Solution Brief

SDSN: Dynamic, Adaptive Multicloud Security

Evolving from firewall to user-intent policies for flexible security in the cloud

Challenge

Legacy security policies, which do not dynamically adapt to different security workflows, must be individually configured and managed for every possible deployment option (physical deployment, private cloud, and public cloud).

Solution

User-intent based policies that leverage unified metadata across deployment realms dynamically adapt to different security workflows, helping enterprises effectively handle the security challenges of the multicloud era.

Benefits

- Provides a single policy model across clouds
- Simplifies management and user readability
- Reduces operational expenses
- Enhances application deployment workflows by eliminating the need to commit configurations
- Improves security workflows for security administrators
- Enables administrators to rapidly take corrective action under critical situations

Because traditional firewall policies rely heavily on IP addresses, they are fairly ineffective in today's cloud-based world. To handle the security challenges of the multicloud era, enterprises must be able to adapt security policies dynamically across different workflows and deployment options, employing a single yet flexible multicloud policy model. When a network is under attack, the need to rapidly isolate the threat and take corrective action is critical.

Juniper Networks offers just such a solution with its Software-Defined Secure Network. With powerful security workflows facilitated by dynamic access groups and a unified and intuitive metadata-based policy model that can be ported across clouds, SDSN gives security admins complete command and control over their multicloud deployments.

The Challenge

Traditional firewall policies, which rely heavily on IP addresses, have not seen significant improvement in more than a decade. Unfortunately, these policies are not terribly effective in today's cloud-based world, where IP addresses are always changing as dynamic virtual instances are constantly being created and terminated. Furthermore, each environment has its own native elements such as security groups in Amazon Web Services (AWS) or VMware NSX, with tags that are hard to accommodate using traditional firewall policies, limiting business agility.

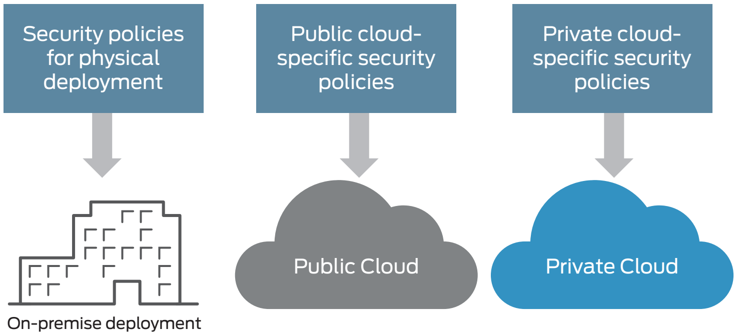
Security policies for physical deployment
Public cloud-specific security policies
Private cloud-specific security policies
On-premise deployment
Public Cloud
Private Cloud

Figure 1: Security policies need to be customized for each deployment realm

To fully understand the impact this evolving environment will have on network security, it's important to first know how an application is deployed in an enterprise.

Once an application is developed, the app team asks the infrastructure team to allocate the necessary resources. The infrastructure team begins building the infrastructure, including compute, storage, and network resources, and then asks the security team to allow these applications. The security team analyzes each application, creates new or modifies existing security policies as needed, and commits the changes.

While this whole process can take anywhere from four to six weeks in a physical deployment, it can be accomplished in a matter of minutes in a cloud environment. In a highly distributed environment, however, where the infrastructure is spread across multiple clouds, it gets more complicated for security administrators tasked with reviewing the consistency of policies across the enterprise and committing the configuration changes. As the number of applications being deployed in a distributed system grows, the security team itself becomes an obstacle to business agility.

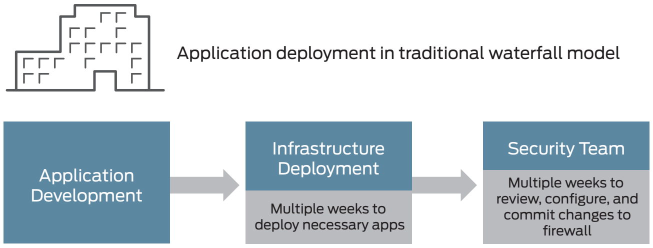
Application deployment in traditional waterfall model
Application Development
Infrastructure Deployment
Multiple weeks to deploy necessary apps
Security Team
Multiple weeks to review, configure, and commit changes to firewall

Figure 2: Typical application deployment timeframes

The static nature of security policy actions presents another challenge to responding effectively to different security workflows. With traditional firewall policies, when the network is under attack, the security admin must sift through literally thousands of rules, disabling unnecessary ones in order to isolate the source of the event. Similarly, security admins need to constantly experiment with policy settings to determine the best security posture for specific network performance conditions, switching between them based on business needs. Making these changes manually across a large rule base is time-consuming and extremely inefficient—not ideal in time-sensitive or repetitive situations.

