VMDC Architecture with Citrix NetScaler VPX and SDX

This document describes design recommendations, configurations and validation results for utilizing Citrix NetScaler VPX and NetScaler SDX load-balancing appliances in the Cisco Virtual Multiservice Data Center (VMDC) 2.2/2.3 and VMDC 3.0 architectural models. VMDC is the Cisco reference architecture for IaaS Cloud deployments. Citrix offers a range of appliances designed to perform server load balancing (SLB) and offloading for certain applications. These appliances are offered in both physical (MPX, SDX) and virtual (VPX) form factors.

The goal of this document is to provide guidance for Cisco field and service teams using NetScaler VPX and NetScaler SDX in cloud designs and data center solutions for Cisco customers. These recommendations and results were gathered while evaluating VPX in a VMware environment, and on a SDX multi-tenant appliance. Most of the validation was performed using the publicly available NetScaler OS 10.0 version, while some validation was performed on an early beta of NetScaler OS 10.1. The beta code became publicly available on May 31, 2013.

This document also serves as a guide for replacing Cisco ACE modules or appliances in VMDC architectures. The following list shows where readers can find guidance in this document for configuring specific commonly-applied ACE server load balancing features on the NetScaler appliance:

- For high availability (active/active or active/standby), see High Availability, page -11
- For virtual MAC (VMAC), see High Availability with High Availability with VMAC, page -16
- For SLB, see VPX Installation on VMware Hypervisor, page -9
- For server health monitoring, see Server Health Check, page -10
- For SSL offload, see SSL Offload, page -10
- For access control lists (ACLs), see L4-L7 Access Control Lists, page -10
- For virtual contexts, see SDX Evaluation in VMDC, page -14

Product screen shots and similar materials in this document are used for illustrative purposes only and may show trademarks of VMware, Inc. and Citrix. All other marks and names mentioned herein may be trademarks of their respective companies.

Use of the word “partner” or “partnership” does not imply a legal partner relationship between Cisco and any other company.

NetScaler VPX and SDX Overview

The Citrix NetScaler products offer SLB and content switching, along with application acceleration Layer 4 - Layer 7 (L4 - L7) traffic management, data compression, Secure Socket Layer (SSL) acceleration, network optimization, and application security. For the purposes of this validation and documentation, the focus was on using SLB, SSL Offload and related features on the NetScaler, as a means of replacing the ACE SLB functionality in VMDC designs.

NetScaler VPX has the same features as the NetScaler MPX physical appliance, but is a virtual form of the NetScaler product. VPX is installed as a virtual machine (VM) on a hypervisor. Currently, VPX can be installed on XenServer, VMware, and Hyper-V. The NetScaler VPX can handle up to 3 Gbps of HTTP traffic when deployed on VMware or Hyper-V. Performance is based on the type of license installed on the VPX instance. For more information about performance for various VPX instances, refer to:

http://www.citrix.com/products/netscaler-application-delivery-controller/features/platforms/vpx.html

VPX on VMware was used for this validation, so installation, configuration requirements, and examples refer to the VMware hypervisor unless otherwise noted.

NetScaler SDX is the multi-tenant NetScaler appliance (Figure 1).Multiple fully isolated, fully independent NetScaler instances can run on a single NetScaler SDX device. The SDX appliance comes with 10 Gbps Ethernet (10GE) and 1 Gbps Ethernet (1GE) ports (type and number of ports depends on the SDX model) that can form an EtherChannel bundle, which is desirable for an appliance-based service design in the VMDC architecture. This evaluation used the SDX 20500, which provides four 10GE ports and eight 1GE ports. This model also has 16 SSL cores to handle SSL hardware acceleration. The SDX 20500 can support up to 20 NetScaler instances.

Code versions earlier than NetScaler 10.1 do not support sharing EtherChannel among multiple NetScaler instances on SDX. Because EtherChannel sharing is a requirement for VMDC appliance deployments, testing was done using a beta of version 10.1. Throughput capacity depends upon which SDX platform is used, and which license is installed on the appliance. Using the SDX 20500, a single NetScaler instance can handle up to 18 Gbps of HTTP traffic. The SDX 20500 appliance can handle an aggregate throughput of 42 Gbps of HTTP traffic (across multiple NetScaler instances).

Figure 1 Citrix SDX Appliance

For more information about the SDX, refer to:

http://www.citrix.com/products/netscaler-application-delivery-controller/features/platforms/sdx.html

Both VPX and SDX can be configured using CLI or a browser-based interface. CLI configurations are used as examples in this technical paper.

NetScaler OS uses three different types of IP addresses:

- **NetScaler IP (NSIP)**—This IP address refers to the management address. Although a separate management subnet is not required, this design was used for this evaluation.
- **Subnet IP (SNIP)**—This IP address is a subnet IP that represents an interface the NetScaler device uses to pass traffic between the server farm and clients. By default, this address type is used as the source IP address for Source-NATing packets from outside clients.
- **Virtual IP (VIP)**—Outside clients use this shared address to connect to the server farm.
- **Mapped IP (MIP)**—Server-side connections use MIP addresses when no SNIP is configured, and when the USNIP option is disabled (by default, this option is enabled).

