



ASA Clustering within VMDC Architecture

The adaptation of an enterprise-wide security framework is essential part of the overall enterprise network architecture. Within the data center new application rollouts, virtualization, the adaptation of various cloud services and an increasingly transparent perimeter are creating radical shifts in the data center security requirements. The need for stackable scalable high capacity firewalls at the data center perimeter is becoming indispensable. Adaptive Security Appliance (ASA) clustering feature on the ASA family of firewalls satisfies such a requirement. The clustering feature allows for an efficient way to scale up the throughput of a group of ASAs, by having them all work in concert to pass connections as one logical ASA device. Using up to 8 ASA appliances, the clustering feature allows the scaling of up to 100Gbps of aggregate throughput within the data center perimeter. This paper provides the basic feature description, design considerations for the ASA clustering feature and its integration within the Virtual Multi-Services Data Center architecture

Document Goal

This document describes basic functionality of ASA clustering, the Layer 2 clustering operational models, and design considerations when integrating ASA clustering within the VMDC architecture, along with benefits and use cases.

This document describes the basic outline of this document

- **ASA Clustering Overview**—The functional overview of ASA Clustering Feature.
- **Use Cases and Benefits**—The benefits and various use cases of the ASA clustering Features within the data center architecture.
- **VMDC Architecture Overview**—The VMDC architectural components, overall framework, and functional capabilities.
- **ASA Clustering Design Considerations**—The design considerations and implementation example
- **Validation Results**—Validation of basic configuration and various failure cases
- **Management**—Management tools that are available for configuration and monitoring of the ASA clustering system is outlined

Audience

The target audience for this document includes sales engineers, field consultants, professional services, IT managers, Cisco channel partner engineering staff, and customers who have requirements for a cloud-ready data center and wish to develop a seamless end-to-end security architecture.

ASA Clustering Overview

The clustering feature for the ASA family of firewalls allows for an efficient way to scale up the throughput of traffic that needs to be inspected by firewalls. The higher throughput is achieved by grouping a number of ASAs into one logical construct. The different ASA devices increased aggregate traffic by having them all work in concert to pass connections as one logical ASA device. Using up to 8 ASA appliances, clustering can allow up to 100Gbps of aggregate traffic.

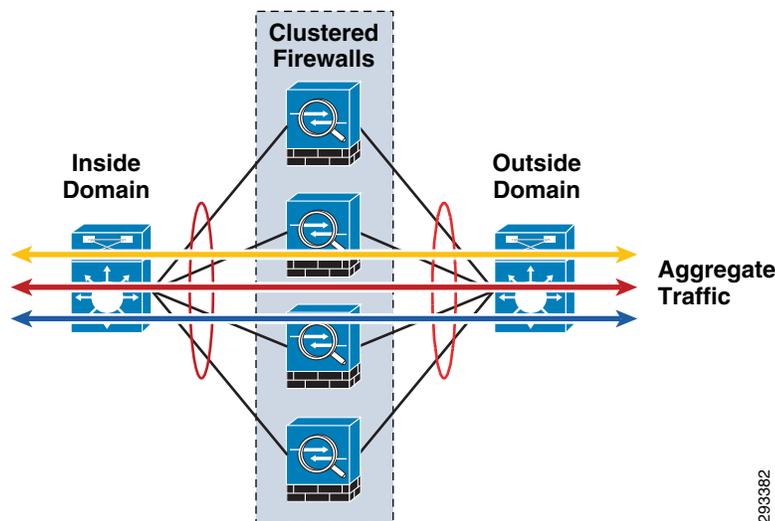
The ASA cluster does not actively load balance flows. It assumes that external mechanisms are in place to ensure that packets and traffic flows hitting each member of the cluster are properly load-balanced. This stateless load balancing can be achieved by external means such as ECLB provided by external switch or ECMP/PBR performed by routers. Within cluster, proprietary clustering control protocol is used to realign asymmetrical flows and to offload some flows to other units if one ASA unit in the cluster is becoming over utilized.

Clustering is different than your traditional active-active model. In the active-active model, a multi-context firewall configuration is assumed. In active-active model two firewalls are connected together by a control link. Some of the contexts are active on the first firewall and standby on the second firewalls, while the rest of the contexts are active on the second firewall and standby on the first. In contrast the ASA clustering features, every member of the cluster is capable of forwarding every traffic flow and can be active for all flows.

All units in the cluster have the same configuration as others and actively pass traffic. In the event of a unit failure, connectivity is maintained through the cluster thanks to connection information being replicated to at least one other unit in the cluster. Each connection has a replicated connection residing on a different cluster unit and takes over in case of a failure.

Clustering supports single and multiple context, as well as transparent mode. A single configuration is maintained across all units in the cluster using automatic configuration sync. Cluster-wide statistics are provided to track resource usage. [Figure 1](#) show the basic clustering construct.

