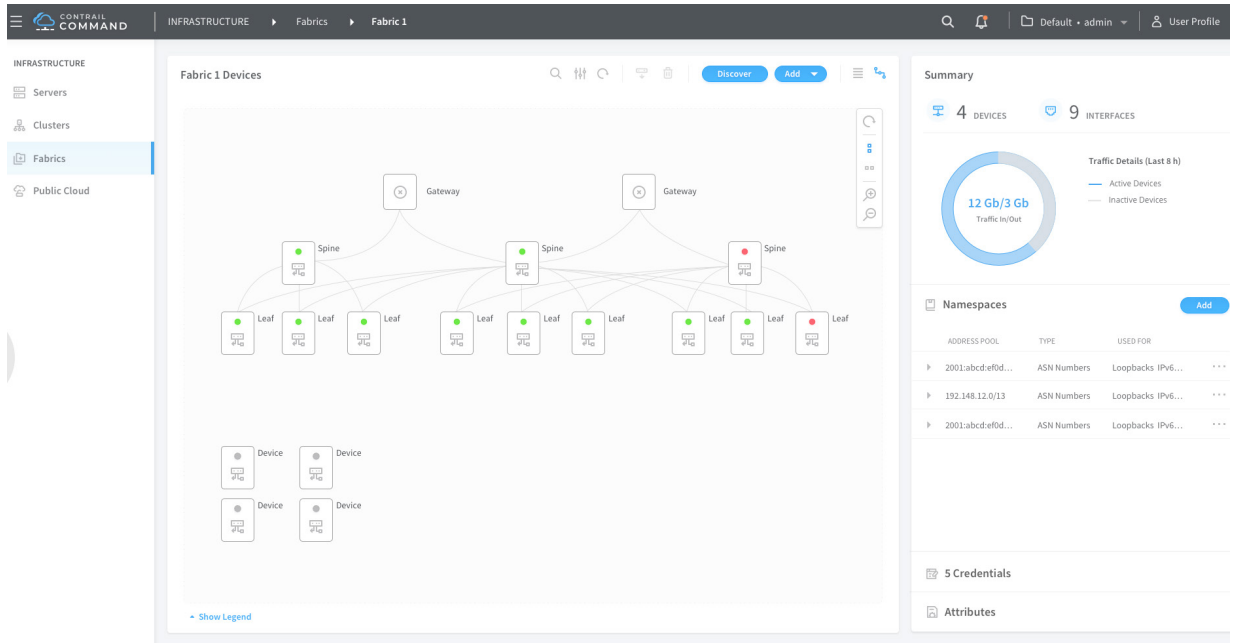


Figure 5: Unified management for multicloud

The multicloud fabric is comprised of three primary elements, outlined in Table 3.

Table 3: Fabric Elements

Element	Function	Description
Control Plane	How information is exchanged between endpoints	<p>An important aspect of implementing a common fabric is that devices utilize their own local control plane, allowing them to connect using EVPN as the mechanism for exchanging control intelligence. As an extension of BGP, with its inherent characteristics, EVPN deployments enjoy limitless scale and high reliability.</p> <p>As with the Internet, each BGP-EVPN-supported device operates independently while learning and announcing locally discovered end systems (or other devices) through the control plane.</p> <p>One of the advantages of BGP-EVPN is the ability to combine multiple fabrics (see Figure 6). Because the control plane is distributed through BGP-EVPN, any information around it can be shared with other devices, and new tools for fabric management automate its operation.</p>
Data Plane	How traffic is transported	<p>The transport mechanism within the fabric is IP, enabling the creation of a simple and open underlay and overlay due to its widespread adoption.</p> <p>In terms of the overlay, encapsulation techniques implement end-to-end connectivity. Standards-based and widely adopted encapsulation protocols such as VXLAN (or MPLS over UDP/GRE and IPsec) should be considered to preserve flexibility in vendor choice for economic leverage in new contracts.</p> <p>Proprietary data plane mechanisms create integration problems, causing incompatibilities in the fabric. Therefore, leveraging IP as a data plane protocol is important.</p>
Management Plane	How everything is glued together to scale and federate	<p>The role of the management plane is to unite everything, allowing for automated operations, visibility, and system integration throughout the multicloud environment.</p> <p>For this to occur, all devices must be part of the fabric, enabling data center and private cloud environments to be extended to the public cloud and creating a unified architecture across disparate domains. This creates a homogeneous management environment where resources can be consumed the same way everywhere, regardless of their location.</p>

