TECHNICAL DESCRIPTION

Nuage Networks Virtualized Services Platform: Service chaining

nuagenetworks
From Nokia

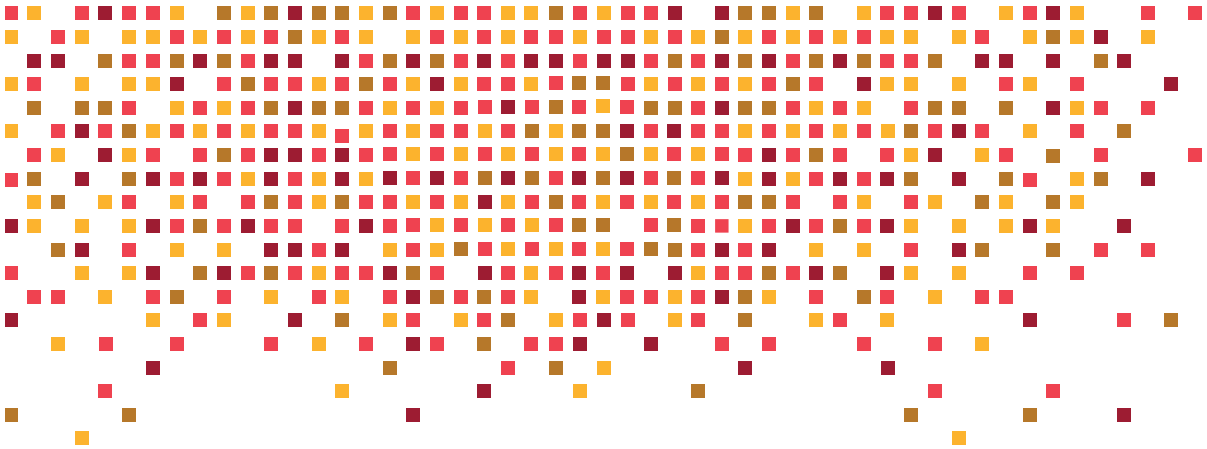

CONTENTS

Nuage Networks Virtualized Services Platform

The following section describes the Nuage Networks Virtualized Services Platform (VSP) solution for Software-Defined Networking (SDN).

VSP solution overview

The Nuage Networks Virtualized Services Platform (VSP) is an industry-leading SDN solution that leverages Network Virtualization Overlay (NVO) technologies to make the network as readily consumable as the compute resources in a cloud environment. The VSP achieves this by ensuring rapid and efficient delivery of highly customizable application services in and across multitenant datacenters. The VSP enables the deployment of massively scalable cloud-based services over an existing IP network fabric with the agility and performance demanded by highly dynamic application environments.

The main components of the VSP solution are:

- Virtualized Services Directory (VSD)
- Virtualized Services Controller (VSC)
- Virtual Routing and Switching (VRS)

FIGURE 1. Nuage Networks VSP solution components

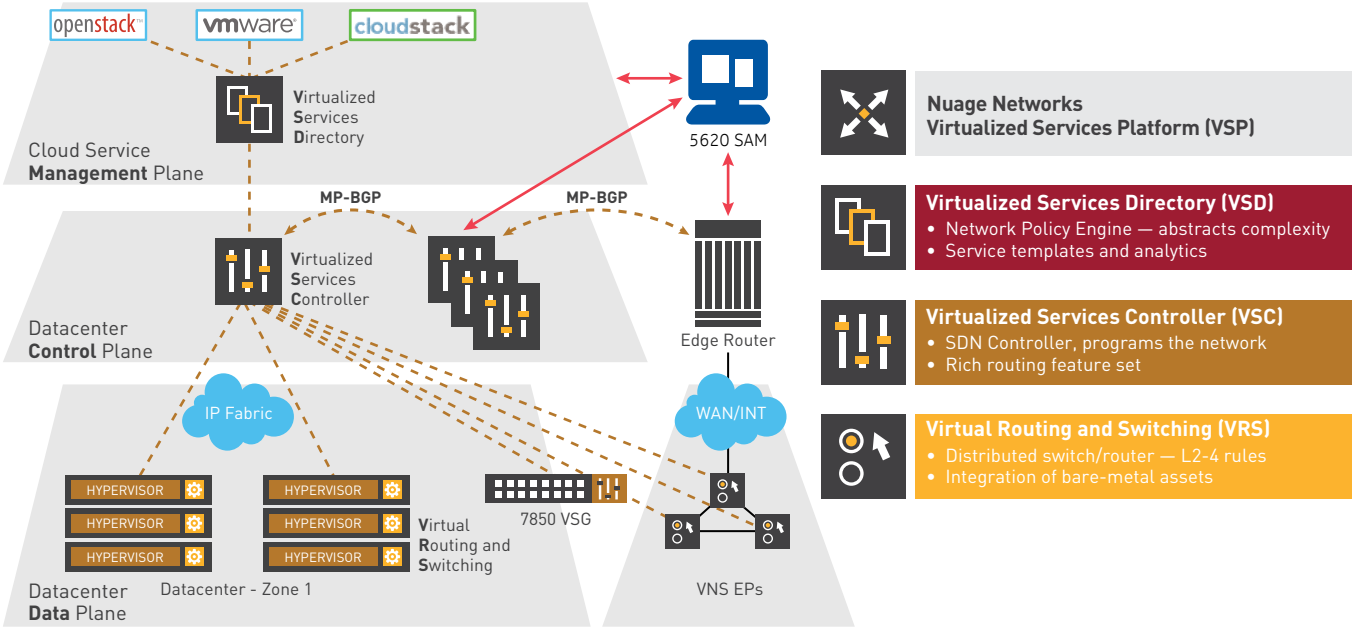

The VSP solution is based on the three proven networking principles depicted in the following figure.

FIGURE 2. Guiding network principles

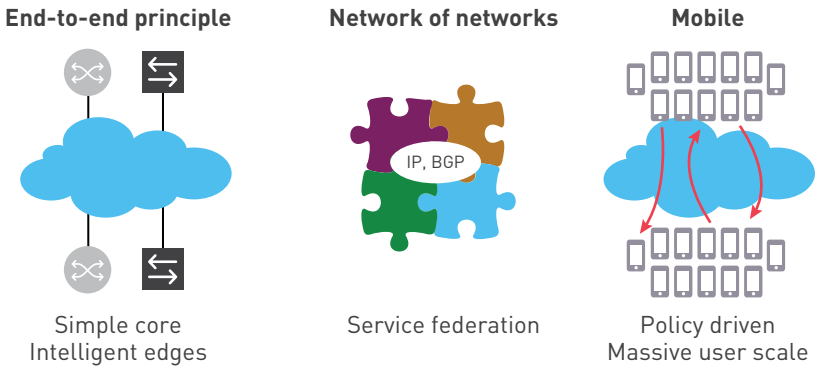

The "simple core, intelligent edge" principle has been successfully applied for more than 10 years in VPN technology deployed in mission-critical networks around the world. The core of the network handles only IP tunneling and does not maintain per-Virtual Machine (VM) state information, enabling better scaling and a longer hardware refresh cycle.

The datacenter/service provider PoP network is not an island; it must plug easily into a "network of networks", including existing networks, the public internet and VPNs. IP and BGP are currently being used to achieve ubiquitous connectivity at massive scale.

Finally, the VSP uses policy-driven auto-instantiation of network connectivity, which allows customers to establish network connectivity at the same speed as the virtualized compute while supporting a large number of attached VMs. Administrators may need to auto-instantiate hundreds or even thousands of compute nodes using APIs and be able to relocate them for disaster recovery or for proximity to end users. Policy-driven auto-instantiation has allowed millions of cellular phones to be connected to the existing networks while on the move without any operator intervention and is now being applied to connectivity for mobile VM endpoints.

Virtualized Services Directory (VSD) – management plane

The Virtualized Services Directory (VSD) is a programmable policy and analytics engine. It provides a flexible and hierarchical network policy framework that enables administrators to define and enforce resource policies in a user-friendly manner.

The VSD contains a multitenant service directory which supports role-based administration of users, compute, and network resources. It also manages network resource assignments such as IP addresses and ACLs.

For service assurance, the VSD allows administrators to define sophisticated statistics rules such as collection frequencies, rolling averages and samples, and Threshold Crossing Alerts (TCAs). A TCA triggers an event that can be exported to external systems through a generic messaging bus. Statistics are aggregated over hours, days and months and stored in a Hadoop® analytics cluster to facilitate data mining and performance reporting. The VSD runs as a number of processes in a VM environment.

Virtualized Services Controller (VSC) – control plane

The Virtualized Services Controller (VSC) is the industry's most powerful SDN controller. The VSC functions as the robust network control plane for the SDN network, maintaining a full view of per-tenant network and service topologies. Through the VSC, virtual routing and switching constructs are established to program the network forwarding plane (the Virtual Routing and Switching (VRS) component) using the OpenFlow™ protocol. Multiple VSC instances can be federated within and across datacenters/service provider PoPs by leveraging MP-BGP, a proven and highly scalable network technology. The VSC leverages the industry-leading, carrier-grade, and field-proven Alcatel-Lucent Service Router OS (SROS) platform that is widely deployed in major carrier networks globally.

The VSC communicates with the VSD policy engine using eXtensible Messaging and Presence Protocol (XMPP). An ejabberd XMPP server/cluster is used to distribute messages between the VSD and VSC entities. Multiple VSC instances can be federated within and across DCs by leveraging the BGP protocol.

Virtual Routing and Switching (VRS) – data plane

The Virtual Routing and Switching (VRS) component is an enhanced Open vSwitch (OVS) implementation that constitutes the network forwarding plane. The VRS user space module is installed directly into the server hypervisor and supports both Layer 2 and Layer 3 networking capabilities. The VRS supports multiple hypervisor types in virtualized server environments and can operate as a gateway for bare metal servers or service appliances.

Leveraging NVO tunneling technologies, the VRS encapsulates and de-encapsulates user traffic, enforcing L2L4 traffic policies as defined by the VSD. When using Virtual eXtensible Local Area Network (VXLAN), the VRS originates and terminates VXLAN tunnels by acting as the VXLAN Tunnel Endpoint (VTEP). By pushing the networking intelligence directly into the hypervisor, the VRS delivers the most efficient network forwarding solution while simultaneously eliminating unnecessary tromboning of traffic.

Additionally, the VRS tracks VM creation, migration and deletion events to dynamically adjust network connectivity.

The following figure illustrates the logical topology of the VRS:

FIGURE 3. VRS data plane logical topology

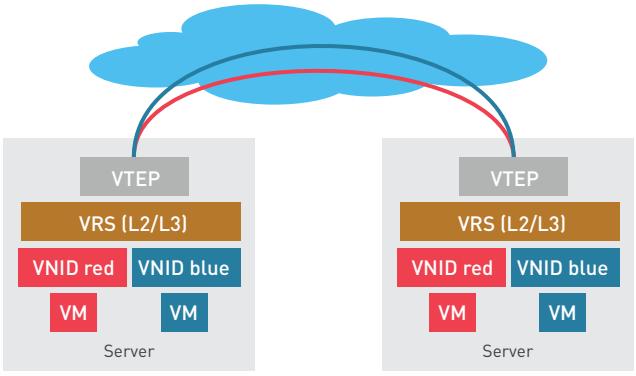

The VSP supports service chaining in both virtual and physical environments within a datacenter. Additionally, Nuage Networks also offers the Virtualized Network Services (VNS) solution that can extend SDN technology to the customer premises. Service chaining can be used to insert services between endpoints, services, or within the traffic flow of a service.

Service auto-instantiation

The Nuage Networks VSP provides automatic instantiation of VRS network services using a policy driven approach similar with the procedures used in mobile networks. The case of a cloud network implemented with the VSP is depicted in the following figure.

FIGURE 4. Use case for VRS service auto-instantiation

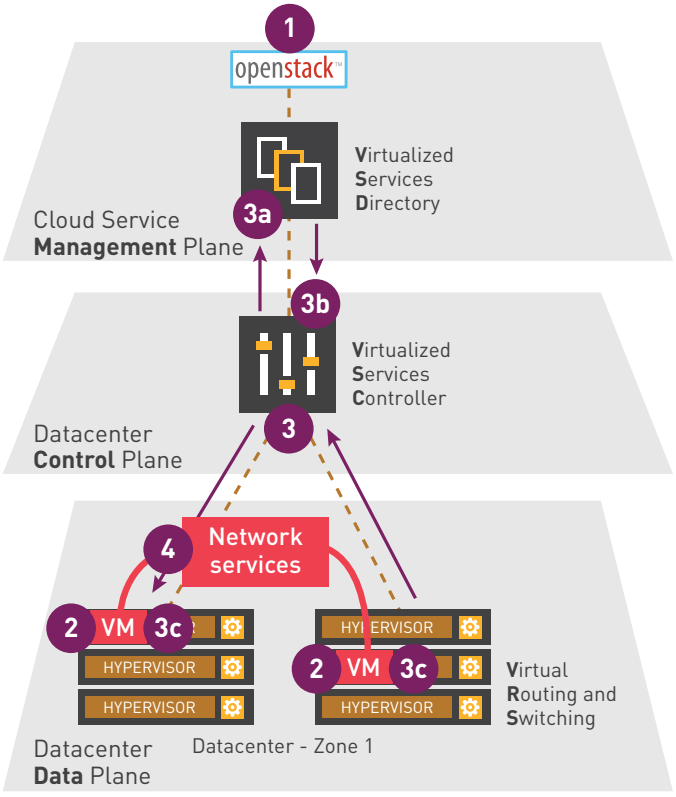

1. OpenStack receives request for compute assets
2. VM instantiated on hypervisors
3. Event triggers VRS which informs VSC of VM placement
   a. VSC queries VSD on policy
   b. VSD issues VSC with network service template
   c. VSC deploys policy to applicable VRSs
4. Network services are created based on policy from VSD

The red VMs are enabled in the servers located in DC1, which could be a service provider PoP with compute nodes. A cloud management system such as OpenStack is then used to create a VM. The local VRS agent intercepts the event and informs the VSD through its local VSC. The VM profile is included in the report sent to the VSD. The VSD authenticates the request to identify which VSD service template should be used and downloads the required network attributes to the VSC.

As new VMs are attached to the servers, the VRS agent intercepts the event and distributes the VM profile to the VSD through its associated VSC as indicated by the red arrows. Upon reception, the VSD authenticates the request, identifies the required service profile and sends the required service creation command together with related attributes (service id, RT, RD, QoS, stats, etc.) to the VSC. The VSC service manager instantiates new service instances as required, recalculates the FIBs and downloads the required information to the VRS agent using the OpenFlow protocol.

The process is repeated for every new VM to enable data plane forwarding between the red VMs located potentially in different racks throughout the datacenter/service provider PoP. If external connectivity is required, the VSC uses BGP to exchange IP routes with the gateways located at the datacenter edge. BGP is also used to exchange information between multiple VSCs enabling easy expansion of VRS services across datacenter zones.

Logical service view

The VRS implements L2 and L3 multi-tenancy using a distributed architecture as shown in the following figure .

FIGURE 5. VRS logical view

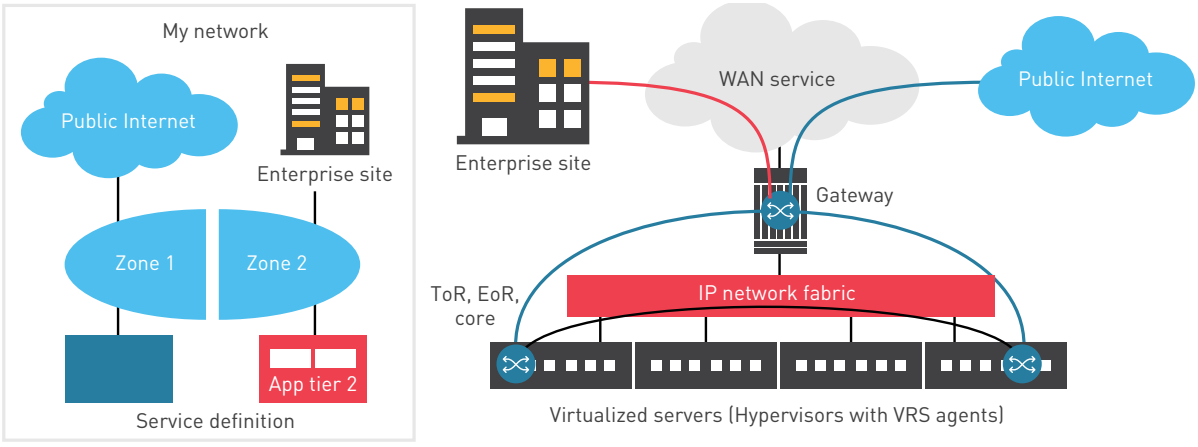

ToR - Top of Rack switch
EoR - End of Row switch

The enterprise or tenant topology in this example consists of two application tiers (red and green), each belonging to two different policy zones. The blue networking domain interconnects the two application tiers and provides external connectivity to the Internet and to the VPN domain. It also provides enforcement of policies (ACLs, QoS, stats) for each individual domain on a per-VM basis.