Figure 1 Basic Clustering Construct



293382

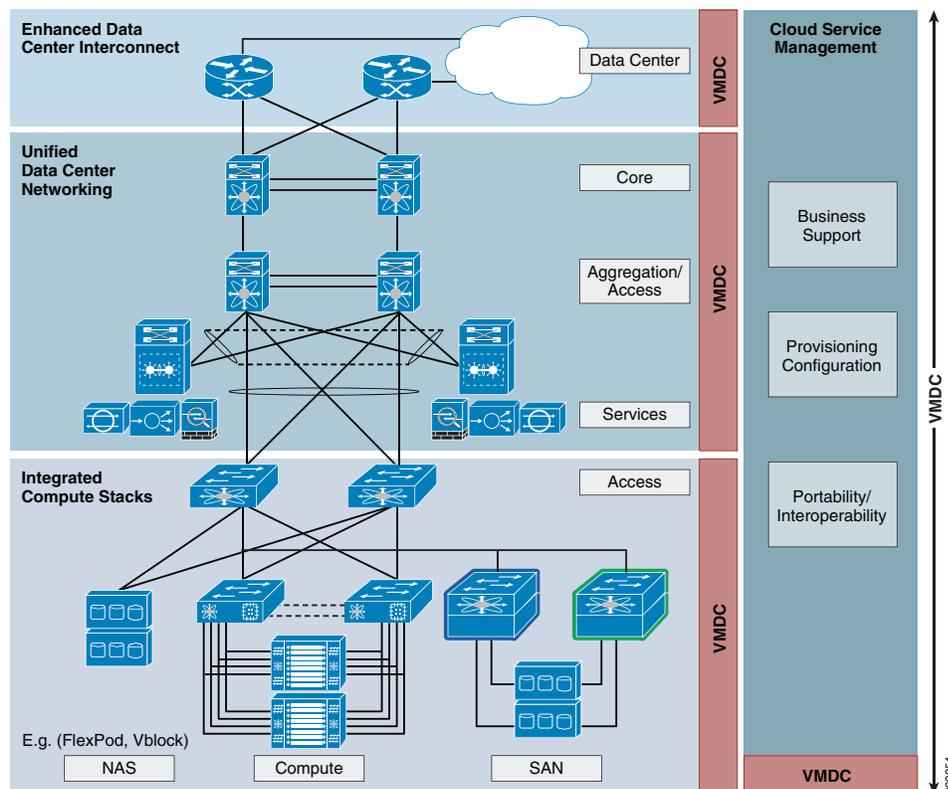
Clustering within VMDC Architecture

Clustering with its ability to provide high aggregate throughput is a natural fit within The Cisco Virtualized Multi-Tenant Data Center (VMDC). This section summarizes the basic framework and features of VMDC architecture.

VMDC System Overview

The Cisco Virtualized Multi-Tenant Data Center solution provides design and implementation guidance for enterprises deploying private cloud services and service providers building virtual private and public cloud services. The Cisco VMDC solution integrates various Cisco and third-party products that are part of the cloud computing ecosystem. Cisco's VMDC system defines an end-to-end architecture, which an organization may reference for the migration or build out of virtualized, multi-tenant data centers for new cloud-based service models such as Infrastructure as a Service (IaaS). Figure 2 shows the basic architectural framework for VMDC.

Figure 2 Basic VMDC Architecture Framework



VMDC is an end to end system that integrates compute, network and storage components with an architectural frame work. The various design paradigms and functional layer are defined within this framework.

Hierarchical Network Layers

The data center within the VMDC reference architecture is based on the classic multi-layer hierarchical network model. Hierarchical model benefits include scalability, resilience, performance, maintainability, and manageability and its design represents a structured approach to building the infrastructure, allowing for relatively easy expansion in modular increments. Redundant nodes and links at each level insure no single point of failure, while link aggregation can be engineered for optimal bandwidth and performance through the aggregation and core layers. In general, this hierarchical model uses three layers:

- **Core Layer**—Characterized by a high degree of redundancy and bandwidth capacity and thus optimized for availability and performance.
- **Aggregation Layer**—Characterized by a high degree of high-bandwidth port density capacity and thus optimized for traffic distribution and link fan-out capabilities to access layer switches. Functionally, the nodes in the aggregation layer typically serve as the Layer 2/Layer 3 boundary.
- **Access Layer**—Serves to connect hosts to the infrastructure, providing network access, typically at Layer 2 (L2) (i.e., LANs or VLANs).

VMDC Functional Layers

The VMDC architecture can also be functionally classified into the following categories.

- Network
- Services
- Compute
- Storage
- Management

The **Network** layer includes the WAN/PE router, which forms the data center perimeter to the Enterprise wide area or provider IP backbone, and to the public Internet. These perimeter nodes may be dedicated to Layer 3 routing functions, or may be multi-service in nature, providing Layer 2 interconnects between data centers as well as Layer 3 services. The VMDC topologies support two variants of the three-layer hierarchical model: a collapsed core/aggregation version, and a collapsed aggregation/access version. These allow for fine-tuning of port capacity and bandwidth to the level of aggregation or access density required to accommodate current and anticipated scale requirements.

The **Services** layer comprises network and security services such as firewalling, server load balancing, SSL offload, intrusion prevention, network analysis, and gateway functions. Within the VMDC reference architecture, the Data Center Services Node (DSN) provides firewalling and server load balancing services, in a service module form factor; alternatively, these are available in appliance form-factors. This layer also serves as the termination point for remote access IPsec or SSL VPNs; within the VMDC architecture, the Cisco physical appliances connected to the DSN fulfill this function, securing remote tenant access to cloud resources.

The **Compute** layer includes several sub-systems. The first is a virtual access switching layer, which allows for extension of the Layer 2 network across multiple physical compute systems. This virtual access switching layer is of key importance in that it also logically extends the Layer 2 network to individual virtual machines within physical servers. The feature-rich Cisco Nexus 1000V generally fulfills this role within the architecture. A second sub-system is that of virtual services. These may include security, load balancing, and optimization services. Services implemented at this layer of the infrastructure will complement more centralized service application, with unique applicability directly to a specific tenant or workgroup and their applications. Specific application based services validated within the VMDC architecture currently include the Cisco Virtual Security Gateway (VSG), providing

a security policy enforcement point within the tenant virtual data center. The third sub-system within the Compute layer is the computing resource that includes the Cisco Unified Compute System consisting of physical servers, hypervisor software providing compute virtualization abilities, and the virtual machines thus enabled.

The **Storage** layer provides storage resources. Data stores reside in SAN (block-based) or NAS (file-based) storage systems. SAN switching nodes use an additional level of resiliency, interconnecting multiple SAN storage arrays to the compute resources, via redundant FC or Ethernet links.