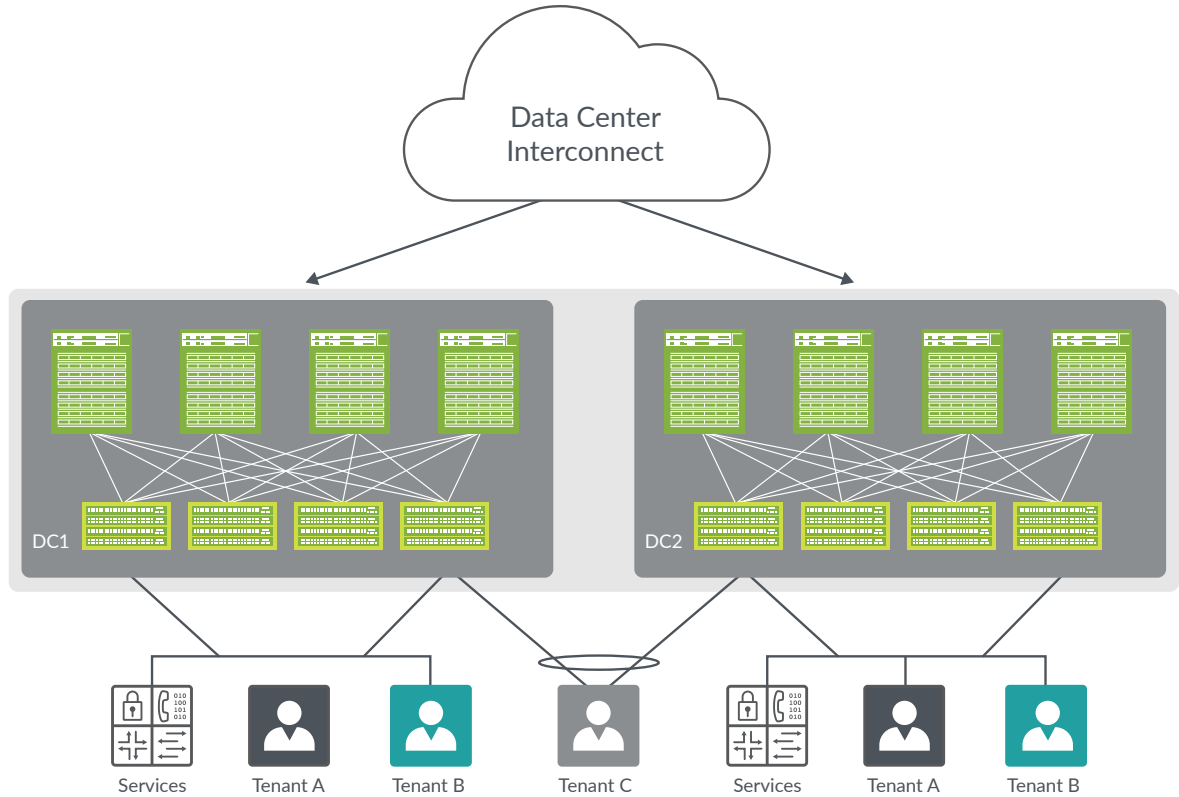


Figure 6: Federated fabric

Intent-Based Fabric Operations

Fabric operations typically focus on the workflows needed to run the network—set up, scale out, maintenance, assurance, and so on. For instance, deploying a new service might require modifying edge policy across a set of devices. Rather than changing the policy configuration on each device, operators can use fabric automation capabilities such as in Contrail Enterprise Multicloud to execute the workflow centrally and push the new policy to all required devices across the network.

The Contrail Enterprise Multicloud controller, deployed centrally on premises or in the cloud, serves as the main point of control for the underlay and overlay. Contrail Enterprise Multicloud transmits user intent submitted via templated workflows in the GUI and translates it to manage the devices directly using standard output such as JavaScript Object Notation (JSON) or Network Configuration Protocol (NETCONF). This way, the platform can manage and automate heterogeneous networks composed of multivendor devices from a single, central location, significantly simplifying fabric operations. Such platforms do not replace element management systems (EMS); EMS will continue to support point configurations for very specific requests, if or when desired.

Unifying Data Center and Campus Fabrics

If simplifying operations is critical to a successful multicloud deployment, driving a common architecture wherever possible is just as important. Otherwise, administrators have no choice but to account for and manually traverse the different environments where operational uniqueness exists.

Juniper’s solutions allow enterprises to leverage the same architecture across both campus and data center. In the short term, these domains likely exist as separate islands. However, as enterprises move towards multicloud, it creates the opportunity to manage campuses and data centers that have a common framework with over-the-top policy and control enabled by Contrail Enterprise Multicloud.

Workload Management

Orchestration provides the network abstraction needed to meet the dynamic needs of applications, often without requiring reconfiguration of the underlying network. As such, network orchestration must work in concert with the automated operations for launching and tearing down new workload instances.

