Deploying Enhanced Secure Multi-Tenancy into Virtualized Data Centers

Last Updated: July 7, 2011

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Aeisha Duncan

Aeisha Duncan, Technical Marketing Engineer, Systems Architecture and Strategy, Cisco Systems

Aeisha Duncan, CCIE #13455, is a Technical Marketing Engineer for data center technologies in Cisco's Systems Architecture and Strategy group. Prior to joining the SASU team, Aeisha spent 4 years as a Customer Support Engineer in Cisco's Technical Assistance Center where she supported LAN switching, VPN and Firewall technologies. She earned a B.S. in Computer Science from the University of Maryland at Baltimore County and an M.S. in Computer Networking from North Carolina State University.

Alex Nadimi

Alex Nadimi, Solutions Architect, Systems Architecture and Strategy, Cisco Systems

Alex has been with Cisco for the past 15 years and is currently working as a Solutions Architect in Cisco Systems Architecture and Strategy group. Prior to this role, he worked as a Technical Marketing Engineer in the Cisco Central Marketing Organization. He has developed solutions and technical guidance on various technologies such as security, VPN networks, WAN transport technologies, data center solutions, and virtualization. Prior to Cisco, he has worked at Hughes LAN Systems and Northern Telecom. He holds a masters of science in electrical engineering from Louisiana State University.

John George

John George, Reference Architect, Infrastructure and Cloud Engineering, NetApp

John George is a Reference Architect in NetApp's Infrastructure and Cloud Engineering team and is focused on developing, validating, and supporting cloud infrastructure solutions that include NetApp products. Before his current role, he supported and administered Nortel's worldwide training network and VPN infrastructure. John holds a master's degree in computer engineering from Clemson University.

Lindsey Street

Lindsey Street, Reference Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a systems architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Master's of Science in Information Security from East Carolina University.

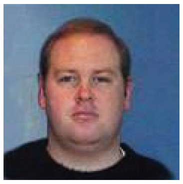

Mike Zimmerman

Mike Zimmerman, Reference Architect, Infrastructure and Cloud Engineering, NetApp

Mike Zimmerman is a reference architect in NetApp's Infrastructure and Cloud Engineering team. He focuses on the implementation, compatibility, and testing of various vendor technologies to develop innovative end-to-end cloud solutions for customers. Zimmerman started his career at NetApp as an architect and administrator of Kilo Client, NetApp's internal cloud infrastructure, where he gained extensive knowledge and experience building end-to-end shared architectures based upon server, network, and storage virtualization.

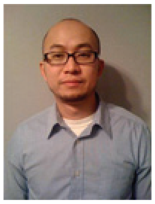

Wen Yul

Wen Yu, Senior Infrastructure Technologist, VMware

Wen Yu is a Sr. Infrastructure Technologist at VMware, with a focus on partner enablement and evangelism of virtualization solutions. Wen has been with VMware for six years during which time four years have been spent providing engineering level escalation support for customers. Wen specializes in virtualization products for continuous availability, backup recovery, disaster recovery, desktop, and vCloud. Wen Yu is VMware, Red Hat, and ITIL certified.

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit http://www.cisco.com/go/designzone.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at http://www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Deploying Enhanced Secure Multi-Tenancy into Virtualized Data Centers

Deploying Enhanced Secure Multi-Tenancy into Virtualized Data Centers

Introduction

Enhanced Secure Multi-Tenancy Overview

The biggest obstacle to adoption of IT as a service (ITaaS) has been lack of confidence that data and applications are securely isolated in a cloud-based infrastructure, where servers, networks, and storage are all shared resources. Cisco®, VMware®, and NetApp® have jointly designed a best-in-breed Enhanced Secure Multi-Tenancy (ESMT) architecture and have validated this design in a lab environment. For more information, see the Enhanced Secure Multi-Tenancy Design Guide (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldg_V2.html).

This document provides detailed implementation information and examples from a lab-validated reference design. It is structured to provide server, network, and storage architects and engineers with the implementation details to deploy and secure multi-tenant environments on four pillars:

- Secure separation
- Service assurance
- Availability
- Manageability

Note This deployment guide assumes an existing FlexPod™ for VMware Infrastructure is in place and has been properly configured.

FlexPod Overview

FlexPod for VMware is a validated "POD-like" configuration built with technologies and best practices from Cisco, NetApp, and VMware. FlexPod for VMware serves as a base infrastructure platform for many of the applications and solutions that customers deploy today. ESMT is one example of a solution

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

that can be layered on top of FlexPod for VMware. For more information on FlexPod for VMware and the other solutions that can be built on the infrastructure, see FlexPod for VMware Deployment Model (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html).

Note NetApp TR-3892, FlexPod for VMware Implementation Guide, is available only to Cisco, NetApp, and VMware sales representatives or qualified FlexPod partners.

Security Framework

Secure separation within this architecture is implemented at all layers and within most devices. This "defense in depth" approach follows the methodology outlined in Cisco SAFE: A Security Reference Architecture (http://www.cisco.com/en/US/solutions/collateral/ns170/ns896/ns954/white_paper_c11-527476.html).

Deployment Procedures

This section outlines the deployment procedures for the various components and devices of the ESMT architecture. The lab implementation used to validate the ESMT deployment included four tenants:

- Tenant 1—SharePoint
- Tenant 2—Exchange
- Tenant 3—SQL
- Tenant 50—Infrastructure

These applications were used to create tenants, however this document does not describe application deployment procedures. The Infrastructure tenant includes commonly shared resources such as Microsoft Active Directory® (AD) and a centralized syslog server.

Deploying the Cisco Nexus 7000

Figure 1 Layer 2 Deployment

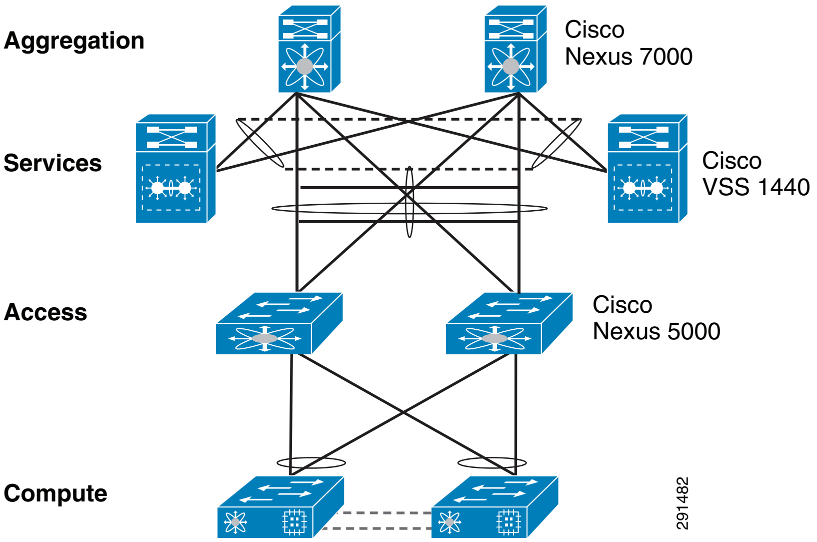

Layer 2 Deployment

The Cisco Nexus® 7000 acts as the aggregation layer in the ESMT deployment, so it needs trunked links toward the access layer and the services chassis (Figure 1). The two Cisco Nexus 7000s deployed should be configured as virtual port channel (vPC) peers, with an EtherChannel acting as the vPC peer link. The following configuration snippet is for one peer in the vPC domain:

```
vpc domain 10
  role priority 50
  peer-keepalive destination 10.1.30.102 source 10.1.30.101 vrf vpc-mgmt
  peer-gateway

interface port-channel1
  description *** Port Channel Between Aggregation - VPC Peer-Link; E1/17, E1/18 ***
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2-3967,4048-4093
  spanning-tree port type network
  mtu 9000
  vpc peer-link
```

Once vPC peering is established, a vPC should be configured going toward the access layer and the services chassis:

```
interface port-channel45
  description *** Port Channel Between Aggregation and Access Switches E1/25, E1/26 ***
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-111,113,115,119-124,129-130
  switchport trunk allowed vlan add 132,140,161-162,481,490-493,499-500
  switchport trunk allowed vlan add 581-586,588-592,599-600,900-902
  switchport trunk allowed vlan add 905
  mtu 9000
  service-policy type queuing input ingress-queuing
  service-policy type queuing output egress-policy
  vpc 45


interface port-channel78
  description *** Port Channel Between Agg and Services - E1/9, E/10 ***
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,51-53,57,90,99-102,111,116,125
  switchport trunk allowed vlan add 131-132,161,166,494,586-587,594
  switchport trunk allowed vlan add 597,902
  logging event port link-status
  logging event port trunk-status
  mtu 9000
  service-policy type queuing input ingress-queuing
  service-policy type queuing output egress-policy
  vpc 78
```

Layer 3 Deployment

Figure 2 Tenant 1 Topology

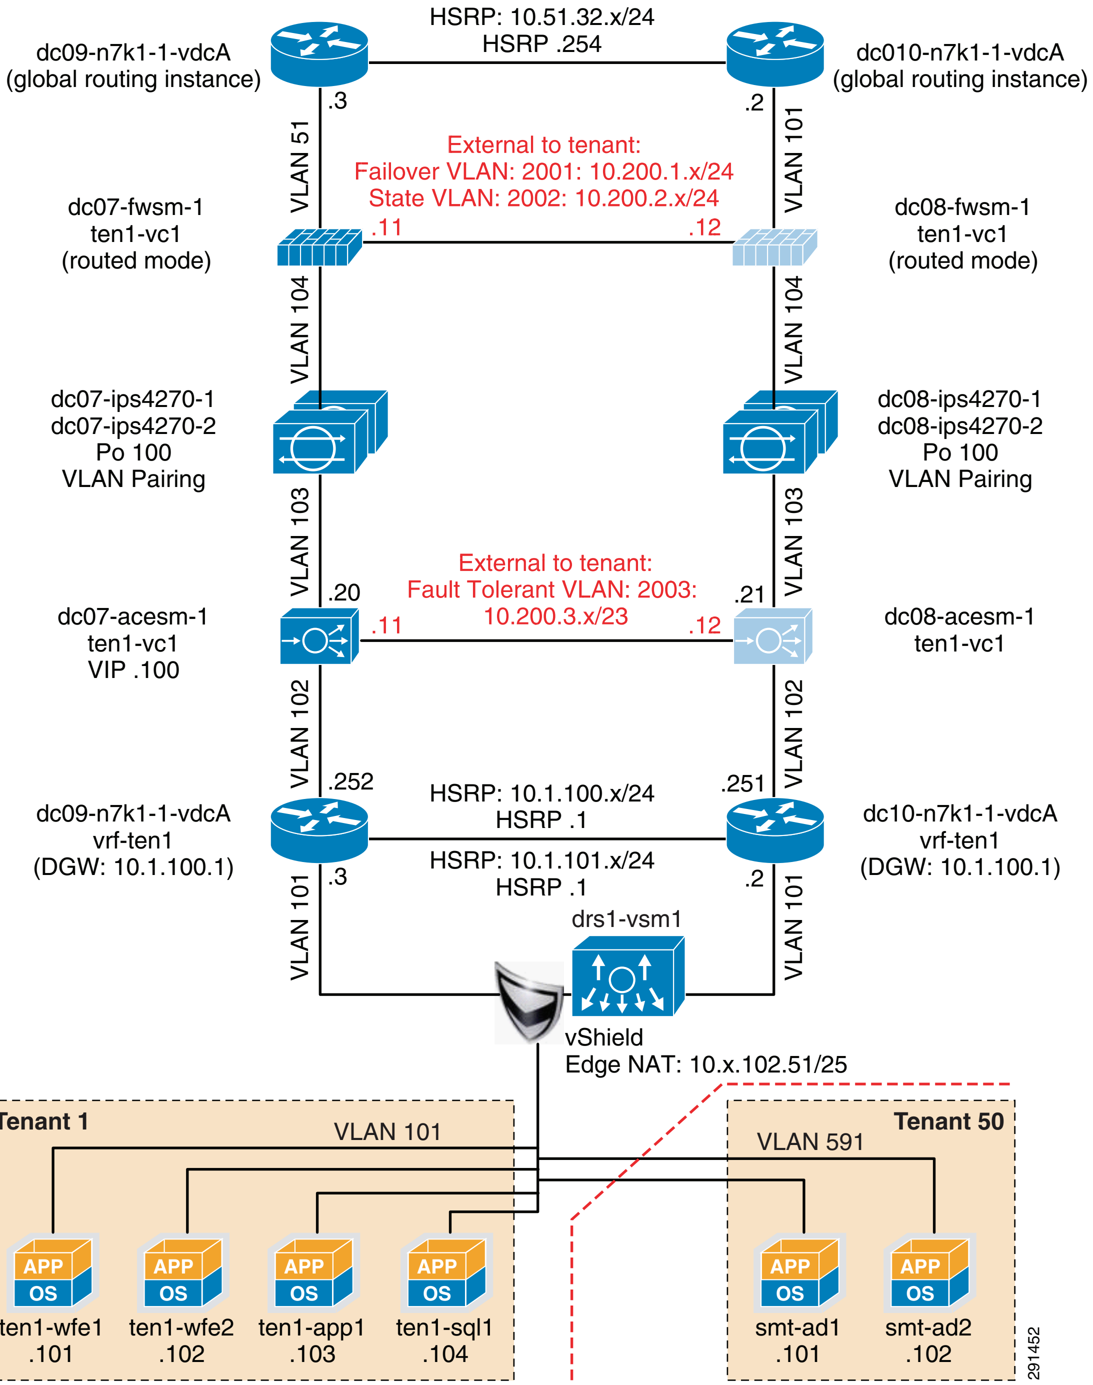

The Cisco Nexus 7000 acts as the Layer 3 gateway for all of the tenants. To ensure secure separation between the tenants, a Virtual Routing and Forwarding (VRF) instance is configured for each tenant. There may be multiple Layer 3 interfaces in a VRF instance to service different services and VLANs within a tenant.

```
vrf context vrf-ten1
  ip route 0.0.0.0/0 10.1.100.1 ---' Inter-tenant routing goes through FWSM in services
chassis


 interface Vlan101
  vrf member vrf-ten1
  no ip redirects
  ip address 10.1.101.3/24
  ip router ospf 1 area 0.0.0.0
  hsrp version 2
  hsrp 101
    authentication text c1sc0
    preempt delay minimum 180
    priority 110
    timers  1  3
    ip 10.1.101.1
  no shutdown
  mtu 9216
  description **SVI for ten1**

interface Vlan102
  vrf member vrf-ten1
  no ip redirects
  ip address 10.1.100.252/24
  hsrp version 2
  hsrp 102
    authentication text c1sc0
    preempt delay minimum 180
    timers  1  3
    ip 10.1.100.254
  no shutdown
  mtu 9216
  description ** FireWall Inside ***

interface Vlan106
  vrf member vrf-ten1
  no ip redirects
  ip address 10.1.102.3/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp version 2
  hsrp 106
    authentication text c1sc0
    preempt delay minimum 180
    timers  1  3
    ip 10.1.102.1
  no shutdown
  description ** SVI for Tenant 1 Storage Services **

interface Vlan107
  vrf member vrf-ten1
  no ip redirects
  ip address 10.1.107.3/25
  hsrp version 2
  hsrp 107
    authentication text c1sc0
    preempt delay minimum 180
    priority 110
    timers  1  3
    ip 10.1.107.1
  no shutdown
  description ** SVI for ten1 mgmt **
```

Figure 3 Tenant Model External Traffic Patterns

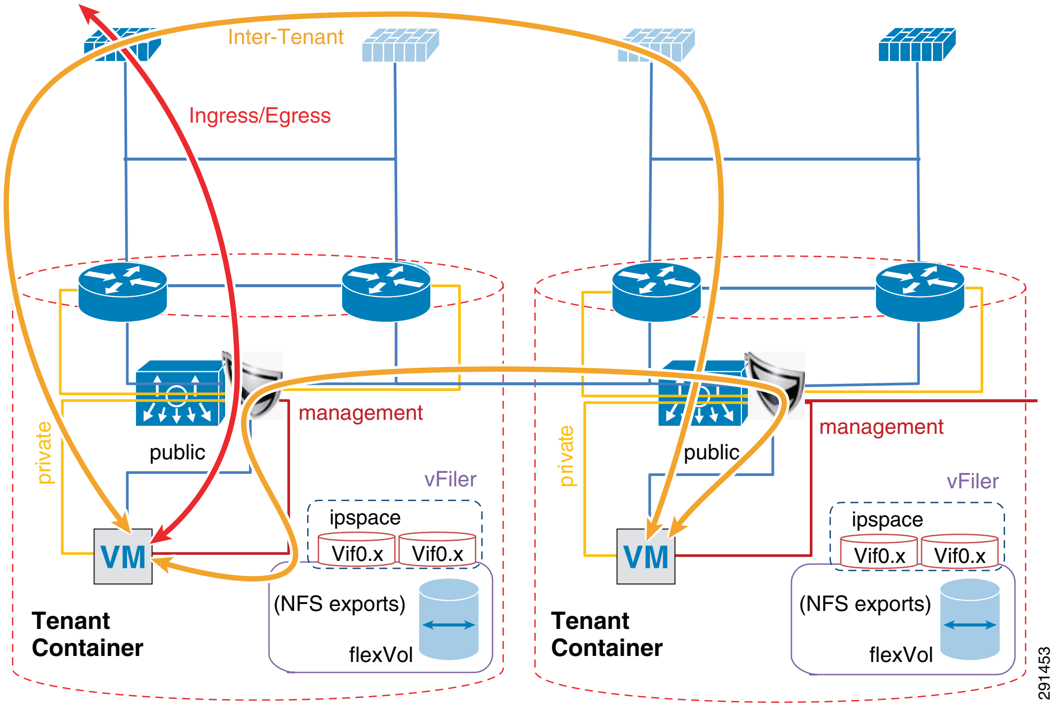

Inter-tenant routes are not propagated to the tenant VRF instances. Inter-tenant routing is handled by the Cisco Firewall Switching Module (FWSM) deployed in the services chassis. Packets destined to another VRF instance or tenant are routed to the FWSM which, in turn, routes them to the global routing table. Static routes to each VRF instance are added to the global routing table. The routes point to the outside interface of the virtual context on the FWSM that belongs to that particular tenant.

```
ip route 10.1.101.0/25 10.51.32.1   ---' 10.51.32.1 is the outside interface of the
ten1-vc on the FWSM
```

Figure 4 Core-Aggregation Connections

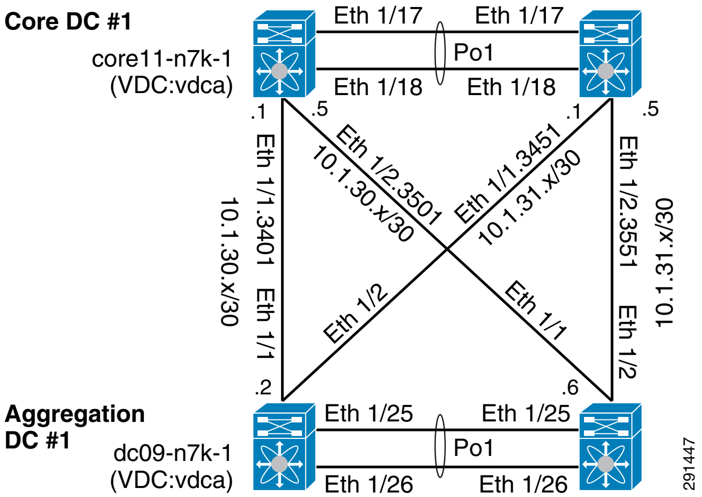

To reach clients that are in the VRF instance but not in the local subnets, routing to the core is needed. A trunk with Layer 3 point-to-point links is configured on the links going to both cores. Each point-to-point link belongs to a particular VRF instance/tenant and is used as the gateway to these clients.

```
interface Ethernet1/1.3401
  description *** Connection to CORE11-N7K-1 vrf-ten1 ***
  encapsulation dot1q 3401
  vrf member vrf-ten1
  ip flow monitor NETFLOW input
  ip flow monitor NETFLOW output
  ip address 10.1.30.2/30
  ip ospf authentication
  ip ospf authentication key-chain RoutingAuth
  ip router ospf 1 area 0.0.0.0
  no shutdown
```

For most cases, static routing is sufficient for intra- and inter-tenant traffic. For tenants requiring a more robust routing protocol, OSPF is used.

```
router ospf 1
  vrf vrf-ten1
    area 0.0.0.0 range 10.1.101.0/25
```

Hardening the Cisco Nexus 7000

In the following example, the Cisco Nexus 7000 can be hardened to allow access only through SSH, FTP, and ICM in the inbound access list. The outbound access list allows access only to the OOB management network and the inside network. More information on the Cisco Nexus 7000 security capabilities and general hardening guidance can be found at: http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_2/nx-os/security/configuration/guide/Cisco_Nexus_7000_NX-OS_Security_Configuration_Guide__Release_4.2.pdf.

```
ip access-list OBB-inbound
  3 permit tcp 172.26.162.216/32 172.26.162.68/32
  9 permit ip 172.26.162.214/32 172.26.162.68/32
  10 permit icmp 172.26.162.0/16 172.26.162.68/32 ttl-exceeded
  20 permit icmp 172.26.162.0/16 172.26.162.68/32 port-unreachable
  30 permit icmp 172.26.162.0/16 172.26.162.68/32 echo-reply
  40 permit icmp 172.26.162.0/16 172.26.162.68/32 echo
  50 permit tcp 172.26.0.0/16 172.26.162.68/32 eq 22
  60 permit tcp 64.102.0.0/16 172.26.162.68/32 eq 22
  61 permit tcp 64.0.0.0/8 172.26.162.68/32 eq 22
  62 permit tcp 10.0.0.0/8 172.26.162.68/32 eq 22
  70 permit tcp 172.26.162.10/32 eq ftp 172.26.162.68/32 gt 1023 established
  80 permit tcp 172.26.162.10/32 eq ftp-data 172.26.162.68/32 gt 1023
  90 permit tcp 172.26.162.0/32 gt 1023 172.26.162.68/32 gt 1023 established
  100 permit udp 172.26.162.0/16 gt 1023 172.26.162.68/32 gt 1023
  102 permit udp 172.26.162.6/32 any eq ntp
  103 permit udp 172.26.162.9/32 any eq ntp
  110 deny ip any any log
ip access-list OBB-outbound
  10 permit ip 172.26.162.68/32 172.26.0.0/16
  20 permit ip 172.26.162.68/32 64.102.0.0/16
  21 permit ip 172.26.162.68/32 64.0.0.0/8
  22 permit ip 172.26.162.68/32 10.0.0.0/8
```

The access list can be applied to the management interface as follows:

```
interface mgmt0
  ip access-group OBB-inbound in
  ip access-group OBB-outbound out
  vrf member management
  ip address 172.26.162.68/16
```

Service Assurance Deployment

Queuing and bandwidth control on the Cisco Nexus 7000 is implemented in much the same way as it is implemented on the Cisco Nexus 5000. Egress queuing is used, however ingress queuing can be employed where it is understood that the traffic coming in is marked correctly and trusted. This is the case for this deployment. For bridged traffic CoS is untouched, while for routed traffic DSCP is copied to the CoS value.

You may also have to change the default mapping of cos-to-queue for this platform, as is the case in this deployment. It is changed to more closely resemble the mapping in the Cisco Unified Computing System^{TM} (UCS) and Cisco Nexus 5000 so bandwidth allocation can be consistent among the different platforms. The cos-to-queue map must be configured in the default VDC of the Cisco Nexus 7000. Queue names beginning with 8q2t-in-q correspond to ingress queues for 10Gbps ports, while queue names beginning with 1p7q4t-out-q correspond to egress queues for these ports:

```
DC09-N7K-1# show class-map type queuing
Type queuing class-maps
  ========================
class-map type queuing match-any 8q2t-in-q2
      Description: Classifier for ingress queue 2 of type 8q2t
      match cos 2

    class-map type queuing match-any 8q2t-in-q3
      Description: Classifier for ingress queue 3 of type 8q2t

    class-map type queuing match-any 8q2t-in-q4
      Description: Classifier for ingress queue 4 of type 8q2t
      match cos 4

    class-map type queuing match-any 8q2t-in-q5
      Description: Classifier for ingress queue 5 of type 8q2t
      match cos 5

    class-map type queuing match-any 8q2t-in-q6
      Description: Classifier for ingress queue 6 of type 8q2t
      match cos 6

    class-map type queuing match-any 8q2t-in-q7
      Description: Classifier for ingress queue 7 of type 8q2t

    class-map type queuing match-any 8q2t-in-q-default
      Description: Classifier for ingress default queue of type 8q2t
      match cos 0-1,3
   class-map type queuing match-any 1p7q4t-out-q2
      Description: Classifier for egress queue 2 of type 1p7q4t
      match cos 2

    class-map type queuing match-any 1p7q4t-out-q3
      Description: Classifier for egress queue 3 of type 1p7q4t

    class-map type queuing match-any 1p7q4t-out-q4
      Description: Classifier for egress queue 4 of type 1p7q4t
      match cos 4

    class-map type queuing match-any 1p7q4t-out-q5
      Description: Classifier for egress queue 5 of type 1p7q4t
      match cos 5
    class-map type queuing match-any 1p7q4t-out-q6
      Description: Classifier for egress queue 6 of type 1p7q4t
      match cos 6

    class-map type queuing match-any 1p7q4t-out-q7
      Description: Classifier for egress queue 7 of type 1p7q4t

    class-map type queuing match-any 1p7q4t-out-q-default
      Description: Classifier for egress default queue of type 1p7q4t
      match cos 0-1,3
```

These class maps can then be used for queuing and bandwidth control service polices to be applied to selected ports:

```
DC09-N7K-1-vdcA# sh policy-map type queuing ingress-queuing


  Type queuing policy-maps
  ========================

  policy-map type queuing ingress-queuing
    class type queuing 8q2t-in-q2
      bandwidth percent 25
    class type queuing 8q2t-in-q4
      bandwidth percent 15
    class type queuing 8q2t-in-q5
      bandwidth percent 20
    class type queuing 8q2t-in-q6
      bandwidth percent 15
    class type queuing 8q2t-in-q-default
      bandwidth percent 25
DC09-N7K-1-vdcA# sh policy-map type queuing egress-policy


  Type queuing policy-maps
  ========================

  policy-map type queuing egress-policy
    class type queuing 1p7q4t-out-q2
      bandwidth percent 25
    class type queuing 1p7q4t-out-q4
      bandwidth percent 15
    class type queuing 1p7q4t-out-q5
      bandwidth percent 20
    class type queuing 1p7q4t-out-q-default
      bandwidth percent 25
```

These policies are then applied to ports attaching to the access layer and the core layer:

```
interface port-channel45
  description *** Port Channel Between Aggregation and Access Switches E1/25, E1/26 ***
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-111,113,115,119-124,129-130
  switchport trunk allowed vlan add 132,140,161-162,481,490-493,499-500
  switchport trunk allowed vlan add 581-586,588-592,599-600,900-902
  switchport trunk allowed vlan add 905
  mtu 9000
  service-policy type queuing input ingress-queuing
  service-policy type queuing output egress-policy
  vpc 45
```

```
DC09-N7K-1-vdcA# sh run int e1/1

!Command: show running-config interface Ethernet1/1
!Time: Wed May 25 23:32:49 2011

version 5.1(3)

interface Ethernet1/1
  description *** DC09-N7K-1 E1/1 To CORE11-N7K-1 E1/1 ***
  rate-mode dedicated force
  udld disable
  service-policy type queuing output egress-policy
  service-policy type queuing input ingress-queuing
  no shutdown
```

Deploying the Services Chassis

In a multi-tenant environment, different tenants may require different levels of network services security protection. Tenants with very stringent security requirements require a host of virtual and physical appliances to satisfy their needs, while other tenants may require basic network protection. This architecture implements Cisco FWSM, Cisco Application Control Engine (ACE), and Cisco Intrusion Prevention System (IPS) within the services segment of the architecture. It is important to note the flexible nature of this architecture, where the network architect can use any combination of security appliances to create their own security offerings. Some of these offerings are foundational, while the others are optional. Cisco ACE load balancer is considered optional, while it is recommended that IPS and FWSM firewall module be implemented for every tenant, as shown in Figure 5.

Figure 5 Tenant Container Logical Topology

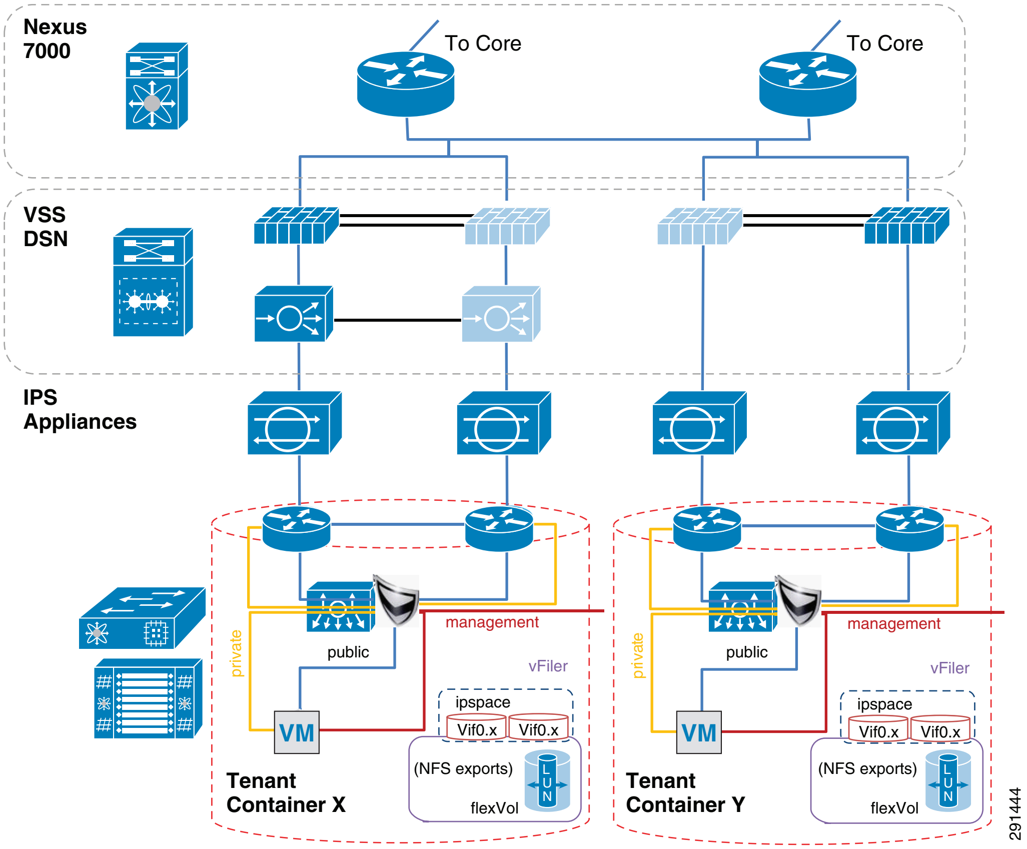

The ACE/FWSM and IPS services are employed external to each tenant's VRF instances. All these services are deployed in Layer 2 transparent mode. Additional information about services integration includes:

- Global routing instance resides on the Cisco Nexus 7000.
- Virtual Switching System (VSS) Domain Services Node (DSN) houses service modules.
- Virtual IPS appliances are positioned inline.
- IPS allows each tenant to create and enforce their own security policy and IPS is tuned to reduce false positives.

More information on the specific deployment steps for FWSM, IPS, and ACE are outlined in Deploying the Cisco Firewall Switching Module in the Network, Deploying the Cisco Intrusion Prevention System, and Deploying the Cisco Application Control Engine.

Deploying Services with Redundancy

Services can be deployed in a redundant manner. ACE and FWSM can be deployed by using a separate logical link between them within a VSS domain. A VSS domain is used to provide a single services domain with redundant chassis, as shown in Figure 6.

Figure 6 Single Services Domain with VSS

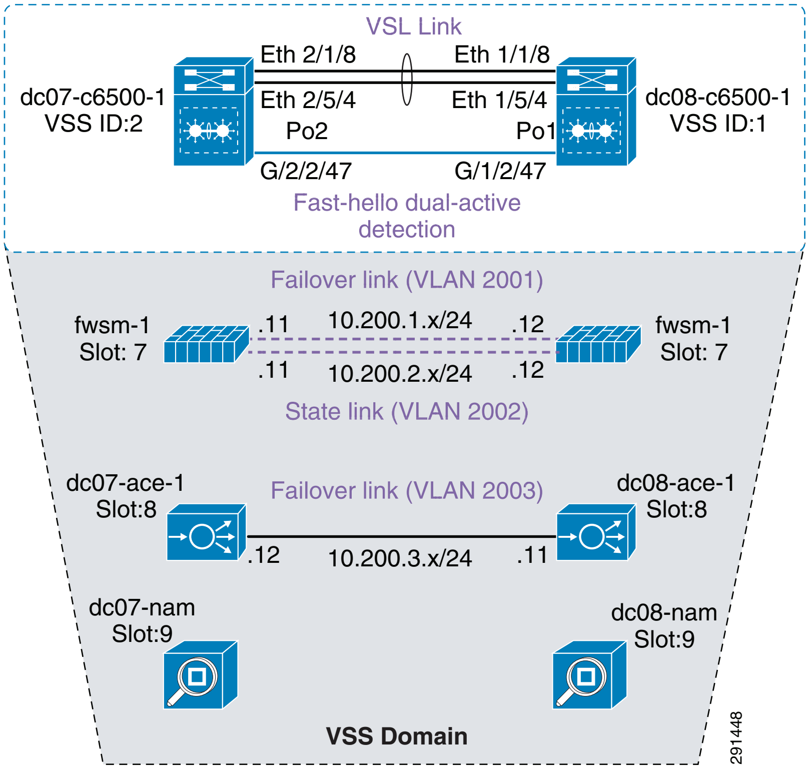

- Service modules leverage Virtual Switch Link (VSL) as data and control paths.
- VLAN interfaces are assigned to the modules through the supervisor.

  This allows multiple VLAN interfaces to support multiple virtual contexts.
- The autostate feature alerts modules to a state change of the VLAN on the supervisor.

  Use this feature when multiple paths exist for the VLAN.

The relevant configuration is:

```
svclc autostate
svclc multiple-vlan-interfaces
svclc switch 1 module 8 vlan-group 8,116,162
svclc switch 2 module 8 vlan-group 8,116,162
svclc vlan-group 8  2,101-103,111,2003
svclc vlan-group 116  116,117
svclc vlan-group 162  162
firewall autostate
firewall multiple-vlan-interfaces
firewall switch 1 module 7 vlan-group 7,9,116,162
firewall switch 2 module 7 vlan-group 7,9,116,162
firewall vlan-group 7  51-53,90,99,100,104,125,2001,2002
firewall vlan-group 9  57,166,494,586,587,594,597
analysis switch 1 module 9 management-port access-vlan 162
analysis switch 2 module 9 management-port access-vlan 162
```

Integration of Services Chassis

The services chassis is connected to the aggregation layer, as shown in Figure 7.

Figure 7 Aggregation-Services Chassis Connectivity

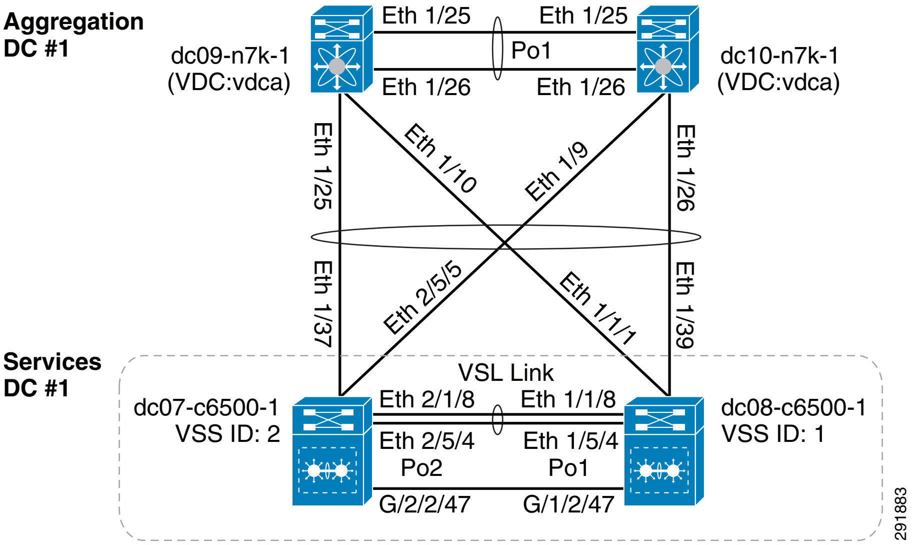

The specific network topology used for validation of the network and services chassis is shown in Figure 8.

Figure 8 SharePoint Tenant Logical Topology

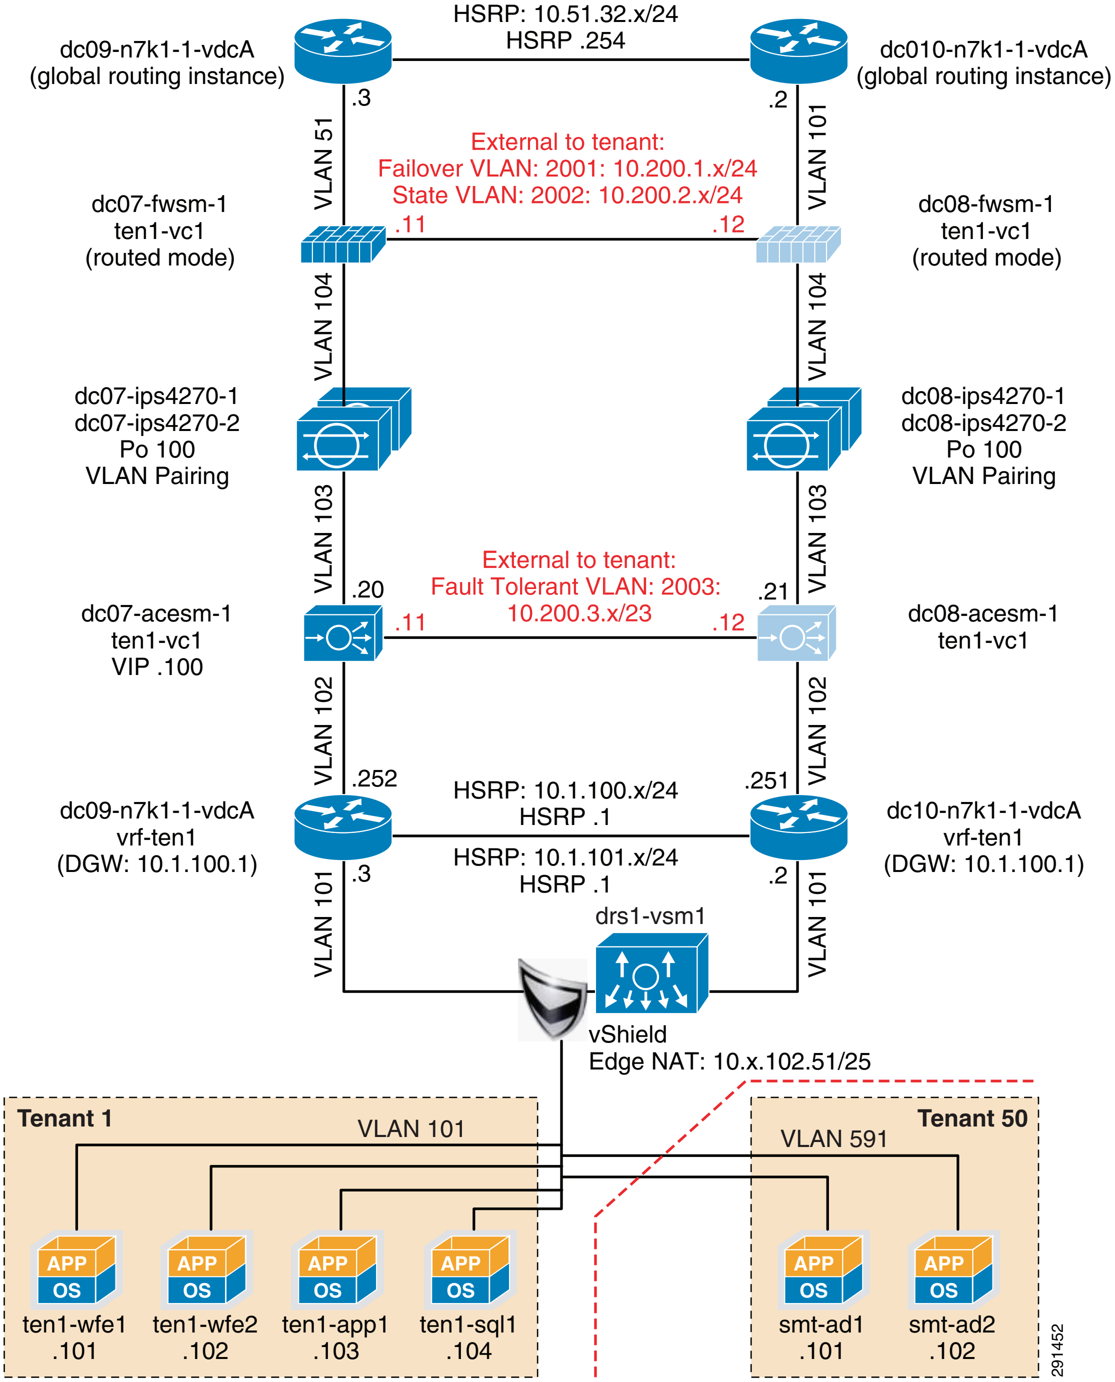

Hardening the Services Chassis

The security hardening deployment steps and general hardening guidelines are defined in:

- Cisco SAFE: A Security Reference Architecture
  http://www.cisco.com/en/US/solutions/collateral/ns170/ns896/ns954/white_paper_c11-527476.html
- Network Security Baseline
  http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/Baseline_Security/securbase.pdf

Deploying the Cisco Nexus 5000

The basic Cisco Nexus 5000 configuration can be found in FlexPod for VMware Deployment Model (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html). Cisco Nexus 5000 configurations for QoS and security hardening are shown in the sections that follow.

