
cisco
You make possible

NGFW Clustering Deep Dive

Andrew Ossipov, Distinguished Engineer

BRKSEC-3032

cisco Live!
Barcelona | January 27-31, 2020

Cisco Webex Teams

Questions?

Use Cisco Webex Teams to chat with the speaker after the session

How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

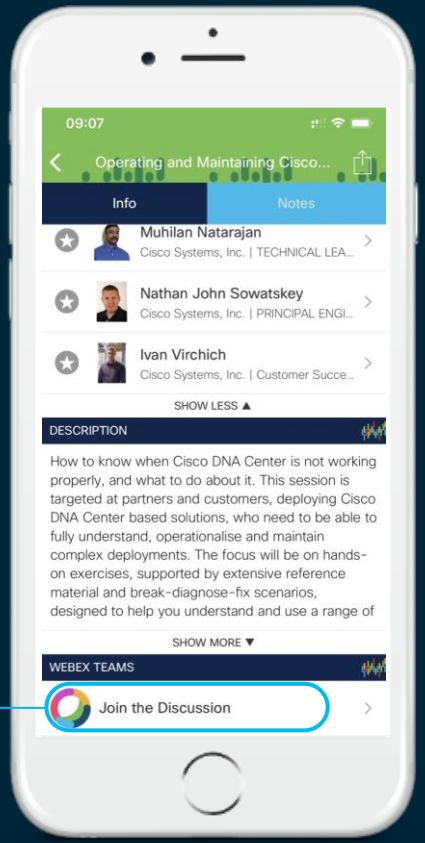

Your Speaker

Andrew Ossipov

aeo@cisco.com

Distinguished Engineer

NGFW, Solution Architecture, Hybrid Cloud DC

IETF: OpSec and TLS Working Groups

Agenda

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

Clustering Overview

High Availability on ASA and FTD

- A pair of identical ASA or FTD devices can be configured in Failover/HA
  - Managed as a single entity
  - Data interface connections must be mirrored between the units with L2 adjacency
  - Virtual IP and MAC addresses on data interfaces move with Active unit
  - Stateful connection table is replicated to Standby in real time
- Failover/HA deliver high availability rather than scalability
  - Limited to two physical appliances/modules or virtual instances
  - Active/Standby for asymmetry avoidance in ASA or FTD
  - Active/Active with multiple contexts in ASA is impractical for scaling

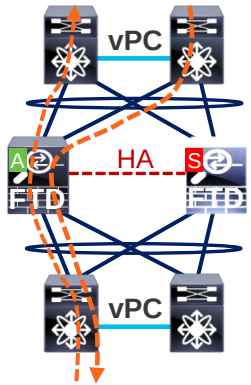

ASA and FTD Clustering

- Up to 16 appliances or modules combine in one traffic processing system
- Preserve failover benefits by configuring and operating as a single entity
  - Virtual IP and MAC addresses for first-hop redundancy
  - Connection states are preserved after a single member failure
- Implement true scalability in addition to high availability
  - Fully distributed data plane for new and existing connections
  - Elastic scaling of throughput and maximum concurrent connections
  - Stateless external load-balancing through standard Etherchannel
  - Out-of-band Cluster Control Link for asymmetry normalization
  - No member-to-member communication on data interfaces

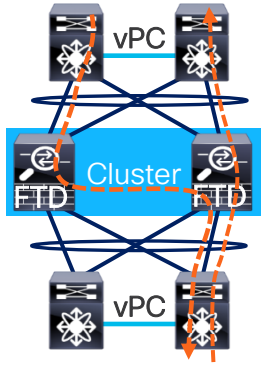

System Requirements

- ASA scales up to 16 identical appliances or modules
  - Up to 16 Firepower 4100 or 9300 modules with matching Export Compliance
  - Up to 16 ASA5585-X with Cluster and same 3DES and 10GE I/O licenses
  - Up to 2 ASA5500-X with Security Plus and matching 3DES licenses
- FTD scales up to 6 identical appliances or modules as documented
  - Up to 16 Firepower 4100 appliances or 9300 modules is configurable
  - Multi-instance capability in FTD 6.6 will no longer require identical hardware
  - Some advanced cluster settings must use FlexConfig
- Any standard-based switch is supported, some are explicitly validated