There are three basic workload management strategies used by enterprises migrating to multicloud today, plus one emergent technology. Most enterprises are likely using a combination of bare-metal servers (sometimes part of a platform-as-a-service offering), VMs (typically managed by VMware or OpenStack products), and containers (typically Docker runtime managed by Kubernetes and OpenShift). These technologies are available both on premises and in public cloud offerings. Enterprises also can employ the emerging serverless approach, where fragments of applications can be executed on demand in different cloud servers rather than running as a monolith or in a few tiers.

The challenge is that, even in moderately sized environments, the application landscape likely uses a mix of these technologies. As an increasing number of virtualization and cloud computing technologies and services are introduced, networks and network security must change the way they support workloads.

Instance-Specific Application Environments

Applications are heavily dependent on edge design and policy as they relate to networking. That is to say, the application experience, security, and topology are all dependent on policy constructs like VLANs, quality of service (QoS), and firewall filters (or ACLs), along with encryption.

Networking has been a fairly static discipline for a number of years. As a result, keeping pace with the fast-moving changes in workload development and management enabled by modern application practices has proven particularly difficult for some.

While the mechanisms of spinning up a new VM or container within one environment is not particularly overwhelming, the prospect of treating multiple environments as fungible resources and spinning up various compute vehicles on any of these environments is daunting. Delivering connectivity and applying consistent policies to a new environment are often manual processes today. Provisioning the connection requires specifying edge policy, which is handled differently in each environment. For example, provisioning on premises or in a private cloud is different than provisioning in AWS or Azure; provisioning for a VM is different than provisioning for a container.

Additionally, when applying traditional security models to clouds or containers (and to edge compute or serverless processes that live for just a few seconds before reappearing elsewhere), security constructs like perimeters, enforcement points, correlation, and so on, will need to evolve from what we have done in the past. Factoring in the notion of location, as in “my data center” versus a “public provider,” the attack surface is suddenly much wider, more diversified, and featuring more entry points than anything IT has had to secure so far.

With the nature of threats constantly changing and evolving, security must be promptly deployed and maintained from workload launch to teardown. Delays between workload life-cycle management activities and the application of appropriate security policies can't be tolerated. Whether it is access control, authentication and encryption, intrusion detection service/intrusion prevention system (IDS/IPS), or some other security service, policies must always be enabled and up to date.

Complexity can also influence security. If administration is operationally complex, it can impact security effectiveness. Enterprises cannot rely on burdensome workflows manually executed by operators over an extended period. Solutions must be simple, multidomain, and dynamically managed if they are to offer adequate protection for the enterprise.

In designing the network and security to support workload management, architects need to consider three aspects: workload abstraction, overlay control, and unified underlay and overlay management.

Workload Abstraction

Workload abstraction is the first step toward automating workload management in the network. The multicloud orchestration tools must be able to express policy (for both application experience and security) in ways that are agnostic to the underlying infrastructure.

The foundational functionality for abstracted policy is to express requirements as intent. Juniper Contrail Enterprise Multicloud provides workload abstraction through its controller, using its underlying workflow engine to translate intent into device- or system-specific commands. Users interact with a workflow template (in this case, policy specification), the controller transforms the expressed intent into southbound calls to the supporting infrastructure for each type of system, EMS, and device: physical or virtual; on premises or in the cloud. Defining intent once, with translations in place, allows operators to consistently apply policy throughout the network from one platform.

Contrail Enterprise Multicloud does this in virtualized environments by utilizing its virtual router (vRouter) as a policy enforcement point for each specific workload, regardless of where it resides. In the case of applications running on bare-metal servers, the controller manages the top-of-rack switch ports to set and apply policies.

While the vRouter is installed natively on the server kernel along with its host OS (Figure 7), it can also be deployed in a VM, as a virtual private cloud (VPC) gateway, or in a public cloud instance. The vRouter evaluates ingress traffic from the resource (or resource pool, in the case of a VPC) and enforces the defined policies. This provides a location for enacting both application experience controls as well as application security in a distributed firewall model. The approach supports microsegmentation, the ability to define in-depth policies and protections for individual workloads.