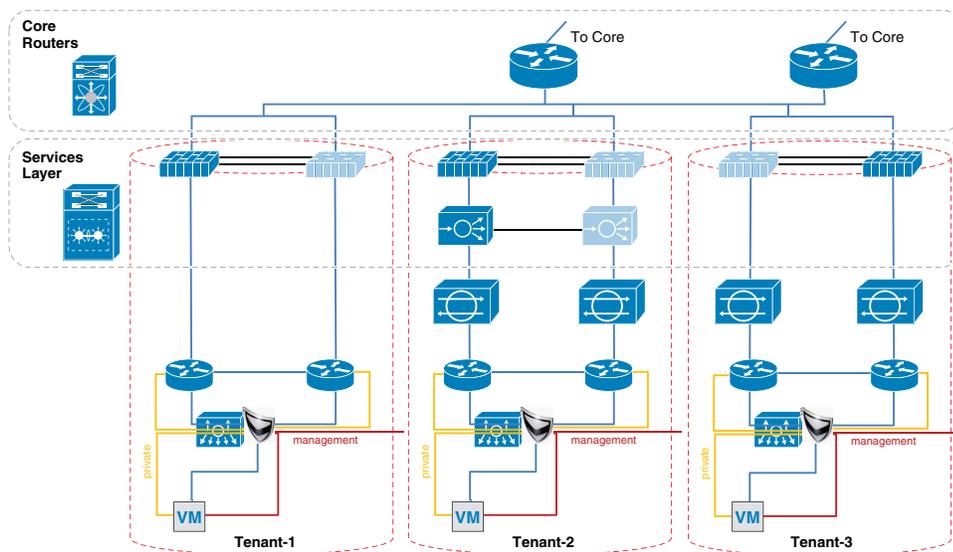
The **Management** layer consists of the “back-end” hardware and software resources required to manage the multi-tenant infrastructure. Such infrastructure include Active Directory, logging collection applications, and various device management software applications.

Multi-Tenancy Architecture

Virtualization of compute and storage resources enables sharing across an organizational entity. In contrast, virtualized multi-tenancy, a concept at the heart of the VMDC reference architecture, refers to the logical isolation of shared virtual compute, storage, and network resources. In essence, this is “bounded” or compartmentalized sharing. A tenant is a user community with some level of shared affinity. For example, within an enterprise, a tenant may be a business unit, department, or workgroup. Depending upon business requirements or regulatory policies, a tenant “compartment” may stretch across physical boundaries, organizational boundaries, and even between corporations.

A tenant container may reside wholly within their private cloud or may extend from the tenant's enterprise to the provider's facilities within a public cloud. The VMDC architecture addresses all of these tenancy use cases through a combination of secured data path isolation and a tiered security model which leverages classical security best practices and updates them for the virtualized multi-tenant environment. [Figure 3](#) shows the implementation of multi-tenancy within the VMDC architecture.

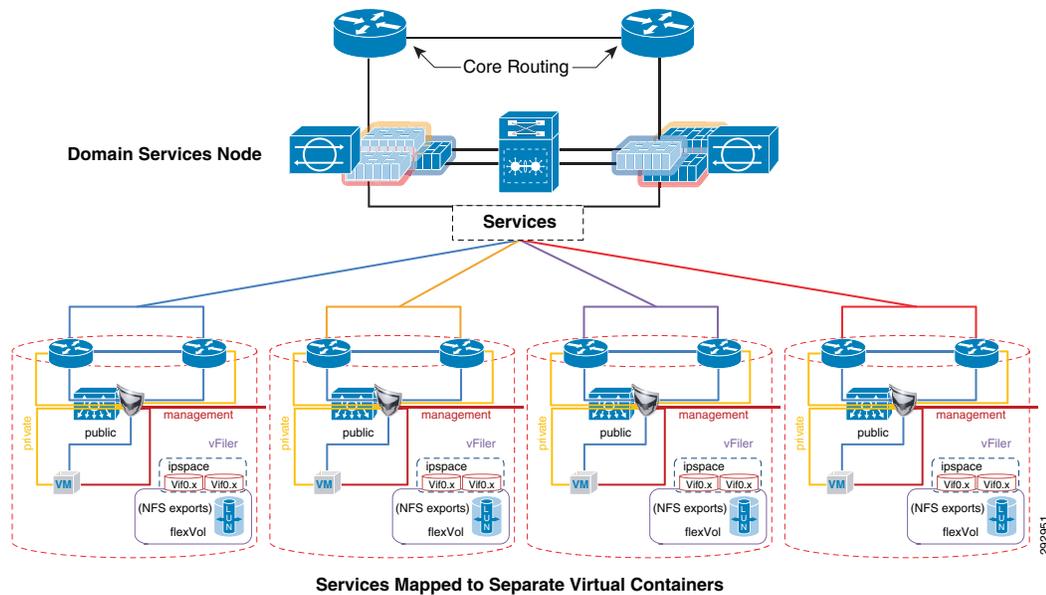
Figure 3 Multi-Tenancy Implementation within VMDC Architecture



Services Overlay within VMDC

VMDC framework facilitates seamless overlay and integration of various services securely and reliably by providing Infrastructure as a Service (IaaS). Such services include physical security, collaboration and IP Telephony. The container model within VMDC's architecture can be leveraged to overlay various services. [Figure 4](#) shows the integration of services within VMDC.

Figure 4 Services Integration within VMDC

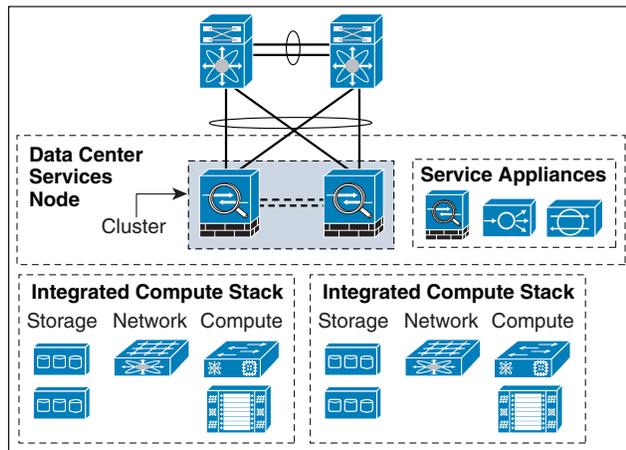


As seen, a number of services can be overlaid on VMDC by placing the various virtualized software components associated with each service within its own separate container. In case of Unified Communication, all of its components can be placed in its own separate container where it can be securely and reliably deployed. Each services-container can access the various common infrastructure service and the various tenants through the access layer or edge layer firewall. Firewall, load balancing and intrusion prevention services can be enabled for each services container as necessary.