The VSP maps the tenant topology on the left to specific network primitives and attributes. As VMs get created it uses the procedure described in the previous section to program the data path to support packet forwarding.

In the example on the right, there is a pair of green and red VMs instantiated in two servers located in different areas of the datacenter. In each server for this tenant logical topology, a blue VRS instance gets auto-instantiated as VMs are added to one of the two application containers. The blue VRS instances represent the L2 and L3 intelligence required to handle the packets forwarded between VMs and to the gateway to external domains. They connect themselves using server-friendly encapsulation and to external datacenter gateways using standards-based VXLAN overlay tunnel technology. Optionally, the VSP platform supports VPN over GRE to connect externally using legacy datacenter gateways that do not support VXLAN. Both encapsulations require only standard IP in the datacenter network. The resulting IP fabric enables reuse of existing datacenter network and eliminates vendor lock-in. Inner encapsulation details are explained in the next section.

Data plane

Virtual eXtensible Local Area Network (VXLAN) has become the predominant data center encapsulation technology. Most NIC vendors are implementing VXLAN hardware-assisted processing on the servers. As a result, the performance on hypervisors will be significantly improved as compared to other encapsulation methods, such as MPLS/GRE. The Nuage Networks VSP solution has supported VXLAN since release 1.0. It is important to support the same data plane encapsulation on the datacenter gateway to allow for a seamless interconnect between hypervisors and the WAN.

Since release 12.0 of the industry-leading Alcatel-Lucent SROS, the Alcatel-Lucent 7750 Service Router (SR), 7450 Ethernet Service Switch (ESS), and 7950 Extensible Routing System (XRS) datacenter gateway products have provided support for VXLAN data plane transport tunnels that can be terminated on VPRN or VPLS services. Those VPRN and VPLS services will also provide a data plane interworking function between the VXLAN data plane supported in the datacenter and the MPLS data plane supported in the WAN.

The following figure provides an example of traffic flow forwarding from the VRS (installed on a server hypervisor) in the data plane.

FIGURE 6. VRS data plane

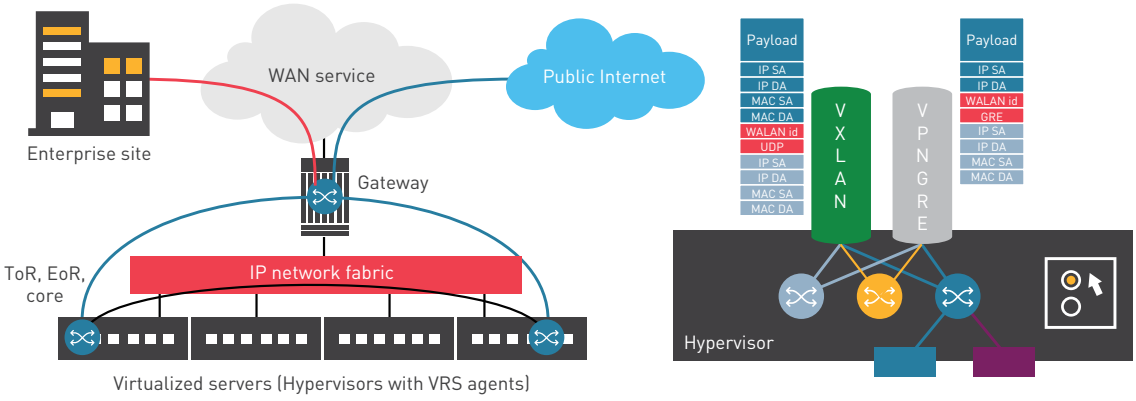

The VRS service on the left, used to network the green and red VMs, is implemented using the service primitives shown on the right. In each of the two hypervisors there is a blue VRS instance used to isolate the tenant quintuple flows from other tenant instances and to provide internal and external connectivity for the green and red VMs. Two other tenant instances share tunnels with the blue VRS: one or more VXLAN tunnels to other hypervisors and one more VPN tunnels to the gateway(s). Each VRS instance is instantiated as a combination of L3 and L2 forwarding entities:

- one distributed L3 Virtual Routing and Forwarding (VRF) instance
- one VXLAN virtual bridge for every VM subnet implementing the MAC FIB
- every VXLAN is represented in VRF as an IRB interface

As packets are received from the VM VNICs they are processed by the associated VRS instance: ACLs are evaluated, L2 and optionally L3 lookups are performed to determine how the rest of the packets in the flow should be treated. The resulting next-hop could be either a local VM VNIC or a tunnel:

- If the next-hop is another hypervisor IP, a VXLAN header is added to the packet (see IETF VXLAN draft: http://tools.ietf.org/html/draft-mahalingam-dutt-dcops-vxlan-01)). The VXLAN context might change if the destination is in a different subnet.

- If the next-hop is a gateway, an IP VPN over GRE header is used (see RFC 4797, "Using IP/GRE tunneling in IP VPNs": http://tools.ietf.org/html/rfc4797)
- The additional encapsulation consists of a regular MAC plus IP followed by UDP (VXLAN) or by GRE (VPN/GRE) and a tenant identification field (VXLAN id or VPN label).

In the reverse direction, packets received from the tunnels are mapped to the blue VRS instance using the VXLAN id or VPN label field. The FIB lookup determines then the local VNIC(s) to which the packet needs to be forwarded after the tunnel's encapsulation is removed.

The datacenter physical network forwards the packets using the IP router header and it does not need to understand the rest of the encapsulation or keep individual VM state.

The forwarding knowledge for each VRS is programmed in the VRS agent of each hypervisor by the VRS policy and control plane intelligence as described in the following section.

Control plane

Regular IP routing is required to provide support for core tunneling. In addition to the core control plane, a service management and control plane is used to perform VM auto-discovery, ACLs, RIB and FIB population for each service instance. The required components are shown in the following figure.

FIGURE 7. Management and control plane

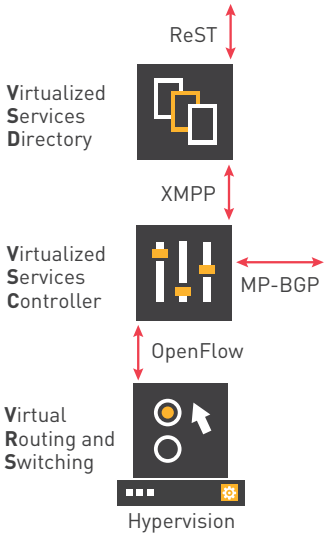

The three components of the VSP work together using a number of protocols:

- XMPP is used for VSD-VSC event communication and policy exchanges.
- OpenFlow is used for VSC-VRS event communication and FIB download.
- MP-BGP is used to exchange information between VSCs and external networks.

The VSP system offers a northbound interface based on REST APIs that can be used by external cloud management or cloud orchestrator systems to consume the networking services.

The VRS service is instantiated with no operator intervention using the high-level procedure described in the initial section. As VMs are instantiated or removed, the VRS agent sends the VM profile associated with the event to the VSD through the VSC. The VSD authenticates the VM user and sends the VM service attributes to the VSC.

The VSC service manager has a complete view of the tenant local topology and takes the following actions:

- Exchanges information using MP-BGP with other VSCs and with the gateways
  - MP-BGP IP VPN SAFI (see IP VPN as described in IETF RFC4364) is used for gateway communication and programming of VPN/GRE encapsulation and tunnels.
  - MP-BGP EVPN SAFI (see IETF RFC7432: "BGP MPLS-Based Ethernet VPN": http://datatracker.ietf.org/doc/rfc7432 and IETF: "A Network Virtualization Overlay Solution using EVPN": http://datatracker.ietf.org/doc/draft-ietf-bess-evpn-overlay/) is used for exchanging information with other VSCs and programming of VXLAN encapsulation.
- Generates a complete topology view for each VRS
- Downloads ACL, L2 and L3 FIB updates and other policy information (QoS, statistic collection) to the VRS agent in each hypervisor that has at least one VM belonging to the tenant

As VMs get added or removed, the related VRS agent(s) are fully programmed as a result with all the information required for local VMs:

- ACLs, L2 and L3 FIBs
- ARP tables
- QoS marking and policing
- Required frequency of statistics-gathering

As new flows get activated, the VRS agent programs the flow tables in the kernel without VSC involvement. It also handles DHCP and ARP requests from the local VMs.

VM mobility

In many scenarios, VMs need to be moved with minimal disruption to a new location. The VM's related state is copied over to ensure the previous communication sessions are maintained. The copied state also includes the VM's networking attributes: VM IP and MAC addresses, ARP table content, and service attributes (QoS, statistics).

The Nuage Networks VSP is able to automatically track VM movement. As soon as the VM move event is intercepted, all three components (VRS, VSC and VSD) become aware of the VM state. The service is instantiated at the new location and, as soon as the VM is re-activated, the local VRS agent moves the service profile to the new hypervisor. If the new location is on a new VSC, MP-BGP is used to advertise the new location to other VSCs or to external networks.

The VRS agent and the VSC at the former location will remove all stale information from the local tables. The VSC will use an MP-BGP withdraw message to remove the old entries from other VSC tables or from external networks.

Service chaining in the Nuage Networks VSP

To generate additional revenue from Value-Added Services (VAS), service providers need to be able to steer traffic to a number of service functions such as firewalls, load balancers, NAT, and IPS/IDS systems within their datacenter or service provider PoP networks. Organizations want the ability to specify Virtual Network Functions (VNFs) or Physical Network Functions (PNFs) and their sequence, so service functions can be added or removed seamlessly without requiring changes to the underlying network infrastructure.

This network sequencing of service functions is known as service chaining. It is accomplished using NVO technologies (e.g. VXLAN) and Policy-Based Routing (PBR) technologies. Service chains have a number of use cases:

- To steer traffic to VNFs or PNFs in the network
- To attach separate functions to an application, especially when providing these functions as services to the cloud environment, for example Firewall-as-a-service (FWaaS), or Load-Balancer-as-a-Service (LBaaS)
- To manage application compliance with security rules
- To enable separation of concerns between the team doing the VM and application design, and the team managing the security and policy surrounding the application
- To provide a single interface to manage function attachment and configuration
- To provide datacenter-based functions (e.g. firewalls) for connecting branches in the WAN

The Nuage Networks VSP supports service chaining in both virtual and physical environments within a datacenter, as well as Layer 3 and Layer 2 SDN networking.

Use case: Service chaining in a multitier application

This section describes an example of a basic use case supported by the Nuage Networks VSP service chaining solution.

A multitier application may involve a service chain, as shown in the following figure.

FIGURE 8. Service chain design

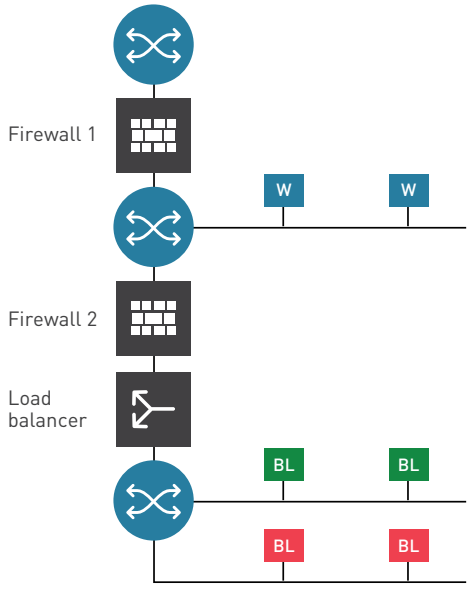

The service chain specifies that traffic between the internet and web (W) tiers will traverse firewall 1, and traffic from the web to business logic (BL) tiers will traverse firewall 2 and a load balancer.

Using templates for service chaining

The VSP service chaining solution is designed to enable the development of templates by security and networking teams where appliance functions running on dedicated hardware or virtual machine resources may be inserted between compute nodes running different applications.

The use of templates for service chaining is shown in the following figure.

FIGURE 9. Use of service chain template

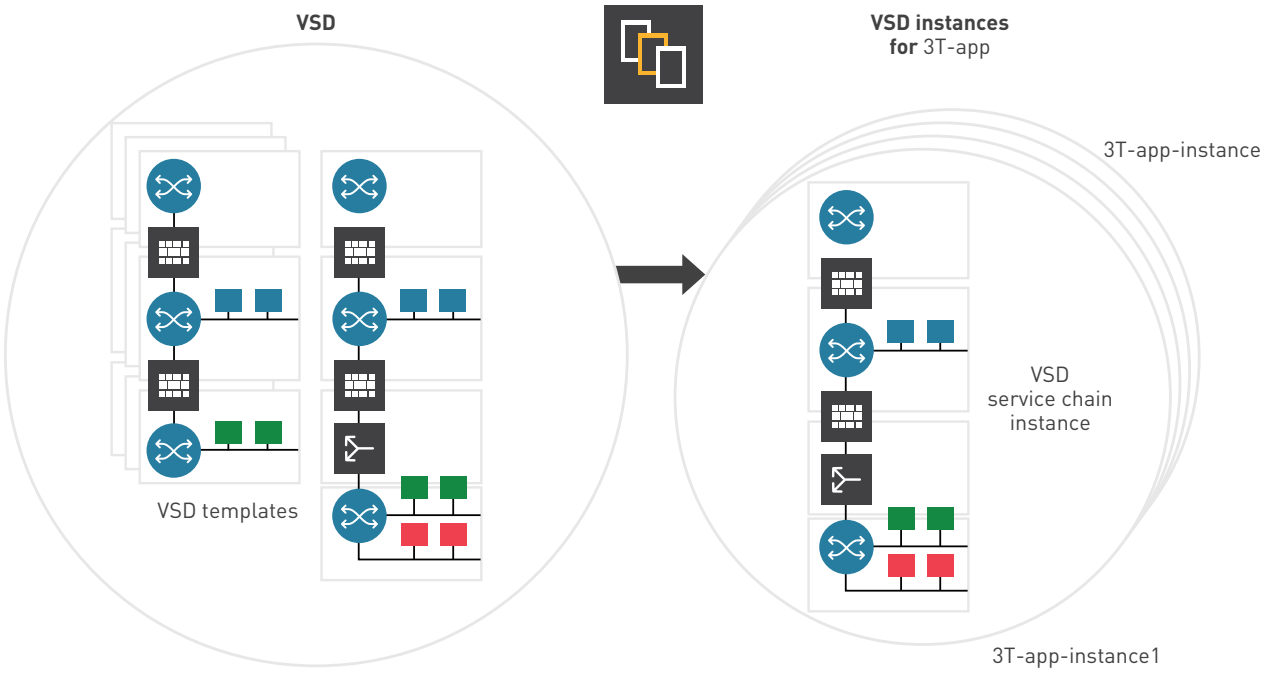

The Virtualized Services Directory (VSD) allows the cloud service provider and end-customer administrators to define service chain templates like the 3T-app made available to different groups of IT users based on their assigned permissions.

These templates can be used to instantiate multiple instances as shown in the figure (3T-app-inst to 3T-app-instn). The template system enables decoupling of responsibilities, providing to security and networking teams tools for increasing service velocity and enforcing policy compliance. It also offers end users an intuitive abstraction-based interface designed to increase service consumption.

