VMware NSX DFW Policy Rules Configuration Technical White Paper

VMware NSX for vSphere, Release 6.x

Sept 23, 2014

Contents

NOTE: To obtain the latest information about NSX for vSphere, please visit

http://www.vmware.com/products/nsx

Introduction

VMware NSX Distributed Firewall (DFW) provides the capability to enforce firewalling functionality directly at the Virtual Machines (VM) vNIC layer. It is a core component of the micro-segmentation security model where east-west traffic can now be inspected at near line rate processing, preventing any lateral move type of attack.

This technical brief gives details about DFW policy rule configuration with NSX. Both DFW security policy objects and DFW consumption model will be discussed in this document.

We assume reader has already some knowledge on DFW and Service Composer functions. Please refer to the appropriate collateral if you need more information on these NSX components.

Distributed Firewall Object Grouping Model

NSX provides the capability to micro-segment your SDDC to provide an effective security posture. To implement micro-segmentation in your SDDC, NSX provides you various ways of grouping VMs and applying security policies to them. This document specifies in detail different ways groupings can be done and details on when you should use one over the other.

Security policy rules can be written in various ways as shown below:

Network Based Policies: This is the traditional approach of grouping based on L2 or L3 elements. Grouping can be based on MAC addresses or IP addresses or a combination of both. NSX supports this approach of grouping objects. The security team needs to aware of networking infrastructure to deploy network-based policies. There is a high probability of security rule sprawl as grouping based on dynamic attributes is not used. This method of grouping works great if you are migrating existing rules from a different vendor's firewall.

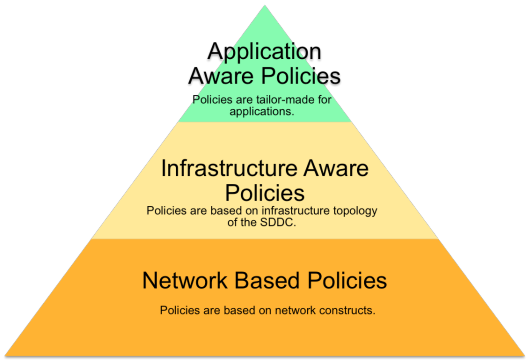

When not to use this: In dynamic environments, e.g. Self-Service IT; Cloud automated deployments, where you are adding/deleting of VMs and application topologies at a rapid rate, MAC addressed based grouping approach may not be suitable as there will be delay between provisioning a VM and adding the MAC addresses to the group. If you have an environment with high mobility like vMotion and HA, L3/IP based grouping approaches may not be adequate either.

Infrastructure Based Policies: In this approach, grouping is based on SDDC infrastructure like vCenter clusters, logical switches, distributed port groups, etc. An example of this would be, clusters 1 to cluster 4 are earmarked for PCI kind of applications. In such a case, grouping can be done based on cluster names and rules can be enforced based on these groups. Another example would be, if you know which logical switches in your environment are connected to which applications. E.g. App Tier Logical switch contains all VMs pertaining to application 'X'. The security team needs to work closely with the vCenter administration team to understand logical and physical boundaries.

When not to use this: If there are no physical or logical boundaries in your SDDC environment then this type of approach is not suitable. Also, you need to be very careful where you can deploy your applications. For example, if you would like to deploy a PCI workload to any cluster that has adequate compute resources available; the security posture cannot be tied to a cluster but should move with the application.

Application Based Policies: In this approach, grouping is based on the application type (e.g: VMs tagged as "Web_Servers"), application environment (e.g: all resources tagged as "Production_Zone") and application security posture. The advantage of this approach is that the security posture of the application is not tied down to either network constructs or SDDC infrastructure. Security policies can move with the application irrespective of network or infrastructure boundaries. Policies can be templated and reusable

across instances of same types of applications and workloads. You can use variety of mechanisms to group. The security team needs to be aware of only the application that it is trying to secure based on the policies. The security policies follow the application life cycle, i.e. comes alive when the application is deployed and is destroyed when the application is decommissioned.

When not to use this: If the environment is pretty static without mobility and infrastructure functions are properly demarcated. You do not need to use application-based policies.

Application-based policy approach will greatly aid in moving towards a Self-Service IT model. The Security team needs to be only aware of how to secure an application without knowing the underlying topology. Concise and reusable security rules will require application awareness. Thus a proper security posture can be developed via application based policies.

VMware recommends using Infrastructure based policies or Applications based policies when deploying NSX.

There are 2 ways to configure DFW rules:

1. Option 1: Using Firewall menu (Networking & Security -> Firewall). In this menu, administrator enter security policy rule as needed using the standard rule schema: rule name / source / destination / service / action. Option 1 encompasses both Network based policies and Infrastructure Based Policies.
2. Option 2: Using Service Composer menu (Networking & Security -> Service Composer). In this menu, a Security Policy (SP) must be created. Within the SP, DFW policy rules can be defined (with some differences compared to previous case). Option 2 maps with Application based policies.

We are going to detail each option but before that, let's review Security-Groups notion.

NSX Security-Groups

Security-Groups is a container-construct which allows to group vCenter objects into a common entity.