The Juniper Networks User-Intent Policy Security Solution

Juniper Networks addresses these challenges with a comprehensive security solution based on user-intent policies that features a metadata-based policy model and powerful dynamic policy action capabilities. Part of Juniper's Software-Defined Secure Networking (SDSN) framework, this solution consists of the following components:

- Juniper Networks® SRX Series Services Gateways and vSRX Virtual Firewall with integrated next-generation firewall (NGFW) and unified threat management (UTM). These products deliver:
    - Core firewall functionality with IPsec VPN and feature-rich networking services such as Network Address Translation (NAT) and routing
    - Intrusion prevention system (IPS) 2.0 to detect and block network intrusions
    - User-based firewalls to analyze, log, and enforce access control based on user roles and groups
    - Application control and visibility with integrated Juniper Networks AppSecure 2.0 to provide application-level analysis, prioritization, and blocking to safely enable applications
    - Antivirus, antispam, and Web and content filtering with UTM to protect against viruses, spam, and malicious URLs and content
    - Support for Linux KVM, VMware, AWS, and Azure platforms (vSRX)
    - Support for dynamic address groups to facilitate a "no-commit" architecture
- Juniper Networks Junos Space® Security Director, which provides centralized, single-pane-of-glass management to deploy, monitor, and configure security features and policies across all SRX Series physical and vSRX virtual firewalls in the network.
    - Policy Enforcer, a component of Security Director, provides an additional level of centralized intelligence for deploying and enforcing security policies on multivendor network elements such as switches, routers, Wi-Fi access points, and the like.
    - With the introduction of user-intent policies, it also provides the ability to use metadata in security policies and take dynamic actions on security policies using the Dynamic Policy Actions (DPA) feature.
    - Security Director includes a customizable dashboard with detailed drill-downs, threat maps, and event logs, providing unprecedented visibility into network security measures.
    - It is also available as a mobile app for Google's Android and Apple's iOS systems to enable remote mobile monitoring.

Metadata-Based Policy Model

Metadata is defined as a dictionary of key value pairs gleaned from a list of possible values that associate attributes with endpoints. Most cloud deployments (VMware NSX, Juniper Contrail® Platform, Amazon AWS, etc.) use metadata native to their respective clouds.

The ability to leverage metadata in security policies significantly enhances traditional security policies that already feature static identifiers such as IP addresses or address objects. By following common metadata-based tagging, the SDSN Policy Enforcer component of Security Director can configure all firewalls throughout the enterprise with the same policies—without requiring any major customization to make it cloud-compatible. A network or DevOps admin simply assigns the right metadata to an endpoint, and a preconfigured security policy can utilize it immediately.