For more information about NetScaler IP addressing, refer to:

http://support.citrix.com/proddocs/topic/ns-system-10-map/ns-nw-ipaddrssng-confrng-ns-ownd-ip-address-con.html

VMDC Overview

VMDC is the Cisco reference architecture for IaaS cloud deployments, and there have been multiple VMDC designs as platforms and technologies evolve. The VMDC 2.x architecture is VRF-Lite and Virtual Port Channel (vPC) based, and VMDC 2.2 is the large-scale version of this architecture, with VMDC 2.3 being an optimized and smaller-footprint version. The VMDC 3.x architecture is VRF-Lite and FabricPath based.

The VMDC IaaS cloud architecture is designed around a set of modular DC components comprised of building blocks of resources called pods, which comprise:

- Cisco Unified Computing System (UCS)
- Storage area network (SAN) and network attached storage (NAS) storage arrays
- Access (switching) layers and aggregation (switching and routing) layers connecting to the Data Center Services Node (DSN) or Appliance based services layer
- Multiple 10 GE fabrics using highly scalable Cisco network switches and routers

VMDC is built around Cisco UCS, Cisco Nexus 1000V, Nexus 5000 and Nexus 7000 switches, Cisco Multilayer Director Switch (MDS), Cisco Aggregation Services Router (ASR) 9000, ASR 1000, Cisco Adaptive Security Appliance (ASA) 5585-X or ASA Services Module (ASASM), Cisco Catalyst 6500 DSN, Cisco ACE, Nexus 1000V Virtual Security Gateway (VSG), VMware vSphere, EMC VMAX, and NetApp FAS storage arrays. Cloud service orchestration is currently provided by the BMC Cloud Lifecycle Management (CLM) suite. Refer to VMDC system releases.

Note NetScaler products were not evaluated in the VMDC 2.3 architecture because of time and lab availability constraints. However, results and observations are expected to be similar to those for VMDC 2.2.

VMDC 2.2 Architecture

VMDC 2.2 uses a hierarchical network design for high availability and scalability. The hierarchical (layered) data center design uses redundant switches at each network layer in the network topology for device-level failover that creates a highly available transport between end nodes using the network. Modules populating a slot in a network switching node, or standalone service appliances, provide additional services such as SLB, firewall, and intrusion prevention.

Each service approach supports redundant hardware deployment to preserve the high availability standards set by the network topology. This layered approach, the basic foundation of the VMDC design, provides scalability, performance, flexibility, resiliency, and service assurance.

Virtual LANs (VLANs) and virtual routing and forwarding (VRF) instances provide tenant isolation in the VMDC architecture. Routing protocols in the VRF instances interconnect the various networking and service devices. The VMDC 2.2 architecture is based on a VRF-Lite design end-to-end through the data center routing platforms, with routing between the VRF instances provided by BGP peering.

Note For detailed descriptions of the VMDC 2.2 architecture, refer to the following documents:
VMDC 2.2 Design Guide
VMDC 2.2 Implementation Guide

The layered VMDC 2.2 architecture comprises core, aggregation, services, and access layers. This architecture enables data center modules (pods and compute blocks) to be added as demand and load increases. The architecture also provides flexibility to create different logical topologies using device virtualization, service devices insertion, and traditional L2 and L3 network configurations. Figure 2 provides a logical representation of the VMDC 2.2 architecture, with the services layer comprised of the Catalyst 6500 DSN, ACE30, and ASA 5585-X (or ASASM).

In this design, the Load Balancing services are provided by the Cisco ACE30 module residing in the Cisco Catalyst 6500 Data Center Services Node (DSN). As described later in this document, the ACE30 can be replaced by the Citrix SDX multi-tenant appliance connecting to the Cisco Nexus 7000 Aggregation layer, or by the Citrix VPX per-tenant virtual appliance installed in the UCS compute layer.