When defining a Security-Groups, multiple inclusion and exclusion can be used as shown in the diagram below:

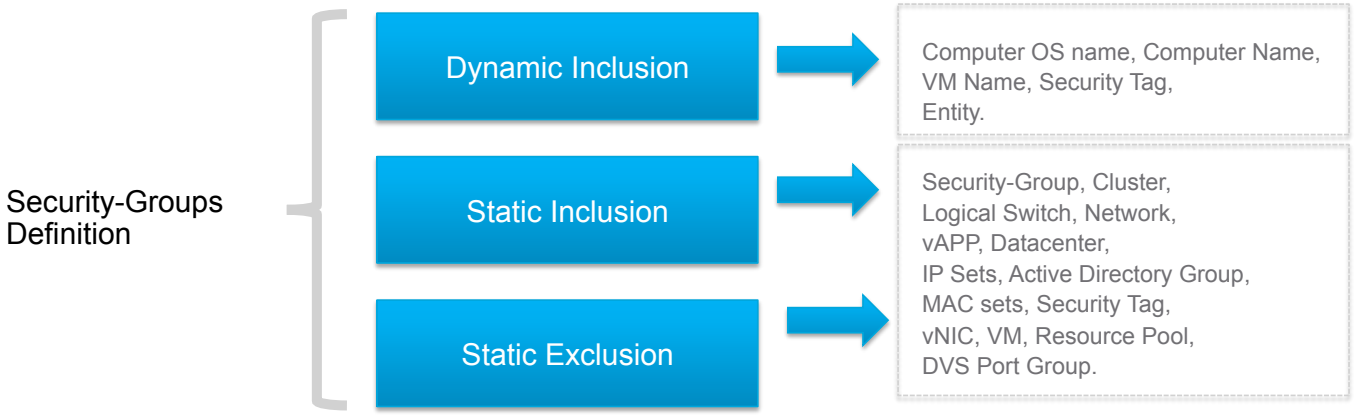

Figure 1 – Security-Groups definition.

Dynamic Inclusion allows to use criteria like VM name or Security Tag to automatically include a VM in the Security Group. For instance, a dynamic inclusion policy can state all VM with name starting by "WEB-VM" to be included in Security-Group named SG-ALL-WEB-VM. An another example is

Static Inclusion provides capability to manually include particular objects into the Security Group. List of objects are shown in the diagram above.

Static Exclusion provides capability to manually exclude particular objects from the Security Group. List of objects are shown in the diagram above.

Security-Groups result is based on this calculation:
Security Group Members = (Dynamic Inclusion + Static Inclusion) – Static Exclusion

Security-Groups is an extremely important concept because it can be leveraged in a very efficient way when writing DFW security policy rules or when defining Service Composer/Security Policy rules.

Distributed Firewall Policy Rule Configuration Using Firewall Menu (Option 1)

Click on Networking & Security -> Firewall to access the DFW policy rule table:

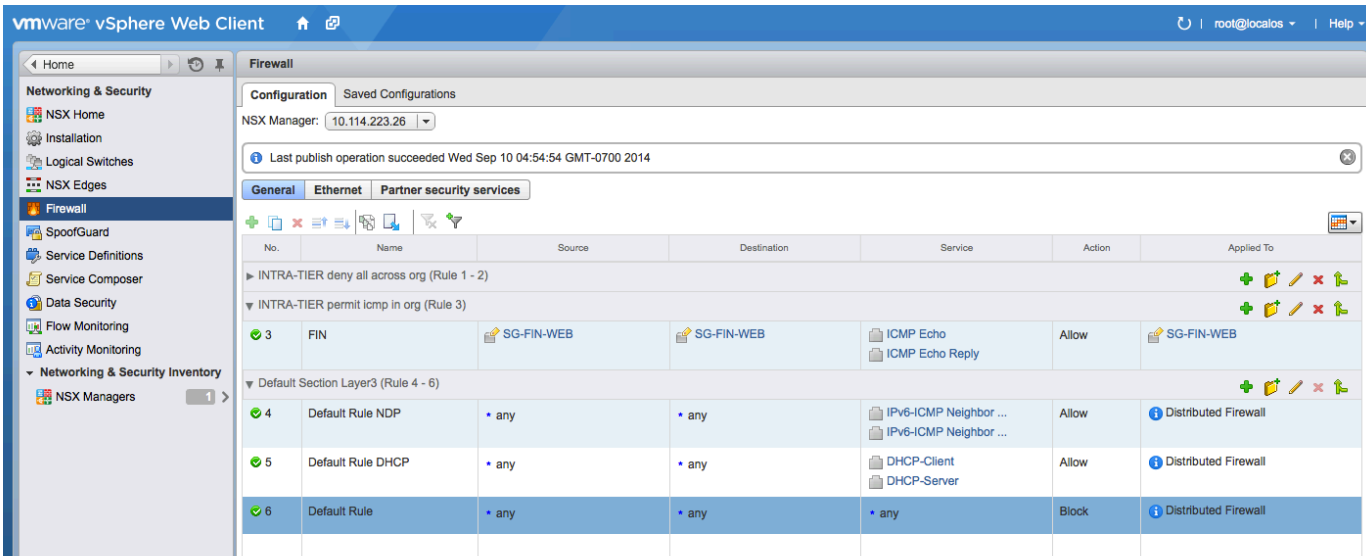

Figure 2 – DFW policy rule table.

The window displays all policy rules configured for DFW and packet lookup will be performed from top to bottom. Packet not matching any explicit rule will be enforced by default rule which is always the last one on the table. NSX comes with a DFW default rule set to Allow action. User can change it to Block if desired. VMware recommends using DFW with default rule set to Block and then create explicit rules for allowed traffic.

A policy rule is composed of the following fields:

Rule ID	Rule Name	Source	Destination	Service	Action	Applied To

- **Rule ID**: Number with 4 digit (e.g 1013) automatically allocated by DFW.
- **Rule Name**: User field which support up to 30 characters.
- **Source and Destination**: Source and Destination fields (respectively) of the packet. Possible entries are:
    1. IPv4/IPv6 addresses or subnets (e.g 192.168.200.1, 192.168.200.1/24, 192.168.200.1-192.168.200.24). The following window allows to enter IP address information:

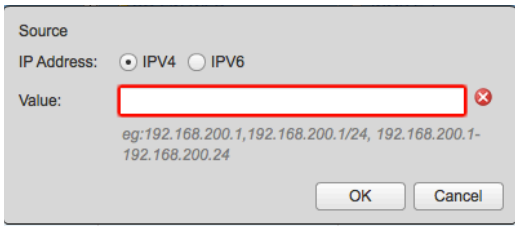

Figure 3 – Source/Destination: IPv4/IPv6 addresses or subnets.

