Cisco live!

July 10-14, 2016 • Las Vegas, NV

Your Time Is Now

ASA Clustering Deep Dive

Andrew Ossipov, Principal Engineer

BRKSEC-3032

Your Speaker

Andrew Ossipov

aeo@cisco.com

Principal Engineer

8 years in Cisco TAC

19+ years in Networking

Agenda

- Clustering Overview
- Unit Roles and Functions
- Packet Flow
- Control and Data Interfaces
- Configuring Clustering on ASA Appliances
- Multi-Site Clustering
- Closing Remarks

Clustering Overview

ASA Failover

- A **pair** of identical ASA devices can be configured in Failover
  - Licensed features are aggregated except 3DES in **ASA 8.3+**
  - Data interface connections must be mirrored between the units **with** L2 adjacency
  - Active/Standby or Active/Active deployment with multiple contexts
  - Virtual IP and MAC addresses on data interfaces move with the active unit
  - Centralized management from the active unit or context
  - Stateful failover “mirrors” stateful conn table between peers
- Failover delivers high availability rather than scalability
  - Cannot scale beyond two physical appliances/modules or virtual instances
  - Active/Active failover requires manual traffic separation with contexts
  - Stateful failover makes Active/Active impractical for scaling

ASA Clustering

- **Up to 16** identical ASA appliances combine in one traffic processing system
- Preserve the benefits of failover
  - Feature license aggregation across entire cluster
  - Virtual IP and MAC addresses for first-hop redundancy
  - Centralized configuration mirrored to all members
  - Connection state preserved after a single member failure
- Implement true scalability in addition to high availability
  - Stateless load-balancing via IP Routing or Spanned Etherchannel with LACP
  - Out-of-band Cluster Control Link to compensate for external asymmetry
  - Elastic scaling of throughput and maximum concurrent connections
  - All units **should** be connected to the same subnet on each logical interface

System Requirements

- All cluster members must have an identical hardware configuration
  - Up to 16 ASA5585-X, Firepower 4110, or Firepower 9300 modules
  - Up to 2 ASA5500-X in **ASA 9.1(4)+**
  - Chassis types, application modules, and interface cards must match precisely
- Each ASA5580/5585-X member must have Cluster license installed
  - Enabled by default on ASA5500-X except ASA5512-X without Security Plus
  - 3DES and 10GE I/O licenses must match on all members for ASA
- Limited switch chassis support for control and data interfaces
  - Catalyst 3750-X, **3850**, **4500**, **4500-X**, 6500, and 6800 with Sup2T
  - Nexus **3000**, 5000, 6000, 7000, 9300, and 9500
  - ASR 9000

Unsupported Features

- Remote Access VPN
  - SSL VPN, Clientless SSL VPN, and IPSec
- DHCP Functionality
  - DHCP client, DHCPD server, DHCP Proxy
- Advanced Application Inspection and Redirection
  - GTP, and Diameter over TCP until **ASA 9.5(2)**
  - SCTP and Diameter over SCTP until **ASA 9.6(1)**
  - CTIQBE, WAAS, MGCP, MMP, RTSP, Skinny, H.323
  - Cloud Web Security, Botnet Traffic Filter, and WCCP
  - ASA CX module
- TLS Proxy until **ASA 9.6(1)** with Diameter inspection **only**

Scalability

- Throughput scales at 70% of the aggregated capacity **on average**
  - 16 ASA5585-X SSP-60 at 40Gbps → 448Gbps of Maximum UDP Throughput
  - 16 ASA5585-X SSP-60 at 20Gbps → 224Gbps of Real World TCP Throughput
  - Scales at ~**100%** with no traffic asymmetry between members (up to **640Gbps**)
- Concurrent connections scale at 60% of the aggregated capacity
  - 16 ASA5585-X SSP-60 at 10M → 96M concurrent connections
- Connections rate scales at 50% of the aggregated capacity
  - 16 ASA5585-X SSP-60 at 350K CPS → 2.8M CPS
  - Optionally delay short-lived connection replication in **ASA 9.4(2)+**

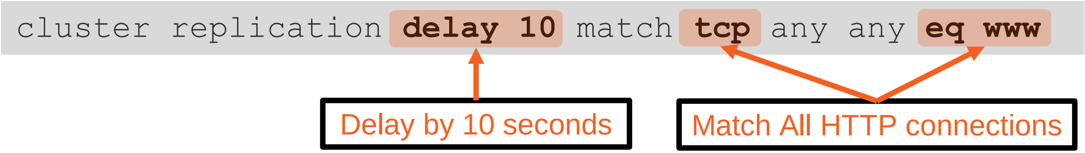

```
cluster replication delay 10 match tcp any any eq www
```

Delay by 10 seconds

Match All HTTP connections

Cisco live!

Centralized Features

- Not all features are distributed, some are **centralized**
  - Control and management connections
  - Non-Per-Session Xlates with PAT (e.g. ICMP)
  - DCERPC, ESMTP, IM, Netbios, PPTP, RADIUS, RSH, SNMP, SQLNet, SunRPC, TFTP, and XDMCP inspection engines
  - Site-to-site VPN
  - Multicast in some scenarios
- Any connections matching these features always land on one cluster member
  - Switchover of such connections is not seamless

Unit Roles and Functions

Master and Slaves

- One cluster member is elected as the **Master**; others are **Slaves**
  - First unit joining the cluster or based on configured priority
  - New master is elected **only** upon departure
- Master unit handles all management and centralized functions
  - Configuration is blocked on slaves
  - Virtual IP address ownership for to-the-cluster connections
- Master and slaves process all regular transit connections equally
  - Management and some centralized connections must re-establish upon Master failure
  - Disable or reload Master to transition the role; **do not** use **cluster master** command

State Transition

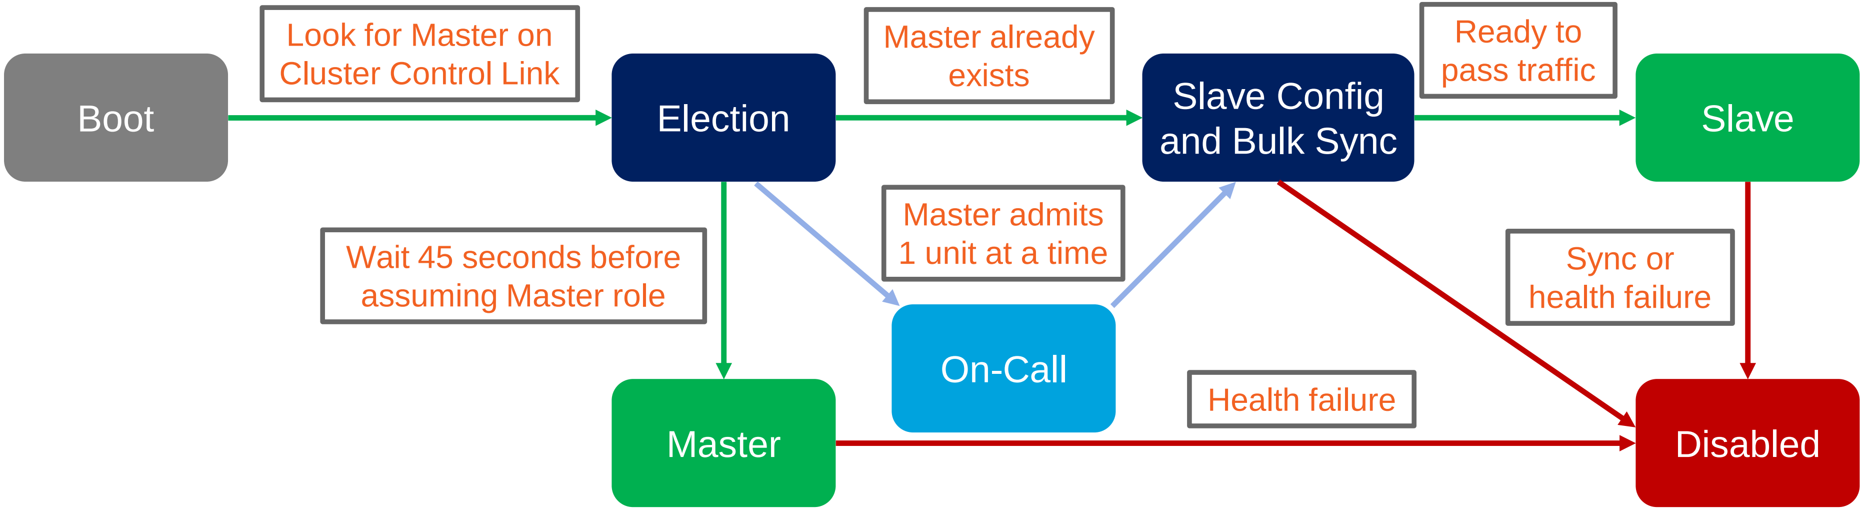
Boot
Look for Master on Cluster Control Link
Election
Master already exists
Slave Config and Bulk Sync
Ready to pass traffic
Slave
Wait 45 seconds before assuming Master role
Master admits 1 unit at a time
On-Call
Sync or health failure
Master
Health failure
Disabled

```
ASA/master# show cluster history
==========================================================================
From State             To State              Reason
==========================================================================
15:36:33 UTC Dec 3 2013
DISABLED               DISABLED              Disabled at startup
15:37:10 UTC Dec 3 2013
DISABLED               ELECTION              Enabled from CLI
15:37:55 UTC Dec 3 2013
ELECTION               MASTER                Enabled from CLI
==========================================================================
```

```
ASA/master# show cluster info
Cluster sjfw: On
    Interface mode: spanned
    This is "A" in state MASTER
        ID         : 0
        Version    : 9.1(3)
        Serial No.: JAF1434AERL
        CCL IP     : 1.1.1.1
        CCL MAC    : 5475.d029.8856
        Last join : 15:37:55 UTC Dec 3 2013
        Last leave: N/A
```

Flow Owner

- All packets for a single **stateful** connection must go through a single member
  - Unit receiving the first packet for a new connection typically becomes **Flow Owner**
  - Ensures symmetry for state tracking purposes and FirePOWER NGIPS inspection

```
ASA/master# show conn
18 in use, 20 most used
Cluster stub connections: 0 in use, 0 most used
TCP outside  10.2.10.2:22 inside  192.168.103.131:35481, idle 0:00:00, bytes 4164516, flags UIO
```

- Another unit will become Flow Owner if the original one fails
  - Receiving packet for an existing connection with no owner
- The **conn-rebalance** feature should be enabled with caution
  - An overloaded member may work even harder to redirect new connections
- Existing connections move **only** on unit departure or with Flow Mobility

Flow Director

- Flow Owner for each connection must be discoverable by all cluster members
  - Each possible connection has a deterministically assigned Flow Director
  - Compute hash of {SrcIP, DstIP, SrcPort, DstPort} for a flow to determine Director
  - Hash mappings for all possible flows are evenly distributed between cluster members
  - All members share the same hash table and algorithm for consistent lookups
  - SYN Cookies reduce lookups for TCP flows with Sequence Number Randomization
- **Flow Director** maintains a backup stub connection entry
  - Other units may query Director over Cluster Control Link to determine Owner identity
  - New Owner can recover connection state from director upon original Owner failure

```
TCP outside  172.18.254.194:5901 inside  192.168.1.11:54397, idle 0:00:08, bytes 0, flags  Y
```

  - Create Backup Flow when Director and Owner are the same or in the same chassis

```
TCP outside  172.18.254.194:5901 inside  192.168.1.11:54397, idle 0:00:08, bytes 0, flags  y
```

Flow Forwarder

- External stateless load-balancing mechanism does not guarantee symmetry
  - Only TCP SYN packets can reliably indicate that the connection is new
- Cluster member receiving a non-TCP-SYN packet must query Flow Director
  - No existing connection → Drop if TCP, become Flow Owner if UDP
  - Existing connection with no Owner → Become Flow Owner
  - Existing connection with active Owner → Become **Flow Forwarder**
- Flow Forwarder maintains stub connection entry to avoid future lookups
  - Asymmetrically received packets are redirected to Owner via Cluster Control Link
  - Slave units become Flow Forwarders for any centralized connections

```
ASA/slave# show conn detail
[…]
TCP inside: 192.168.103.131/52033 NP Identity Ifc: 10.8.4.10/22,
   flags  z, idle 0s, uptime 8m37s, timeout -, bytes 0,
   cluster sent/rcvd bytes 25728/0, cluster sent/rcvd total bytes 886204/0, owners (1,255)
```