QoS Configuration on the Cisco Nexus 5000

For storage traffic originating on the Cisco Nexus 5000, classification is required to guarantee an acceptable service level for the traffic. This traffic can be classified with a global QoS policy on the Cisco Nexus 5000 and assigned to a system class:

```
ip access-list classify_COS_5
  10 permit ip 192.168.98.100/32 any
  20 permit ip 10.49.32.10/32 any
  30 permit ip any 192.168.98.100/32
  40 permit ip any 10.49.32.10/32
class-map type qos match-any Platinum_Traffic
  match access-group name classify_COS_5

policy-map type qos Global_Classify_NFS_Application
  class Platinum_Traffic
    set qos-group 2
```

The other CoS values that are configured via port-profiles on the Cisco Nexus 1000V must be put into system classes on the Cisco Nexus 5000 to allow for bandwidth control and queuing within the Cisco Nexus 5000 switching fabric.

```
class-map type qos match-any Platinum_Traffic
  match access-group name classify_COS_5
class-map type qos match-any Gold_Transactional
  match cos 6
class-map type qos match-any Bronze_Transactional
  match cos 2
class-map type qos match-any Silver_Transactional
  match cos 4
class-map type qos match-any Platinum_Transactional
  match cos 5
policy-map type qos Global_Classify_NFS_Application
  class Platinum_Traffic
    set qos-group 2
  class Platinum_Transactional
    set qos-group 2
  class Gold_Transactional
    set qos-group 3
  class Silver_Transactional
    set qos-group 4
  class Bronze_Transactional
    set qos-group 5
```

On the Cisco Nexus 5000, a queuing policy is applied globally to the output of all interfaces on the switch. Classification for each qos-group or system class is done by the aforementioned policy-map Global_Classify_NFS_Application.

```
class-map type queuing Gold_Traffic_Q
  match qos-group 3
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing Bronze_Traffic_Q
  match qos-group 5
class-map type queuing Silver_Traffic_Q
  match qos-group 4
class-map type queuing Platinum_Traffic_Q
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type queuing Global_BW_Queuing
  class type queuing Platinum_Traffic_Q
    priority
    bandwidth percent 20
  class type queuing Gold_Traffic_Q
    bandwidth percent 20
  class type queuing Silver_Traffic_Q
    bandwidth percent 20
  class type queuing Bronze_Traffic_Q
    bandwidth percent 15
  class type queuing class-fcoe
    bandwidth percent 15
  class type queuing class-default
    bandwidth percent 10

system qos
  service-policy type queuing output Global_BW_Queuing
```

Hardening the Cisco Nexus 5000

In the following example, the Cisco Nexus 5000V can be hardened to allow access only through SSH, FTP, and ICM in the inbound access list. The outbound access list allows access only to the OOB management network and the inside network. More information on the Cisco Nexus 5000 security capabilities and general hardening guidance can be found at: http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/security/502_n2_1m/b_Cisco_n5k_security_config_gd_rel_502_n2_1.pdf.

```
ip access-list OBB-inbound
  3 permit tcp 172.26.162.216/32 172.26.162.28/32
  9 permit ip 172.26.162.214/32 172.26.162.28/32
  10 permit icmp 172.26.162.0/16 172.26.162.28/32 ttl-exceeded
  20 permit icmp 172.26.162.0/16 172.26.162.28/32 port-unreachable
  30 permit icmp 172.26.162.0/16 172.26.162.28/32 echo-reply
  40 permit icmp 172.26.162.0/16 172.26.162.28/32 echo
  50 permit tcp 172.26.0.0/16 172.26.162.28/32 eq 22
  60 permit tcp 64.102.0.0/16 172.26.162.28/32 eq 22
  61 permit tcp 10.0.0.0/8 172.26.162.28/32 eq 22
  70 permit tcp 172.26.162.10/32 eq ftp 172.26.162.28/32 gt 1023 established
  80 permit tcp 172.26.162.10/32 eq ftp-data 172.26.162.28/32 gt 1023
  90 permit tcp 172.26.162.0/32 gt 1023 172.26.162.28/32 gt 1023 established
  100 permit udp 172.26.162.0/16 gt 1023 172.26.162.28/32 gt 1023
  120 permit tcp 172.26.162.214/32 eq tacacs 172.26.162.28/32
ip access-list OBB-outbound
  10 permit ip 172.26.162.28/32 172.26.0.0/16
  20 permit ip 172.26.162.28/32 64.102.0.0/16
  30 permit ip 172.26.162.28/32 10.0.0.0/8
```

The access list can be applied to the management interface as follows

```
interface mgmt0

  ip access-group OBB-inbound in
  ip access-group OBB-outbound out
  ip address 172.26.162.28/16
```

Deploying the Cisco Nexus 1010 and 1000V

In this deployment the combination of Cisco Nexus 1010 and Cisco Nexus 1000V are used to implement virtual switching at the access layer. For redundancy a pair of Cisco Nexus 1010 virtual appliances were configured. An active/passive pair of Cisco Nexus 1000V virtual service modules were configured within the Cisco Nexus 1010. All vSphere^{TM} networking is configured to use a Cisco Nexus 1000V virtual switch, including VM, NFS, and vMotion traffic. The Cisco Nexus 1010s are logically connected to the Cisco Nexus 5000 access layer switches, as shown in Figure 9.

Figure 9 Cisco Nexus 1010 Integration and Logical Connectivity

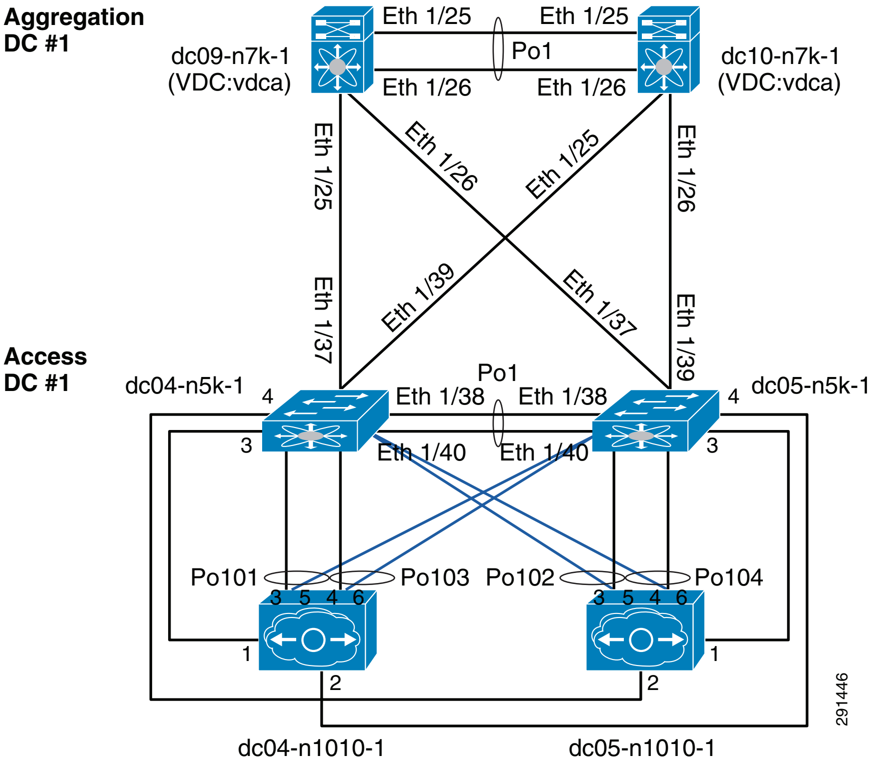

Classification Using the Cisco Nexus 1000V

In this deployment, classification is done on the Cisco Nexus 1000V. Certain port-profiles are marked with a designated CoS value through a service policy attached to the service profile. Any marking done by the host is ignored. This CoS value is in turn trusted by the virtual interface card on the Cisco UCS server blade. CoS values were assigned in the following fashion:

```
drs1-vsm1# show policy-map Silver_CoS_4
```

```
  Type qos policy-maps
  ====================

  policy-map type qos Silver_CoS_4
    class  class-default
      set cos 4

port-profile type vethernet vMotion_192_168_1
  vmware port-group
  switchport mode access
  switchport access vlan 901
  ip flow monitor NFMonitor input
  ip flow monitor NFMonitor output
  service-policy type qos input Silver_CoS_4
  no shutdown
  state enabled
```

Deploying Traffic Engineering with MAC-Pinning

Mac-pinning is implemented on the Cisco Nexus 1000V to fully utilize the redundant fabric uplinks. Port-profiles assigned to the same CoS value are split evenly between the two uplinks with static mac-pinning.

```
port-profile type vethernet Exch_NFS_192_168_120
  vmware port-group
  switchport mode access
  switchport access vlan 120
  ip flow monitor NFMonitor input
  ip flow monitor NFMonitor output
  service-policy type qos input Platinum_CoS_5
  pinning id 0         ?------ Fabric A
  no shutdown
  state enabled

port-profile type vethernet Infra_NFS_192_168_100
  vmware port-group
  switchport mode access
  switchport access vlan 600
  ip flow monitor NFMonitor input
  ip flow monitor NFMonitor output
  service-policy type qos input Platinum_CoS_5
  pinning id 1       ?-------- Fabric B
  no shutdown
  state enabled
```

The uplink port-profile must be configured to honor the mac-pinning and implement dynamic mac-pinning when necessary:

```
port-profile type ethernet VIC-vm-data-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 101,105-107,111,113,115,122,124,132,161,592-593
  channel-group auto mode on mac-pinning   ?------- Mac-pinning enabled
  no shutdown
  state enabled
```

Deploying the Cisco Firewall Switching Module in the Network

The FWSM is logically located between the global interface and the VRF interfaces for each tenant. FWSM is deployed in a bridged mode with different contexts for each tenant. Figure 10 shows the FWSM implemented within the services layer.

Figure 10 Integration of FWSM within Services Layer

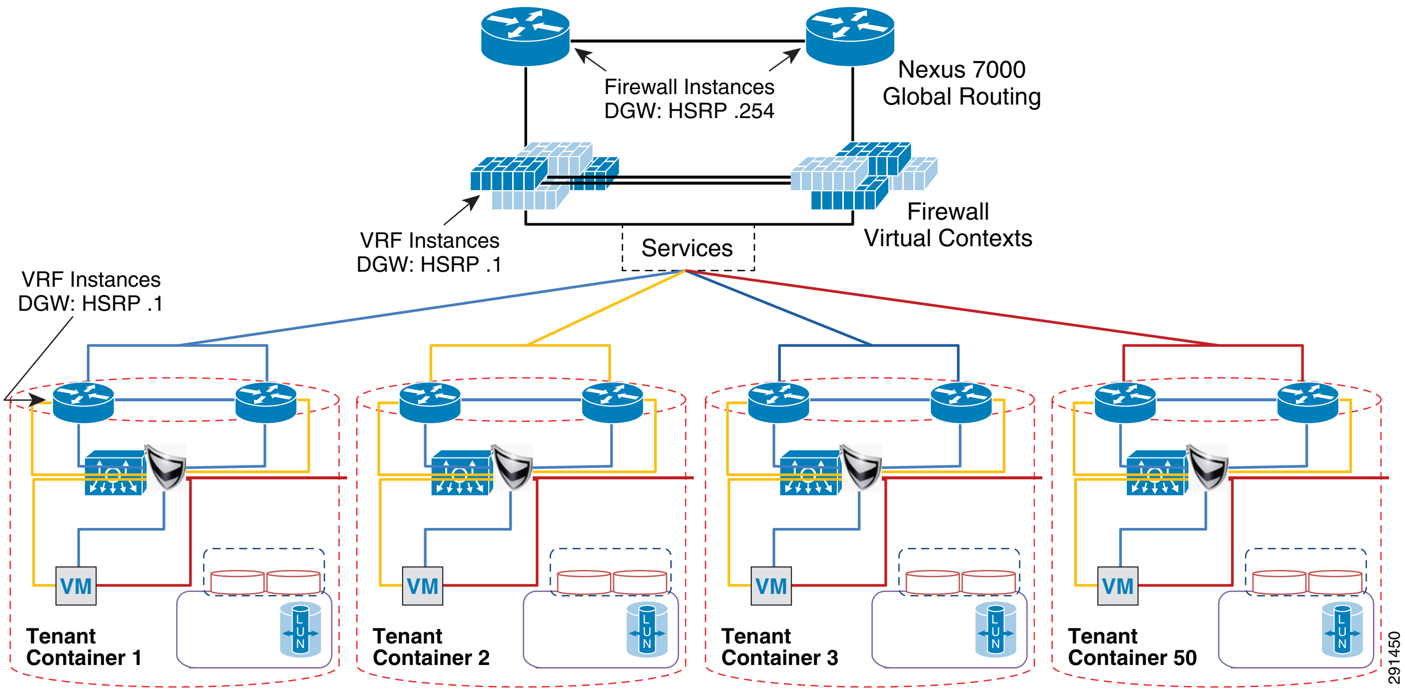

The following design principles are implemented in FWSM in the ESMT topology:

- The default gateway for tenant VMs is the Cisco Nexus 7000 VRF instance HSRP address.

  Each tenant has its own VLAN and associated subnet assignments. For example:
  - Tenant 1: 10.1.x.x/25
  - Tenant 2: 10.2.x.x/25
  - Tenant 50: 10.50.x.x/25
- A static route from VRF points to a firewall virtual context "inside" address.

  Additional services may be positioned but not detailed.
- There is a static route from each firewall instance to the Cisco Nexus 7000 global routing table.
- OSPF prefix-list injects routes into area.

The basic configuration for tenant-1 (SharePoint) is:

```
hostname ten1-vc1
!
interface Vlan104
 nameif inside
 security-level 100
 ip address 10.1.100.1 255.255.255.0 standby 10.1.100.101
!
interface Vlan51
```

```
 nameif outside
 security-level 0
 ip address 10.51.32.1 255.255.255.0 standby 10.51.32.2
!
mtu outside 1500
mtu inside 1500
icmp permit any outside
icmp permit any inside
no asdm history enable
arp timeout 14400
nat-control
static (inside,outside) 10.1.101.0 10.1.101.0 netmask 255.255.255.128
static (inside,outside) 10.1.100.0 10.1.100.0 netmask 255.255.255.0
route outside 0.0.0.0 0.0.0.0 10.51.32.254 1
route inside 10.1.101.0 255.255.255.128 10.1.100.254 1
```

Deploying the Cisco Adaptive Security Device Manager for Management of Firewall Module

For information on the steps necessary to deploy the Cisco Adaptive Security Device Manager (ASDM) to mange the firewall module, see: http://www.cisco.com/en/US/docs/security/asdm/6_2/user/guide/asdmug.pdf.

Deploying the Cisco Intrusion Prevention System

The Cisco IPS is deployed as an inline device and a separate context is used for each tenant. The IPS location within the services network is shown in Figure 8. The IPS can be configured using the CLI or the Cisco IPS Device Manager (IDM). In the configuration below, VLAN 103 and 104 are used by the IPS to provide inline intrusion protection within the SharePoint tenant. The CLI configuration is:

```
subinterface-type inline-vlan-pair
subinterface 1
description to dc07-c6500-1 ten1/2
vlan1 103
vlan2 104
exit
exit
exit
physical-interfaces TenGigabitEthernet7/1
description to dc08-c6500-1 ten1/2
admin-state enabled
duplex auto
speed auto
default-vlan 1
alt-tcp-reset-interface none
subinterface-type inline-vlan-pair
subinterface 1
description to dc08-c6500-1 ten1/2
vlan1 103
vlan2 104
exit

service analysis-engine
virtual-sensor vs0
description Tenant1 Virtual Sensor
anomaly-detection
operational-mode detect
exit
physical-interface TenGigabitEthernet7/0 subinterface-number 1
physical-interface TenGigabitEthernet7/1 subinterface-number 1
```

```
inline-TCP-session-tracking-mode virtual-sensor
inline-TCP-evasion-protection-mode strict
```

Once basic IPS networking parameters are configured, the IPS signatures can be tuned. Figure 11 shows how signatures can be tuned using Cisco IPS Manager Express (IME).

1. Go to the Configurations/Policies pane and click **All Signatures**.
2. Select the signature you want to tune, right click, and click **Enable**.
3. You can choose the actions and severity levels using the same procedure.

Figure 11 Tuning Signatures Using Cisco IPS Manager Express

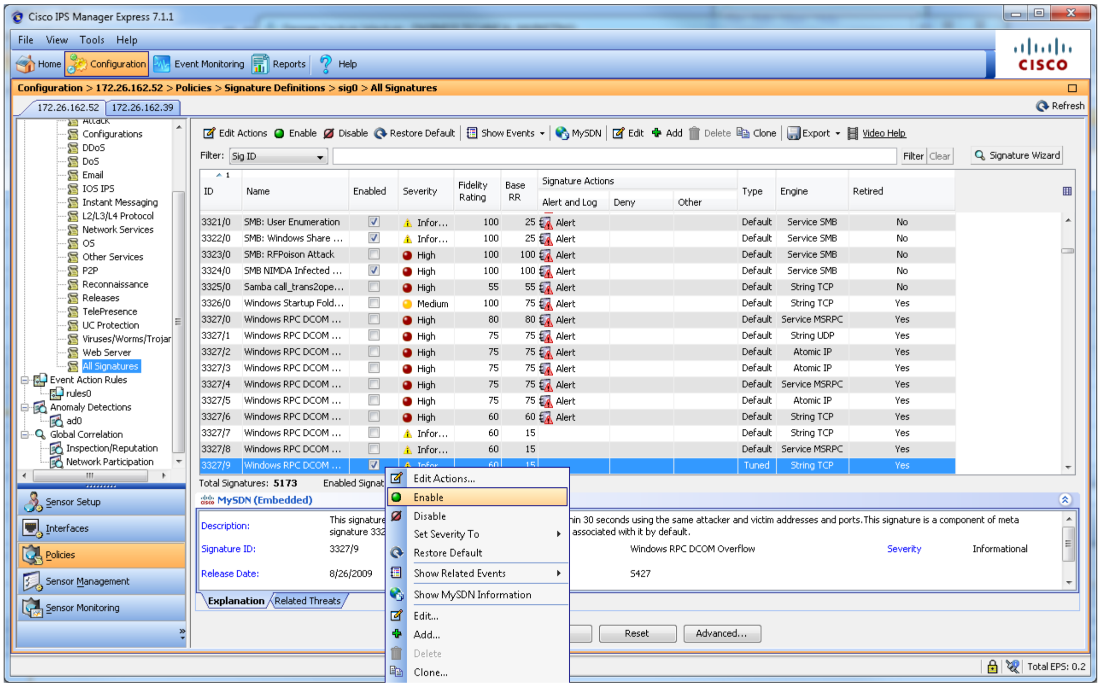

ID	Name	Enabled	Severity	Fidelity Rating	Base RR	Signature Actions			Type	Engine	Retired
						Alert and Log	Deny	Other			
3321/0	SMB: User Enumeration	☑	Infor...	100	25	Alert			Default	Service SMB	No
3322/0	SMB: Windows Share ...	☑	Infor...	100	25	Alert			Default	Service SMB	No
3323/0	SMB: RFPoison Attack	☐	High	100	100	Alert			Default	Service SMB	No
3324/0	SMB NIMDA Infected ...	☑	High	100	100	Alert			Default	Service SMB	No
3325/0	Samba call_trans2ope...	☐	High	55	55	Alert			Default	String TCP	No
3326/0	Windows Startup Fold...	☐	Medium	100	75	Alert			Default	String TCP	Yes
3327/0	Windows RPC DCOM ...	☐	High	80	80	Alert			Default	Service MSRPC	Yes
3327/1	Windows RPC DCOM ...	☐	High	75	75	Alert			Default	String UDP	Yes
3327/2	Windows RPC DCOM ...	☐	High	75	75	Alert			Default	Atomic IP	Yes
3327/3	Windows RPC DCOM ...	☐	High	75	75	Alert			Default	Atomic IP	Yes
3327/4	Windows RPC DCOM ...	☐	High	75	75	Alert			Default	Service MSRPC	Yes
3327/5	Windows RPC DCOM ...	☐	High	75	75	Alert			Default	Atomic IP	Yes
3327/6	Windows RPC DCOM ...	☐	High	60	60	Alert			Default	String TCP	Yes
3327/7	Windows RPC DCOM ...	☐	Infor...	60	15				Default	String TCP	Yes
3327/8	Windows RPC DCOM ...	☐	Infor...	60	15				Default	Service MSRPC	Yes
3327/9	Windows RPC DCOM ...	☑	Infor...	60	15				Tuned	String TCP	Yes

Total Signatures: **5173** Enabled Signat

MySDN (Embedded)

Description: This signature ... in 30 seconds using the same attacker and victim addresses and ports.This signature is a component of meta signature 332 ... associated with it by default.

Signature ID: 3327/9 Windows RPC DCOM Overflow Severity Informational

Release Date: 8/26/2009 S427

Explanation Related Threats

For more information on configuring IPS using Cisco IME or CLI, see: http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/ime/ime_interfaces.html.

Hardening the Cisco Intrusion Prevention System

Telnet can be disabled on the management interface and management access can be limited to certain subnets, as shown below.

```
network-settings
host-ip 172.26.162.52/23,172.26.162.1
host-name dc08-ips4270-1
telnet-option disabled
access-list 10.0.0.0/8
access-list 64.0.0.0/8
access-list 172.0.0.0/8
```

SSH should be configured as a means to access the CLI on IPS4270. The following link outlines the steps to configure SSH and other security parameters: http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_setup.html#wp1035869.

Installing and Using Cisco IPS Manager Express

For information on installing and using Cisco IPS Manager Express, see: http://www.cisco.com/en/US/docs/security/ips/trash/book_files/CLI7_1.pdf.

Deploying the Cisco Application Control Engine

The Cisco ACE module provides load balancing capability. In this deployment, ACE is configured to load balance between two front end servers 10.1.101.101 and 10.1.101.102 within the tenant-1 SharePoint tenant. The virtual interface address (VIP) is 10.1.100.100. ACE is also bridging between VLAN 102 and 103. The configuration for context ten1-vc-1 is shown in Figure 12.

Figure 12 Integration of ACE within Services Layer

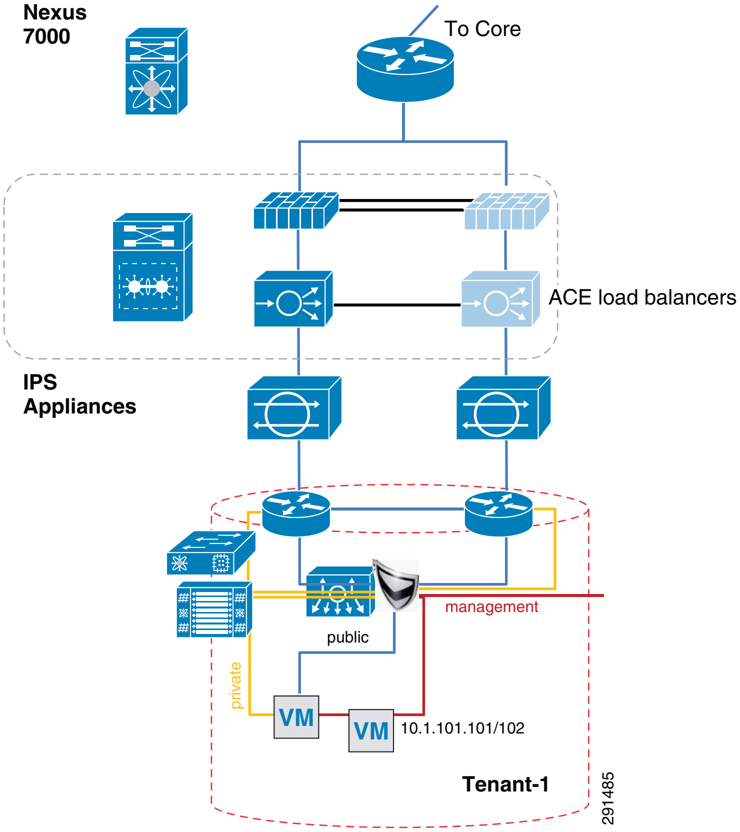

The CLI for the ACE configuration is:

```
dc08-ace-1/Admin# changeto ten1-vc1
dc08-ace-1/ten1-vc1# sh running-config
Generating configuration....


logging enable
logging buffered 7


access-list BPDU ethertype permit bpdu
```

```
access-list IPANYANY line 5 extended permit ip any any
access-list IPANYANY line 6 extended permit icmp any any


probe tcp IIS
  interval 2
  faildetect 2
  passdetect interval 10
  passdetect count 2
probe icmp PING
  interval 2
  faildetect 2

rserver host real1
  ip address 10.1.101.101
  inservice
rserver host real2
  ip address 10.1.101.102
  inservice

serverfarm host farm1
  predictor leastconns
  probe IIS
  probe PING
  rserver real1
    inservice
  rserver real2
    inservice

parameter-map type http PM-REUSE
  server-conn reuse
  case-insensitive
parameter-map type connection timeouts
  set tcp timeout embryonic 10
  set tcp ack-delay 400

sticky ip-netmask 255.255.255.255 address both group1
  timeout 720
  timeout activeconns
  replicate sticky
  serverfarm farm1

class-map match-all sp-vip
  2 match virtual-address 10.1.100.100 any

policy-map type loadbalance first-match lbpol
  class class-default
    sticky-serverfarm group1

policy-map multi-match LBPOL
  class sp-vip
    loadbalance vip inservice
    loadbalance policy lbpol
    loadbalance vip icmp-reply active
    connection advanced-options timeouts

service-policy input LBPOL
interface vlan 102
  bridge-group 1
  access-group input BPDU
  access-group input IPANYANY
  no shutdown
interface vlan 103
  bridge-group 1
  access-group input BPDU
  access-group input IPANYANY
  no shutdown

interface bvi 1
  ip address 10.1.100.21 255.255.255.0
  alias 10.1.100.22 255.255.255.0
  peer ip address 10.1.100.20 255.255.255.0
  no shutdown

ip route 10.1.101.0 255.255.255.0 10.1.100.254
ip route 0.0.0.0 0.0.0.0 10.1.100.1
```

Hardening the Cisco Application Control Engine

The ACE management interface can be hardened by implementing a policy map and assigning it to the interface, as shown below. More information on security hardening commands on the ACE module can be found in the Cisco Application Control Engine Administration Guide (http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/vA2_3_0/configuration/administration/guide/ace_adgd.pdf).

```
class-map type management match-any MANAGEMENT

  2 match protocol ssh source-address 172.26.0.0 255.255.0.0

  3 match protocol ssh source-address 10.0.0.0 255.0.0.0

  4 match protocol ssh source-address 64.0.0.0 255.0.0.0

  5 match protocol icmp source-address 172.26.0.0 255.255.0.0

  6 match protocol https source-address 64.0.0.0 255.0.0.0

  7 match protocol https source-address 172.0.0.0 255.0.0.0

  8 match protocol https source-address 10.0.0.0 255.0.0.0

  9 match protocol snmp source-address 172.26.0.0 255.255.0.0

class-map type management match-all class-Query

  2 match protocol icmp source-address 10.8.99.0 255.255.255.0

policy-map type management first-match MANAGEMENT

  class MANAGEMENT

    permit

policy-map type management first-match QUERY

  class class-Query

    permit


interface vlan 162

  ip address 172.26.162.56 255.255.0.0
  peer ip address 172.26.162.43 255.255.0.0

  service-policy input MANAGEMENT

  no shutdown
```

Additional commands configured on the Admin console that pertain to peering and defining tenant contexts are:

```
interface vlan 162
  ip address 172.26.162.56 255.255.0.0
  peer ip address 172.26.162.43 255.255.0.0
  access-group input IPANYANY
  service-policy input MANAGEMENT
  no shutdown

ft interface vlan 2003
  ip address 10.200.3.12 255.255.254.0
  peer ip address 10.200.3.11 255.255.254.0
  no shutdown

ft peer 1
  heartbeat interval 100
  heartbeat count 10
  ft-interface vlan 2003
ft group 1
  peer 1
  priority 150
  peer priority 50
  associate-context Admin
  inservice

role Operator

ip route 0.0.0.0 0.0.0.0 172.26.162.1

resource-class dc-gold
  limit-resource all minimum 0.00 maximum unlimited
  limit-resource sticky minimum 10.00 maximum unlimited
resource-class dc-silver
  limit-resource all minimum 0.00 maximum unlimited
resource-class rc-ten1
  limit-resource all minimum 0.00 maximum unlimited


context ten1-vc1
  description Tenant 1
  allocate-interface vlan 101-103
  member dc-gold
context ten2-vc1
  description Tenant 2
  allocate-interface vlan 116-117

snmp-server contact "ANM"
snmp-server location "ANM"
snmp-server community public group Network-Monitor

snmp-server host 172.26.165.36 traps version 2c public


ft group 2
  peer 1
  priority 150
  peer priority 50
  associate-context ten1-vc1
  inservice
ft group 3
  peer 1
  priority 150
  peer priority 50
  associate-context ten2-vc1
  inservice
username admin password 5<>.  role Admin domain
default-domain
username www password 5 <> role Admin domain de
fault-domain
ssh key rsa 1024 force
```

Managing Cisco ACE with Cisco Application Networking Manager

For information on Cisco Application Networking Manager (ANM) 4.2, see: http://www.cisco.com/en/US/docs/app_ntwk_services/data_center_app_services/application_networking_manager/4.2/user/guide/ug-book.pdf.

Deploying the Cisco Network Access Module

The Cisco Network Access Module (NAM) provides the capability to terminate Encapsulated Remote Switch Port Analyzer (ERSPAN) sessions and NetFlow connections. The Cisco Nexus 1000V can be configured to initiate ERSPAN sessions to the NAM, where the NAM can be used to capture packets and look at packet statistics. ERSPAN/NAM traffic flow is shown in Figure 13.

Figure 13 ERSPAN/NAM Traffic Flow

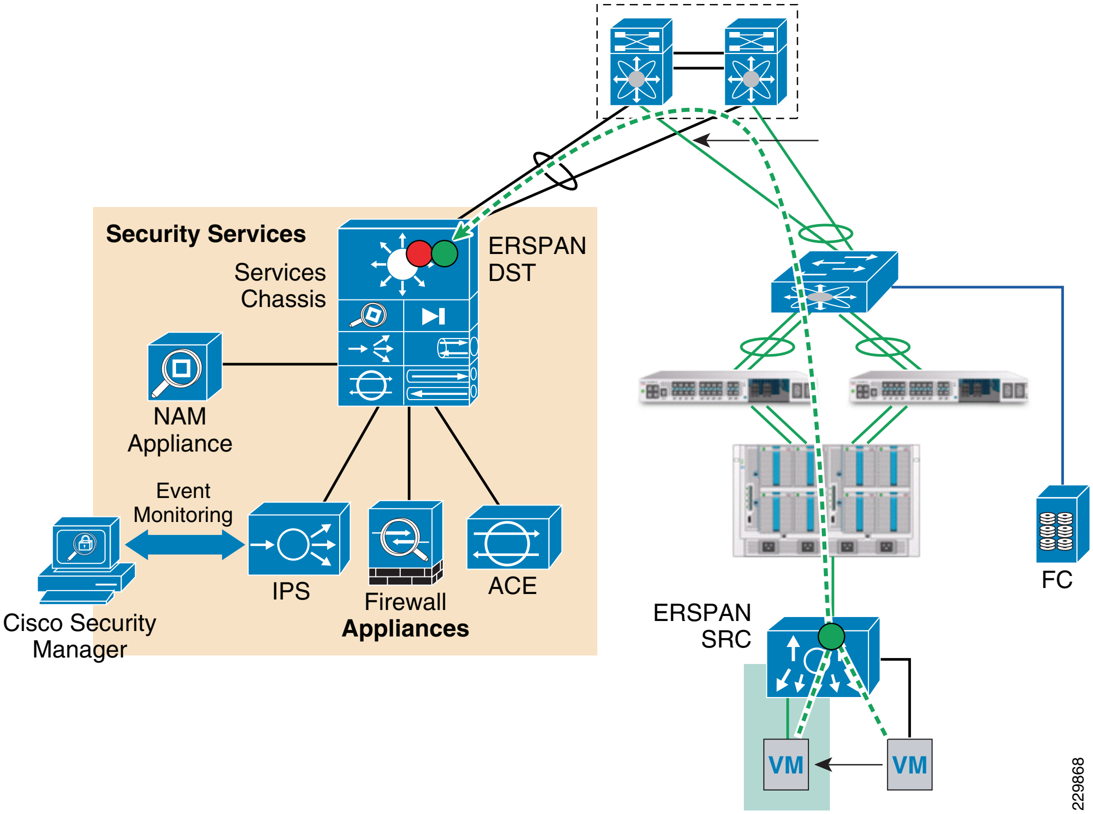

To configure ERSPAN between the Nexus 1000V and the services chassis:

1. For ERSPAN, create sessions for different range of VLANs. In this example, five ERSPAN sessions were defined as follows (only two ERSPAN sessions are shown):

```
monitor session 1 type erspan-source
  description ** to NAM dc08-namsm-1 **
  source vlan 101-110 rx
  destination ip 10.202.101.200
  erspan-id 1
  ip ttl 64
  ip prec 0
  ip dscp 0
  mtu 1500
  header-type 2
monitor session 50 type erspan-source
  description ** INFRA to NAM dc08-namsm-1 **
  source vlan 590-600 rx
  destination ip 10.202.101.200
  erspan-id 50
  ip ttl 64
  ip prec 0
  ip dscp 0
  mtu 1500
  header-type 2
```

2. The default setting for the ERSPAN sessions is shut to save CPU power on the Cisco Nexus 1000V; you should apply the no shut command on the appropriate ERSPAN session.

```
drs1-vsm1(config)# monitor session 50 type erspan-source
drs1-vsm1(config-erspan-src)# no shut
```

3. The command show monitor session should show the session is operational:

```
drs1-vsm1# sh monitor
Session  State        Reason                  Description
-------  -----------  ----------------------  --------------------------------
1        down         Session admin shut      ** to NAM dc08-namsm-1 **
2        down         Session admin shut      ** to NAM dc07-namsm-1 **
3        down         Session admin shut      ** to NAM dc08-namsm-1 **
4        down         Session admin shut      ** to NAM dc07-namsm-1 **
50       up           The session is up       ** INFRA to NAM dc08-namsm-1 **
```

The following steps must be configured on the Catalyst 6500 for each ERSPAN session (only sessions 3, 4, and 50 are shown below):

```
monitor session 50 type erspan-destination
 description ** N1k ERSPAN #50 - Tenant 50 INFRA **
 destination switch 1 analysis-module 9 data-port 2
 source
  erspan-id 50
  ip address 10.202.101.200

monitor session 3 type erspan-destination
 description ** N1k ERSPAN #1 - Tenanat 1 **
 destination switch 1 analysis-module 9 data-port 1
 source
  erspan-id 1
  ip address 10.202.101.200
!
!
monitor session 4 type erspan-destination
 description ** N1k ERSPAN #2 - Tenanat 2 **
 destination switch 2 analysis-module 9 data-port 1
 source
  erspan-id 2
  ip address 10.202.101.200
```

Configuring the Cisco Virtual Network Access Module as a NetFlow Collector

The Cisco Virtual Network Access Module (vNAM) is a virtual service blade on the Cisco Nexus 1010 that acts as a NetFlow collector. The traffic flow for the NAM and vNAM is shown in Figure 14.

Figure 14 NAM and vNAM Integration and Traffic Flow

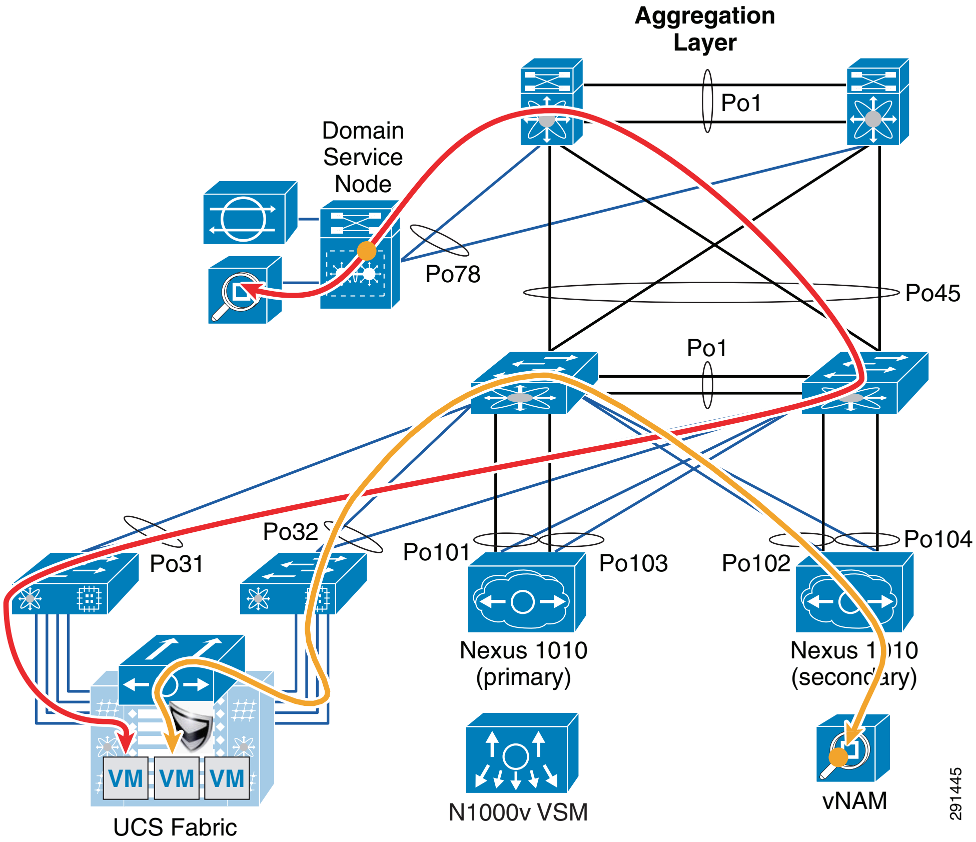

You can use the vNAM as a NetFlow collector by performing the following steps:

1. On the Cisco Nexus 1000, configure the following:

```
flow exporter vNAM
  destination 10.202.101.201
  transport udp 3000
  source mgmt0
  version 9
    template data timeout 1800
```

2. Verify NetFlow is operational by looking at statistics:

```
drs1-vsm1# show flow exporter

flow monitor NFMonitor
  record netflow-original
  exporter namsm-1
  exporter vNAM
  timeout active 1800
  cache size 4096
```

3. Configure NetFlow on all the required port-profiles as shown below:

```
port-profile type vethernet Management_172_26_162
  vmware port-group
  vmware max-ports 64
  switchport mode access
  switchport access vlan 162
  ip flow monitor NFMonitor input
  ip flow monitor NFMonitor output
  service-policy type qos input Gold_CoS_6
  no shutdown
  system vlan 162
  state enabled
```

Flow exporter vNAM:

```
    Destination: 10.202.101.201
    VRF: default (1)
    Destination UDP Port 3000
    Source Interface mgmt0 (172.26.163.109)
    Export Version 9
    Exporter Statistics
        Number of Flow Records Exported 50186128
        Number of Templates Exported 4313
        Number of Export Packets Sent 2239180
        Number of Export Bytes Sent 2262224948
```

The following steps must be performed to configure vNAM:

```
smt-n1010-1(config)# virtual-service-blade NAM
smt-n1010-1(config-vsb-config)# enable primary
Enter vsb image: [nam-4-2-1.iso]
Enter Management IPV4 address: 10.202.101.201
Enter Management subnet mask: 255.255.255.0
IPv4 address of the default gateway: 10.202.101.200
Enter HostName: smt-1010-nam
Setting Web user/passwd will enable port 80. Enter[y|n]: [n] y
Web User name: [admin]
Web User password: **********
smt-n1010-1(config-vsb-config)# interface data vlan 902
```

The configuration in the Cisco Nexus 1010 looks like this:

```
virtual-service-blade NAM
  virtual-service-blade-type name NAM-1.0
  interface data vlan 902
  ramsize 2048
  disksize 53
  no shutdown secondary
```

Deploying Cisco Unified Computing System

The basic Cisco UCS configuration can be found in FlexPod for VMware Deployment Model (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html). Details of the queuing and bandwidth control configuration used for ESMT are shown in Queuing and Bandwidth Control.

Queuing and Bandwidth Control

Queuing and bandwidth control are implemented within the Cisco UCS and at the access layer (Cisco Nexus 5000). Within the Cisco UCS, CoS values are assigned to a system class and given a certain percentage of the effective bandwidth. This is configured under the LAN tab in the Cisco UCS Manager (Figure 15).

Figure 15 Cisco UCS Manager—LAN Tab

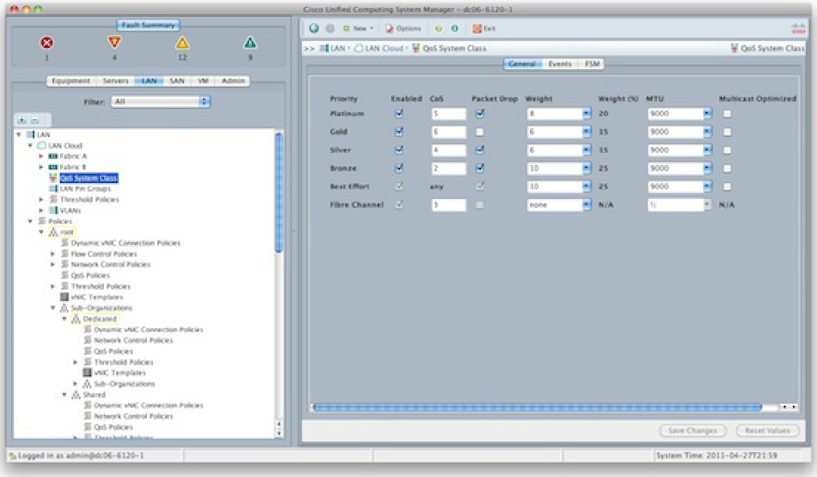

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	☑	5	☑	8	20	9000	□
Gold	☑	6	□	6	15	9000	□
Silver	☑	4	☑	6	15	9000	□
Bronze	☑	2	☑	10	25	9000	□
Best Effort	☑	any	☑	10	25	9000	□
Fibre Channel	☑	3	□	none	N/A	fc	N/A

The settings for the system classes are applied to traffic leaving the Cisco UCS 6100 fabric interconnects toward users and storage.

Deploying NetApp Storage

Deploying NetApp Data ONTAP

Each NetApp controller shares a unified storage architecture based on the Data ONTAP® 7G operating system and uses an integrated suite of application-aware manageability software. This provides an efficient consolidation of storage area network (SAN), network-attached storage (NAS), primary storage, and secondary storage on a single platform while allowing concurrent support for block and file protocols using Ethernet and Fibre Channel interfaces. These interfaces include Fibre Channel over Ethernet (FCoE), Network File System (NFS), Common Internet File System protocol (CIFS), and iSCSI. This common architecture allows businesses to start with an entry-level storage platform and easily migrate to the higher-end platforms as storage requirements increase, without learning a new OS, management tools, or provisioning processes.

To provide resilient system operation and high data availability, Data ONTAP 7G is tightly integrated into the hardware systems. The FAS systems use redundant, hot-swappable components. Combined with the patented dual-parity RAID-DP® (high-performance RAID 6), the net result can be superior data protection with little or no performance loss. For a higher level of data availability, Data ONTAP provides optional mirroring, backup, and disaster recovery solutions.

Active/Active Configuration

A NetApp active/active configuration keeps data available in the unlikely event of a controller failure or when controller maintenance might be necessary. Although this NetApp feature is highlighted in the “FlexPod for VMware Implementation Guide,” it is important to understand some of the rules pertaining to a properly configured active/active controller pair and how the configuration can be tested to enable proper failover if an event should occur.

The Data ONTAP 7.3 Active/Active Configuration Guide details the complete rules and procedures that should be followed to keep the controller configuration consistent with proper active/active controller operation (https://now.netapp.com/NOW/knowledge/docs/ontap/rel7351/pdfs/ontap/aaconfig.pdf).

A component of NetApp Data Fabric® Manager (DFM), Operations Manager can be used to check the active/active configuration on both controllers in the high-availability (HA) pair. This tool provides a detailed list of items that are not in compliance with a proper active/active configuration as well as the actions that should be taken to correct such issues.

To run the active/active configuration check tool within Operations Manager:

1. Navigate to the Operations Manager Web GUI.
2. Login as a user with Administrator privileges.
3. Click the **Control Center** tab.
4. Click the **Home** tab.
5. Click the **Member Details** tab.
6. Click the link for either storage controller in the Storage System column.
7. In the menu on the left under Storage Controller Tools, click **Check Configuration**.

Several checks are performed during the active/active configuration check. Potential issues or configuration errors are presented in list format and appropriate actions should be taken to mitigate these items.

Figure 16 shows an active-active configuration status.