2. vCenter objects. The following table list all possibilities:

Object	Description
Cluster	All VM/vNIC within this ESXi cluster will be selected.
Datacenter	All VM/vNIC within this Datacenter cluster will be selected.
Distributed Port Group	All VM/vNIC connected to this DVS port-group will be selected.
IP Sets	Selected IP Sets container will be used. IP Sets contains individual IP address or IP subnet or range or IP addresses.
Legacy port group	All VM/vNIC connected to this VSS port-group will be selected.
Logical Switch	All VM/vNIC connected to this Logical Switch (or VXLAN) segment will be selected.
Resource Pool	All VM/vNIC defined within the Resource Pool will be selected.
Security Group	All VM/vNIC defined within the Security Group will be selected.
vAPP	All VM/vNIC defined within the vAPP will be selected.
Virtual Machine	All VM/vNIC will be selected.
vNIC	This particular vNIC instance will be selected.

Window to enter vCenter object for Source/Destination field is displayed below:

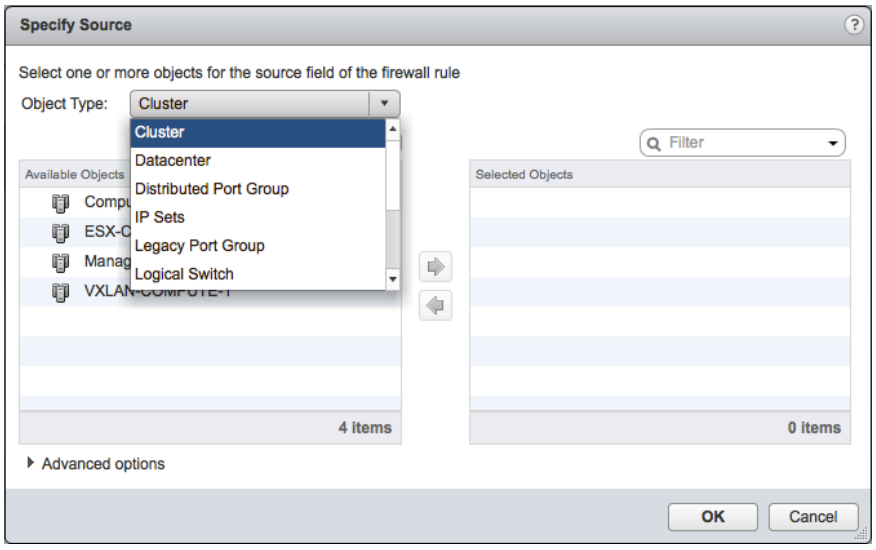

Figure 4 – Source/Destination: vCenter objects.

All permutations are possible for Source/Destination field: IP Address/Subnet and vCenter objects can be used individually or simultaneously.

Important note:
When using vCenter objects in Source or Destination field, it is mandatory to have VMtools installed on guest VM. VMtools enables DFW to retrieve IP address(es) of guest VM in order to enforce properly security rule.
If VMtools cannot be installed on a guest VM, use explicit IP address in Source or Destination field to secure traffic for this VM.

- **Service**: Protocols (TCP,UDP,..) or Pre-Defined Services or Pre-Defined Services Group can be selected.

  Protocols

  Protocols window is displayed below:

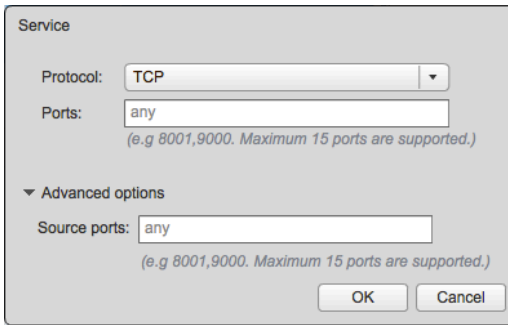

Figure 5 – Source/Destination: Protocols.

When selecting Protocols like TCP or UDP, it is possible to define individual ports number (up to a maximum of 15). Port range is not supported.

User can pick other protocols like FTP, ICMP, ORACLE_TNS, …

In the advanced options, it is possible to define source ports (up to a maximum of 15). Port range is not supported.

Pre-Defined Services

Pre-Defined Services window is displayed below:

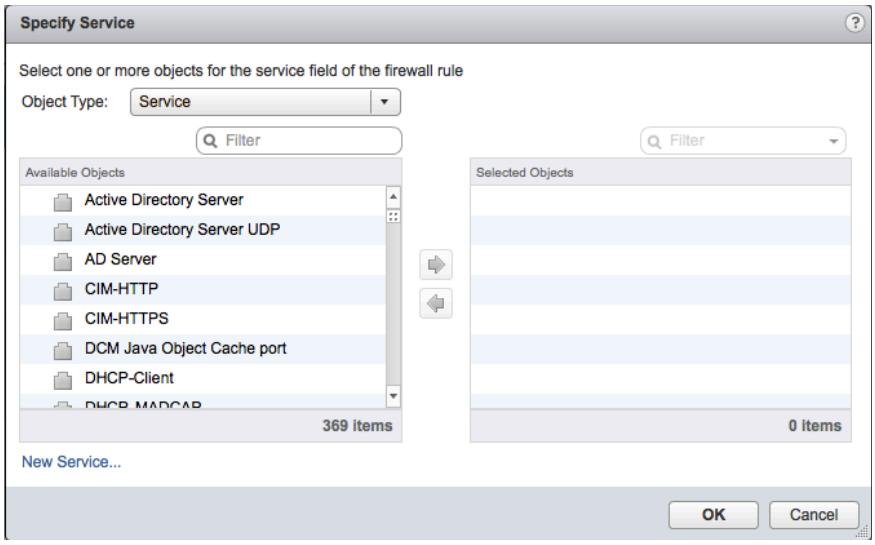

Figure 6 – Source/Destination: Pre-defined services.