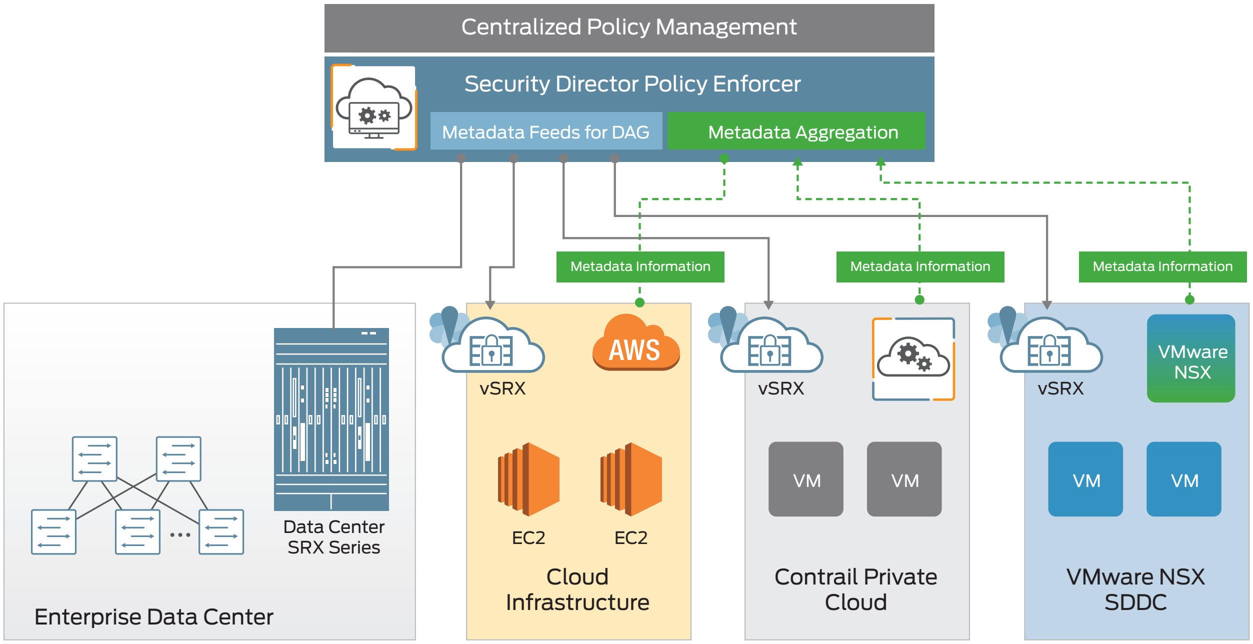

Figure 3: Metadata aggregation and feeds for dynamic address groups on SRX Series and vSRX NGFWs

Dynamic Address Groups on SRX Series and vSRX NGFWs

Dynamic address groups on SRX Series physical and vSRX virtual NGFWs facilitate adaptive security policies—elements comprised of lists of IP addresses. Once a dynamic address group is used in a security policy, IP addresses can be added or removed without requiring a configuration commit—a tremendous benefit considering endpoint IP addresses change constantly due to the dynamic nature of creation and termination, and appending thousands of addresses is incredibly time-consuming. As an added benefit, the SRX Series and vSRX firewalls also dynamically populate and update IP addresses in the dynamic address group based on feeds from external sources.

Metadata Mapping to Dynamic Address Groups

Security Director Policy Enforcer serves as an aggregation server, learning the native metadata from different realms (Contrail, VMware NSX, AWS, etc.) and their IP associations. This information is then fed to the SRX Series and vSRX NGFWs in different realms, acting as a metadata feed server as well as a policy management system.

When applications migrate from physical on-premise hardware to the cloud, the firewall in the cloud will use the metadata from the application endpoint to determine the right action for traffic moving to and from that endpoint. This facilitates a much smoother and agile workflow for migrating applications without compromising security.

Using meaningful metadata also captures the user's intent, making endpoints easy to identify and work with in a security policy.

The solution also introduces the ability to use logical operators in security policies, such as an "AND" and "OR" operation that can be performed on endpoint metadata. This level of flexibility simplifies the overall policy configuration.

For example, if configuring a source field for all endpoints with metadata "Type" as "DB" and metadata "PCI" as "no," it can be specified in the security policy's source or destination address.

The following tables offer a quick glimpse of the possibilities available to admins to reduce the security burden on the enterprise and create a higher level of operational flexibility.

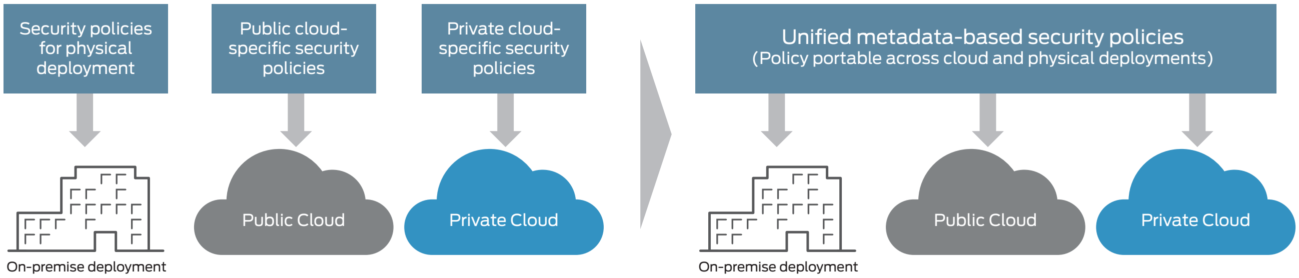

Figure 4: Metadata-based policy model

The security policy using a static IP looks like this:

#	Source	Destination	Service	Firewall	IPS
1	Web server object	Database address object	http	Permit	None

The backend address object-to-static IP mapping looks like this (requires the security team):

#Address Object	Mapping
HR_SR_object	1.1.1.1
SQL_FIN_OBJECT	2.2.2.1

When a new module needs to be added to an application, the DevOps or network admin provisions the necessary application and asks the security team to add the IP address to the allowed list. The administrator makes the necessary changes to the address object or static address and commits the configuration. Configuration commits are typically a time-consuming process, adding an additional layer of overhead.