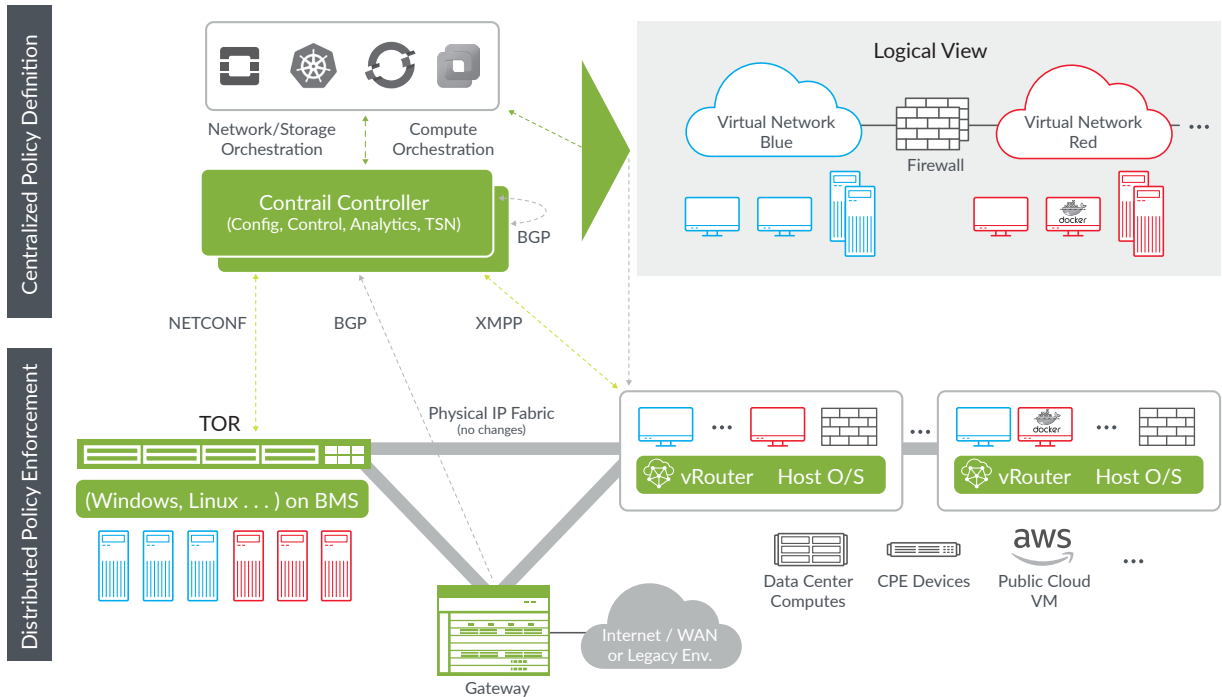


Figure 7: Contrail architecture

Contrail Enterprise Multicloud lets operators uniformly orchestrate and manage policies across all environments where the application might execute—all from a single, central location. In this way, the policy becomes automatically leashed to the workload wherever it launches. The solution supports workloads running in VMs or containers on premises, in private clouds, or in a public cloud (AWS, Azure, Google Cloud, and so on). Contrail Enterprise Multicloud provides virtual networking with consistent and highly granular services (routing, firewall, NAT, load balancer, and so on), all extended to the workload level. This includes the following virtualization environments, enabling a level of dynamic control required for the modern multicloud (see Figure 8):

- A private cloud VM running a VMware ESXi/vCenter environment
- A private cloud VM running OpenStack
- A private cloud Docker container
- A public cloud VM or container