Integration of Clustering within VMDC

The ASA physical firewalls are located within the services node of VMDC. It primarily an edge firewall, providing secure connectivity between the data center and the rest of the enterprise network, as well as providing network isolation between different zones or tenants within the data center. [Figure 5](#) shows the adaptation of clustering within VMDC services node is given below.

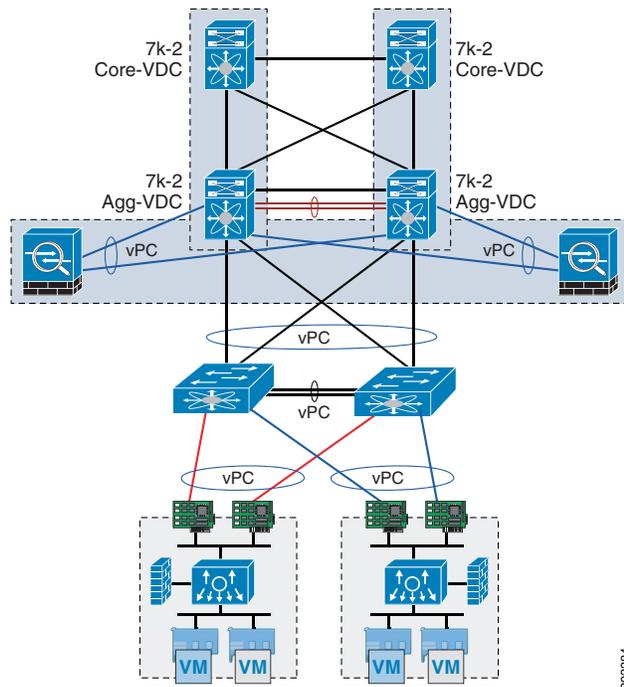
Figure 5 Clustering Adaptation within VMD Services Node



ASA Clustering Use Cases and Benefits

As mentioned, the most obvious use case and benefit of clustering is the ability to aggregate traffic to achieve higher throughput. The scaling of a number of ASA appliances into one logical firewall is a good fit within the datacenter architecture. Consolidation of data centers and the virtualization of network services are increasing bandwidth usage and are necessitating the adaptation of higher bandwidth firewalls. [Figure 6](#) shows the logical clustering of firewalls within a data center.

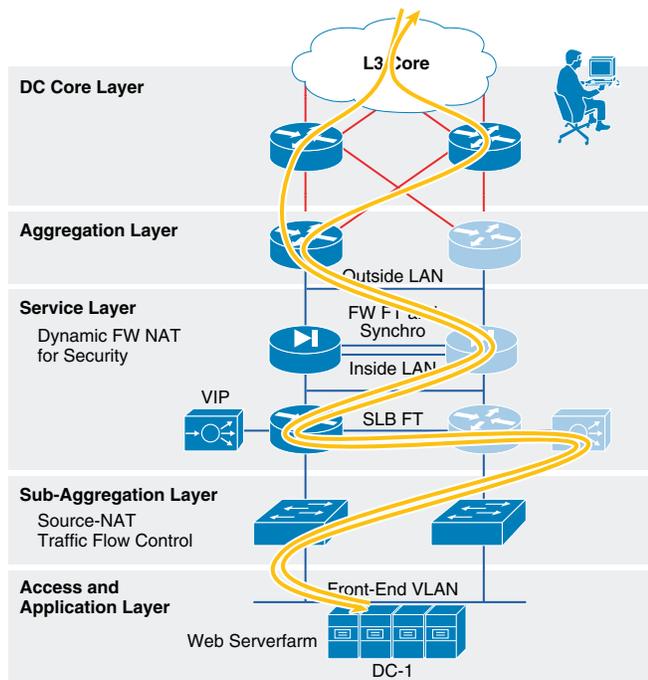
Figure 6 Logical Clustering of Firewalls within a Data Center



Symmetrical Flow Load-Balancing Inside a Data Center

In a data center there is a variety of appliances that inspect and forward state-full traffic. Firewalls, Load Balancers, Inspection/Detection Systems, SSL-Off loader and WAAS are examples of such appliances. Many of these appliances are implemented in a redundant fashion in active-standby or active-active mode for redundancy, scalability, full convergence and recovery of state-full traffic. Since state-full flows imply establishing a one-way symmetrical session, traffic flow through all devices can bounce between different paths, as shown in [Figure 7](#).

Figure 7 State Full Traffic Flow through Data Center



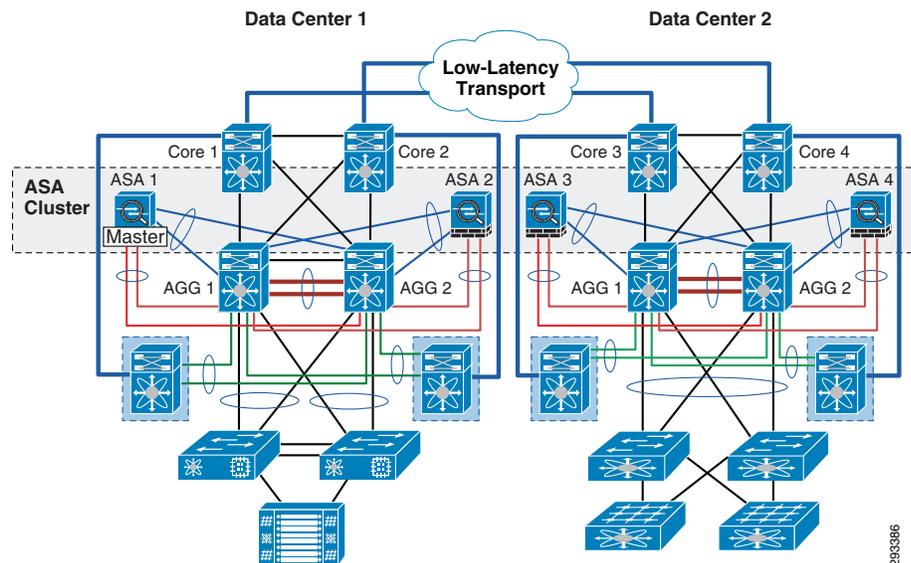
By implementing the clustering feature one can force state-full flows to take a more symmetrical path which improves predictability improved efficiency and session consistency of traffic flows.