Figure 2 VMDC 2.2 Logical Diagram

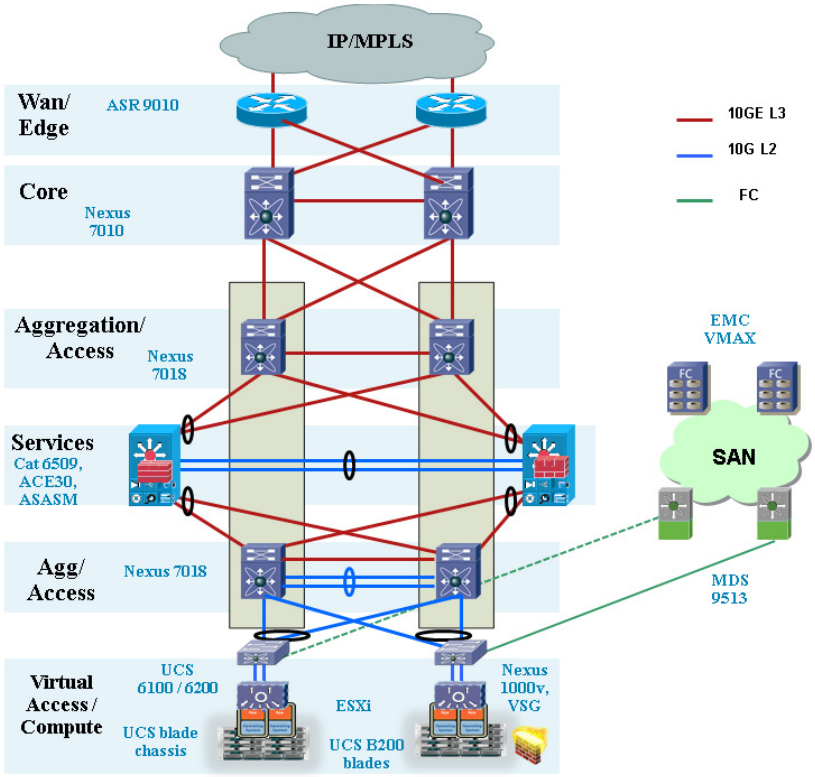

VMDC 2.3 Architecture

The VMDC 2.3 system leverages the end-to-end architecture defined in VMDC 2.2, with some optimizations in the platforms and tenancy models, to reduce the cost and footprint of the solution, while increasing the tenant scale. The VMDC 2.3 design comprises of WAN, aggregation and services layers, with the services being provided by appliances connecting directly to the aggregation layer. Figure 3 provides a logical representation of the VMDC 2.3 architecture, with the services layer comprised of the Cisco ACE 4710 and ASA 5585-X appliances connecting to the Cisco Nexus 7004 aggregation nodes.

Note For detailed descriptions of the VMDC 2.3 architecture, refer to the following documents:
VMDC 2.3 Design Guide
VMDC 2.3 Implementation Guide

In this design, the Load Balancing services are provided by the multi-tenant ACE 4710 appliance connecting through vPC to the Nexus 7004. As described later in this document, the ACE 4710 can be replaced by the Citrix SDX multi-tenant appliance connecting to the Nexus 7004 Aggregation layer, or by the Citrix VPX per-tenant virtual appliance installed in the UCS compute layer.

Figure 3 VMDC 2.3 Logical Diagram

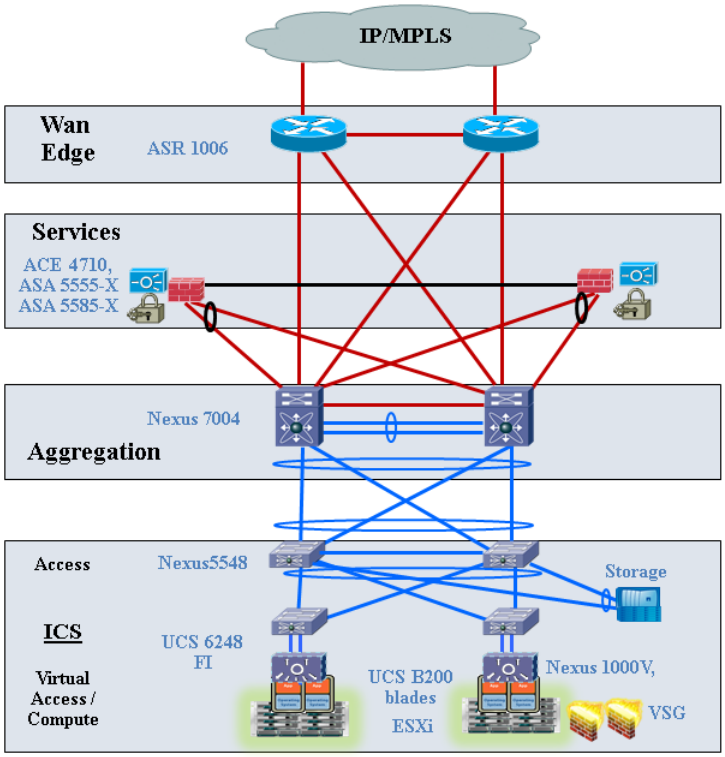

VMDC 3.0 Architecture

VMDC 3.0 introduces Cisco FabricPath as an optional L2 alternative to a hierarchical vPC-based design for the intra-data center network. FabricPath simplifies and expands L2 network design, removing the complexities of Spanning Tree Protocol (STP) and thus enabling more extensive, flexible, and scalable

L2 designs. Other VMDC releases will follow as Cisco develops and evolves FabricPath. While FabricPath comprises an improved L2 multipathing technology, vPC-based resiliency remains a valid option in the VMDC portfolio. Customers will continue to be able to choose between vPC-based and FabricPath designs to meet their requirements. This design modified only the Unified Fabric and Data Center Networking layer of the architecture, leveraging existing design guidance for Unified Computing and Integrated Systems (UCIS) and DCI layers.

Figure 4 provides a logical representation of the VMDC 3.0 Typical DC architecture, with FabricPath providing the L2 fabric within a VMDC Pod. The Pod consists of Nexus 7000 / 5500 switches as FabricPath leaf nodes, and Nexus 7000 as the FabricPath spine and L3 Aggregation node. The services within the pod are provided by the Cisco ACE 4710 and ASA 5585-X appliances connecting to the Cisco Nexus 7000 spine/aggregation nodes.

Note For detailed descriptions of the VMDC 3.0 architecture, refer to the following documents:
VMDC 3.0 Design Guide
VMDC 3.0 Implementation Guide

In this design, the Load Balancing services are provided by the multi-tenant ACE 4710 appliance connecting through vPC to the Nexus 7000. As described later in this document, the ACE 4710 can be replaced by the Citrix SDX multi-tenant appliance connecting to the Nexus 7000 Spine/Aggregation layer, or by the Citrix VPX per-tenant virtual appliance installed in the UCS compute layer.