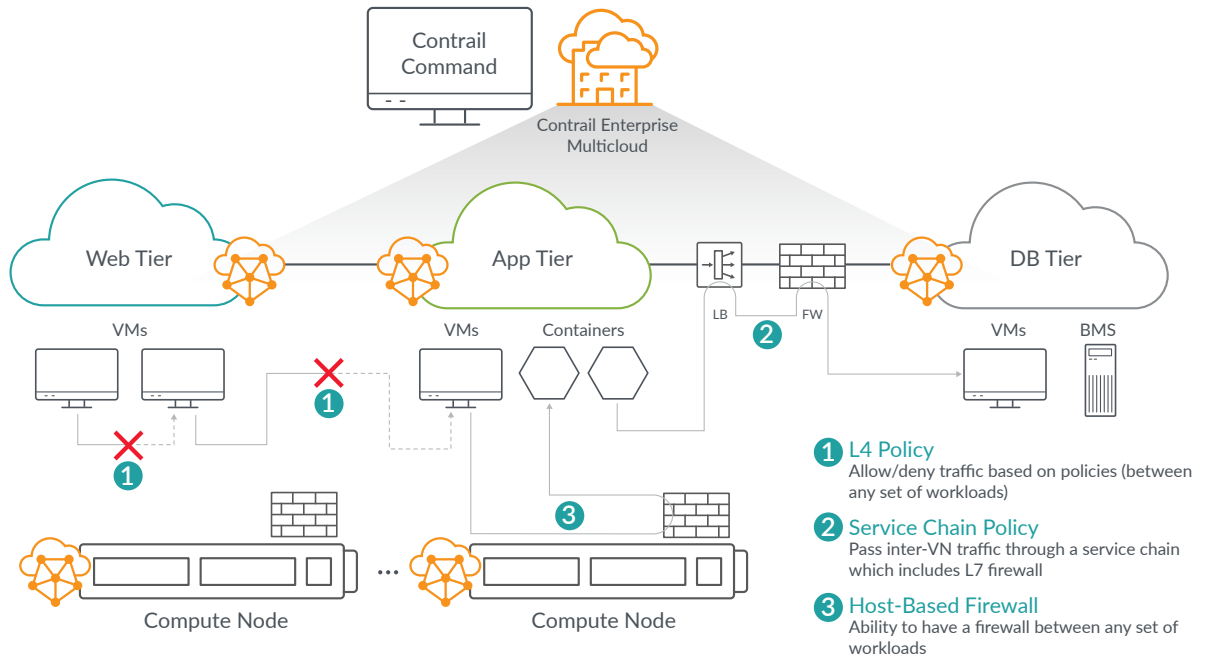


Figure 8: Policy and control

Overlay Control

The second major requirement for workload management in the multicloud is overlay control to use the established EVPN-VXLAN fabric everywhere. If application instances are being spun up and down across the multicloud based on demand, the network must be able to dynamically set up new secure paths for those applications.

The overlay network is key to allowing applications to operate this way. Knowing where workload, data, and users are, along with whatever role-based permissions are relevant, should be enough for an application to function. To enable this, the overlay must be automatically configured based on the secure connectivity requirements of the application.

Using overlays, edge policies are applied to the overlay endpoints, including the vRouters associated with the workloads (refer to Figure 8). Managing these connectivity and security policies is the responsibility of the overlay controller. Using intent-based workflows, Contrail Enterprise Multicloud can provision these policies across any number of endpoints, no matter where they reside, giving operators a unified approach to their multicloud.

The design also allows traffic to be encrypted at the overlay endpoints, close to the workload itself, providing encryption at the source. The secured overlay environment means traffic is sent to the underlay fully encrypted—especially important when using public cloud infrastructure, which is owned by third-party cloud providers.

Unified Underlay and Overlay Management

Lastly, to ease the day-to-day operations, architects should consider a single platform like Contrail Enterprise Multicloud that combines underlay and overlay management and correlation. Contrail Enterprise Multicloud supports fabric management along with heterogeneous compute environments, including bare-metal servers, VMs, containers, and networking devices; private and public clouds; and orchestration of networking and security policies, including microsegmentation, meeting the broad set of requirements for multicloud management.

Services Consumption

The ultimate measure of success for any multicloud deployment is whether the underlying infrastructure is transparent to the user. The goal of multicloud is to allow workloads to be deployed anywhere based on business and functional needs such as cost. The user should not be able to tell whether a workload is served out of a private or a public cloud.

For this to be possible, the network must ultimately integrate into the application layer, both in terms of connectivity and security, as well as how new applications and services are deployed and consumed. This top layer of the multicloud architecture decouples infrastructure and services by transparently abstracting the lower layers into the set of services required for each of the applications.

By decoupling services from the underlying infrastructure, the design establishes secure multitenancy, isolating services from each other. The isolation mechanism only allows services to communicate based on the operator’s specified intent, simultaneously enforcing both connectivity and security policies.

The following sections describe ways the multicloud network and applications can be tightly integrated during the application development process, in response to real-time conditions and as a part of service-aware constructs.

Application Development Life Cycle

Most modern applications follow a development-staging-production life cycle. Where the application runs, along with its network and security needs, will change as each application moves through its life cycle.

Contrail Enterprise Multicloud treats these stages separately by using labels, each with its own networking and security environment and a separation of policy for the various application stages. Because policy can be bound to each of these stages, it establishes a powerful multicloud construct where developers can create and validate applications for a certain network setup with specific security policies and then deploy that environment as part of a continuous integration and continuous delivery (CI/CD) application pipeline (see Figure 9). Labels also allow a higher level of abstraction that leads to location independence. By labeling objects, operators can now have any stage anywhere with specific policy associations.