Packet Flow

New TCP Connection

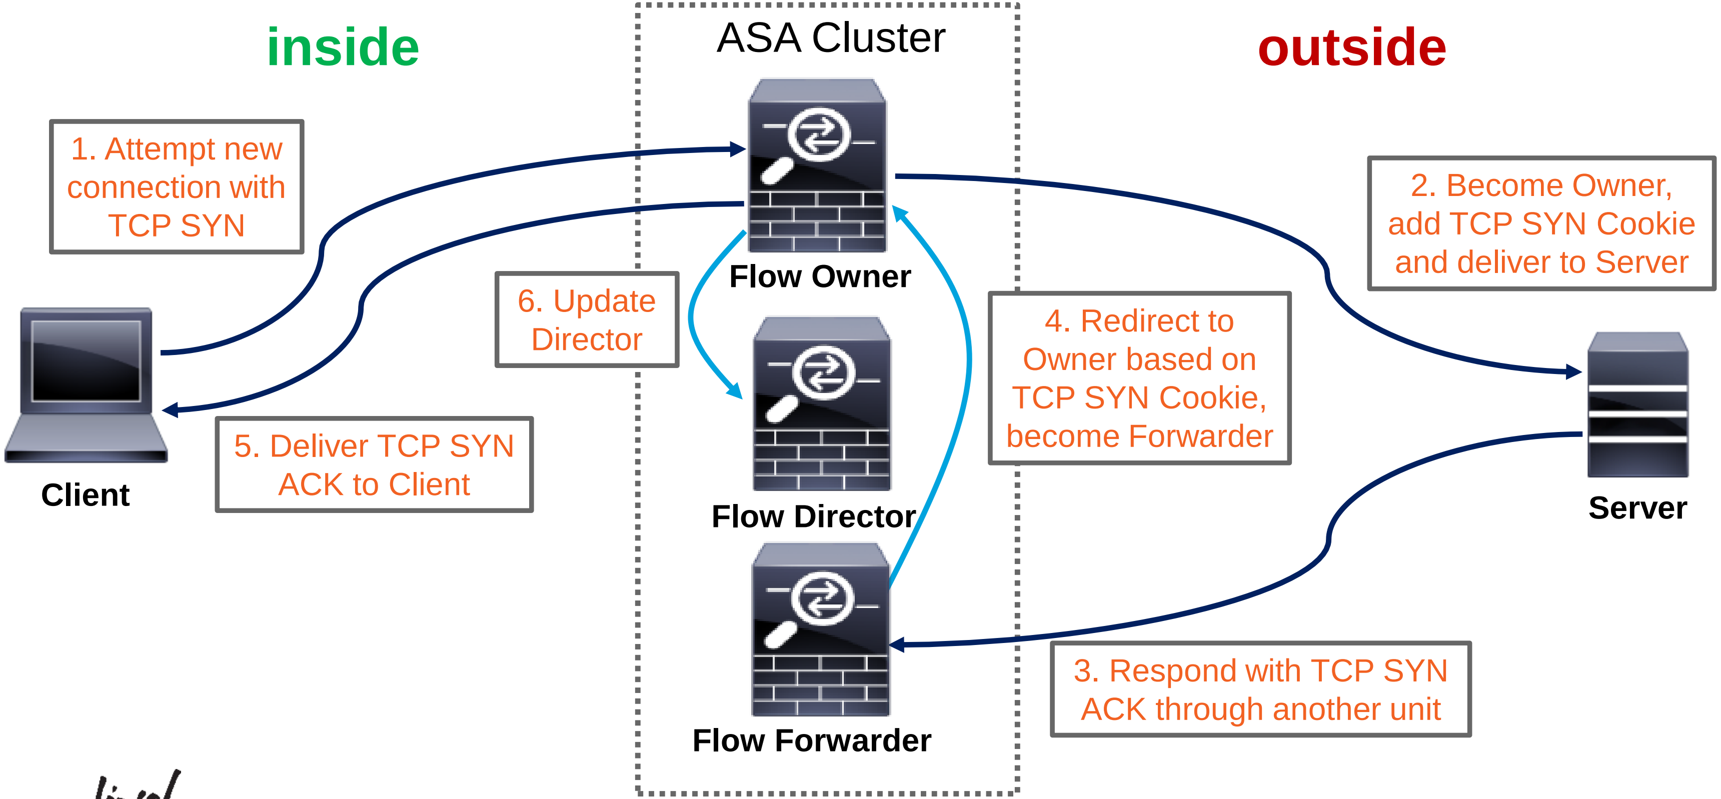

1. Attempt new connection with TCP SYN

2. Become Owner, add TCP SYN Cookie and deliver to Server

3. Respond with TCP SYN ACK through another unit

4. Redirect to Owner based on TCP SYN Cookie, become Forwarder

5. Deliver TCP SYN ACK to Client

6. Update Director

New UDP-Like Connection

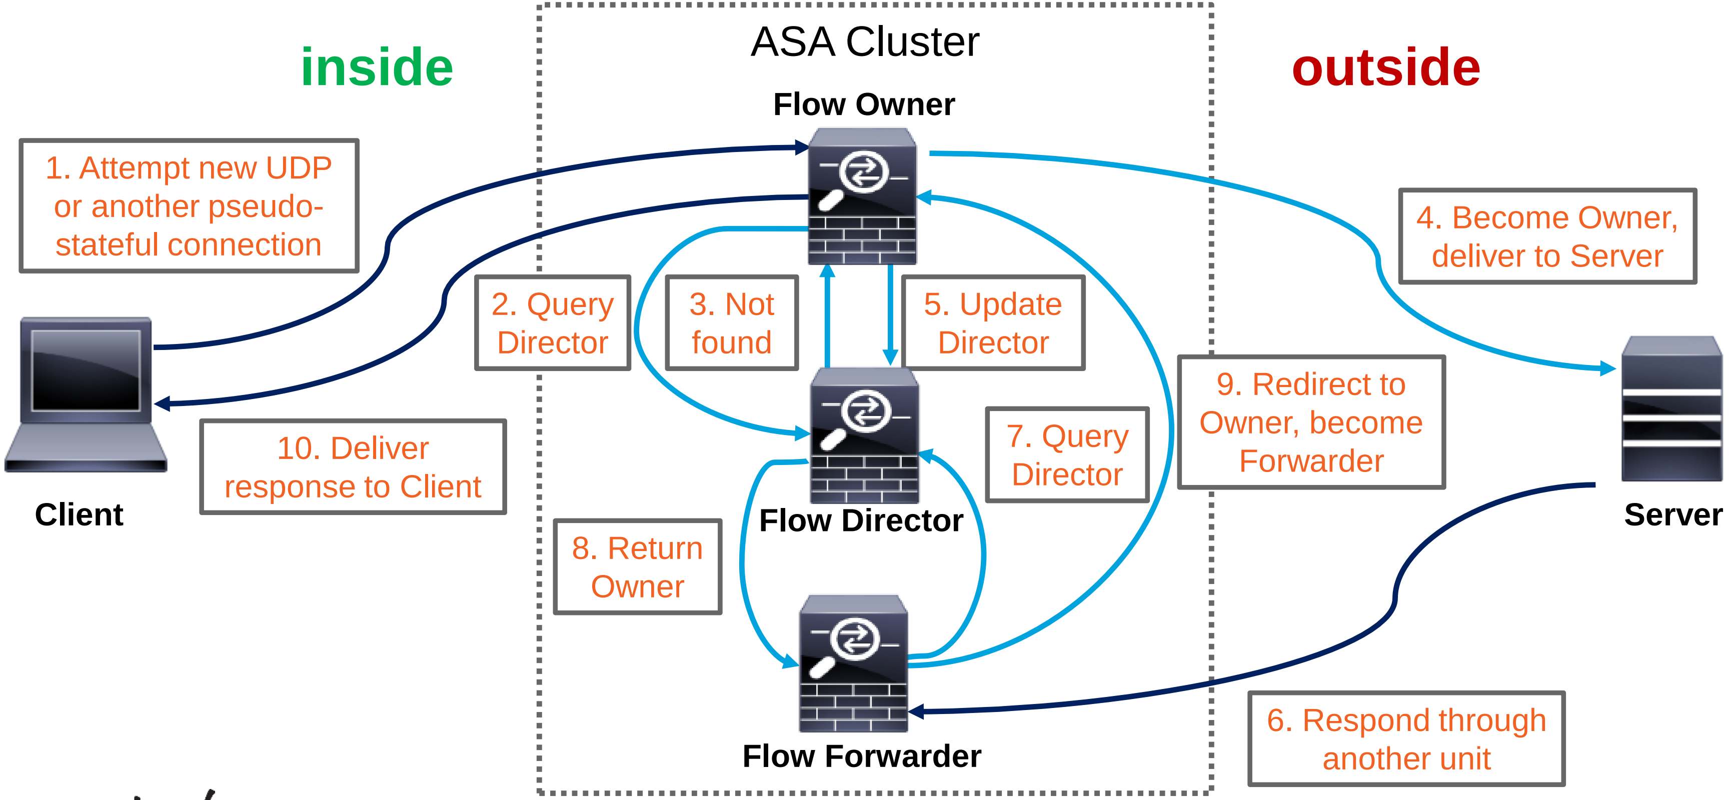

1. Attempt new UDP or another pseudo-stateful connection

2. Query Director

3. Not found

4. Become Owner, deliver to Server

5. Update Director

6. Respond through another unit

7. Query Director

8. Return Owner

9. Redirect to Owner, become Forwarder

10. Deliver response to Client

Reference

New Centralized Connection

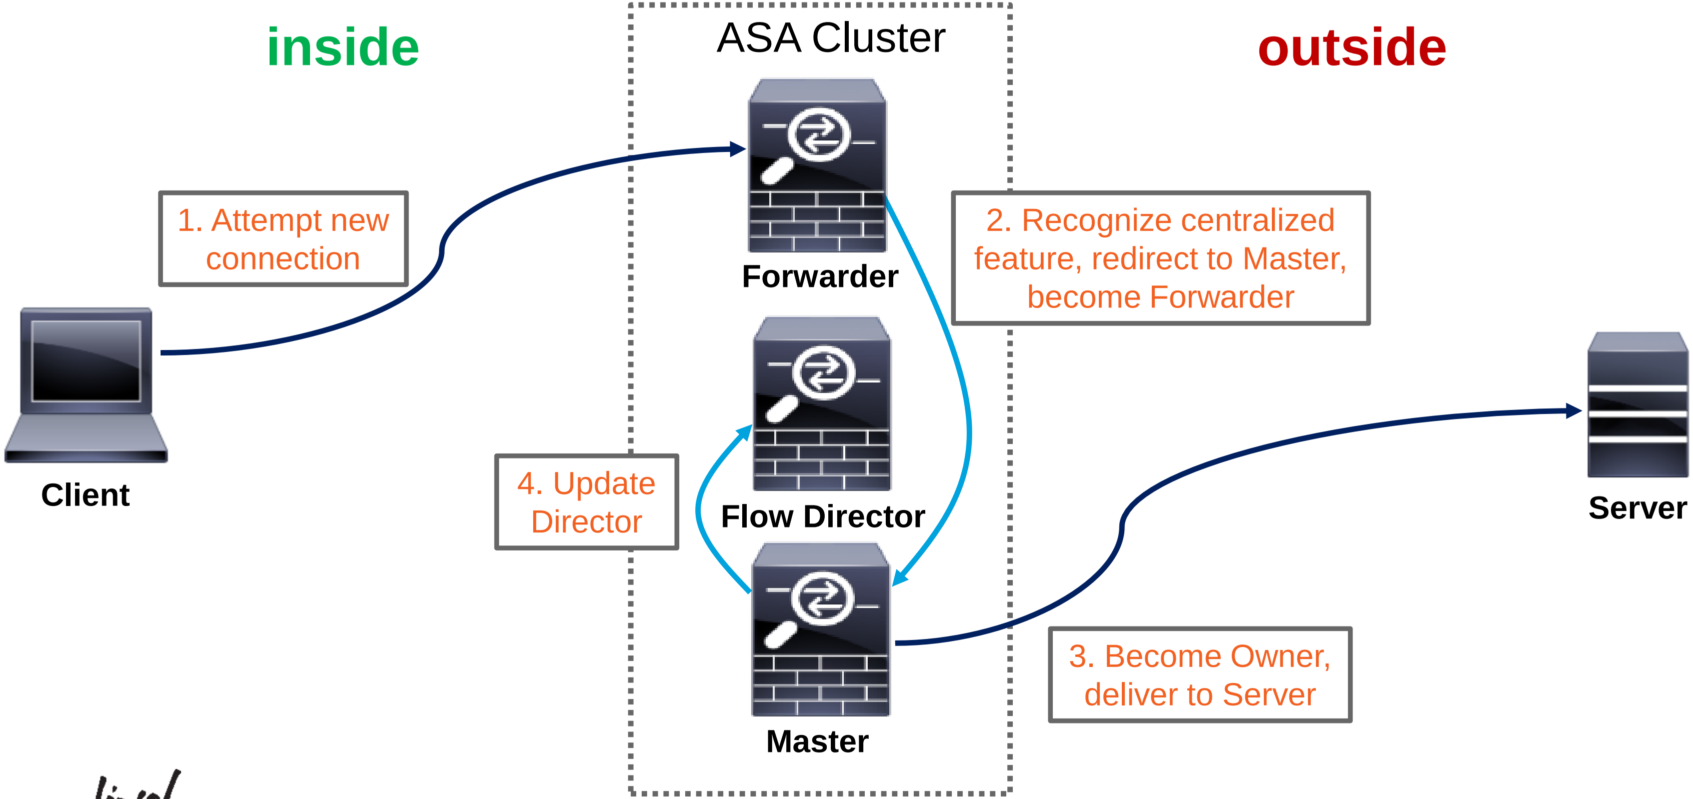

inside
ASA Cluster
outside
Forwarder
Flow Director
Master
Client
Server
1. Attempt new connection
2. Recognize centralized feature, redirect to Master, become Forwarder
3. Become Owner, deliver to Server
4. Update Director

Owner Failure

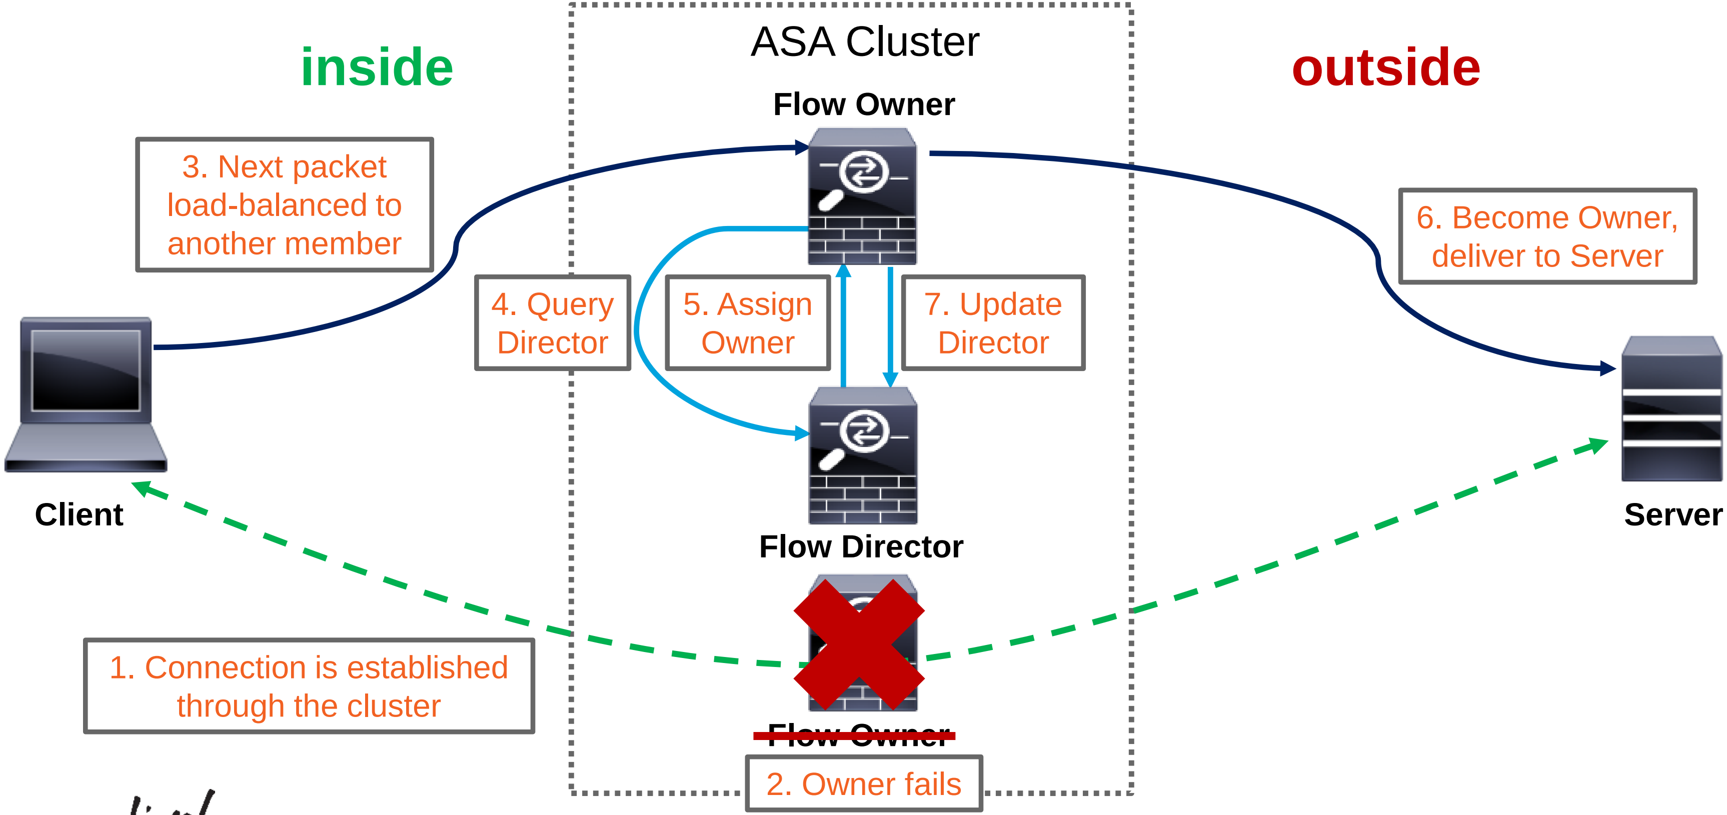

1. Connection is established through the cluster

2. Owner fails

3. Next packet load-balanced to another member

4. Query Director

5. Assign Owner

6. Become Owner, deliver to Server

7. Update Director

Application Inspection

- Centralized
  - All packets for control and associated data connections are redirected to Master
  - Examples: ESMTP, SQLNet, TFTP
- Fully Distributed
  - Control and associated data connections are processed independently by all units
  - Examples: HTTP, FTP, GTP
- Semi Distributed with **ASA 9.4(1)+**
  - Control connections are processed independently by all units
  - Data connections are redirected to the associated control connections' Owners
  - Examples: SIP, SCTP

Per-Session Port Address Translation (PAT)

- By default, dynamic PAT xlates have a 30-second idle timeout
  - Single global IP (65535 ports) allows about 2000 conn/sec for TCP and UDP
- **ASA 9.0** Per-Session Xlate feature allows immediate reuse of the mapped port
  - Enabled by default for all TCP and DNS connections

```
asa# show run all xlate
xlate per-session permit tcp any4 any4
xlate per-session permit tcp any4 any6
xlate per-session permit tcp any6 any4
xlate per-session permit tcp any6 any6
xlate per-session permit udp any4 any4 eq domain
xlate per-session permit udp any4 any6 eq domain
xlate per-session permit udp any6 any4 eq domain
xlate per-session permit udp any6 any6 eq domain
```

  - TCP Reset is generated to force immediate termination

Network Address Translation (NAT)

- Static NAT is performed by all cluster members based on configuration
- One-to-one dynamic NAT xlates are created by Master and replicated to Slaves
- Dynamic PAT is distributed to individual members
  - Master evenly allocates PAT addresses from the configured pools to each member
  - Provision **at least** as many pool IPs as cluster members to avoid centralization
  - Per-session xlates are local to the Owner with an Xlate backup
  - Some connections require non-per-session xlates which are centralized to Master

```
asa(config)# xlate per-session deny tcp any4 any4 eq 5060
```

- NAT limits clustering scalability with nearly guaranteed flow asymmetry
  - NAT and PAT pools are not advertised
  - No interface PAT or Proxy ARP in Individual mode

Control and Data Interfaces

Cluster Control Link (CCL)

- Carries all data and control communication between cluster members
  - Master discovery, configuration replication, keepalives, interface status updates
  - Centralized resource allocation (such as PAT/NAT, pinholes)
  - Flow Director updates and Owner queries
  - Centralized and asymmetric traffic redirection from Forwarders to Owners
- Must use same dedicated interfaces on each member
  - Separate physical interface(s), no sharing or VLAN sub-interfaces
  - An isolated non-overlapping subnet with a switch in between members
  - No packet loss or reordering; up to 10ms one-way latency in **ASA 9.1(4)+**
- CCL loss **forces** the member out of the cluster
  - No direct back-to-back connections **except** Firepower 4100 and 9300

CCL Best Practices

- Size and protect CCL appropriately
  - Bandwidth should match maximum forwarding capacity of each member
  - Use an LACP Etherchannel for redundancy and bandwidth aggregation
  - 20Gbps of Real World traffic with ASA5585-X SSP-60 → 2x10GE CCL
  - Dual-connect to different physical switches in vPC/VSS
  - Use I/O cards for extra 10GE ports in **ASA 9.1(2)+,** not IPS/SFR SSP
- Set L2 MTU 100 bytes above largest data interface SVI/L3 MTU
  - Avoids fragmentation of redirected traffic due to extra trailer
- Ensure that CCL switches do not verify L4 checksums
  - TCP and ICMP checksums for redirected packets look “invalid” on CCL
- Enable Spanning Tree Portfast and align MTU on the switch side