Figure 4 VMDC 3.0 Typical Data Center Design

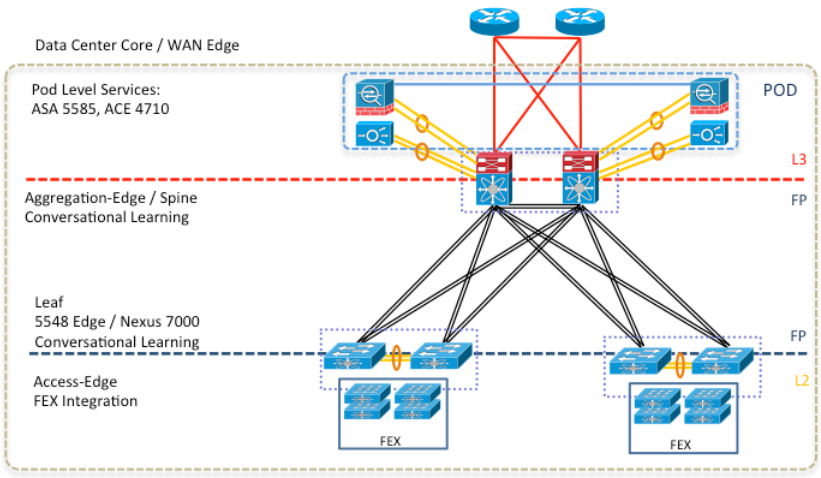

Note The NetScaler products were not evaluated in the VMDC 3.0.1 architecture. However, results and observations are expected to be similar to those seen with VMDC 3.0.

NetScaler VPX in VMDC

The NetScaler VPX supports application load balancing and optimization in the compute layer, at rates of up to 3 Gbps. It is inserted in the VMDC architecture in the compute layer, installed as a virtual machine (VM) on the VMware hypervisor. Each consumer or tenant can be assigned their own VPX

instance. Those requiring multiple server segments can be serviced by the same VPX, as it can be installed with one or more interfaces (one interface can be configured as a trunk, making it capable of servicing multiple server segments.

Alternatively, multiple VPX instances can be installed in a consumer container to accommodate each server segment. The VMDC design prescribes for the server load balancing to be done after security checks, therefore the VPX connection should exist on the inside network of any firewall within the container. In one-armed mode, this means that traffic between the load balancer and tenant VM's will no longer have to traverse the access layer or aggregation layer in the infrastructure, as this traffic needs only traverse the Nexus 1000V in the compute layer ()Figure 5).

Figure 5 VPX in VMDC 3.0 Palladium Network Container

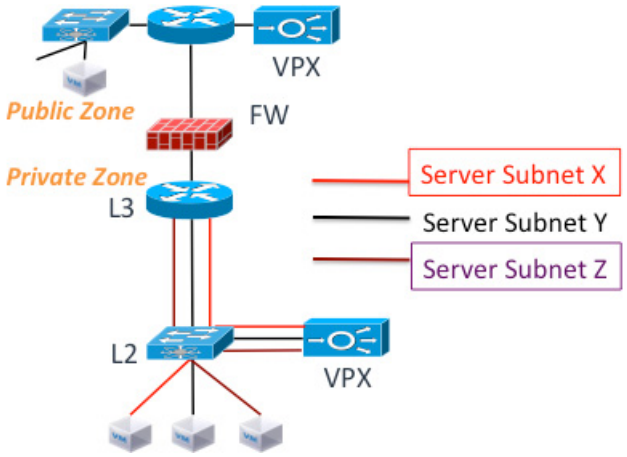

Note While this validation and document covers the use of Citrix VPX in a VMDC 3.0 based Palladium Network Container, the VPX can be used in a similar fashion in the VMDC 2.2 and 2.3 Gold and Silver Network Containers.

Evaluation Topology and Methodology

The VPX was inserted in a network container with one server segment, and deployed in one-armed mode as this is how SLB is typically configured in the VMDC reference design (Figure 6). The aggregation Nexus 7000 was used as the default gateway for the tenant VMs and the VPX.

Figure 6 Evaluation Logical Topology

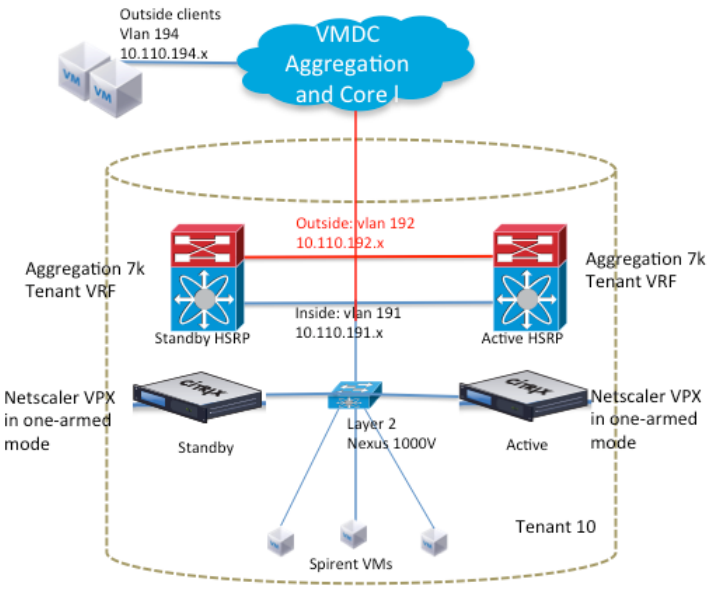

Installation requirements for the VPX VM on VMware include:

- 1 vCPU
- 2 GB RAM
- 20 GB HD
- At least one vNIC (a separate interface is not required for High Availability)

The VPX VM used for testing was allocated the following resources:

- 2 vCPU
- 2 GB RAM
- 20 GB HD
- 2 vNICs (one for management and one for data traffic)