Unsupported Features

- Remote Access VPN: TLS VPN, Clientless SSL VPN, and IPSec
- S2S VPN on FTD only until 6.2.3.3
- DHCP client, DHCP server, DHCP Proxy
- Advanced Application Inspection and Redirection
  - CTIQBE, WAAS, MGCP, MMP, RTSP, Skinny/SCCP, H.323
  - Dead Connection Detection (DCD) until ASA 9.13, Botnet Traffic Filter, and WCCP
- Interfaces: Integrated Routing/Bridging (IRB), Virtual Tunnel Interface (VTI)
- Intermediate System-to-Intermediate System (IS-IS)
- Firepower Multi-Instance Capability until FTD 6.6

Scalability

- Throughput scales at 70-80% of the aggregated capacity on average
  - ASA: 16 Firepower 4145 at 50Gbps → 640Gbps of Multiprotocol Throughput
  - FTD: 6 Firepower 9300 SM-44 at 50Gbps → 240Gbps of NGFW AVC Throughput
- Replicated concurrent conn(ection)s scale at 60% of aggregated capacity
  - FTD: 6 Firepower 4150 at 35M → 126M concurrent conns
  - Firepower 9300 supports 120M (ASA) or 60M (FTD) conns per clustered chassis
- Conn rate with full replication scales at 50% of the aggregated capacity
  - ASA: 16 ASA5585-X SSP-60 at 350K CPS → 2.8M CPS
  - Short-lived connections may scale at 100% with delayed replication

```
asa(config)# cluster replication delay 10 match tcp any any eq www
```

Delay by 10 seconds

Match All HTTP connections

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 until ASA 9.9(1) with optional distribution on Firepower 9300
  - Multicast in rare scenarios
- Any connections with 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 a departure of the existing one
- Master unit handles all management and centralized functions
  - Configuration is blocked on all other members
  - Virtual IP address ownership for to-the-cluster connections
- Master and slaves process all regular transit connections equally
  - Management and centralized connections must reestablish upon Master failure
  - Disable or reload Master to transition the role