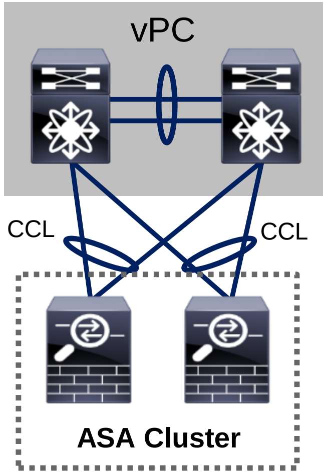

Data Interface Modes

- Recommended data interface mode is **Spanned Etherchannel “L2”**
  - Multiple physical interfaces across all members bundle into a single Etherchannel

```
asa(config)# interface Port-Channel1
asa(config-if)# port-channel span-cluster
```

  - External Etherchannel load-balancing algorithm defines per-unit load
  - All units use the same virtual IP and MAC on each logical data interface
- Each member has a separate IP on each data interface in **Individual “L3”** mode
  - Use Nexus ITD or PBR or dynamic routing protocols to load-balance traffic
  - All Etherchannels are local to each member
  - Virtual IPs are owned by Master, interface IPs are assigned from configured pools

```
asa(config)# ip local pool INSIDE 192.168.1.2-192.168.1.17
asa(config-if)# interface Port-Channel1
asa(config-if)# ip address 192.168.1.1 255.255.255.0 cluster-pool INSIDE
```

Spanned Etherchannel Interface Mode

- Create transparent and routed firewalls on per-context basis
- Must use Etherchannels: “firewall-on-a-stick” VLAN trunk or separate
- Use symmetric Etherchannel hashing algorithm with different switches
- Seamless load-balancing and unit addition/removal with cLACP

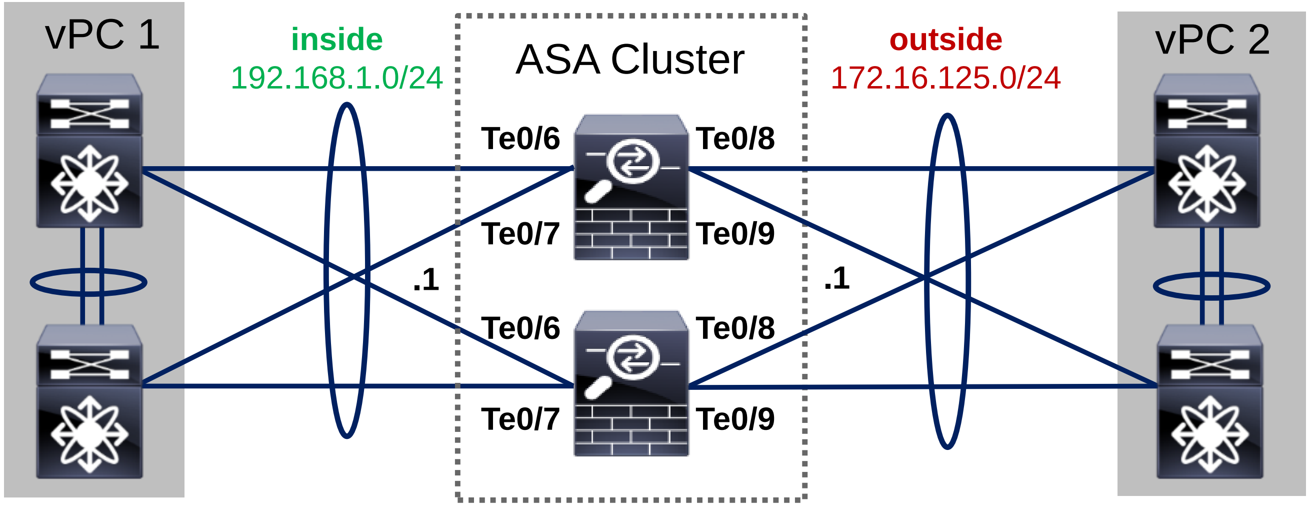

Clustering LACP (cLACP)

- Spanned Etherchannel is recommended for **data** interfaces on ASA appliances
  - Up to 8 active and 8 standby links in **9.0/9.1** with dynamic port priorities in vPC/VSS

```
asa(config)# interface TenGigabitEthernet 0/8
asa(config-if)# channel-group 1 mode active vss-id 1
```

  - Up to 32 active total (up to 16 per unit) links with global static port priorities in **9.2(1)+**

```
asa(config)# cluster group DC_ASA
asa(cfg-cluster)# clacp static-port-priority
```

  - Always configure virtual MAC addresses for each Etherchannel to avoid instability
  - **Disable** LACP Graceful Convergence and Adaptive Hash on adjacent NX-OS
- Supervisor bundles data and CCL interfaces on Firepower 4100 and 9300
- cLACP **assumes** each Spanned Etherchannel connects to one logical switch
  - LACP actor IDs between member ports are not strictly enforced, allowing creativity

Individual Interface Mode

- **Not supported** on Firepower 4100 or 9300; **routed** firewalls only elsewhere
- Master owns virtual IP on data interfaces for management purposes only
- All members get data interface IPs from the pools in the order of admittance
- Per-unit Etherchannels support up to 16 members in **9.2(1)+**

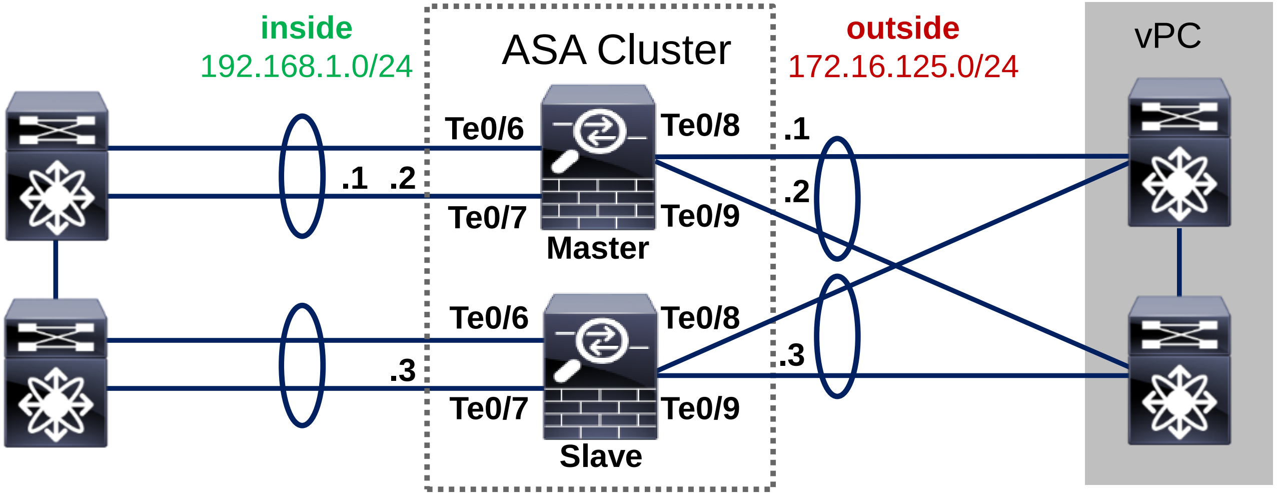

Traffic Load Balancing in Individual Mode

- Each unit has a separate IP/MAC address pair on its data interfaces
  - Traffic load-balancing is not as seamless as with Spanned Etherchannel mode
- **Policy Based Routing** (**PBR**) with route maps is very static by definition
  - Simple per-flow hashing or more elaborate distribution using ACLs
  - Difficult to direct return connections with NAT/PAT
  - Must use SLA with Object Tracking to detect unit addition and removal
  - Nexus **Intelligent Traffic Director** (**ITD**) simplifies configuration process
- Dynamic routing with **Equal Cost Multi Path** (**ECMP**)
  - Per-flow hashing with no static configuration
  - Easier to detect member addition and removal
  - Preferred approach with some convergence caveats

Dynamic Routing

- Master unit runs dynamic routing in Spanned Etherchannel mode
  - RIP, EIGRP, OSPFv2, OSPFv3, and PIM
  - BGPv4 in **ASA 9.3(1)+** and BGPv6 in **ASA 9.3(2)+**
  - Routing and ARP tables are synchronized to other members, like in failover
  - Possible external convergence impact **only** on Master failure
- Each member forms independent adjacencies in Individual mode
  - Same protocols as in Spanned Etherchannel, but multicast data is **centralized** as well
  - Higher overall processing impact from maintaining separate routing tables
  - Slower external convergence on any member failure

Non Stop Forwarding (NSF)

- Routing Information Base (RIB) is replicated in Spanned Etherchannel mode
  - Master establishes dynamic routing adjacencies and keeps Slaves up-to-date
  - When Master fails, the cluster continues traffic forwarding based on RIB
  - New Master re-establishes the dynamic routing adjacencies and updates the RIB
  - Adjacent routers flush routes and cause momentary traffic blackholing
- Non Stop Forwarding (NSF) and Graceful Restart (GR) support in **ASA 9.3(1)+**
  - New Master notifies compatible peer routers in Spanned Etherchannel clustering
  - Master acts as a helper to support a restart of the peer router in all modes

1. Cluster Master fails; new Master initiates adjacency with the peer router indicating that traffic forwarding should continue.

4. ASA cluster continues normal traffic forwarding until the primary RP restarts or the backup takes over or the timeout expires.

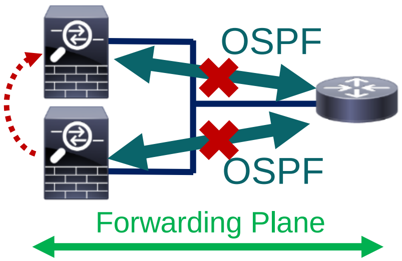

2. Router re-establishes adjacency with Master while retaining the stale routes; these routes are refreshed when the adjacency reestablishes.

3. Primary Route Processor undergoes a restart, signals the peer cluster to continue forwarding while the backup re-establishes adjacencies.

NSF and GR Configuration

Reference

- Feature has to be enabled on all adjacent devices to work
- Use Cisco with all Cisco peers (default) or IETF NSF with third-party for OSPFv2

```
router ospf 1
 nsf cisco enforce-global
 nsf cisco helper
```

(Optional) Disable NSF if any adjacent device is incompatible.

(Default) Help other NSF devices restart gracefully.

```
router ospf 1
 nsf ietf restart-interval 260
 nsf ietf helper strict-lsa-checking
```

Default graceful restart time is 120 seconds.

(Optional) Helper aborts peer’s NSF restart on impactful LSA changes

- Common Graceful Restart configuration for OSPFv3

```
router ospf 1
 graceful-restart restart-interval 180
 graceful-restart helper strict-lsa-checking
```

- BGPv4 Graceful Restart is enabled globally and configured for each neighbor

```
! System Context
router bgp 65001
 bgp graceful-restart restart-time 180 stalepath-time 720
! Context A
router bgp 65001
 address-family ipv4 unicast
  neighbor 192.168.1.101 ha-mode graceful-restart
```

Default maximum wait time for a restarting peer is 120 seconds.

Default wait time before flushing routes toward a GR capable peer is 360 seconds.

Enable GR for each neighbor.

Dynamic Routing Convergence Optimization

- Reduce protocol timers on **all connected segments** to speed up convergence
  - OSPF timers **must** match between peers
  - **Do not** lower dead interval in Spanned Etherchannel mode with NSF/GR
- **ASA 9.1 and earlier** software uses higher minimum timers

```
asa(config)# interface GigabitEthernet0/0
asa(config-if)# ospf hello-interval 1
asa(config-if)# ospf dead-interval 3
asa(config-if)# router ospf 1
asa(config-router)# timers spf 1 1
```

Generate OSPF hello packets every 1 second

Declare neighbor dead with no hello packets for 3 seconds

Delay before and between SPF calculations for 1 **second**

- **ASA 9.2(1)+** provides faster convergence

```
asa(config)# interface GigabitEthernet0/0
asa(config-if)# ospf dead-interval minimal hello-multiplier 3
asa(config-if)# router ospf 1
asa(config-router)# timers throttle spf 500 1000 5000
```

Generate 3 OSPF FastHello packets per second; 1 second to detect a dead neighbor

Delay SPF calculation by 500 **ms**, delay between calculations for 1 second and no more than 5 seconds

Verifying Load Distribution

- Uneven Owner connection distribution implies a load-balancing issue
    - Use a more granular Etherchannel hashing algorithm on connected switches
- High Forwarder connection count implies flow asymmetry
    - Always match Etherchannel hashing algorithms between all connected switches
    - Cannot avoid asymmetry with NAT/PAT

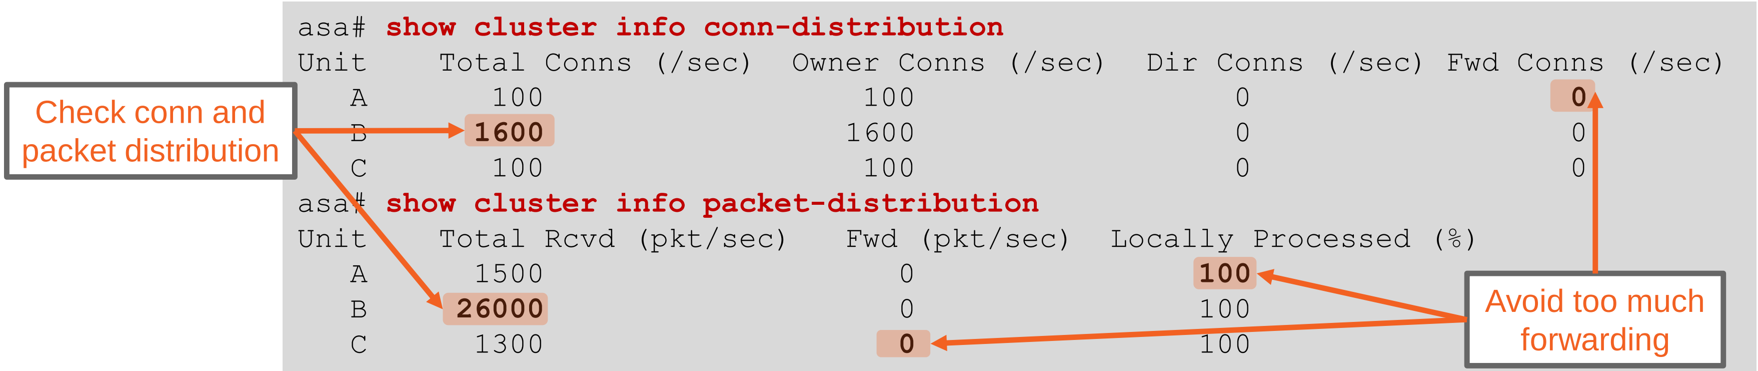

```
asa# show cluster info conn-distribution
Unit     Total Conns (/sec)   Owner Conns (/sec)   Dir Conns (/sec)  Fwd Conns (/sec)
   A           100                   100                  0                 0
   B          1600                  1600                  0                 0
   C           100                   100                  0                 0
asa# show cluster info packet-distribution
Unit     Total Rcvd (pkt/sec)    Fwd (pkt/sec)  Locally Processed (%)
   A          1500                     0                 100
   B         26000                     0                 100
   C          1300                     0                 100
```

Check conn and packet distribution

Avoid too much forwarding

Management Interface

- Any regular data interface can be used for managing the cluster
  - Always connect to virtual IP to reach the Master and make configuration changes
  - **cluster exec** allows to execute non-configuration commands on all members

```
asa/master# cluster exec show version | include Serial
A(LOCAL):*************************************************************
Serial Number: JAF1434AERL

B:********************************************************************
Serial Number: JAF1511ABFT
```

  - Units use same IP in Spanned Etherchannel mode for syslog and NSEL
- Dedicated management interface is recommended to reach all units
  - **management-only** allows MAC/IP pools even in Spanned Etherchannel mode
  - Some monitoring tasks requires individual IP addressing (such as SNMP polling)
  - No dynamic routing support, only static routes

Health Monitoring

- CCL link loss causes unit to shut down all data interfaces and disable clustering
  - Clustering **must** be re-enabled manually after such an event until **ASA 9.5(1)**
- Each member generates keepalives on CCL every 1 second by default
  - Master will remove a unit from the cluster after 3 missed keepalives (holdtime)
  - Member leaves cluster if its interface/SSP is "down" and another member has it "up"
  - Re-join attempted 3 times (after 5, 10, 20 minutes), then the unit disables clustering
- Disable health check during changes and tune other parameters

```
a/master# cluster group sjfw
a/master(cfg-cluster)# no health-check
a/master(cfg-cluster)# health-check holdtime 1
a/master(cfg-cluster)# no health-check monitor-interface Management0/0
a/master(cfg-cluster)# health-check cluster-interface auto-rejoin 5 1 1
a/master(cfg-cluster)# health-check data-interface auto-rejoin 10 2 1
```

Keepalive is always 1/3 of the configured holdtime

Added in **ASA 9.4(1)**

Configurable **re-join attempts**, **interval**, and **interval multiplier** in **9.5(1)**

Configuring Clustering on ASA Appliances

Preparation Checklist

- Get **serial console** access to all future cluster members
- Clear the existing configuration and configure appropriate boot images
- Switch to the multiple-context mode if desired
- Install Cluster (ASA5580/5585-X) and matching 3DES/10GE I/O licenses
- Designate a dedicated management interface (same on all members)
- Designate one or more physical interfaces per unit for CCL
- Assign an isolated subnet for CCL on a separate switch or VDC
- Configure **jumbo-frame reservation** command and reload each ASA
- Pick Spanned Etherchannel or Individual interface mode for the entire cluster