User can pick any of the pre-defined service (from the long list of available objects).

It is also possible to define custom services by clicking on New Service link (in the above window):

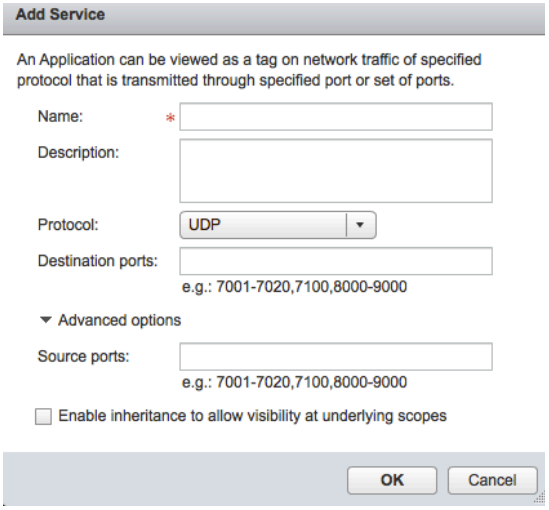

Figure 7 – Source/Destination: Add custom services.

When selecting protocol like TCP or UDP, it is possible to define individual destination ports or a range of destination ports. This is also true for source ports when expanding advanced options link.

Pre-Defined Services Group

Pre-Defined Services Group window is displayed below:

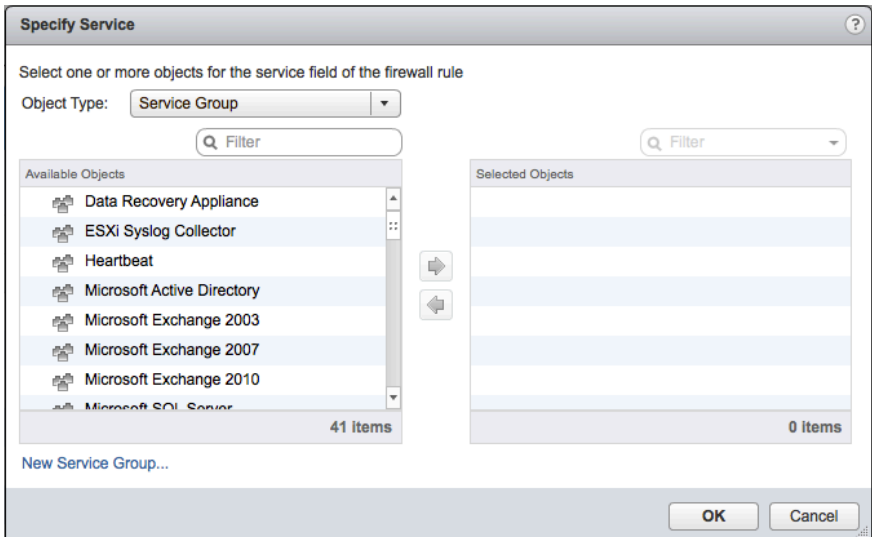

Figure 8 – Source/Destination: Pre-defined services group.

User can pick any of the pre-defined services group (from the long list of available objects).

It is also possible to define custom services group by clicking on New Service Group link (in the above window):

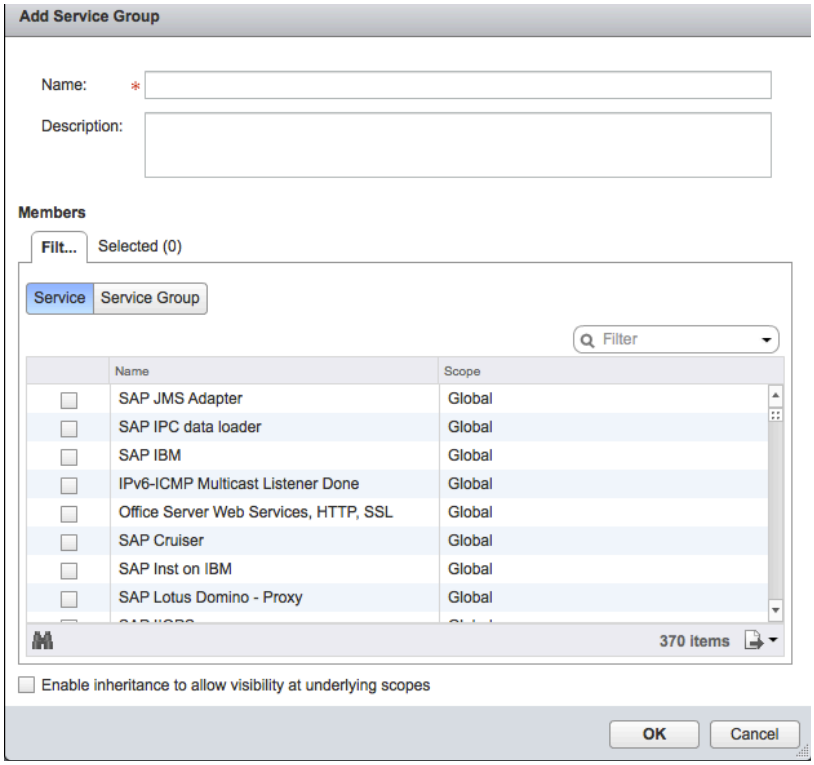

Figure 9 – Source/Destination: Add custom services group.

Custom services group is a collection of services or services group.

- **Action**: Define enforcement method for this policy rule. Available options are:

Action	Description
Block	Block silently the traffic.
Allow	Allow the traffic.
Reject (introduced since NSX 6.1)	Reject action will send back to initiator: • RST packets for TCP connections. • ICMP unreachable with network administratively prohibited code for UDP, ICMP and other IP connections.