#	Source	Destination	Service	Firewall	IPS
1	PCI ==NO	TYPE == DB && PCI ==YES	Any	Deny	None

The metadata-to-address object mapping looks like this:

#Address Object	IP Mapping	Metadata Deployment Status	Metadata PCI	Metadata Type
HR_SR_ object	1.1.1.1	Production	No	Web
SQL_FIN_ OBJECT	2.2.2.1	Production	Yes	DB

The security admin creates security policies based on metadata logic for the source and destination fields. When the DevOps or network admin provisions a new endpoint to the application, only the appropriate metadata needs to be associated with the new endpoint in order to deploy it in the application pool.

Using security templates or predefined metadata-based policies approved by the security team can eliminate the 3-4 week process of configuring and committing configurations. This dramatically reduces application deployment times, ensures greater business agility, and significantly lowers OpEx. Similarly, endpoints can be removed from the application pool by simply deleting the associated metadata.

Dynamic Policy Actions

Security Director's Policy Enforcer includes a Dynamic Policy Action (DPA) feature based on configurable global NetSec environmental variables that allow preconfigured security policies to take different actions under various conditions. This facilitates a proactive approach to threat response. The ability to take alternative actions across a large set of policies dramatically simplifies the security posture, allowing the system to dynamically adjust its response under different attack conditions.

When a network is under attack, the need to rapidly isolate the threat and take corrective action is critical. A typical workflow requires security admins to manually disable all noncritical business applications until a diagnosis is made. In a normal enterprise, there could be thousands of security policies governing noncritical applications such as video or audio streaming, private e-mails, and mobile apps. If the network is believed to be under attack, using a DPA can allow a single NetSec variable to perform customized policy actions on a large set of policies to protect the infrastructure.

Similarly, when optimizing security policy settings for performance and security or mode-switch to prioritize one application over another, DPA achieves this with a single NetSec variable change. For example, a security admin could set a NetSec variable to skip additional IPS settings on internal traffic when large data backups occur, switching back to "Full Security" mode once the backup is complete. Admins could also create an "Experimental" mode to find the right balance between security and performance settings, then preconfigure the policies to take appropriate actions based on this setting.

These fast-switch workflows not only allow security teams to easily collaborate with other groups throughout the organization, it also allows security admins to respond effectively in high-pressure situations.

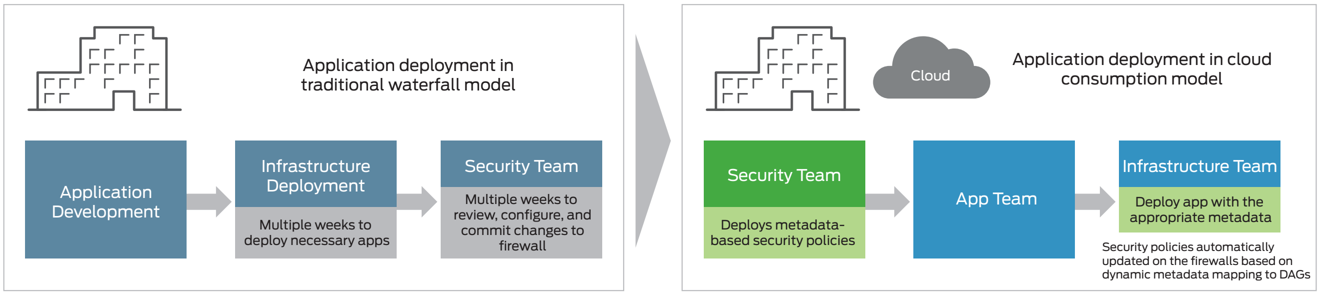
Application deployment in traditional waterfall model
Application Development
Infrastructure Deployment
Multiple weeks to deploy necessary apps
Security Team
Multiple weeks to review, configure, and commit changes to firewall
Cloud
Application deployment in cloud consumption model
Security Team
Deploys metadata-based security policies
App Team
Infrastructure Team
Deploy app with the appropriate metadata
Security policies automatically updated on the firewalls based on dynamic metadata mapping to DAGs