Setting Interface Mode

- Use **cluster interface-mode** command before configuring clustering
  - The running configuration is checked for incompatible commands
  - Interface mode setting is stored outside of the startup configuration
  - Use **show cluster interface-mode** to check current mode
  - Use **no cluster interface-mode** to return to standalone mode
- Clearing the interface configuration and reloading each ASA is **recommended**
  - You can display the list of conflicts and resolve them manually

```
asa(config)# cluster interface-mode spanned check-details
ERROR: Please modify the following configuration elements that are incompatible with
'spanned' interface-mode.
 - Interface Gi0/0 is not a span-cluster port-channel interface, Gi0/0(outside)
cannot be used as data interface when cluster interface-mode is 'spanned'.
```

  - It is **not recommended** to bypass the check and force the mode change

Establishing Management Access

- Start clustering configuration on the Master unit
- ASDM High Availability and Scalability Wizard simplifies deployment
  - Only set the interface mode on Master, then add Slaves automatically over HTTPS
  - Requires basic management connectivity to all members

```
ip local pool CLUSTER_MANAGEMENT 172.16.162.243-172.16.162.250
!
interface Management0/0
 description management interface
 management-only
 nameif mgmt
 security-level 0
 ip address 172.16.162.242 255.255.255.224 cluster-pool CLUSTER_MANAGEMENT
!
route mgmt 0.0.0.0 0.0.0.0 172.16.162.225 1
http server enable
http 0.0.0.0 0.0.0.0 mgmt
aaa authentication http console LOCAL
username cisco password cisco privilege 15
```

Master: Management IP address pool for all units; do **not** configure on Slaves

Dedicated management interface allows individual IP addressing in all modes

Master: Configure the IP pool under management interface
Slaves: Use individual IP addresses from the pool (starting from **.244** in this example) on the same management interfaces

ASDM High Availability and Scalability Wizard

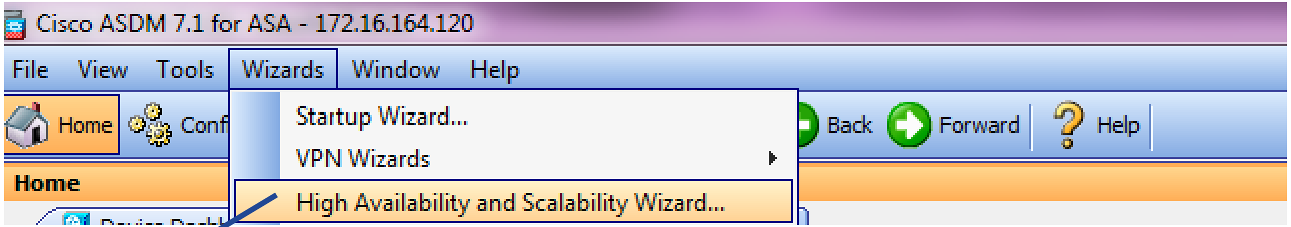

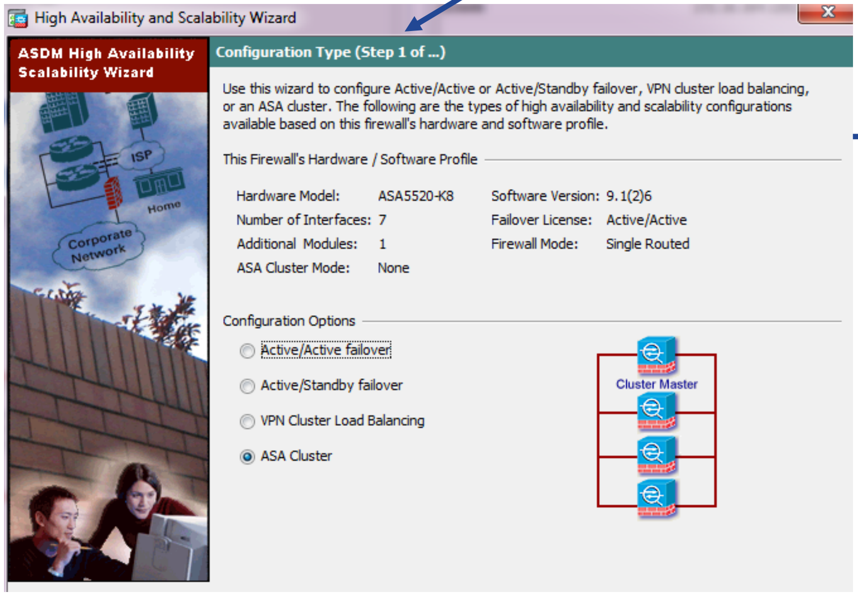
High Availability and Scalability Wizard
ASDM High Availability Scalability Wizard
ISP
Home
Corporate Network
Cluster Master

Configuration Type (Step 1 of ...)

Use this wizard to configure Active/Active or Active/Standby failover, VPN cluster load balancing, or an ASA cluster. The following are the types of high availability and scalability configurations available based on this firewall's hardware and software profile.

This Firewall's Hardware / Software Profile

Hardware Model: ASA5520-K8
Number of Interfaces: 7
Additional Modules: 1
ASA Cluster Mode: None
Software Version: 9.1(2)6
Failover License: Active/Active
Firewall Mode: Single Routed

Configuration Options

○ Active/Active failover
○ Active/Standby failover
○ VPN Cluster Load Balancing
● ASA Cluster

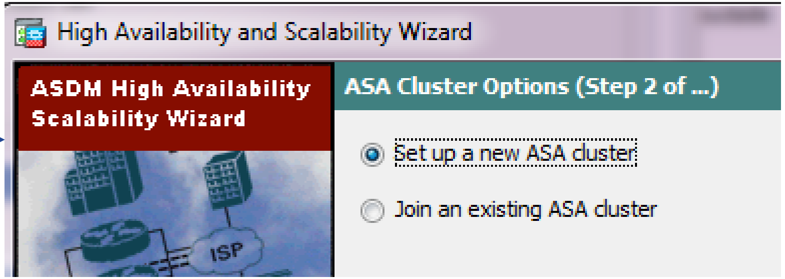
High Availability and Scalability Wizard
ASDM High Availability Scalability Wizard
ISP

ASA Cluster Options (Step 2 of ...)

● Set up a new ASA cluster
○ Join an existing ASA cluster

Fully configure Master in 4 easy steps, then have ASDM add Slaves one by one over basic HTTPS management connection.

… or use good old CLI ;-)

CLI Configuration: CCL Etherchannel

- Create an Etherchannel interface for CCL on each member separately
  - Same physical interface members across all units
  - Use LACP for quicker failure detection or static **on** mode for less complexity
  - Use system context in the multiple-context mode
  - Connect one physical interface to each logical switch in VSS/vPC

```
ciscoasa(config)# interface TenGigabitEthernet 0/6
ciscoasa(config-if)# channel-group 1 mode on
INFO: security-level, delay and IP address are cleared on TenGigabitEthernet0/6.
ciscoasa(config-if)# no shutdown
ciscoasa(config-if)# interface TenGigabitEthernet 0/7
ciscoasa(config-if)# channel-group 1 mode on
INFO: security-level, delay and IP address are cleared on TenGigabitEthernet0/7.
ciscoasa(config-if)# no shutdown
```

CLI Configuration: Cluster Group

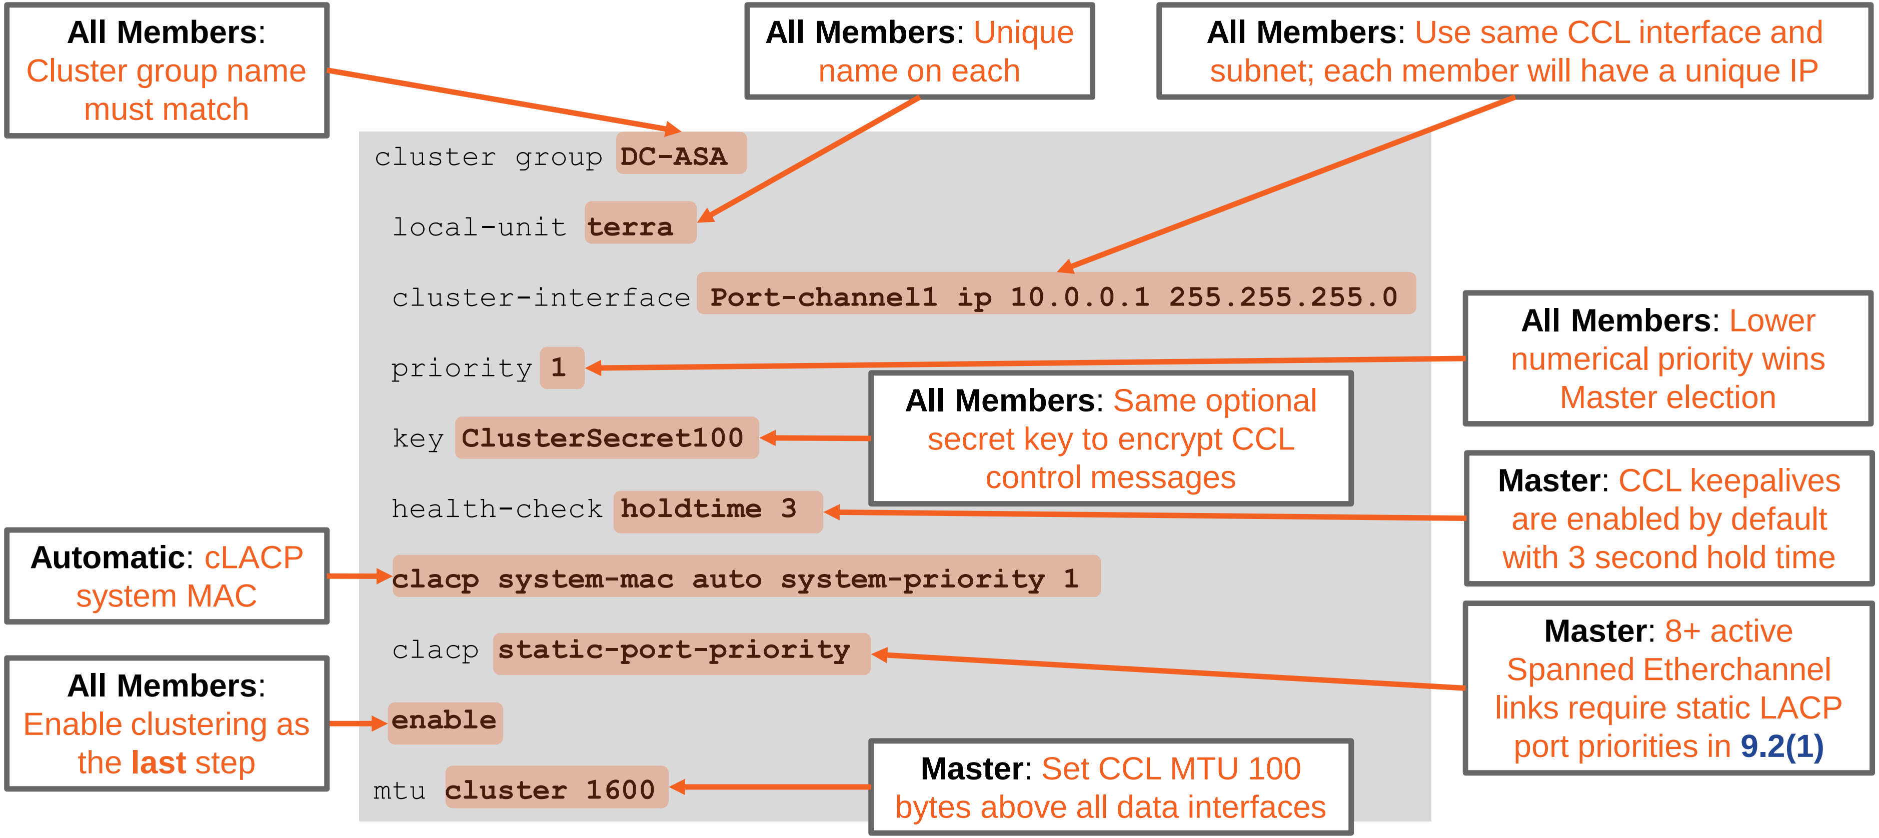

```
cluster group DC-ASA
 local-unit terra
 cluster-interface Port-channel1 ip 10.0.0.1 255.255.255.0
 priority 1
 key ClusterSecret100
 health-check holdtime 3
 clacp system-mac auto system-priority 1
 clacp static-port-priority
 enable
mtu cluster 1600
```

All Members: Cluster group name must match

All Members: Unique name on each

All Members: Use same CCL interface and subnet; each member will have a unique IP

All Members: Lower numerical priority wins Master election

All Members: Same optional secret key to encrypt CCL control messages

Master: CCL keepalives are enabled by default with 3 second hold time

Automatic: cLACP system MAC

Master: 8+ active Spanned Etherchannel links require static LACP port priorities in **9.2(1)**

All Members: Enable clustering as the **last** step

Master: Set CCL MTU 100 bytes above all data interfaces

Cisco live!

CLI Configuration: Data Interfaces on Master

Spanned Etherchannel Mode

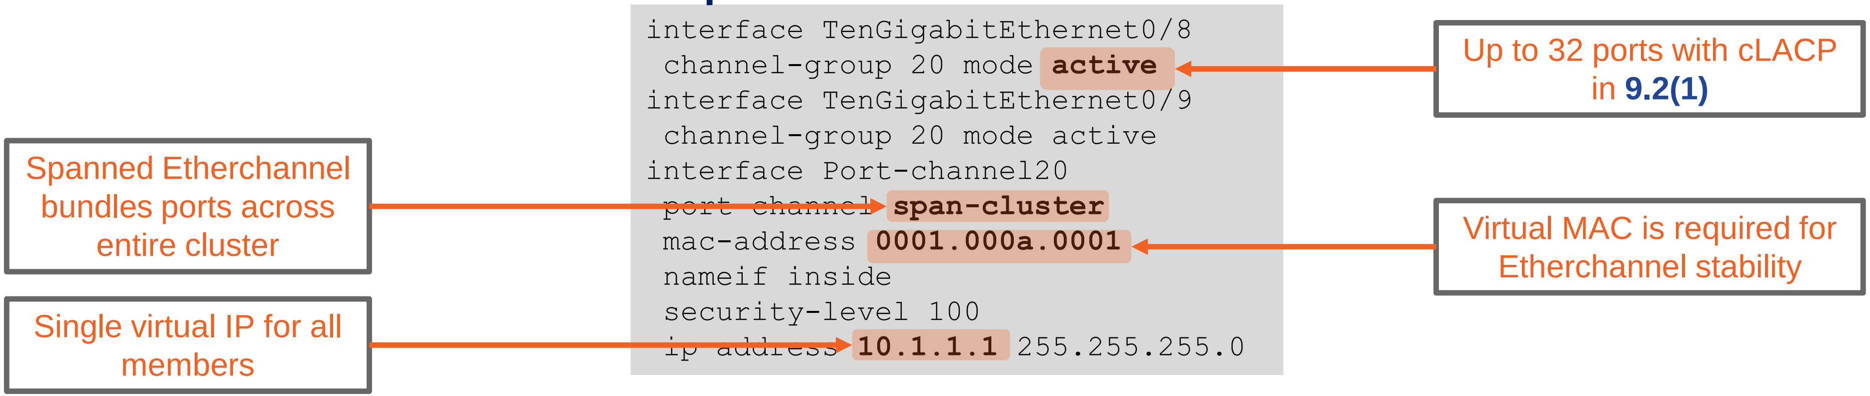

```
interface TenGigabitEthernet0/8
 channel-group 20 mode active
interface TenGigabitEthernet0/9
 channel-group 20 mode active
interface Port-channel20
 port-channel span-cluster
 mac-address 0001.000a.0001
 nameif inside
 security-level 100
 ip address 10.1.1.1 255.255.255.0
```

Up to 32 ports with cLACP in **9.2(1)**

Spanned Etherchannel bundles ports across entire cluster

Virtual MAC is required for Etherchannel stability

Single virtual IP for all members

Individual Mode

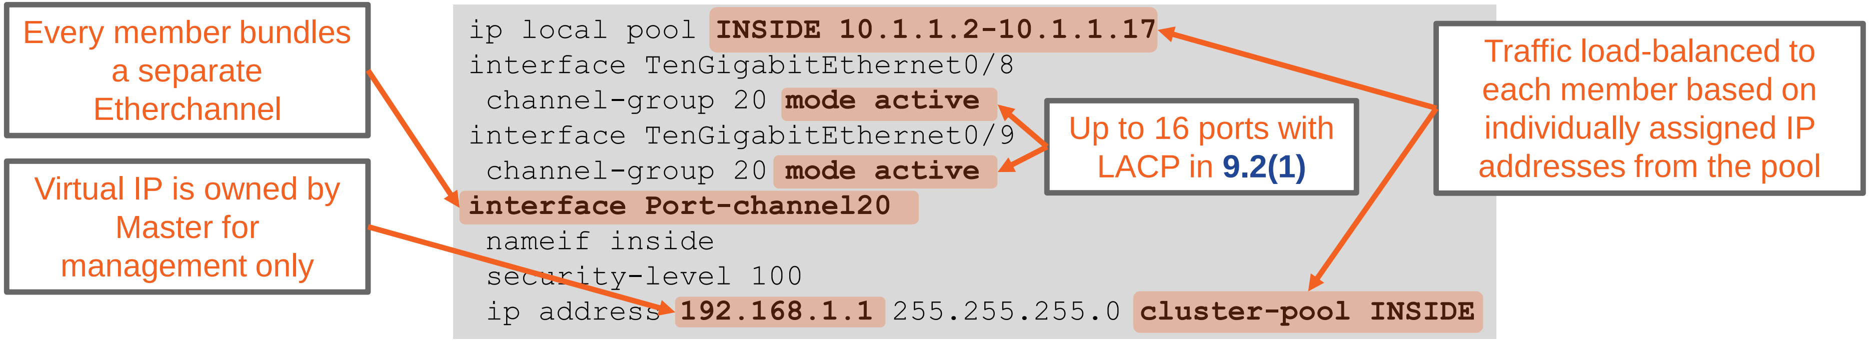

```
ip local pool INSIDE 10.1.1.2-10.1.1.17
interface TenGigabitEthernet0/8
 channel-group 20 mode active
interface TenGigabitEthernet0/9
 channel-group 20 mode active
interface Port-channel20
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0 cluster-pool INSIDE
```