Workflow for instantiating service chaining

Service chain instantiation involves the following workflow:

1. Cloud service provider or customer administrators:
   - Create Nuage domain templates.
   - Add virtual port tags (Redirection-Targets) for appliance attachment.
   - Add appropriate advanced forwarding redirect rules to direct traffic to appropriate appliances.
   - Assign permissions for users to use templates.
   - At this point, templates like 3T-app are available for instantiation.
2. Customer administrators:
   - Instantiate the service chains from (assigned) templates. This creates service chain instances like 3T-app-inst1.
   - Create VPorts for appliance resources, and map them to Redirection-Targets.
   - Instantiate appliances on VMs or physical resources. This is a mandatory step for certain types of appliances such as firewalls, which must be in place before the applications are enabled on compute nodes. The VSP will automatically associate the appliance attachment points to VPorts and instantiate the required service connectivity for appliance attachment.
3. Users:
   - Create application VMs and map them to network services provided by the VSP.

The VSP detects VM creation and automatically instantiates the required PBR rules to direct the new VM traffic through the service chain.

The Nuage Networks solution allows additional functions and appliance instances to be inserted in the chain at a later time by repeating steps 1 and 2. This may be required because a new function (e.g. DPI) needs to be added to the chain or to expand the resources used to implement a certain appliance function (e.g. a new firewall instance). The VSP automatically adjusts the PBR rules across hundreds of other attachment points for devices already in the chain, without operator intervention.

Note that the above workflow decouples the responsibilities of administrators and users. Users are not involved in the design and instantiation of service chains. Instead, they simply create the compute resource required for their applications. The VSP takes care of inserting the new compute node in the service chain, automatically instantiating the required PBR match criteria.

Auto-instantiation of service networking in a service chain using the VSP

This section describes how the VSP auto-instantiates service networking primitives to provide steering through the service chain based on abstractions designed in a cloud management system (e.g. OpenStack) or in the VSD Architect. In addition to the regular service networking, the service chain involves special steering of some of the traffic as shown in the following figure.

FIGURE 10. Service chain networking

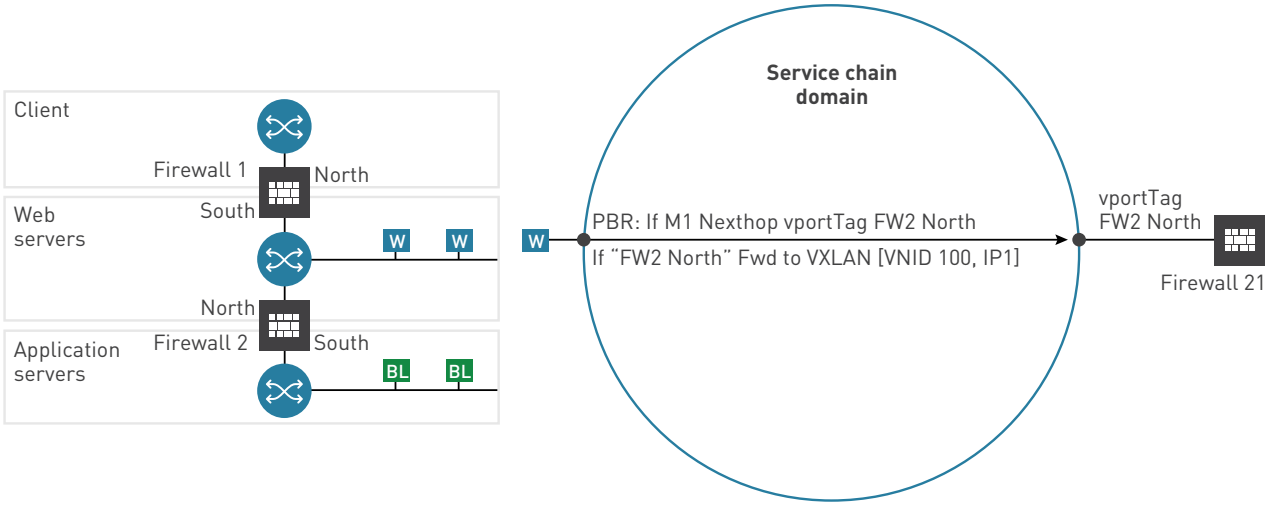

The blue domain in the diagram is configured in the VSD to provide the service chaining topology depicted on the left side. Regular VRS forwarding rules are employed to emulate the routers in the 3-tier topology. Whenever a certain type of traffic is required to pass through firewalls FW1 or FW2, a set of PBR rules is configured in the VSD (e.g. "If web traffic matches M1 criteria, forward to FW2 North Redirection-Target.") The use of Redirection-Targets is described in the section called "Using the VSD Architect" below.

In order to implement these kinds of rules, a special type of forwarding needs to be instantiated in the service chain domain: if ingress traffic from web (W) servers matches M1 criteria, PBR forwarding needs to send the flows straight to the FW2 north interface, skipping all the regular forwarding in the blue domain. This section describes how the related networking is instantiated using the VSP for the traffic from a web server to the FW2 North interface. The same principles apply to the other traffic flows.

As soon as the service chain design is instantiated from the VSD template by the administrator, the VSD automatically assigns a number of networking parameters required to instantiate the blue service chain domain.

When the appliance functions (FW1, FW2) are instantiated, the VSD automatically assigns networking identifiers for the associated Redirection-Targets. Specifically, VXLAN VNID 100 is associated with Redirection-Target "FW2 North". The appliance attachment point or location on the hypervisor or gateway identified by IP1 is discovered using regular VSP procedures as soon as FW2 becomes active on a VM or as a gateway attachment.

The networking parameters for both domain and Redirection-Target are then pulled from the VSD by the local VSC and used to instantiate the required VRS service instance. In this service chaining example, the Redirection-Target "FW2 North" is associated with VXLAN VNID100 and IP tunnel IP1. BGP EVPN is then used to advertise the association to remote VSCs interested in the blue VRS domain.

When applications are instantiated on VMs, the networking service parameters including the PBR rules are pulled and instantiated by the local VSC. In addition to the regular VRS procedures, the PBR action "forward to Redirection-Target FW2 North" is translated to the following data plane actions: "forward (and encapsulate) to VXLAN VNID 100 and tunnel IP1."

The end result is that whenever match criteria M1 are met, traffic from the web server is forwarded straight into the VXLAN tunnel identified by VNID 100 and IP1 destination. At the receiving end, on hypervisor IP1, this traffic is forwarded in the blue domain instance without any additional lookup, straight to the FW2 North interface.

Using the VSD Architect to set up service chaining

This section explains how to set up service chaining with the VSD Architect user interface.

Create a multitier L3 domain template

1. From the Networks menu, select the Layer 3 Domains submenu, and click + to create a new L3 domain.
2. Modify the Name and Description fields to the desired values, then click Create.

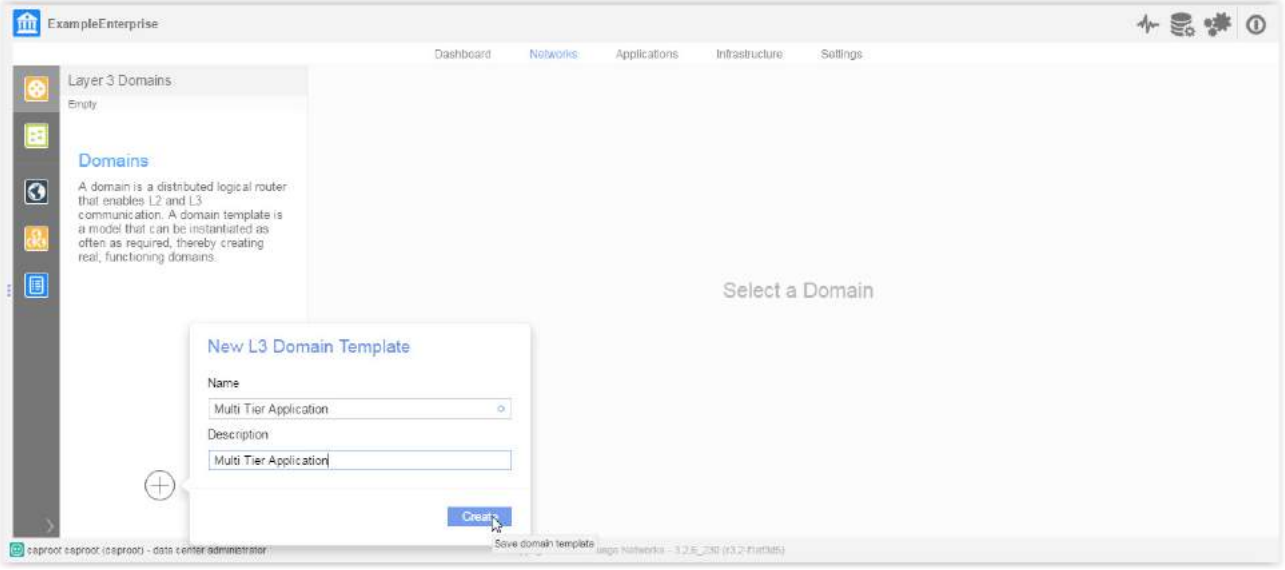

3. Confirm that the L3 domain template has been created.

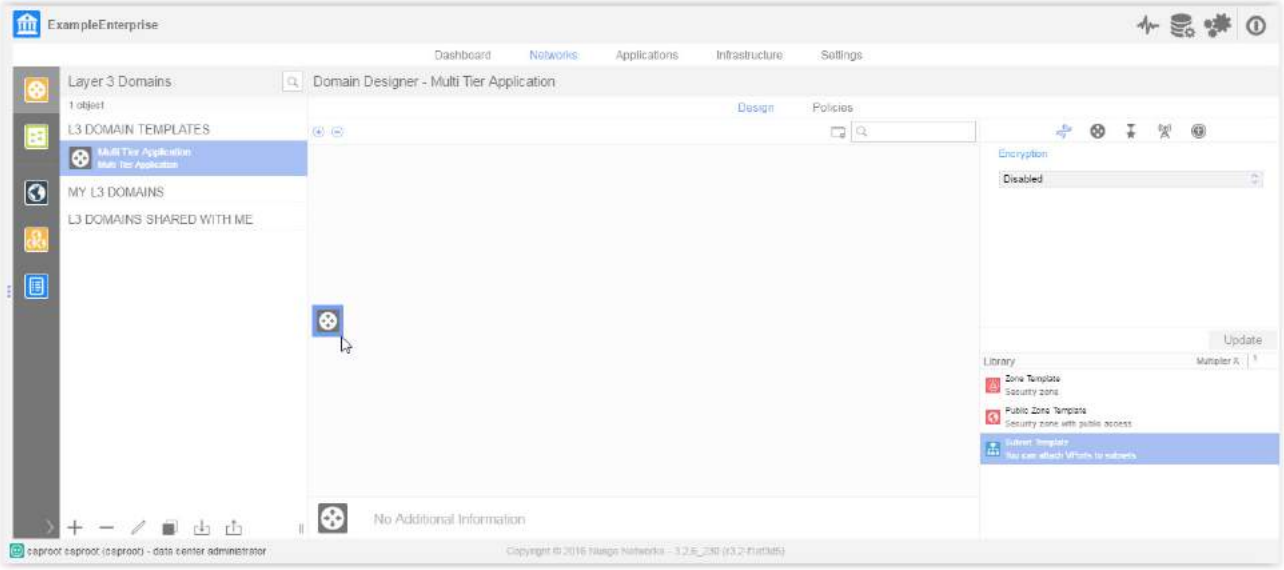

4. Add a Nuage zone object to the L3 domain by dragging the Zone Template object onto the L3 Domain object.

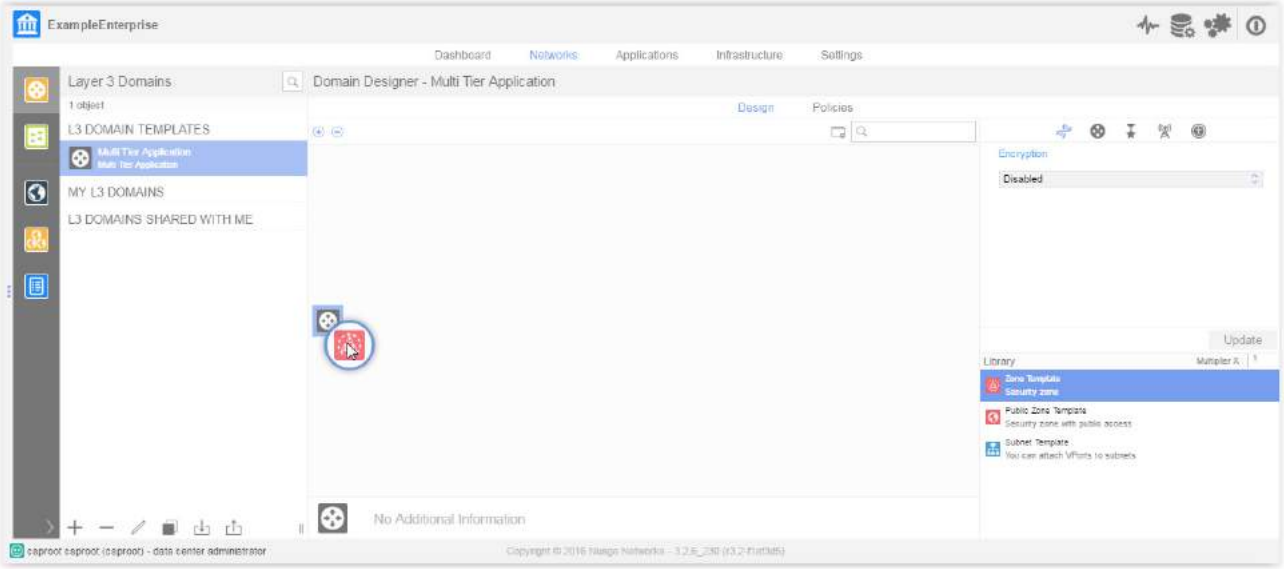

5. Select the zone object, modify the Name and Description fields to the desired values, then click Update.

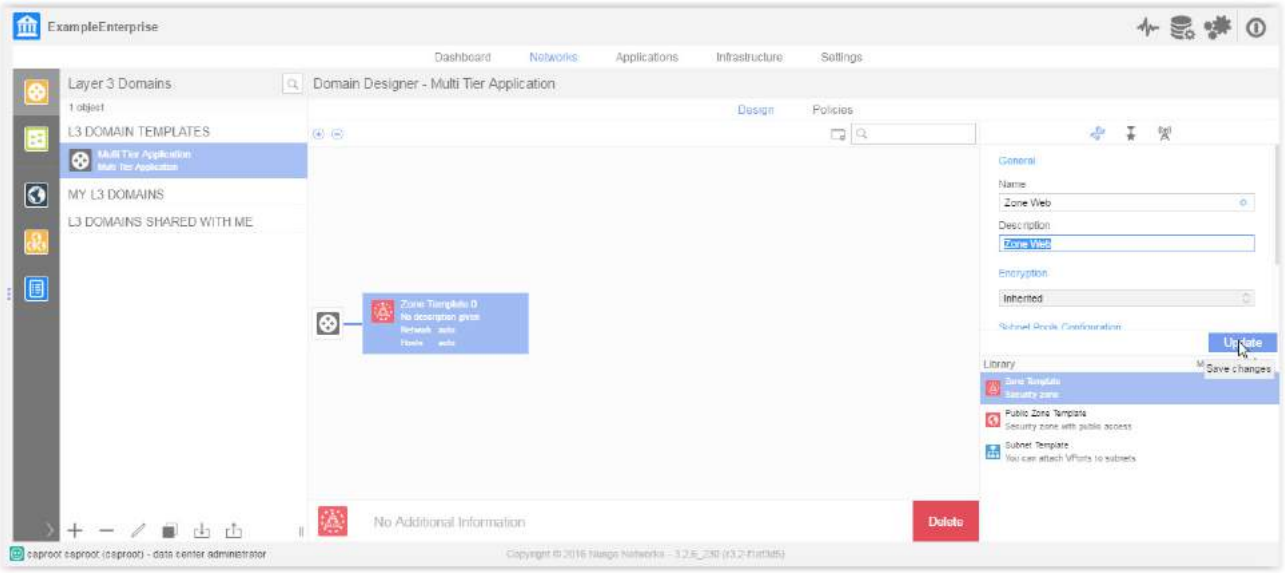

6. Repeat steps 4 and 5 to create additional zone objects. Note that one subnet will be used for web clients, one for application (web) services, and one for database services in the multitier application.