The testing scope was similar to the Cisco ACE testing in VMDC. NetScaler VPX has many features, but only SLB and SSL acceleration were tested for this evaluation, which validated the following functions:

- L4 server load balancing
- L7 server load balancing
- Server Health Check
- L4-7 ACLs
- SSL Offload
- Syslog
- High availability
- IPv6
- Dynamic Routing

VPX Evaluation Observations and Deployment Considerations

The following sections provide detailed VPX evaluation observations and deployment considerations:

VPX Installation on VMware Hypervisor

As stated previously, the VPX is installed as a VM on the VMware hypervisor. The .ovf file and other VMware-related files can be obtained from the Citrix website with a valid myCitrix account. No license is needed to install the VM, but a license is required for operation. After the files are obtained, installation is performed using the .ovf file. Up to 10 interfaces, and more than one CPU, can be allocated to the VPX VM. See the previous section for installation requirements.

Console access through VMware must be used to configure out-of-band management access using the CLI. The default user name and password are nsroot. To configure management access, the following configuration is needed:

```
set ns config -IPAddress 172.26.162.225 -netmask 255.255.0.0
```

This IP address maps to the first port on the VPX VM. It must be connected to the management VLAN for OOB management to be possible. Once the NSIP is configured, all management access methods are available by default (SSH, telnet, HTTP, FTP). To configure the other data interface IP addresses, SSH access can be used to get to the CLI. IP addresses for data must be entered in the order they were created on the VM. The command to add a SNIP is as follows:

```
add ns ip 10.110.191.250 255.255.255.0 -vServer DISABLED -dynamicRouting ENABLED
```

Server Load Balancing

The VPX was tested as both an L4 (TCP) and L7 (HTTP) load balancer. It can be configured to use many different load balancing algorithms, such as least connections and round robin. The following configuration, which was used in the evaluation, uses a round robin algorithm:

```
add lb vserver vip-1-http-80 HTTP 10.110.191.200 80 -persistenceType NONE -lbMethod
ROUNDROBIN -cltTimeout 180 -icmpVsrResponse ACTIVE
```

In this example, the VIP uses port 80 for the IP 10.110.191.200.

To bind servers to this VIP, server objects must be created and then bound to a service that is eventually bound to the VIP:

```
add server server-1 10.110.191.100
add service service-1-http-80 server-1 HTTP 80 -gslb NONE -maxClient 0 -maxReq 0 -cip
DISABLED -usip NO -useproxyport YES -sp ON -cltTimeout 180 -svrTimeout 360
-CustomServerID "\"None\"" -CKA YES -TCPB YES -CMP NO
bind lb vserver vip-1-http-80 service-1-http-80
```

For layer 4 testing, the VIP is configured as a TCP VIP, while the service bound to the VIP is a layer 7 service.

```
add lb vserver vip-1-tcp TCP 10.110.191.200 * -persistenceType NONE -lbMethod
ROUNDROBIN -cltTimeout 180 -icmpVsrResponse ACTIVE
bind lb vserver vip-1-tcp service-1-http-80
```

Server Health Check

VPX can be configured to determine whether a server is down using many different monitors, including ICMP and TCP/UDP. The following configuration snippet from the evaluation uses the default ICMP monitor:

```
bind service service-1-http-80 -monitorName ping
```

VPX can detect server failure as it was configured.

SSL Offload

VPX was configured to perform SSL transactions for a group of HTTP servers. In this configuration, clients initiate an SSL connection to the VIP. VPX then creates an HTTP connection to the real server. A generic certificate was generated on VPX for testing purposes. The certificate for the server was installed on VPX so that it can act as a proxy for the real server. It is also possible to configure VPX to make a secure connection to the real server using SSL, but that configuration was not used.

```
set ssl vserver vip-1-ssl-443 -eRSA ENABLED
bind lb vserver vip-1-ssl-443 service-1-http-80
bind ssl service nshttps-::1l-443 -certkeyName ns-server-certificate
```

For more information about generating SSL certificates on VPX, refer to URL:

http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-ssl-generate-server-test-cert-tsk.html

For more information SSL about offload configuration, refer to:

http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-ssl-config-ssloffloading-con.html

L4-L7 Access Control Lists

VPX can protect the VIP from unauthorized connections using an ACL. These lists can be applied to any VPX segment, and can even be used to protect the VPX VM management interface. These ACLs can be standard (IP) or extended (L4 - L7). In this configuration snippet, the VIP is being protected from a certain client:

```
add ns acl block_outside_client DENY -srcIP = 10.110.192.100 -destIP =
10.110.191.1-10.110.191.254 -destPort = 80 -protocol TCP -interface 1/1 -priority 10
-state ENABLED -kernelstate APPLIED
```

High Availability

VPX supports active/standby failover and clustering. In high availability (HA) mode, heartbeat packets are sent on all active interfaces; therefore, there is no need to dedicate an interface for this purpose. Failover in HA mode during normal operations takes at least three seconds. This is because failover depends on the configured dead-interval. By default, this is three seconds, which is also the shortest time that can be configured for this parameter. The dead-interval indicates how long the secondary VPX will wait for heartbeat packets before it considers the primary VPX to be down. During failover, connections must be reestablished on the new primary instance of the VPX.

To create an HA pair, the following needs to be configured on the primary VPX first:

```
add HA node 1 172.26.162.227
```

The IP address refers to the VPX NSIP that is the secondary in the pair. The number 1 represents the node ID that is local to VPX and must be unique for each node added. For this evaluation, the default value for the dead interval was used.