ASA Cluster over Different Data Centers

In theory clustering can be implemented across different data centers.¹ The determining factor for implementing this use case is the latency between the data centers. The most desirable scenario is the situation where dark-fibre is used as a means of transport between the data centers. A possible scenario is shown in [Figure 8](#).

1. This use-case is not currently supported unless a dark-fibre is used between data centers and the support for this use-case is slated for future releases.

Figure 8 Dark Fibre Data Center Transport



The benefits of clustering over multiple data centers are as follows:

- Some of the traffic flows within a multi-site data centers can be asymmetrical. Clustering features forces asymmetrical flows to be symmetrical.
- Provides seamless state-full live migration with automatic redirection of flows.
- Consistent firewall configurations between data centers.
- Offloading of new connection to other members of the cluster, if the firewall is too busy.
- Solid redundancy and disaster recovery within the cluster in case of appliance or link failures.

ASA Clustering Operational Details

This section gives a brief overview of how ASA clustering functionality and how it is implemented. The cluster of ASAs operates in either Layer 2 mode or Layer 3 mode.

Layer 2 Mode

ASA interfaces are grouped in EtherChannel, and LACP must be used in conjunction. The same layer 2 interface on all ASA share a system IP and MAC, and together they act as a single logical link. Switch use ECLB to balance load to each ASA. Each interface also has its own private MAC address which is used in LACP auto negotiation if enabled.

Layer 3 Mode

ASA interfaces on each ASA have its own IP and MAC. Upstream router can use PBR or ECMP to balance packets to individual units in the cluster. At this time Layer 3 mode is not supported.

Data path Packet Flow Through the Cluster

To support the stateful operation of ASA data path, we centralize the state for each connection on one unit called the “owner”; all packets belonging to the same connection must be processed at the owner, and if packets arrive at a non-owner unit, they are forwarded to the owner over the Cluster Control Link for processing.

The first ASA to receive traffic for a connection is designated as the owner.

To provide fault-tolerance, the connection states are backed-up (replicated) on a different ASA unit called the “director”, in a similar method as stateful failover. In addition, the director unit also provides owner info of this connection to all other units. This is possible because director is selected by performing a hash over connection's 5-tuple. Any unit would find the same director using 5-tuple, and would get the owner unit by querying the director.

Types of Cluster Connections

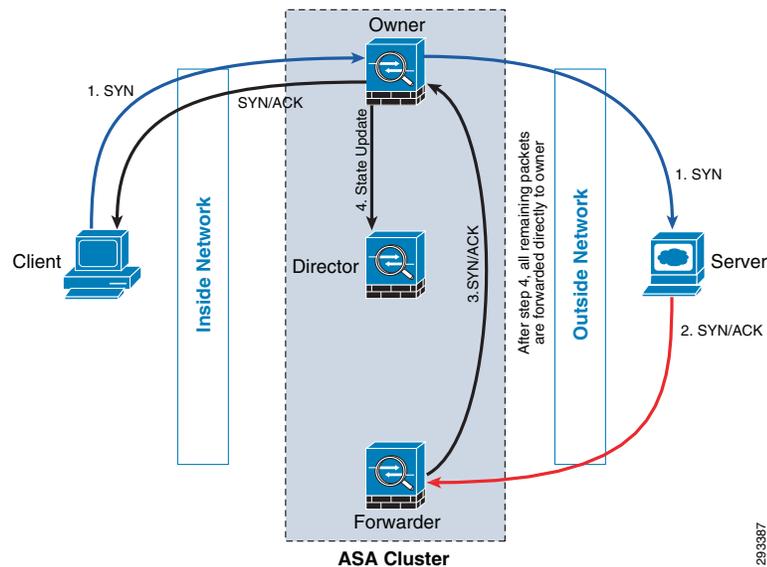
The following show the conn flags for cluster connections.

- **Owner flow**—The actual connection flow that is passing the traffic. We can't know for sure which unit in the cluster will “own” the flow since whichever ASA receives the first packet in the flow will become the owner. Only TCP and UDP flows send logical flow updates to the stub flow (and possibly the director stub flow).
- **Forwarding Stub Flow**—If a unit receives a packet for a flow that it does not own, it will contact the director of that flow to learn which unit owns the flow. Once it knows this, it will create and maintain a forwarder flow, which it will then be used to forward any packets it receives on that connection directly to the owner, bypassing the director. Forwarder flows do not receive Link Updates (LUs) (since they're just forwarding the packets and don't care about state). Short lived flows such as DNS and ICMP will not have forwarder flows; the unit receiving the packets for those conns will simply forward them to the director, which will forward them to the owner, and the director will not reply back to the forwarder unit asking it to create a forwarder flow.
- **Backup Stub Flow**—Based on the flow's characteristics, all units can derive the Director unit for the flow. The director unit typically maintains the stub (or backup) flow, which can become the full flow in the case the flow's owner unit fails, and also be used to redirect units towards the flow's owner unit if they receive packets for the flow. Backup flows receive conn updates to keep them up-to-date in case the owner fails and the stub flow needs to become the full flow.
- **Stub (or Backup) Director Flow**—If the director chosen for the flow is also the owner (meaning the director received the first packet in the flow) then it can't be its own backup. Therefore a 'director backup' flow will be created, and a second hash table will be used to track this. Obviously this director backup flow will receive LUs, since it needs to be ready to take over if the director/owner fails.

The following example outlines operational details of clustering by illustrating a TCP session connection buildup through the cluster.

Assuming asymmetric traffic, here's what would happen to establish a TCP connection through ASA, as shown in [Figure 9](#).