1. **Reduced complexity** (*less # of policies*)
2. **Simplified management** (*change control, etc. is much easier*)
3. **Improved scalability**
4. **Define / review / approve once -> Use everywhere**

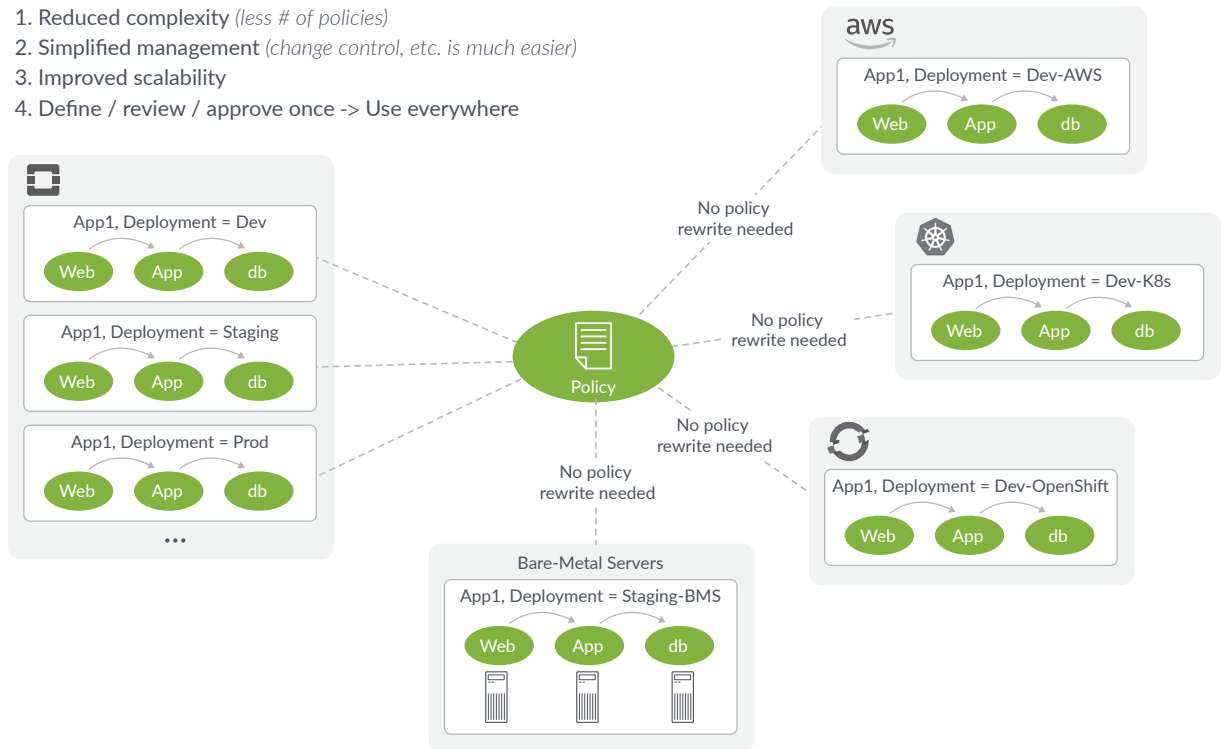


Figure 9: Simplified policy enforcement

By integrating the network and security into the application environment, Juniper breaks down the traditional separation between the workflow of the application and network teams, resulting in a more seamless process that produces greater velocity and supports broader enterprise agility objectives.

Service-Aware Networking

When determining how the network and applications work together, architects must consider how to service dynamic application needs based on the real-time conditions of the infrastructure. Traditionally, the network and served applications are treated independently. Applications typically assume zero-loss, low-latency connectivity, while networks are mostly unaware of what an application is doing.

The rise of technologies like software-defined WAN (SD-WAN) present an opportunity to dynamically tune policy in an effort to optimize network performance for specific application and user needs. For example, offloading traffic during critical or bandwidth-intensive activities (such as disaster recovery) can be an effective way of minimizing impact on precious WAN resources. Or an anomalous network event

might trigger the isolation of key resources and the quarantining of a particular workload. In both cases, the key is feeding real-time telemetry from the infrastructure to a controlling entity for event correlation and dynamic remediation or alerting.

Architecturally, these use cases require some means of collecting and distributing data. Modern distributed architectures use a message bus with listeners like a collector capable of performing real-time analysis, which puts a premium on real-time streaming of event and state information within the infrastructure (see Figure 10).