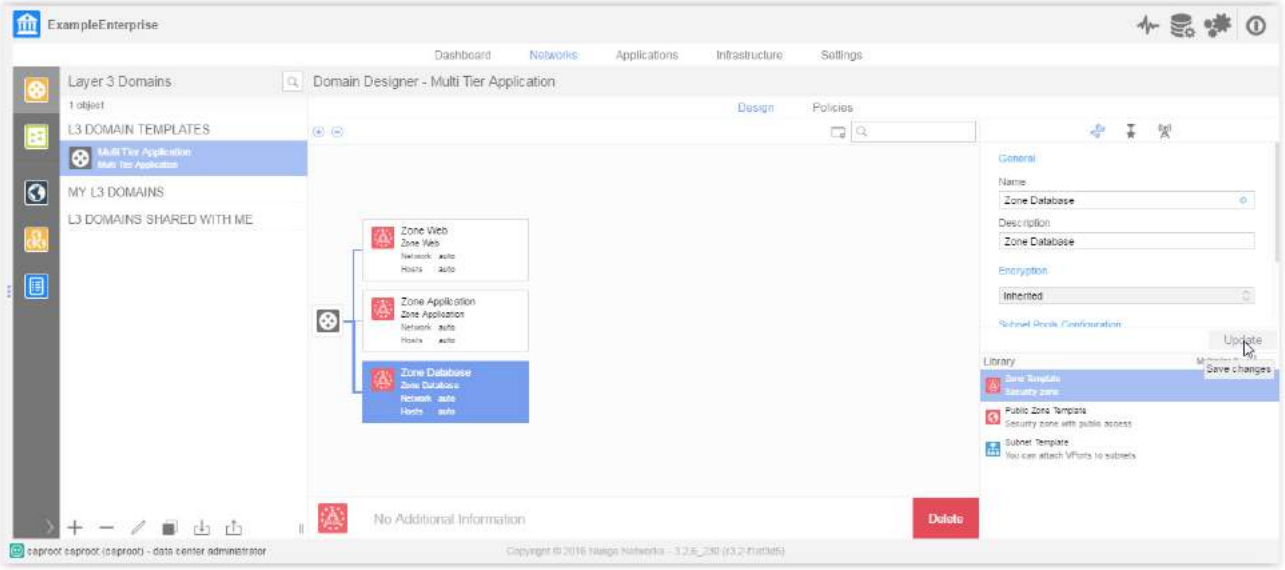

7. Add a subnet object to a zone by dragging the Subnet Template object onto the Zone object.

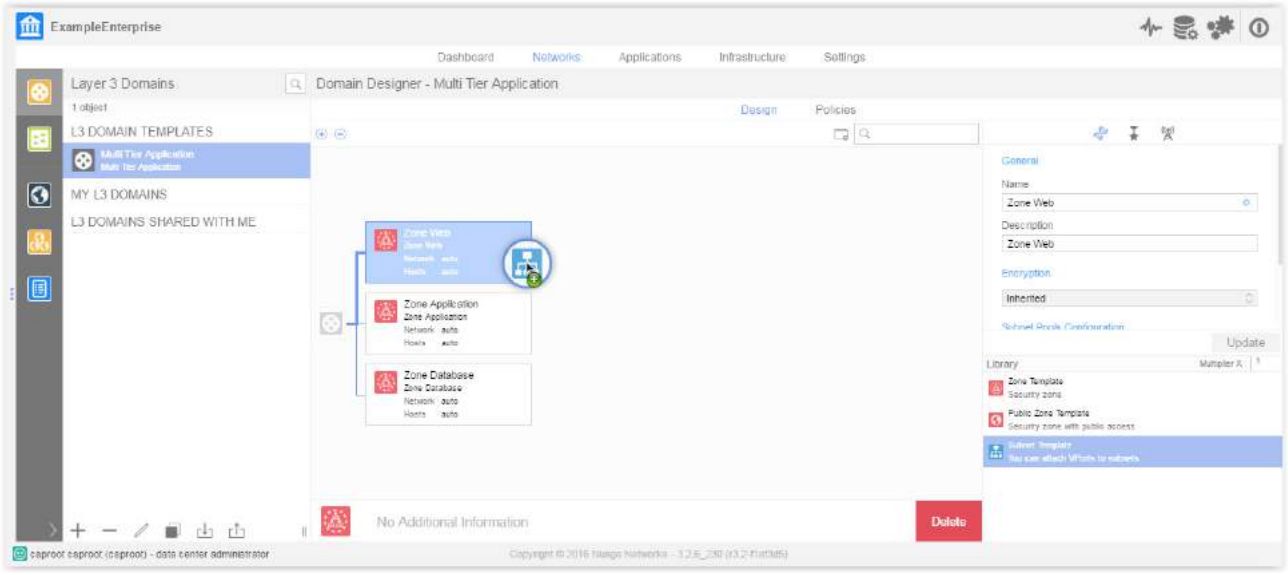

8. Modify the Name, Network, and Gateway fields to the desired values, then click Create.

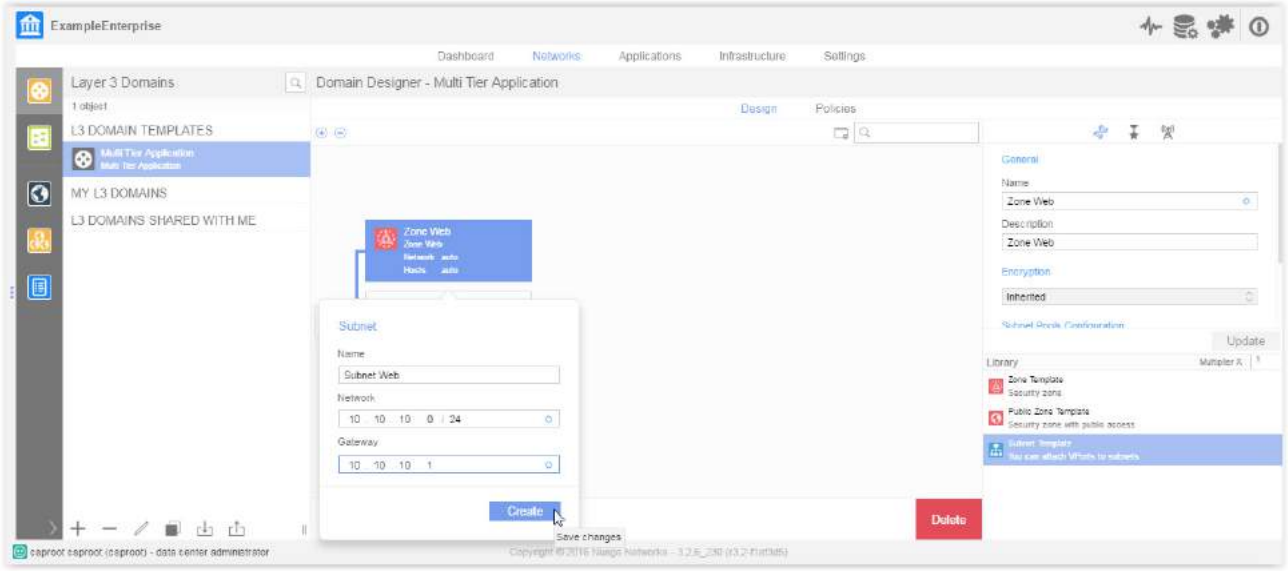

9. Modify the Description field to the desired value, then click Update.

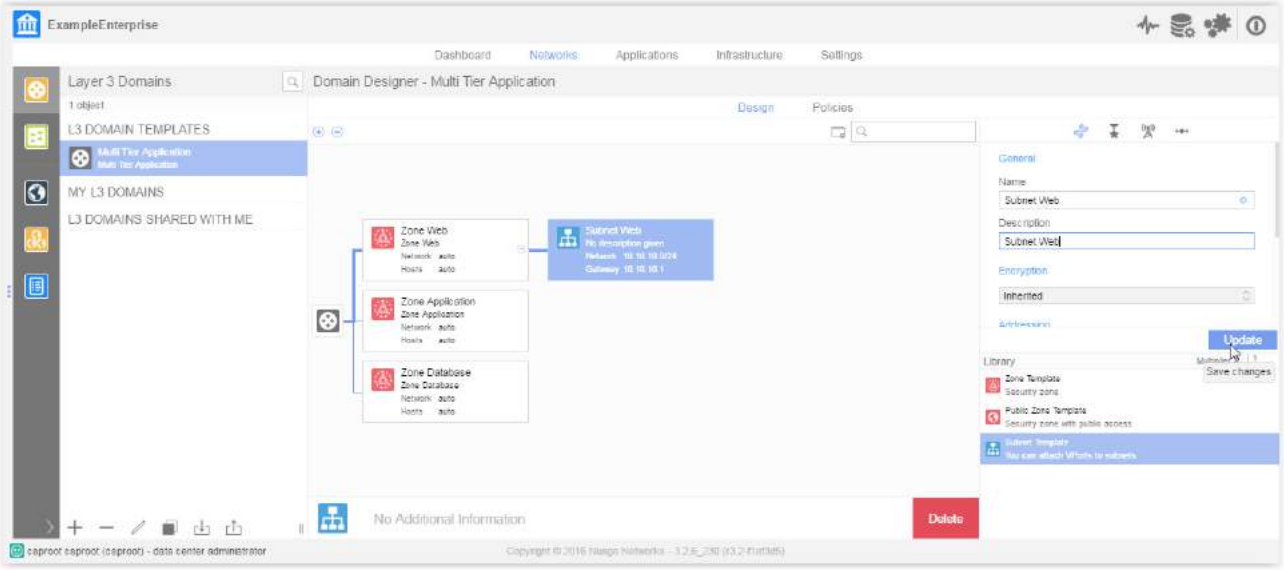

10. Repeat steps 7 through 9 to create additional subnet objects.

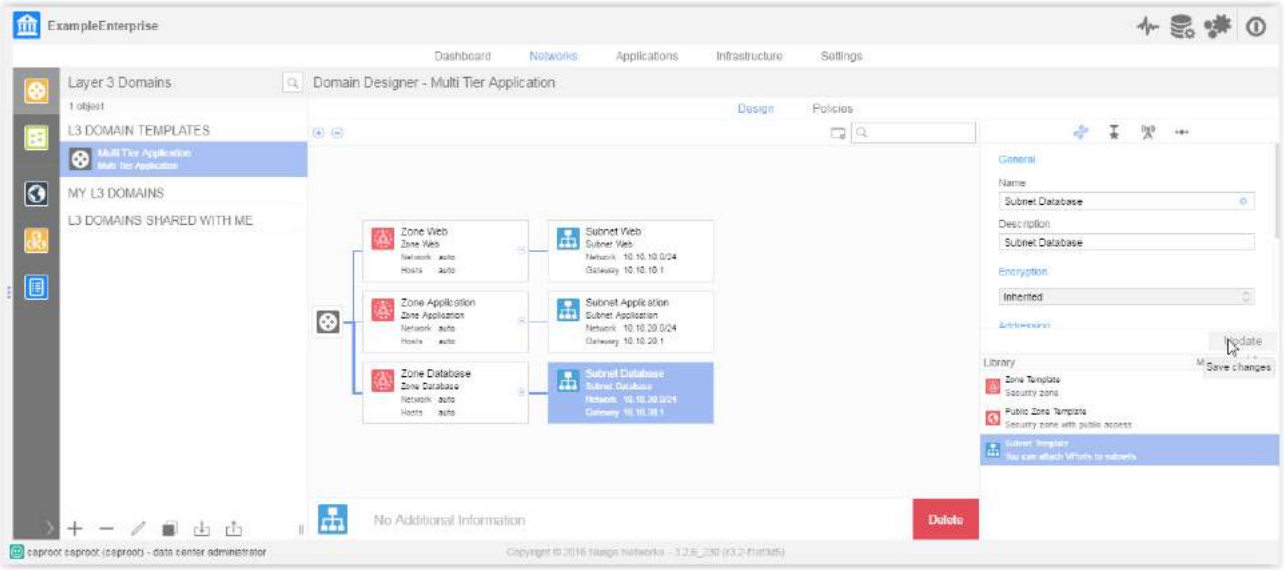

11. Configure DHCP IP Pool by selecting the subnet object, then clicking the Address Ranges option.

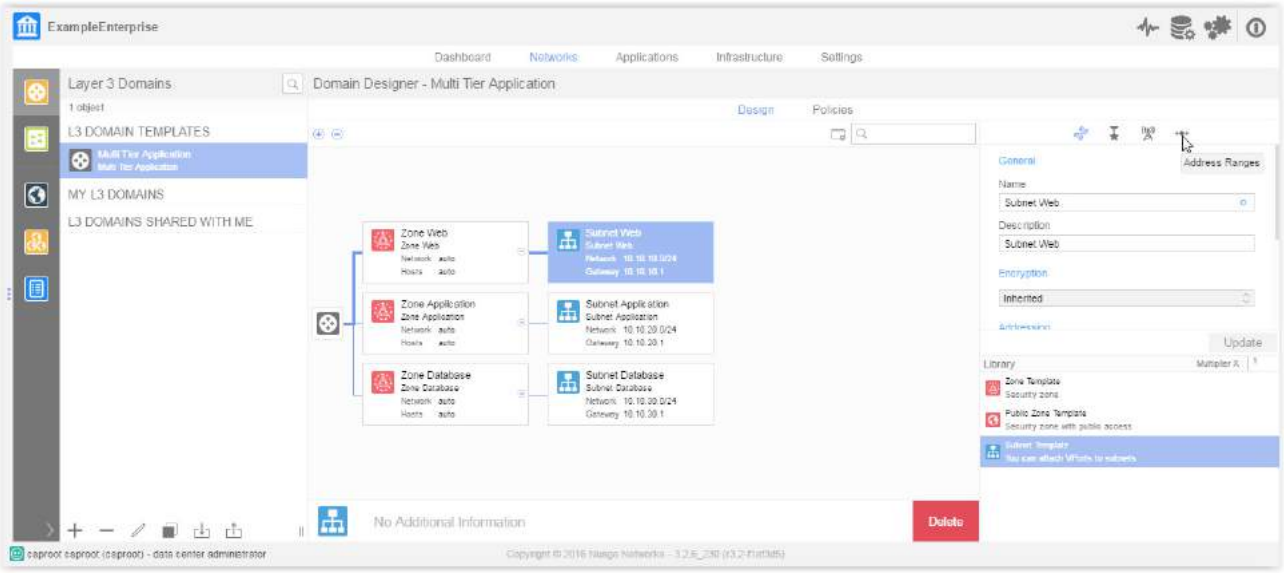

12. Click + to create a new address range.

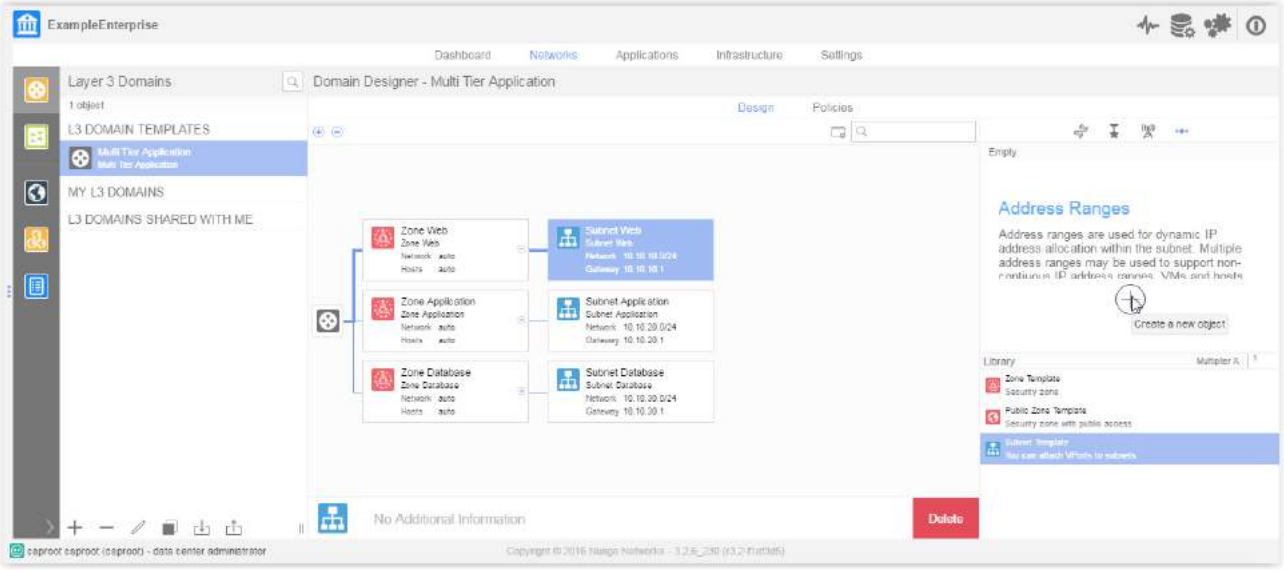

13. Configure the subnet address range information, then click Create.

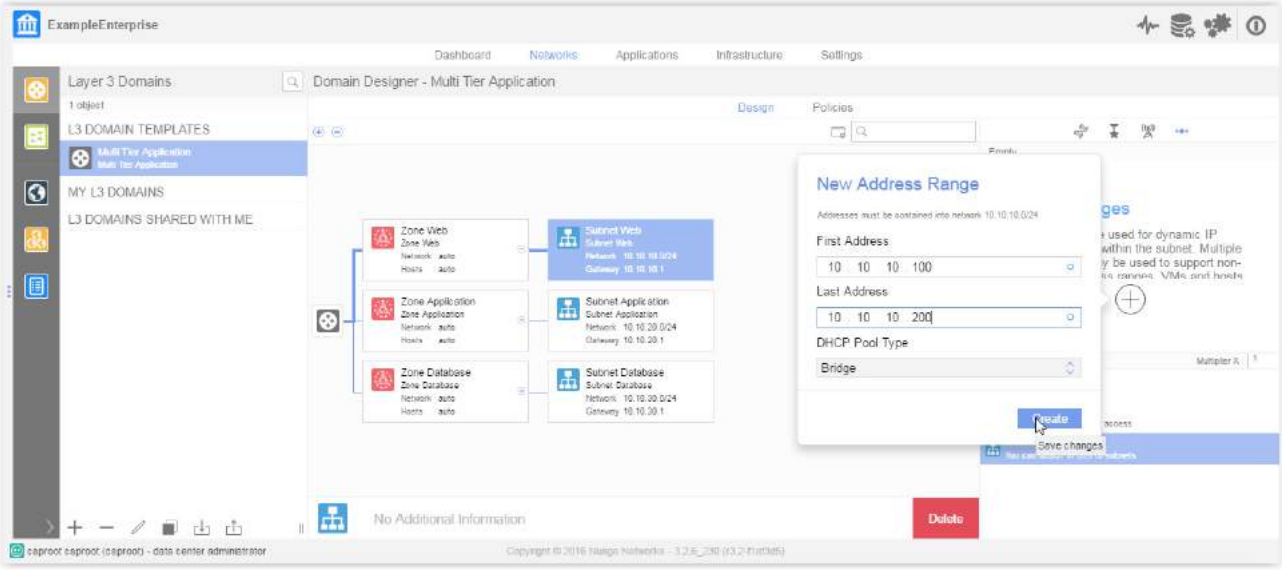

14. Confirm that the address range was created successfully.