Figure 9 TCP Connection through ASA



TCP Connection through ASA

-
- Step 1** SYN packet originates from client and is delivered to owner unit. Owner unit creates flow; encode owner information into SYN cookie, and forward packet to server.
 - Step 2** SYNACK packet originates from server and is delivered to forwarder unit.
 - Step 3** Forwarder unit decodes owner information from SYN cookie, create forwarding flow to record owner, and forward SYNACK to owner unit.
 - Step 4** Owner unit sends state update to director unit, and forward SYNACK to client.
 - Step 5** Director unit receives state update from owner, creates director flow and records TCP state information as well as owner.
-

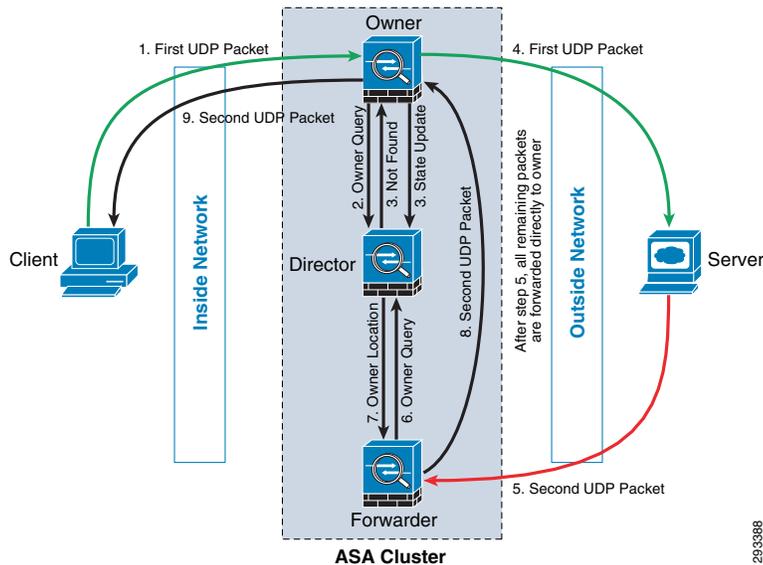
After the above process, the TCP connection is established and is backed up at director unit. Any subsequent packets delivered to director unit or forwarder will be forwarded to owner. If packets are delivered to any other cluster unit, it will query the director to get owner unit. Owner information is recorded via a forwarding flow for future forwarding. Any state change for this flow will result in a state update from owner to director.

The above sequence assumes TCP sequence number randomization is turned on (the default configuration). If it is turned off, owner information will not be encoded into SYN cookie. As a result the forwarder unit in step (3) will need to query director to obtain owner information. This also means the owner unit in step (1) must also send a state update to director, thus allowing the director to possess owner info. The overall impact is thus reduced efficiency and throughput.

UDP connection buildup steps through the cluster outlined below:

Assuming asymmetric traffic, here's what would happen to establish a UDP connection through ASA, as shown in [Figure 10](#).

Figure 10 UDP Connection through ASA



UDP Connection through ASA

-
- Step 1** UDP packet originates from client and is delivered to owner unit.
- Owner unit query director for any existing owner.
 - Director finds no existing flow; create director flow and forward packet back to owner unit. In other words director has elected owner unit for this flow.
 - Owner unit creates flow, sends state update to director unit, and forward packet to server.
- Step 2** UDP packet originates from server and is delivered to forwarder unit.
- Forwarder unit queries director for owner.
 - Director replies with owner info.
 - Forwarder creates forwarding flow to record owner info, and forward packet to owner.
 - Owner forward packet to client.
-

Centralized, Distributed and Unsupported Clustering Features

The TCP and UDP examples highlight the distributed data path clustering supports. Certain features on the ASA have been modified to be “distributed”, meaning that each ASA cluster member can individually complete these tasks. If a feature is “centralized”, the tasks must be completed by the Master (for example, packets for a non-distributed protocol inspection would have to all be forwarded to the cluster master for processing). Obviously this increases the processing and packet forwarding required to complete the operation, and decreases overall cluster performance. Any centralized feature will have flows always residing on the master unit. For list of distributed, unsupported, and centralized ASA features refer to [Clustered Firewall Features, page -17](#).

Cluster Link Aggregation Protocol

ASA uses a logical link aggregation construct called the Cluster Link Aggregation Control Protocol (cLACP). It is designed to extend standard LACP to multiple devices so that it can support span-cluster. Furthermore, cLACP also provides following functionality:

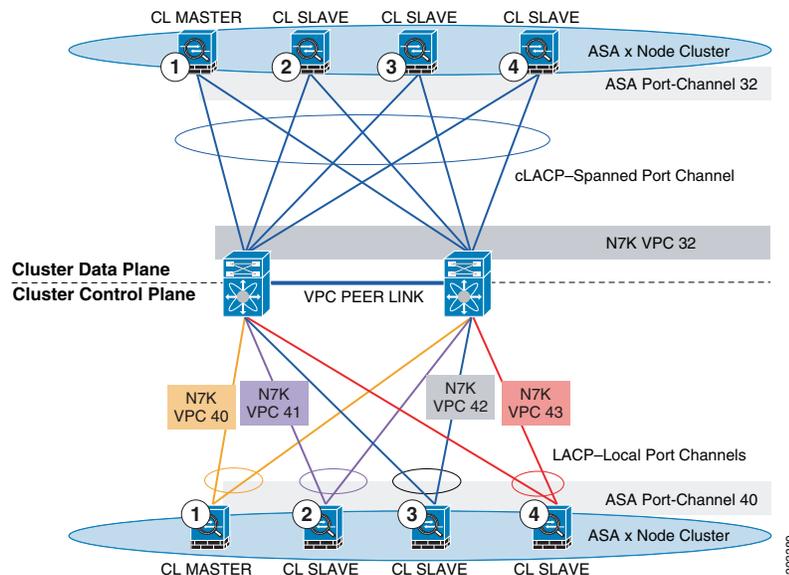
- Inter-operate with standard LACP device
- Load-balance at both link and device level
- Failure recovery during link failure or unit failure

LACP, as defined as a link-level control protocol that allows the dynamic negotiation and establishment of Link Aggregation Groups only among two systems. In ASA Clustering, EtherChannels need to be span across the cluster. CLACP allows link aggregation between one switch, or pair of switches, to multiple (more than two) ASAs in a cluster.