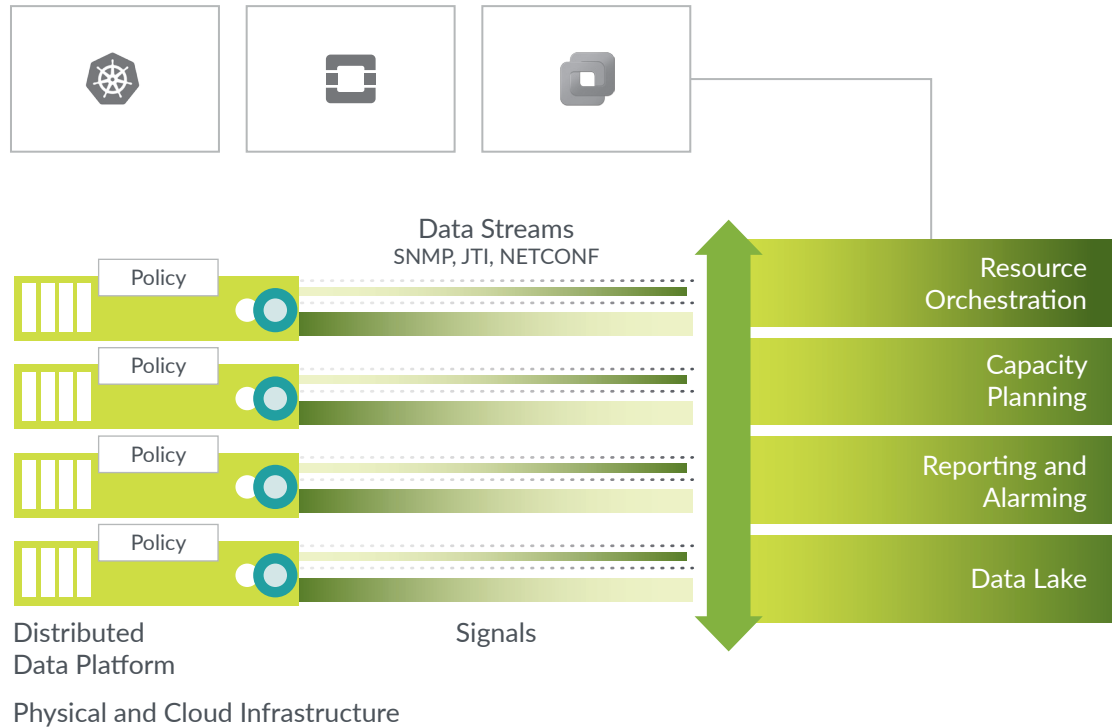


Figure 10: Advanced analytics

In the underlay, devices might use a technology like Google's open source gRPC, a feature common across Juniper's portfolio. In higher layer elements, infrastructure performance monitoring (IPM) and application performance monitoring (APM) tools provide insight into things like performance and utilization across key physical and virtual resources. The key to determining effectiveness is the richness of available data, along with the extent of its real-time availability. Legacy protocols like SNMP have gaps both in the breadth of state available and in polling performance. Whatever the streaming mechanism, the data is typically distributed via the message bus, which feeds the information to the listeners for actions to be performed.

Things like topology, addressing, location, traffic type, rates, sessions, name space, and so on are collected and then distributed to the analytics application, providing the context necessary for making informed decisions about path selection or security actions. Juniper Networks Contrail Insights software provides comprehensive and real-time visibility into the environment, along with intent-based analytics. Contrail Insights enables automated preventative remediation in real time, covering the layers of the stack and operating across infrastructure types (see Figure 11).

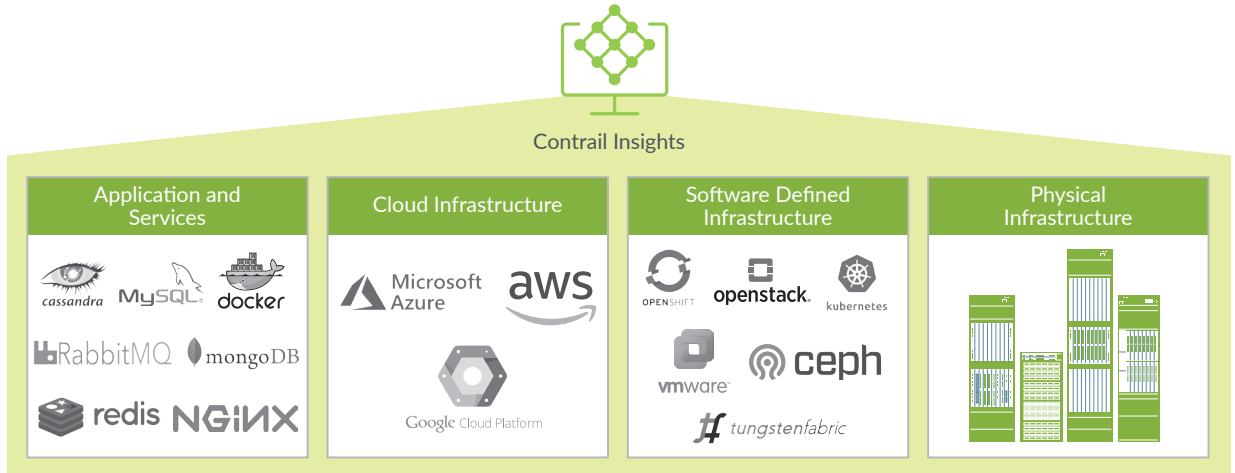


Figure 11: Cross-layer visibility and analytics

Service-Aware Constructs

Over time, multicloud deployments will feature service-aware constructs that offer consistent network services across all technology silos, both current and future (see Figure 12). These constructs specify service details that are fundamentally decoupled from the underlying infrastructure, allowing them to be applied through policy translation administered dynamically by the multicloud controller.