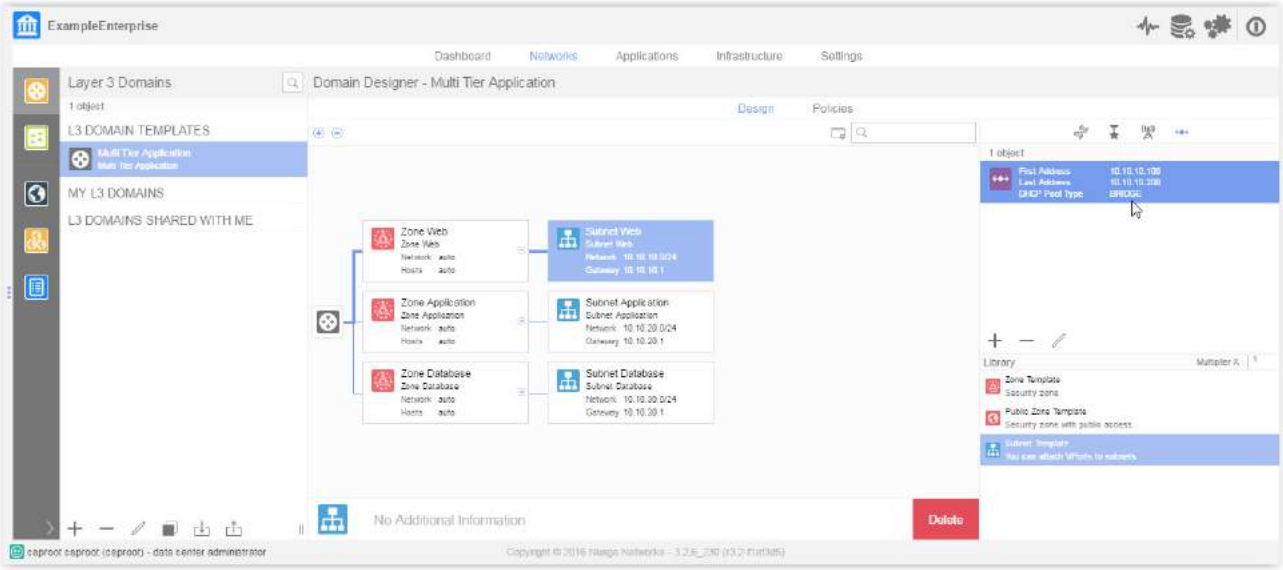

15. Repeat the procedure for all subnets.

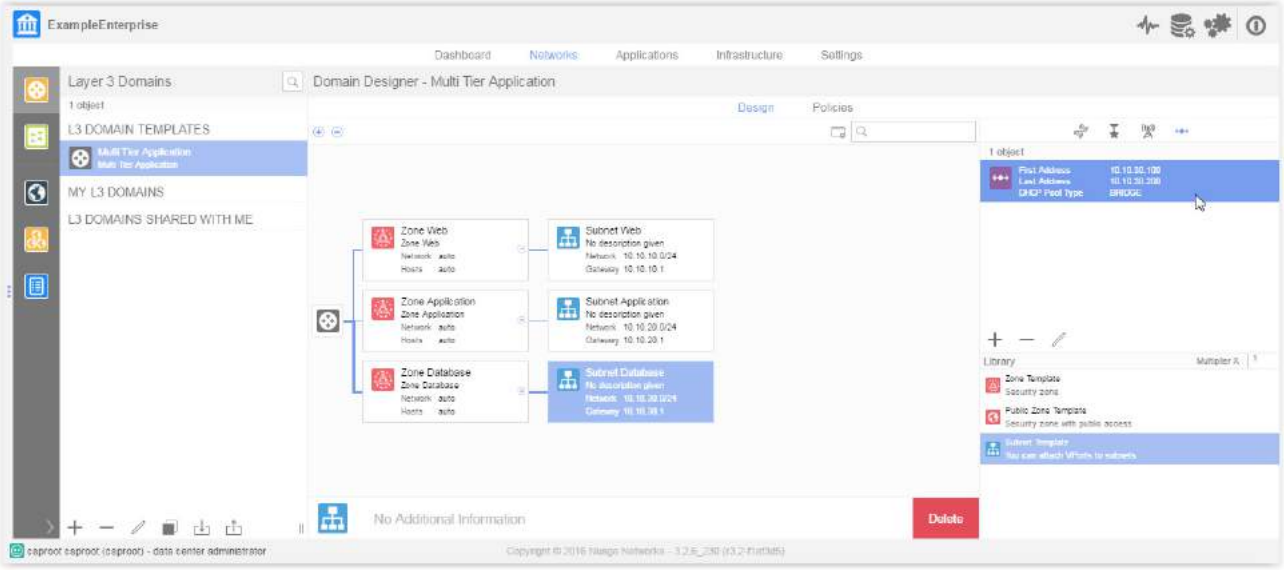

Instantiate a multitier template

1. Right-click on the template, and select Instantiate.

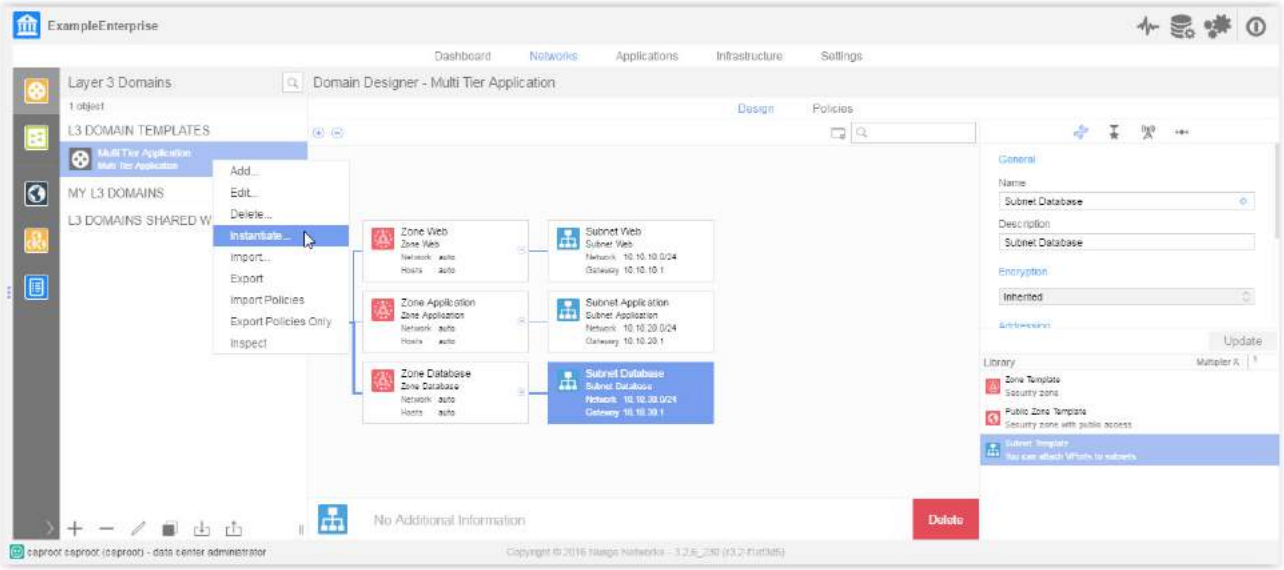

2. Modify the Name and Description fields to the desired values, then click Instantiate.

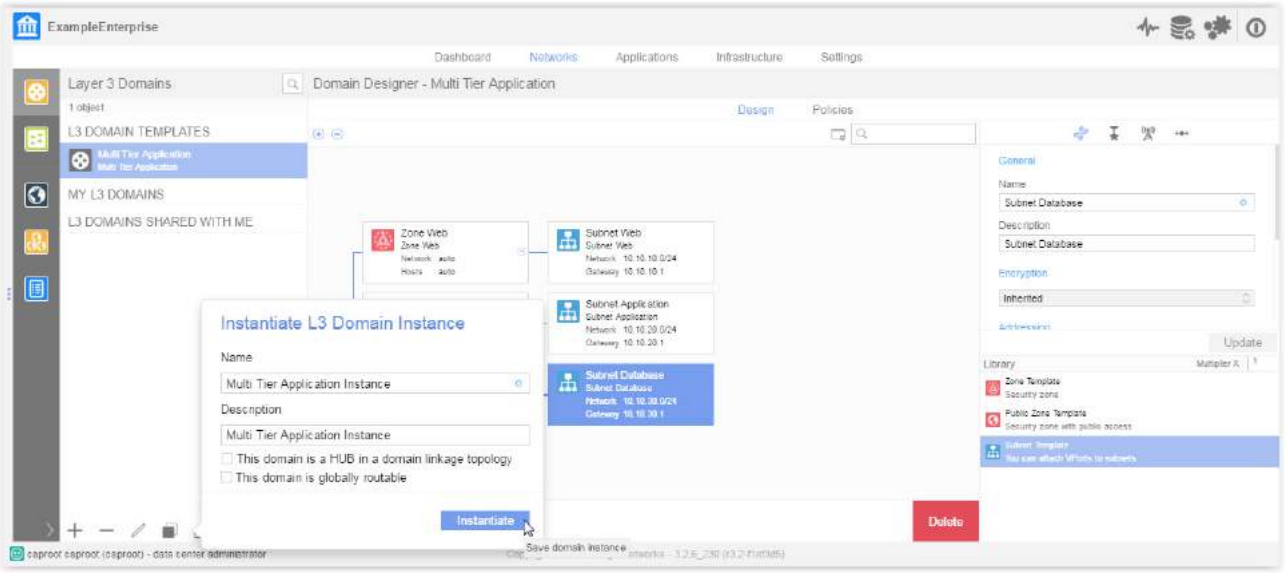

3. Confirm that the L3 Domain Template instance has been successfully created.

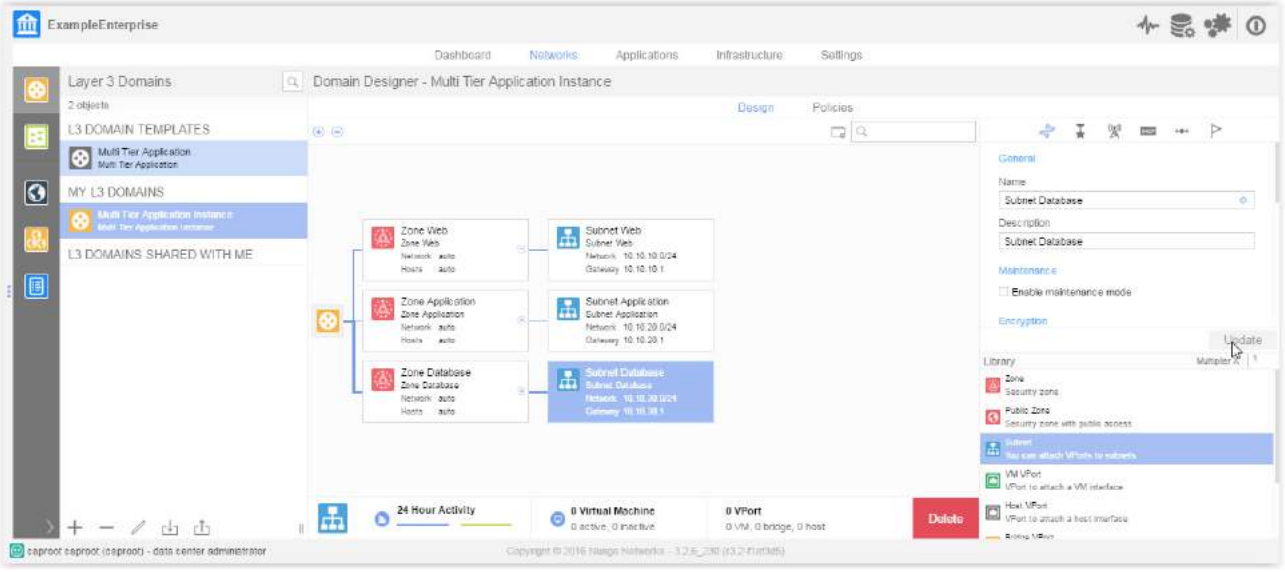

Create Web Client, Application Server, and Database Server VMs

1. Using a VM manager (e.g. KVM Virtual Machine Manager/VMWare vCenter) or a cloud management system (e.g. OpenStack), create the following VMs (procedure not shown):
   - A Web Client VM attached to the Subnet Web subnet object
   - An Application Server VM attached to the Subnet Application subnet object
   - A Database Server VM attached to the Subnet Database subnet object