ASA Clustering Implementation Example and Design Considerations

To implement an ASA clustering, one needs to perform two basic tasks. The first is to configure the Cluster-Control-Link (CCL) and Data Link. The CCL provides control plane information between the different cluster members. Also the flows are redirected within the CCL. [Figure 11](#) shows an example of Cluster Control Plane and Cluster Data Plane configuration that was implemented within the VMDC framework using Layer 2 Transparent mode.

Figure 11 Layer 2 Transparent Mode Cluster Control and Data Planes Configuration



Topology

The topology for the scenario discussed is shown in [Figure 11](#). There are four firewalls connected to two nexus 7000. Each ASA has two 10GE interfaces connected to the respective Nexus 7000s. Port-channels on the ASA side and vPCs on the Nexus7000 side provide link aggregation functionality. The data links and the CCL links are separate. In the diagram above the four firewalls at the top and bottom are the same devices. The separation in the topology is only for illustration purposes.

CCL Configuration

It is recommended that port-channels be used instead of single links to provide redundancy. CCL is an integral part of the cluster, and a failure in communication between different members of the cluster will result in ejection of that firewall from the cluster. Using redundant links within CCL greatly reduces the possibility of total link failure between members.

To configure the CCL, one configures local port channels with the same channel identifier (40 in this case) on each firewall and connect them to separate vPCs on the corresponding Nexus7000s. The same port-channel identifier has to be used on each firewall in the cluster. This is required because the clustering feature expects identical port-channel configuration on each member of the cluster. On the Nexus 7000 side different identifiers are used (identifiers 40-43 in example above) since the VPCs are not spanned across the two Nexus switches. The VPCs are local within each switch and each VPC can operate separately even if the ASAs on the other side use the same LACP identifier.

Cluster Data Link Configuration

The most important distinction in implementing the cluster data plane is the configuration of a “spanned port channel (cLACP)” on the firewall. This is necessary because only one Port-Channel/vPC pair is used in the data plane. To provide channel consistency and seamless operation between both sides, it is necessary to configure a logical port-channel construct across all the members of the ASA cluster members. In the topology shown above channel identifier 32 is used on the ASA cluster side and on the Nexus 7000 side.

Other Clustering Design Considerations

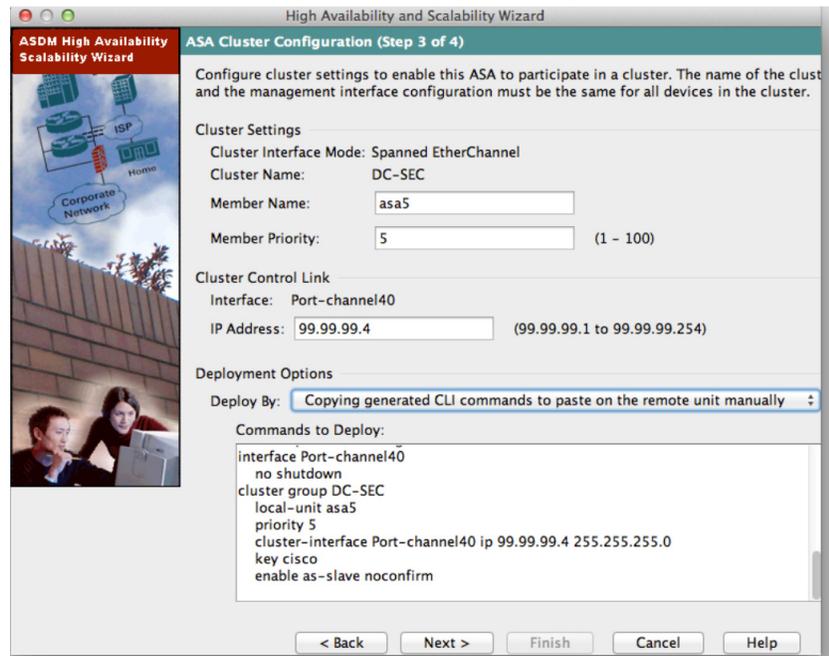
Some other design considerations are as follows:

- It is recommended to use an even number of firewalls in the cluster. This is necessary because Layer 2 load-balancing mechanisms do not redistribute traffic symmetrically if odd numbers of links are aggregated within a port-channel.
- It is recommended that the same link-speed to be used for CCL and Data plane interfaces. If one is using a 10GE interface on the Data interfaces, one has to use a 10GE on the CCL interfaces.
- The total real-life aggregate throughput of traffic through the cluster is normally less than the theoretical maximum. If four 10GE ASAs are used within a cluster, depending on the traffic type, one may see 20% reduction in total theoretical 40G aggregate traffic capacity of the cluster*.
- It is recommended not to use Spanned port-channel cLACP for CCL links.

Managing ASA Clustering

Like most ASA features, clustering can be configured via the ASDM Graphical User Interface. ASDM provides a more enhanced cluster-wide graphing for a high-level overview of cluster statistics and traffic parameters. [Figure 12](#) shows an example of the configuration wizard used for ASA clustering feature.

Figure 12 ASA Clustering Configuration Wizard



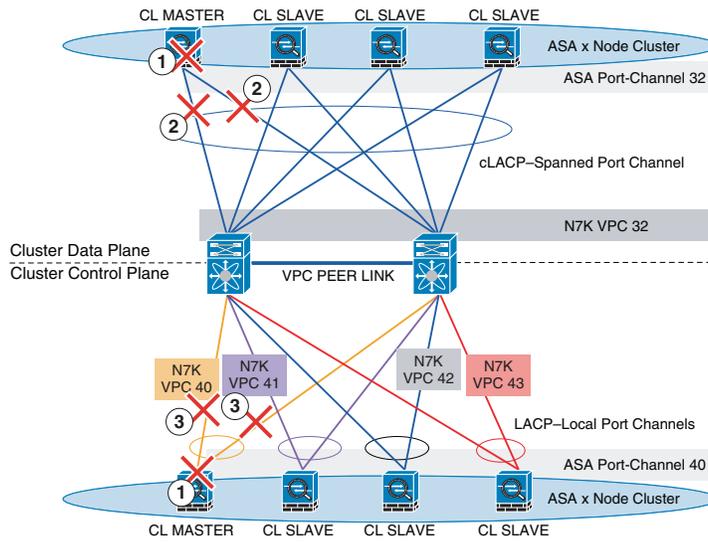
Validation Results for ASA Clustering

The ASA clustering was configured within VMDC framework. The topology used is shown at Figure X. Stateless and state full traffic was used and various design failures was validated. The following failures scenarios were simulated and validated on the physical setup. In all cases the firewall or links that were failed were actively passing traffic.