Figure 16 Active/Active Configuration Status

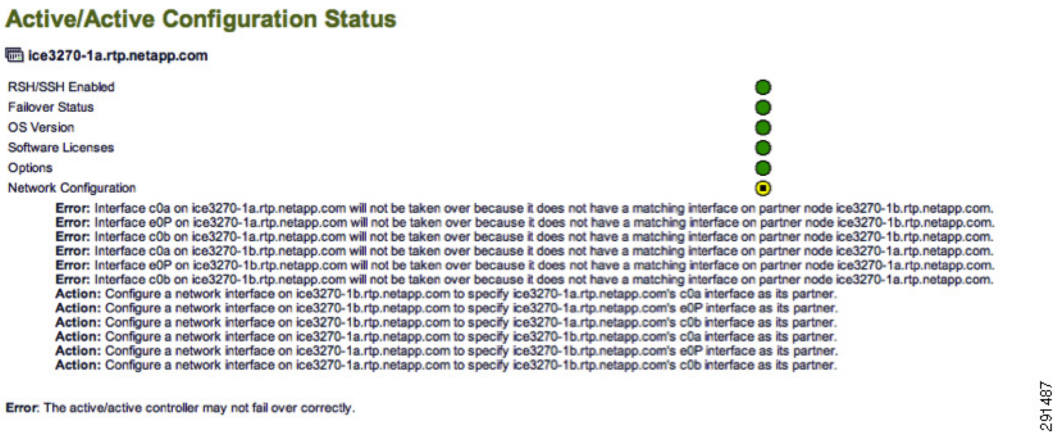

Active/Active Configuration Status

ice3270-1a.rtp.netapp.com

RSH/SSH Enabled ●
Failover Status ●
OS Version ●
Software Licenses ●
Options ●
Network Configuration ⊙

```
Error: Interface c0a on ice3270-1a.rtp.netapp.com will not be taken over because it does not have a matching interface on partner node ice3270-1b.rtp.netapp.com.
Error: Interface e0P on ice3270-1a.rtp.netapp.com will not be taken over because it does not have a matching interface on partner node ice3270-1b.rtp.netapp.com.
Error: Interface c0b on ice3270-1a.rtp.netapp.com will not be taken over because it does not have a matching interface on partner node ice3270-1b.rtp.netapp.com.
Error: Interface c0a on ice3270-1b.rtp.netapp.com will not be taken over because it does not have a matching interface on partner node ice3270-1a.rtp.netapp.com.
Error: Interface e0P on ice3270-1b.rtp.netapp.com will not be taken over because it does not have a matching interface on partner node ice3270-1a.rtp.netapp.com.
Error: Interface c0b on ice3270-1b.rtp.netapp.com will not be taken over because it does not have a matching interface on partner node ice3270-1a.rtp.netapp.com.
Action: Configure a network interface on ice3270-1b.rtp.netapp.com to specify ice3270-1a.rtp.netapp.com's c0a interface as its partner.
Action: Configure a network interface on ice3270-1b.rtp.netapp.com to specify ice3270-1a.rtp.netapp.com's e0P interface as its partner.
Action: Configure a network interface on ice3270-1b.rtp.netapp.com to specify ice3270-1a.rtp.netapp.com's c0b interface as its partner.
Action: Configure a network interface on ice3270-1a.rtp.netapp.com to specify ice3270-1b.rtp.netapp.com's c0a interface as its partner.
Action: Configure a network interface on ice3270-1a.rtp.netapp.com to specify ice3270-1b.rtp.netapp.com's e0P interface as its partner.
Action: Configure a network interface on ice3270-1a.rtp.netapp.com to specify ice3270-1b.rtp.netapp.com's c0b interface as its partner.

Error: The active/active controller may not fail over correctly.
```

Active/Active Configuration Example

The following example details a proper active/active configuration as it pertains to the implementation of vFiler® units and their associated attributes. Note that other advanced settings are also either required or recommended to be identical on both controllers. The list of advanced settings can be found in the Data ONTAP 7.3 Active/Active Configuration Guide (https://now.netapp.com/NOW/knowledge/docs/ontap/rel7351/pdfs/ontap/aaconfig.pdf).

Table 1 NetApp1 (Physical Controller)

vFiler	IPspace	VLANs	IP Addresses
Infrastructure_vfiler	Infrastructure_ipspace	592	10.50.102.100
		599	192.168.99.101
		600	192.168.100.100
SQL_vfiler	SQL_ipspace	124	10.3.102.100
		129	192.168.129.101
		130	192.168.130.100

Table 2 NetApp2 (Physical Controller)

vFiler	IPspace	VLANs	IP Addresses
Exchange_vfiler	Exchange_ipspace	115	10.2.102.100
		119	192.168.119.101
		120	192.168.120.100
Sharepoint_vfiler	Sharepoint_ipspace	106	10.1.102.100
		109	192.168.109.101
		110	192.168.110.100

```
NetApp1> ifconfig -a

<subsection of output>

vif0: flags=0xa3d0a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM,VLAN> mtu 9000
        ether 02:a0:98:0b:b0:81 (Enabled virtual interface)
vif0-106: flags=0x310a862<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-106 (not in use)
        ether 02:a0:98:0b:b0:82 (Enabled virtual interface)
vif0-109: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-109 (not in use)
        ether 02:a0:98:0b:b0:83 (Enabled virtual interface)
vif0-110: flags=0x310a862<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-110 (not in use)
        ether 02:a0:98:0b:b0:84 (Enabled virtual interface)
vif0-115: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-115 (not in use)
        ether 02:a0:98:0b:b0:85 (Enabled virtual interface)
vif0-119: flags=0x310a862<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-119 (not in use)
        ether 02:a0:98:0b:b0:86 (Enabled virtual interface)
vif0-120: flags=0x310a862<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-120 (not in use)
        ether 02:a0:98:0b:b0:87 (Enabled virtual interface)
vif0-124: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 10.3.102.100 netmask-or-prefix 0xffffff00 broadcast 10.3.102.255
        partner vif0-124 (not in use)
        ether 02:a0:98:0b:b0:88 (Enabled virtual interface)
vif0-129: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.129.101 netmask-or-prefix 0xffffff00 broadcast 192.168.129.255
        partner vif0-129 (not in use)
        ether 02:a0:98:0b:b0:89 (Enabled virtual interface)
vif0-130: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.130.100 netmask-or-prefix 0xffffff00 broadcast 192.168.130.255
        partner vif0-130 (not in use)
        ether 02:a0:98:0b:b0:8a (Enabled virtual interface)
vif0-592: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 10.50.102.100 netmask-or-prefix 0xffffff00 broadcast 10.50.102.255
        partner vif0-592 (not in use)
        ether 02:a0:98:0b:b0:8b (Enabled virtual interface)
vif0-599: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.99.101 netmask-or-prefix 0xffffff00 broadcast 192.168.99.255
        partner vif0-599 (not in use)
        ether 02:a0:98:0b:b0:8c (Enabled virtual interface)
vif0-600: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.100.100 netmask-or-prefix 0xffffff00 broadcast 192.168.100.255
        partner vif0-600 (not in use)
        ether 02:a0:98:0b:b0:8d (Enabled virtual interface)

</subsection of output>


NetApp2> ifconfig -a

<subsection of output>

vif0: flags=0xa3d0a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM,VLAN> mtu 9000
        ether 02:a0:98:0b:b0:91 (Enabled virtual interface)
vif0-106: flags=0x310a862<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 10.1.102.100 netmask-or-prefix 0xffffff00 broadcast 10.1.102.255
        partner vif0-106 (not in use)
        ether 02:a0:98:0b:b0:92 (Enabled virtual interface)
vif0-109: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.109.101 netmask-or-prefix 0xffffff00 broadcast 192.168.109.255
        partner vif0-109 (not in use)
        ether 02:a0:98:0b:b0:93 (Enabled virtual interface)
vif0-110: flags=0x310a862<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.110.100 netmask-or-prefix 0xffffff00 broadcast 192.168.110.255
        partner vif0-110 (not in use)
        ether 02:a0:98:0b:b0:94 (Enabled virtual interface)
vif0-115: flags=0x394a863<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 10.2.102.100 netmask-or-prefix 0xffffff00 broadcast 10.2.102.255
        partner vif0-115 (not in use)
        ether 02:a0:98:0b:b0:95 (Enabled virtual interface)
vif0-119: flags=0x310a862<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.119.101 netmask-or-prefix 0xffffff00 broadcast 192.168.119.255
        partner vif0-119 (not in use)
        ether 02:a0:98:0b:b0:96 (Enabled virtual interface)
vif0-120: flags=0x310a862<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        inet 192.168.120.100 netmask-or-prefix 0xffffff00 broadcast 192.168.120.255
        partner vif0-120 (not in use)
        ether 02:a0:98:0b:b0:97 (Enabled virtual interface)
vif0-124: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-124 (not in use)
        ether 02:a0:98:0b:b0:98 (Enabled virtual interface)
vif0-129: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-129 (not in use)
        ether 02:a0:98:0b:b0:99 (Enabled virtual interface)
vif0-130: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-130 (not in use)
        ether 02:a0:98:0b:b0:9a (Enabled virtual interface)
vif0-592: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-592 (not in use)
        ether 02:a0:98:0b:b0:9b (Enabled virtual interface)
vif0-599: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
```

```
        partner vif0-599 (not in use)
        ether 02:a0:98:0b:b0:9c (Enabled virtual interface)
vif0-600: flags=0x394a863<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 9000
        partner vif0-600 (not in use)
        ether 02:a0:98:0b:b0:9d (Enabled virtual interface)
</subsection of output>


NetApp1> ipspace list
Number of ipspaces configured: 4
default-ipspace                 (e0M e0P e0a e0b)
Infrastructure_ipspace          (vif0-592 vif0-599 vif0-600)
SQL_ipspace                     (vif0-124 vif0-129 vif0-130)
Exchange_ipspace                (vif0-115 vif0-119 vif0-120)
Sharepoint_ipspace              (vif0-106 vif0-109 vif0-110)



NetApp2> ipspace list
Number of ipspaces configured: 4
default-ipspace                 (e0M e0P e0a e0b)
Infrastructure_ipspace          (vif0-592 vif0-599 vif0-600)
SQL_ipspace                     (vif0-124 vif0-129 vif0-130)
Exchange_ipspace                (vif0-115 vif0-119 vif0-120)
Sharepoint_ipspace              (vif0-106 vif0-109 vif0-110)




NetApp1> rdfile /etc/rc

<subsection of output>

vif create lacp vif0 -b ip e6a e6b
vlan create vif0 106 110 119 120 124 130 583 589 590 592 600
ifconfig vif0 partner vif0
ifconfig e0a `hostname`-e0a netmask 255.255.0.0 mtusize 1500 mediatype auto flowcontrol
full partner e0a
ifconfig vif0-106 `hostname`-vif0-106 netmask 255.255.255.0 mtusize 9000 partner vif0-106
ifconfig vif0-109 `hostname`-vif0-109 netmask 255.255.255.0 mtusize 9000 partner vif0-109
ifconfig vif0-110 `hostname`-vif0-110 netmask 255.255.255.0 mtusize 9000 partner vif0-110
ifconfig vif0-115 `hostname`-vif0-115 netmask 255.255.255.0 mtusize 9000 partner vif0-115
ifconfig vif0-119 `hostname`-vif0-119 netmask 255.255.255.0 mtusize 9000 partner vif0-119
ifconfig vif0-120 `hostname`-vif0-120 netmask 255.255.255.0 mtusize 9000 partner vif0-120
ifconfig vif0-124 `hostname`-vif0-124 netmask 255.255.255.0 mtusize 9000 partner vif0-124
ifconfig vif0-129 `hostname`-vif0-129 netmask 255.255.255.0 mtusize 9000 partner vif0-129
ifconfig vif0-130 `hostname`-vif0-130 netmask 255.255.255.0 mtusize 9000 partner vif0-130
ifconfig vif0-592 `hostname`-vif0-592 netmask 255.255.255.0 mtusize 9000 partner vif0-592
ifconfig vif0-599 `hostname`-vif0-599 netmask 255.255.255.0 mtusize 9000 partner vif0-599
ifconfig vif0-600 `hostname`-vif0-600 netmask 255.255.255.0 mtusize 9000 partner vif0-600

</subsection of output>


NetApp2> rdfile /etc/rc

<subsection of output>

vif create lacp vif0 -b ip e6a e6b
vlan create vif0 106 110 119 120 124 130 583 589 590 592 600
ifconfig vif0 partner vif0
ifconfig e0a `hostname`-e0a netmask 255.255.0.0 mtusize 1500 mediatype auto flowcontrol
full partner e0a
ifconfig vif0-106 `hostname`-vif0-106 netmask 255.255.255.0 mtusize 9000 partner vif0-106
ifconfig vif0-109 `hostname`-vif0-109 netmask 255.255.255.0 mtusize 9000 partner vif0-109
ifconfig vif0-110 `hostname`-vif0-110 netmask 255.255.255.0 mtusize 9000 partner vif0-110
ifconfig vif0-115 `hostname`-vif0-115 netmask 255.255.255.0 mtusize 9000 partner vif0-115
ifconfig vif0-119 `hostname`-vif0-119 netmask 255.255.255.0 mtusize 9000 partner vif0-119
ifconfig vif0-120 `hostname`-vif0-120 netmask 255.255.255.0 mtusize 9000 partner vif0-120
ifconfig vif0-124 `hostname`-vif0-124 netmask 255.255.255.0 mtusize 9000 partner vif0-124
ifconfig vif0-129 `hostname`-vif0-129 netmask 255.255.255.0 mtusize 9000 partner vif0-129
ifconfig vif0-130 `hostname`-vif0-130 netmask 255.255.255.0 mtusize 9000 partner vif0-130
ifconfig vif0-592 `hostname`-vif0-592 netmask 255.255.255.0 mtusize 9000 partner vif0-592
ifconfig vif0-599 `hostname`-vif0-599 netmask 255.255.255.0 mtusize 9000 partner vif0-599
ifconfig vif0-600 `hostname`-vif0-600 netmask 255.255.255.0 mtusize 9000 partner vif0-600

</subsection of output>


NetApp1> rdfile /etc/hosts

<subsection of output>

127.0.0.1 localhost
# 0.0.0.0       NetApp1-vif0
172.26.162.10   NetApp1 NetApp1-e0a
# 0.0.0.0       NetApp1-e0b
# 0.0.0.0       NetApp1-e0c
# 0.0.0.0       NetApp1-e0d
# 0.0.0.0       NetApp1-e0e
# 0.0.0.0       NetApp1-e0f
0.0.0.0         NetApp1-vif0-106
10.1.102.100    NetApp2-vif0-106
0.0.0.0         NetApp1-vif0-109
192.168.109.101 NetApp2-vif0-109
0.0.0.0         NetApp1-vif0-110
192.168.110.100 NetApp2-vif0-110
0.0.0.0         NetApp1-vif0-115
10.2.102.100    NetApp2-vif0-115
0.0.0.0         NetApp1-vif0-119
192.168.119.101 NetApp2-vif0-119
0.0.0.0         NetApp1-vif0-120
192.168.120.100 NetApp2-vif0-120
0.0.0.0         NetApp1-vif0-124
0.0.0.0         NetApp2-vif0-124
192.168.129.101 NetApp1-vif0-129
0.0.0.0         NetApp2-vif0-129
192.168.130.100 NetApp1-vif0-130
0.0.0.0         NetApp2-vif0-130
10.50.102.100   NetApp1-vif0-592
0.0.0.0         NetApp2-vif0-592
192.168.99.101  NetApp1-vif0-599
0.0.0.0         NetApp2-vif0-599
192.168.100.100 NetApp1-vif0-600
0.0.0.0         NetApp2-vif0-600

</subsection of output>


NetApp2> rdfile /etc/hosts

<subsection of output>

127.0.0.1 localhost
# 0.0.0.0       NetApp2-vif0
172.26.162.11   NetApp2 NetApp2-e0a
# 0.0.0.0       NetApp2-e0b
# 0.0.0.0       NetApp2-e0c
# 0.0.0.0       NetApp2-e0d
# 0.0.0.0       NetApp2-e0e
# 0.0.0.0       NetApp2-e0f
0.0.0.0         NetApp1-vif0-106
10.1.102.100    NetApp2-vif0-106
0.0.0.0         NetApp1-vif0-109
192.168.109.101 NetApp2-vif0-109
0.0.0.0         NetApp1-vif0-110
192.168.110.100 NetApp2-vif0-110
0.0.0.0         NetApp1-vif0-115
10.2.102.100    NetApp2-vif0-115
0.0.0.0         NetApp1-vif0-119
192.168.119.101 NetApp2-vif0-119
0.0.0.0         NetApp1-vif0-120
192.168.120.100 NetApp2-vif0-120
0.0.0.0         NetApp1-vif0-124
0.0.0.0         NetApp2-vif0-124
192.168.129.101 NetApp1-vif0-129
0.0.0.0         NetApp2-vif0-129
192.168.130.100 NetApp1-vif0-130
0.0.0.0         NetApp2-vif0-130
10.50.102.100   NetApp1-vif0-592
0.0.0.0         NetApp2-vif0-592
192.168.99.101  NetApp1-vif0-599
0.0.0.0         NetApp2-vif0-599
192.168.100.100 NetApp1-vif0-600
0.0.0.0         NetApp2-vif0-600

</subsection of output>
```

Security Hardening

NetApp recommends that you configure and enable an administrative user other than root immediately after initially setting up Data ONTAP. NetApp also recommends enabling secure protocols (refer to Table 3 for the default and recommended settings for each protocol) and disabling unsecure and unused protocols.

Data ONTAP has two distinct types of access: user data access through the NAS and SAN modules and administrative access through the storage controller's administrative module. Use caution when assigning elevated administrative access for any user. Data ONTAP has many security-related options that can be set to meet particular requirements. NetApp strongly recommends the use of secure administration methods for Data ONTAP and the disabling of any administrative protocols deemed to be of high risk. All management protocols used should be secure and encrypted whenever possible. For example, https should be used instead of http whenever possible.

Several services should be considered for disabling. Depending on the enterprise security structure, the state of any service depends on where the service is deployed and how deep it is in the infrastructure. The services in Table 3 do not require the purchase of additional licensing from NetApp. All of these settings are configurable through the options command within the Data ONTAP CLI.

For more information on securing NetApp storage controllers, refer to:

- TR3649: Best Practices for Secure Configuration of Data ONTAP 7G (http://media.netapp.com/documents/tr-3649.pdf)
- TR-3358: Role-Based Access Control for Data ONTAP 7G (http://media.netapp.com/documents/tr-3358.pdf)

Table 3 Data ONTAP services

Service	Default State Data ONTAP	Recommended Setting
File Transfer Protocol (FTP)	Off	Off
File Transfer Protocol over SSH (SFTP)	Off	On
File Transfer Protocol over SSL (FTPS)	Off	On
FilerView https://<filer_IP>/na_admin (httpd.admin.ssl.enable)	Off	On
FilerView http://<filer_IP>/na_admin (httpd.admin.enable)	On	Off
Network Data Management Protocol (NDMP)	Off	On
Remote Shell (rsh)	On	Off
RIP - routed (RIPv1)	On	Off
Secure Shell Service (ssh)	Off	On
Secure Shell v1 (SSHv1)	Off	Off
Secure Shell v2 (SSHv2)	Off	On
Secure Sockets Service (ssl)	Off	On
Secure Sockets Layer v2 (SSLv2)	On	Off
Secure Sockets Layer v3 (SSLv3)	On	On
Simple Network Management Protocol (SNMPv1) ("Public" as a community string)	On	Off
Simple Network Management Protocol (SNMPv3)	Off	On
Telnet	On	Off
Transport Layer Security v1 (TLSv1)	Off	On
Trivial File Transfer Protocol (TFTP)	Off	Off

Centralized Logging

NetApp Data ONTAP has the ability to transmit logs to a centralized syslog server. The storage system receives information about the syslog server and what logs to send from the user configurable `syslog.conf` file. There are several pieces to Data ONTAP log messages including the "level" of logging and the "facility" of logging. The "level" of logging describes the severity of the message; the "facility" of logging describes the part of the system generating the message.

For further information on configuring centralized logging through the `syslog.conf` file, refer to the Data ONTAP 7.3 System Administration Guide. For Data ONTAP 7.3.5.1, which is configured in the current version of FlexPod for VMware, see: (https://now.netapp.com/NOW/knowledge/docs/ontap/rel7351/pdfs/ontap/sysadmin.pdf).

To enable logging:

1. Mount the `/etc` directory from an administrative machine.
2. Create the `/etc/syslog.conf` file from the provided `/etc/syslog.conf.sample` file that exists by default in the `/etc` directory.

The following is an example /etc/syslog.conf file:

```
NetApp> rd1file /etc/syslog.conf
# $Id: //depot/prod/ontap/R7.3.3x/files/syslog.conf.sample#1 $
# Copyright (c) 1994-1996 Network Appliance.
# All rights reserved.
# Sample syslog.conf file.  Copy to /etc/syslog.conf to use.
# You must use TABS for separators between fields.

# Log messages of priority info or higher to the console and to /etc/messages
*.info                                  /dev/console
*.info                                  /etc/messages
*.info                                  @smt-splunk-1.smt.com

# Edit and uncomment following line to log all messages of priority
# err or higher and all kernel messages to a remote host, e.g. adminhost
# *.err;kern.*                          @adminhost

# Edit and uncomment following line to log all messages of priority
# err or higher and all kernel messages to the local7 facility of the
# syslogd on a remote host, e.g. adminhost.
# *.err;kern.*                          local7.*@adminhost

# Edit and uncomment following line to log all messages of priority
# err or higher and all kernel messages to a remote host, e.g. adminhost,
# at priority debug.
# *.err;kern.*                          *.debug@adminhost

# Edit and uncomment following line to log all messages of priority
# err or higher and all kernel messages to the local5 facility of the
# syslogd on a remote host, e.g. adminhost, at priority info.
# *.err;kern.*                          local5.info@adminhost
```

Centralized Authentication

NetApp Data ONTAP supports centralized authentication for both data access and storage controller administration. Supported centralized authentication and directory services include Microsoft Active Directory®, NIS, LDAP, and Kerberos. Microsoft Active Directory is the most commonly used and widely supported directory service and is therefore recommended in an ESMT environment.

Each physical or virtual (vFiler) storage controller can be added to the same Active Directory domain or added to separate independent Active Directory domains depending on the controller's purpose and owner (administrator or tenant). The procedure for joining a physical or virtual storage controller to an Active Directory domain is exactly the same and differs only in which context (physical or vFiler unit) the commands are executed.

For information on integrating NetApp storage into a Microsoft Active Directory domain as well as other directory services for both Windows and UNIX/Linux clients, see:

- Data ONTAP 7.3 System Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel7351/pdfs/ontap/sysadmin.pdf
- Data ONTAP 7.3 File Access and Protocols Management Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel7351/pdfs/ontap/filesag.pdf
- TR-3771: Windows File Services Best Practices with NetApp Storage Systems
  http://media.netapp.com/documents/tr-3771.pdf
- TR-3458: Unified Windows and Linux Authorization Using Microsoft Active Directory LDAP as a Directory Store
  http://media.netapp.com/documents/tr-3458.pdf

- TR-3457: Unified Windows and UNIX Authentication Using Microsoft Active Directory Kerberos http://media.netapp.com/documents/tr-3457.pdf
- TR-3464: Integration of a NetApp Storage System with a UNIX Based LDAP Server http://media.netapp.com/documents/tr-3464.pdf

Deploying NetApp Operations Manager and Provisioning Manager

Security Hardening

Secure management interfaces should be used with NetApp Operations Manager and Provisioning Manager. Operations Manager is managed with a secure Web page (https on TCP port 8443). Provisioning Manager and Protection Manager are managed using the NetApp Management Console, which is a small client-based application installed on a management workstation. This workstation could also contain other management applications such as the VMware vSphere Client. Communication between the NetApp Management Console and Operations Manager also uses https on TCP port 8448. Both of these secure interfaces are enabled as part of the FlexPod for VMware deployment, along with disabling of the non-secure interfaces. Also, all communication between Operations Manager and Provisioning Manager and the NetApp storage systems should be secure. Storage system management should be setup with ssh (TCP port 22), https (TCP port 443), and snmpV3, which in this implementation encrypts the user login password. The vFiler units are managed with http (TCP port 80), but this management is all done within non-routed VLANs with private address spaces.

Centralized Logging

The Microsoft Windows version of NetApp Operations Manager and Provisioning Manager does not send logs to a centralized syslog server. These logs are stored locally on the server hosting Operations Manager. However, all actions initiated by NetApp Operations Manager and Provisioning Manager on the NetApp storage systems generate appropriately labeled storage system logs, which can be forwarded to a centralized syslog server.

Centralized Authentication

Since the versions of NetApp Operations Manager and Provisioning Manager included in FlexPod run on a Windows Server, centralized authentication can be set up in the Active Directory domain. RBAC can be used to provide appropriate capabilities to users on different resources within the environment.

Deploying NetApp Virtual Storage Console

NetApp Virtual Storage Console (VSC) is a VMware vCenter™ plug-in that provides four functions:

- Proper timeout and adapter settings on the VMware ESXi hosts managed by the vCenter
- Rapid cloning of virtual machines and templates utilizing NetApp FlexClone® technology
- Provisioning of tenant NFS datastores
- Scheduled backups of virtual machines utilizing NetApp Snapshot® copies with the capability of triggering NetApp SnapMirror® updates as part of a disaster recovery plan

NetApp recommends installing VSC on a Windows Server 2008 R2 virtual machine along with NetApp Operations Manager and Provisioning Manager. On this virtual machine, in addition to the management network interface, network interfaces with jumbo frames setup should be configured in each storage network in the infrastructure for both infrastructure and tenant vFiler unit direct access. This direct access to the vFiler management interface is needed both to clone virtual machines and to back up virtual machines within a mounted datastore that is contained within a vFiler unit.

For more information on installing and configuring the NetApp VSC, see the NetApp Virtual Storage Console 2.1 for VMware vSphere Installation and Administration Guide (https://now.netapp.com/NOW/knowledge/docs/virtual_storage_console/relvsc21/pdfs/install.pdf).

Security Hardening

Secure interfaces should be used between NetApp VSC and VMware vCenter, for which VSC serves as a plug-in. In this case, a secure interface is the only option because both directions of communication between these two components occur using https on TCP port 443. Also, all communication between VSC and the NetApp storage systems should be secure. Storage system management should be set up with ssl (TCP port 443). vFilers are managed with http (tcp port 80), but this management is all done within non-routed VLANs with private address spaces.

Centralized Logging

VSC does not directly generate logs. However, all actions initiated by VSC on the NetApp storage systems generate appropriately labeled storage system logs that can be forwarded to a centralized syslog server. Additionally, all tasks utilizing the VSC vCenter Plug-in generate events and logs within vCenter.

Centralized Authentication

Since VSC runs on a Windows Server, centralized authentication can be set up in the Active Directory domain. Tasks performed by users using the VSC vCenter Server Plug-in are authenticated through vCenter.

Deploying NetApp SANscreen

Security Hardening

Because Netapp SANscreen® Server and Data Warehouse Server run on Microsoft Windows VMs, it is important to always follow Microsoft Windows security best practices whenever possible. This may include running the native Windows firewall or disabling any un-used protocols, and so on, to ensure the most secure environment possible. Microsoft provides a Security Compliance Toolkit to aid in the proper configuration of Microsoft Windows 2008 systems in the environment. Best practice documentation and guidelines can also be obtained for other operating system versions and platforms.

In order for NetApp SANscreen to work properly in a secure manner within the environment, it is important to create a specific SANscreen administrative user (that can be a cloud admin or storage admin) who possesses the capabilities to perform the necessary SANscreen operations and functions on the NetApp controllers. Once the SANscreen administrative user is created in Active Directory, the following procedure should be followed to provide the minimum capabilities to the user on each NetApp controller:

1. Create a SANscreen Admin role on the storage controller(s) that includes the necessary capabilities.

   Syntax:

   ```
   useradmin role add <role_name> -a capability1[,capability2…]
   ```

   Example:

   ```
   NetApp1> useradmin role add SANScreen_admin -a 
   api-*,login-http-admin,security-api-vfiler
   ```

   Where:

   - api-*—Allows the NetApp SANscreen server to execute the appropriate API commands on the NetApp controller(s)
   - login-http-admin—Allows the NetApp SANscreen server to connect to the NetApp controller(s) through HTTP(s)
   - security-api-vfiler—Allows the NetApp SANscreen server to execute the appropriate API commands to retrieve vFiler unit information

2. Assign the newly created SANscreen_admin role to the group to which the SANscreen administrator belongs on the storage controller.

   Syntax:

   ```
   useradmin group modify <group_name> -r <custom_role>
   ```

   Example:

   ```
   NetApp1> useradmin group modify SANScreen_admin_grp -r SANScreen_admin
   ```

In securing the SANscreen Server installation, it is also important to choose the correct secure protocols to use for communication to and from the local server. Always make sure to choose the secure protocol option, which, for example, might be using HTTPS versus HTTP. In this example, HTTPS is the preferred secure protocol. This methodology should apply to all aspects of the environment, not just the NetApp SANscreen Server configuration.

Centralized Authentication

Outside of typical Windows Active Directory logon authentication, NetApp SANscreen Server leverages two methods for user authentication: LDAP version 2 or 3 and a local SANscreen user database. SANscreen first tries to authenticate the user against LDAP if LDAP has been enabled. Otherwise the user is authenticated against the local SANscreen user database. To enable LDAP:

1. Login as a user with Administrative privileges.
2. Navigate to the User Management section from the Main Menu.
3. Click **LDAP Configuration** and enable LDAP by clicking the checkbox labeled **Enable LDAP**.
4. Fill in the appropriate values as desired to fully configure LDAP authentication.

Figure 17 shows the LDAP configuration screen.

Figure 17 LDAP Configuration

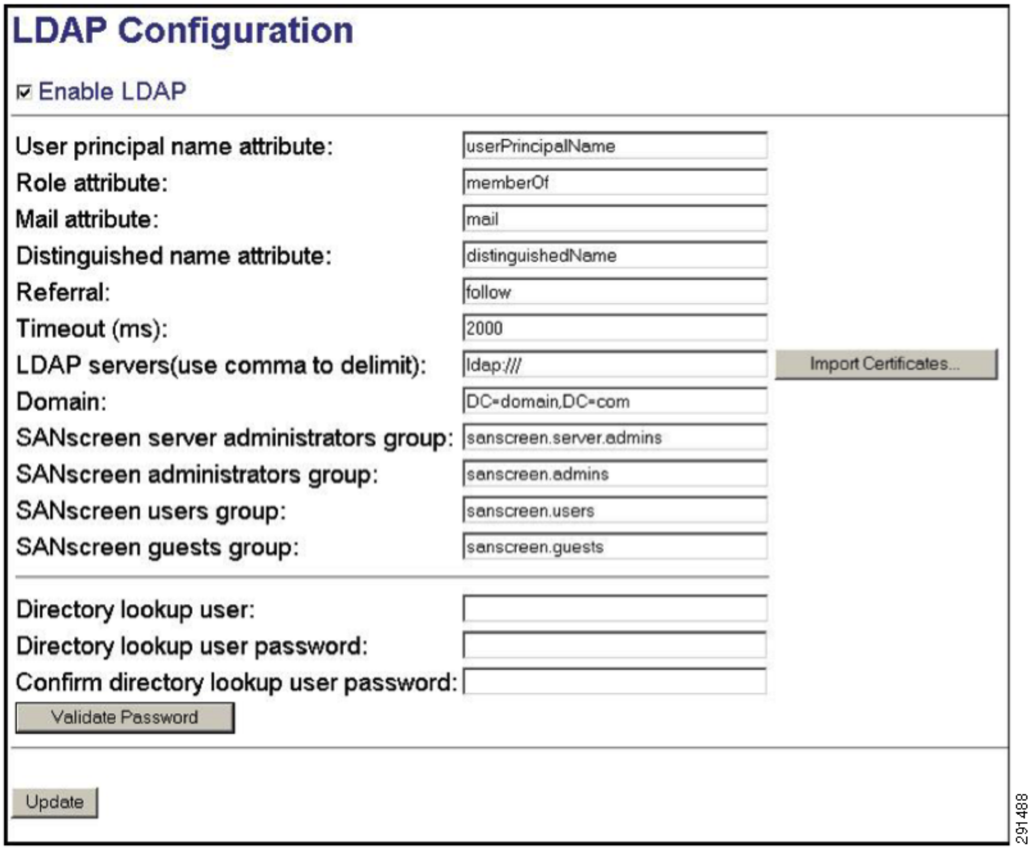

LDAP Configuration

☑ Enable LDAP

User principal name attribute:	userPrincipalName
Role attribute:	memberOf
Mail attribute:	mail
Distinguished name attribute:	distinguishedName
Referral:	follow
Timeout (ms):	2000
LDAP servers(use comma to delimit):	ldap:///
Domain:	DC=domain,DC=com
SANscreen server administrators group:	sanscreen.server.admins
SANscreen administrators group:	sanscreen.admins
SANscreen users group:	sanscreen.users
SANscreen guests group:	sanscreen.guests
Directory lookup user:	
Directory lookup user password:	
Confirm directory lookup user password:	

Deploying VMware

Installing and Configuring VMware vSphere

Installation and configuration procedures for VMware vSphere components (ESXi and vCenter Server) are identical to those documented in the FlexPOD deployment guide (NetApp Technical Report/TR-3892). The ESMT architecture runs on top of FlexPOD, therefore the standard procedures from FlexPOD deployment need to be completed.

After installation of vSphere, follow the steps below to configure the ESX clusters:

- Option 1—Create two separate ESX Clusters, one for management/infrastructure and the other for pure compute resources dedicated to tenants. This option was used with FlexPod validation.
- Option 2—Create one ESX Cluster and create sub-resource pools for management/infrastructure and sub-resource pool for individual tenants. This option was used in ESMT validation.

In the ESMT validation, option 2 was used. See Figure 18 for a reference of Cluster/Resource Pool layout for both management/infrastructure and individual tenants (Exchange, SQL, and SharePoint).

Figure 18 Cluster/Resource Pool Layout Example

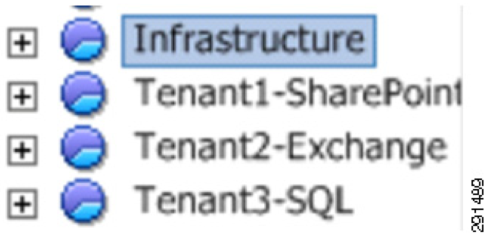

⊞ Infrastructure

⊞ Tenant1-SharePoint

⊞ Tenant2-Exchange

⊞ Tenant3-SQL

1. Create a new data center.
2. Create new ESX Cluster by adding all ESXi servers to the cluster.
3. Configure the settings in Figure 19 for HA and DRS.

Note The values here serve as a reference only. Refer to the Enhanced Secure Multi-Tenancy Design Guide (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldg_V2.html) for design considerations on recommended values based on high availability protection and individual tenant criticality.

Figure 19 HA Cluster Level Settings

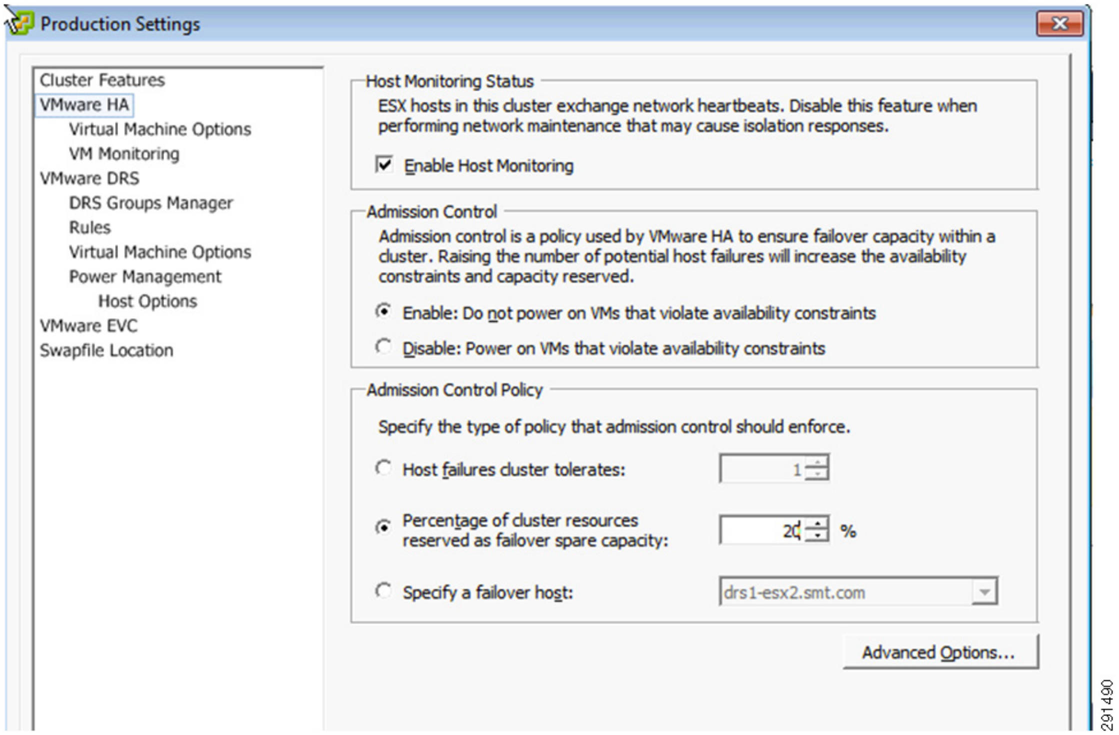

Host Monitoring Status

ESX hosts in this cluster exchange network heartbeats. Disable this feature when performing network maintenance that may cause isolation responses.

☑ Enable Host Monitoring

Admission Control

Admission control is a policy used by VMware HA to ensure failover capacity within a cluster. Raising the number of potential host failures will increase the availability constraints and capacity reserved.

● Enable: Do not power on VMs that violate availability constraints

○ Disable: Power on VMs that violate availability constraints

Admission Control Policy

Specify the type of policy that admission control should enforce.

○ Host failures cluster tolerates: 1

● Percentage of cluster resources reserved as failover spare capacity: 20 %

○ Specify a failover host: drs1-esx2.smt.com

Note Virtual Machine restart priority (on ESX host failover)—Ensure that the infrastructure and critical production tenant VMs are configured with High VM restart priority. Figure 20 illustrates the infrastructure and critical tenant VM configured with High VM restart priority in the ESMT environment: vShield Manager, vShield Edge, vShield App, vCenter Server Primary/Standby, Virtual Storage Console, and Tenant 2 and Tenant 3 (Exchange and SQL Server).

Figure 20 VMware HA Virtual Machine Restart Priority

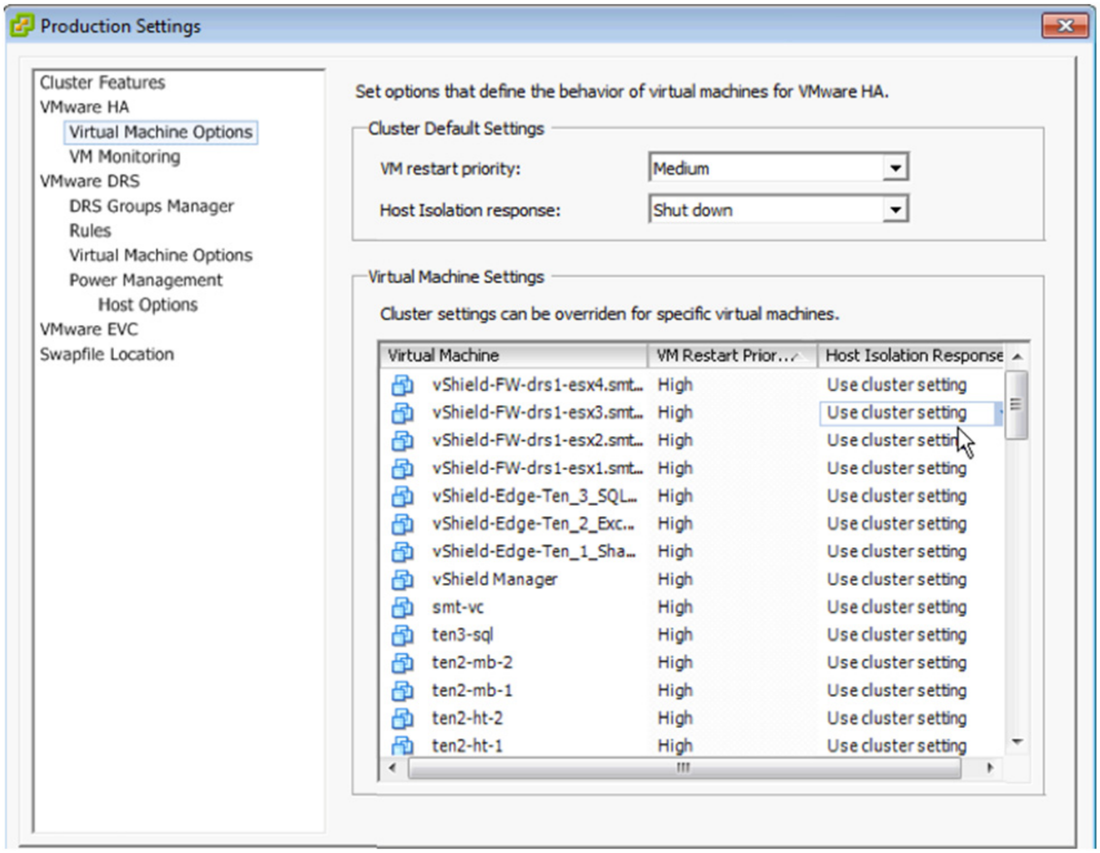

Set options that define the behavior of virtual machines for VMware HA.

Cluster Default Settings

VM restart priority: Medium

Host Isolation response: Shut down

Virtual Machine Settings

Cluster settings can be overriden for specific virtual machines.

Virtual Machine	VM Restart Prior...	Host Isolation Response
vShield-FW-drs1-esx4.smt...	High	Use cluster setting
vShield-FW-drs1-esx3.smt...	High	Use cluster setting
vShield-FW-drs1-esx2.smt...	High	Use cluster setting
vShield-FW-drs1-esx1.smt...	High	Use cluster setting
vShield-Edge-Ten_3_SQL...	High	Use cluster setting
vShield-Edge-Ten_2_Exc...	High	Use cluster setting
vShield-Edge-Ten_1_Sha...	High	Use cluster setting
vShield Manager	High	Use cluster setting
smt-vc	High	Use cluster setting
ten3-sql	High	Use cluster setting
ten2-mb-2	High	Use cluster setting
ten2-mb-1	High	Use cluster setting
ten2-ht-2	High	Use cluster setting
ten2-ht-1	High	Use cluster setting

For VM monitoring, ensure cluster-wide setting is set to Enable, with threshold level set to High. Given that a large number of virtual machines in an ESMT environment are serving infrastructure management and production tenants, it is easier to set the cluster-wide setting to accommodate the most common setting. For individual non-production/less critical tenant virtual machines, the individual VM setting can override the cluster-level setting.

Figure 21 VMware HA VM Monitoring Sensitivity

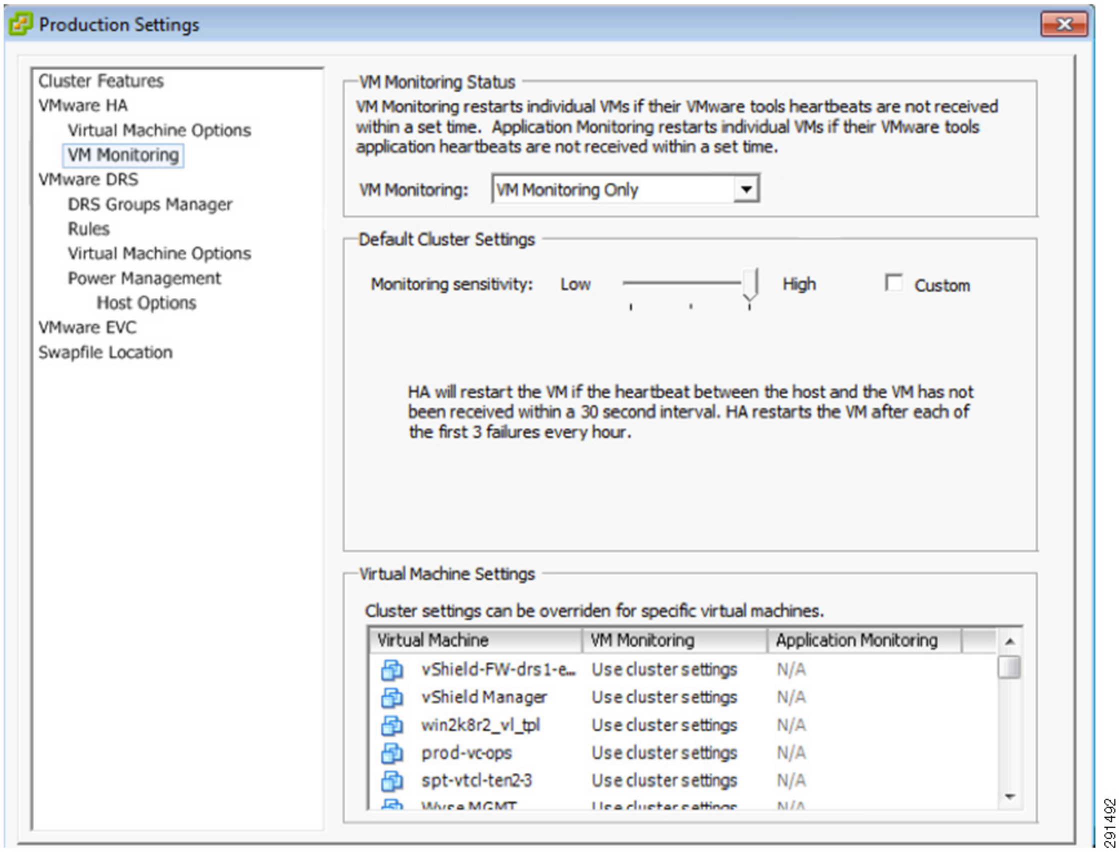

VM Monitoring Status

VM Monitoring restarts individual VMs if their VMware tools heartbeats are not received within a set time. Application Monitoring restarts individual VMs if their VMware tools application heartbeats are not received within a set time.