For more information about configuring HA, please refer to the following url:

http://support.citrix.com/proddocs/topic/ns-system-10-map/ns-nw-ha-intro-wrppr-con.html

Clustering is a NetScaler feature that supports making multiple VPX instances appear as one device. This feature was not tested in this evaluation. For more information on clustering, refer to:

http://support.citrix.com/article/CTX132840

Dynamic Routing

VPX supports the following routing protocols: RIP, OSPF, ISIS, and BGP. Dynamic routing must be enabled on the IP interface used to connect to the routing protocol. VMDC 3.0 uses OSPF as the internal routing protocol within the DC, so VPX was configured to participate in this OSPF process:

```
router ospf 1
 redistribute static
 area 0 range 10.110.191.0/24
 network 10.110.191.0/24 area 0
```

This configuration must be entered in the Virtual Teletype Shell interface (VTYSH). This shell can be accessed by connecting to the VPX CLI and entering “vtysh”.

Alternately, BGP can also be used as the DC routing protocol for a VMDC implementation (as in the VMDC 2.2 and 2.3 designs). The configuration for BGP on VPX must also be entered in VTYSH. An example follows:

```
router bgp 13
 bgp router-id 10.110.191.250
 network 10.110.191.0/24
 neighbor 10.110.191.251 remote-as 13
```

For more information on configuring dynamic routing protocols on VPX, refer to:

http://support.citrix.com/proddocs/topic/ns-system-10-map/ns-nw-iprouting-config-dyna-rout-con.html

IPv6

VPX can perform SLB and SSL offloading in IPv6 mode, and can simultaneously service IPv4 and IPv6 traffic flows. This evaluation tested basic load balancing and SSL offloading using IPv6. The following configuration snippet shows a basic load balancing and SSL offloading configuration using IPv6:

```
add server server-1-ipv6 2001:db8:c18:1::4
add server server-2-ipv6 2001:db8:c18:1::5

add service service-1-ipv6-http-80 server-1-ipv6 HTTP 80 -gslb NONE -maxClient 0
-maxReq 0 -cip DISABLED -usip NO -useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout
360 -CustomServerID "\"None\"" -CKA YES -TCPB YES -CMP YES
add service service-2-ipv6-http-80 server-2-ipv6 HTTP 80 -gslb NONE -maxClient 0
-maxReq 0 -cip DISABLED -usip NO -useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout
360 -CustomServerID "\"None\"" -CKA YES -TCPB YES -CMP YES

add lb vserver vip-1-ipv6-http-80 HTTP 2001:db8:c18:1::10 80 -persistenceType NONE
-cltTimeout 180
add lb vserver vip-1-ipv6-ssl-443 SSL 2001:db8:c18:1::10 443 -persistenceType NONE
-cltTimeout 180

bind lb vserver vip-1-ipv6-http-80 service-1-ipv6-http-80
bind lb vserver vip-1-ipv6-http-80 service-2-ipv6-http-80
bind lb vserver vip-1-ipv6-ssl-443 service-1-ipv6-http-80
bind lb vserver vip-1-ipv6-ssl-443 service-2-ipv6-http-80
```

VPX Performance Results

VPX throughput depends upon which license is purchased for a particular VPX instance. For this evaluation, VPX was licensed for 3 Gbps. Traffic generation was done using a Spirent 9000A chassis equipped with a CV-10G-S8 card. Table 1 shows the results gathered from the performance tests for this evaluation.

Table 1 VPX on VMware Performance Results

Test	IPv4
L4 throughput	3.2 Gbps
L4 connections per second	77,000
L4 concurrent connections	253,000
L7 throughput	2.5 Gbps
L7 connections per second	400,000
L7 concurrent connections	112,000
SSL transactions per second	2300
SSL throughput	1.0 Gbps

The Spirent chassis emulated both clients and real servers. The VIP was configured on the VPX instance. Throughput tests involved a mix of traffic, including FTP and DNS, and a mix of packet sizes. Connection and transaction tests used one protocol, HTTP, in most cases. Figure 7 shows the logical topology.

Figure 7 Logical Topology for Performance Testing

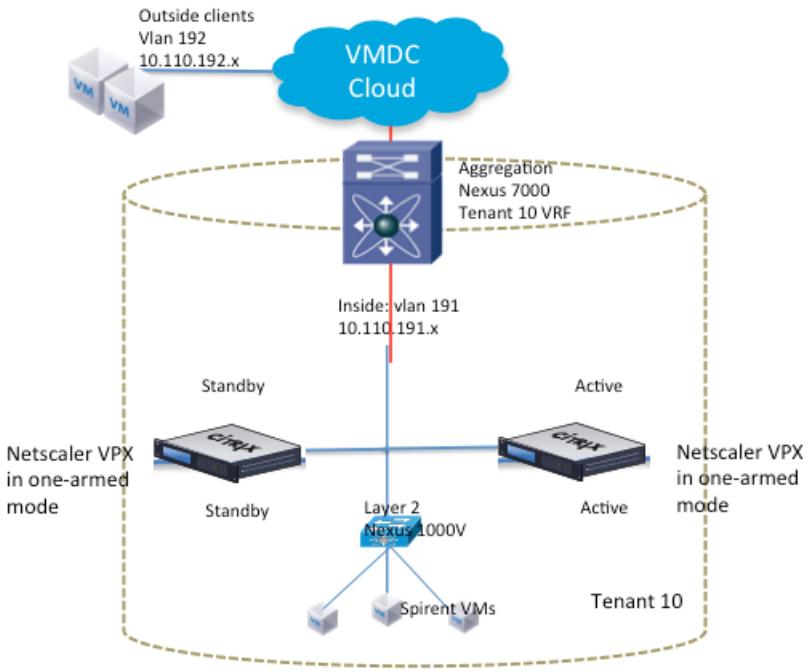

Testing Environment Components

Table 2 lists the hardware and software versions used to evaluate VPX installed on VMware.