State Transition

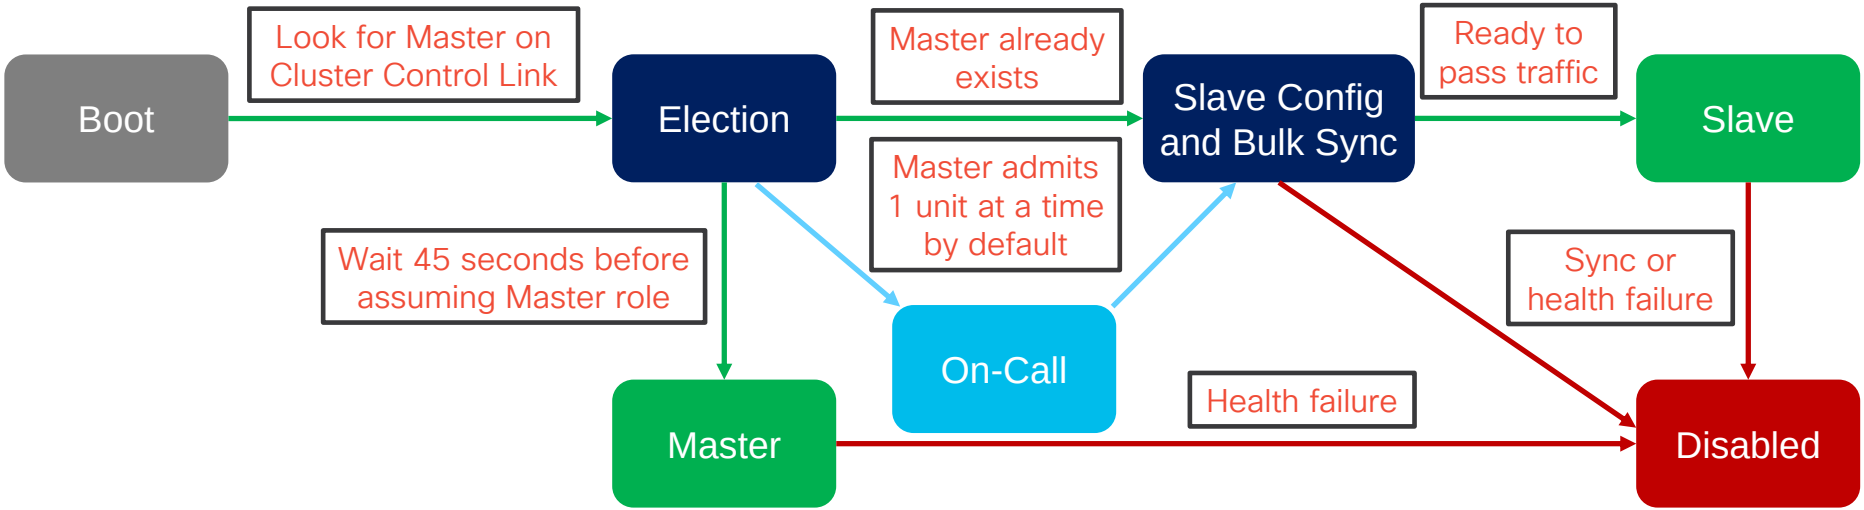


```
ASA/master# show cluster history
==========================================================================
From State            To State              Reason
==========================================================================
15:36:33 UTC Dec 3 2018
DISABLED              DISABLED              Disabled at startup
15:37:10 UTC Dec 3 2018
DISABLED              ELECTION              Enabled from CLI
15:37:55 UTC Dec 3 2018
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   : 9.10(1)
        Serial No.: JAF1434AERL
        CCL IP    : 1.1.1.1
        CCL MAC   : 5475.d029.8856
        Last join : 15:37:55 UTC Dec 3 2018
        Last leave: N/A
```

Member Admission Optimization

- ASA 9.10(1) and FTD 6.3 allow parallel cluster join on Firepower 9300
  - Each chassis optionally bundles data interfaces only when all modules are ready

```
asa(cfg-cluster)# unit parallel-join 3 max-bundle-delay 5
```

How many modules must replicate configuration and state before enabling chassis data plane