VM Monitoring: VM Monitoring Only

Default Cluster Settings

Monitoring sensitivity: Low — High □ Custom

HA will restart the VM if the heartbeat between the host and the VM has not been received within a 30 second interval. HA restarts the VM after each of the first 3 failures every hour.

Virtual Machine Settings

Cluster settings can be overriden for specific virtual machines.

Virtual Machine	VM Monitoring	Application Monitoring
vShield-FW-drs1-e...	Use cluster settings	N/A
vShield Manager	Use cluster settings	N/A
win2k8r2_vl_tpl	Use cluster settings	N/A
prod-vc-ops	Use cluster settings	N/A
spt-vtcl-ten2-3	Use cluster settings	N/A
Wyse MGMT	Use cluster settings	N/A

Leave all VMware DRS and VMware EVC cluster-level settings at deafult.

For Swapfile Location, ensure the swap file location is set to store in a specific datastore in the ESX host. This is the best practice recommendation when SRM is used for disaster recovery, as the dedicated datastore storing virtual machine swap file is not replicated.

Figure 22 Swapfile Location

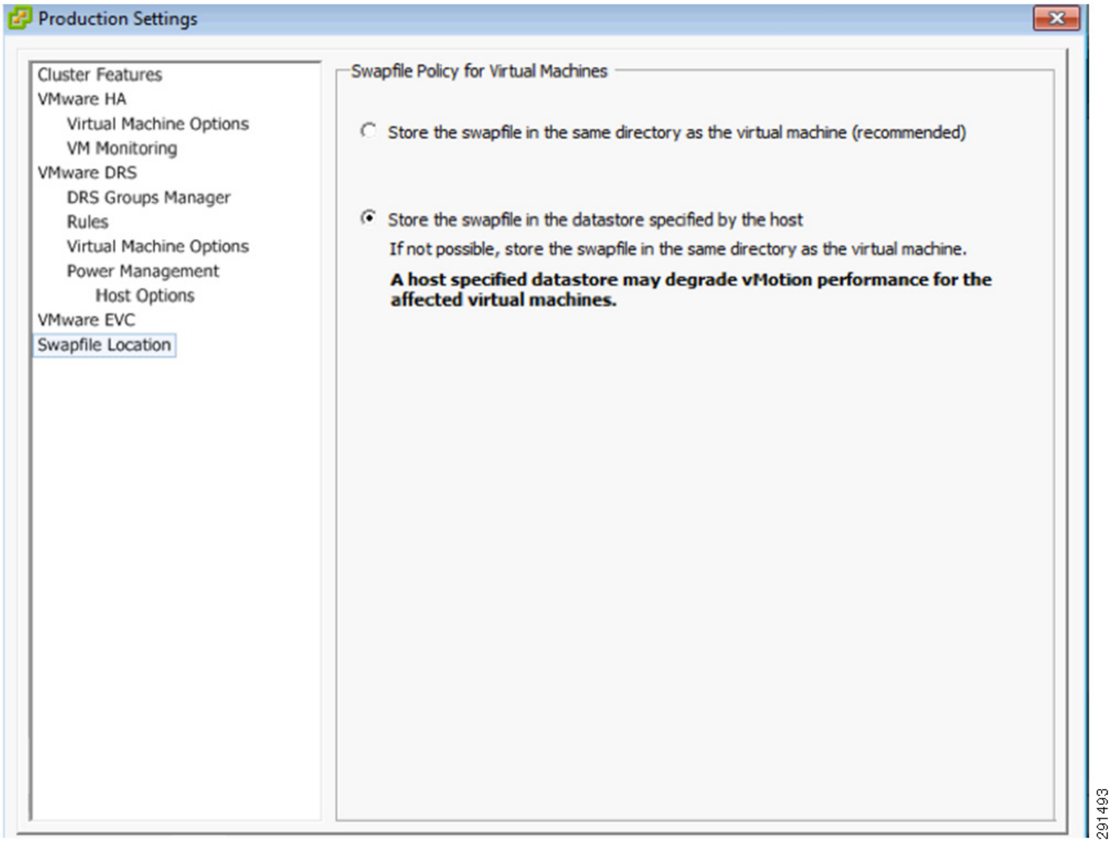

Swapfile Policy for Virtual Machines

○ Store the swapfile in the same directory as the virtual machine (recommended)

● Store the swapfile in the datastore specified by the host
If not possible, store the swapfile in the same directory as the virtual machine.
A host specified datastore may degrade vMotion performance for the affected virtual machines.

Resource Pool settings for Infrastructure, Exchange, SQL, and SharePoint Tenants:

- Each tenant is allocated a sub-resource pool under the Production cluster.
- Resource Pool share values are configured on the criticality of the tenant. In the validation:
  - The Exchange and SQL tenants are most critical, therefore CPU and memory share values are set to High, with Expandable Reservation enabled.
  - The SharePoint tenant CPU and memory share values are set to Normal, without Expandable Reservation.
  - The Infrastructurre tenant CPU and memory share values are set to Normal, with Expandable Reservation enabled.
- In terms of resource pool Limits, the infrastructure, Exchange, and SQL tenants have setting of Unlimited; the SharePoint tenant has a resource cap for both CPU and memory resources, restricting the maximum amount of CPU and memory resource allocation per virtual machine.

vCenter Server Secure Separation

With compute resources separated at the sub-resource pool level, vCenter Role Based Access Control (RBAC) needs to be configured to ensure end users have visibility into and control of their own resources.

The following steps are required to configure vCenter RBAC based on the Enhanced Secure Multi-Tenancy Design Guide recommendations (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldg_V2.html):

1. Create Cloud Admin, Tenant Admin, and Tenant User roles.

Note The exact privileges to select for each role can be found in the Enhanced Secure Multi-Tenancy Design Guide.

Figure 23 Cloud Admin Role—Access and Control to All Resources

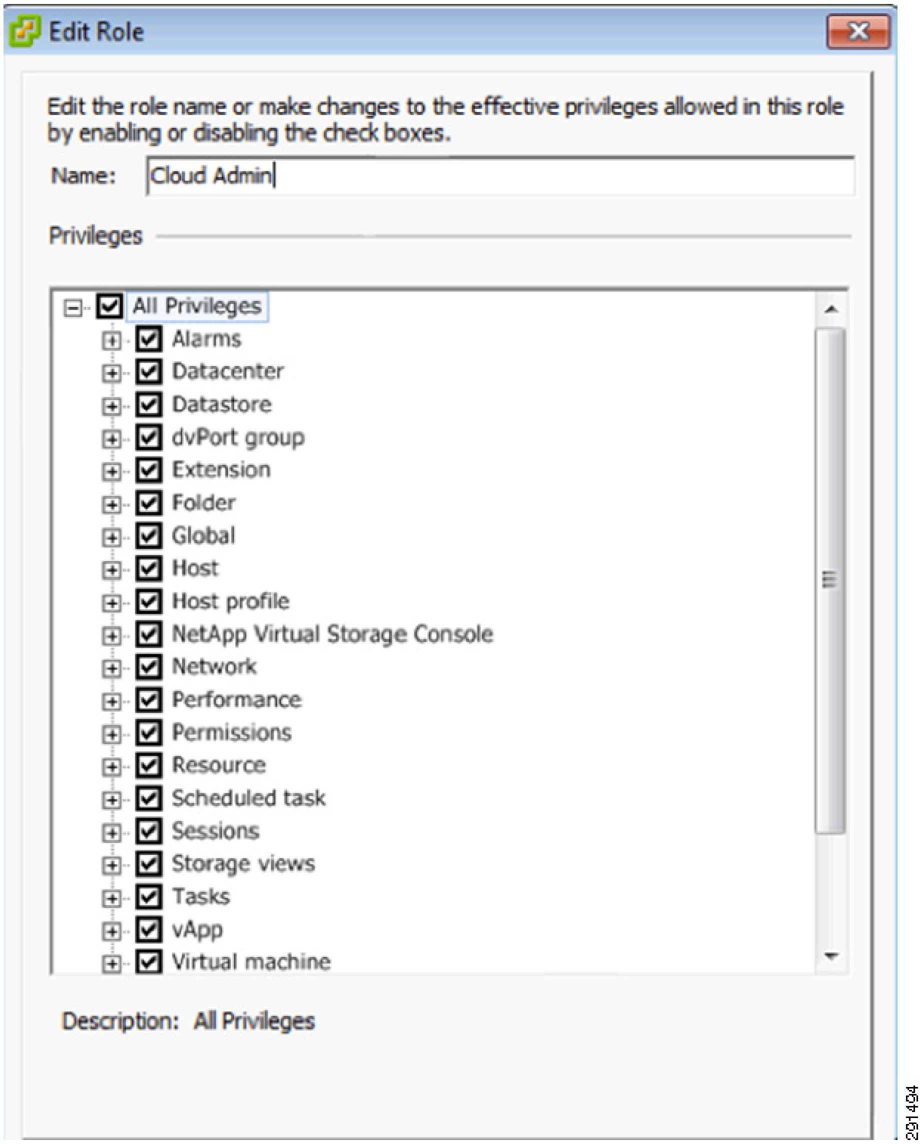

Edit Role

Edit the role name or make changes to the effective privileges allowed in this role by enabling or disabling the check boxes.

Name: Cloud Admin

Privileges

- ☑ All Privileges
  - ☑ Alarms
  - ☑ Datacenter
  - ☑ Datastore
  - ☑ dvPort group
  - ☑ Extension
  - ☑ Folder
  - ☑ Global
  - ☑ Host
  - ☑ Host profile
  - ☑ NetApp Virtual Storage Console
  - ☑ Network
  - ☑ Performance
  - ☑ Permissions
  - ☑ Resource
  - ☑ Scheduled task
  - ☑ Sessions
  - ☑ Storage views
  - ☑ Tasks
  - ☑ vApp
  - ☑ Virtual machine

Description: All Privileges

Tenant Admin has limited set of privileges compared to Cloud Admin and Tenant User has the most restrictive set of privileges (Figure 24 and Figure 25).

Figure 24 Tenant Admin Privileges

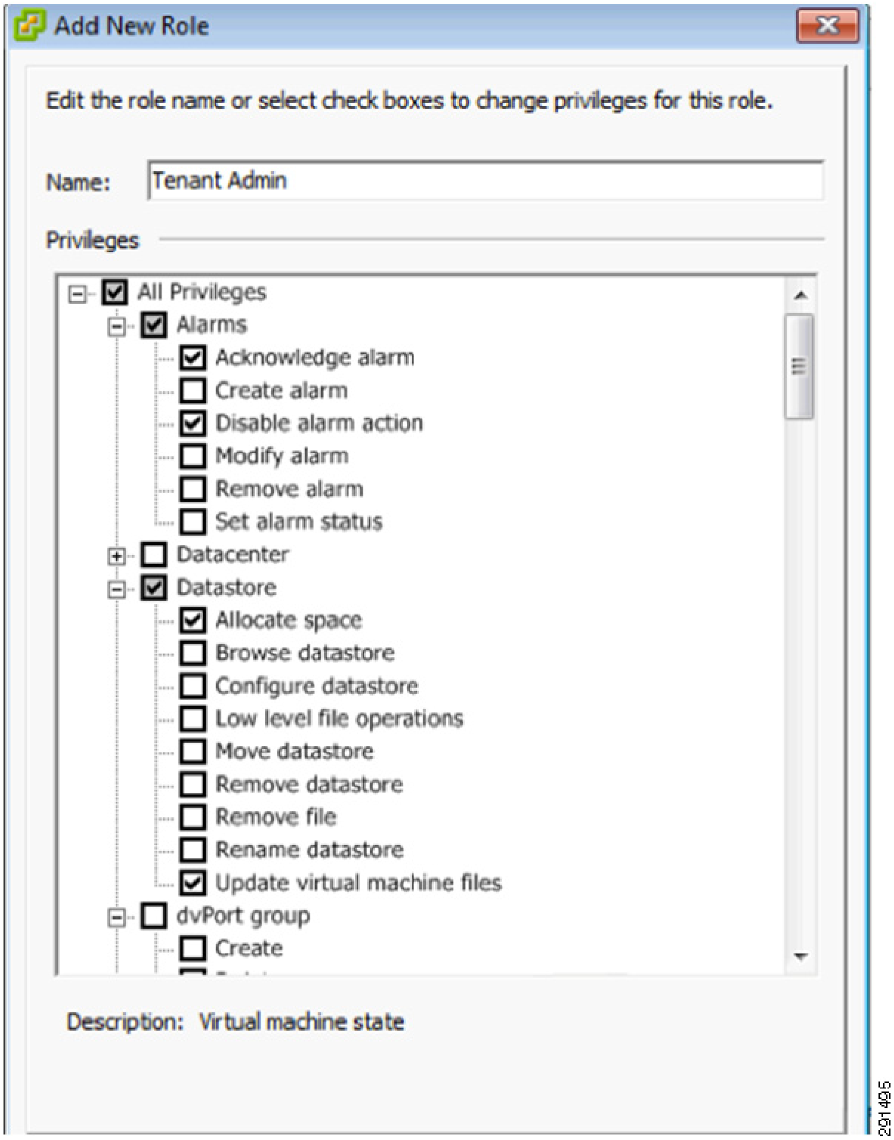

Add New Role

Edit the role name or select check boxes to change privileges for this role.

Name: Tenant Admin

Privileges

- ☑ All Privileges
  - ☑ Alarms
    - ☑ Acknowledge alarm
    - □ Create alarm
    - ☑ Disable alarm action
    - □ Modify alarm
    - □ Remove alarm
    - □ Set alarm status
  - □ Datacenter
  - ☑ Datastore
    - ☑ Allocate space
    - □ Browse datastore
    - □ Configure datastore
    - □ Low level file operations
    - □ Move datastore
    - □ Remove datastore
    - □ Remove file
    - □ Rename datastore
    - ☑ Update virtual machine files
  - □ dvPort group
    - □ Create

Description: Virtual machine state

Figure 25 Tenant User Privileges

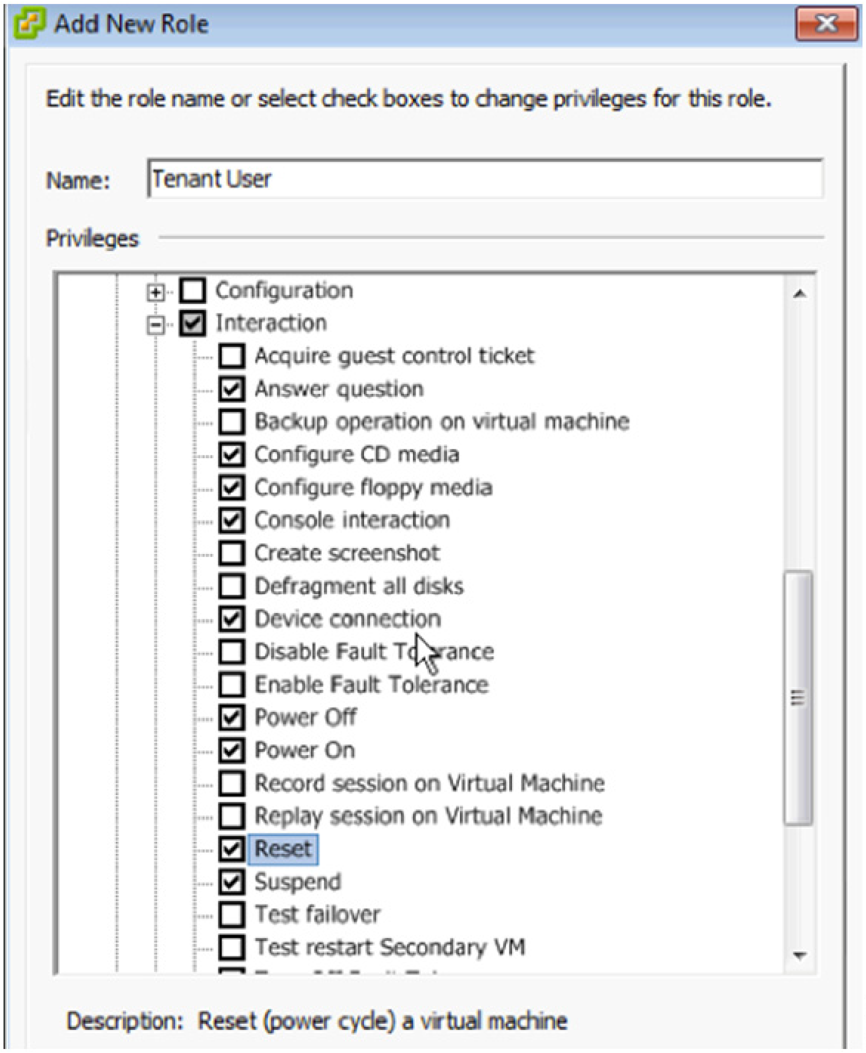

Edit the role name or select check boxes to change privileges for this role.

Name: Tenant User

Privileges

- □ Configuration
- ☑ Interaction
  - □ Acquire guest control ticket
  - ☑ Answer question
  - □ Backup operation on virtual machine
  - ☑ Configure CD media
  - ☑ Configure floppy media
  - ☑ Console interaction
  - □ Create screenshot
  - □ Defragment all disks
  - ☑ Device connection
  - □ Disable Fault Tolerance
  - □ Enable Fault Tolerance
  - ☑ Power Off
  - ☑ Power On
  - □ Record session on Virtual Machine
  - □ Replay session on Virtual Machine
  - ☑ Reset
  - ☑ Suspend
  - □ Test failover
  - □ Test restart Secondary VM

Description: Reset (power cycle) a virtual machine

291496

2. After the roles have been created, associate the role with the respective user for each of the tenant sub-resource pools (for example, Figure 26 and Figure 28).

Figure 26 Infrastructure Resource Pool Permission Assignment for Cloud Admin Role—Part 1

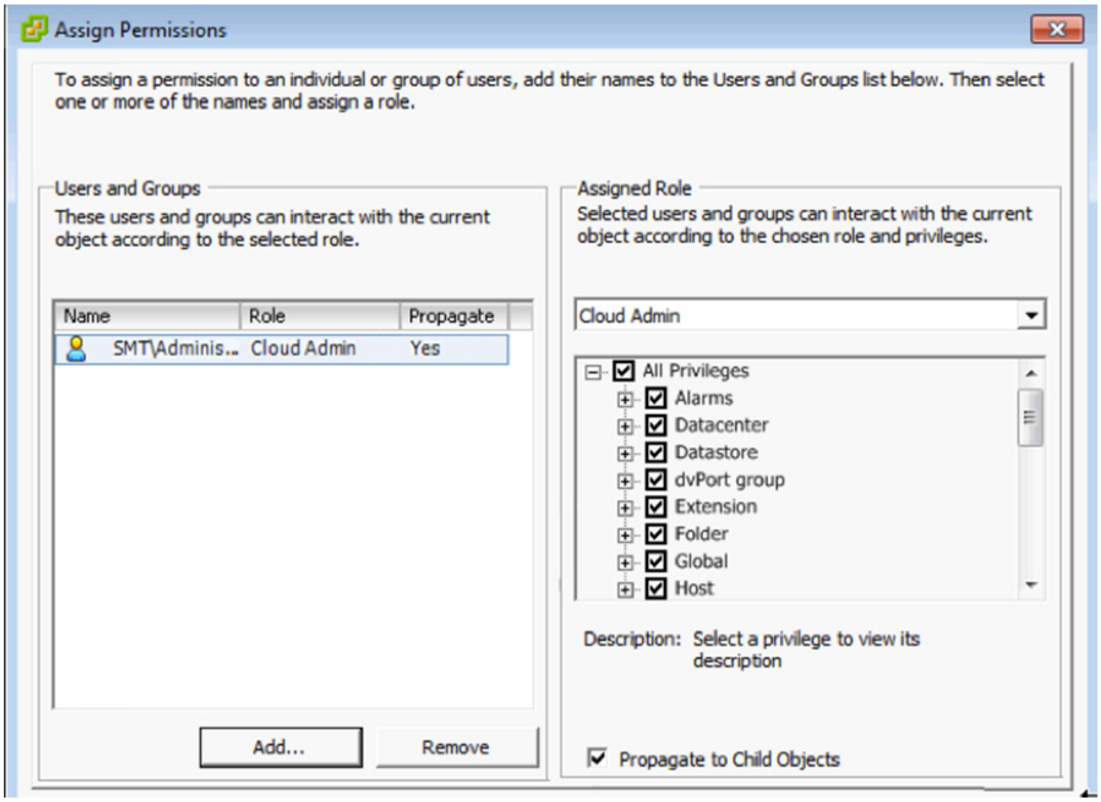

Assign Permissions

To assign a permission to an individual or group of users, add their names to the Users and Groups list below. Then select one or more of the names and assign a role.

Users and Groups

These users and groups can interact with the current object according to the selected role.

Name	Role	Propagate
SMT\Adminis...	Cloud Admin	Yes

Add... Remove

Assigned Role

Selected users and groups can interact with the current object according to the chosen role and privileges.

Cloud Admin

- ☑ All Privileges
  - ☑ Alarms
  - ☑ Datacenter
  - ☑ Datastore
  - ☑ dvPort group
  - ☑ Extension
  - ☑ Folder
  - ☑ Global
  - ☑ Host

Description: Select a privilege to view its description

☑ Propagate to Child Objects

Figure 27 Infrastructure Resource Pool Permission Assignment for Cloud Admin Role—Part 2

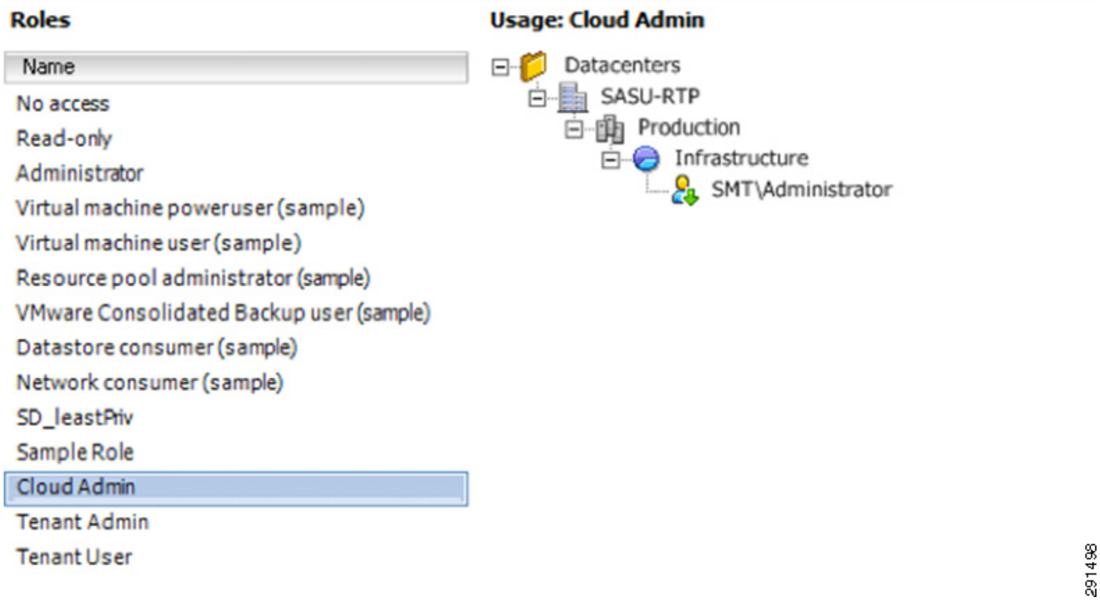

Roles

Name
No access
Read-only
Administrator
Virtual machine poweruser (sample)
Virtual machine user (sample)
Resource pool administrator (sample)
VMware Consolidated Backup user (sample)
Datastore consumer (sample)
Network consumer (sample)
SD_leastPriv
Sample Role
Cloud Admin
Tenant Admin
Tenant User

Usage: Cloud Admin

- Datacenters
  - SASU-RTP
    - Production
      - Infrastructure
        - SMT\Administrator

Figure 28 Tenant Resource Pool Permission Assignment for Tenant Admin Role—Part 1

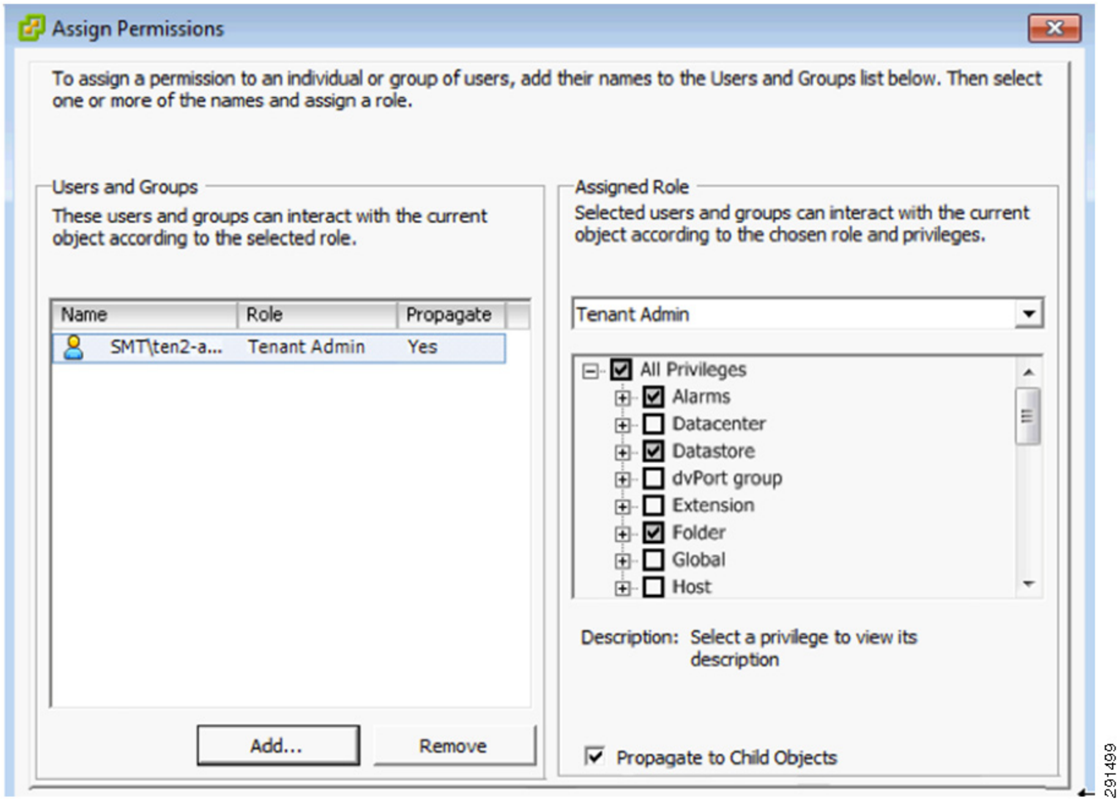

Assign Permissions

To assign a permission to an individual or group of users, add their names to the Users and Groups list below. Then select one or more of the names and assign a role.

Users and Groups

These users and groups can interact with the current object according to the selected role.

Name	Role	Propagate
SMT\ten2-a...	Tenant Admin	Yes

Assigned Role

Selected users and groups can interact with the current object according to the chosen role and privileges.

Tenant Admin

- ☑ All Privileges
  - ☑ Alarms
  - □ Datacenter
  - ☑ Datastore
  - □ dvPort group
  - □ Extension
  - ☑ Folder
  - □ Global
  - □ Host

Description: Select a privilege to view its description

☑ Propagate to Child Objects

Figure 29 Tenant Resource Pool Permission Assignment for Tenant Admin Role—Part 2

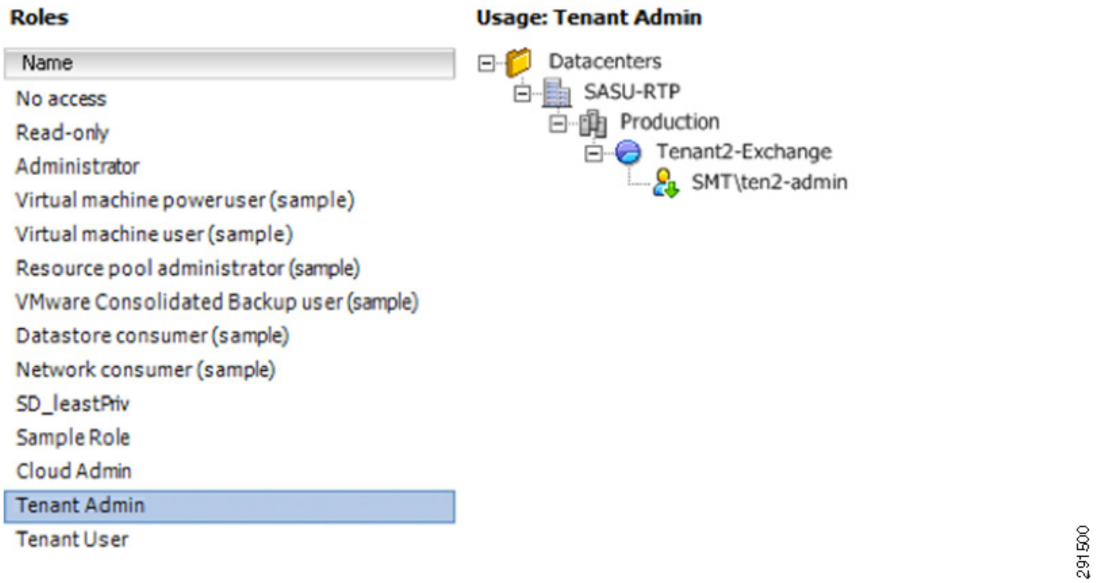

Roles

Name
No access
Read-only
Administrator
Virtual machine poweruser (sample)
Virtual machine user (sample)
Resource pool administrator (sample)
VMware Consolidated Backup user (sample)
Datastore consumer (sample)
Network consumer (sample)
SD_leastPriv
Sample Role
Cloud Admin
Tenant Admin
Tenant User

Usage: Tenant Admin

- Datacenters
  - SASU-RTP
    - Production
      - Tenant2-Exchange
        - SMT\ten2-admin

Follow similar steps for all tenants in the ESMT environment.

Installing and Configuring Additional VMware Components for ESMT

The following additional products are not in the base FlexPod and are required for ESMT:

- vCenter Heartbeat (for installation instructions, see: http://vmware.com/pdf/vcenter-server-heartbeat-63-u1-installation-guide.pdf)
- vShield suite—vShield Manager, vShield Edge, and vShield App (for installation instructions, see: http://vmware.com/pdf/vshield_410U1_quickstart.pdf)

vCenter Server Heartbeat (additional configuration steps with NetApp Virtual Storage Console):

- Option 1—NetApp VSC is installed on the same VM as vCenter Server
- Option 2—NetApp VSC is installed on a separate VM

In the ESMT deployment, option 2 is used. The NetApp VSC is installed on a separate virtual machine. No additional steps are needed for VSC to function in the event of vCenter Server failover (from primary to secondary).

If option 1 were used in the deployment, the following steps are required to ensure VSC is protected and functions correctly on failover of vCenter Server from primary to secondary.

Add Protection for Virtual Storage Console

Note The following steps will work **only** if the NetApp Virtual Storage Console has already been installed on vCenter Server. If not, install the VSC before performing the steps below.

1. Click **vCenter Server Heartbeat\Manage Server** and log into the Management GUI.
2. Click **Advanced Configuration**.
3. Click **Application** > **Services**.
4. Click **Add**.
5. Select the individual service listed, NetApp vSphere Plugin Framework (NVPF).

   Leave all service configurations to the default, “recover service”.
6. Click **Data** > **Configuration**.
7. Click **Add Inclusion Filter**.
8. Navigate to and select <ProgramFiles>\NetApp\Virtual Storage Console.

The vCenter Server Heartbeat now monitors and protects the NetApp Virtual Storage Console services and synchronizes the application data/configuration files so that both servers are able to fully function as vCenter Servers and NetApp Virtual Storage Console Servers. For additional information, refer to KB 1036507 (http://kb.vmware.com/kb/1036507).

vShield Configuration

Refer to Layer 2 Deployment for the steps needed to configure vShield for the protection of tenants.

Securing VMware vSphere ESXi

Security Hardening

- NTP—Configure an NTP server by following the steps below for each ESX Server in the cluster:
  a. In the vSphere Client, select the host in the inventory.
  b. Click the **Configuration** tab and click **Time Configuration**.
  c. Click the **Properties** link at the top right of the panel.
  d. Click **Options**.
  e. Under NTP Settings, click **Add** to specify the NTP servers for the environment.
- Management network—Configure the management network on its own private VLAN, not tied to the infrastructure management, storage access, or tenant VM network.

Figure 30 Management Network

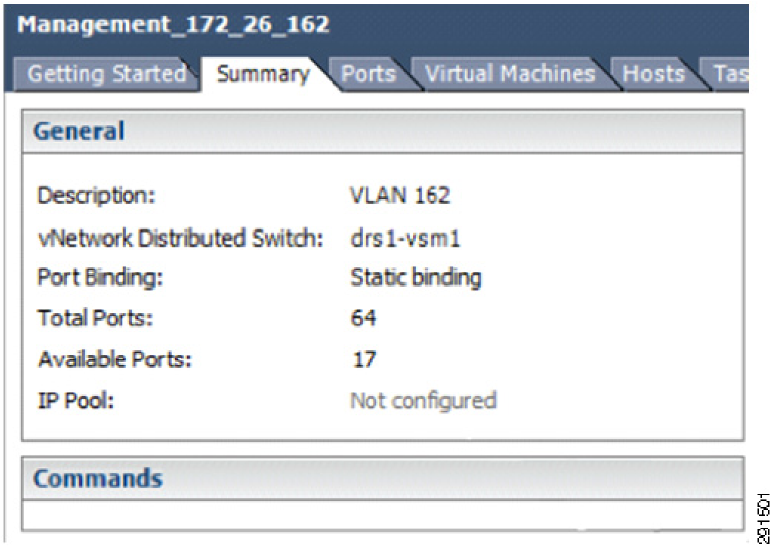

Management_172_26_162

General

Description: VLAN 162
vNetwork Distributed Switch: drs1-vsm1
Port Binding: Static binding
Total Ports: 64
Available Ports: 17
IP Pool: Not configured

Commands

- vMotion network—Configure vMotion network on its own private VLAN, not tied to the infrastructure management, storage access, or tenant VM network.

Figure 31 vMotion Network

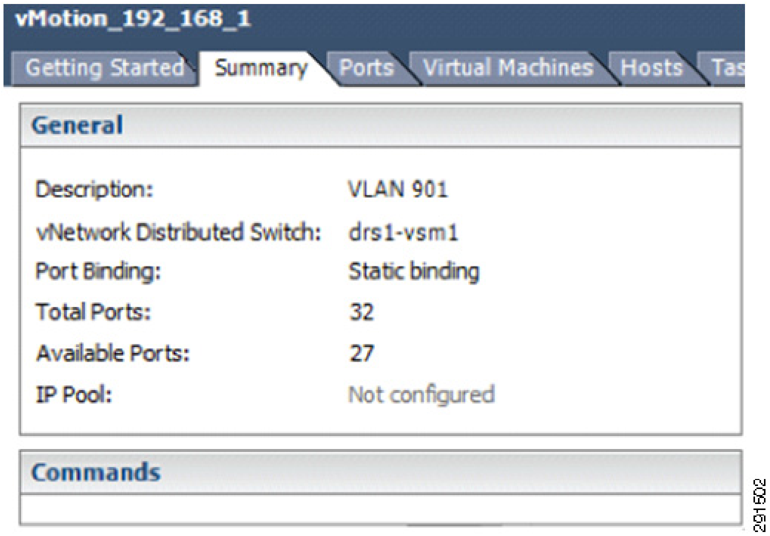

Description: VLAN 901
vNetwork Distributed Switch: drs1-vsm1
Port Binding: Static binding
Total Ports: 32
Available Ports: 27
IP Pool: Not configured

- SSL certificate—Replace default self-signed certificates with those from a trusted certificate authority, either commercial or organizational. Follow the instructions in the ESX Configuration Guide (http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esx_serv er_config.pdf), in the "Security" chapter of the "Authentication and User Management" section.

Centralized Logging

For each ESXi host, follow the steps below to configure remote syslog:

1. From the vSphere Client, select a host in the inventory.
2. Select the Configuration tab and click **Advanced Settings**.
3. Click **Syslog**.
4. For Syslog.Remote.Hostname, enter a host name to which log files can be forwarded.

Centralized Authentication

Given that the tenants are not granted access to ESXi server directly, no AD authentication is needed for the ESXi host.

Securing VMware vSphere vCenter Server

Security Hardening

- vCenter Server default certification—Follow the instructions in "Replacing vCenter Server 4.1 Certificates" (http://www.vmware.com/resources/techresources/10124) to replace the default certificate for vCenter server.
- Web Service Datastore Browser—Disable the datastore browser functionality built-in within vCenter Web service by following instructions below:

  To disable the datastore browser, edit the `vpxd.cfg` file and ensure that the following element is set:

  ```
  <enableHttpDatastoreAccess>false</enableHttpDatastoreAccess>
  ```

This should be within the `<vpxd>... </vpxd>` element in the `vpxd.cfg` file.

Centralized Authentication

vCenter Server is part of the SMT domain, therefore users and groups defined in the central AD server can be used to access vCenter Server.

Securing VMware vShield

Security Hardening

- vShield Manager—The default password must be changed for vShield Manager and vShield agents after the first logon. Follow the steps below to change the password for vShield Manager:

  To change the admin account password:

  a. Log in to the vShield Manager user interface.

  b. Click **Settings & Reports** from the vShield Manager inventory panel.

  c. Click the **Users** tab.

  d. Select the admin account.

  e. Click **Update User**.

  f. Enter a new password.

  g. Confirm the password by typing it a second time in the Retype Password field.

  h. Click **OK** to save your changes.

- NTP—To set the date and time configuration of the vShield Manager:

  a. Click **Settings & Reports** from the vShield Manager inventory panel.

  b. Click the **Configuration** tab.

  c. Click **Date/Time**.

  d. In the Date and Clock field, type the date and time in the format YYYY-MM-DD HH:MM:SS.

  e. In the NTP Server field, type the IP address/hostname of your NTP server.

  f. From the Time Zone drop-down menu, select the appropriate time zone.

  g. Click **Save**.

- SSL certificate—Follow the instruction in the vShield Administration Guide (http://vmware.com/pdf/vshield_410U1_admin.pdf), in the “Add an SSL Certificate to Identify the vShield Manager Web Service” section, to import a pre-existing SSL certificate or generate a new one.

Centralized Logging

Configure remote syslog server for vShield Edge.

Figure 32 Centralized Logging Configuration for vShield Edge

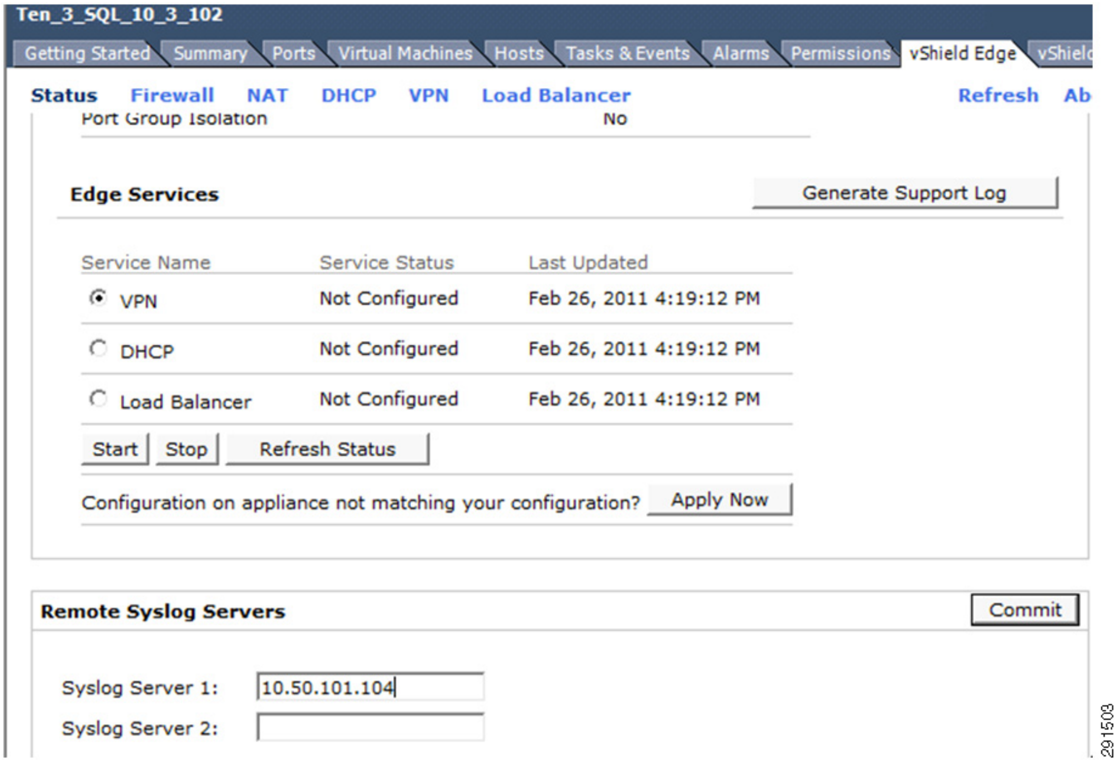

Port Group Isolation No

Edge Services

	Service Name	Service Status	Last Updated
●	VPN	Not Configured	Feb 26, 2011 4:19:12 PM
○	DHCP	Not Configured	Feb 26, 2011 4:19:12 PM
○	Load Balancer	Not Configured	Feb 26, 2011 4:19:12 PM

Configuration on appliance not matching your configuration?

Remote Syslog Servers

Syslog Server 1: 10.50.101.104

Syslog Server 2:

Centralized Authentication

In the ESMT deployment, the Admin account is solely owned and controlled by the Cloud Admin. No access is given to the tenants.

Local Backup and Recovery

Local backup and recovery in an ESMT environment comes in the form of NetApp Snapshot copies. Whether the data being stored is a virtual machine C: drive or application binaries within a vmdk or application-specific data on raw device mappings (RDMs), all are initially backed up locally on the NetApp storage controller using NetApp Snapshot copies. Depending on the data type however, Snapshot copies are captured using different methods. For complete local virtual machine and application backups, several features are used together, including NetApp SnapDrive® and NetApp SnapManager® software.

Virtual Machine OS and Application Binaries

To capture Snapshot copies of virtual machines and application binaries (excluding application data on RDMs), NetApp SnapManager for Virtual Infrastructure (SMVI) is used. SMVI is now included as the Backup and Recovery tab within the NetApp Virtual Storage Console or VSC plug-in for VMware vCenter Server. Snapshot copies are triggered either by a user-configured schedule or manually within the SMVI GUI.

Figure 33 highlights hourly, daily, and weekly backup jobs created in the NetApp Backup and Recovery (SMVI) tab of VMware vCenter Server for the virtual machine named `ten3_SQL` that is stored in the NFS datastore named `tenant_3_SQL_NFS_1`.