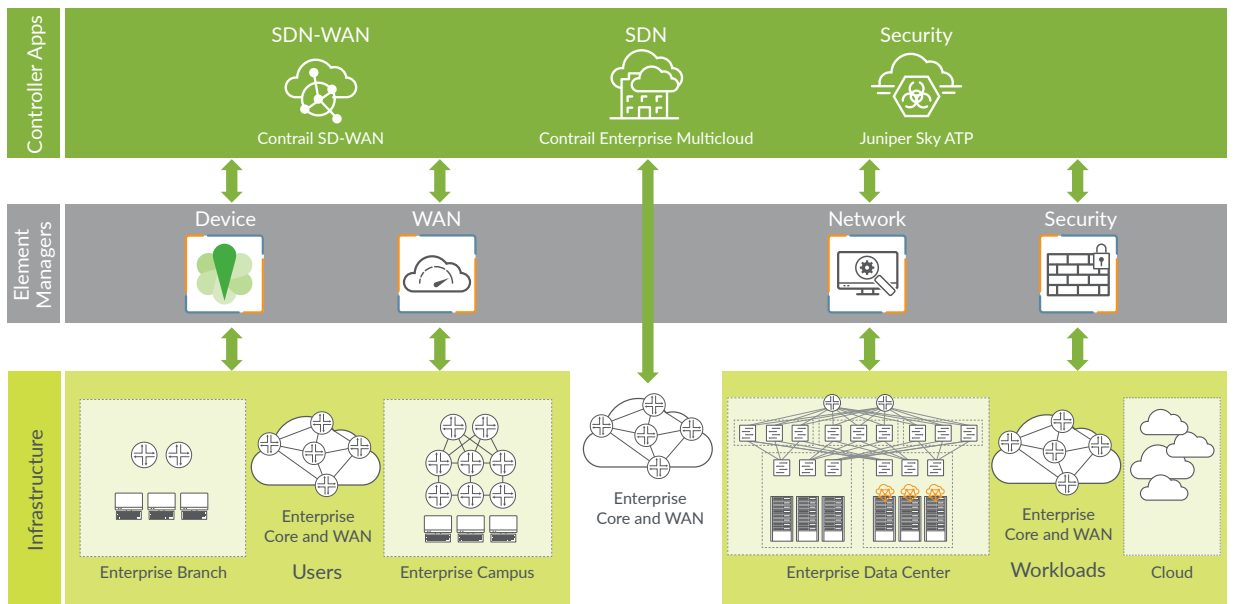


Figure 12: Service-aware constructs

Authenticated users authorized to be on the network should be able to access any application that exists within their domain, provided it falls within the necessary role- and user-based permissions. When the request is made, the multicloud controller learns the user's needs through the predefined service-aware construct, which then defines the specific network services required to establish optimal connectivity between the user and the application. This includes path setup as well as policy to enforce appropriate application experience (like QoE at the gateway) and security (including encrypted transport or firewall rules).

Since the construct is not bound to the infrastructure, if the workload moves or is instantiated in a different resource pool, the construct follows it to the new resource.

Conclusion

Application innovation is powering today's digital business. The acceleration of the application development and updating process, coupled with the flexibility of where workloads can run, introduces significant complexity—particularly operational—into the network. Challenges arise when organizations begin migrating their application workloads from cloud to multicloud, as well as to more virtualized technologies. Not only must the network support these new operational environments, it must securely support communications with the existing applications and data.

With application environments and requirements evolving rapidly, enterprises require new approaches to network design, security, and operations. A multicloud capable of connecting and securing applications end to end across many clouds, as simply as if they were one, lets organizations optimize resources as a single, cohesive infrastructure with consistent operations throughout.

This means operators can holistically manage the network for workloads, running on VMs, containers, or bare-metal servers, on premises and in the public cloud, while managing the overlay along with the underlay. They can provision, execute workflows, and monitor everything end to end based on intent-driven direction relevant to their role. This move to multicloud is more than just technological; it requires enterprises to evolve their architectures, processes, and people.

Juniper's blueprint for multicloud features multidomain connectivity, multivendor orchestration, end-to-end visibility, and pervasive security. Using a single orchestration platform directing common policy and operations across a multivendor environment, Juniper solutions help you manage the complexity of multicloud, while a multicloud framework provides a methods-driven migration map for your multicloud transition.

Juniper's unique approach lets you retain control while providing application developers and users with the agility they must have to compete.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

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