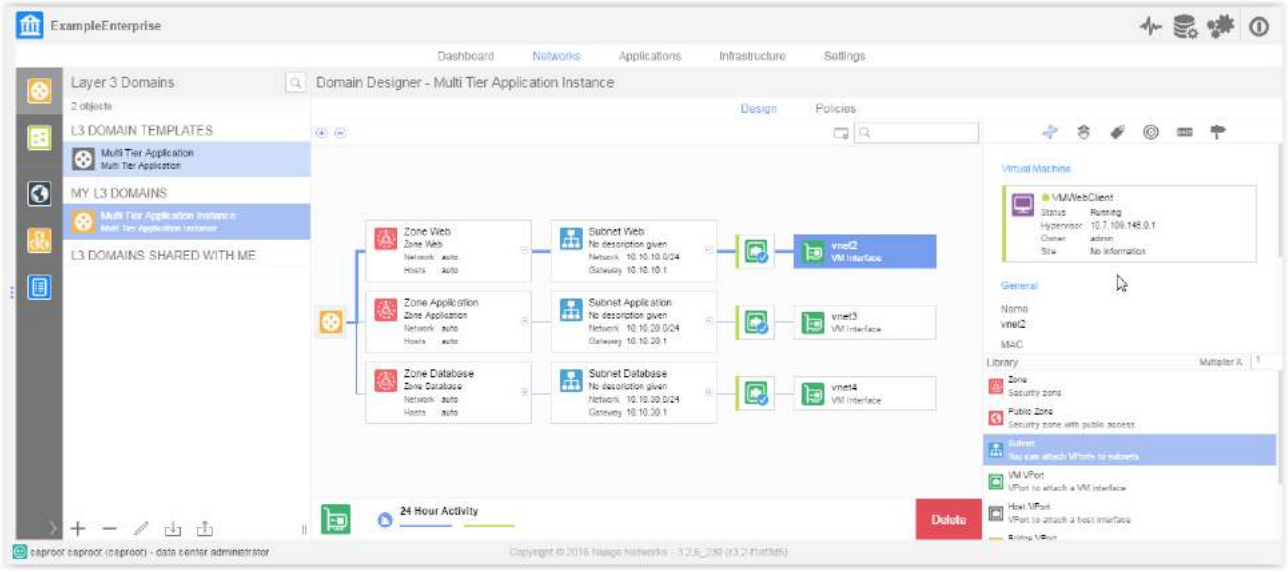

Create a firewall VM

Using a VM manager (e.g. KVM Virtual Machine Manager/VMWare vCenter) or a cloud management system (e.g. OpenStack), create a firewall VM with one interface in the Subnet Web and a second interface in the Subnet Application (procedure not shown).