Figure 33 Hourly, Daily, and Weekly Backup Jobs in NetApp Backup and Recovery Tab

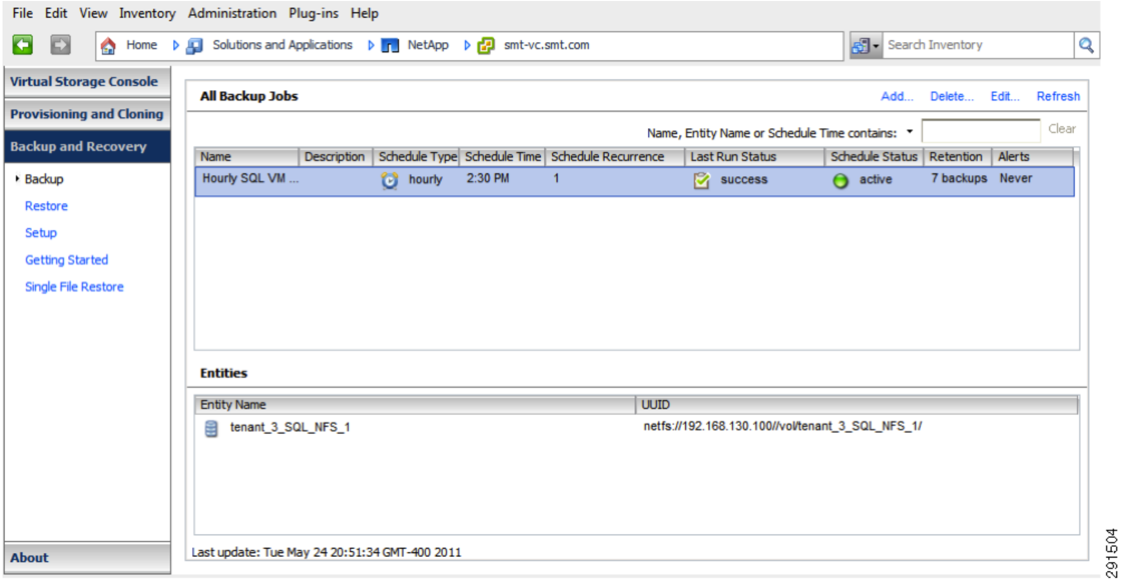

All Backup Jobs

Name, Entity Name or Schedule Time contains:

Name	Description	Schedule Type	Schedule Time	Schedule Recurrence	Last Run Status	Schedule Status	Retention	Alerts
Hourly SQL VM ...		hourly	2:30 PM	1	success	active	7 backups	Never

Entities

Entity Name	UUID
tenant_3_SQL_NFS_1	netfs://192.168.130.100//vol/tenant_3_SQL_NFS_1/

Last update: Tue May 24 20:51:34 GMT-400 2011

For more information on installing and configuring SMVI as included in the NetApp VSC, see:

- NetApp Virtual Storage Console 2.1 for VMware vSphere Installation and Administration Guide: https://now.netapp.com/NOW/knowledge/docs/virtual_storage_console/relvsc21/pdfs/install.pdf
- NetApp Virtual Storage Console 2.1 for VMware vSphere Backup and Recovery Administration Guide: https://now.netapp.com/NOW/knowledge/docs/virtual_storage_console/relvsc21/pdfs/backup.pdf

Application Specific Data

To capture Snapshot copies of application-specific data stored on VMware RDMs, the application-specific SnapManager software (for example, SnapManager for MS SQL) is used in conjunction with NetApp SnapDrive. NetApp SnapDrive is installed directly on the application server and allows for the creation, mapping, and management of block-level storage objects as they pertain to that instance of SnapDrive. SnapDrive communicates directly with both VMware vCenter Server and the NetApp storage controller (physical or virtual) to perform these functions as well as initiate Snapshot copies. NetApp SnapManager software is written by design to be application-specific in order to enable application consistent backups using NetApp Snapshot copies. NetApp SnapManager software collaborates with the application itself to prepare data in the file system for Snapshot capture. Once the application is prepared, NetApp SnapManager communicates with NetApp SnapDrive and initiates a Snapshot capture.

Use the Create Disk action within NetApp SnapDrive to add new RDM logical unit numbers (LUNs) to the virtual machine. NetApp SnapDrive follows this workflow to create new iSCSI attached RDM disks in an ESMT environment:

1. Creates LUN(s) on storage controller.
2. Creates igroup(s) on storage controller.

Note SnapDrive by default configures only the igroup with the iSCSI iqn name for the host on which the virtual machine currently resides. If an ESX(i) cluster exists, manually add the additional iSCSI iqn names for the other ESX(i) servers in the cluster.

3. Maps the igroup to the LUN(s) on the storage controller.
4. Initiates a storage rescan on the ESX(i) hosts in VMware vCenter Server.
5. Adds a small .vmdk file to the virtual machine that maps directly to the iSCSI RDM disk or LUN seen by the ESX(i) hosts.

Figure 34 and Figure 35 highlight the iSCSI LUNs that have been created and mapped to the TEN3-SQL virtual machine using NetApp SnapDrive.

Figure 34 LUNs Mapped to Virtual Machine

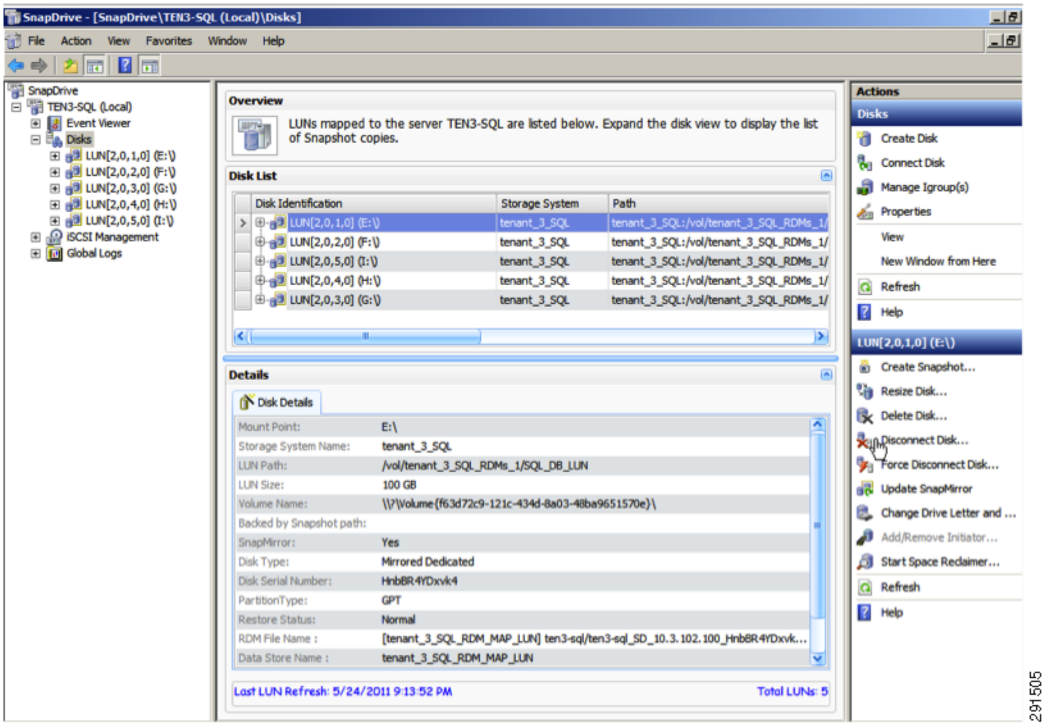

Overview

LUNs mapped to the server TEN3-SQL are listed below. Expand the disk view to display the list of Snapshot copies.

Disk List

Disk Identification	Storage System	Path
LUN[2,0,1,0] (E:\)	tenant_3_SQL	tenant_3_SQL:/vol/tenant_3_SQL_RDMs_1/
LUN[2,0,2,0] (F:\)	tenant_3_SQL	tenant_3_SQL:/vol/tenant_3_SQL_RDMs_1/
LUN[2,0,5,0] (I:\)	tenant_3_SQL	tenant_3_SQL:/vol/tenant_3_SQL_RDMs_1/
LUN[2,0,4,0] (H:\)	tenant_3_SQL	tenant_3_SQL:/vol/tenant_3_SQL_RDMs_1/
LUN[2,0,3,0] (G:\)	tenant_3_SQL	tenant_3_SQL:/vol/tenant_3_SQL_RDMs_1/

Details

Disk Details

Mount Point:	E:\
Storage System Name:	tenant_3_SQL
LUN Path:	/vol/tenant_3_SQL_RDMs_1/SQL_DB_LUN
LUN Size:	100 GB
Volume Name:	\\?\Volume{f63d72c9-121c-434d-8a03-48ba9651570e}\
Backed by Snapshot path:	
SnapMirror:	Yes
Disk Type:	Mirrored Dedicated
Disk Serial Number:	HnbBR4YDxvk4
PartitionType:	GPT
Restore Status:	Normal
RDM File Name :	[tenant_3_SQL_RDM_MAP_LUN] ten3-sql/ten3-sql_SD_10.3.102.100_HnbBR4YDxvk...
Data Store Name :	tenant_3_SQL_RDM_MAP_LUN

Last LUN Refresh: 5/24/2011 9:13:52 PM

Total LUNs: 5

Figure 35 Tenant 3 Storage View in Control Panel

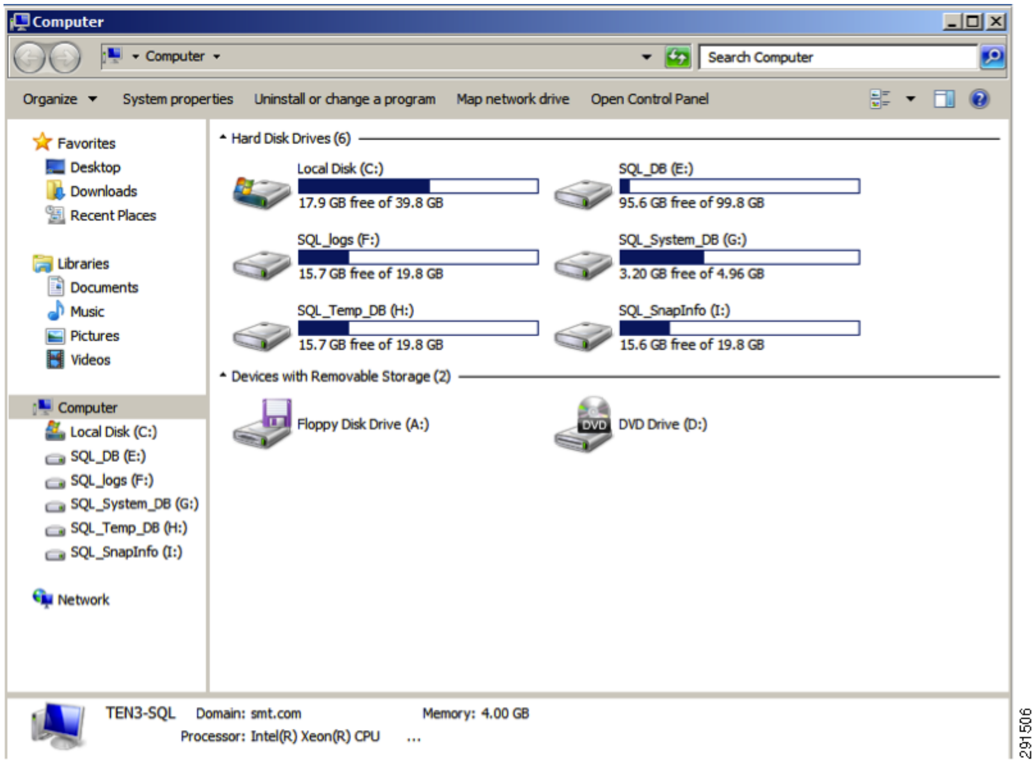

Hard Disk Drives (6)

Local Disk (C:)
17.9 GB free of 39.8 GB

SQL_DB (E:)
95.6 GB free of 99.8 GB

SQL_logs (F:)
15.7 GB free of 19.8 GB

SQL_System_DB (G:)
3.20 GB free of 4.96 GB

SQL_Temp_DB (H:)
15.7 GB free of 19.8 GB

SQL_SnapInfo (I:)
15.6 GB free of 19.8 GB

Devices with Removable Storage (2)

Floppy Disk Drive (A:)

DVD Drive (D:)

TEN3-SQL Domain: smt.com Memory: 4.00 GB
Processor: Intel(R) Xeon(R) CPU ...

For more information on installing and configuring NetApp SnapDrive, see the SnapDrive 6.3 for Windows; Installation and Administration Guide (https://now.netapp.com/NOW/knowledge/docs/snapdrive/relsnap631/pdfs/admin.pdf).

Once new RDMs have been added to the virtual machine leveraging SnapDrive, the appropriate NetApp SnapManager should be installed depending on the application. In this example, NetApp SnapManager for SQL is used. Once installed, SnapManager for SQL allows the migration of existing databases to the newly created and attached disks as well as the ability to create backup jobs within MS SQL for those databases. Once created, backup jobs are exported and available to the Microsoft SQL Server Management Studio under the SQL Server Agent Jobs folder. Figure 36 shows Tenant 3 backup options.

Figure 36 Tenant 3 Backup Options

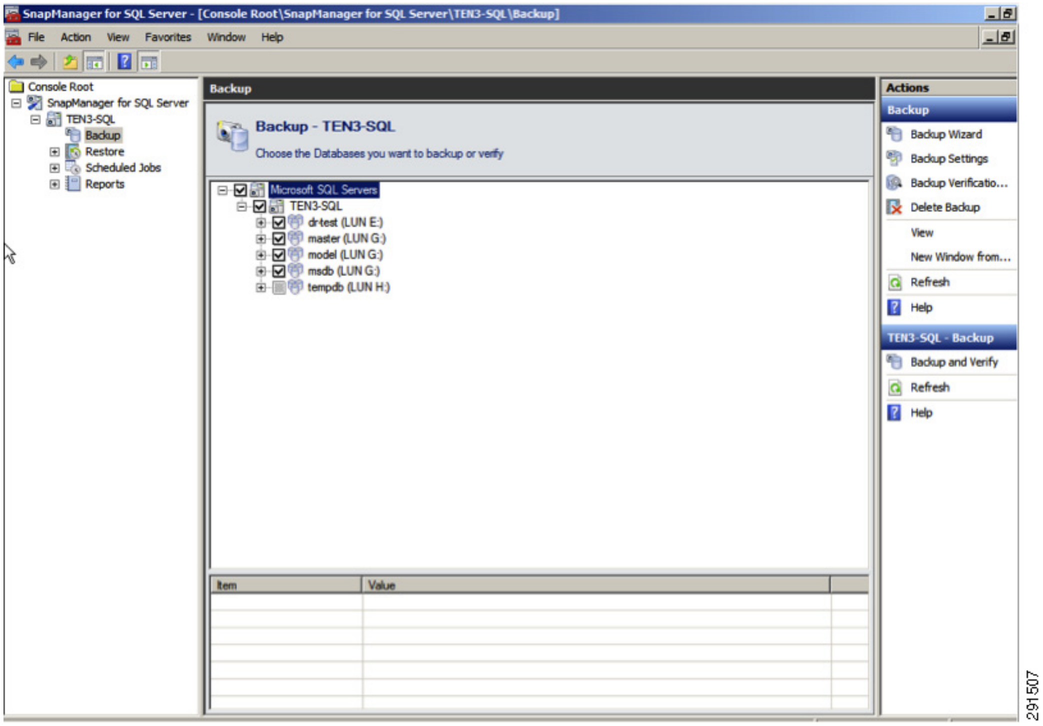

In Figure 37, the backup jobs named `Full_Backup` and `Transaction_Log_Backup` have been created using NetApp SnapManager for SQL and automatically imported into the SQL Server Agent Jobs folder.

Figure 37 Backup Jobs Imported into the Jobs Folder

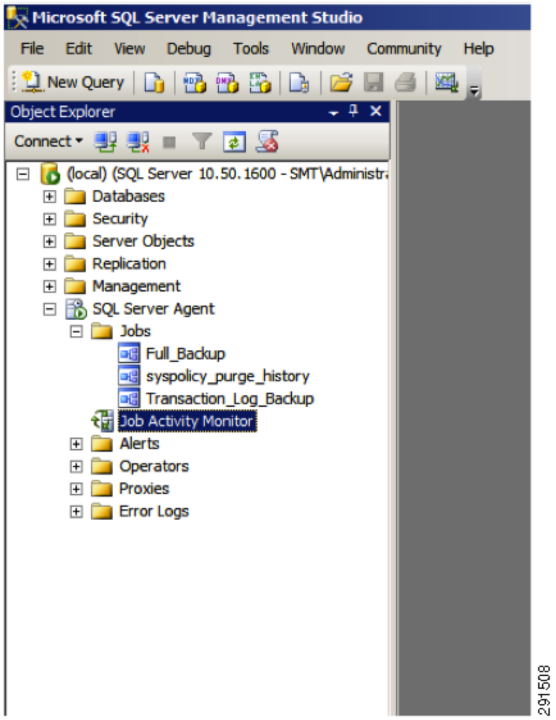

For more information on installing and configuring NetApp SnapManager products, see the application-appropriate SnapManager collateral on the NetApp Support (formerly NOW®) site (http://now.netapp.com/).

Disaster Recovery

SnapMirror Relationships

Once the appropriate data has been captured in NetApp Snapshots, the Snapshot copies can be replicated locally or to another data center for disaster recovery (DR) purposes. NetApp SnapMirror is a feature built into NetApp Data ONTAP that allows for the efficient replication or mirroring of data across two storage controllers. A baseline transfer is initially required to create a consistent mirror, but once the baseline is complete only data that has changed is transferred from that point forward.

The same software tools that were used to capture local Snapshot copies can also be used to initiate a SnapMirror update for disaster recovery purposes. When creating the local backup jobs using the Backup Wizards within both the NetApp VSC Backup and Recovery tab and NetApp SnapManagers, look for the appropriate checkbox to enable SnapMirror updates as an additional process to the backup workflow.

Note SnapMirror relationships between the source and destination must be manually configured. The SnapMirror update process does not create the SnapMirror relationship.

Note NetApp MultiStore is used in the ESMT environment; therefore, in most cases the data being replicated lives in a vFiler unit on NetApp storage. As such, the source and destination of the SnapMirror relationship should both be vFiler units.

Figure 38 and Figure 39 show the checkboxes to enable SnapMirror update, first within the NetApp VSC Backup and Recovery Backup Wizard and second within the NetApp SnapManager for SQL Backup Wizard.

Figure 38 Enabling SnapMirror Update in NetApp VSC Backup and Recovery Backup Wizard

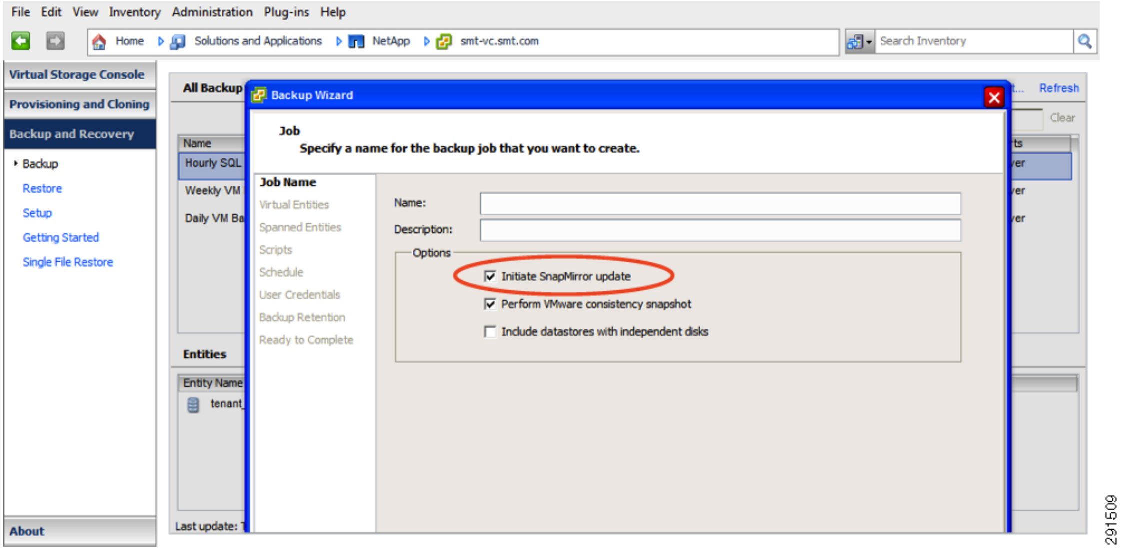

Job

Specify a name for the backup job that you want to create.

Name:

Description:

Options

☑ Initiate SnapMirror update

☑ Perform VMware consistency snapshot

☐ Include datastores with independent disks

Figure 39 Enabling SnapMirror Update in NetApp SnapManaer for SQL Backup Wizard

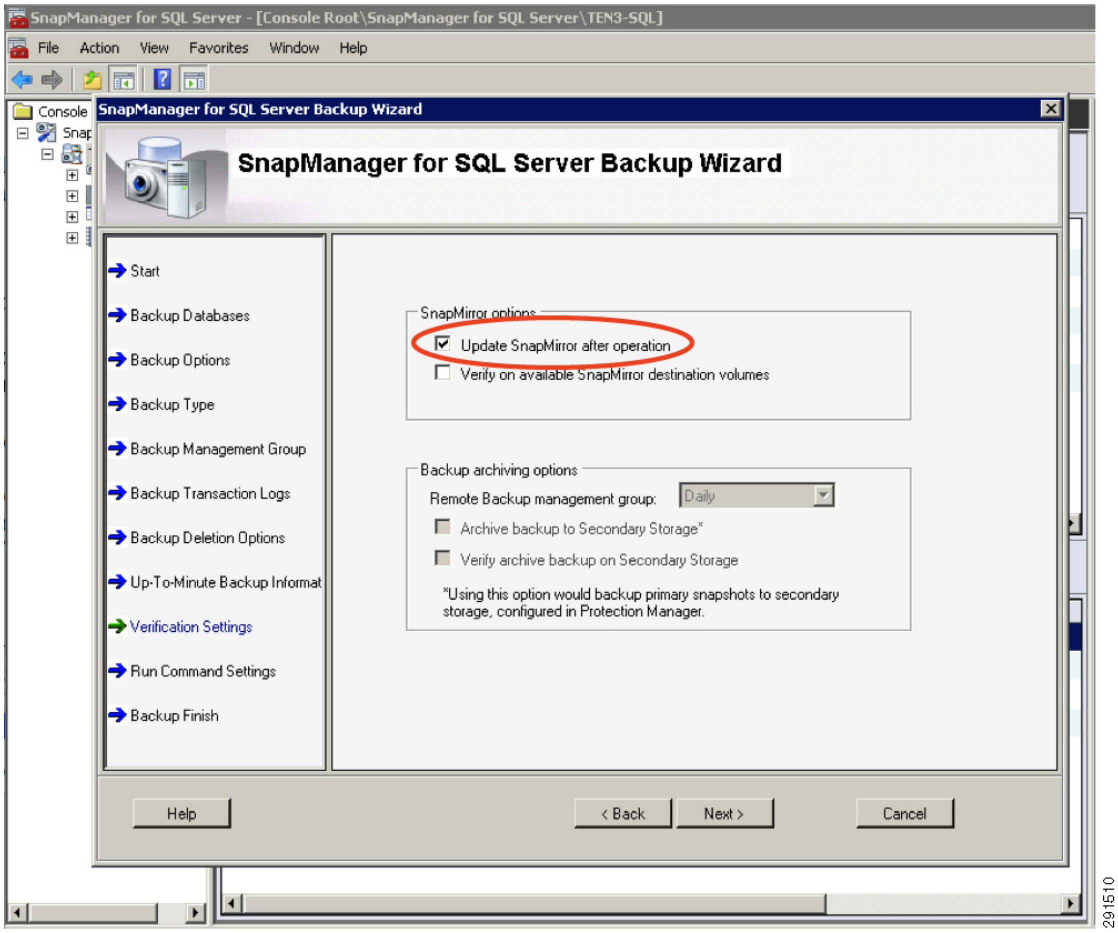

SnapMirror options

☑ Update SnapMirror after operation

□ Verify on available SnapMirror destination volumes

Backup archiving options

Remote Backup management group: Daily

□ Archive backup to Secondary Storage*

□ Verify archive backup on Secondary Storage

*Using this option would backup primary snapshots to secondary storage, configured in Protection Manager.

VMware Site Recovery Manager (SRM)

VMware SRM works in conjunction with NetApp SnapMirror to provide an automated disaster recovery workflow if an unforeseen event should occur. VMware SRM server is installed at both the primary and the DR sites and communicates with the respective VMWare vCenter Servers at each location as well as the SnapMirror source and destination vFiler units. With the use of the NetApp Adapter for SRM, VMware SRM discovers existing SnapMirror relationships, the corresponding storage objects, and the virtual machines that are associated with these objects. These virtual machines are then added to protection groups within VMware SRM and protection workflows and priorities can be configured on an individual protection group basis depending on importance, etc. Once the protection groups are in place, failovers can be initiated from within the SRM portion of vCenter Server, which begins the configured disaster recovery plan(s). Test functionality is also built into SRM and the NetApp Adapter for SRM. This capability allows the administrator to test a failover scenario to ensure proper configuration without effecting the production VMs, applications, etc. Figure 40 provides an overview of disaster recovery.

Figure 40 Disaster Recovery Overview

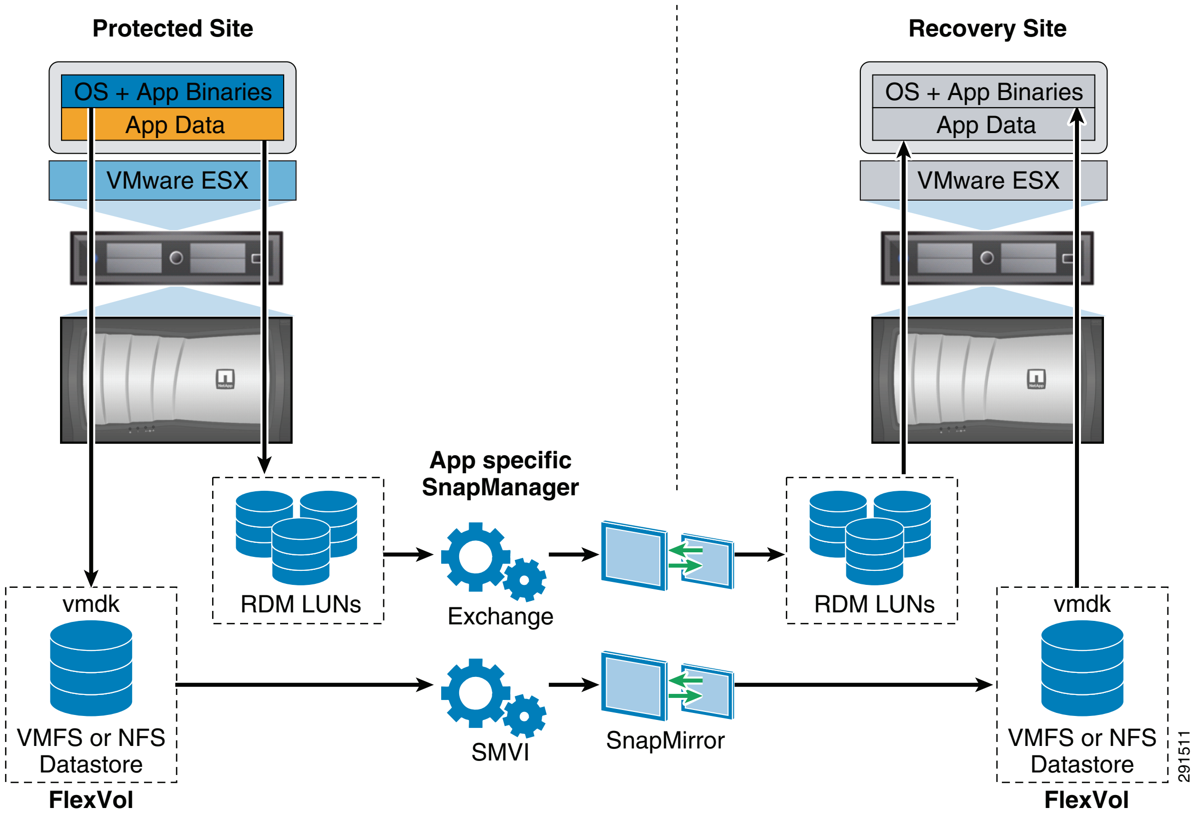

The following documents provide more details on the installation and configuration of both NetApp SnapMirror and VMware Site Recovery manager in a DR scenario:

- TR-3822: Disaster Recovery of Microsoft Exchange, SQL Server, and SharePoint Server Using VMware vCenter Site Recovery Manager, NetApp SnapManager and SnapMirror, and Cisco Nexus Unified Fabric: http://media.netapp.com/documents/tr-3822.pdf
- TR-3671: Vmware vCenter Site Recovery Manager in a NetApp Environment: http://media.netapp.com/documents/tr-3671.pdf
- VMware vCenter Site Recovery Manager 4.0 Performance and Best Practices for Performance: http://www.vmware.com/files/pdf/VMware-vCenter-SRM-WP-EN.pdf

System Monitoring

System Monitoring Capabilities

Deploying a Centralized Syslog Server Monitoring

This architecture employs a hierarchical syslog monitoring system. Each tenant can be configured to have its own syslog server that only monitors events within a tenant. In turn each of the tenant's syslog servers forwards events to a centralized syslog repository, as shown in Figure 41.

Figure 41 Syslog Implementation within a Tenant

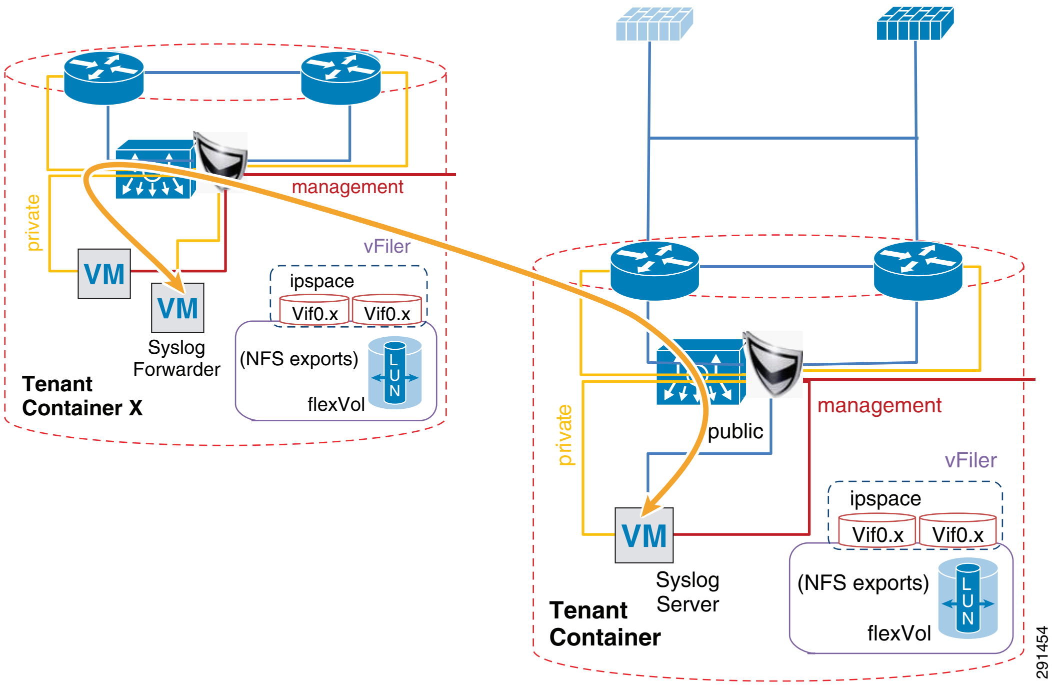

To deploy syslog servers:

1. Use the same NAT functionality of vShield that was used to facilitate application-tenant to infrastructure-tenant communication.
2. Create a static route on the tenant-syslog server to point to the tenant's IP address used by vShield as a internal NAT address.

Deploying Centralized Authentication and Access Monitoring

Active Directory and Cisco Access Control System (ACS) are the two major components that provide authentication capability within this architecture, as shown in Figure 42.

Figure 42 Implementing ACS within ESMT

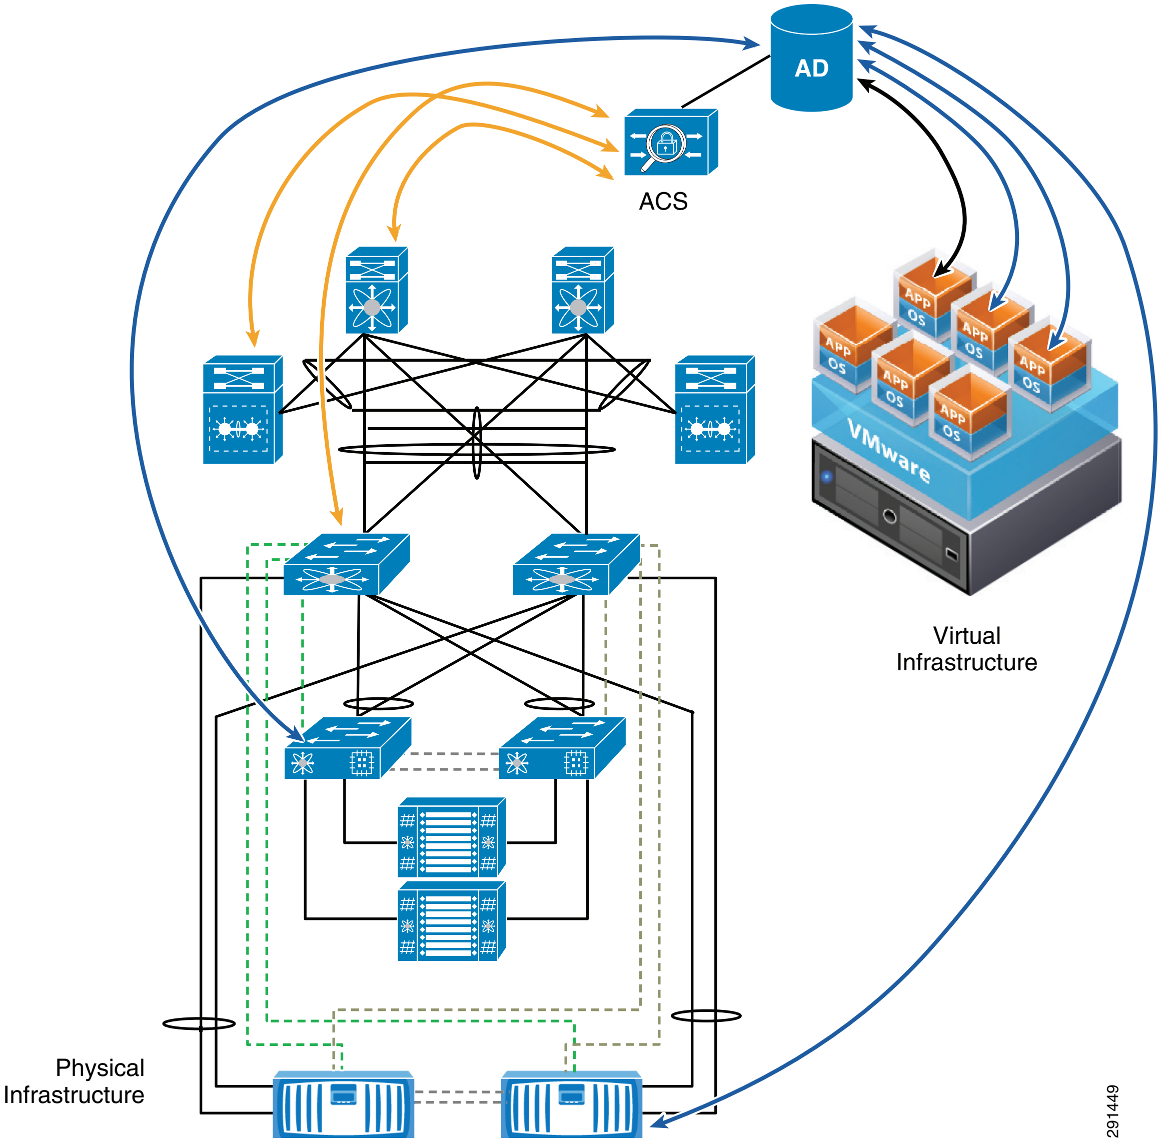

Centralized authentication provides a unified approach that can reduce complexity and security concerns and provide operational efficiency. The following deployment framework was adapted in this architecture:

- All active directory servers reside in the separate Infrastructurre tenant.
- All tenants can access the active directory by creating a NAT rule (by using vShield edge) between the tenant and Infrastructurre tenant.
- Cloud admin creates the global authentication policy, users, and groups.
- Active Directory contains data.
- All devices within the environment integrate with AD through LDAP or TACACS+ through Cisco ACS. The configuration and deployment for each device is outlined in the respective product deployment sections below.

Table 4 outlines some of the user grouping and policy mapping between different categories of users and its associated assigned privileges.

Table 4 User Grouping and Policy Mapping

Group Name	User(s)	Privileges
Cloud Admins	Cloud Admin	Full Access and privileges
	Infrastructure Admin	Rights to physical infrastructure
	Cisco UCS Admin	Rights to Cisco UCS fabric
	Server Virtualization Admin	Rights to VMware vSphere
	Network Admin	Rights to Cisco Nexus/VSS (including Cisco Nexus 1000V)
	Storage Admin	Rights to NetApp controllers
Tenant X Admins	Admin	Full rights to Tenant X container including VM(s) and VSC
	Server Admin	Rights to Tenant X VM(s)
	Application Admin	Rights to Tenant X applications
	Services User	Rights to Tenant X services such as SnapDrive
Tenant X User(s)	Tenant “X” User	Rights to specific applications associated with the user

Figure 43 shows the system-wide AD and hierarchy structure within the ESMT architecture.

Figure 43 Active Directory and Hierarchy Structure within ESMT

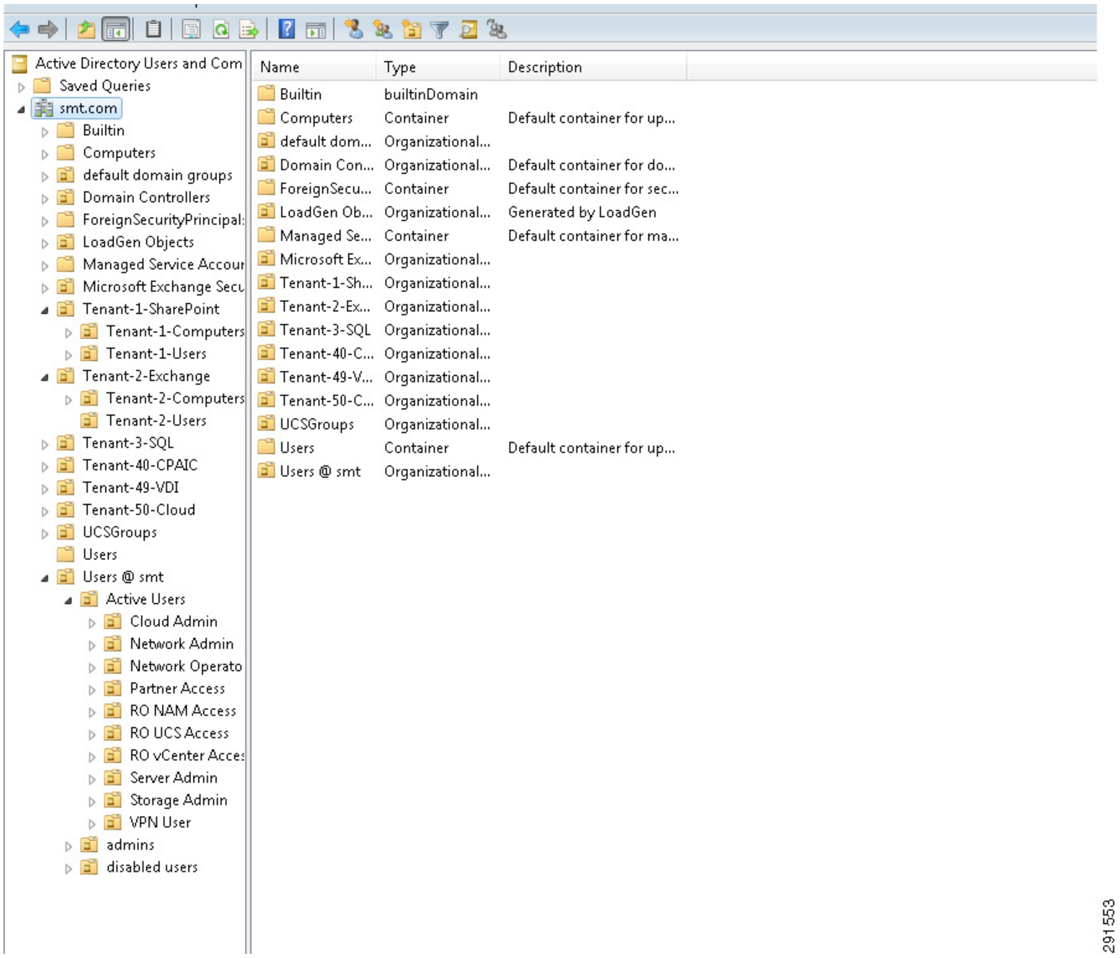

- Active Directory Users and Com
  - Saved Queries
  - smt.com
    - Builtin
    - Computers
    - default domain groups
    - Domain Controllers
    - ForeignSecurityPrincipal:
    - LoadGen Objects
    - Managed Service Accour
    - Microsoft Exchange Secu
    - Tenant-1-SharePoint
      - Tenant-1-Computers
      - Tenant-1-Users
    - Tenant-2-Exchange
      - Tenant-2-Computers
      - Tenant-2-Users
    - Tenant-3-SQL
    - Tenant-40-CPAIC
    - Tenant-49-VDI
    - Tenant-50-Cloud
    - UCSGroups
    - Users
    - Users @ smt
      - Active Users
        - Cloud Admin
        - Network Admin
        - Network Operato
        - Partner Access
        - RO NAM Access
        - RO UCS Access
        - RO vCenter Acce:
        - Server Admin
        - Storage Admin
        - VPN User
      - admins
      - disabled users

Name	Type	Description
Builtin	builtinDomain	
Computers	Container	Default container for up...
default dom...	Organizational...	
Domain Con...	Organizational...	Default container for do...
ForeignSecu...	Container	Default container for sec...
LoadGen Ob...	Organizational...	Generated by LoadGen
Managed Se...	Container	Default container for ma...
Microsoft Ex...	Organizational...	
Tenant-1-Sh...	Organizational...	
Tenant-2-Ex...	Organizational...	
Tenant-3-SQL	Organizational...	
Tenant-40-C...	Organizational...	
Tenant-49-V...	Organizational...	
Tenant-50-C...	Organizational...	
UCSGroups	Organizational...	
Users	Container	Default container for up...
Users @ smt	Organizational...	

Deploying the Cisco Access Control System

In this deployment Active Directory is used to define the different groups of users, while ACS is used to give each group a scope of permissions within the cloud. All the devices that will be using ACS must be entered as a network device. Like devices are grouped together for ease of administration.

Figure 44 *ACS Network Device Entry*

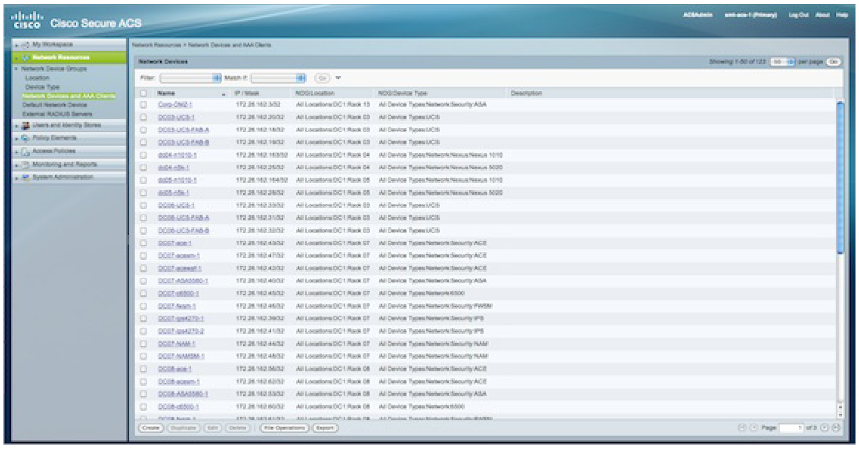

Network Resources > Network Devices and AAA Clients

Network Devices

Showing 1-50 of 123 50 per page Go

Filter: Match if: Go

Name	IP / Mask	NDG:Location	NDG:Device Type	Description
Corp-DN2-1	172.26.162.3/32	All Locations:DC1:Rack 13	All Device Types:Network:Security:ASA	
DC03-UCS-1	172.26.162.20/32	All Locations:DC1:Rack 03	All Device Types:UCS	
DC03-UCS-FAB-A	172.26.162.18/32	All Locations:DC1:Rack 03	All Device Types:UCS	
DC03-UCS-FAB-B	172.26.162.19/32	All Locations:DC1:Rack 03	All Device Types:UCS	
dc04-n1010-1	172.26.162.163/32	All Locations:DC1:Rack 04	All Device Types:Network:Nexus:Nexus 1010	
dc04-n5k-1	172.26.162.25/32	All Locations:DC1:Rack 04	All Device Types:Network:Nexus:Nexus 5020	
dc05-n1010-1	172.26.162.164/32	All Locations:DC1:Rack 05	All Device Types:Network:Nexus:Nexus 1010	
dc05-n5k-1	172.26.162.28/32	All Locations:DC1:Rack 05	All Device Types:Network:Nexus:Nexus 5020	
DC06-UCS-1	172.26.162.33/32	All Locations:DC1:Rack 03	All Device Types:UCS	
DC06-UCS-FAB-A	172.26.162.31/32	All Locations:DC1:Rack 03	All Device Types:UCS	
DC06-UCS-FAB-B	172.26.162.32/32	All Locations:DC1:Rack 03	All Device Types:UCS	
DC07-ace-1	172.26.162.43/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:ACE	
DC07-acesm-1	172.26.162.47/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:ACE	
DC07-acesall-1	172.26.162.42/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:ACE	
DC07-ASA5580-1	172.26.162.40/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:ASA	
DC07-c6500-1	172.26.162.45/32	All Locations:DC1:Rack 07	All Device Types:Network:6500	
DC07-fwsm-1	172.26.162.46/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:FWSM	
DC07-ips4270-1	172.26.162.39/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:IPS	
DC07-ips4270-2	172.26.162.41/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:IPS	
DC07-NAM-1	172.26.162.44/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:NAM	
DC07-NAMSM-1	172.26.162.48/32	All Locations:DC1:Rack 07	All Device Types:Network:Security:NAM	
DC08-ace-1	172.26.162.56/32	All Locations:DC1:Rack 08	All Device Types:Network:Security:ACE	
DC08-acesm-1	172.26.162.62/32	All Locations:DC1:Rack 08	All Device Types:Network:Security:ACE	
DC08-ASA5580-1	172.26.162.53/32	All Locations:DC1:Rack 08	All Device Types:Network:Security:ASA	
DC08-c6500-1	172.26.162.60/32	All Locations:DC1:Rack 08	All Device Types:Network:6500	
DC08-fwsm-1	172.26.162.61/32	All Locations:DC1:Rack 08	All Device Types:Network:Security:FWSM	

The AD instance is used as an identity store in ACS. Once linked, ACS uses the AD user list to determine if the user is valid.