Window to define policy rule action is displayed below:

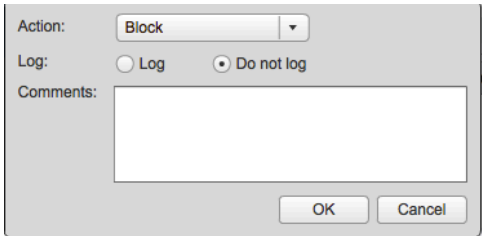

Figure 10 –Action for policy rule.

On the same window, user can decide to enable packet logging or not.

- **Applied To**: define scope of rule publishing. User can decide to publish policy rule to all clusters where DFW was enabled or restrict publication to a specific object as listed below:

Object	Description
Cluster	Selecting Cluster will push the rule down to all VM/vNIC on the ESXi cluster.
Datacenter	Selecting Datacenter will push the rule down to all VM/vNIC on the Datacenter.
Distributed Port Group	Selecting DVS port-group will push the rule down to all VM/vNIC on the Datacenter.
Host	Selecting Host will push the rule down to all VM/vNIC on the ESXi host.
Legacy port group	Selecting Legacy port group will push the rule down to all VM/vNIC on the VSS port-group.
Logical Switch	Selecting Logical Switch will push the rule down to all VM/vNIC connected on this Logical Switch (or VXLAN) segment .
Security Group	Selecting Security Group will push the rule down to all VM/vNIC defined within the Security Group.

Virtual Machine	Selecting Virtual Machine will push the rule down to all vNIC of this VM.
vNIC	Selecting vNIC will push the rule down to this particular vNIC instance.

Capability to define Security Group in Applied To field is introduced since NSX 6.1.

(Note: it is also possible to publish the policy rule to all Edge Service Gateways (ESG) or specific ESG using the appropriate option. This capability is provided since NSX 6.1.)

Window to specify Applied To field is displayed below:

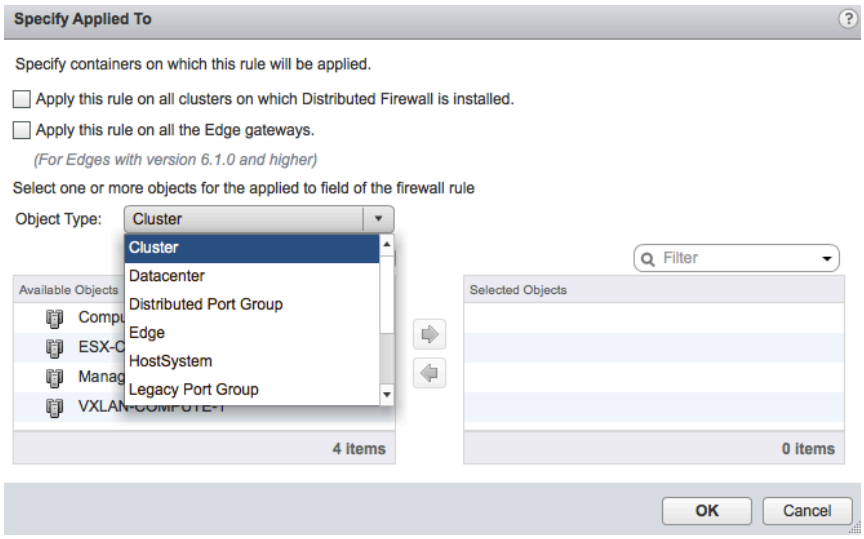

Figure 11 –Applied To scope for the policy rule.

Examples

Let's take different policy rule constructs and let's see how DFW behave for each case.

We will use the same and unique logical network topology for this purpose:

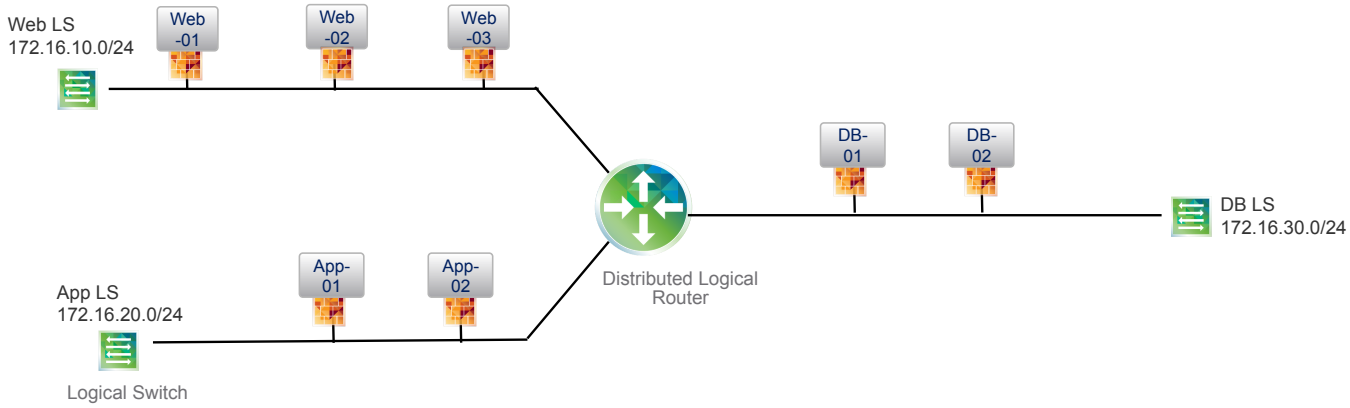

Figure 12 – Logical network base topology.

This is a standard 3-tier topology with Web/App/DB segmentation. 3 web servers are connected to Web Logical Switch (or VXLAN), 2 applications servers are connected to App LS and 2 DB servers connected to DB LS. A Distributed Logical Router is used to interconnect the 3 tiers together by providing inter-tier routing. DFW has been enabled on the ESXi cluster and as a result, each VM has a dedicated instance of DFW attached to its vNIC.