Figure 5: Reduced application deployment timeframes with metadata-based policy

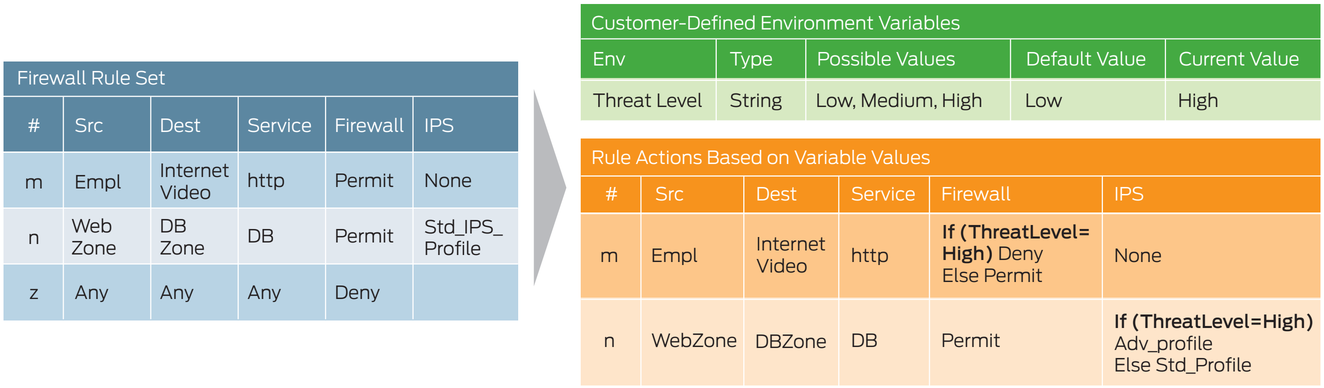

Firewall Rule Set

#	Src	Dest	Service	Firewall	IPS
m	Empl	Internet Video	http	Permit	None
n	Web Zone	DB Zone	DB	Permit	Std_IPS_ Profile
z	Any	Any	Any	Deny	

Customer-Defined Environment Variables

Env	Type	Possible Values	Default Value	Current Value
Threat Level	String	Low, Medium, High	Low	High

Rule Actions Based on Variable Values

#	Src	Dest	Service	Firewall	IPS
m	Empl	Internet Video	http	**If (ThreatLevel= High)** Deny Else Permit	None
n	WebZone	DBZone	DB	Permit	**If (ThreatLevel=High)** Adv_profile Else Std_Profile

Figure 6: Metadata aggregation and feeds for a dynamic access group on SRX Series and vSRX NGFWs

Features and Benefits

The Juniper intent-driven policy solution delivers the following features and benefits:

- Ensures agile application deployments by eliminating the need for manual configuration changes and commits
- Reduces operational expenses
- Improves user readability and flexibility when using metadata logic in security policies
- Allows security policies to be ported across different realms, including physical data centers, private clouds, and public clouds
- Dynamically adapts to network threat levels, facilitating efficient debugging and threat isolation workflows

Summary—Juniper Delivers Unified Security Solution for the Multicloud Era

By supporting advanced L4-L7 security features such as user firewall, application firewall, and IPS, coupled with powerful underlying features such as dynamic access group to facilitate a “no-commit” architecture, the SRX Series Services Gateways and vSRX Virtual Firewalls deliver the ideal cloud-ready NGFW.

With the powerful security workflows facilitated by dynamic access groups, and a unified and intuitive metadata-based policy model that can be ported across clouds, Juniper’s SDSN-based user-intent policy solution gives security admins complete command and control over their multicloud deployments.

Next Steps

For more information about Juniper Networks security solutions, please visit us at www.juniper.net/us/en/products-services/security or contact your Juniper Networks sales representative.

About Juniper Networks

Juniper Networks challenges the status quo with products, solutions and services that transform the economics of networking. Our team co-innovates with customers and partners to deliver automated, scalable and secure networks with agility, performance and value. Additional information can be found at Juniper Networks or connect with Juniper on Twitter and Facebook.

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

3510634-001-EN Dec 2017