Figure 45 *ACS Active Directory Identity Store Entry*

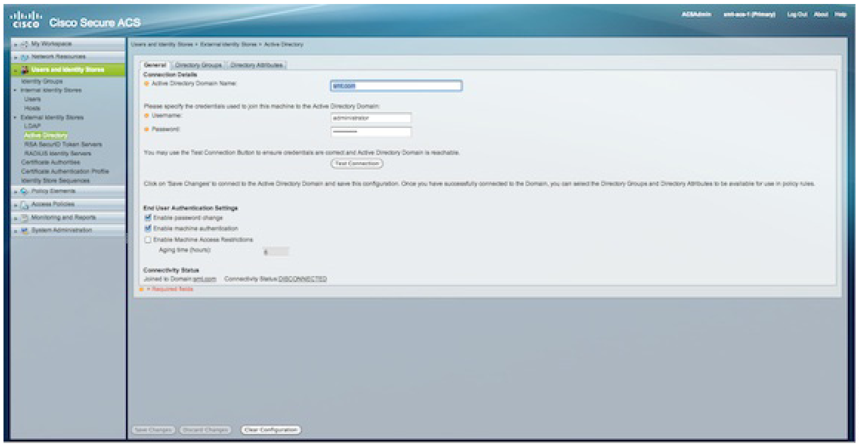

Users and Identity Stores > External Identity Stores > Active Directory

Connection Details

Active Directory Domain Name: smt.com

Please specify the credentials used to join this machine to the Active Directory Domain:

Username: administrator

Password: ••••••••••

You may use the Test Connection Button to ensure credentials are correct and Active Directory Domain is reachable.

Click on 'Save Changes' to connect to the Active Directory Domain and save this configuration. Once you have successfully connected to the Domain, you can select the Directory Groups and Directory Attributes to be available for use in policy rules.

End User Authentication Settings

☑ Enable password change

☑ Enable machine authentication

☐ Enable Machine Access Restrictions

Aging time (hours): 6

Connectivity Status

Joined to Domain: smt.com Connectivity Status: DISCONNECTED

= Required fields

Use Shell Profiles in ACS to define the different scopes of permissions on the different devices (for example, network operator access on Cisco Nexus devices, privilege level 15 on IOS boxes, and so on).

Figure 46 Shell Profile Definition

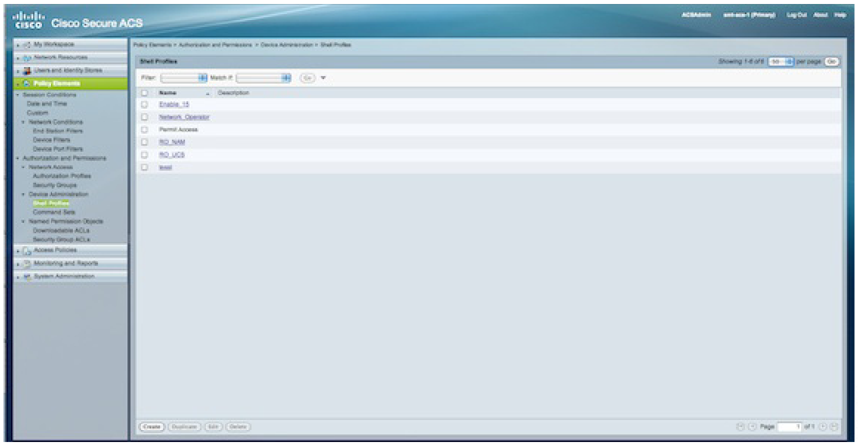

Policy Elements > Authorization and Permissions > Device Administration > Shell Profiles

Shell Profiles

Name	Description
Enable_15	
Network_Operator	
Permit Access	
RO_NAM	
RO_UCS	
test	

Once these elements are configured, they can be used to build access policies in ACS.

Figure 47 Access Policy for Default Network Administration

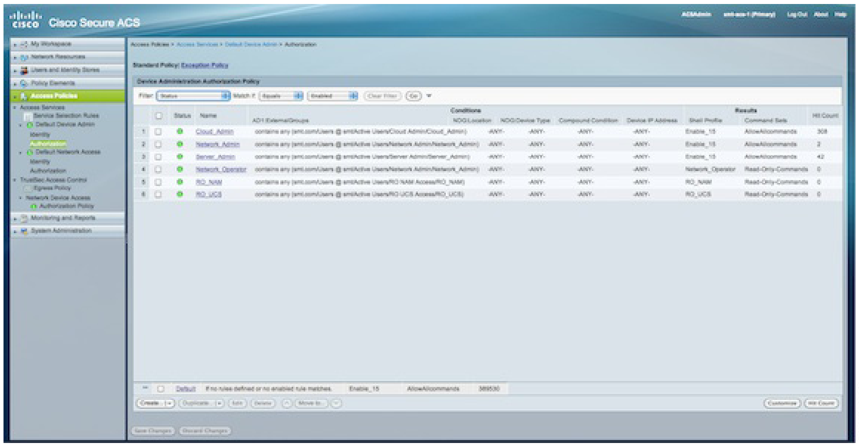

Access Policies > Access Services > Default Device Admin > Authorization

Standard Policy: Exception Policy

Device Administration Authorization Policy

Filter: Status Match if: Equals Enabled

	Status	Name	Conditions AD1:ExternalGroups	NDG:Location	NDG:Device Type	Compound Condition	Device IP Address	Results Shell Profile	Command Sets	Hit Count
1	●	Cloud_Admin	contains any (aml.com/Users @ amlActive Users/Cloud Admin/Cloud_Admin)	-ANY-	-ANY-	-ANY-	-ANY-	Enable_15	AllowAllcommands	308
2	●	Network_Admin	contains any (aml.com/Users @ amlActive Users/Network Admin/Network_Admin)	-ANY-	-ANY-	-ANY-	-ANY-	Enable_15	AllowAllcommands	2
3	●	Server_Admin	contains any (aml.com/Users @ amlActive Users/Server Admin/Server_Admin)	-ANY-	-ANY-	-ANY-	-ANY-	Enable_15	AllowAllcommands	42
4	●	Network_Operator	contains any (aml.com/Users @ amlActive Users/Network Admin/Network_Admin)	-ANY-	-ANY-	-ANY-	-ANY-	Network_Operator	Read-Only-Commands	0
5	●	RO_NAM	contains any (aml.com/Users @ amlActive Users/RO NAM Access/RO_NAM)	-ANY-	-ANY-	-ANY-	-ANY-	RO_NAM	Read-Only-Commands	0
6	●	RO_UCS	contains any (aml.com/Users @ amlActive Users/RO UCS Access/RO_UCS)	-ANY-	-ANY-	-ANY-	-ANY-	RO_UCS	Read-Only-Commands	0
		Default	If no rules defined or no enabled rule matches.					Enable_15	AllowAllcommands	389530

As users log in to devices, they are put in the appropriate groups according to the AD implementation and given the appropriate permissions according to the ACS configuration.

For more information on configuring ACS, see: http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_system/5.2/user/guide/acsuserguide.html.

Monitoring Capability with Network Access Module

Once the NAM is configured, the following steps enable one to monitor traffic using the NAM user interface.

1. Using a browser, connect to the NAM module and create a capture session, as shown in Figure 48.

Figure 48 NAM Capture Session

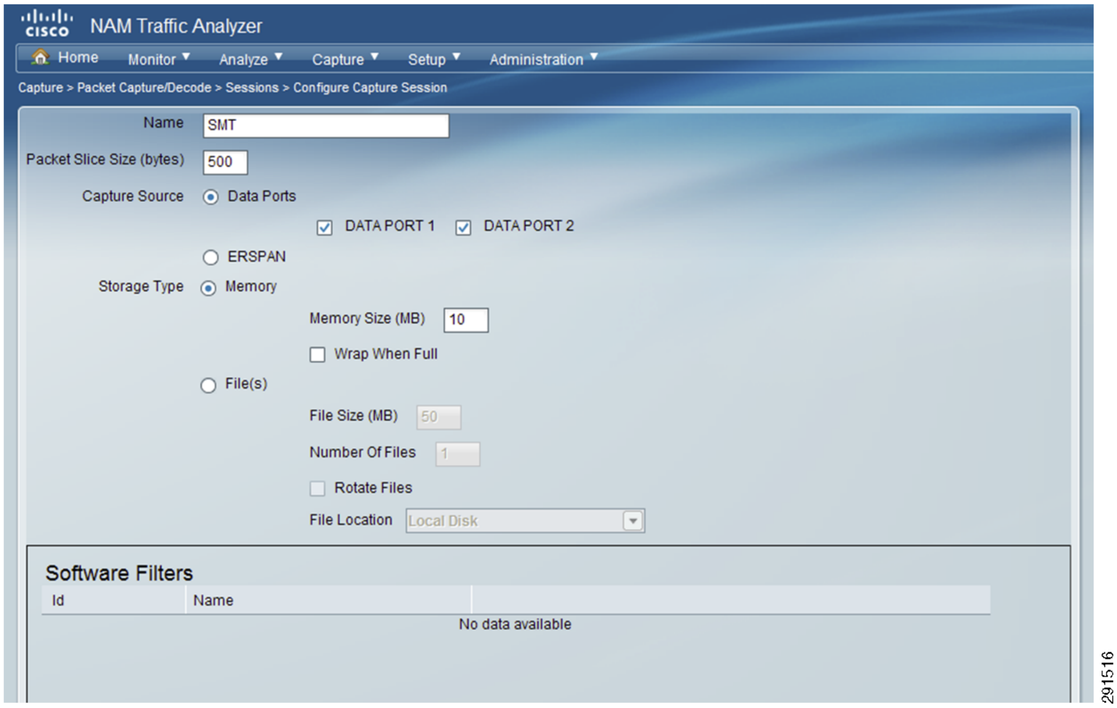
NAM Traffic Analyzer
Home
Monitor
Analyze
Capture
Setup
Administration
Capture > Packet Capture/Decode > Sessions > Configure Capture Session
291516

Name: SMT

Packet Slice Size (bytes): 500

Capture Source: ● Data Ports

☑ DATA PORT 1 ☑ DATA PORT 2

○ ERSPAN

Storage Type: ● Memory

Memory Size (MB): 10

□ Wrap When Full

○ File(s)

File Size (MB): 50

Number Of Files: 1

□ Rotate Files

File Location: Local Disk

Software Filters

Id	Name

No data available

2. One can view and decode the captured traffic by clicking **Start** and **Stop** and viewing the packets by clicking **Decode**.

Figure 49 Starting Capture

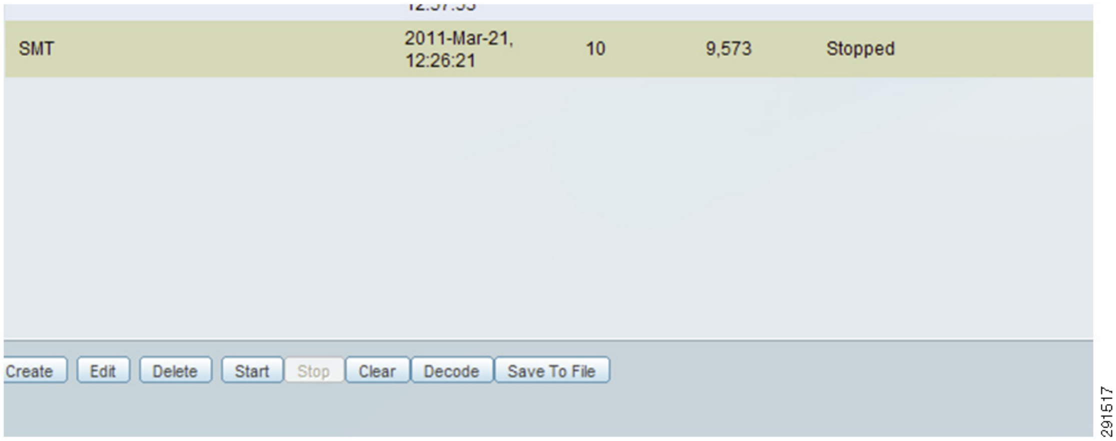

SMT	2011-Mar-21, 12:26:21	10	9,573	Stopped

3. The packet capture view is shown in Figure 50.

Figure 50 Packet Decoder

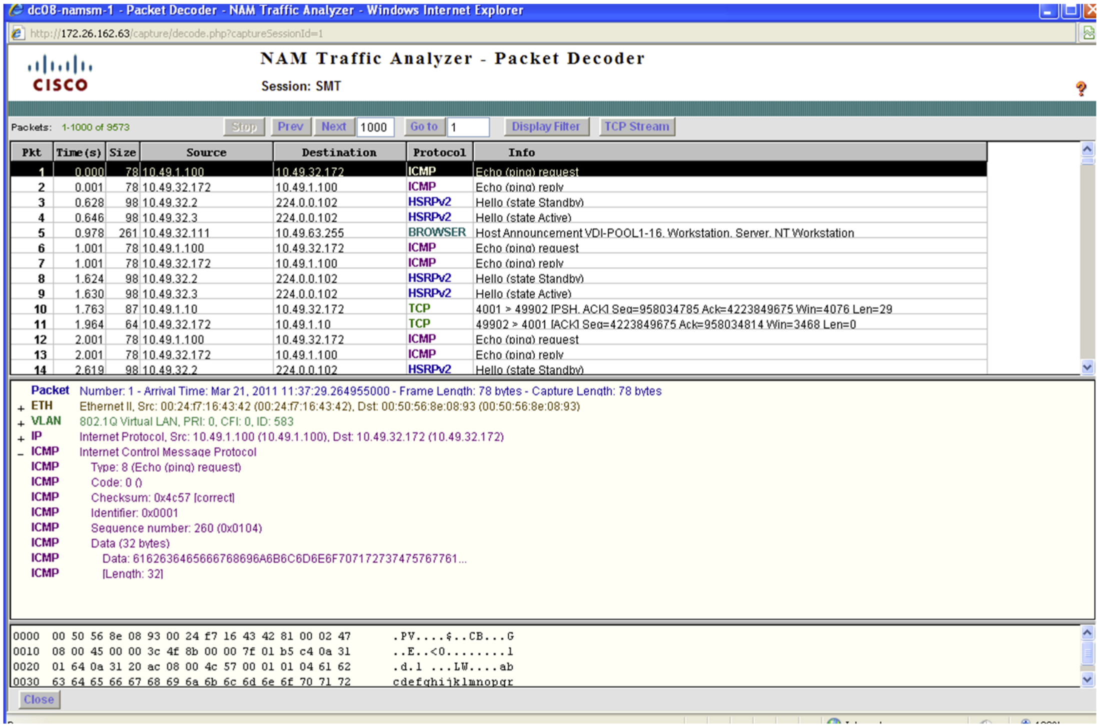
dc08-namsm-1 - Packet Decoder - NAM Traffic Analyzer - Windows Internet Explorer
http://172.26.162.63/capture/decode.php?captureSessionId=1
cisco
NAM Traffic Analyzer - Packet Decoder
Session: SMT
Stop
Prev
Next
Go to
Display Filter
TCP Stream
Close
291518

Packets: 1-1000 of 9573   Next: 1000   Go to: 1

Pkt	Time(s)	Size	Source	Destination	Protocol	Info
1	0.000	78	10.49.1.100	10.49.32.172	ICMP	Echo (ping) request
2	0.001	78	10.49.32.172	10.49.1.100	ICMP	Echo (ping) reply
3	0.628	98	10.49.32.2	224.0.0.102	HSRPv2	Hello (state Standby)
4	0.646	98	10.49.32.3	224.0.0.102	HSRPv2	Hello (state Active)
5	0.978	261	10.49.32.111	10.49.63.255	BROWSER	Host Announcement VDI-POOL1-16, Workstation, Server, NT Workstation
6	1.001	78	10.49.1.100	10.49.32.172	ICMP	Echo (ping) request
7	1.001	78	10.49.32.172	10.49.1.100	ICMP	Echo (ping) reply
8	1.624	98	10.49.32.2	224.0.0.102	HSRPv2	Hello (state Standby)
9	1.630	98	10.49.32.3	224.0.0.102	HSRPv2	Hello (state Active)
10	1.763	87	10.49.1.10	10.49.32.172	TCP	4001 > 49902 [PSH, ACK] Seq=958034785 Ack=4223849675 Win=4076 Len=29
11	1.964	64	10.49.32.172	10.49.1.10	TCP	49902 > 4001 [ACK] Seq=4223849675 Ack=958034814 Win=3468 Len=0
12	2.001	78	10.49.1.100	10.49.32.172	ICMP	Echo (ping) request
13	2.001	78	10.49.32.172	10.49.1.100	ICMP	Echo (ping) reply
14	2.619	98	10.49.32.2	224.0.0.102	HSRPv2	Hello (state Standby)

```
  Packet  Number: 1 - Arrival Time: Mar 21, 2011 11:37:29.264955000 - Frame Length: 78 bytes - Capture Length: 78 bytes
+ ETH     Ethernet II, Src: 00:24:f7:16:43:42 (00:24:f7:16:43:42), Dst: 00:50:56:8e:08:93 (00:50:56:8e:08:93)
+ VLAN    802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 583
+ IP      Internet Protocol, Src: 10.49.1.100 (10.49.1.100), Dst: 10.49.32.172 (10.49.32.172)
- ICMP    Internet Control Message Protocol
  ICMP      Type: 8 (Echo (ping) request)
  ICMP      Code: 0 ()
  ICMP      Checksum: 0x4c57 [correct]
  ICMP      Identifier: 0x0001
  ICMP      Sequence number: 260 (0x0104)
  ICMP      Data (32 bytes)
  ICMP        Data: 6162636465666768696A6B6C6D6E6F7071727374757677 61...
  ICMP        [Length: 32]
```

```
0000  00 50 56 8e 08 93 00 24 f7 16 43 42 81 00 02 47     .PV....$..CB...G
0010  08 00 45 00 00 3c 4f 8b 00 00 7f 01 b5 c4 0a 31     ..E..<O........1
0020  01 64 0a 31 20 ac 08 00 4c 57 00 01 01 04 61 62     .d.1 ...LW....ab
0030  63 64 65 66 67 68 69 6a 6b 6c 6d 6e 6f 70 71 72     cdefghijklmnopqr
```

Monitoring Capability with the Cisco Virtual Network Access Module

Once vNAM is configured, one can connect to the vNAM, using the vNAM address, with a Web interface. Using the Web interface one can choose the profile (NDE-drs1-vsm) from the Monitor/Overview pane and view the traffic flow from the Cisco Nexus 1010, as shown in Figure 51.

Figure 51 Monitoring Traffic Using vNAM

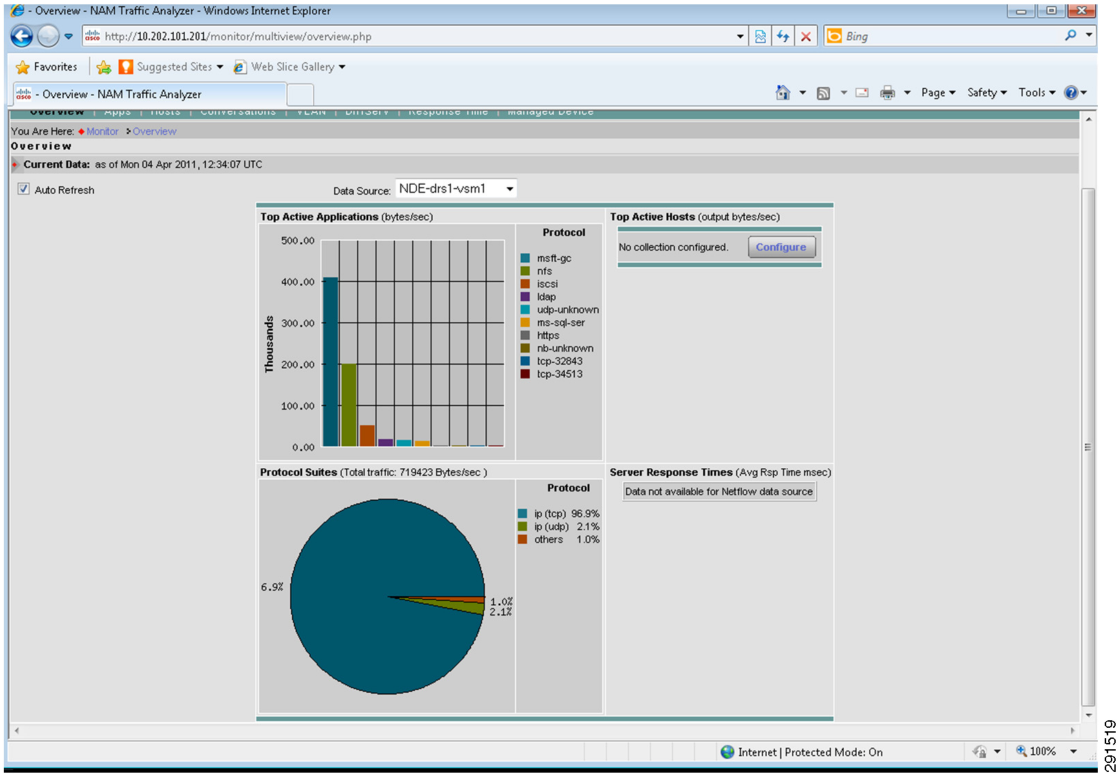

- Overview - NAM Traffic Analyzer - Windows Internet Explorer

http://10.202.101.201/monitor/multiview/overview.php

You Are Here: Monitor > Overview

Overview

Current Data: as of Mon 04 Apr 2011, 12:34:07 UTC

☑ Auto Refresh

Data Source: NDE-drs1-vsm1

Top Active Applications (bytes/sec)

Top Active Hosts (output bytes/sec)

No collection configured. Configure

Protocol Suites (Total traffic: 719423 Bytes/sec)

Server Response Times (Avg Rsp Time msec)

Data not available for Netflow data source

Internet | Protected Mode: On

Monitoring Traffic Flow through Physical Firewall

The ASDM tool for Cisco firewalls is used to monitor both inside and outside interface traffic. Firewall rules, ARP tables, and NAT rules within all contexts can be viewed using the ASDM monitoring capability. The SharePoint tenant, with multiple users accessing a variety of services within the SharePoint environment, is shown in Figure 52.

Figure 52 Monitoring Traffic Using ASDM Monitoring Tool

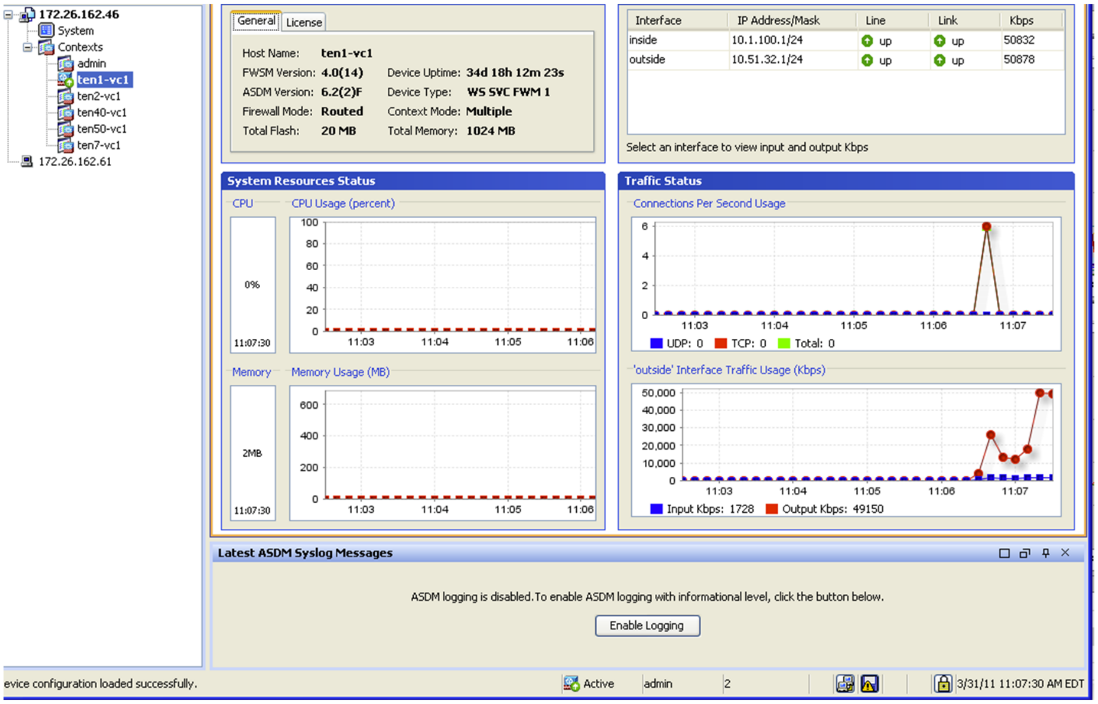

Host Name: **ten1-vc1**
FWSM Version: **4.0(14)** Device Uptime: **34d 18h 12m 23s**
ASDM Version: **6.2(2)F** Device Type: **WS SVC FWM 1**
Firewall Mode: **Routed** Context Mode: **Multiple**
Total Flash: **20 MB** Total Memory: **1024 MB**

Interface	IP Address/Mask	Line	Link	Kbps
inside	10.1.100.1/24	up	up	50832
outside	10.51.32.1/24	up	up	50878

Select an interface to view input and output Kbps

ASDM logging is disabled.To enable ASDM logging with informational level, click the button below.

evice configuration loaded successfully.

Monitoring Capability of ACE Load Balancer

The Cisco Application Network Manager (ANM) displays statistics, such as concurrent connections, connections per second, etc., for the ACE load balancer. These statistics can be monitored and graphed in real time for each physical server, as shown in Figure 53.

Figure 53 Monitoring ACE Using ANM

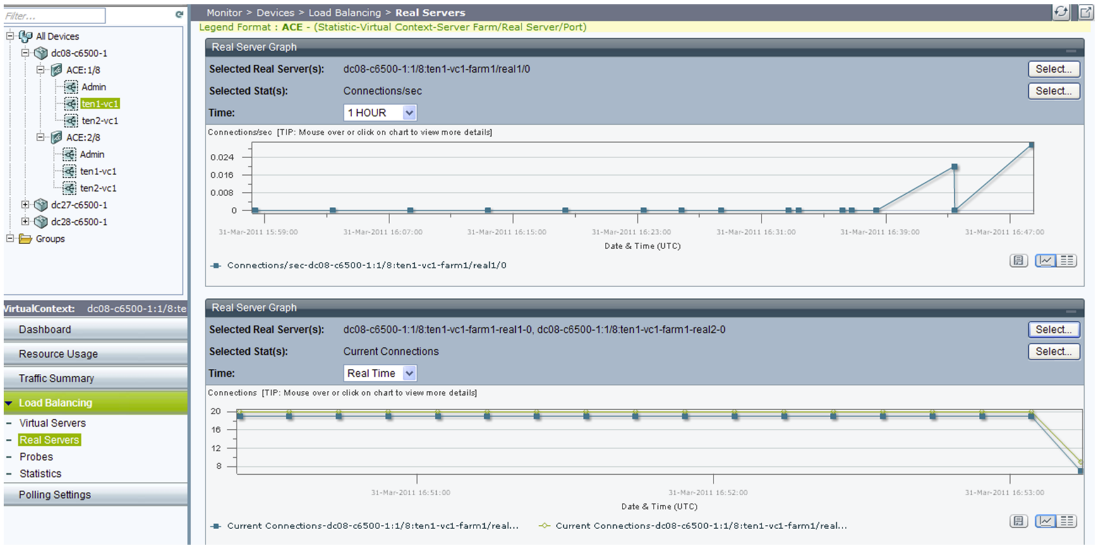

Legend Format : ACE - (Statistic-Virtual Context-Server Farm/Real Server/Port)

Real Server Graph

Selected Real Server(s): dc08-c6500-1:1/8:ten1-vc1-farm1/real1/0
Selected Stat(s): Connections/sec
Time: 1 HOUR

Connections/sec [TIP: Mouse over or click on chart to view more details]

Real Server Graph

Selected Real Server(s): dc08-c6500-1:1/8:ten1-vc1-farm1-real1-0, dc08-c6500-1:1/8:ten1-vc1-farm1-real2-0
Selected Stat(s): Current Connections
Time: Real Time

Connections [TIP: Mouse over or click on chart to view more details]

The various health-probes can also be monitored using the ANM, as shown in Figure 54.

Figure 54 Monitoring Health Probes Using ANM

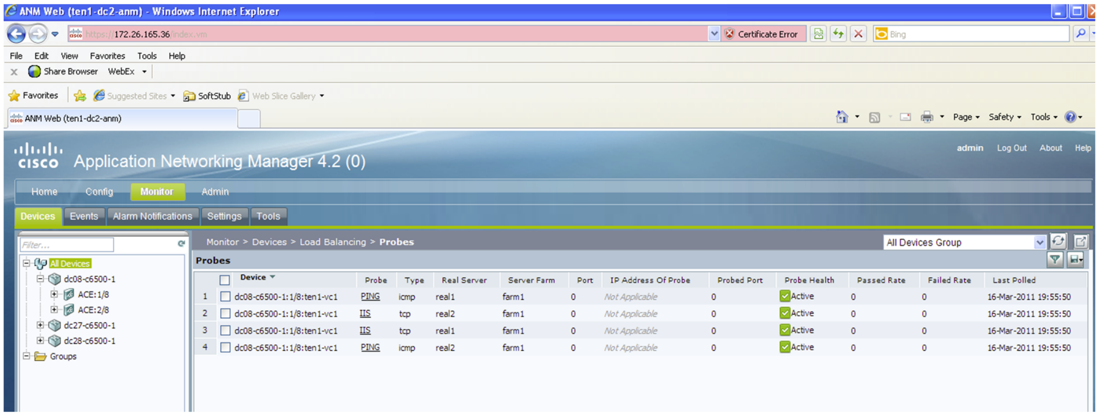

ANM Web (ten1-dc2-anm) - Windows Internet Explorer
https://172.26.165.36/index.vm
Certificate Error
File Edit View Favorites Tools Help
Share Browser WebEx
Favorites
Suggested Sites
SoftStub
Web Slice Gallery
ANM Web (ten1-dc2-anm)
Page Safety Tools
Application Networking Manager 4.2 (0)
admin Log Out About Help
Home Config Monitor Admin
Devices Events Alarm Notifications Settings Tools
Filter...
All Devices
dc08-c6500-1
ACE:1/8
ACE:2/8
dc27-c6500-1
dc28-c6500-1
Groups
Monitor > Devices > Load Balancing > Probes
All Devices Group
291522

Probes

	Device	Probe	Type	Real Server	Server Farm	Port	IP Address Of Probe	Probed Port	Probe Health	Passed Rate	Failed Rate	Last Polled
1	□ dc08-c6500-1:1/8:ten1-vc1	PING	icmp	real1	farm1	0	Not Applicable	0	✓Active	0	0	16-Mar-2011 19:55:50
2	□ dc08-c6500-1:1/8:ten1-vc1	IIS	tcp	real2	farm1	0	Not Applicable	0	✓Active	0	0	16-Mar-2011 19:55:50
3	□ dc08-c6500-1:1/8:ten1-vc1	IIS	tcp	real1	farm1	0	Not Applicable	0	✓Active	0	0	16-Mar-2011 19:55:50
4	□ dc08-c6500-1:1/8:ten1-vc1	PING	icmp	real2	farm1	0	Not Applicable	0	✓Active	0	0	16-Mar-2011 19:55:50

Cisco Intrusion Prevention System Monitoring and Mitigation

The IPS can monitor and mitigate a variety of attacks, such as worms, Trojans, etc. Either the Cisco IPS Manager Express (IME) or Cisco Security Manager (4.2) can be used for monitoring. The IME can import the IPS devices and monitor the health of the network.

The following steps can be taken to mitigate attacks using the Cisco IPS:

1. From the Home/Event dashboard, view the current status of the network and identify events/attacks against the ESMT infrastructure.
2. Identify the event you want to stop by noting the ID number of the signature.
3. In Configuration/Policies, click active signatures.
4. To locate the signature, enter the Signature ID in the Sig ID pane.
5. Right click the signature, click **Edit Actions**, and click **Deny Packet Inline** or any other action deemed appropriate.
6. View the actions taken on the packets in the Event Monitoring/Event View, as shown in Figure 55.

Figure 55 IPS Event Monitoring View

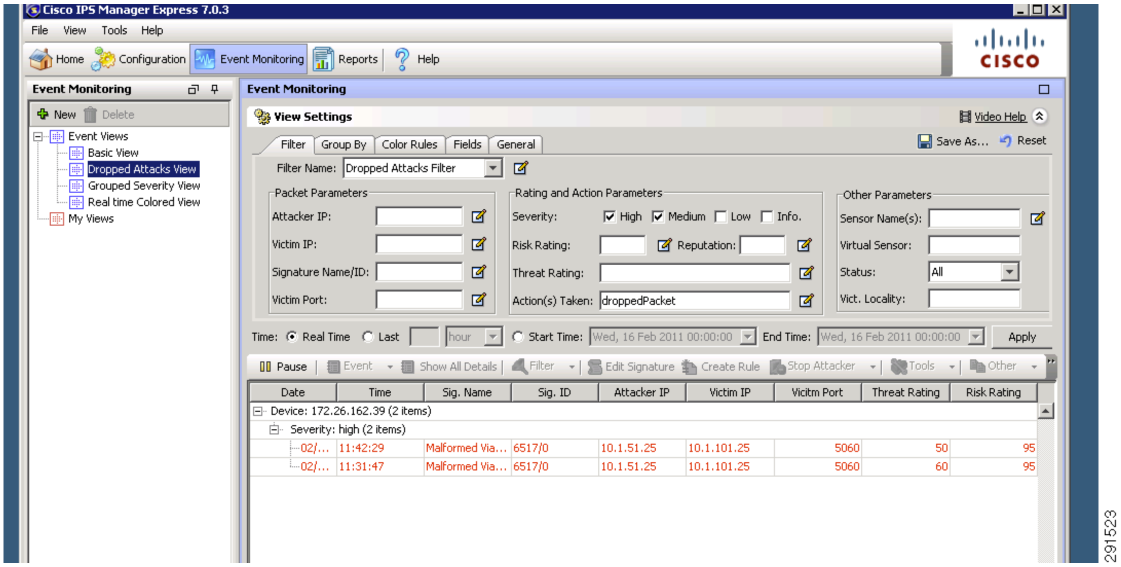

Filter Name: Dropped Attacks Filter

Packet Parameters

Attacker IP:
Victim IP:
Signature Name/ID:
Victim Port:

Rating and Action Parameters

Severity: ☑ High ☑ Medium □ Low □ Info.
Risk Rating: Reputation:
Threat Rating:
Action(s) Taken: droppedPacket

Other Parameters

Sensor Name(s):
Virtual Sensor:
Status: All
Vict. Locality:

Time: ● Real Time ○ Last hour ○ Start Time: Wed, 16 Feb 2011 00:00:00 End Time: Wed, 16 Feb 2011 00:00:00 Apply

Date	Time	Sig. Name	Sig. ID	Attacker IP	Victim IP	Vicitm Port	Threat Rating	Risk Rating
Device: 172.26.162.39 (2 items)								
Severity: high (2 items)								
02/...	11:42:29	Malformed Via...	6517/0	10.1.51.25	10.1.101.25	5060	50	95
02/...	11:31:47	Malformed Via...	6517/0	10.1.51.25	10.1.101.25	5060	60	95

Monitoring Capability Within vShield

vShield can monitor realtime traffic and discover VMs.

Monitoring Realtime Traffic

Monitoring realtime traffic can be important. In the event of a attack, it is necessary to be determine the realtime traffic profile and characteristics to mitigate the attack. Within the VM flow tab, Show Report provides a realtime view of the traffic, as shown in Figure 56. Most malicious attacks are visible through the uncategorized traffic, as the port numbers are random and categorized traffic shows the predefined port-mappings as defined in the port-mapping table under VM flow.

Figure 56 Monitoring Real-time Traffic

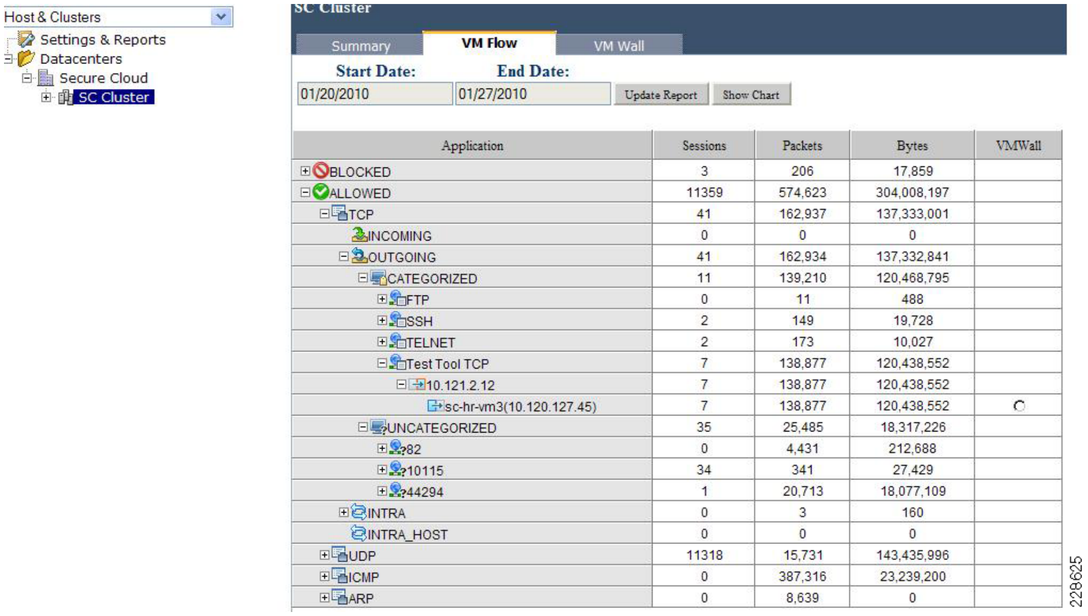

Start Date: 01/20/2010
End Date: 01/27/2010

Application	Sessions	Packets	Bytes	VMWall
⊞ BLOCKED	3	206	17,859	
⊟ ALLOWED	11359	574,623	304,008,197	
⊟ TCP	41	162,937	137,333,001	
INCOMING	0	0	0	
⊟ OUTGOING	41	162,934	137,332,841	
⊟ CATEGORIZED	11	139,210	120,468,795	
⊞ FTP	0	11	488	
⊞ SSH	2	149	19,728	
⊞ TELNET	2	173	10,027	
⊟ Test Tool TCP	7	138,877	120,438,552	
⊟ 10.121.2.12	7	138,877	120,438,552	
sc-hr-vm3(10.120.127.45)	7	138,877	120,438,552	○
⊟ ?UNCATEGORIZED	35	25,485	18,317,226	
⊞ ?82	0	4,431	212,688	
⊞ ?10115	34	341	27,429	
⊞ ?44294	1	20,713	18,077,109	
⊞ INTRA	0	3	160	
INTRA_HOST	0	0	0	
⊞ UDP	11318	15,731	143,435,996	
⊞ ICMP	0	387,316	23,239,200	
⊞ ARP	0	8,639	0	

Figure 57 shows a graphical representation of traffic that is pre-defined in the port-mapping configuration table, as shown in Figure 57.

Figure 57 VM Chart

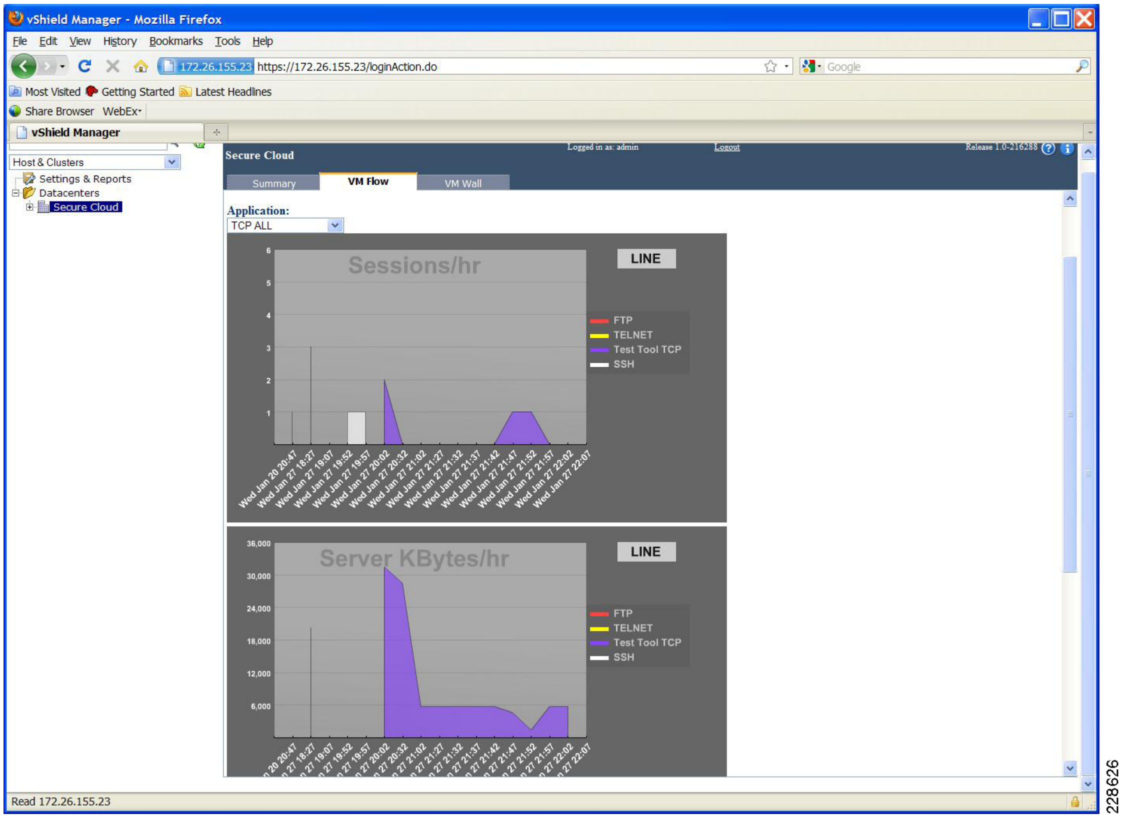

https://172.26.155.23/loginAction.do

Application: TCP ALL

Read 172.26.155.23

VM Discovery

VM Discovery enables one to monitor inter-tenant and intra-tenant steady-state available services. This can be useful in cases where one wants to ascertain which ports are visible and active for a particular VM. One can start the VM discovery process by operating on each vShield agent and staring the VM discovery process. One can perform this discovery continuously or in a scheduled manner, as shown in Figure 58.