NSX offers multilple ways to define DFW policy rule configuration. Let's see how some of them look like.

Example 1: using static IP addresses/subnets in security policy rule

DFW policy configuration may look like this:

Name	Source	Destination	Service	Action	Applied To
Web to App	172.16.10.0/24	172.16.20.0/24	<Enterprise Service Bus>	Allow	All clusters
App to DB	172.16.20.0/24	172.16.30.0/24	SQL	Allow	All clusters
Default	Any	Any	Any	Block	All clusters

When using static IP addresses or subnets in policy rule, there is no need to install VMtools on guest VM. DFW engine is able to enforce network traffic access control based on the provided information.

To use this type of construct, user needs to know the exact IP information and then relay it to policy rule. This construct is quite static and does not fully leverage dynamic capabilities with modern cloud systems.

Example 2: using Logical Switch object in security policy rule

A better way to configure security policy rule is by using dynamic objects provided by vCenter/NSX manager (commonly called vCenter objects or containers).

Name	Source	Destination	Service	Action	Applied To
Web to App	Web LS	App LS	<Enterprise Service Bus>	Allow	All clusters
App to DB	App LS	DB LS	SQL	Allow	All clusters
Default	Any	Any	Any	Block	All clusters

This type of construct necessitates guest VM to have VMTools running.

Reading policy rule table is easier for all teams in the organization ranging from security auditors to architects and operations.

Any new VM connected on any Logical Switch will be automatically enforced with the corresponding security posture (for instance, a new installed Web server will be seamlessly protected by the first policy rule with no human intervention). On the same way, a VM disconnected from a Logical Switch will have no more security policy applied to it.

This type of construct fully leverages the dynamic nature of NSX.

NSX gives capability to check VM connected to a Logical Switch at any point of time. If no more VM are connected to a particular Logical Switch, it then becomes very easy to remove security policy rule dealing with this particular Logical Switch.

EXample 3: using Security Group object in security policy rule

Before writing DFW policy rule, let's first create the appropriate Security-Groups (SG).

SG name	SG definition
SG-WEB	Static inclusion: Web LS
SG-APP	Static inclusion: App LS
SG-DB	Static inclusion: DB LS

Name	Source	Destination	Service	Action	Applied To
Web to App	SG-WEB	SG-APP	<Enterprise Service Bus>	Allow	All clusters

App to DB	SG-APP	SG-DB	SQL	Allow	All clusters
Default	Any	Any	Any	Block	All clusters

All statements given for example 2 still prevail here. In fact, using Security-Groups provides much more flexibility than anything else.

Using properly dynamic inclusion, static inclusion and static exclusion, user can define in a very granular way what objects to include in this container.

Writing DFW policy rules using Security-Groups reduces dramatically number of rules needed in the enterprise and gives the most comprehensible security policy configuration.

Distributed Firewall Policy Rule Configuration Using Service Composer Menu (Option 2)

NSX Service Composer provides an application oriented mechanism to construct security policy. Given an application or a group of VMs, Service Composer enables admins to configure all the services (Firewall, Load Balancing, Traffic Redirection to 3rd party services, etc) in one place. Following discussion is more focused on configuring firewall policies using Service Composer

Click on Networking & Security -> Service Composer to access Security Policy table:

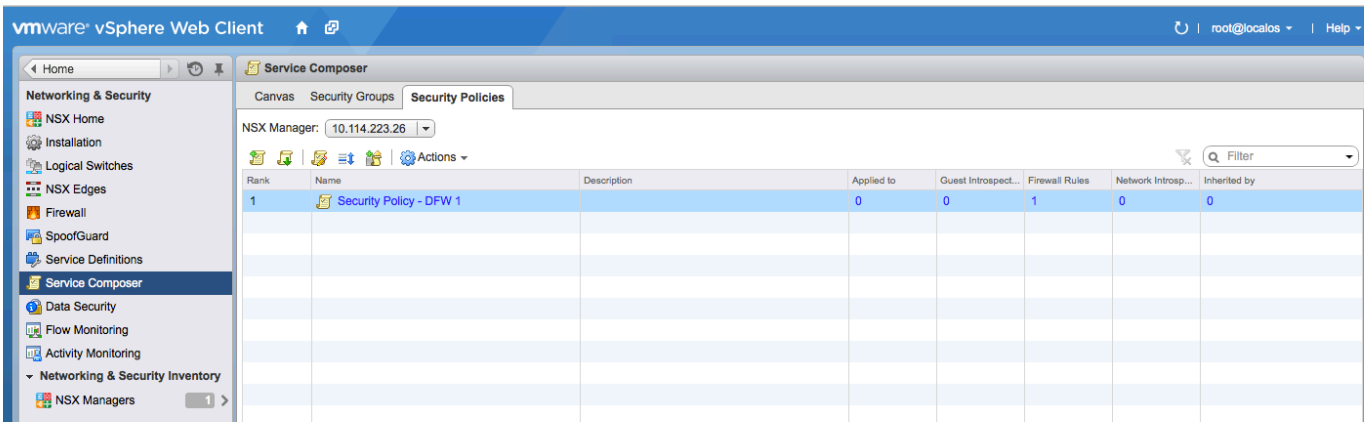

Figure 13 – Service Composer/Security Policy table.