Every member bundles a separate Etherchannel

Up to 16 ports with LACP in **9.2(1)**

Traffic load-balanced to each member based on individually assigned IP addresses from the pool

Virtual IP is owned by Master for management only

CLI Configuration: Adding Slave Units

- Verify that the Master is operational before adding Slave members

```
asa# show cluster info
Cluster DC-ASA: On
    Interface mode: spanned
    This is “terra" in state MASTER
        ID         : 1
        Version    : 9.1(3)
        Serial No.: JAF1511ABFT
        CCL IP     : 10.0.0.1
        CCL MAC    : 5475.d05b.26f2
        Last join : 17:20:24 UTC Sep 26 2013
        Last leave: N/A
```

- Add one Slave at a time by configuring the cluster group

```
cluster group DC-ASA
 local-unit sirius
 cluster-interface Port-channel1 ip 10.0.0.2 255.255.255.0
 priority 100
 key ClusterSecret100
 enable
```

Spanned Etherchannel Verification

- Each cluster member shows only local Etherchannel member ports

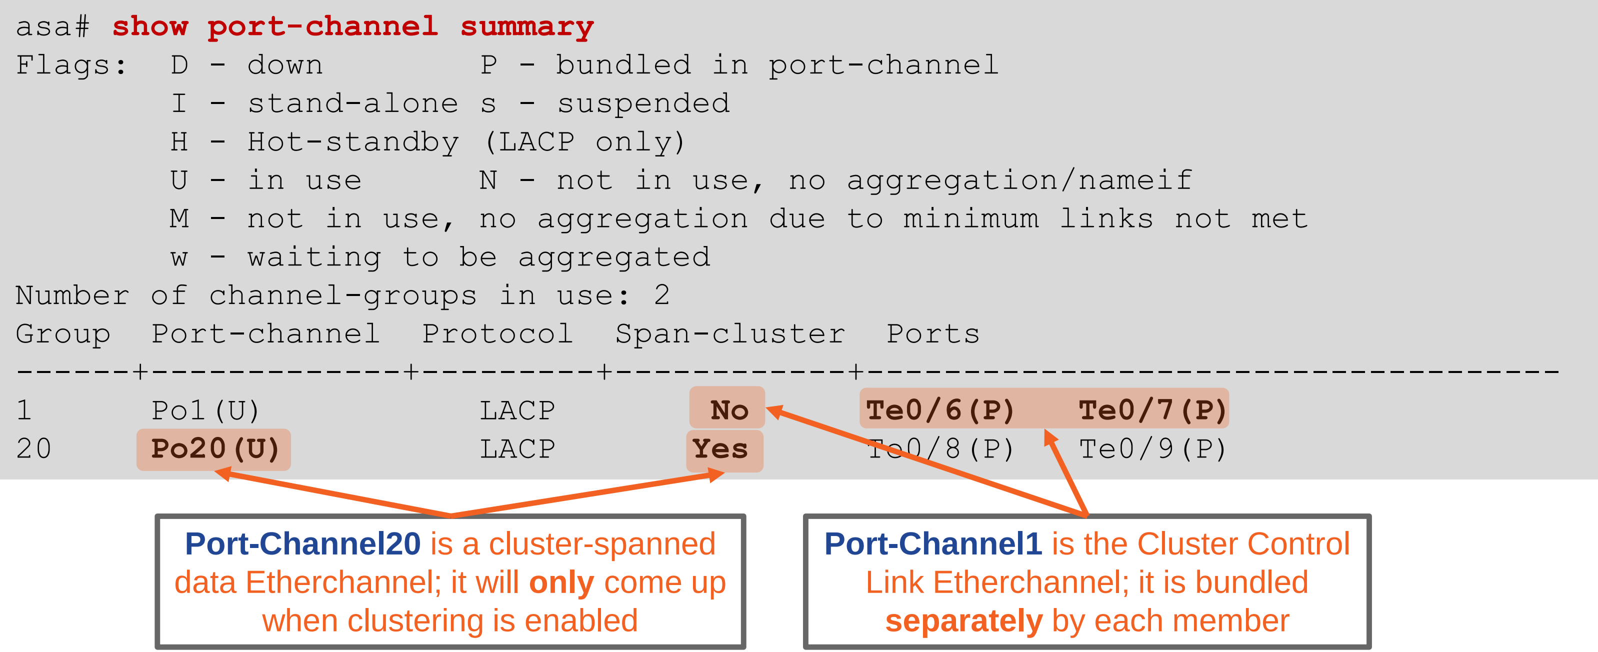

```
asa# show port-channel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        U - in use      N - not in use, no aggregation/nameif
        M - not in use, no aggregation due to minimum links not met
        w - waiting to be aggregated
Number of channel-groups in use: 2
Group  Port-channel  Protocol  Span-cluster  Ports
------+-------------+---------+------------+------------------------------------
1      Po1(U)            LACP          No     Te0/6(P)   Te0/7(P)
20     Po20(U)           LACP         Yes     Te0/8(P)   Te0/9(P)
```

Port-Channel20 is a cluster-spanned data Etherchannel; it will **only** come up when clustering is enabled

Port-Channel1 is the Cluster Control Link Etherchannel; it is bundled **separately** by each member

Monitoring and Troubleshooting Clustering

- ASDM Clustering dashboard shows aggregated health information
- **show cluster** command group displays aggregated traffic and resource data
  - **show cluster history** helps to understand state transitions and failure reasons
  - **show cluster cpu** helps to check CPU utilization across cluster
- **show cluster info** command group displays cluster subsystem information
  - **show cluster info health** helps to monitor aggregated unit health data
  - **show cluster info loadbalance** relates to optional Conn Rebalance feature
  - **show cluster info trace** shows cluster state machine debug data for Cisco TAC
- Leverage syslogs to understand failure reasons

  ```
  %ASA-3-747022: Clustering: Asking slave unit terra to quit because it failed interface health
  check 3 times (last failure on Port-channel1), rejoin will be attempted after 20 min.
  ```

  - Use logging device-id to identity reporting members for connection events

Deploying Clustering on Firepower 4100 and 9300

- **Only** Spanned Etherchannel interface mode is supported
- Supervisor pushes clustering configuration during logical device deployment
  - Per-module cluster unit name and health checks are configurable
  - Inter-site clustering parameters must be configured manually on each module
  - Firewall context mode, 3DES/AES license, SSL ciphers are replicated
- Off-chassis flow backup for N+1 chassis-level fault tolerance on Firepower 9300
- Module- and chassis-level overflow protection syslogs

```
%ASA-6-748008: CPU load 80% of module 1 in chassis 1 (unit-1-1) exceeds overflow
protection threshold CPU 75%. System may be oversubscribed on member failure.
%ASA-6-748009: Memory load 80% of chassis 1 exceeds overflow protection threshold
memory 78%. System may be oversubscribed on chassis failure.
```

Multi-Site Clustering

Inter Data Center (DC) Clustering

- Clustering **assumes rather than requires** data interface adjacency at Layer 2
- Geographically separated clusters supported in **ASA 9.1(4)+**
  - "Dark Media" CCL with up to 10ms of one-way latency
  - No tolerance for packet re-ordering or loss
  - Routed firewall in Individual interface mode **only**
- **ASA 9.2(1)** extends inter-DC clustering support to Spanned Etherchannel mode
  - Transparent firewall **only**
  - Routed firewall support presents design challenges
- **ASA 9.5(1)** adds inter-site Spanned Etherchannel clustering in routed mode

Split or Single Individual Mode Cluster

ASA 9.2(1)

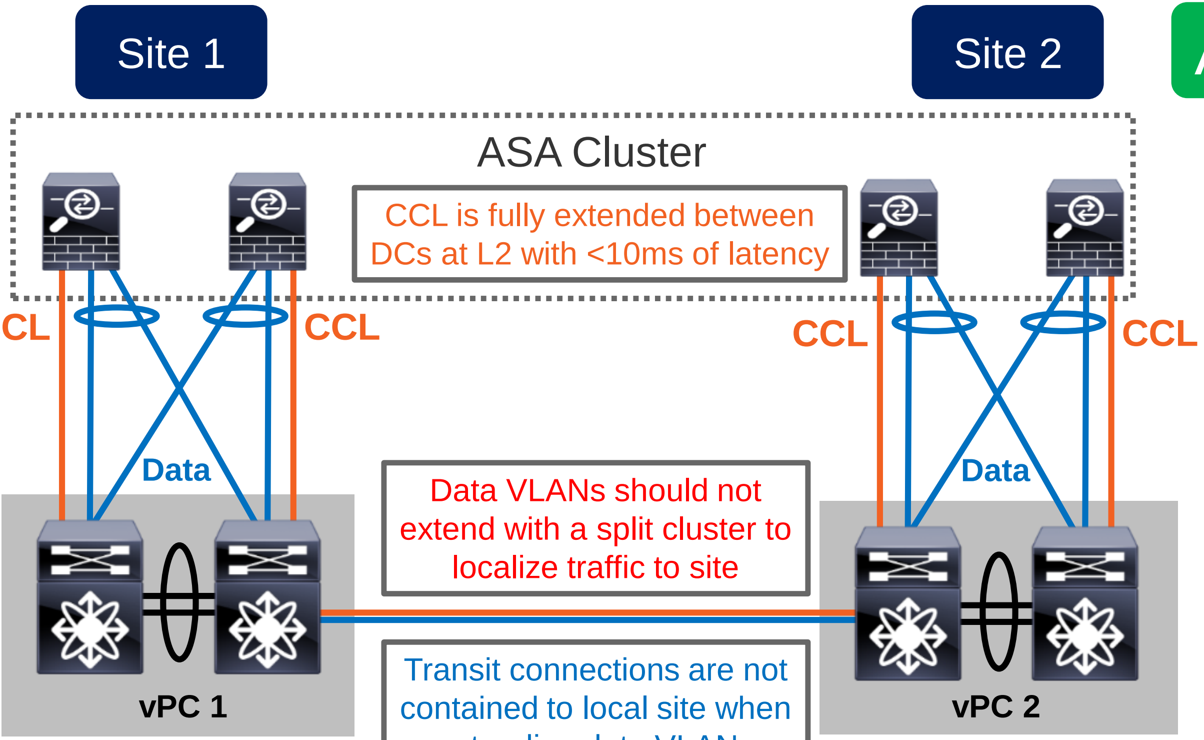

CCL is fully extended between DCs at L2 with <10ms of latency

Data interfaces connect to local switch pair only

Data VLANs should not extend with a split cluster to localize traffic to site

Transit connections are not contained to local site when extending data VLANs

Local vPC/VSS pairs at each site

Local vPC/VSS pairs at each site

Extended Spanned Etherchannel Cluster

Reference

ASA 9.1(4)

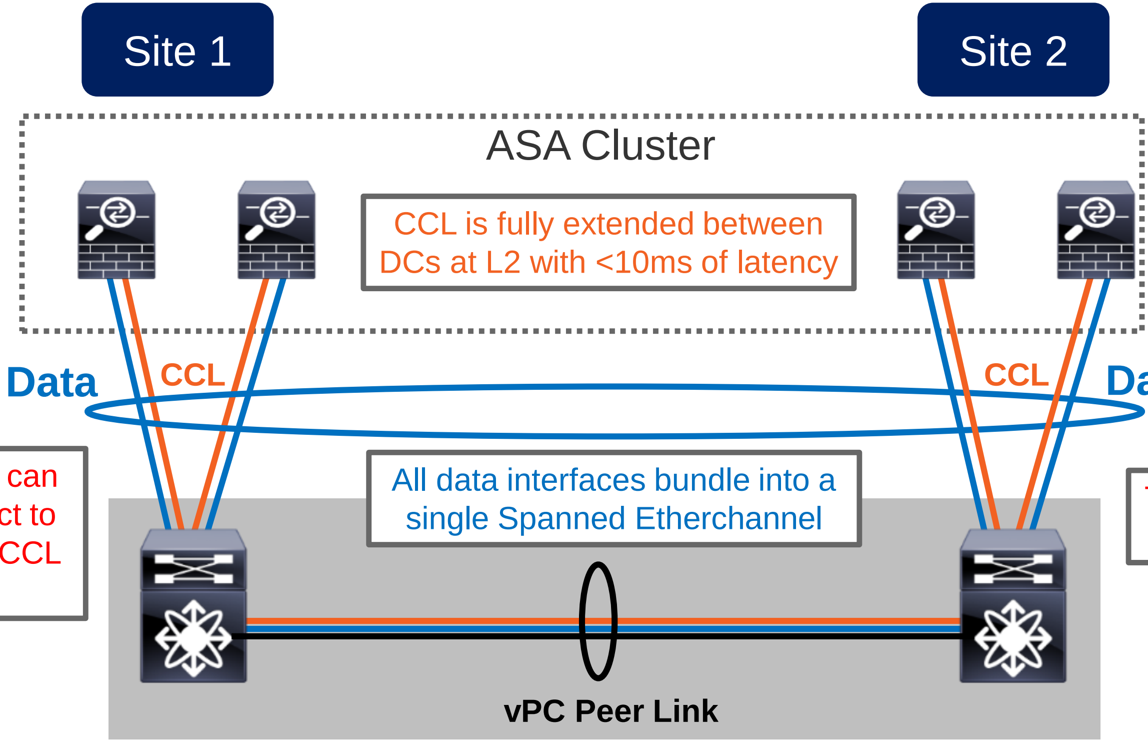

CCL is fully extended between DCs at L2 with <10ms of latency

All data interfaces bundle into a single Spanned Etherchannel

Each cluster member can single- or dual-connect to the VSS/vPC pair for CCL and Data

Transit connections are not contained to the local site

vPC logical switch pair is stretched across sites

Split Spanned Etherchannel Cluster

ASA 9.2(1)

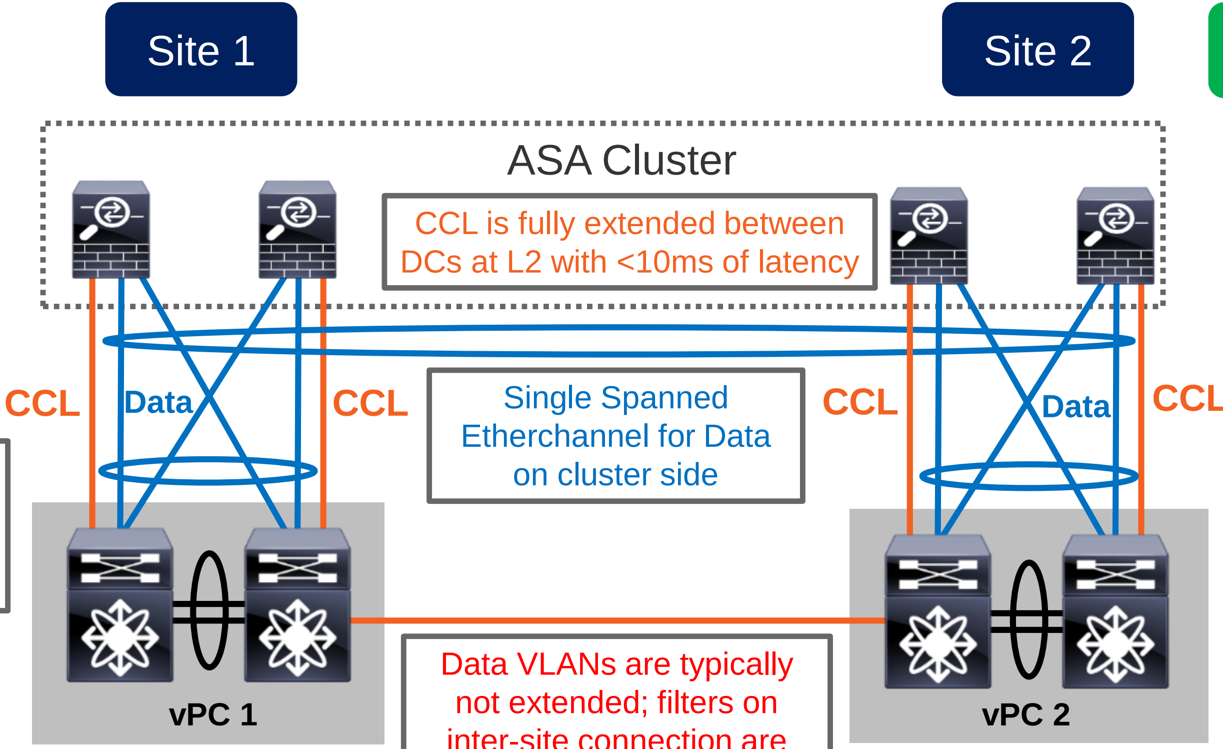

CCL is fully extended between DCs at L2 with <10ms of latency

Single Spanned Etherchannel for Data on cluster side

Local Data Etherchannels on each VPC/VSS switch pair

Local Data Etherchannels on each vPC/VSS switch pair

Local vPC/VSS pairs at each site

Data VLANs are typically not extended; filters on inter-site connection are needed to prevent loops and address conflicts