Figure 58 Scheduling VM Discovery

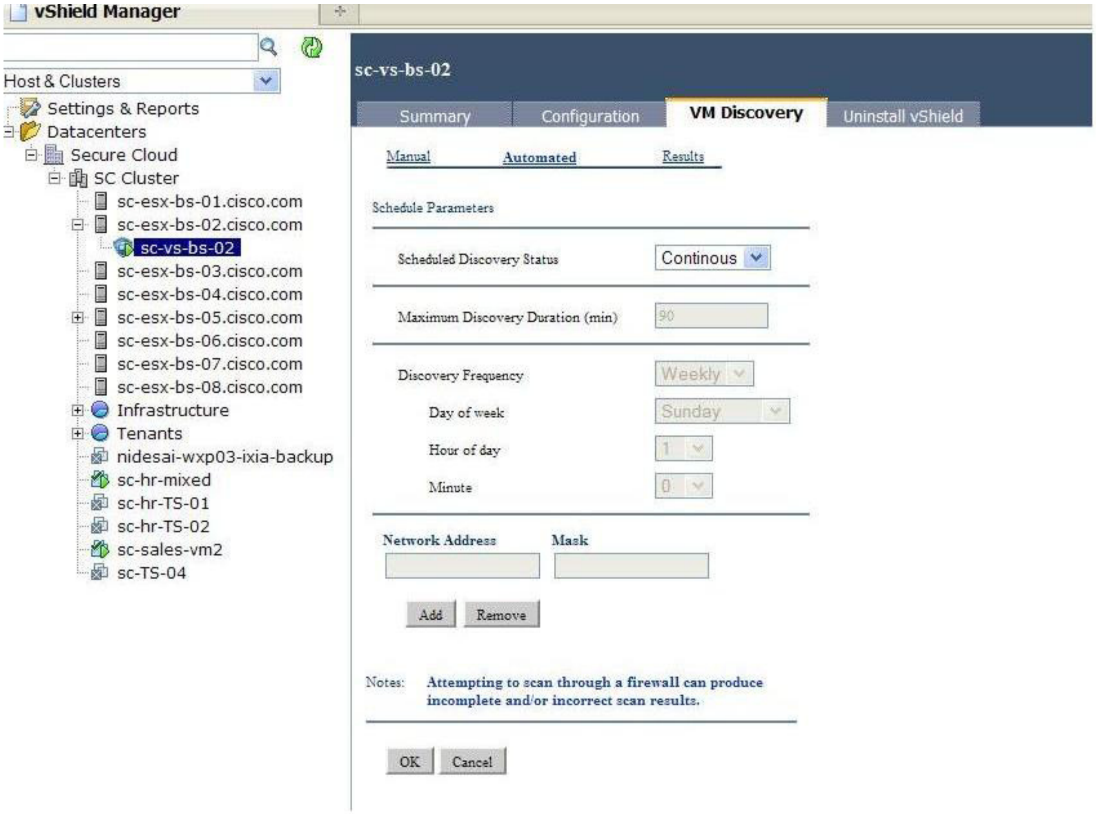

Schedule Parameters

Scheduled Discovery Status: Continous

Maximum Discovery Duration (min): 90

Discovery Frequency: Weekly

Day of week: Sunday

Hour of day: 1

Minute: 0

Network Address

Mask

Notes: Attempting to scan through a firewall can produce incomplete and/or incorrect scan results.

228634

Once the VM discovery process is complete, one can view the steady state traffic and open ports, as shown in Figure 59.

Figure 59 VM Discovery Mode

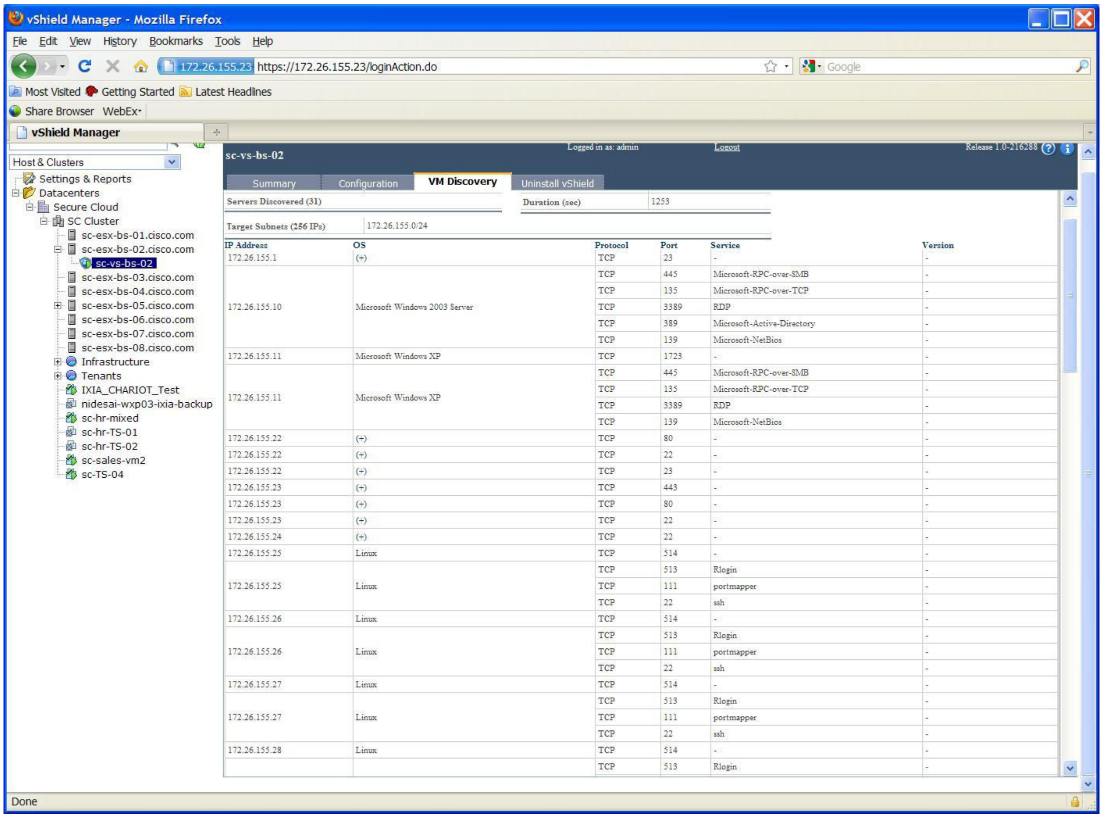

Servers Discovered (31)

Duration (sec) 1253

Target Subnets (256 IPs) 172.26.155.0/24

IP Address	OS	Protocol	Port	Service	Version
172.26.155.1	(+)	TCP	23	-	-
172.26.155.10	Microsoft Windows 2003 Server	TCP	445	Microsoft-RPC-over-SMB	-
		TCP	135	Microsoft-RPC-over-TCP	-
		TCP	3389	RDP	-
		TCP	389	Microsoft-Active-Directory	-
		TCP	139	Microsoft-NetBios	-
172.26.155.11	Microsoft Windows XP	TCP	1723	-	-
172.26.155.11	Microsoft Windows XP	TCP	445	Microsoft-RPC-over-SMB	-
		TCP	135	Microsoft-RPC-over-TCP	-
		TCP	3389	RDP	-
		TCP	139	Microsoft-NetBios	-
172.26.155.22	(+)	TCP	80	-	-
172.26.155.22	(+)	TCP	22	-	-
172.26.155.22	(+)	TCP	23	-	-
172.26.155.23	(+)	TCP	443	-	-
172.26.155.23	(+)	TCP	80	-	-
172.26.155.23	(+)	TCP	22	-	-
172.26.155.24	(+)	TCP	22	-	-
172.26.155.25	Linux	TCP	514	-	-
172.26.155.25	Linux	TCP	513	Rlogin	-
		TCP	111	portmapper	-
		TCP	22	ssh	-
172.26.155.26	Linux	TCP	514	-	-
172.26.155.26	Linux	TCP	513	Rlogin	-
		TCP	111	portmapper	-
		TCP	22	ssh	-
172.26.155.27	Linux	TCP	514	-	-
172.26.155.27	Linux	TCP	513	Rlogin	-
		TCP	111	portmapper	-
		TCP	22	ssh	-
172.26.155.28	Linux	TCP	514	-	-
		TCP	513	Rlogin	-

228627

VM discovery only works with tenants that are routable to the management VLAN.

Load Balancing

100 separate users accessed the SharePoint servers concurrently (this was done using the Microsoft Visual Studio tool). The 100 concurrent connections were between two real servers. The ACE was able to load balance these many concurrent connections, as shown in Figure 60.

Figure 60 Monitoring Conncurrent Load Balancing Sessions

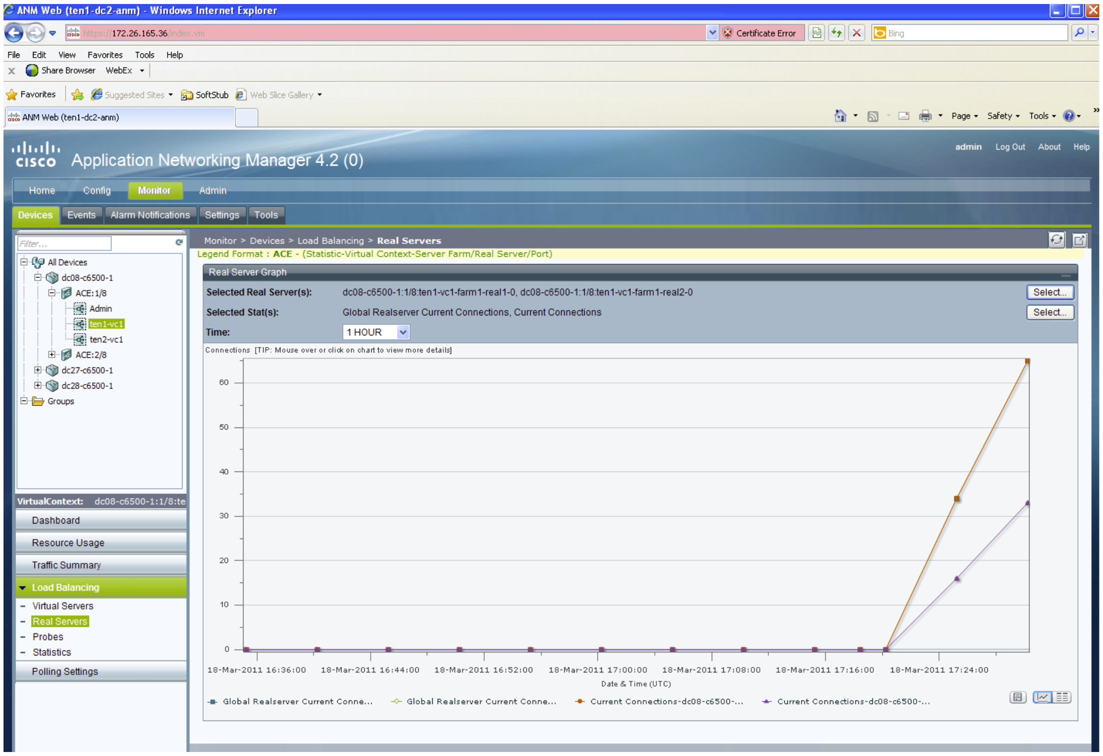

Legend Format : ACE - (Statistic-Virtual Context-Server Farm/Real Server/Port)

Selected Real Server(s): dc08-c6500-1:1/8:ten1-vc1-farm1-real1-0, dc08-c6500-1:1/8:ten1-vc1-farm1-real2-0

Selected Stat(s): Global Realserver Current Connections, Current Connections

Time: 1 HOUR

Connections [TIP: Mouse over or click on chart to view more details]

Deploying Tenant Resources and Security

Tenant Storage Resources

Tenant vFiler Units

NetApp Operations Manager and Provisioning Manager should be used to provision tenant vFiler units and to apply consistent storage efficiency policies to storage within those vFiler units. To provision tenant vFiler units, a tenant storage VLAN with an MTU of 9000—if jumbo frames are being used—must be provisioned in the Cisco Nexus 1000V, in the Cisco UCS fabric and all Cisco UCS vNICs on the ESXi hosts, in the Cisco Nexus 5000s, and on the storage system VIF. Tenant datastores should be provisioned with the Virtual Storage Console (VSC) Provisioning and Cloning Utility, as described in the following section. Tenant storage policies should be created for both NFS and SAN (iSCSI) storage. Storage can then be consistently provisioned using these policies. Storage provisioned with VSC can also be imported into datasets that include the policies in order to apply the policies to that storage.

Tenant NFS Datastores

NetApp recommendeds provisioning tenant NFS datastores using the VSC Provisioning and Cloning Utility. To provisionin tenant NFS datastores in this way, a tenant storage VLAN with an MTU of 9000—if jumbo frames are being used—must be provisioned in the Cisco Nexus 1000V, in the Cisco

UCS fabric and all Cisco UCS vNICs on the ESXi hosts, in the Cisco Nexus 5000s, and on the storage system VIF. The tenant vFiler unit should be provisioned using NetApp Provisioning Manager and should include the VLAN just created. Also, a VMware vmkernel port must be provisioned on each ESXi host from which you want to mount the datastore. This vmkernel port should be provisioned with an MTU of 9000 if jumbo frames are being used. In order to create a vmkernel port with an MTU of 9000, first provision the port on the Cisco Nexus 1000V in vCenter under the ESXi host Configuration tab where Networking, then vNetwork Distributed Switch, then Manage Virtual Adapters is selected. This creates a vmkernel port with an MTU of 1500. It is then necessary to log into the ESXi host console, query and delete the port, then add it back in to the DVS on the same DVS port with an MTU of 9000. This vFiler unit then must be added as a storage system on the VSC Provisioning and Cloning Storage controllers screen.

Once the network, the vmkernel port, and the tenant vFiler unit have been properly setup, the new tenant NFS datastore can be setup by selecting NetApp in the right-click menu on the cluster containing all the ESXi hosts where the datastore is to be mounted. The storage system containing the vFiler unit should be selected along with the vFiler context. Once Finish is selected, VSC creates the volume for the datastore on the storage system, moves the volume into the tenant vFiler unit, creates the NFS exports to all of the ESXi hosts in the cluster, and mounts the datastore on all ESXi hosts in the cluster. The VSC Provisioning and Cloning Utility can then be used to clone VMs within this new datastore. The VSC Backup and Recovery Utility can also be used to take consistent Snapshot copies of the VMs in this datastore. The datastore volume can also be imported into a dataset in Provisioning Manager to automatically apply storage efficiency policies to the volume, such as Deduplication Management.

Intra- and Inter-Tenant Traffic Rules

This section outlines the necessary steps to achieve network isolation between and within tenants.

Deploying vShield—Implementing Security Rules at the Access Layer

A virtual firewall such as VMware's vShield can be used to create separation between tenants and from outside. Within the access layer in the ESMT architecture, virtual firewalls are implemented to segment traffic within the same tenant environment (sometimes referred to as east-west traffic, or optionally across tenants as well). Figure 61 shows how the various traffic flows are protected by virtual and physical firewalls.

Figure 61 Traffic Flows within a Multi-Tenant Architecture

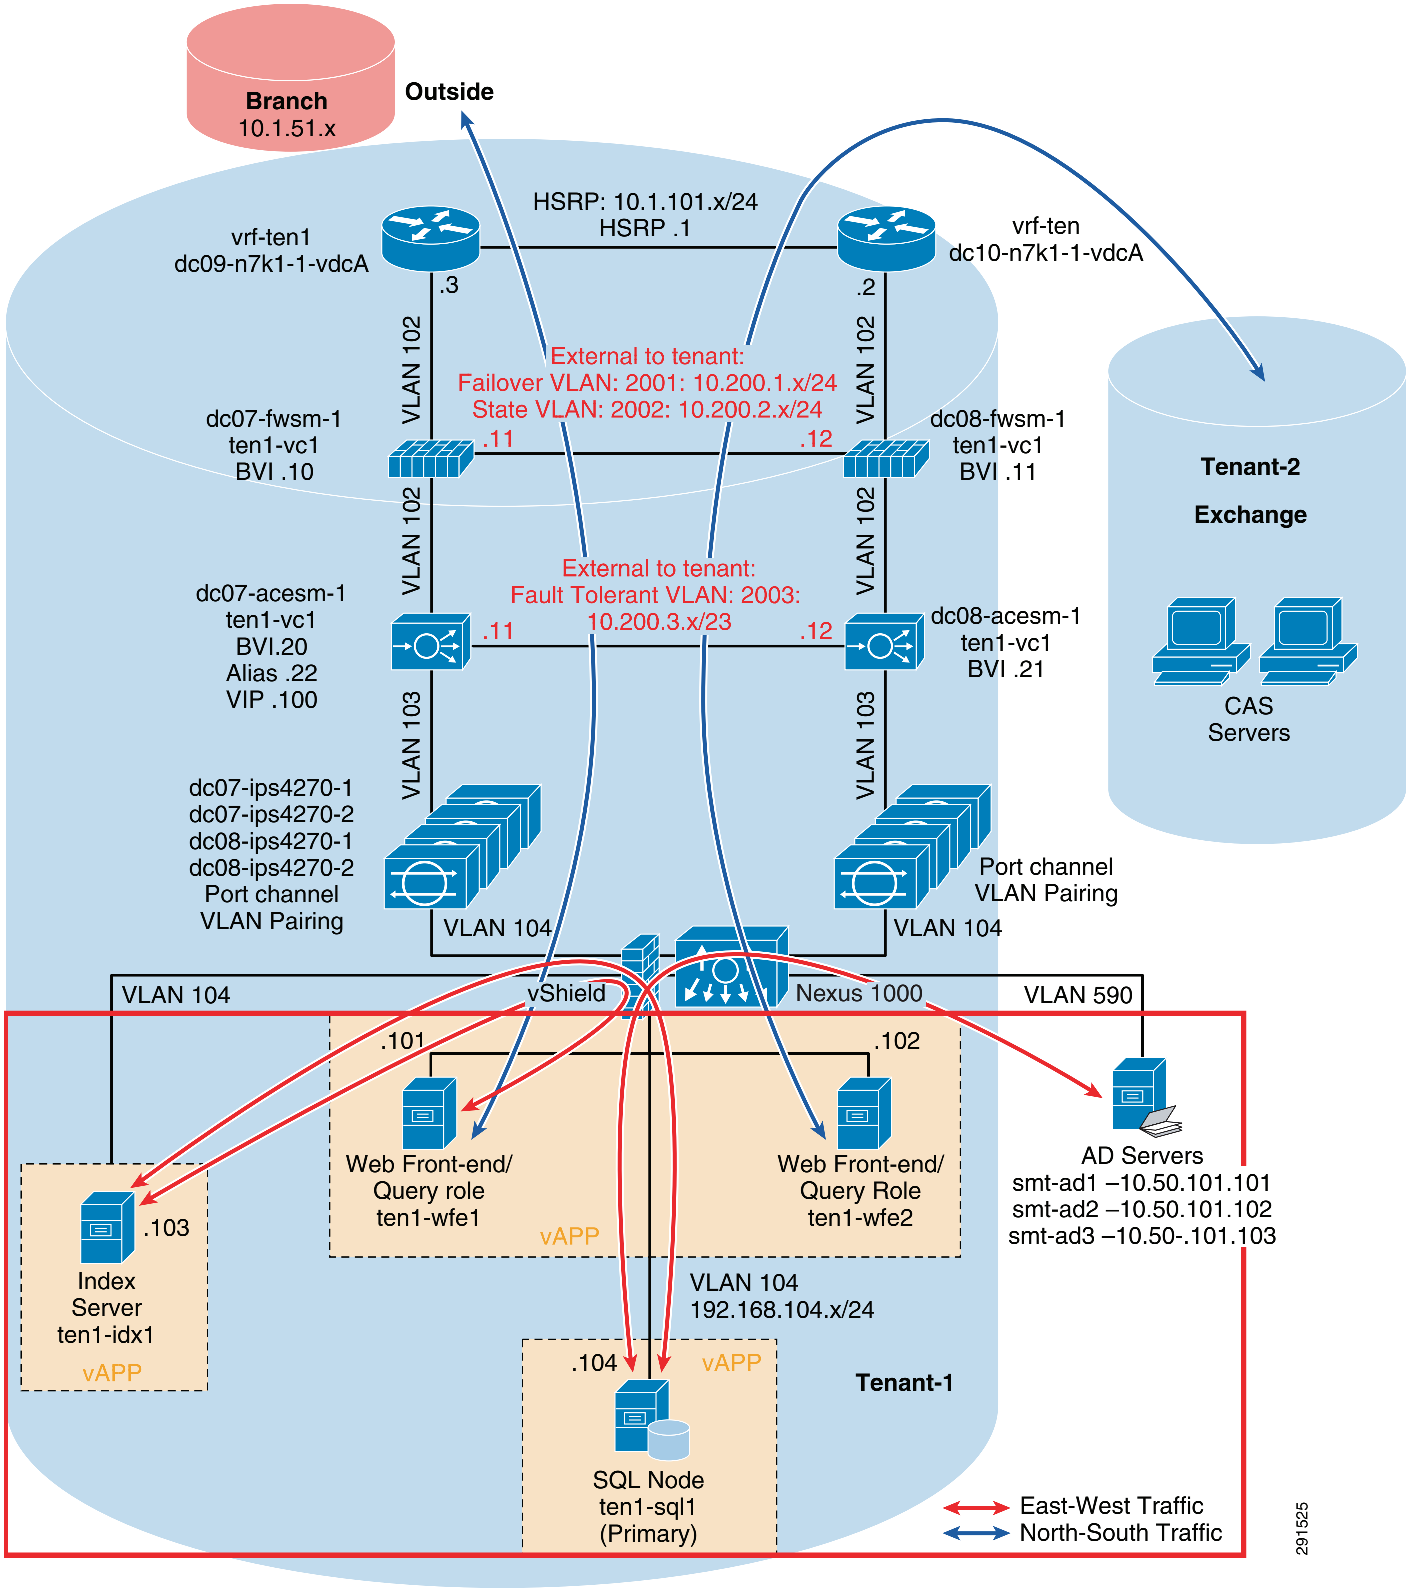

Firewall rules can be easily set to deny or authorize hosts based on an IP address or specific VLANs. It can also lock down traffic to certain ports. The example below shows how the SharePoint tenant is allowed access only from the branch (10.1.151.x subnet).

For example:

1. Choose Host & Clusters view, as shown in Figure 62.

Figure 62 Host & Clusters View

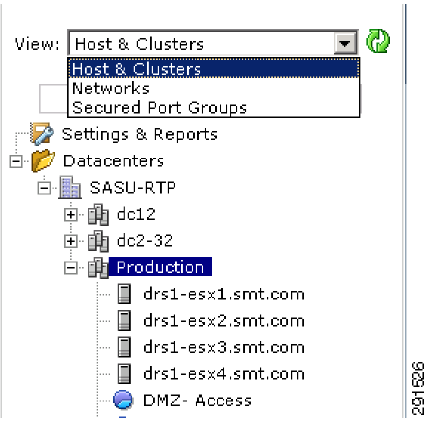

View: Host & Clusters

2. Choose your corresponding data center cluster (Production in the example shown in Figure 63).

Figure 63 Data Center View

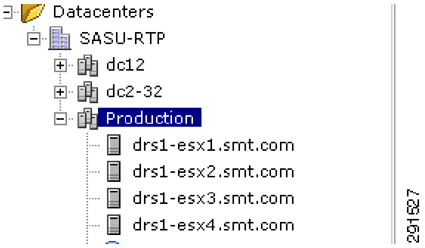

3. Choose the App Firewall pane, as shown in Figure 64.

Figure 64 App Firewall Pane

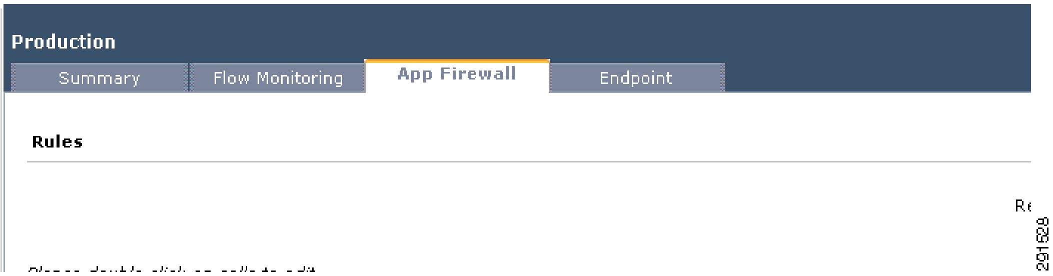

4. Create a rule to allow only a certain host with traffic on certain ports to access the inside network of the tenant. In this case, subnet 10.1.51.0/24 is allowed access on ports 80 and 3112.

Figure 65 Branch-Tenant 1 Firewall Rules

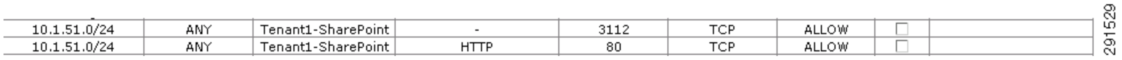

10.1.51.0/24	ANY	Tenant1-SharePoint	-	3112	TCP	ALLOW	□	
10.1.51.0/24	ANY	Tenant1-SharePoint	HTTP	80	TCP	ALLOW	□	

5. Create a rule to deny every other host access to the inside network of the tenant (including Tenant 2).

Figure 66 Firewall Rules Restricting Traffic to Tenant 1

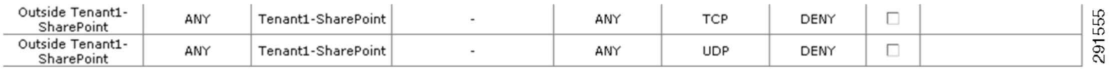

Outside Tenant1-SharePoint	ANY	Tenant1-SharePoint	-	ANY	TCP	DENY	□	
Outside Tenant1-SharePoint	ANY	Tenant1-SharePoint	-	ANY	UDP	DENY	□	

6. Now one needs to set Layer 2/Layer 3 rules to deny ping. Go to the root cluster SASU-RTP in the Host & Clusters view, as shown in Figure 67.

Figure 67 Root Cluster View

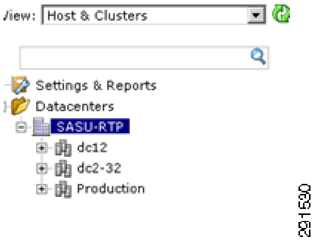

7. Click **L2/L3 Rules** in the App Firewall pane, as shown in Figure 68.

Figure 68 L2/L3 Rules in the App Firewall Pane

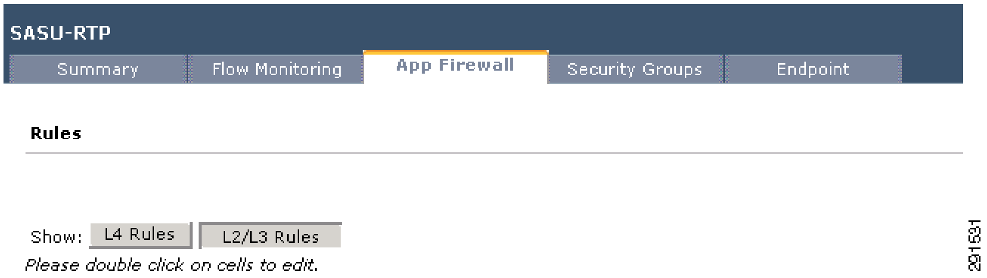

Please double click on cells to edit.

8. Make rules to deny ICMP packets to inside of the SharePoint tenant.

Figure 69 L2/L3 Firewall Rule for Tenant 1

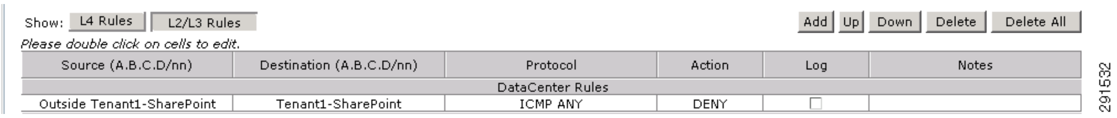

Please double click on cells to edit.

Source (A.B.C.D/nn)	Destination (A.B.C.D/nn)	Protocol	Action	Log	Notes
DataCenter Rules					
Outside Tenant1-SharePoint	Tenant1-SharePoint	ICMP ANY	DENY	□	

Configuring NAT Between SharePoint and Infrastructure Tenant

The SharePoint tenant needs access to the shared Infrastructurre tenant for Active Directory and other services. To perform this, vShield edge can be deployed to provide NAT between SharePoint and AD tenants, as shown below. The back-end subnet 10.1.102.X is NATed to the Infrastructurre tenant.

1. Choose Networks in the View pane and choose the VLAN/subnet for the SharePoint backend.

Figure 70 vShield Edge View

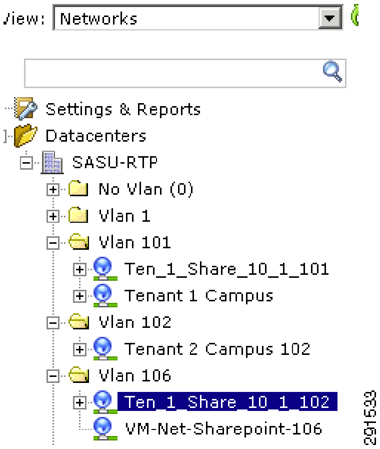

2. Choose the Edge Tab and start filling out appropriate information for the vShield Edge instance.
   - You have to choose an internal and external IP address for the NAT. The internal address designates the internal IP address within the SharePoint tenant that is used to access Infrastructure tenants. External IP address is the address within the Infrastructure tenant to which all packets are translated.
   - Enter the default gateway of the Infrastructure tenant and choose the appropriate data store where you want to store the vShield Edge instance.
3. Click **Install**. After installation Figure 71 shows what should be shown on the Status tab on the vShield Instance.

Figure 71 vShield Edge Configuration Pane

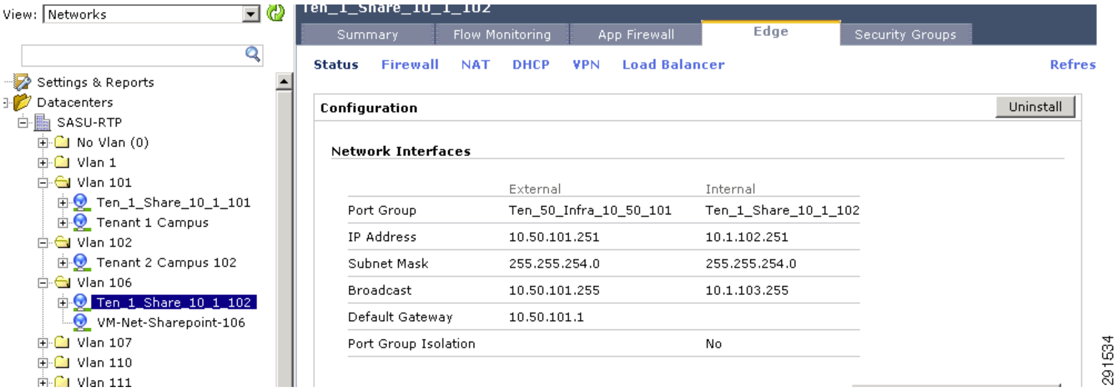

Configuration

Network Interfaces

	External	Internal
Port Group	Ten_50_Infra_10_50_101	Ten_1_Share_10_1_102
IP Address	10.50.101.251	10.1.102.251
Subnet Mask	255.255.254.0	255.255.254.0
Broadcast	10.50.101.255	10.1.103.255
Default Gateway	10.50.101.1	
Port Group Isolation		No

4. Install NAT rules, as shown in Figure 72, to allow NATing between the SharePoint and Infrastructure tenants.

Figure 72 NAT Configuration Pane

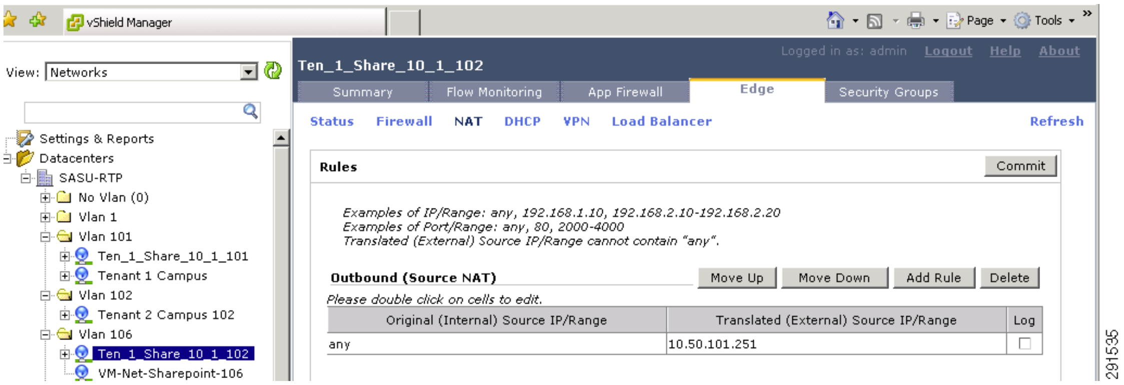

Ten_1_Share_10_1_102

Rules

Examples of IP/Range: any, 192.168.1.10, 192.168.2.10-192.168.2.20
Examples of Port/Range: any, 80, 2000-4000
Translated (External) Source IP/Range cannot contain "any".

Outbound (Source NAT)

Please double click on cells to edit.

Original (Internal) Source IP/Range	Translated (External) Source IP/Range	Log
any	10.50.101.251	☐

5. Configure static/persistent routes between hosts on SharePoint tenants and the Infrastructure subnet to point to the internal NAT IP address (10.1.102.251 in Figure 72) as the next hop.

Configuring Rules in the Backend to Securely Deploy Database Applications

Optionally one can use the vShield to create rules between different hosts in the same tenant. There may be a case where the different database components coexist with other virtual machines in the same tenant. The vApp feature within vShield lets you set rules based on host IP address within the same tenant. In Figure 73, the ten1-media1 server (IP address 10.1.101.106) is only accessed from outside and is denied access to all hosts within the SharePoint infrastructure.

1. From the host/cluster view, choose Production and App Firewall.
2. Create a rule to deny 10.1.101.106 to all hosts within the SharePoint tenant, as shown in Figure 73.

Figure 73 vApp Firewall Rules within Tenant 1

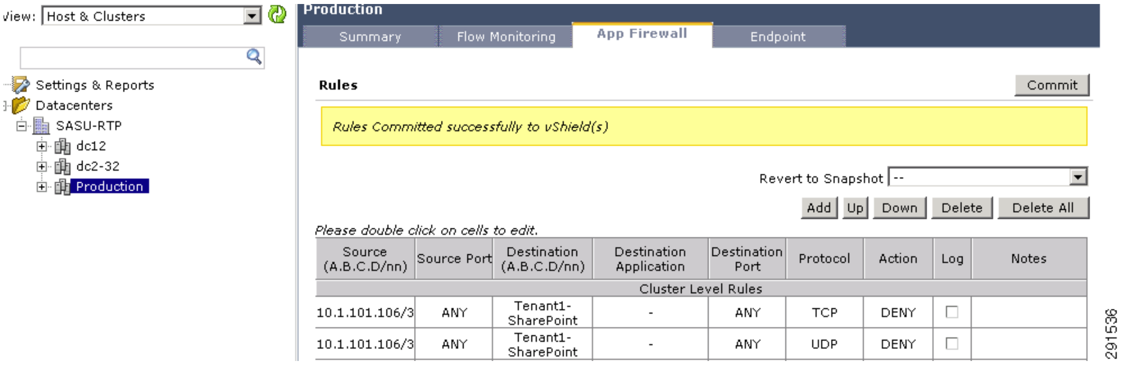

Production

Rules

Rules Committed successfully to vShield(s)

Revert to Snapshot --

Please double click on cells to edit.

Source (A.B.C.D/nn)	Source Port	Destination (A.B.C.D/nn)	Destination Application	Destination Port	Protocol	Action	Log	Notes
Cluster Level Rules								
10.1.101.106/3	ANY	Tenant1-SharePoint	-	ANY	TCP	DENY	☐	
10.1.101.106/3	ANY	Tenant1-SharePoint	-	ANY	UDP	DENY	☐	

3. Create a similar rule for Layer 2/Layer 3 to prevent ICMP connectivity (follow 7. above to get to the appropriate configuration window).

Figure 74 vApp L2/L3 Firewall Example

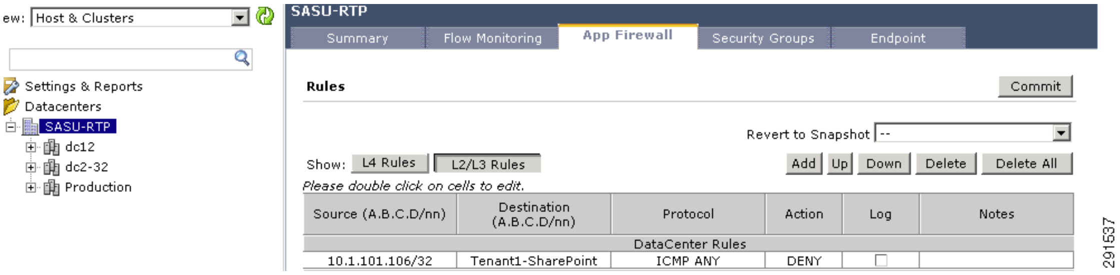

Source (A.B.C.D/nn)	Destination (A.B.C.D/nn)	Protocol	Action	Log	Notes
DataCenter Rules					
10.1.101.106/32	Tenant1-SharePoint	ICMP ANY	DENY	□	

Deploying the Cisco Firewall Switching Module—Implementing Security Rules at Services Layer

Optionally the firewall at the services layer can be configured to isolate tenants from outside or from each other. The choice is up to the network administrator, where he/she may decide to isolate high-performance tenants using a physical firewall at the services layer. In the example below the exchange tenant is isolated from other tenants and outside traffic is restricted to a predefined subnet (10.2.51.0). One can use the ASDM configuration tool to perform this task:

1. Select the Configuration tab.
2. Go to your Context, click **Access Rules**, then click **Add**.

Figure 75 Configuring Firewall Rules with FWSM

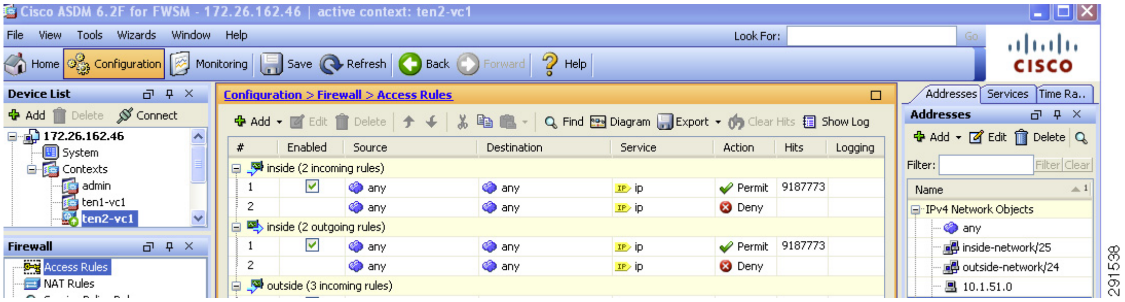

#	Enabled	Source	Destination	Service	Action	Hits	Logging
inside (2 incoming rules)							
1	☑	any	any	ip	✓ Permit	9187773	
2		any	any	ip	☒ Deny		
inside (2 outgoing rules)							
1	☑	any	any	ip	✓ Permit	9187773	
2		any	any	ip	☒ Deny		
outside (3 incoming rules)							

3. Choose rule for outside, add the subnet you want to allow, click **OK**, then click **Apply**.

Figure 76 FWSM Firewall Rule Configuration Window

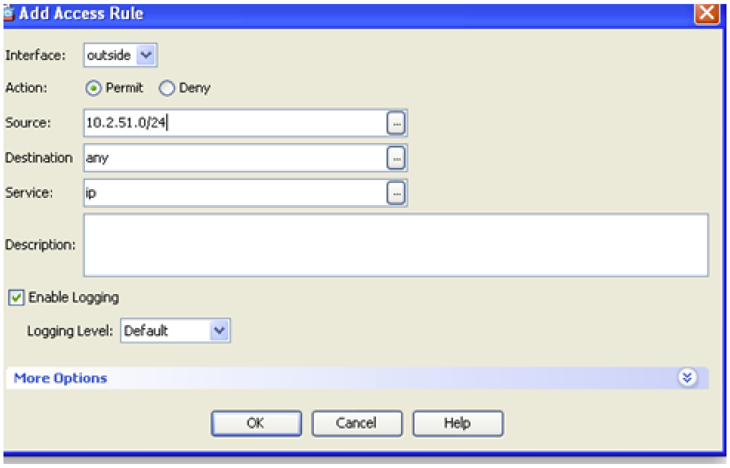

Interface: outside

Action: ● Permit ○ Deny

Source: 10.2.51.0/24

Destination: any

Service: ip

Description:

☑ Enable Logging

Logging Level: Default

Attack Mitigation

A specific attack was initiated from outside the network and seen on the main dashboard.

Figure 77 IPS Dashboard Events View

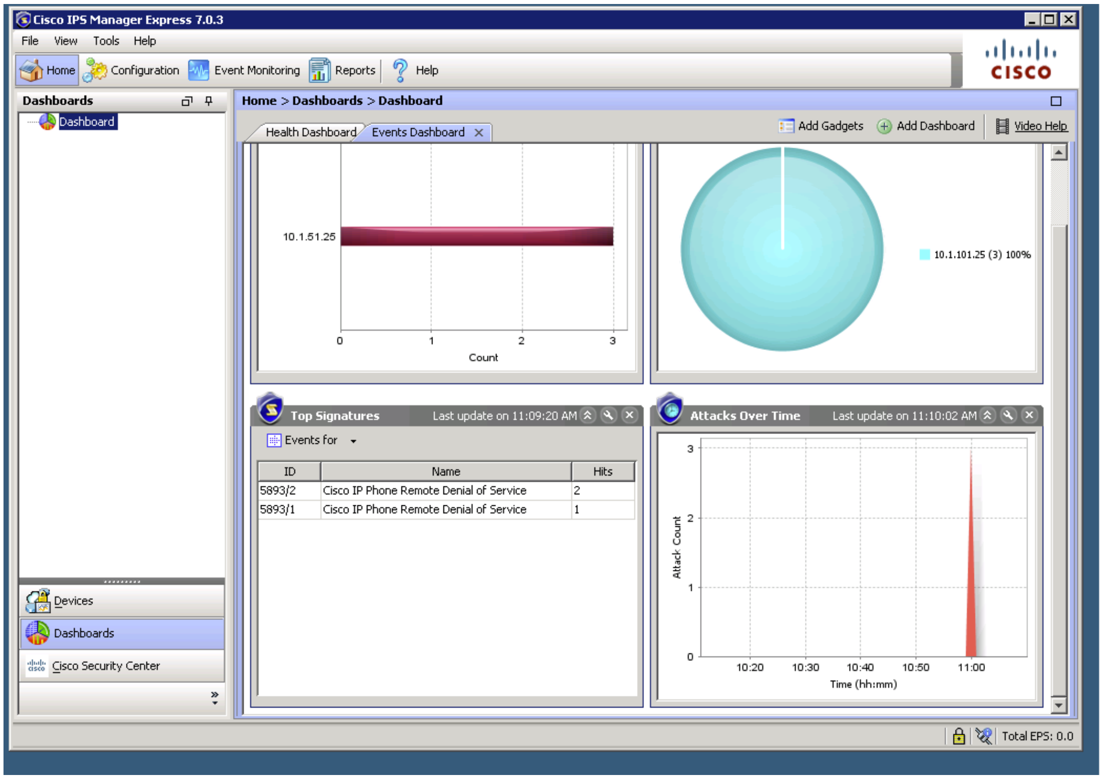
Cisco IPS Manager Express 7.0.3
File View Tools Help
Home Configuration Event Monitoring Reports Help
CISCO
Dashboards
Dashboard
Home > Dashboards > Dashboard
Health Dashboard Events Dashboard
Add Gadgets Add Dashboard Video Help
10.1.51.25
0 1 2 3
Count
10.1.101.25 (3) 100%
Top Signatures Last update on 11:09:20 AM
Events for
Attacks Over Time Last update on 11:10:02 AM
3 2 1 0
Attack Count
10:20 10:30 10:40 10:50 11:00
Time (hh:mm)
Devices
Dashboards
Cisco Security Center
Total EPS: 0.0
291540

ID	Name	Hits
5893/2	Cisco IP Phone Remote Denial of Service	2
5893/1	Cisco IP Phone Remote Denial of Service	1

The Signature ID was used to locate the attack within the available list of signatures and appropriate actions were taken by right clicking the signature to drop these packets and mitigate the attack, as shown in Figure 78.