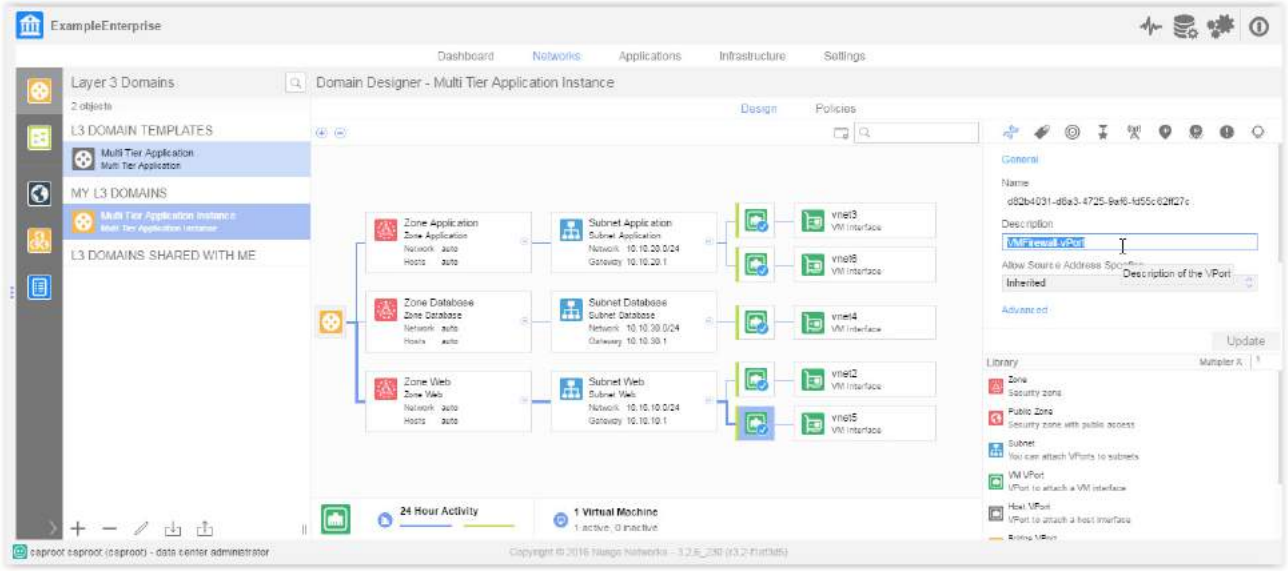

Create attachment for firewalls

1. Create a Redirection-Target to group external app firewall interfaces.

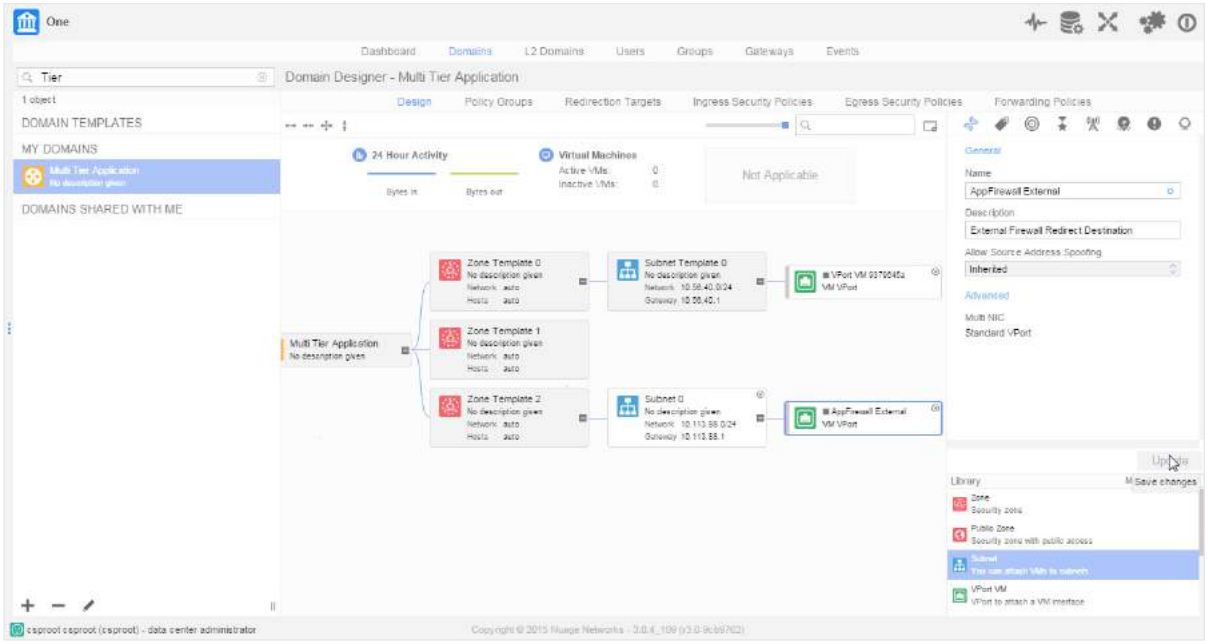

2. Create a Redirection-Target to group internal app firewall interfaces.

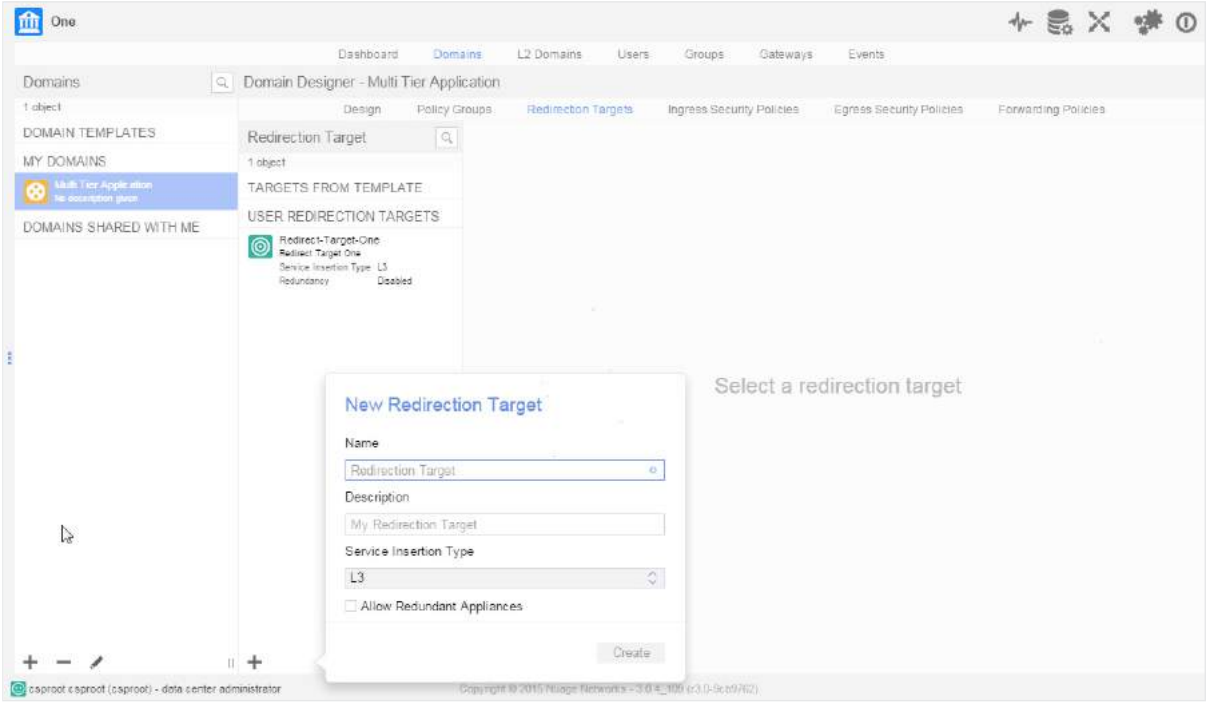

Add firewall VIP

1. Add a virtual IP address for the redundant firewall.

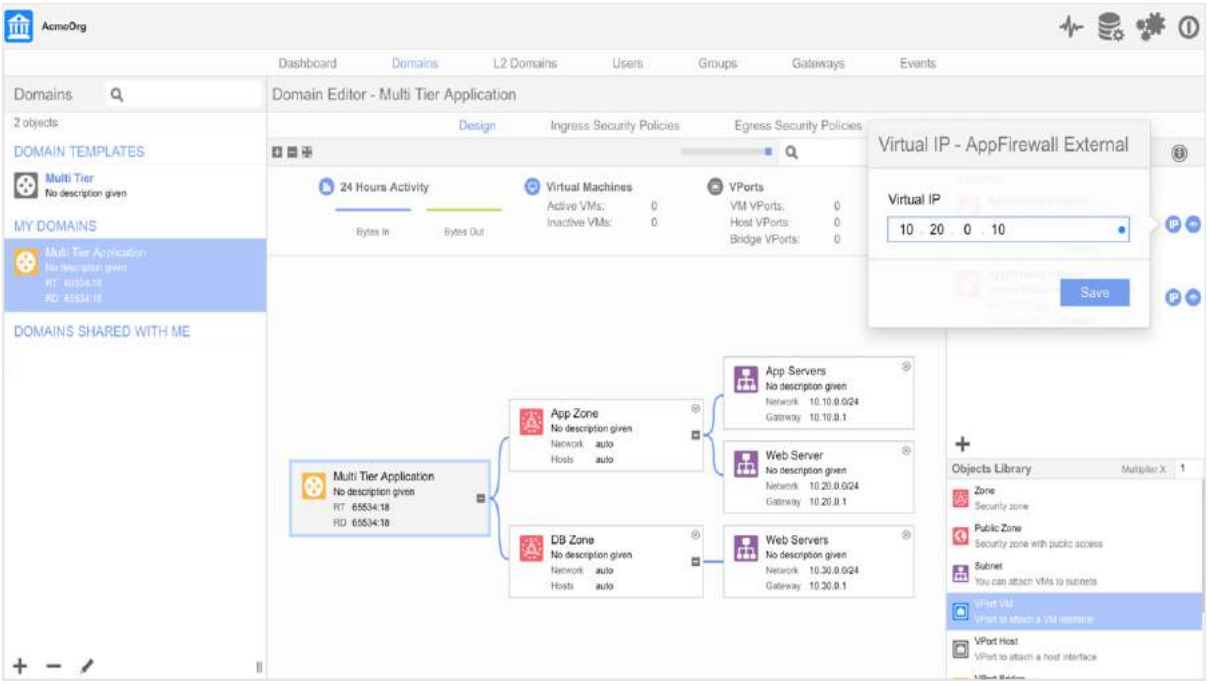

Add firewall internal VIP

1. Add a virtual IP address for the internal firewall.

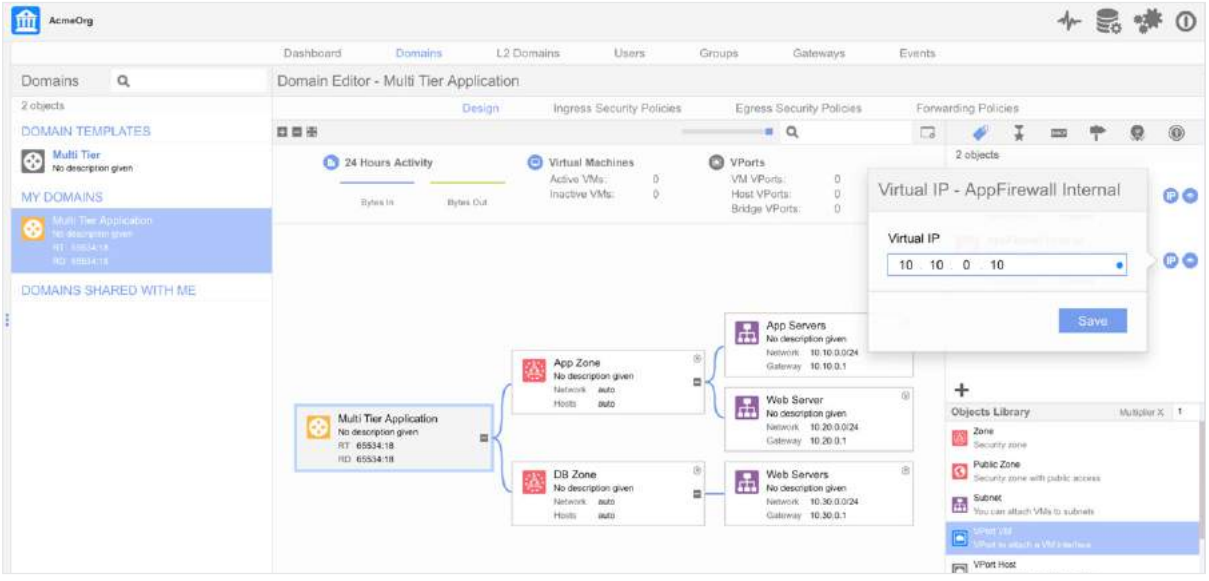

Add firewall external interfaces

1. Create the first VPort for a firewall external interface.

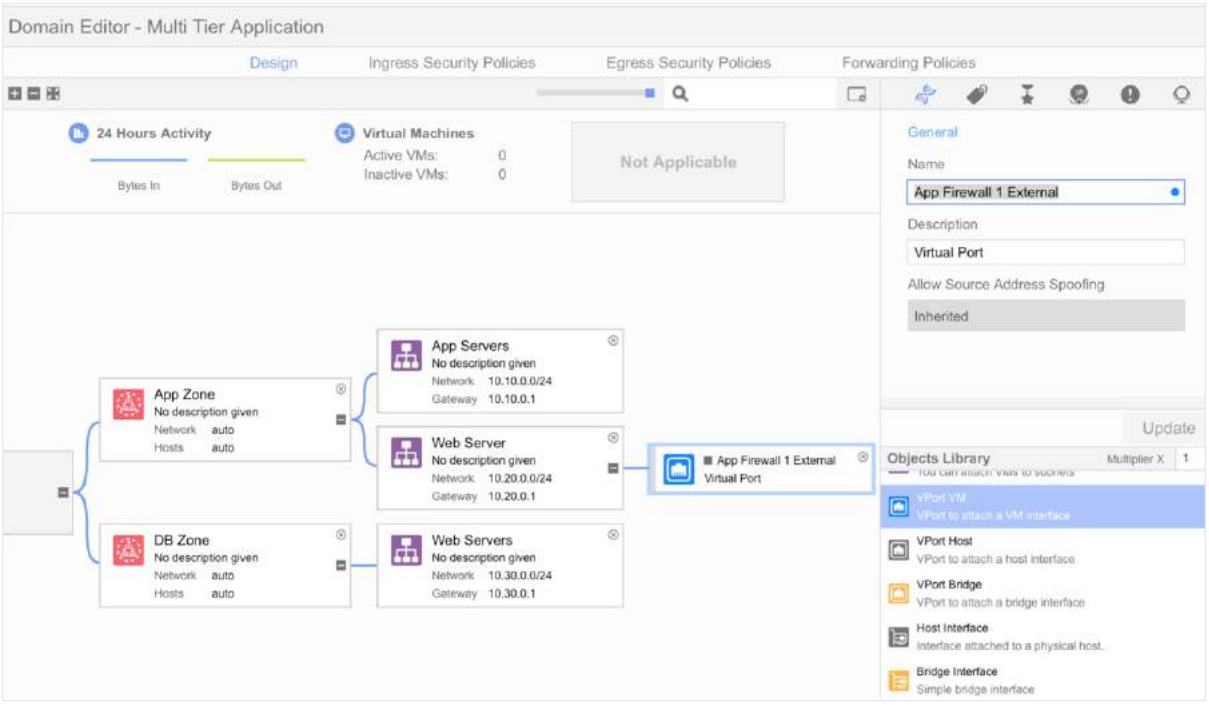

2. Create VPorts for redundant firewall external interfaces.

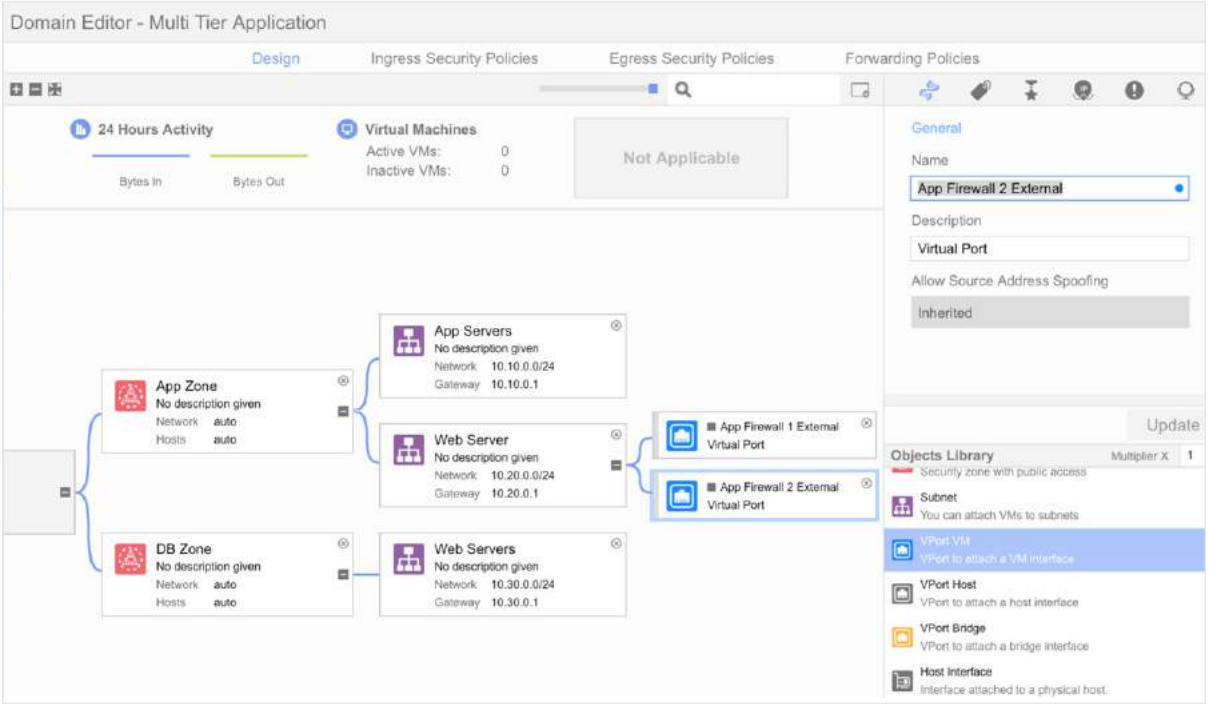

Add firewall external interface redundancy

1. Add a Redirection-Target for redundant firewall external interfaces.

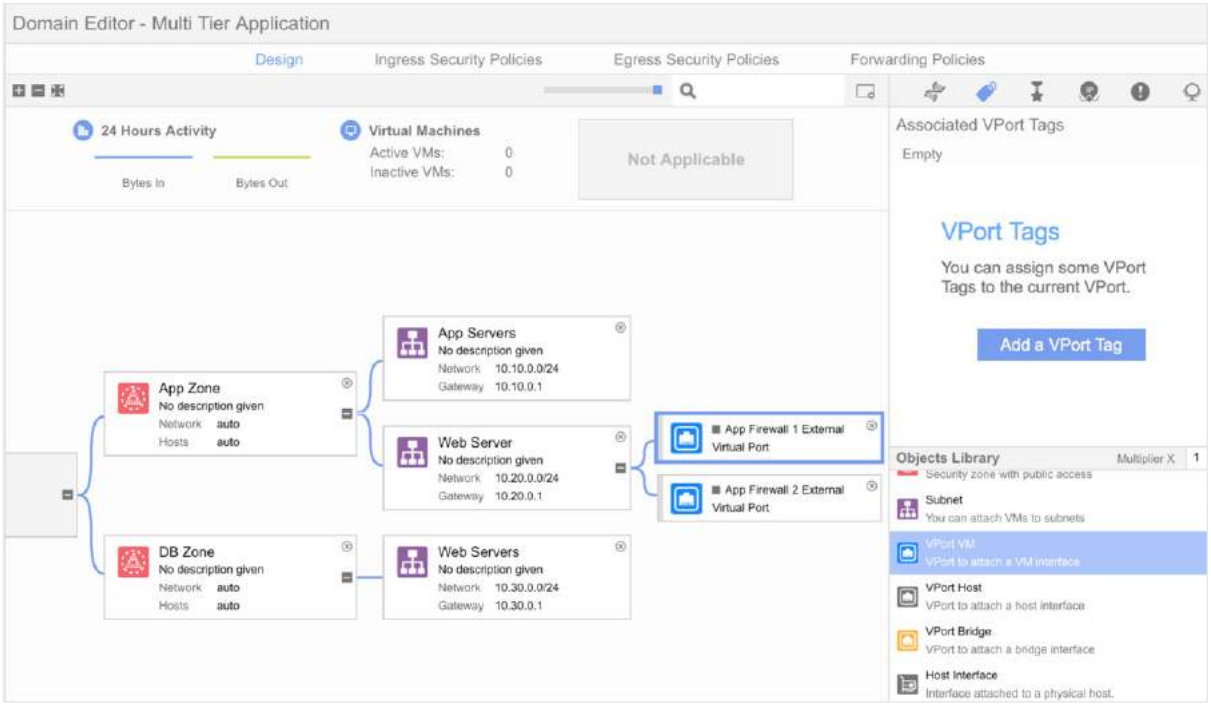

Add firewall internal interfaces and redundancy

1. Add VPorts and Redirection-Targets for firewall internal interfaces.

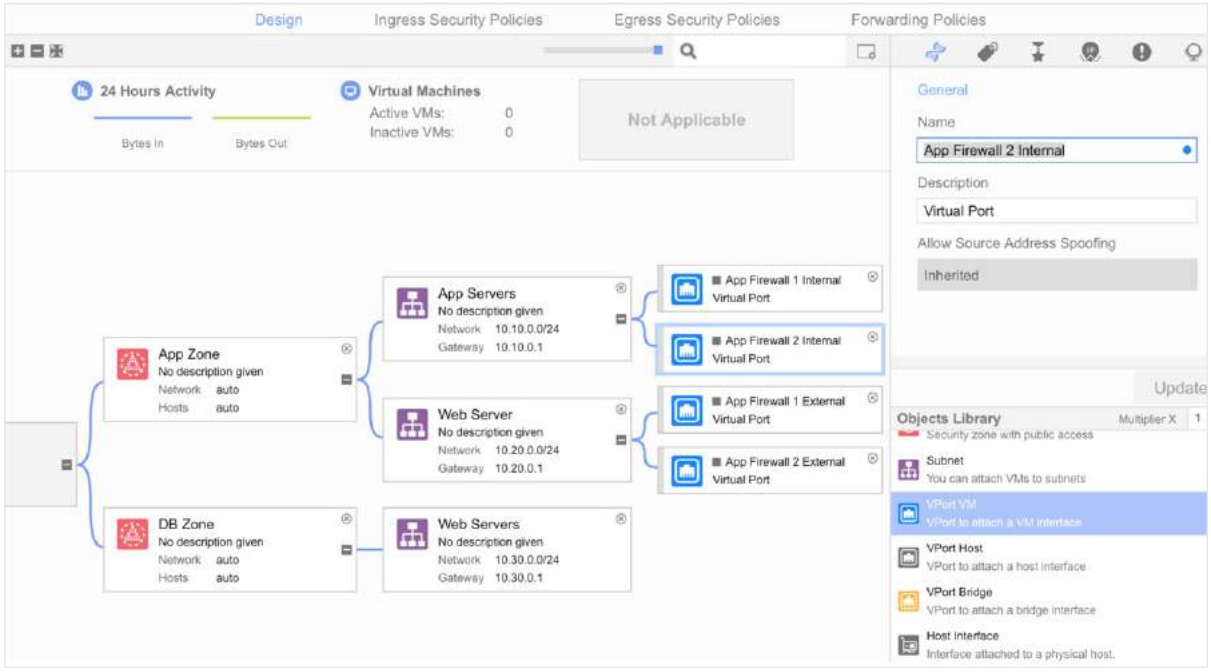

Create an ingress security policy

The ingress security policy is used to forward interesting traffic to the firewall and to drop all other traffic.