1. Failure of an existing firewall within the cluster
2. Failure of Cluster Data Links within one firewall
3. Failure of CCLs within one firewall
4. Failure of three out of four firewall appliances within the cluster where only a single link was left standing

Figure 13 shows failure scenarios 1-3.

Figure 13 Failure Scenarios 1-3

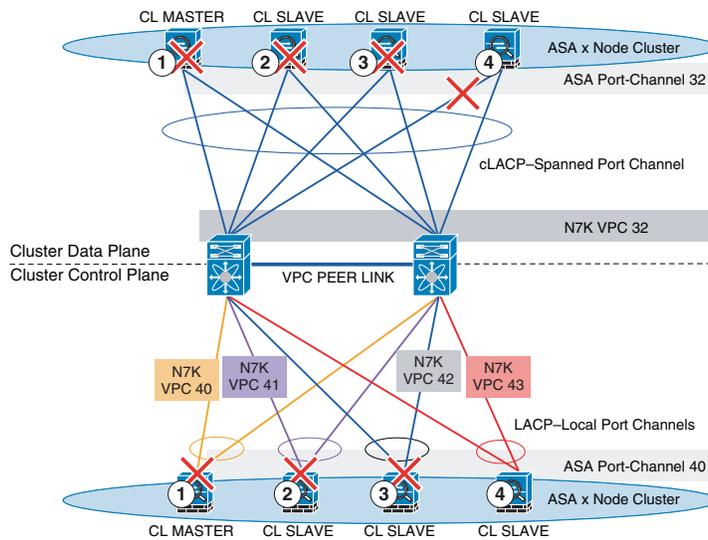


Scenario 1: Single Appliance Failure
Scenario 2: Data Link Failure
Scenario 3: CCL Link Failure

293381

Figure 14 shows failure scenario 4.

Figure 14 Failure Scenario 4



Scenario 4: Fail ASA1, ASA2 and ASA 3 and one of the data plane Links on ASA4

293382

In all cases cluster recovered correctly from failures and none or very minimal packet loss was detected.

Conclusion

This document describes the integration of the ASA clustering feature within the Virtual Multiservice Datacenter (VMDC) architecture. In addition, this documents the benefits, and use cases for this feature as well design considerations and functional description of the ASA clustering feature.

Clustered Firewall Features

The following clustered firewall feature set lists are provided:

- [Distributed ASA Features, page -17](#)
- [Centralized ASA Features, page -18](#)
- [Unsupported ASA Features, page -19](#)

Distributed ASA Features

Distributed ASA features are listed below.

- AAA for admin access
- Access List
- AIC
- AIP Midland card for Spyker
- App inspection: DNS
- App inspection: FTP (unless cut-through proxy is enabled, see CSCuc38180)
- App inspection: HTTP
- App inspection: ICMP
- App inspection: ICMP ERROR
- App Inspection: ILS
- App inspection: ip-optoins
- ARP Inspection
- ASDM Handler
- Card Manager
- Connection Limit
- Date and time
- DNS lookup
- Dynamic DNS
- Layer 2 Unicast Routing NSF
- Layer 3 Dynamic Routing
- Etherchannel
- Identity Firewall
- Interface vlan trunking
- IP Audit
- IP Fragment protection
- IPS
- IPv4 and Ipv6
- Jumbo-frame

- Licensing
- Mac-address
- Management (to the box traffic)
- Management-only
- MPC
- Multicast NSF
- Nameif
- NAT
- Object grouping
- Packet Tracer
- Passwd, enable passwd, host name, domain name
- PKI
- QOS
- Quack
- Redundant interface (Layer 3 only)
- Regex config/matching
- Routed/Transparent mode
- Security-level
- Shun
- Single/Multi mode
- Smart Call Home
- Syslog and netflow
- TCP Intercept
- TCP Normalizer
- TCP state bypass
- Threat Detection (no state sharing)
- Time Ranges
- Top-N Calculations

Centralized ASA Features

Centralized ASA features are listed below.

- AAA for network access
- Filtering Services
- Inspect DCERPC
- Inspect ESMTTP
- Inspect IM
- Inspect NetBios

- Inspect PPTP
- Inspect Radius
- Inspect RSH
- Inspect SNMP
- Inspect SQLNET
- Inspect SUNRPC
- Inspect TFTP
- Inspect XDMCP
- Inspect FTP (only if cut-through proxy is enabled, see CSCuc38180)
- PIM
- IGMP
- Layer 2 Dynamic Routing
- Layer 3 Multicast Data Traffic
- VPN:Layer 3/IKEv1
- VPN:Layer 3/IKEv2
- VPN management access

Unsupported ASA Features

Unsupported ASA features in clustering, which must be disabled for clustering to run, are listed below.

- Auto Update Server
- Inspect CTIQBE
- Inspect WCCP
- Inspect WAAS
- Inspect MGCP
- Inspect MMP
- Inspect RTSP
- Inspect Scansafe Scansafe
- Inspect SIP
- Inspect Skinny
- Inspect H323, H225, and RAS
- Inspect GTP
- UC/IME/Phone Proxy
- TLS Proxy
- BTF
- NAC
- DHCP client, server, relay, proxy
- VPN Remote Access

- VPN Load Balancing
- Failover
- IPSec Passthru

Axel Nadimi



Alex Nadimi, Solutions Architect, Systems Development Unit (SDU), Cisco Systems

Alex has been with Cisco for the past 15 years and is currently working as a Solutions Architect in Cisco’s Systems Development Unit. 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



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.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks.; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARtNet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0805R)