Local vPC/VSS pairs at each site

North-South Inter DC Clustering

ASA 9.1(4)

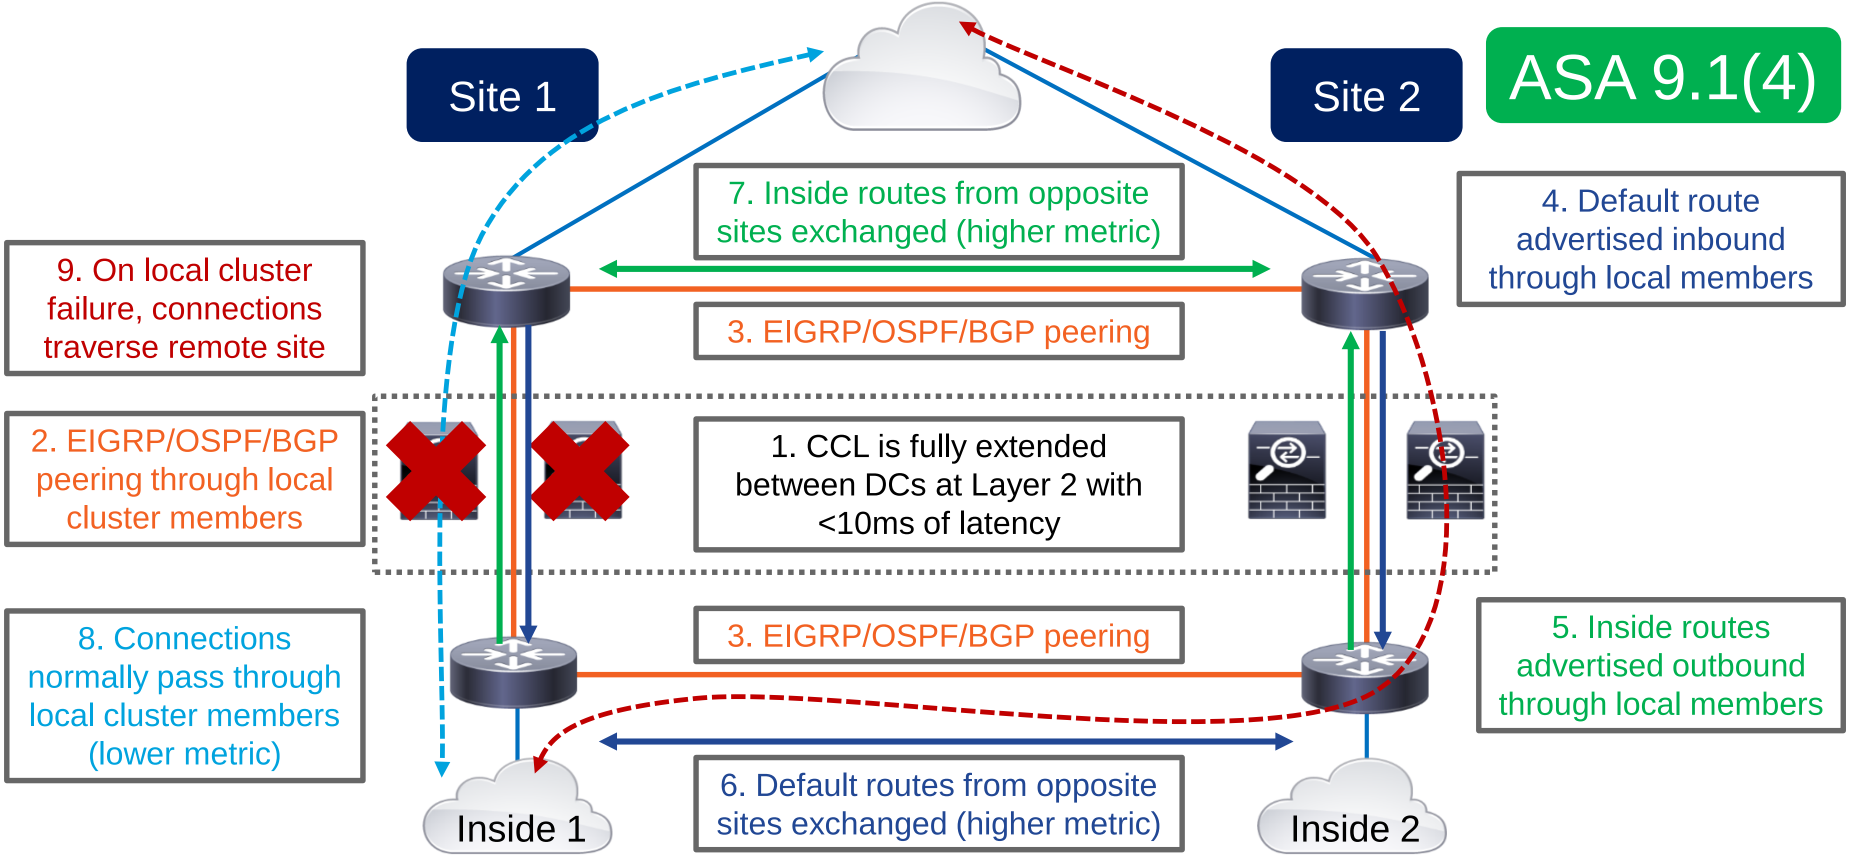

1. CCL is fully extended between DCs at Layer 2 with <10ms of latency

2. EIGRP/OSPF/BGP peering through local cluster members

3. EIGRP/OSPF/BGP peering

3. EIGRP/OSPF/BGP peering

4. Default route advertised inbound through local members

5. Inside routes advertised outbound through local members

6. Default routes from opposite sites exchanged (higher metric)

7. Inside routes from opposite sites exchanged (higher metric)

8. Connections normally pass through local cluster members (lower metric)

9. On local cluster failure, connections traverse remote site

Example: N-S Split Individual Mode Cluster

- A pair of standalone (non-vPC) Nexus switches at each site
  - One Individual mode cluster unit per switch, single attached
  - Routed firewall-on-a-stick VRF sandwich with OSPF
- Inside VLAN is fully extended between sites with OTV
  - Each pair of switches uses localized GLBP as first hop router
  - GLBP traffic is blocked between sites
  - OSPF allows re-routing in case of local cluster unit failure
- Traffic symmetry is achievable without NAT
  - Outbound connections use the directly attached cluster member
  - Inbound traffic requires LISP to eliminate tromboning due to ECMP

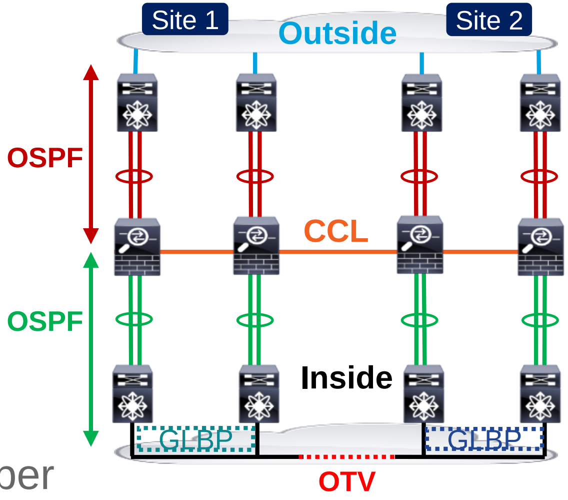

N-S Split Individual Cluster Sample Configuration

```
ip local pool OUTSIDE 192.168.2.2-
 192.168.2.17
interface Port-Channel10.20
 vlan 20
 nameif FW-outside
 ip address 192.168.2.1 255.255.255.0
  cluster-pool OUTSIDE
```

```
ip local pool OUTSIDE 192.168.1.2-
 192.168.1.17
interface Port-Channel10.10
 vlan 10
 nameif FW-inside
 ip address 192.168.1.1 255.255.255.0
  cluster-pool INSIDE
```

```
interface Ethernet3/1
 channel-group 1 mode active
interface Ethernet3/2
 channel-group 1 mode active
interface Port-Channel1
 switchport trunk allowed vlans 10,20
```

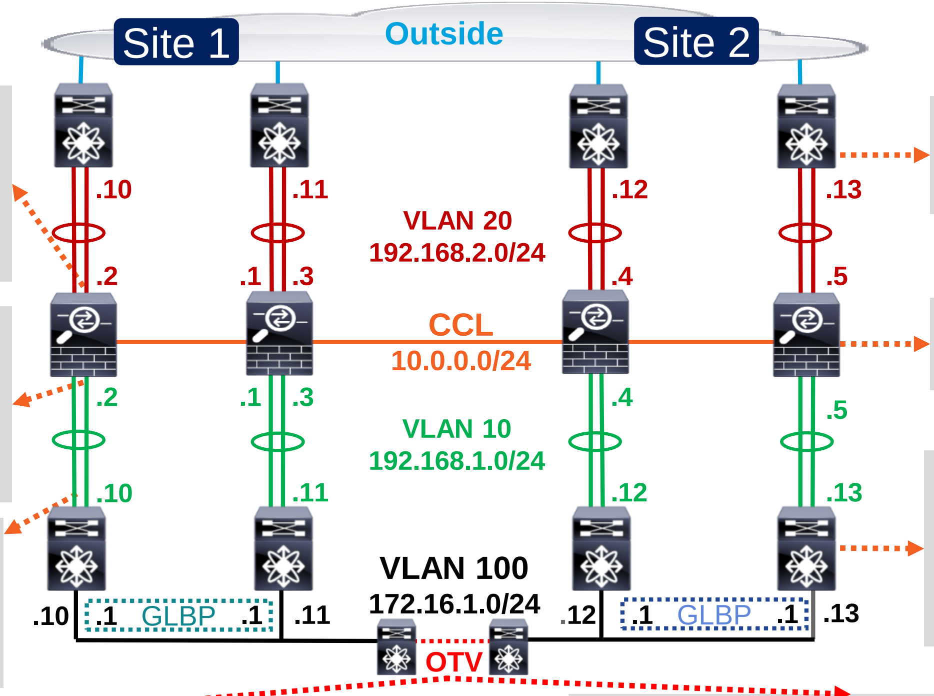

```
interface Vlan20
 vrf member OUTSIDE
 ip address 192.168.2.13/24
 ip router ospf 1 area 0.0.0.0
```

```
router ospf 1
 network 0.0.0.0 0.0.0.0 area
  0.0.0.0
```

```
interface Vlan10
 vrf member INSIDE
 ip address 192.168.1.13/24
 ip router ospf 2 area 0.0.0.0
interface Vlan100
 vrf member INSIDE
 ip router ospf 2 area 0.0.0.0
```

```
mac-list GLBP_FILTER seq 10 deny 0007.b400.0000 ffff.ffff.0000
mac-list GLBP_FILTER seq 20 permit 0000.0000.0000 0000.0000.0000
otv-isis default
 vpn Overlay1
  redistribute filter route-map GLBP_FILTER
```

OTV MAC Filter for GLBP

```
ip access-list NON_GLBP
 10 deny udp any 224.0.0.102/32 eq 3222
 20 permit ip any any
vlan access-map FILTER 10
 match ip address NON_GLBP
 action forward
vlan filter FILTER vlan-list 100
```

GLBP VLAN Filter

Reference

Example: N-S Split Spanned Etherchannel Cluster

- A vPC pair of Nexus switches at each site
  - Split Spanned Etherchannel cluster in transparent mode
  - Separate Etherchannel to local cluster members per vPC pair
  - VRF sandwich “through” the cluster with static PBR and SLA
- Non-overlapping inside subnets between sites
  - Mirrored SVI MAC addresses between two cluster transit VLANs
  - Dual-homed cluster members on each vPC pair localize traffic
  - Inter-site Layer 3 links (higher cost) to re-route traffic on failure
  - Bi-directional connection symmetry without NAT
  - Inbound asymmetry only between same-site members with NAT

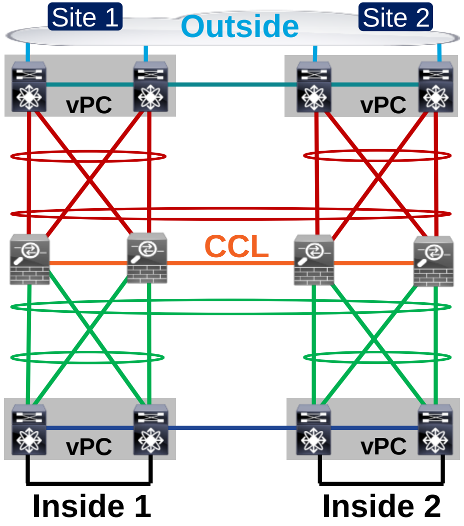

N-S Split Spanned Cluster Sample Configuration

Reference

```
ip sla 1
 icmp-echo 192.168.1.2
ip sla schedule 1 life forever start-
 time now
track 1 ip sla 1 reachability
ip access-list PBR
  permit ip any 172.16.1.0 255.255.255.0
route-map PBR
 match ip address PBR
 set ip next-hop verify-availability
  192.168.1.2 track 1
 set ip next-hop 192.168.4.2
interface Vlan300
 ip policy route-map PBR
```

```
interface Port-Channel10
 port-channel span-cluster
interface Port-Channel10.10
 vlan 10
 nameif FW-inside
 bridge-group 1
interface Port-Channel10.20
 vlan 20
 nameif FW-outside
 bridge-group 1
```

```
interface Ethernet3/1
 channel-group 1 mode active
interface Ethernet3/2
 channel-group 1 mode active
interface Port-Channel1
 switchport trunk allowed vlans 10,20
 vpc 10
```

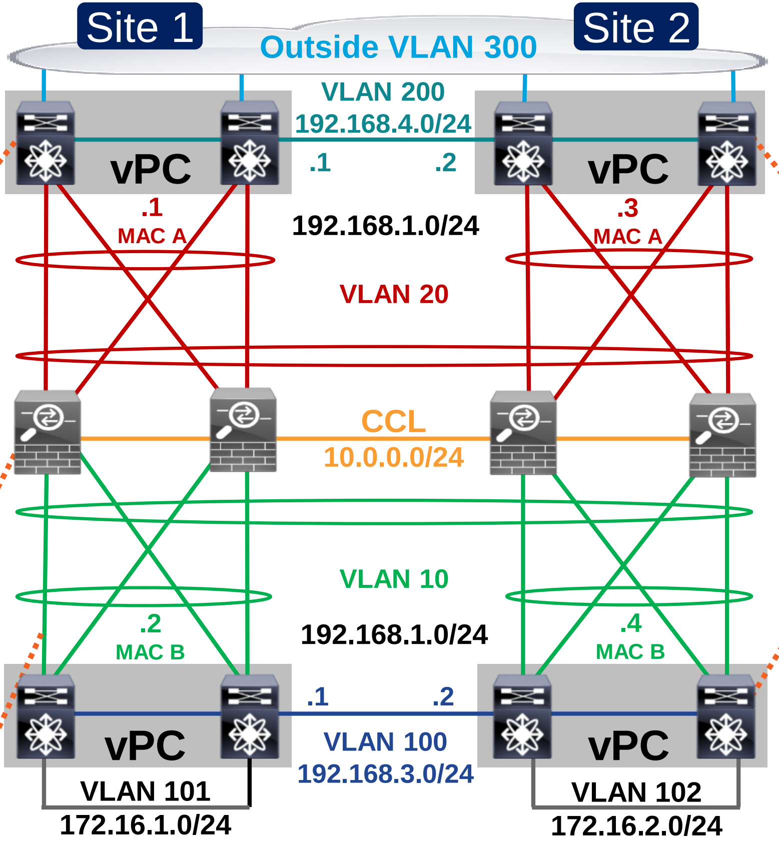

```
ip sla 1
 icmp-echo 192.168.1.4
ip sla schedule 1 life forever start-
 time now
track 1 ip sla 1 reachability
ip access-list PBR
  permit ip any 172.16.2.0 255.255.255.0
route-map PBR
 match ip address PBR
 set ip next-hop verify-availability
  192.168.1.4 track 1
 set ip next-hop 192.168.4.1
interface Vlan300
 ip policy route-map PBR
```

```
ip sla 1
 icmp-echo 192.168.1.3
ip sla schedule 1 life forever start-
 time now
track 1 ip sla 1 reachability
ip access-list PBR
  permit ip any any
route-map PBR
 match ip address PBR
 set ip next-hop verify-availability
  192.168.1.3 track 1
 set ip next-hop 192.168.3.1
interface Vlan102
 ip policy route-map PBR
```

Locator/Identifier Separation Protocol (LISP)