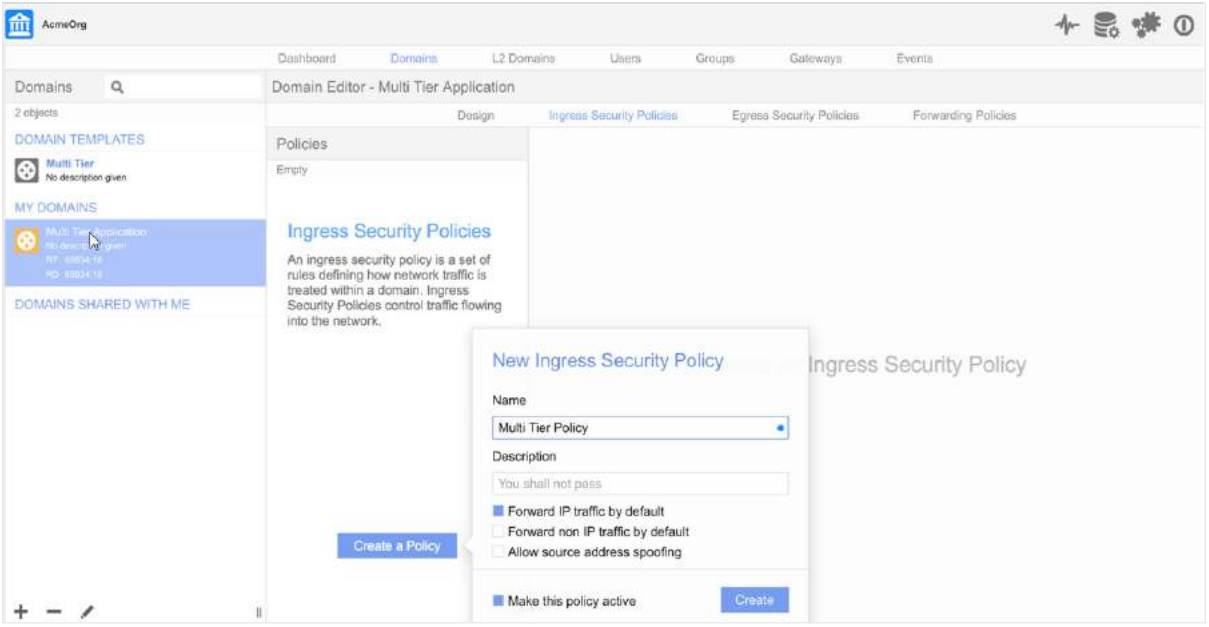

Forward web to app traffic on port 8888 to firewall for inspection

1. Redirect port 8888 traffic to the firewall.

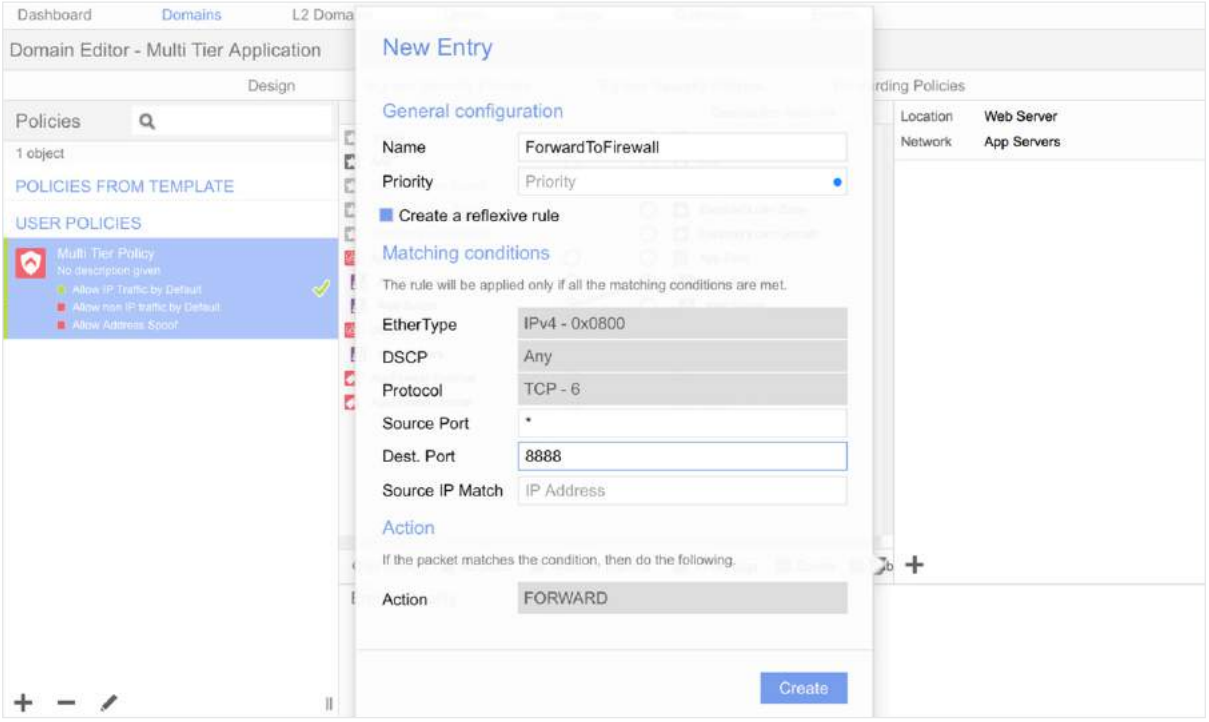

Drop other web to app traffic

1. Drop all other web to app traffic (i.e. not port 8888).

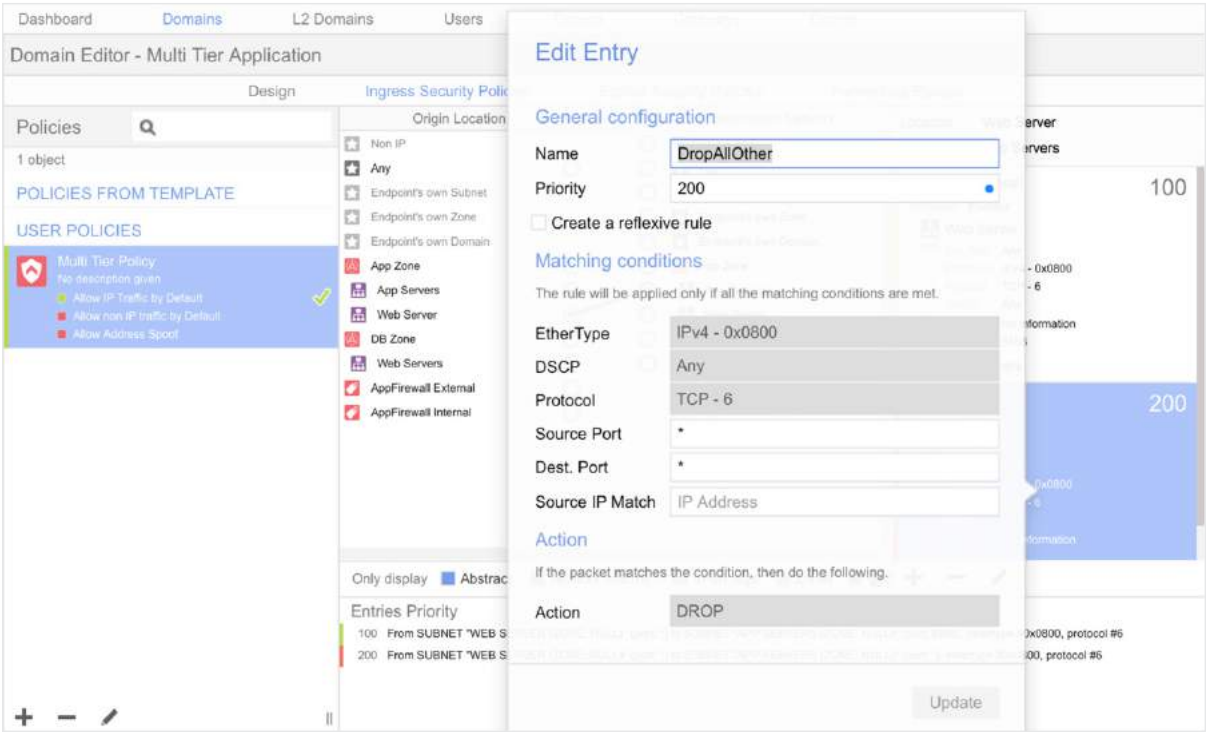

Summary of ingress security policy

This figure summarizes the ingress security policy created in the steps above.

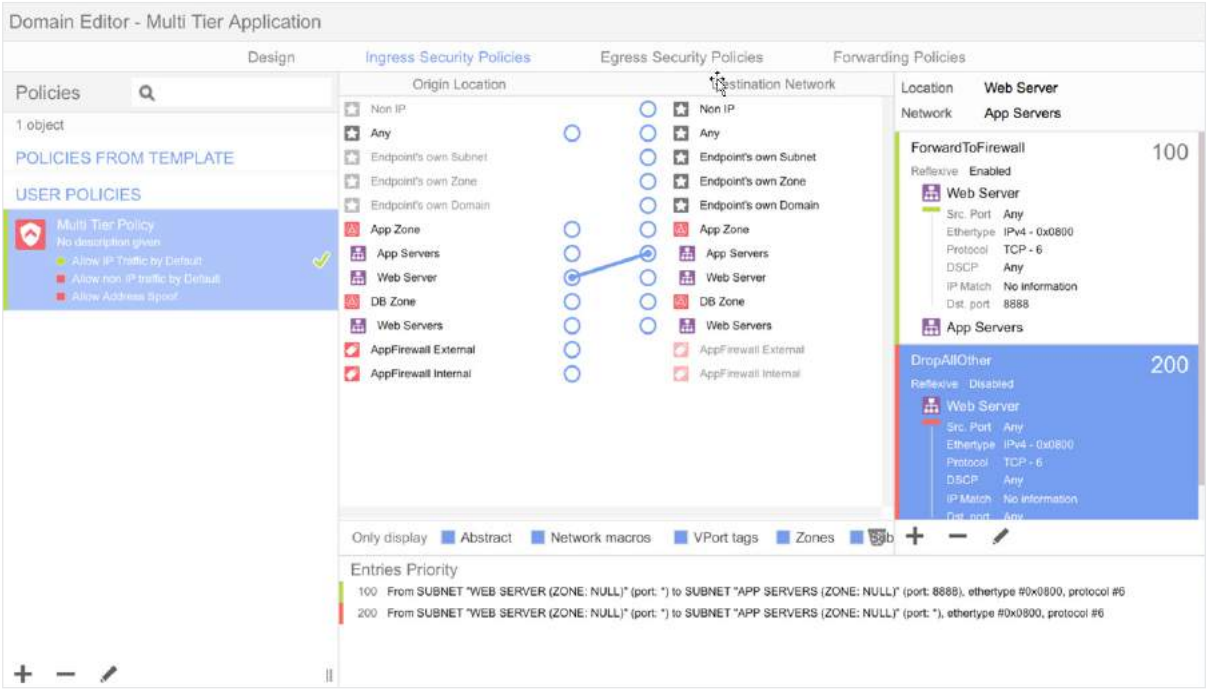

Create appropriate forwarding policies

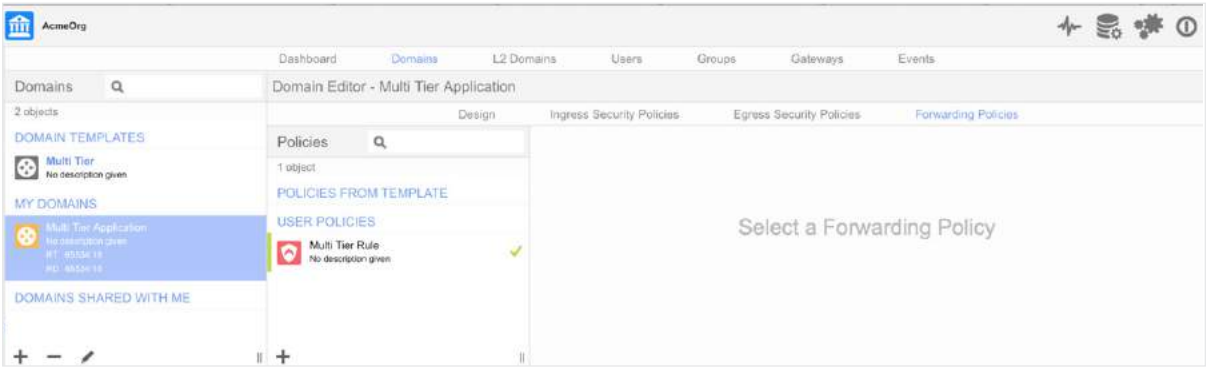

Redirect web to app traffic to firewall

1. Create a forwarding policy for traffic originating from the web server and destined for the app servers on port 8888. The policy required in this case is for the above traffic to be redirected to the firewall.

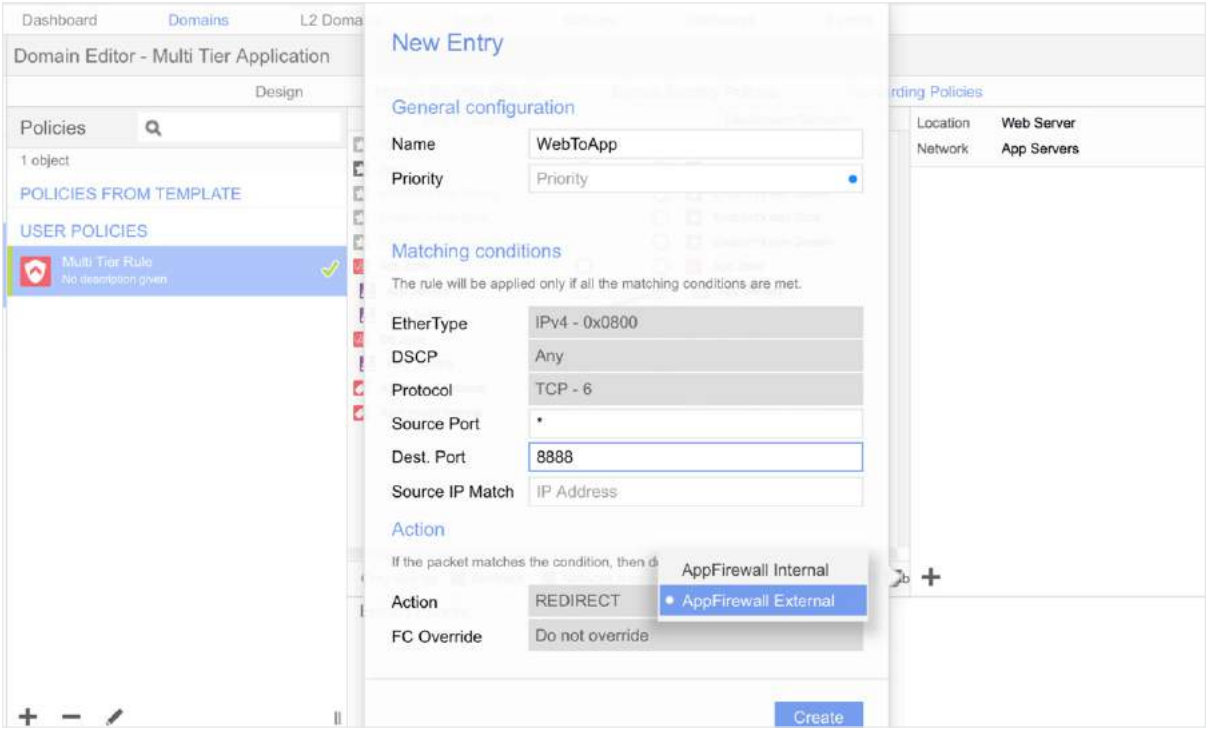

Summary of forwarding policy

This figure summarizes the forwarding policy created in the steps above.

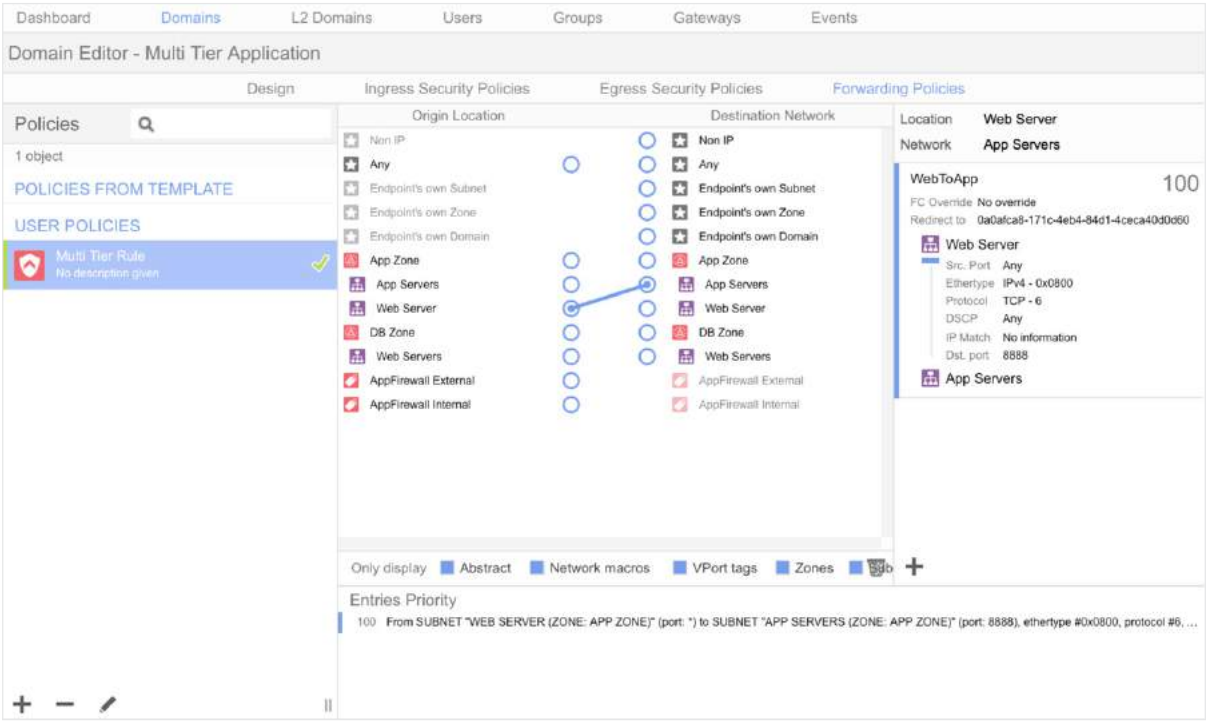

Domain ready for VM instantiation

At this point, VMs can be instantiated for the firewalls and will handle traffic between the subnets. If finer granularity than a subnet is needed, additional Redirection-Targets can be created and applied to individual VMs.

Any VM created on the web server subnet can only communicate with VMs on the app server subnet via the firewall. Routing elsewhere within the domain occurs normally.

Acronyms

ACL	Access Control List
API	Application Programming Interface
ARP	Address Resolution Protocol
BGP	Border Gateway Protocol
DC	Datacenter
DHCP	Dynamic Host Configuration Protocol
DPI	Deep Packet Inspection
EoR	End of Row [switch]
EP	Endpoint
ESS	[Alcatel-Lucent 7450] Ethernet Service Switch
FIB	Forwarding Information Base
FWaaS	Firewall-as-a-Service
GRE	Generic Routing Encapsulation
IPS/IDS	Intrusion Prevention System/Intrusion Detection System
IRB	Integrated Routing and Bridging
LBaas	Load-Balancer-as-a-Service
MP-BGP	Multiprotocol-Border Gateway Protocol
MPLS	Multiprotocol Label Switching
NAT	Network Address Translation
NIC	Network Interface Card
NVO	Network Virtualization Overlay
OVS	Open vSwitch
PBR	Policy-Based Routing
PNF	Physical Network Function
RD	Route Distinguisher
RIB	Routing Information Base
RT	Redirection Target
SDN	Software-Defined Networking
SR	[Alcatel-Lucent 7750] Service Router
SROS	[Alcatel-Lucent] Service Router OS
TCA	Threshold Crossing Alert
ToR	Top of Rack [switch]
UDP	User Datagram Protocol
VM	Virtual Machine
VNF	Virtual Network Function
VNID	Virtual Network IDentifier
VNS	Virtualized Network Services
VPLS	Virtual Private LAN Service
VPRN	Virtual Private Routed Network
VRF	Virtual Routing and Forwarding
VRS	Virtual Routing and Switching
VSC	Virtualized Services Controller
VSD	Virtualized Services Directory
VSG	[Nuage Networks 7850] Virtualized Services Gateway
VSP	Virtualized Services Platform
VTEP	VXLAN Tunnel Endpoint
VXLAN	Virtual eXtensible Local Area Network
XMPP	eXtensible Messaging and Presence Protocol
XRS	[Alcatel-Lucent 7950] Extensible Routing System

nuagenetworks
From Nokia