The window displays all Security Policies (SP) defined under NSX. Selecting a particular Security Policy will open a window showing the content of the SP as shown below:

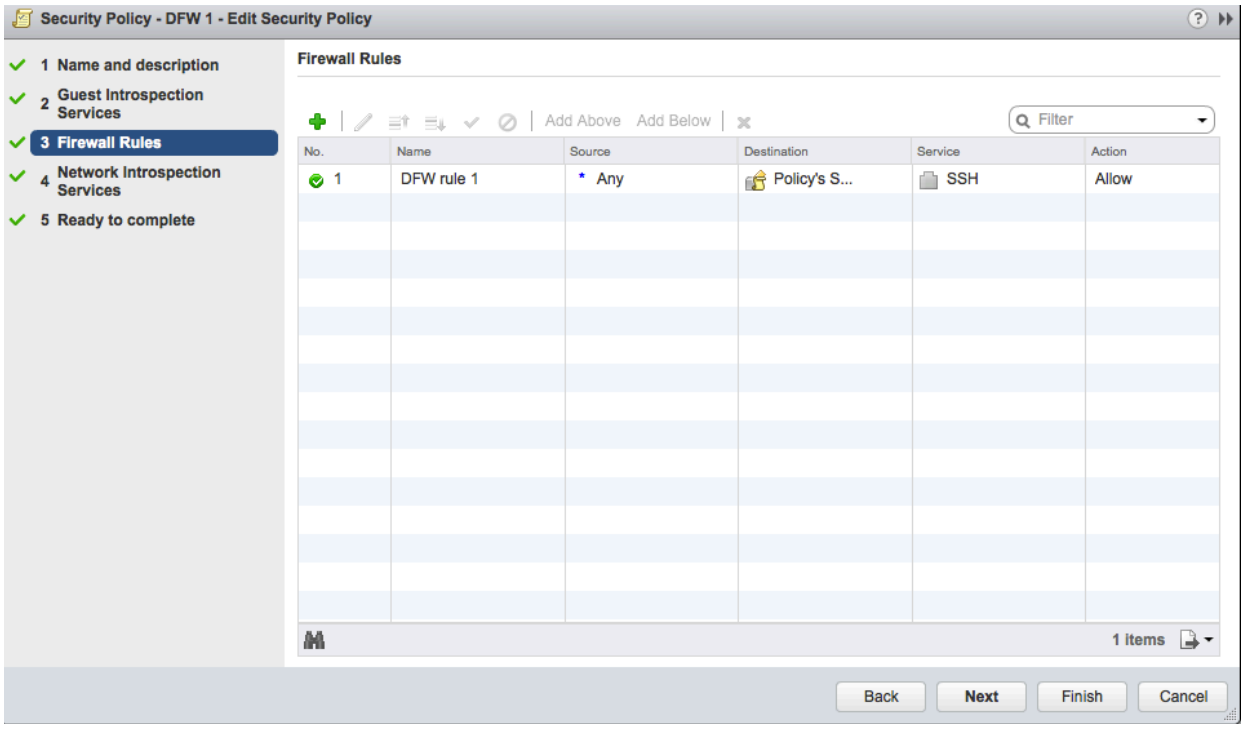

Figure 14 – Security Policy content.

Security Policy is a construct that allows specifying host and networks services enablement. For instance, guest introspection services enables guest Anti-Virus or file integrity check for guest VM; Network introspection services provides capability to specify patterns for traffic redirection to third party security vendor.

Security Policy must be applied to 1 or multiple Security-Group in order to be enforced.

Security Policy permits to define Distributed Firewall rules. Let's see the capabilities here.

To create a new policy rule, click on the + button as show in the diagram above. The following window appears:

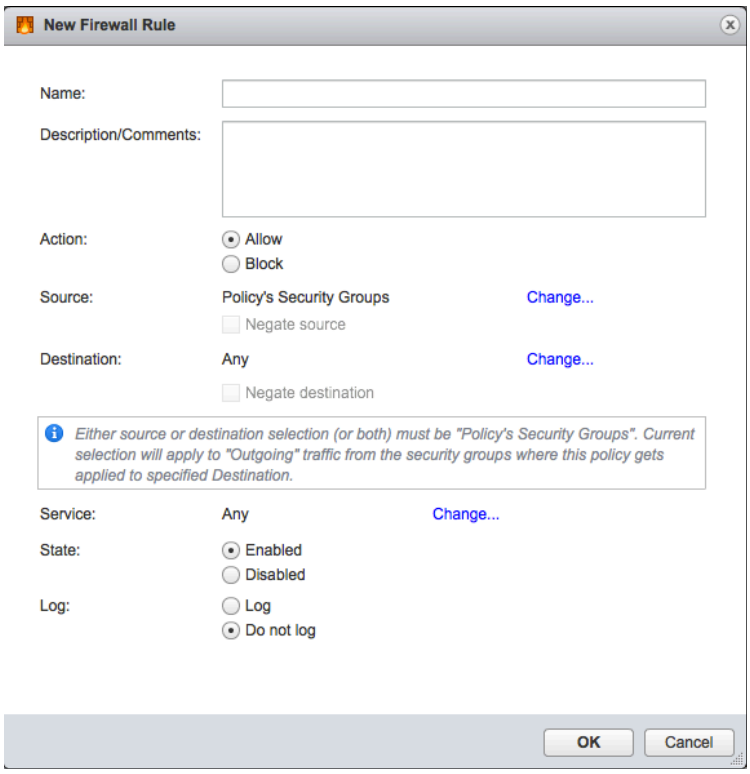

Figure 15 – Security Policy/Firewall rule.

User needs to fill the following fields:

- **Name**: firewall rule name
- **Description**: firewall rule description
- **Action**: Allow or Block
- **Source / Destination**: can take any of these values

Value	Description
Policy's Security Group	When applying the Security Policy to 1 or more Security Group (SG), this field will be internally replaced with the

	applied SG(s).
Any	Any.
Select Security Groups	User can select 1 or more Security Groups (NSX will list all available SG in the system for selection).