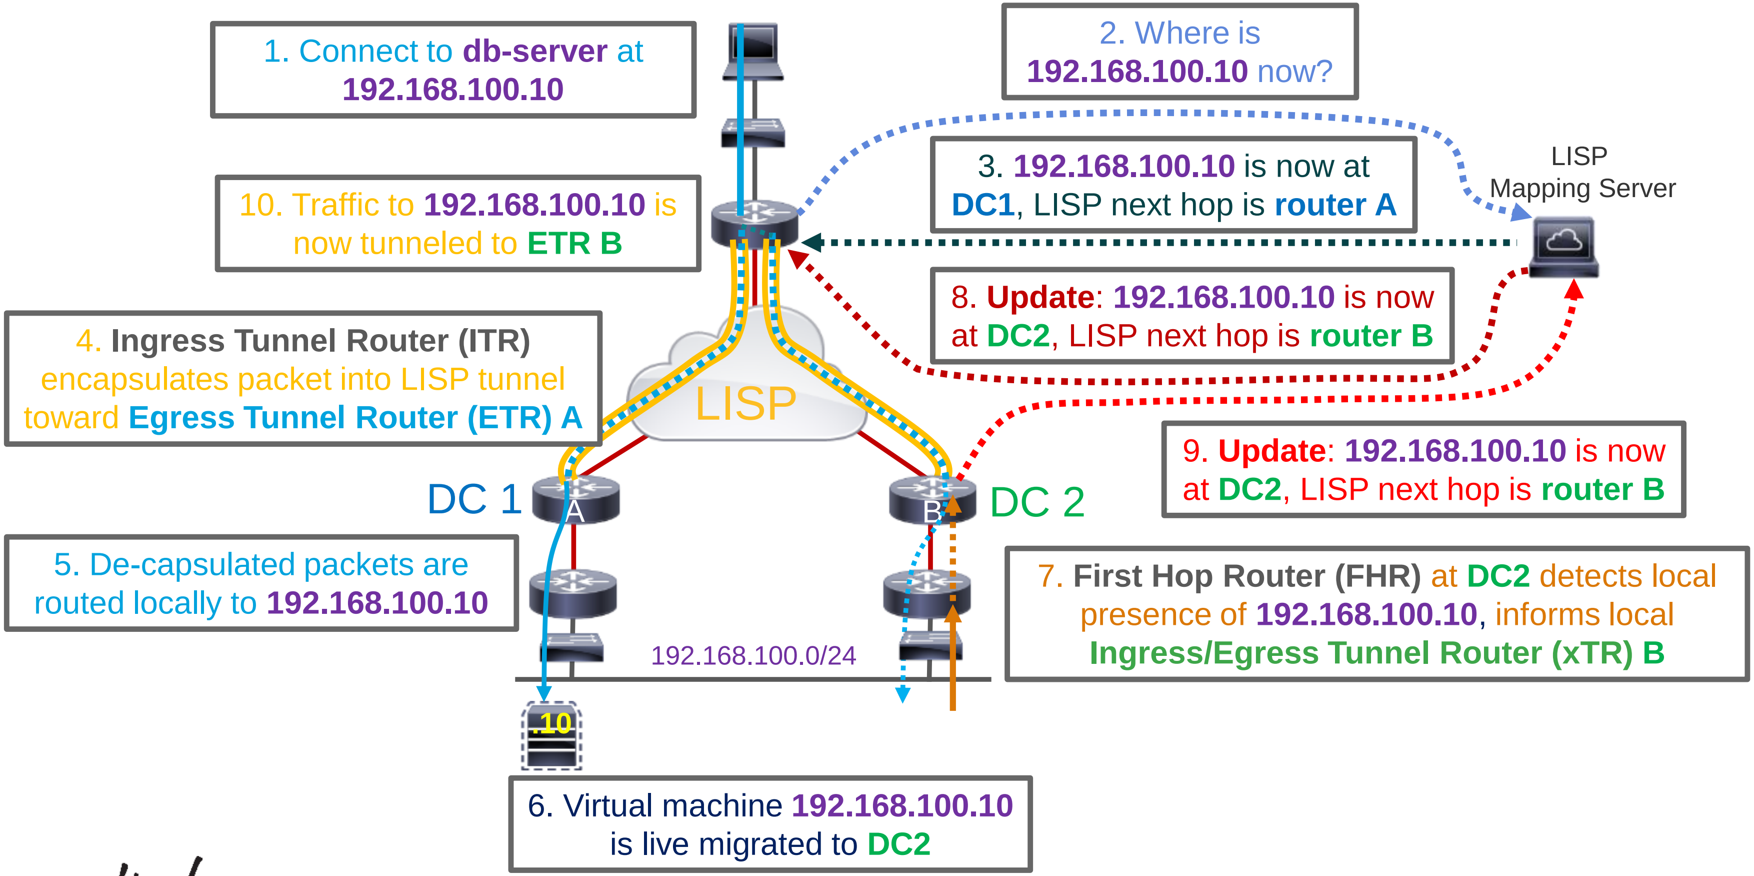

1. Connect to **db-server** at **192.168.100.10**

2. Where is **192.168.100.10** now?

3. **192.168.100.10** is now at **DC1**, LISP next hop is **router A**

4. **Ingress Tunnel Router (ITR)** encapsulates packet into LISP tunnel toward **Egress Tunnel Router (ETR) A**

5. De-capsulated packets are routed locally to **192.168.100.10**

6. Virtual machine **192.168.100.10** is live migrated to **DC2**

7. **First Hop Router (FHR)** at **DC2** detects local presence of **192.168.100.10**, informs local **Ingress/Egress Tunnel Router (xTR) B**

8. **Update**: **192.168.100.10** is now at **DC2**, LISP next hop is **router B**

9. **Update**: **192.168.100.10** is now at **DC2**, LISP next hop is **router B**

10. Traffic to **192.168.100.10** is now tunneled to **ETR B**

Dynamic Owner Reassignment with LISP

- Move flow ownership with VM in **ASA 9.5(2)**
  - Only supported with North-South clustering
  - Based on inspection of LISP FHR→xTR updates

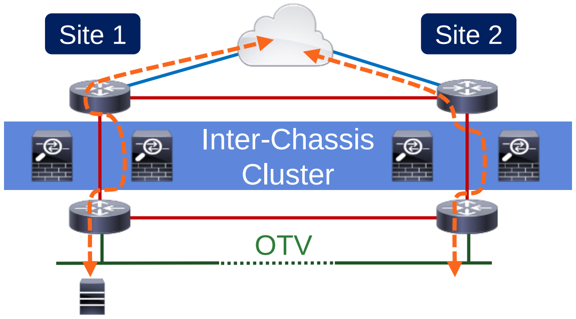

```
access-list MOBILITY_APP permit tcp any any eq 8443
class-map MOBILITY_APP
 match access-list MOBILITY_APP

cluster group DC-ASA
 site-id 2

policy-map global_policy
 class inspection_default
  inspect lisp
 class MOBILITY_APP
  cluster flow-mobility lisp
```

Select specific applications or flows that are eligible for owner reassignment

Up to 8 sites in a single cluster

UDP/4342 traffic is inspected for LISP by default

Other triggers for owner reassignment will be added in the future

Transparent East-West Inter DC Clustering

ASA 9.3(2)

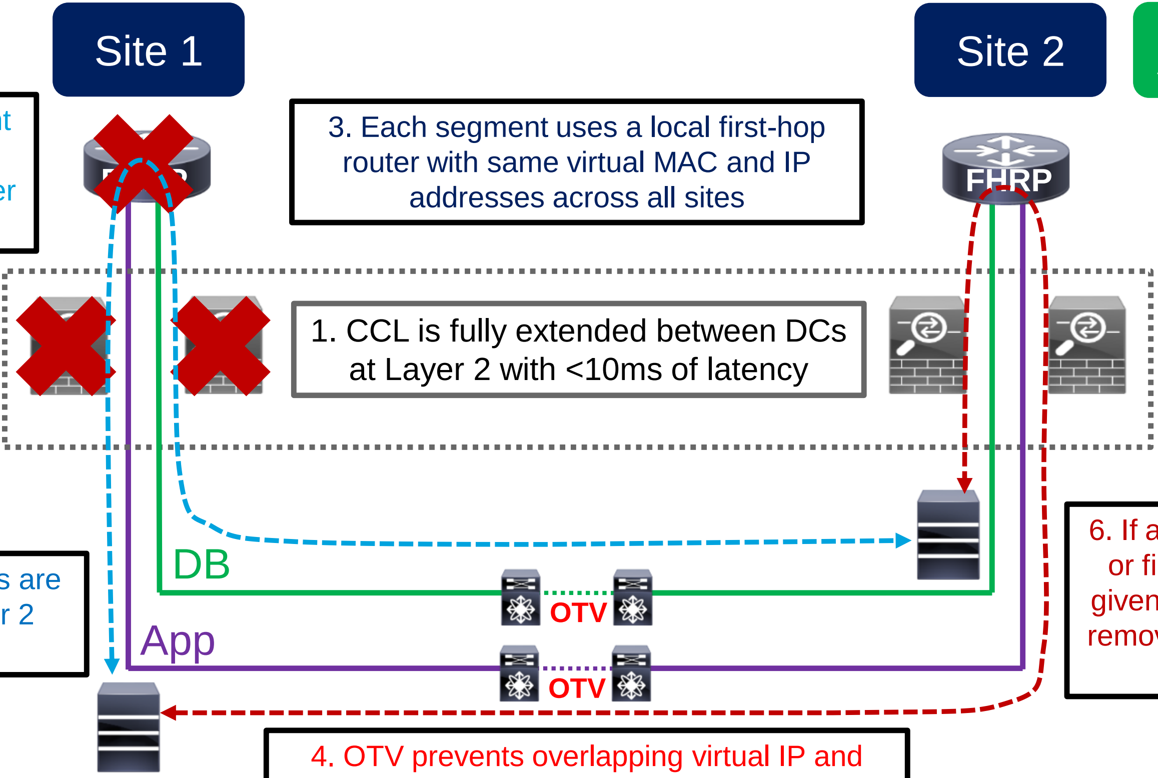

5. ASA cluster in transparent mode inserts between the endpoints and first-hop router on each segment

3. Each segment uses a local first-hop router with same virtual MAC and IP addresses across all sites

1. CCL is fully extended between DCs at Layer 2 with <10ms of latency

2. Protected data VLANs are fully extended at Layer 2 between sites

6. If all local cluster members or first-hop routers fail at a given site, OTV filter must be removed manually to fail over to another site

4. OTV prevents overlapping virtual IP and MAC addresses of the first-hop routers from leaking between sites

Reference

Example: E-W Transparent Spanned Cluster

- A vPC pair of Nexus switches at each site
  - Split Spanned Etherchannel cluster in transparent mode to separate internal segments
  - Separate Etherchannel to local cluster members per vPC pair
  - Acceptable impact from passing ASA twice between segments
- Internal VLANs are fully extended between sites with OTV
  - Each site uses localized HSRP as first hop router
  - HSRP traffic is blocked between sites
  - Full Layer 2 reachability from each router to remote site
  - OTV filters must be manually removed on full upstream path failure
- Must implement LISP to avoid excessive flow redirection

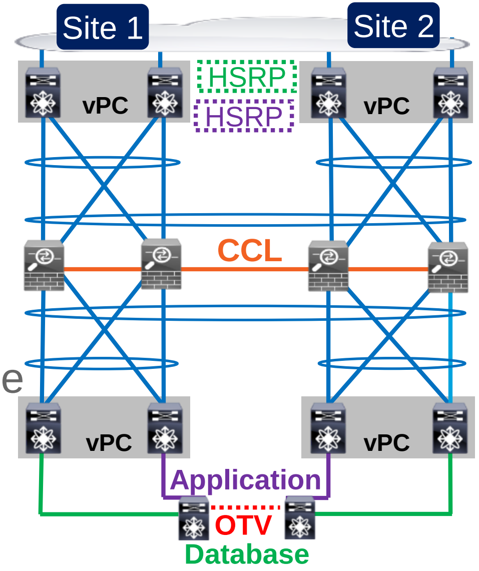

E-W Transparent Spanned Cluster Configuration

Reference

```
interface Vlan101
 ip address 192.168.1.2/24
 hsrp 10
  preempt
  ip 192.168.1.1
interface Vlan201
 ip address 192.168.2.2/24
 hsrp 20
  preempt
  ip 192.168.2.1
```

```
interface Port-Channel10
 port-channel span-cluster
interface Port-Channel10.100
 vlan 100
 nameif DB-inside
 bridge-group 1
interface Port-Channel10.101
 vlan 101
 nameif DB-outside
 bridge-group 1
interface Port-Channel10.200
 vlan 200
 nameif App-inside
 bridge-group 2
interface Port-Channel10.201
 vlan 201
 nameif App-outside
 bridge-group 2
interface BVI1
 ip address 192.168.1.4 255.255.255.0
interface BVI2
 ip address 192.168.2.4 255.255.255.0
```

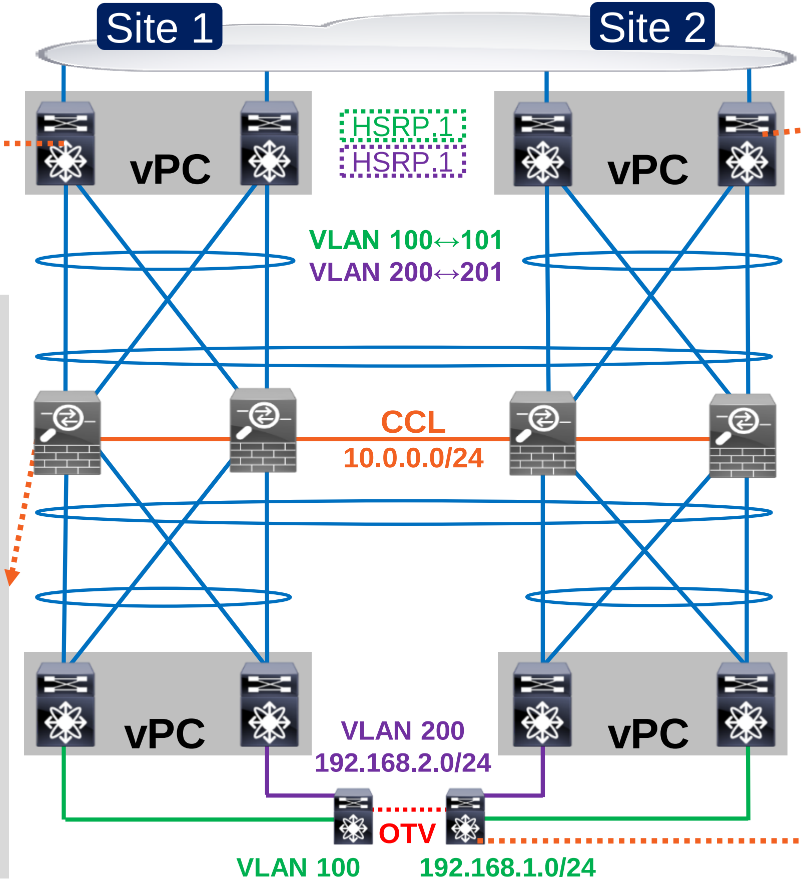

```
interface Vlan101
 ip address 192.168.1.3/24
 hsrp 10
  ip 192.168.1.1
interface Vlan201
 ip address 192.168.2.3/24
 hsrp 20
  ip 192.168.2.1
```

```
mac-list HSRP_FILTER seq 10 deny
    0000.0c07.ac00 ffff.ffff.ff00
mac-list HSRP_FILTER seq 20 deny
    0000.0c9f.f000 ffff.ffff.ff00
mac-list HSRP_FILTER seq 30 permit
    0000.0000.0000 0000.0000.0000
otv-isis default
 vpn Overlay1
  redistribute filter route-map HSRP_FILTER
!
ip access-list HSRP_TRAFFIC
 10 permit udp any 224.0.0.2/32 eq 1985
 20 permit udp any 224.0.0.102/32 eq 1985
ip access-list ALL
 10 permit ip any any
vlan access-map HSRP_FILTER 10
 match ip address HSRP_TRAFFIC
 action drop
vlan access-map HSRP_FILTER 20
 match ip address ALL
 action forward
vlan filter FILTER vlan-list 100, 200
```

Routed East-West Inter DC Scenario

ASA 9.5(1)