Figure 78 Configuring Actions to Mitigate Attacks

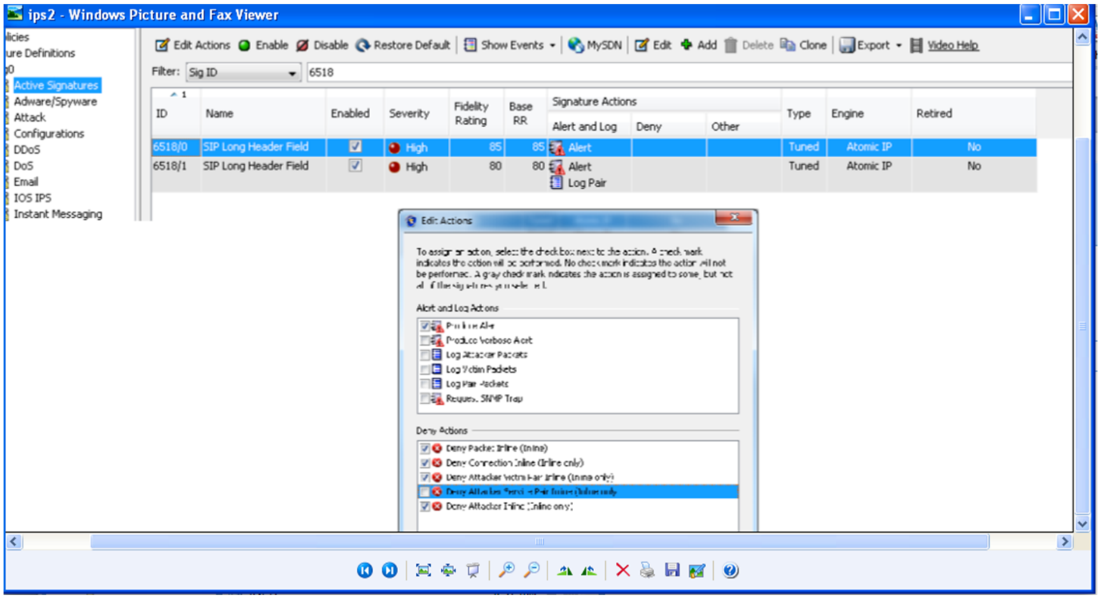

Filter: Sig ID 6518

ID	Name	Enabled	Severity	Fidelity Rating	Base RR	Signature Actions: Alert and Log	Signature Actions: Deny	Signature Actions: Other	Type	Engine	Retired
6518/0	SIP Long Header Field	☑	● High	85	85	Alert			Tuned	Atomic IP	No
6518/1	SIP Long Header Field	☑	● High	80	80	Alert Log Pair			Tuned	Atomic IP	No

To assign an action, select the check box next to the action. A check mark indicates the action will be performed. No check mark indicates the action will not be performed. A gray check mark indicates the action is assigned to some, but not all, of the signatures you selected.

Alert and Log Actions

- ☑ Produce Alert
- ☐ Produce Verbose Alert
- ☐ Log Attacker Packets
- ☐ Log Victim Packets
- ☐ Log Pair Packets
- ☐ Request SNMP Trap

Deny Actions

- ☑ Deny Packet Inline (Inline)
- ☑ Deny Connection Inline (Inline only)
- ☑ Deny Attacker Victim Pair Inline (Inline only)
- ☐ Deny Attacker Service Pair Inline (Inline only)
- ☑ Deny Attacker Inline (Inline only)

One could also go to the Event Monitoring/Dropped View screen to view dropped packets in real time or from a previous time span.

Figure 79 Dropped Attacks Event Statistics

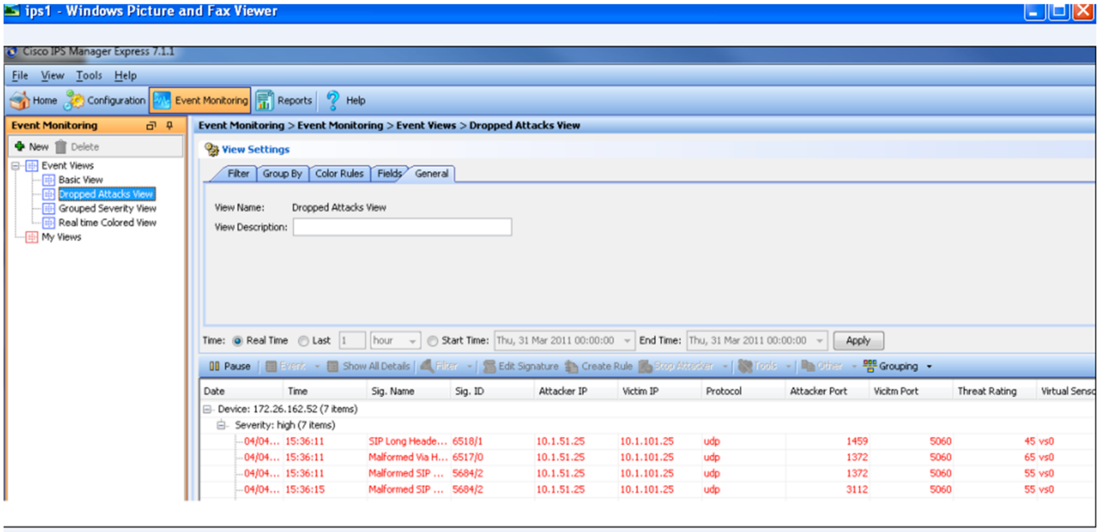

Event Monitoring > Event Monitoring > Event Views > Dropped Attacks View

View Name: Dropped Attacks View

View Description:

Time: ● Real Time ○ Last 1 hour ○ Start Time: Thu, 31 Mar 2011 00:00:00 End Time: Thu, 31 Mar 2011 00:00:00 Apply

```
Date        Time      Sig. Name          Sig. ID   Attacker IP   Victim IP     Protocol   Attacker Port   Vicitm Port   Threat Rating   Virtual Sens
⊟ Device: 172.26.162.52 (7 items)
  ⊟ Severity: high (7 items)
    04/04...  15:36:11  SIP Long Heade...  6518/1    10.1.51.25    10.1.101.25   udp                 1459          5060              45   vs0
    04/04...  15:36:11  Malformed Via H... 6517/0    10.1.51.25    10.1.101.25   udp                 1372          5060              65   vs0
    04/04...  15:36:11  Malformed SIP ...  5684/2    10.1.51.25    10.1.101.25   udp                 1372          5060              55   vs0
    04/04...  15:36:15  Malformed SIP ...  5684/2    10.1.51.25    10.1.101.25   udp                 3112          5060              55   vs0
```

Figure 80 shows the detection of an ICMP sweep by IPS. Event views in the Event Monitoring pane show that IPS detected the ICMP sweep in realtime.

Figure 80 Detecting ICMP Sweep

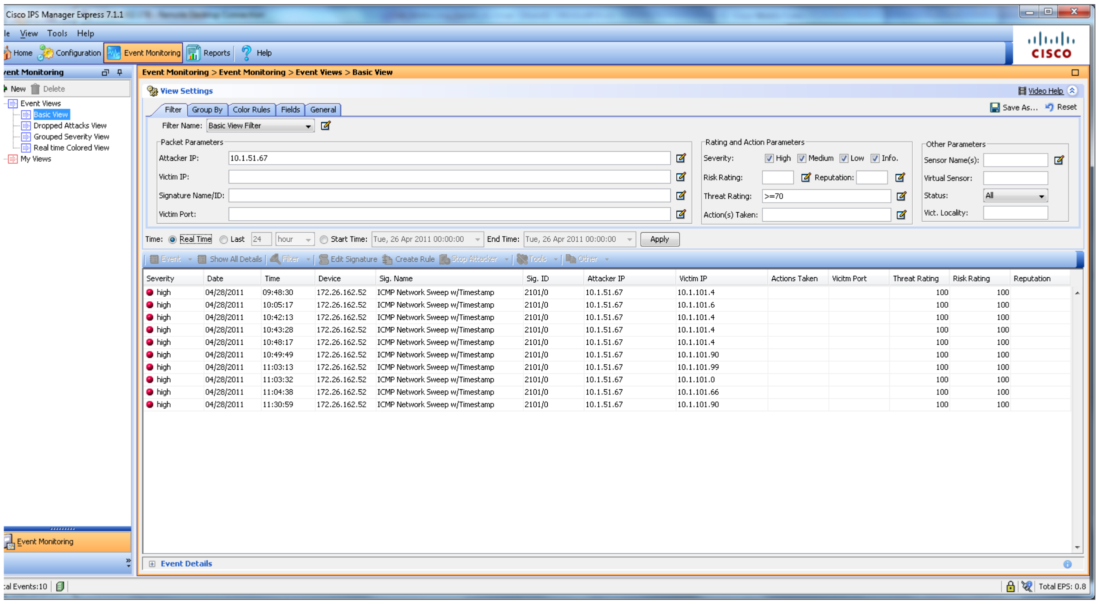

Cisco IPS Manager Express 7.1.1
File View Tools Help
Home Configuration Event Monitoring Reports Help
Event Monitoring
New Delete
Event Views
Basic View
Dropped Attacks View
Grouped Severity View
Real time Colored View
My Views
Event Monitoring
Total Events:10
Event Monitoring > Event Monitoring > Event Views > Basic View
View Settings
Video Help
Filter Group By Color Rules Fields General
Save As... Reset
Event Show All Details Filter Edit Signature Create Rule Stop Attacker Tools Other
Event Details
Total EPS: 0.8
291543

Filter Name: Basic View Filter

Packet Parameters
Attacker IP: 10.1.51.67
Victim IP:
Signature Name/ID:
Victim Port:

Rating and Action Parameters
Severity: ☑ High ☑ Medium ☑ Low ☑ Info.
Risk Rating: Reputation:
Threat Rating: >=70
Action(s) Taken:

Other Parameters
Sensor Name(s):
Virtual Sensor:
Status: All
Vict. Locality:

Time: ● Real Time ○ Last 24 hour ○ Start Time: Tue, 26 Apr 2011 00:00:00 End Time: Tue, 26 Apr 2011 00:00:00 Apply

Severity	Date	Time	Device	Sig. Name	Sig. ID	Attacker IP	Victim IP	Actions Taken	Vicitm Port	Threat Rating	Risk Rating	Reputation
● high	04/28/2011	09:48:30	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.4			100	100	
● high	04/28/2011	10:05:17	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.6			100	100	
● high	04/28/2011	10:42:13	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.4			100	100	
● high	04/28/2011	10:43:28	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.4			100	100	
● high	04/28/2011	10:48:17	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.4			100	100	
● high	04/28/2011	10:49:49	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.90			100	100	
● high	04/28/2011	11:03:13	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.99			100	100	
● high	04/28/2011	11:03:32	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.0			100	100	
● high	04/28/2011	11:04:38	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.66			100	100	
● high	04/28/2011	11:30:59	172.26.162.52	ICMP Network Sweep w/Timestamp	2101/0	10.1.51.67	10.1.101.90			100	100	

vShield could also be used to mitigate suspicious traffic. The Flow Monitoring pane displays the various traffic flows for each VLAN. In Figure 81, traffic flows for the tena2-cas1 VM are shown.

Figure 81 Monitoring Specific Flow Statistics with vShield

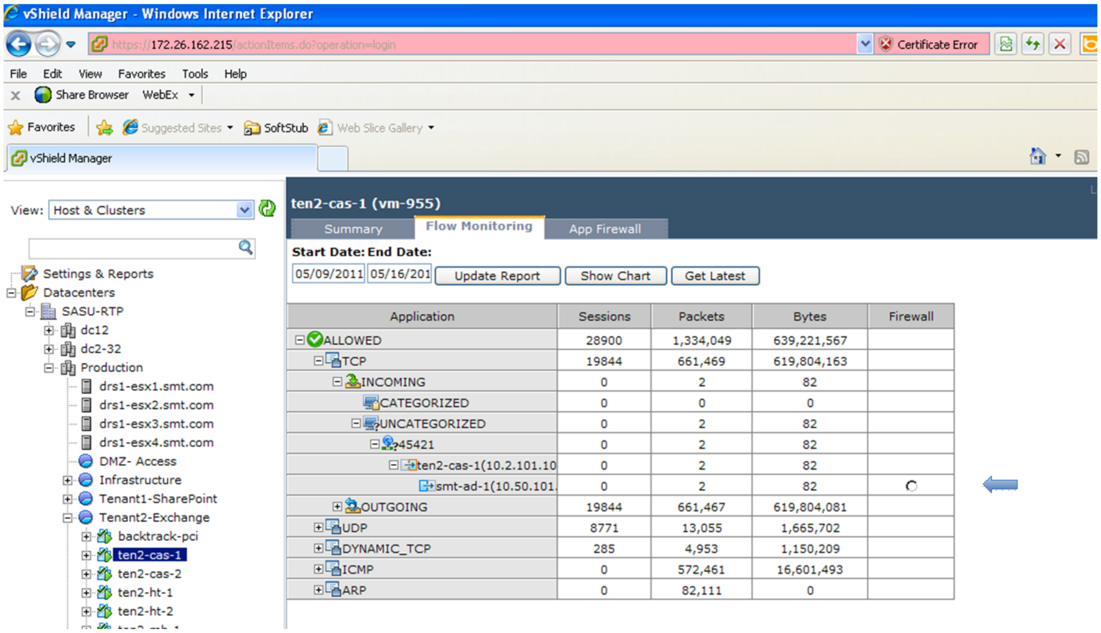

Start Date: End Date:
05/09/2011 05/16/201

Application	Sessions	Packets	Bytes	Firewall
⊟ ALLOWED	28900	1,334,049	639,221,567	
⊟ TCP	19844	661,469	619,804,163	
⊟ INCOMING	0	2	82	
CATEGORIZED	0	0	0	
⊟ UNCATEGORIZED	0	2	82	
⊟ 45421	0	2	82	
⊟ ten2-cas-1(10.2.101.10	0	2	82	
smt-ad-1(10.50.101.	0	2	82	○
⊞ OUTGOING	19844	661,467	619,804,081	
⊞ UDP	8771	13,055	1,665,702	
⊞ DYNAMIC_TCP	285	4,953	1,150,209	
⊞ ICMP	0	572,461	16,601,493	
⊞ ARP	0	82,111	0	

If one of the traffic flows is deemed to be suspicious, one can create a rule automatically by clicking the flow in the Firewall column. This takes the user to the App Firewall pane with the rule automatically created, as shown in Figure 82 for a traffic flow with destination port 45421.

Figure 82 Configuration Window to Stop Malicious Traffic

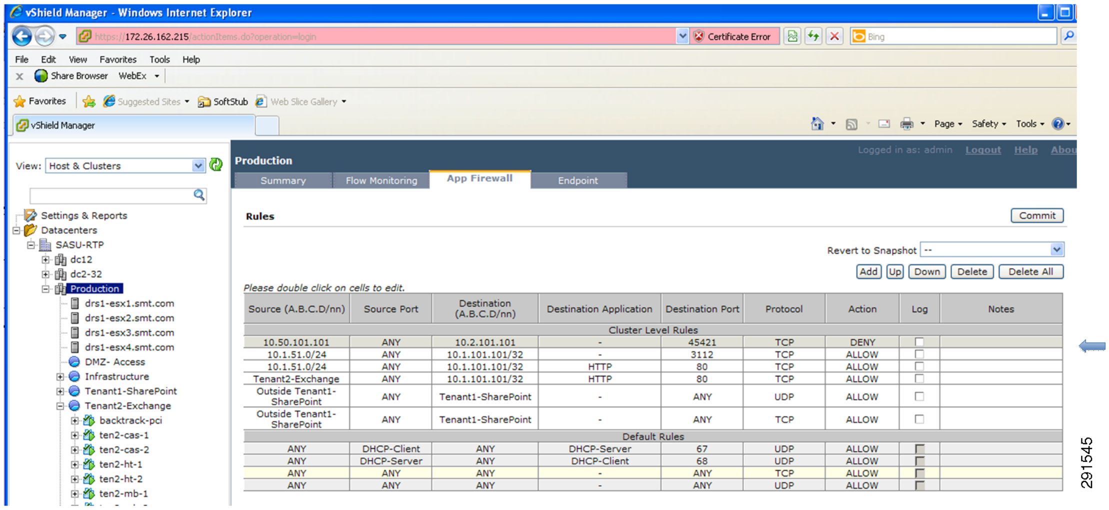

Please double click on cells to edit.

Source (A.B.C.D/nn)	Source Port	Destination (A.B.C.D/nn)	Destination Application	Destination Port	Protocol	Action	Log	Notes
Cluster Level Rules								
10.50.101.101	ANY	10.2.101.101	-	45421	TCP	DENY	□	
10.1.51.0/24	ANY	10.1.101.101/32	-	3112	TCP	ALLOW	□	
10.1.51.0/24	ANY	10.1.101.101/32	HTTP	80	TCP	ALLOW	□	
Tenant2-Exchange	ANY	10.1.101.101/32	HTTP	80	TCP	ALLOW	□	
Outside Tenant1-SharePoint	ANY	Tenant1-SharePoint	-	ANY	UDP	ALLOW	□	
Outside Tenant1-SharePoint	ANY	Tenant1-SharePoint	-	ANY	TCP	ALLOW	□	
Default Rules								
ANY	DHCP-Client	ANY	DHCP-Server	67	UDP	ALLOW	□	
ANY	DHCP-Server	ANY	DHCP-Client	68	UDP	ALLOW	□	
ANY	ANY	ANY	-	ANY	TCP	ALLOW	□	
ANY	ANY	ANY	-	ANY	UDP	ALLOW	□	

Appendix A—PCI Compliance

The Payment Card Industry Data Security Standard (PCI-DSS) defines a set of requirements for the configuration, operation, and security of payment card transactions in your business. If you use ESMT to store, process, or transmit payment card information, these standards and this guide apply to you.

Failure to comply with these standards can result in significant fines should a security breach occur. For more details about PCI DSS, see: https://www.pcisecuritystandards.org/security_standards/pci_dss.shtml.

Summary of PCI DSS Requirements

The following summary provides a basic overview of the PCI DSS requirements and shows how they apply to your business and to the ESMT architecture.

Requirement 1—Install and Maintain a Firewall Configuration to Protect Cardholder Environment

What the Requirement Says:

Firewalls are devices that control computer traffic allowed between an entity’s networks (internal) and untrusted networks (external), as well as traffic into and out of more sensitive areas within an entity’s internal trusted networks. The cardholder data environment is an example of a more sensitive area within an entity’s trusted network.

A firewall examines all network traffic and blocks those transmissions that do not meet the specified security criteria.

All systems must be protected from unauthorized access from untrusted networks, whether entering the system through the Internet as e-commerce, employee Internet access through desktop browsers, employee e-mail access, dedicated connections such as business-to-business connections, through wireless networks, or through other sources. Often, seemingly insignificant paths to and from untrusted networks can provide unprotected pathways into key systems. Firewalls are a key protection mechanism for any computer network.

Other system components may provide firewall functionality, provided they meet the minimum requirements for firewalls as provided in Requirement 1. When other system components are used within the cardholder data environment to provide firewall functionality, these devices must be included within the scope and assessment of Requirement 1.

How the ESMT Architecture and Components Meet These Requirements:

The ESMT architecture includes the ability to completely segment systems, not only between but also within tenants. At the physical layer, physical firewall appliances or modules provide the functionality to segment different tenants. In addition, these firewalls provide a strong line of defense against outside attacks on the data center infrastructure and unauthorized access to cardholder data. Within the access layer in the ESMT architecture, virtual firewalls are implemented to segment traffic within the same tenant environment (sometimes referred to as east-west traffic, or optionally across tenants as well). Figure 61 in Deploying vShield—Implementing Security Rules at the Access Layer (shown in Figure 83 for reference) shows how the various traffic flows are protected by virtual and physical firewalls.

Figure 83 Traffic Flows within a Multi-Tenant Architecture

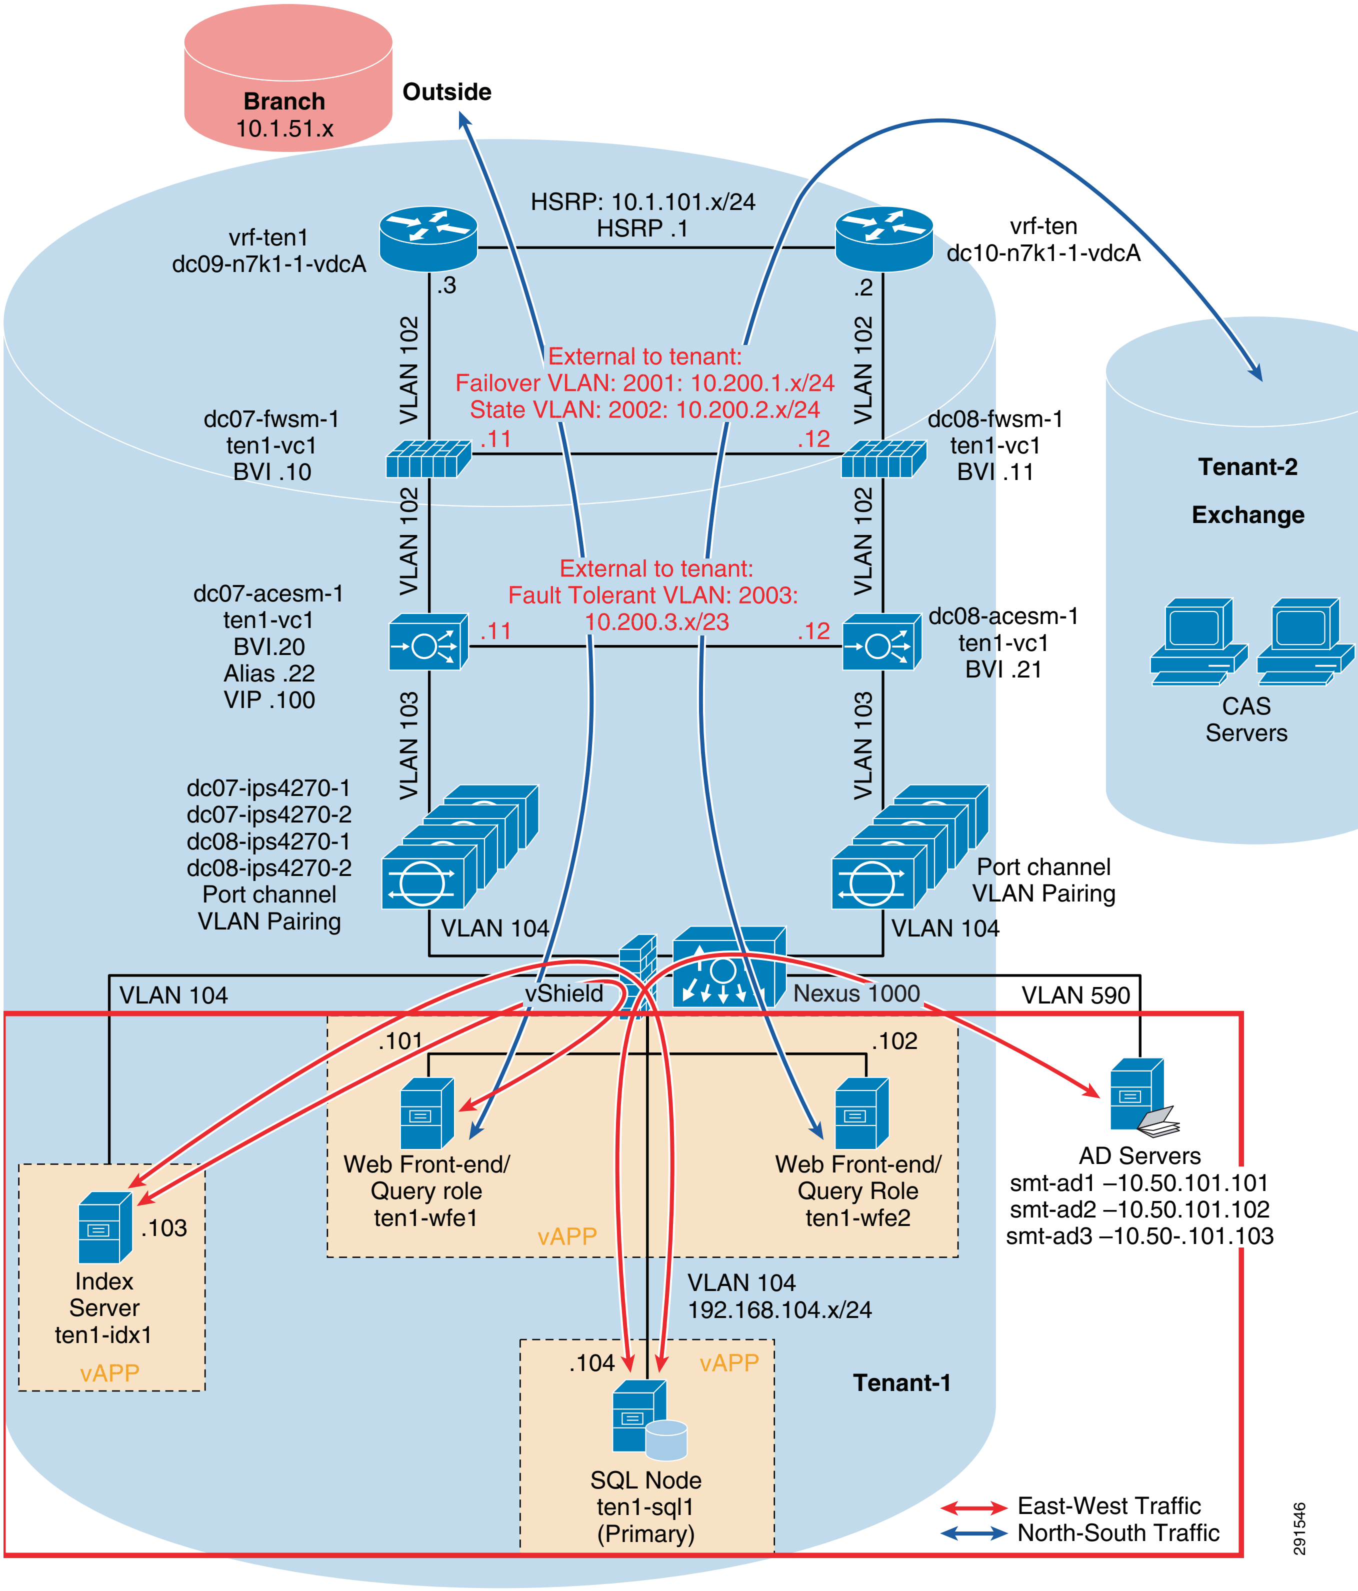

Firewall rules can easily be set to deny or authorize hosts based on IP address or specific VLANS. The firewall can also lock down traffic to certain ports, as shown in Deploying vShield—Implementing Security Rules at the Access Layer.

Example:

In Figure 83, the Tenant 1 environment is composed of the Common Desktop Environment (CDE, also known as Payment Card Industry [PCI] environment), while the Tenant 2 environment is composed of another tenant that the client wants to completely segment from the CDE. The client also wants to allow only certain connections from the branch environment into the PCI zone.

This PCI environment (Tenant 1) is composed of the following systems:

- Two Front-End Web servers
- SQL server
- Index server

The Tenant 2 environment is composed of the:

- Microsoft Exchange Environment

One method to secure the CDE from threats from outside is to grant access only to predefined hosts and to lock down traffic to certain ports. To accomplish this, follow the procedure in Deploying vShield—Implementing Security Rules at the Access Layer.

Requirement 2—Do Not Use Vendor-Supplied Defaults for System Passwords and other Security Parameters

What the Requirement Says:

Malicious individuals (external and internal to an entity) often use vendor default passwords and other vendor default settings to compromise systems. These passwords and settings are well known by hacker communities and are easily determined through public information.

How the ESMT Architecture and Components Meet these Requirements:

The Cisco ESMT architecture is composed of many different foundational components. In order to change vendor supplied defaults and harden systems in accordance with security best practices follow the links referred to below.

Note This does not include any operating system (that is, Windows, Linux, and so on), database (Oracle, MSSQL, and so on), or other highly customizable organizational component. For these items, it is the client’s responsibility to ensure that vendor-supplied defaults are changed and that systems are hardened in accordance with security best practices (that is, disabling telnet, and so on) before being implemented within the PCI environment.

- Cisco Nexus 5000:
  - Password-Initial Setup—Sections 3-7
  - Disable insecure services (for example, telnet)—Chapter 19

  http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/Nexus5000-NX-OS-ConfigurationGuide.pdf
- Cisco Nexus 7000:
  - Password/RBac-Initial Setup—Chapter 8

    - Disable insecure services (for example, telnet)—Chapter 6

  http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_2/nx-os/security/configuration/guide/Cisco_Nexus_7000_NX-OS_Security_Configuration_Guide__Release_4.2.pdf

- Cisco UCS Manager:
    - Password-Initial Setup—Chapter 9
    - Disable insecure services (for example, telnet)—Chapter 6

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/1.4.1/b_CLI_Config_Guide_1_4_1.pdf

- Cisco DC Network Manager Authentication Settings—Chapter 5

  http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_2/dcnm/fundamentals/configuration/guide/fundamentals.pdf

- NetApp Storage Controller

  https://now.netapp.com/NOW/knowledge/docs/ontap/rel7351/pdfs/ontap/sysadmin.pdf

- vShield Configuration:
    - Password-Initial Setup—page 22

  http://www.vmware.com/pdf/vsz_10_admin.pdf

- Cisco Nexus 1000V:
    - Password -Initial Setup—Chapter 2
    - Disable insecure services (for example, telnet)—Chapter 7

  http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/security/configuration/guide/n1000v_security.pdf

- ACE Load Balancing Module
    - Password/Console Setting - Initial Setup—Chapter 1

  http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/vA2_3_0/configuration/administration/guide/ace_adgd.pdf

- C6500 Services Chassis and Telnet

  http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/swcgIX.html

- IPS Configuration

  http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_cfg_sec_4cl.html

Wireless is not a part of the proposed architecture, but could be if the client needs it. In this case, the client is responsible for ensuring that the wireless network is using strong authentication/encryption protocols (e.g., WPA, WPA2) and that only services that are needed are allowed into the CDE environment. Steps provided in requirement 1 would be used to segment any wireless network from the rest of the CDE environment.

Requirement 3—Protect Stored Cardholder Data

What the Requirement Says:

Protection methods such as encryption, truncation, masking, and hashing are critical components of cardholder data protection. If an intruder circumvents other security controls and gains access to encrypted data, without the proper cryptographic keys, the data is unreadable and unusable to that person. Other effective methods of protecting stored data should be considered as potential risk mitigation opportunities. For example, methods for minimizing risk include not storing cardholder data unless absolutely necessary, truncating cardholder data if full PAN is not needed, and not sending unprotected PANs using end user messaging technologies, such as e-mail and instant messaging.

How the ESMT Architecture and Components Meet these Requirements:

Requirement 3 is the responsibility of the client. The Cisco ESMT architecture does not interfere with the client's ability to securely store cardholder data. The client has the option to encrypt data before it is stored. The network connectivity between the application and storage is segmented and inaccessible to unauthorized users. All key handling and management procedures outlined in Requirement 3 are the sole responsibility of the client.

Requirement 4—Encrypt Transmission of Cardholder Data Across Open, Public Networks

What the Requirement Says:

Sensitive information must be encrypted during transmission over networks that are easily accessed by malicious individuals. Misconfigured wireless networks and vulnerabilities in legacy encryption and authentication protocols continue to be targets of malicious individuals who exploit these vulnerabilities to gain privileged access to cardholder data environments.

How the ESMT Architecture and Components Meet these Requirements:

Requirement 4 is the responsibility of the client. The Cisco ESMT architecture does not prohibit the client's ability to securely transmit cardholder information over public networks. The client has the option of using IPsec, SSL, or another secure method to provide secure connectivity over public networks. IPsec tunnels can be safely terminated within the corporate network or VPN technologies can provide data integrity over a public infrastructure. In either case cardholder transactions over public networks can be encrypted in accordance with PCI. For more details on how to implement IPsec and SSL termination, see the references below.

The links to implement encryption over the public network are:

- http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/DMVPNbk.pdf
- http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/Dir_Encap.html

Requirement 5—Use and Regularly Update Anti-Virus Programs

What the Requirement Says:

Malicious software, commonly referred to as malware—including viruses, worms, and Trojans—enters the network during many business-approved activities including employee e-mail and use of the Internet, mobile computers, and storage devices, resulting in the exploitation of system vulnerabilities. Anti-virus software must be used on all systems commonly affected by malware to protect systems from current and evolving malicious software threats.

How the ESMT Architecture and Components Meet These Requirements:

Requirement 5 is the responsibility of the client. The Cisco ESMT architecture does not interfere with the client's ability to regularly update anti-virus software on the systems they deploy in this environment.

Note The client should refer to anti-virus specific documentation for updating and managing anti-virus software on PCI in-scope systems.

Requirement 6—Develop and Maintain Secure Systems and Applications

What the Requirement Says:

Unscrupulous individuals use security vulnerabilities to gain privileged access to systems. Many of these vulnerabilities are fixed by vendor-provided security patches, which must be installed by the entities that manage the systems. All critical systems must have the most recently released, appropriate software patches to protect against exploitation and compromise of cardholder data by malicious individuals and malicious software.

How the ESMT Architecture and Components Meet These Requirements:

The client has the ability to use VMware update manager (VUM) for updating and patching Vcenter system software and guest OS operating systems. The steps can be found at: http://www.vmware.com/pdf/vi3_vum_10_admin_guide.pdf

Any guest operating system that is not supported by VUM must be patched either by another automated means (WSUS, commercial product, etc.) or manually. The Cisco ESMT architecture can be configured to allow for VM s to access a repository of updated images and patches. Network devices can be updated by using secure TFTP/FTP server within the ESMT environment. Relevant inks are provided below.

- Cisco Nexus 5000:
  - Upgrading Switch—Sections 3-4

  http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/Nexus5000-NX-OS-ConfigurationGuide.pdf
- Cisco Nexus 7000:
  - Upgrading Switch guide

  http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_1/nx-os/upgrade/guide/nx-os_upgrade.html
- Cisco Nexus 1000V:
  - Upgrading Switch guide

  http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/upgrade/software/guide/n1000v_upgrade_software.html
- VMware vShield:
  - Upgrading Switch guide—chapter 5

  http://www.vmware.com/pdf/vsz_10_admin.pdf
- ACE Load Balancing Module
  - Upgrading software—Appendix A

  http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/vA2_3_0/configuration/administration/guide/ace_adgd.pdf

- IPS Intrusion Prevention
  - Upgrading software

  http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_system_images.html
- VMware vCenter
  - Upgrading VMware vCenter

  http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1022104

Note The Cisco ESMT was validated using the latest software releases and patches from all vendors.

Requirement 7—Restrict Access to Cardholder Data by Business Need to Know

What the Requirement Says:

To limit access to critical data to authorized personnel only, systems and processes must be in place to limit access based on need to know and according to job responsibilities.

How the ESMT Architecture and Components Meet These Requirements:

Requirement 7 is the responsibility of the client. The Cisco ESMT architecture does not interfere with allowing only those individuals who require such access into the CDE. Cisco Secure Access Control Server (ACS) and Windows Active Directory provide excellent capabilities to log and document any user who accesses devices within the CDE, whether it is VMware's virtual software components or networking and storage devices. the ESMT architecture allows secure connectivity between these devices and ACS/AD servers. Group policies can be created within ACS and AD configurations that can also restrict user access privileges based on the individual's user ID. In addition, group attributes, as well as incorporating tenant-id and global-id or for tenant and cloud system administrators, can be used.

Note If the client chooses to use Active Directory, it is the client's responsibility to ensure that the deployment and management of AD is done in a PCI-compliant manner. For more information about Active Directory configuration best practices, see: http://www.microsoft.com/windowsserver2008/en/us/ad-main.aspx.

Requirement 8—Assign a Unique ID to Each Person with Computer Access

What the Requirement Says:

Assigning a unique identification (ID) to each person with access makes each individual uniquely accountable for his or her actions. When such accountability is in place, actions taken on critical data and systems are performed by, and can be traced to, known and authorized users.

How the ESMT Architecture and Components Meet These Requirements:

Requirement 8 is the responsibility of the client. The client must ensure that all operating systems, applications, and network devices attached to the PCI network are accessed using unique IDs. For Cisco-specific network devices, see Requirement 2 for links to manuals explaining how to assign unique user IDs. All passwords should meet the following requirements:

- Create unique first-time passwords
- Immediately revoke access for any terminated users

- Remove/disable inactive user accounts at least every 90 days
- Do not use group, shared, or generic accounts and passwords
- Change user passwords at least every 90 days
- Require a minimum password length of at least seven characters
- Use passwords containing both numeric and alphabetic characters
- Do not allow an individual to submit a new password that is the same as any of the last four passwords
- Limit repeated access attempts by locking out the user ID after not more than six attempts
- Set the lockout duration to a minimum of 30 minutes or until an administrator has enabled the user ID
- If a session has been idle for more than 15 minutes, require the user to re-authenticate to re-activate the terminal or session

Note If the client chooses to use Active Directory, it is the client's responsibility to ensure that the deployment and management of AD is done in a PCI compliant manner. For more information about Active Directory configuration best practices, see: http://www.microsoft.com/windowsserver2008/en/us/ad-main.aspx.

Requirement 9—Restrict Physical Access to Cardholder Data

What the Requirement Says:

Any physical access to data or systems that house cardholder data provides the opportunity for individuals to access devices or data and to remove systems or hardcopies and should be appropriately restricted. For the purposes of Requirement 9, "onsite personnel" refers to full-time and part-time employees, temporary employees, contractors, and consultants who are physically present on the entity's premises. A "visitor" refers to a vendor, guest of any onsite personnel, service workers, or anyone who needs to enter the facility for a short duration, usually not more than one day. "Media" refers to all paper and electronic media containing cardholder data.

How the ESMT Architecture and Components Meet These Requirements:

Requirement 9 is the responsibility of the client. The Cisco ESMT architecture does not interfere with where the client chooses to implement this solution and how media backups are performed and handled.

Requirement 10—Track and Monitor All Access to Network Resources and Cardholder Data

What the Requirement Says:

Logging mechanisms and the ability to track user activities are critical in preventing, detecting, or minimizing the impact of a data compromise. The presence of logs in all environments allows thorough tracking, alerting, and analysis when something does go wrong. Determining the cause of a compromise is very difficult, if not impossible, without system activity logs.

How the ESMT Architecture and Components Meet These Requirements:

All Cisco device events are forwarded to a syslog server for logging purposes. These logs must be stored for a year based on PCI-DSS requirements. The ESMT architecture incorporates a tiered syslog solution in which each tenant incorporates its own syslog server and events from all tenants are forwarded to a central syslog repository server. For details on implementing syslog servers, see this general link for guidance on configuring syslog on Cisco networking devices: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/Baseline_Security/securebasebook.html.

Requirement 11—Regularly Test Security Systems and Processes

What the Requirement Says:

Vulnerabilities are being discovered continually by malicious individuals and researchers and being introduced by new software. System components, processes, and custom software should be tested frequently to ensure security controls continue to reflect a changing environment.

How the ESMT Architecture and Components Meet These Requirements:

Requirement 11 of the PCI DSS is the responsibility of the client. The Cisco ESMT architecture does not interfere with the client's ability to regularly test CDE systems. ESMT architecture incorporates the use of Cisco Intrusion Prevention systems (IPS) which provides the capability of detecting Denial of Service (DOS) attacks and attacks in the form of malware, worms, and Trojans. Figure 84 is a screenshot for the Cisco IPS interface.

Figure 84 Cisco IPS Interface

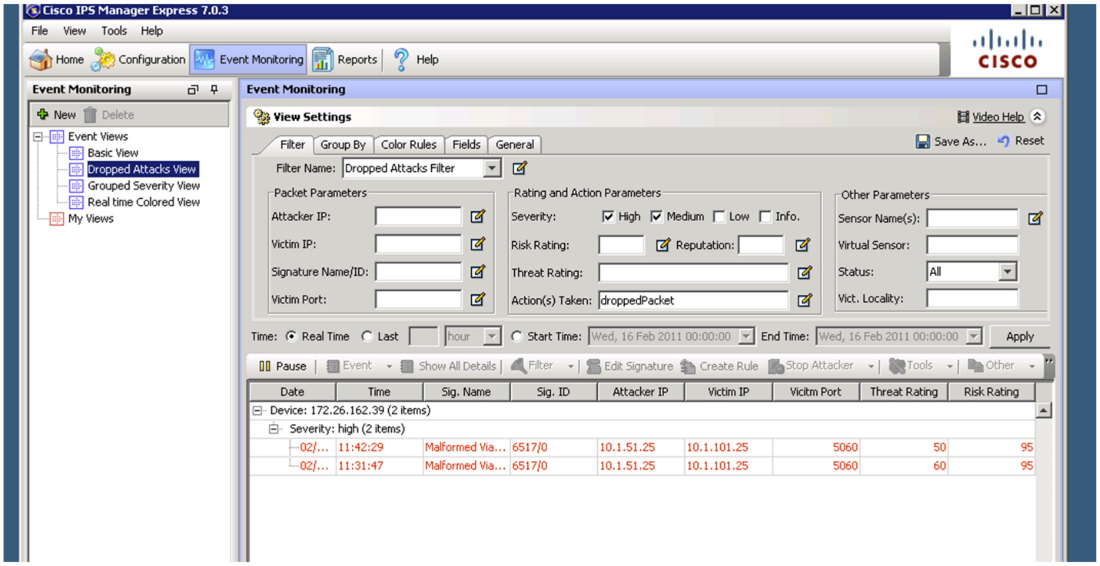

Filter Name: Dropped Attacks Filter

Packet Parameters

Attacker IP:

Victim IP:

Signature Name/ID:

Victim Port:

Rating and Action Parameters

Severity: ☑ High ☑ Medium ☐ Low ☐ Info.

Risk Rating: Reputation:

Threat Rating:

Action(s) Taken: droppedPacket

Other Parameters

Sensor Name(s):

Virtual Sensor:

Status: All

Vict. Locality:

Time: ● Real Time ○ Last hour ○ Start Time: Wed, 16 Feb 2011 00:00:00 End Time: Wed, 16 Feb 2011 00:00:00

Date	Time	Sig. Name	Sig. ID	Attacker IP	Victim IP	Vicitm Port	Threat Rating	Risk Rating
Device: 172.26.162.39 (2 items)								
Severity: high (2 items)								
02/...	11:42:29	Malformed Via...	6517/0	10.1.51.25	10.1.101.25	5060	50	95
02/...	11:31:47	Malformed Via...	6517/0	10.1.51.25	10.1.101.25	5060	60	95

For the steps that can be taken to mitigate attacks using the Cisco IPS, see Cisco Intrusion Prevention System Monitoring and Mitigation.

Requirement 12—Maintain a Policy that Addresses Information Security for All Personnel

What the Requirement Says:

A strong security policy sets the security tone for the whole entity and informs personnel about what is expected of them. All personnel should be aware of the sensitivity of data and their responsibilities for protecting it. For the purposes of Requirement 12, "personnel" refers to full-time and part-time employees, temporary employees, contractors, and consultants who are resident on the entity's site or otherwise have access to the cardholder data environment.

How the ESMT Architecture and Components Meet These Requirements:

Requirement 12 is the responsibility of the client. All policies and procedures with regards to PCI should be created and disseminated by the client.

Appendix B—ICSA Labs

ICSA Labs is an independent auditor that has been providing credible, independent, third-party product assurance for end users and enterprises for the last 20 years. ICSA Labs was invited to perform an audit of SMT built on FlexPod for VMware on behalf of Cisco, NetApp, and VMware. ICSA Labs performed an extensive audit of the SMT architecture in a laboratory at NetApp in RTP, NC.

The audit specifically included sections on

- Management
- Encrypted password authentication
- Role-based access controls (RBAC)
- Persistence testing (unexpected failure)
- Password authentication
- Centralized logging
- Protocol
- Security assessment

The architecture was also subjected to several direct attacks, such as the Targa 3, Denial of Service, SYNFlood, UDPFlood, and Targa 2.

ICSA labs confirmed that SMT built on FlexPod for VMware is indeed a secure architecture. The final report can be accessed at: https://www.icsalabs.com/evaluations.