Table 2 Testbed Validated Components

Product	Description	Hardware	Software
Citrix NetScaler VPX	Virtualized Load Balancer	N/A	10.0
Cisco ASA 1000V	Virtualized ASA, VPN Gateway, Firewall	N/A	8.7(1)
Cisco Nexus 1000V	Distributed Virtual Switch	N/A	NX-OS 4.2(1)SV1(5.2)
Cisco UCS	Compute Blade Servers	UCS 6120XP Fabric Interconnect UCS 5108 chassis UCS 2104XP IOM UCS B200 blades with 48G RAM UCS M81KR Adapter	2.0(2q)
Cisco Nexus 7000	Data Center Aggregation and Core devices	Nexus 7010 Sup-1 N7K-M132XP-12	6.0(4)

Table 2 Testbed Validated Components (continued)

Product	Description	Hardware	Software
Cisco Nexus 5000	Data Center Access device	Nexus 5548UP	5.1(3)N2(1)
VMware vSphere/ESXi	Virtualization/Hypervisor	N/A	5.0.0 Build 623373

SDX Evaluation in VMDC

The following sections provide detailed SDX evaluation considerations:

SDX in VMDC 3.0

Citrix NetScaler SDX is a virtualized physical appliance that supports multiple independent, isolated NetScaler instances running on the same appliance. In either a hierarchical classical Ethernet type Data Center architecture, such as VMDC 2.X systems or CLOS-type FabricPath designs as in VMDC 3.x systems, SDX is typically be placed in a services layer of the infrastructure, connecting to aggregation nodes or FabricPath aggregation-edge nodes.

SDX is equipped with 1 GE and 10 GE ports (type and number vary based on SDX model) that connect to the of the data center infrastructure they service. Each NetScaler instance on the SDX appliance has its own set of resources that are isolated from other NetScaler instances. Each NetScaler instance is functionally identical to a VPX virtual appliance and NetScaler MPX physical appliances. A NetScaler instance on an SDX 20500 (used for this validation) can handle speeds up to 18 Gbps, while the SDX 20500 can handle an aggregate throughput of 42 Mbps

As previously noted, the VMDC 3.0 design is based on a FabricPath-based network fabric. Like other appliance-based service designs, the SDX would typically connect to this fabric using a classical Ethernet vPC trunk connecting it to the aggregation layer of the data center fabric. In this case, the interfaces used to connect SDX to the aggregation layer are configured as an EtherChannel trunk. In NetScaler release 10.0 and earlier, an EtherChannel could be used by only one VPX instance at a time. In release 10.1, an EtherChannel can be shared among the NetScaler instances housed on an SDX appliance.

Figure 8 SDX Network Topology in VMDC 3.0

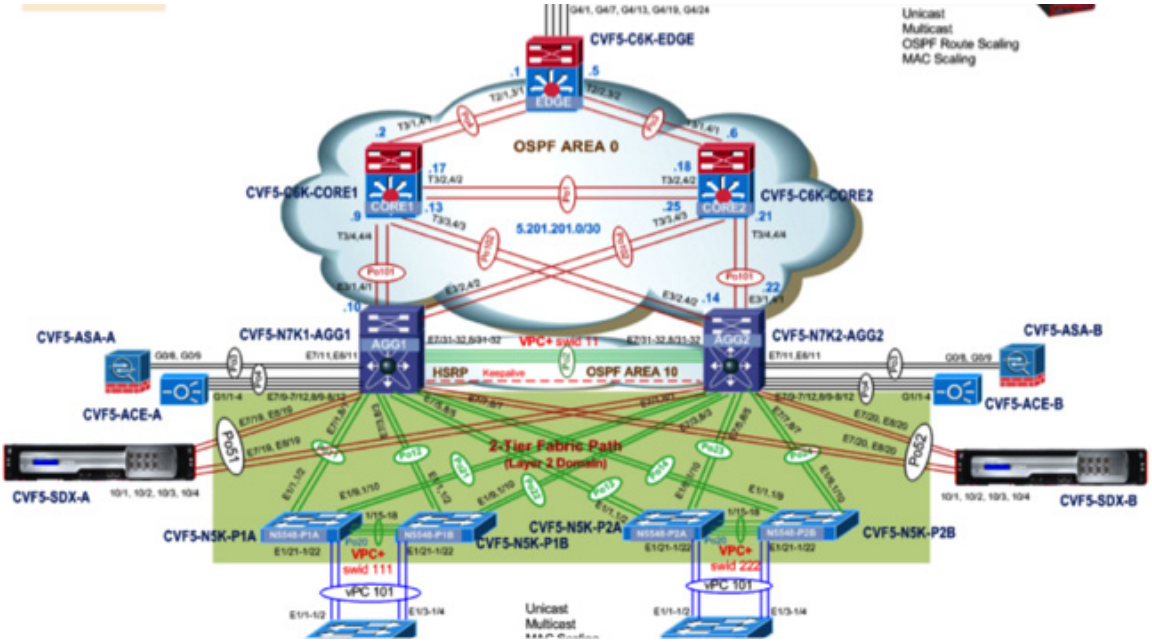

Note While this validation and document covers the use of Citrix SDX in a VMDC 3.0 based Palladium Network Container, the SDX can be used in a similar fashion in the VMDC 2.2 and 2.3 Gold and Silver Network Containers.