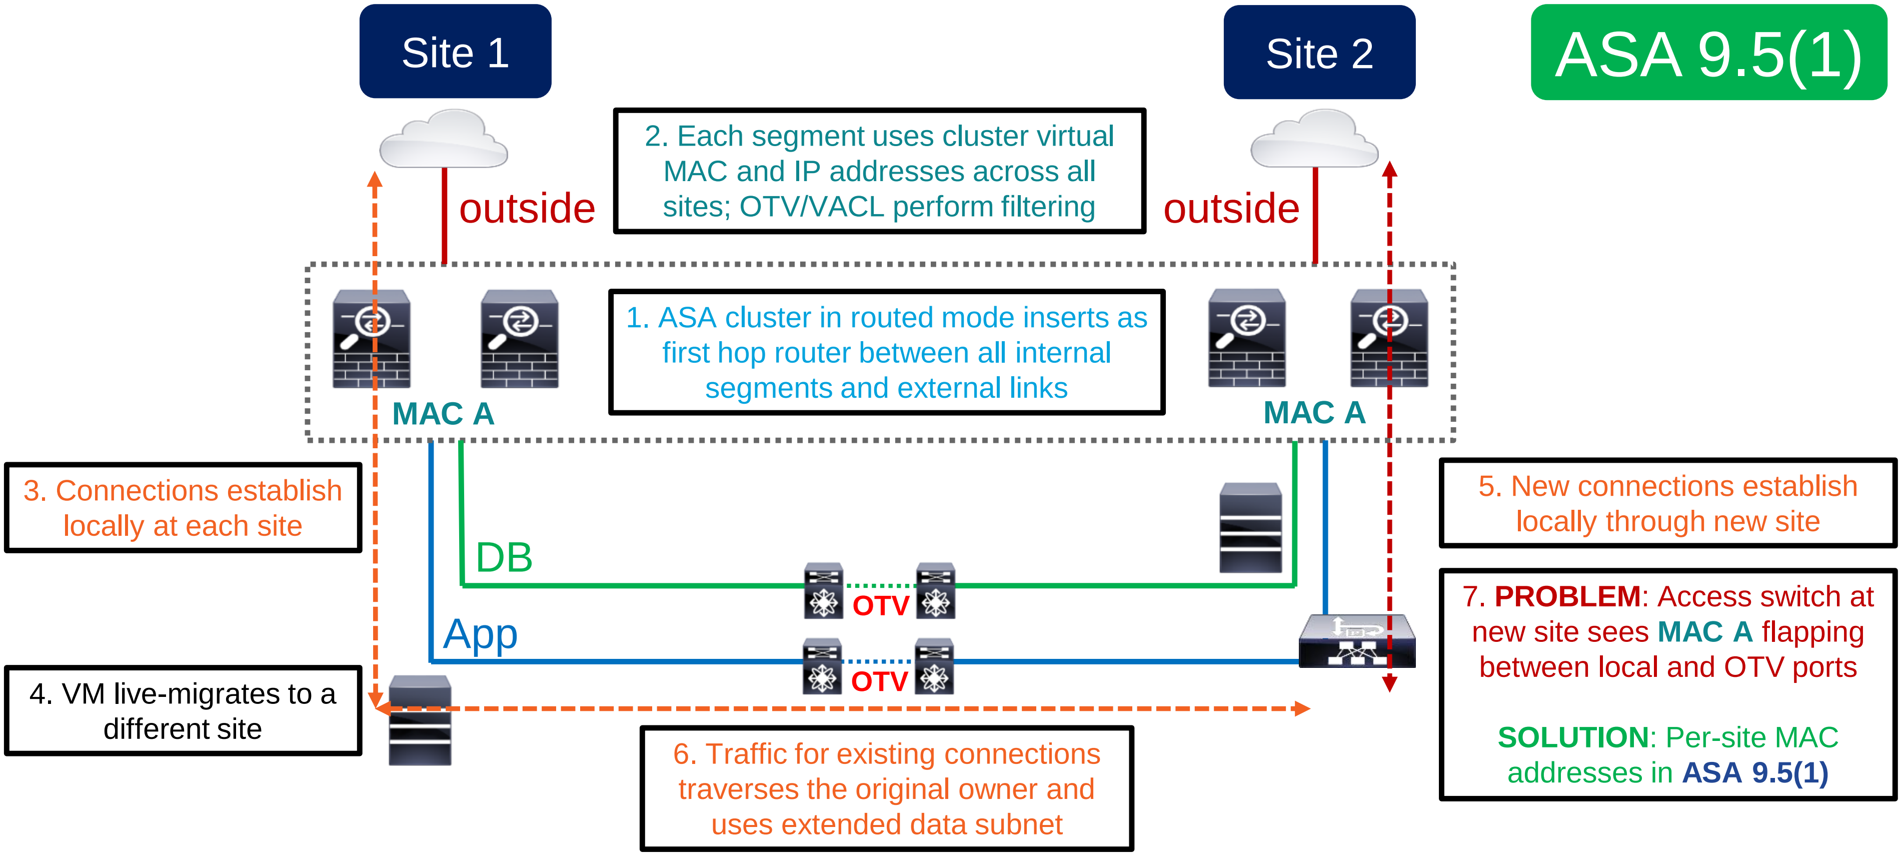

1. ASA cluster in routed mode inserts as first hop router between all internal segments and external links

2. Each segment uses cluster virtual MAC and IP addresses across all sites; OTV/VACL perform filtering

3. Connections establish locally at each site

4. VM live-migrates to a different site

5. New connections establish locally through new site

6. Traffic for existing connections traverses the original owner and uses extended data subnet

7. **PROBLEM**: Access switch at new site sees **MAC A** flapping between local and OTV ports

SOLUTION: Per-site MAC addresses in **ASA 9.5(1)**

Per-Site MAC Addresses

- Routed Spanned Etherchannel cluster uses different MAC addresses in **9.5(1)**
  - Global interface MAC address is used to receive and source frames by default
  - Per-site MAC addresses are optionally used to source frames on extended segments

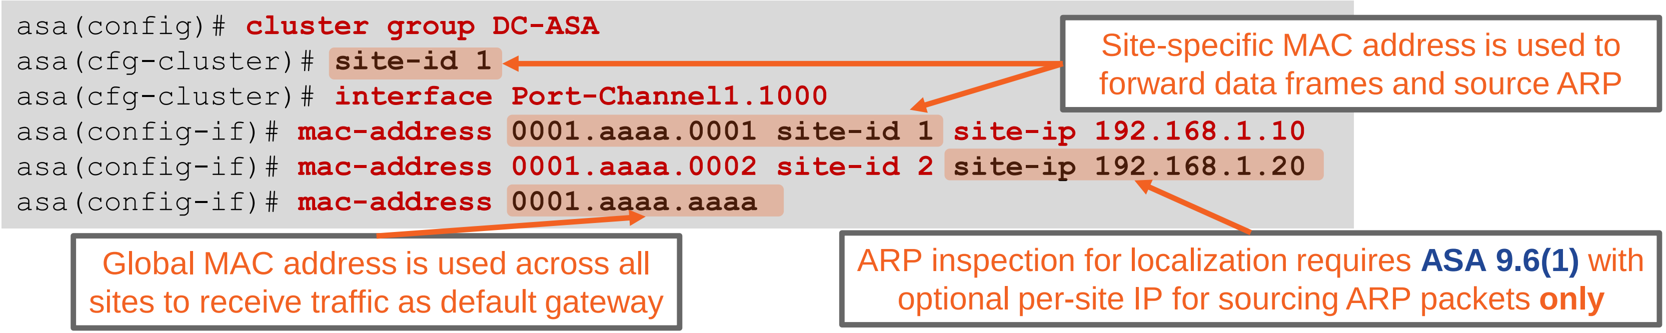

```
asa(config)# cluster group DC-ASA
asa(cfg-cluster)# site-id 1
asa(cfg-cluster)# interface Port-Channel1.1000
asa(config-if)# mac-address 0001.aaaa.0001 site-id 1 site-ip 192.168.1.10
asa(config-if)# mac-address 0001.aaaa.0002 site-id 2 site-ip 192.168.1.20
asa(config-if)# mac-address 0001.aaaa.aaaa
```

Site-specific MAC address is used to forward data frames and source ARP

Global MAC address is used across all sites to receive traffic as default gateway

ARP inspection for localization requires **ASA 9.6(1)** with optional per-site IP for sourcing ARP packets **only**

- Dynamic routing is centralized, so it does not work with split outside networks
- Global MAC address localization is required through OTV or similar mechanisms
  - OTV Unicast Flooding for “silent” hosts is required

Example: E-W Routed Spanned Cluster

- A vPC pair of Nexus switches at each site
  - Split Spanned Etherchannel cluster in routed mode to separate internal segments
  - Separate Etherchannel to local cluster members per vPC pair
  - Static routing between distribution and core is acceptable
- Internal VLANs are fully extended between sites with OTV
  - Each site uses localized ASA cluster as first hop router
  - Traffic to and from global cluster MAC is blocked between sites
  - Nexus F2 line cards allow VACL filtering without ARP Inspection
  - OTV filters must be manually removed on full upstream path failure
  - Only a few "silent" hosts at each site
- Must implement LISP to avoid excessive flow redirection

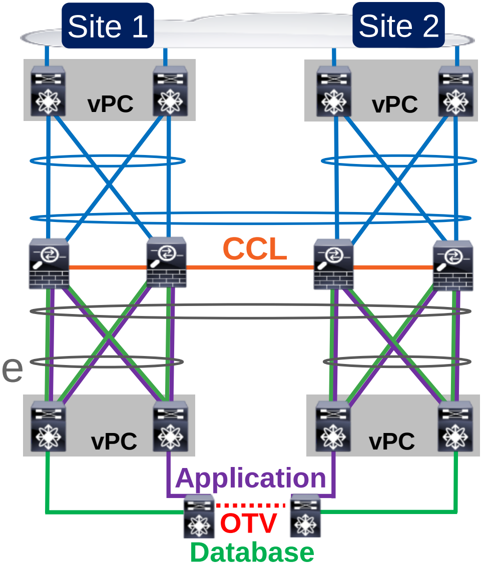

E-W Routed Spanned Cluster Configuration

```
interface Vlan300
 ip address 192.168.3.2/24
hsrp 10
  preempt
  ip 192.168.3.1
ip route 192.168.1.0/24 192.168.3.5
ip route 192.168.2.0/24 192.168.3.5
```

```
cluster-group DC-ASA
 site-id 1
interface Port-Channel10
 port-channel span-cluster
 mac-address 0001.aaaa.aaaa
interface Port-Channel10.100
 vlan 100
 nameif DB
 ip address 192.168.1.1 255.255.255.0
 mac-address 0001.aa01.0001 site-id 1
 mac-address 0001.aa01.0002 site-id 2
interface Port-Channel10.200
 vlan 200
 nameif App
 ip address 192.168.2.1 255.255.255.0
 mac-address 0001.aa02.0001 site-id 1
 mac-address 0001.aa02.0002 site-id 2
interface Port-Channel10.300
 vlan 300
 nameif outside
 ip address 192.168.3.5 255.255.255.0
route outside 0.0.0.0 0.0.0.0
           192.168.3.1
```

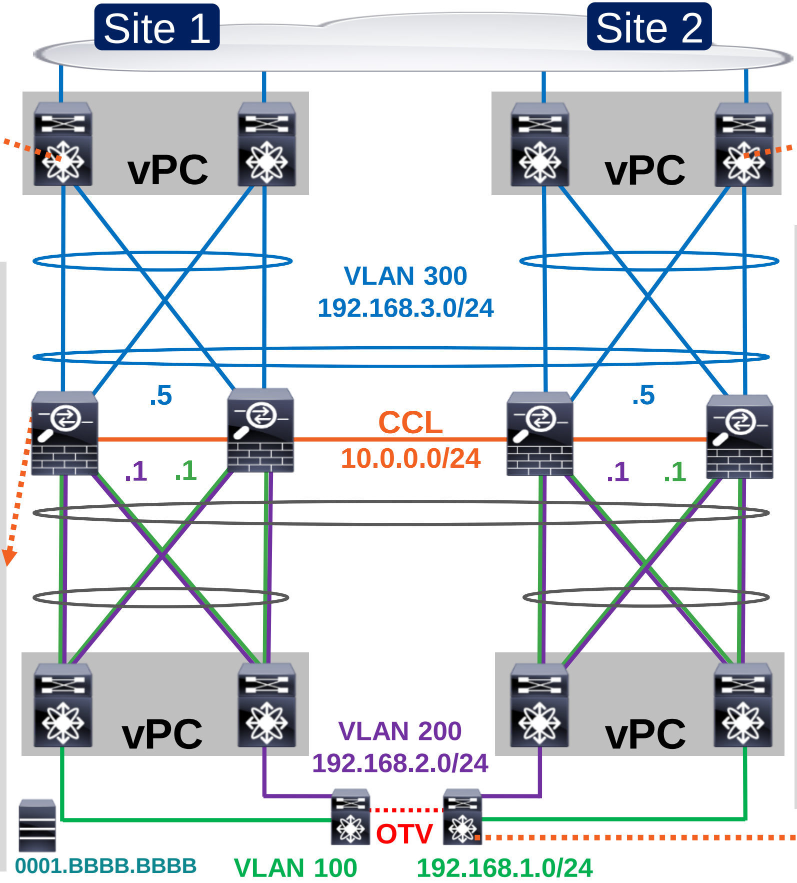

```
interface Vlan300
 ip address 192.168.3.3/24
hsrp 10
  ip 192.168.1.1
ip route 192.168.1.0/24 192.168.3.5
ip route 192.168.2.0/24 192.168.3.5
```

```
mac-list GMAC_FILTER seq 10 deny
           0001.aaaa.aaaa ffff.ffff.ffff
mac-list GMAC_FILTER seq 20 permit
           0000.0000.0000 0000.0000.0000
otv-isis default
 vpn Overlay1
  redistribute filter route-map HSRP_FILTER
!
mac access-list GMAC_TRAFFIC
 10 permit 0001.aaaa.aaaa 0000.0000.0000 any
 20 permit any 0001.aaaa.aaaa 0000.0000.0000
mac access-list ALL
 10 permit any any
vlan access-map FILTER 10
 match mac address GMAC_TRAFFIC
 action drop
vlan access-map FILTER 20
 match mac address ALL
 action forward
vlan filter FILTER vlan-list 100, 200
!
otv flood mac 0001.bbbb.bbbb vlan 100
```

Closing Remarks

Clustering Best Practices

- Only use compatible Catalyst and Nexus switches
- Leverage LACP Etherchannel for CCL and dual-connect to VSS/vPC
  - Match the forwarding capacity of each member
  - Raise CCL MTU to 100 bytes above all data interfaces
- Speed up switching and routing convergence
  - Enable Spanning Tree Portfast on CCL and data interfaces
  - Use NSF/GR or lower dead interval and SPF throttle timers on cluster and peers
- Reduce asymmetry to increase scale
  - Keep TCP Sequence Number Randomization enabled for SYN Cookies
  - Minimize centralized features and NAT/PAT
  - Use Spanned Etherchannel mode with symmetric hashing for best load distribution

Complete Your Online Session Evaluation

- Give us your feedback to be entered into a Daily Survey Drawing. A daily winner will receive a $750 Amazon gift card.
- Complete your session surveys through the Cisco Live mobile app or from the Session Catalog on CiscoLive.com/us.

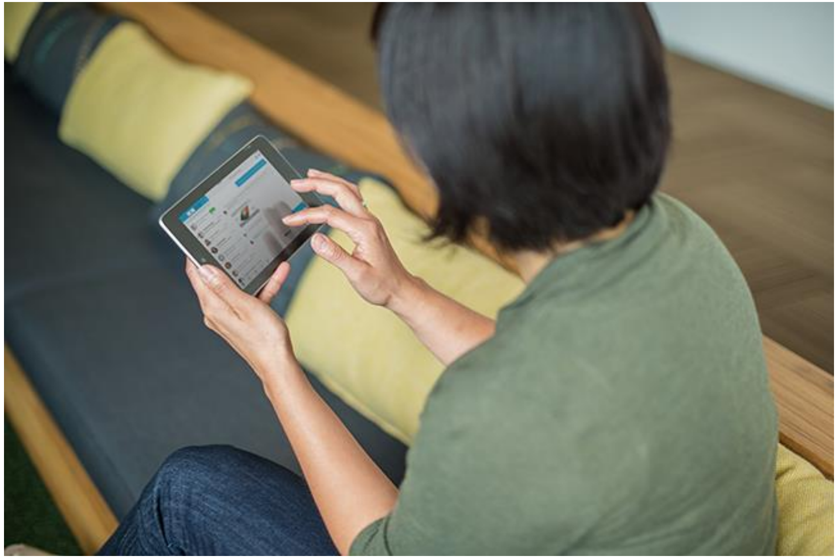

Don’t forget: Cisco Live sessions will be available for viewing on-demand after the event at CiscoLive.com/Online

Continue Your Education

- Demos in the Cisco campus
- Walk-in Self-Paced Labs
- Lunch & Learn
- Meet the Engineer 1:1 meetings
- Related sessions

Security Joins the Customer Connection Program

Customer User Group Program

- **Who can join**: Cisco customers, service providers, solution partners and training partners
- **Private online community** to connect with peers & Cisco's Security product teams
- Monthly **technical & roadmap briefings** via WebEx
- Opportunities to **influence product direction**
- Local **in-person meet ups** starting Fall 2016
- **New member thank you gift* & badge ribbon** when you join in the **Cisco Security booth**
- **Other CCP tracks**: Collaboration & Enterprise Networks

19,000+ Members Strong

Join in World of Solutions

Security zone → Customer Connection stand

- Learn about CCP and Join
- New member thank-you gift*
- Customer Connection Member badge ribbon

Join Online

www.cisco.com/go/ccp

Come to Security zone to get your new member gift* and ribbon

* While supplies last

Please join us for the Service Provider Innovation Talk featuring:

Yvette Kanouff | Senior Vice President and General Manager, SP Business

Joe Cozzolino | Senior Vice President, Cisco Services

Thursday, July 14th, 2016

11:30 am - 12:30pm, In the Oceanside A room

What to expect from this innovation talk

- Insights on market trends and forecasts
- Preview of key technologies and capabilities
- Innovative demonstrations of the latest and greatest products
- Better understanding of how Cisco can help you succeed

Register to attend the session live now or watch the broadcast on cisco.com

Thank you

Cisco live!

July 10-14, 2016 • Las Vegas, NV