Following screenshot shows Source/Destination field selection window:

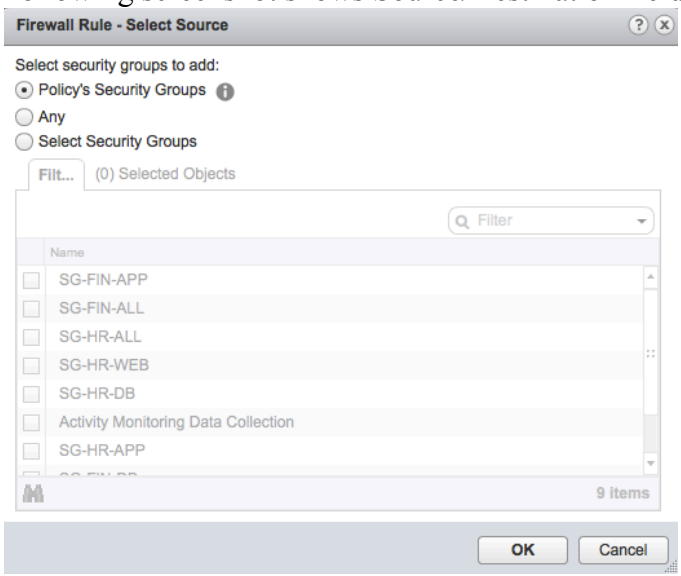

Figure 16 – Source/Destination field selection.

- **Service**: can take any of these values

Value	Description
Any	Any.
Select services and service groups	Pre-defined services and services groups. This is the same list as seen previously in the DFW policy rule section.

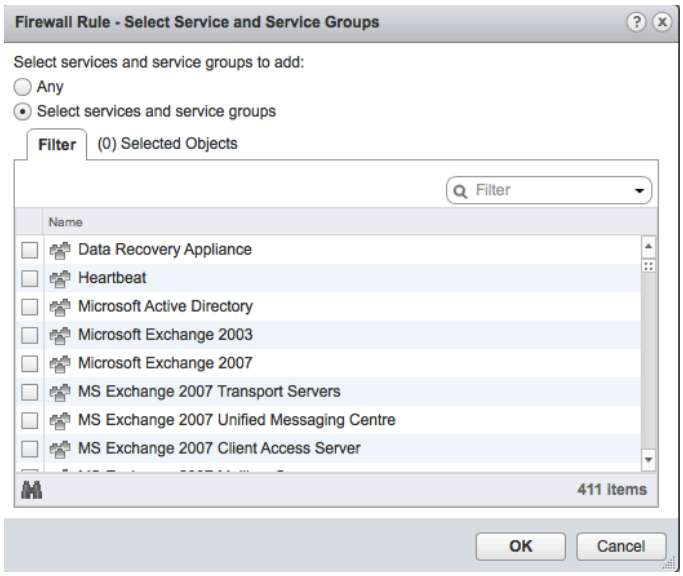

Figure 17 – Service field selection.

- **State**: Enable or disable this rule.
- **Log**: log or do not log packets matching this rule.

Example

Let's re-use the same logical network topology and see how to configure firewall (DFW) rules using Service Composer/Security Policy.

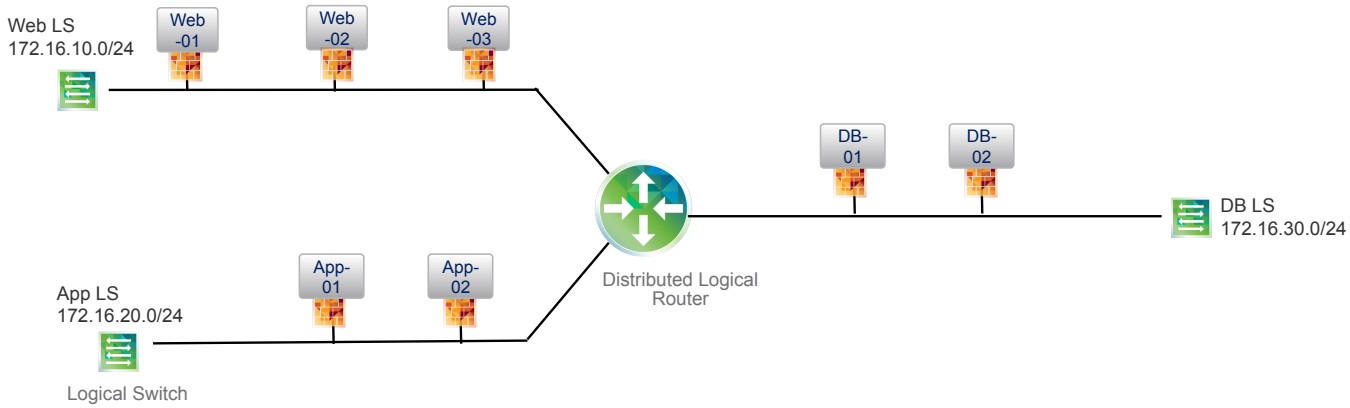

Figure 18 – Logical network base topology for the use case.

First step is to define properly Security-Groups (SG).

SG name	SG definition
SG-WEB	Static inclusion: Web LS
SG-APP	Static inclusion: App LS
SG-DB	Static inclusion: DB LS

Security Policy/Firewall rules will then look like this

Name	Source	Destination	Service	Action
Web to App	SG-WEB	Policy's Security Group	<Enterprise Service Bus>	Allow
App to DB	Policy's Security Group	SG-DB	SQL	Allow
Default	Any	Any	Any	Block

Last step is to apply the Security Policy to SG-APP to make it operational.

References

NSX-v Documentation

In addition to this document, you can read the following documents for help setting up NSX-v. All are available from https://www.vmware.com/support/pubs/nsx_pubs.html:

- NSX for vSphere Installation and Upgrade Guide
- NSX for vSphere Administration Guide
- NSX for vSphere API Reference Guide
- NSX for vSphere Command Line Interface Reference

Supporting Guides

- NSX vSphere Design Guide
- Getting Started with NSX vSphere
- Getting Started with Micro-segmentation with NSX vSphere

Contacting the NSX Technical Services Team

You can reach the NSX technical services team at http://www.vmware.com/support.html.