The installation process involves bringing up the SDX appliance, then bringing up the NetScaler instances. Licensing is different from the VPX on VMware in that a license is needed for the SDX appliance, not the VPX instances. For more information on SDX installation, refer to:

http://support.citrix.com/proddocs/topic/netscaler-getting-started-map-10/ns-instpk-install-ns-wrapper.html

Evaluation Topology and Methodology

This component of the evaluation looked at how SDX can be used in VMDC 3.0, including failure convergence scenarios. We tested the following features, along with performance and convergence, using Spirent Avalanche and IXIA for cases which required traffic generation:

- HA
- L4 SLB and ACL
- Syslog
- DSCP preservation

High Availability

In the case of the Cisco ACE, HA involves a Fault Tolerance partnership between the two ACE appliances. Each SLB context is active on one ACE appliance and standby on the other appliance.

With the Citrix SDX, the HA relationship exists only between two NetScaler instances, not between the SDX appliances. To decrease the chance of losing both NetScaler instances, we recommend placing the HA pair on different SDX chassis.

HA configuration for the NetScaler instances is available in the preceding section describing the VPX evaluation. Convergence numbers for HA are provided in the performance results section. SDX was taken through various failure scenarios, such as shutting down the primary of the failover pair. Observed convergence times indicate that replacing ACE with SDX would have little or no effect on operational recovery in failure scenarios.

High Availability with VMAC

For the HA mode, the Virtual MAC (VMAC) functionality on the NetScaler instances was validated. The VMAC feature is needed for failover scenarios in which upstream devices do not understand GARP (gratuitous ARP) messages sent by the new primary instance to update the Address Resolution Protocol (ARP) table. VMAC enables the failover pair to share a MAC address, eliminating the need for ARP table updates.

Each NetScaler instance pair on SDX was configured with same virtual router ID (VRID) to generate the VMAC. The VMAC is generated automatically using the VRID. The VRID must be configured on the SDX management plane, and on the instance.

In the following example, the VRID is bound to the EtherChannel configured on the SDX management plane and on the NetScaler instance:

```
add vrid 100
bind vrid 100 -ifnum LA/1
```

L4 SLB and ACL

See the preceding VPX evaluation section describing this feature, as this functionality is the same in all versions of VPX.

Syslog

The Syslog functionality on the SDX platform was validated. Logging level names differ from Cisco logging level names. For detailed information about Syslog messages, refer to:

http://support.citrix.com/article/CTX132382

Note that to send link failures to a Syslog server, the NOTICE logging level must be sent to the Syslog server.

The following configuration snippet is for enabling Syslog on VPX:

```
add audit syslogAction SYSLOG 192.168.5.247 -loglevel ALL -logFacility LOCAL5
-timeZone LOCAL_TIME
set audit syslogParams -serverIP 192.168.5.247 -loglevel ALL -timeZone LOCAL_TIME
add audit syslogPolicy sys_pol ns_true SYSLOG
bind system global sys_pol -priority 10
set audit syslogparams -serverip 127.0.0.1
set ns param -timezone GMT-04:00-EDT-America/New_York
```

DSCP Marking Preservation

SDX preserves DSCP packet marking for packets traversing the device for L4 SLB configurations. Preservation varies by application, and depends upon whether the NetScaler instance performs packet manipulation. For UDP connections, the DSCP value is preserved. For TCP connections, it may be preserved for the when transmitting the actual application traffic, but not for opening and closing the connection.

Performance and Convergence Results

The SDX convergence times were evaluated for various failure scenarios in the VMDC 3.0 design. Performance testing was executed on a NetScaler instance running in an SDX appliance.

IXIA IxLoad generated traffic for convergence test results, while Spirent Avalanche generated traffic for performance testing. Traffic generation using Avalanche used a Spirent Test Center chassis equipped with a MX-10G-S8 card. The Spirent chassis emulated both clients and real servers. The VIP was configured on the VPX instance. As was done for VPX performance testing, SDX throughput testing used a mix of traffic and packet sizes. Connection and transaction tests used one protocol, HTTP, in most cases.

For performance testing, the VPX instance was allocated one SSL chip.

Table 3 and Table 4 summarize convergence and performance results for SDX.

Table 3 Convergence Results for SDX in VMDC 3.0

Test	Convergence Time
Active SDX shutdown	.097 sec
Active NetScaler instance shutdown	3.168 sec
vPC port-channel link failure	.70 sec (AGG1 Po51) .06 sec (AGG2 Po51)
vPC Peer-link Shutdown	2.27 sec (AGG1) 34.3 sec (AGG2)

Table 4 Performance Results for One NetScaler Instance on SDX in VMDC 3.0

Test	IPv4
L4 throughput	17.4 Gbps
L4 connections per second	95,000
L4 concurrent connections	417,000
L7 throughput	12 Gbps
L7 connections per second	55,000
L7 concurrent connections	777,000
SSL transactions per second	25,000
SSL throughput	3.3 Gbps

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.