Maximum wait time in minutes

Flow Owner

- All packets for a single stateful connection 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 NGFW/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 ASA 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 a 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 among members
  - All members share the same hash table and algorithm for consistent lookups
  - **SYN Cookies** reduce lookups for TCP flows with **Sequence Number Randomization**
- Other units ask **Flow Director** to identify **Owner** or restore flow from backup
  - 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** is same member or in 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 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 ask **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**

```
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

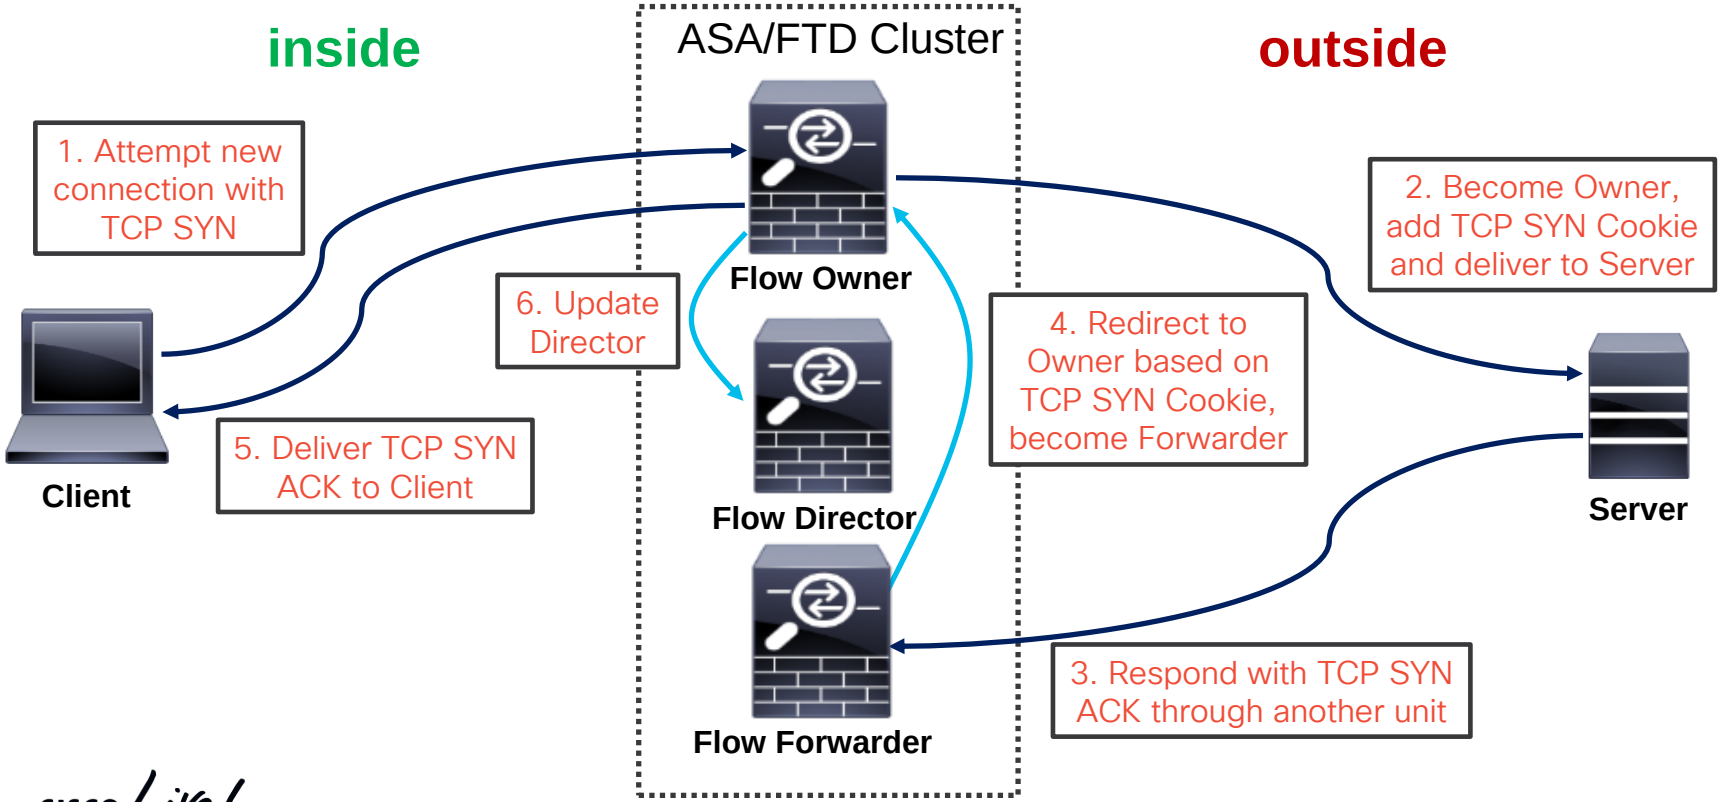

New UDP-Like Connection

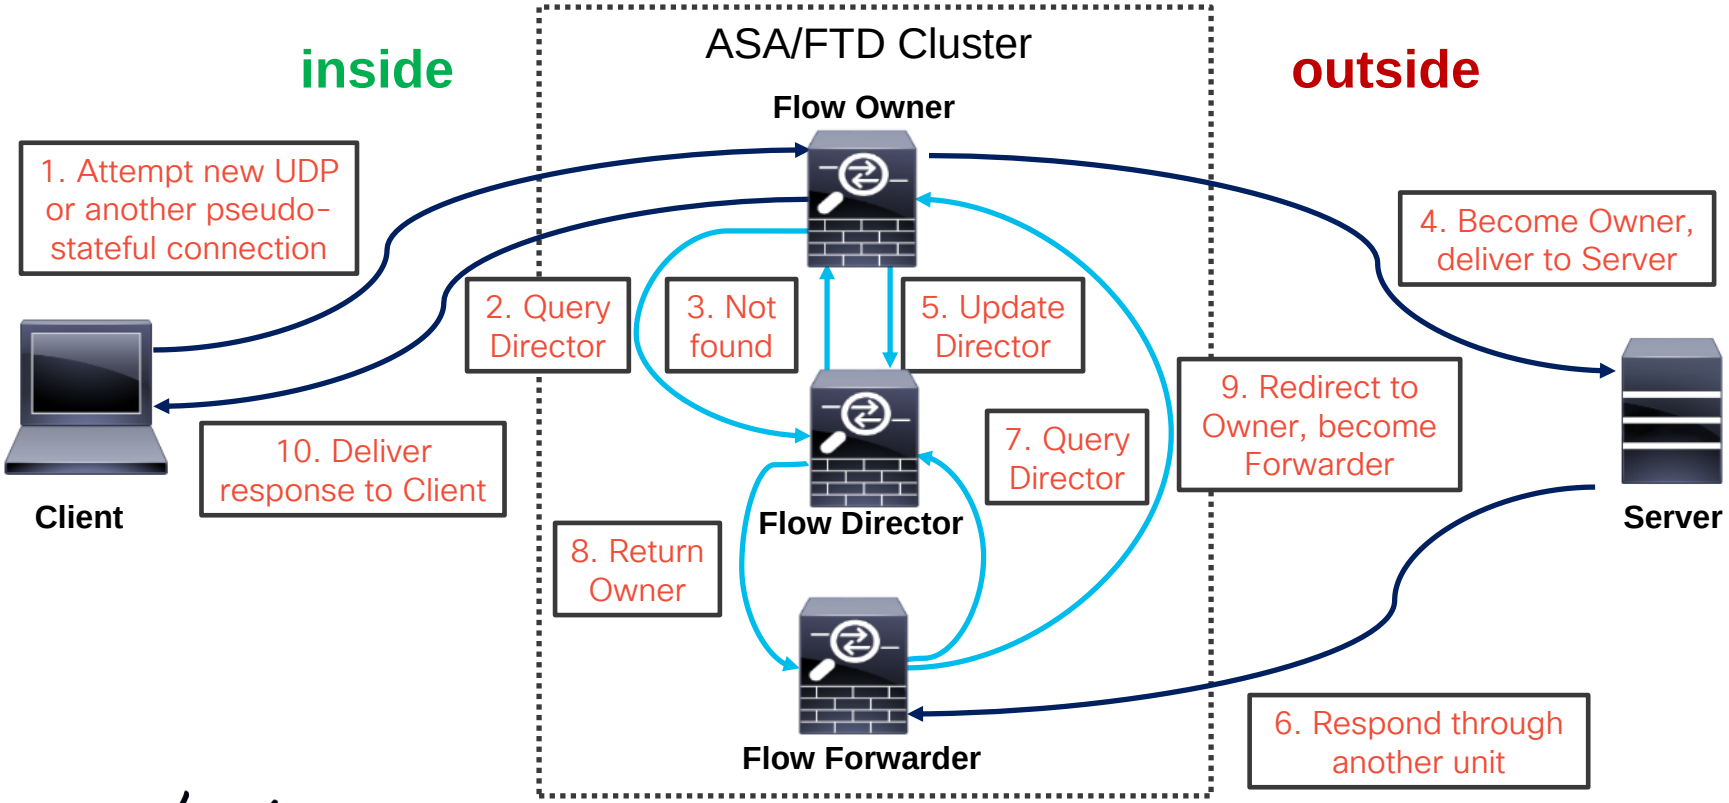

New Centralized Connection

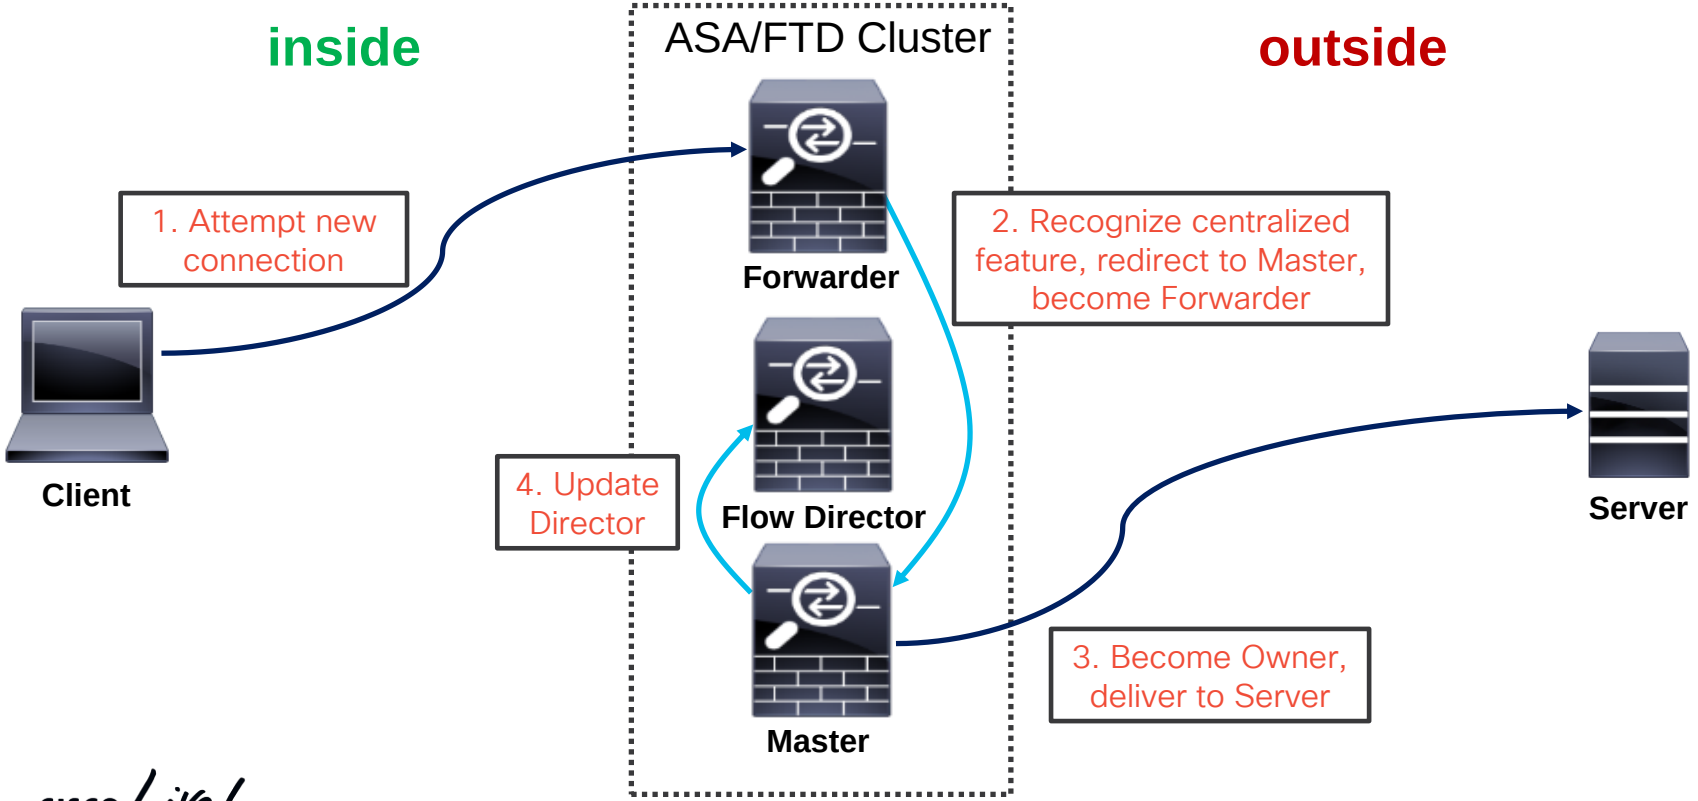

Owner Failure

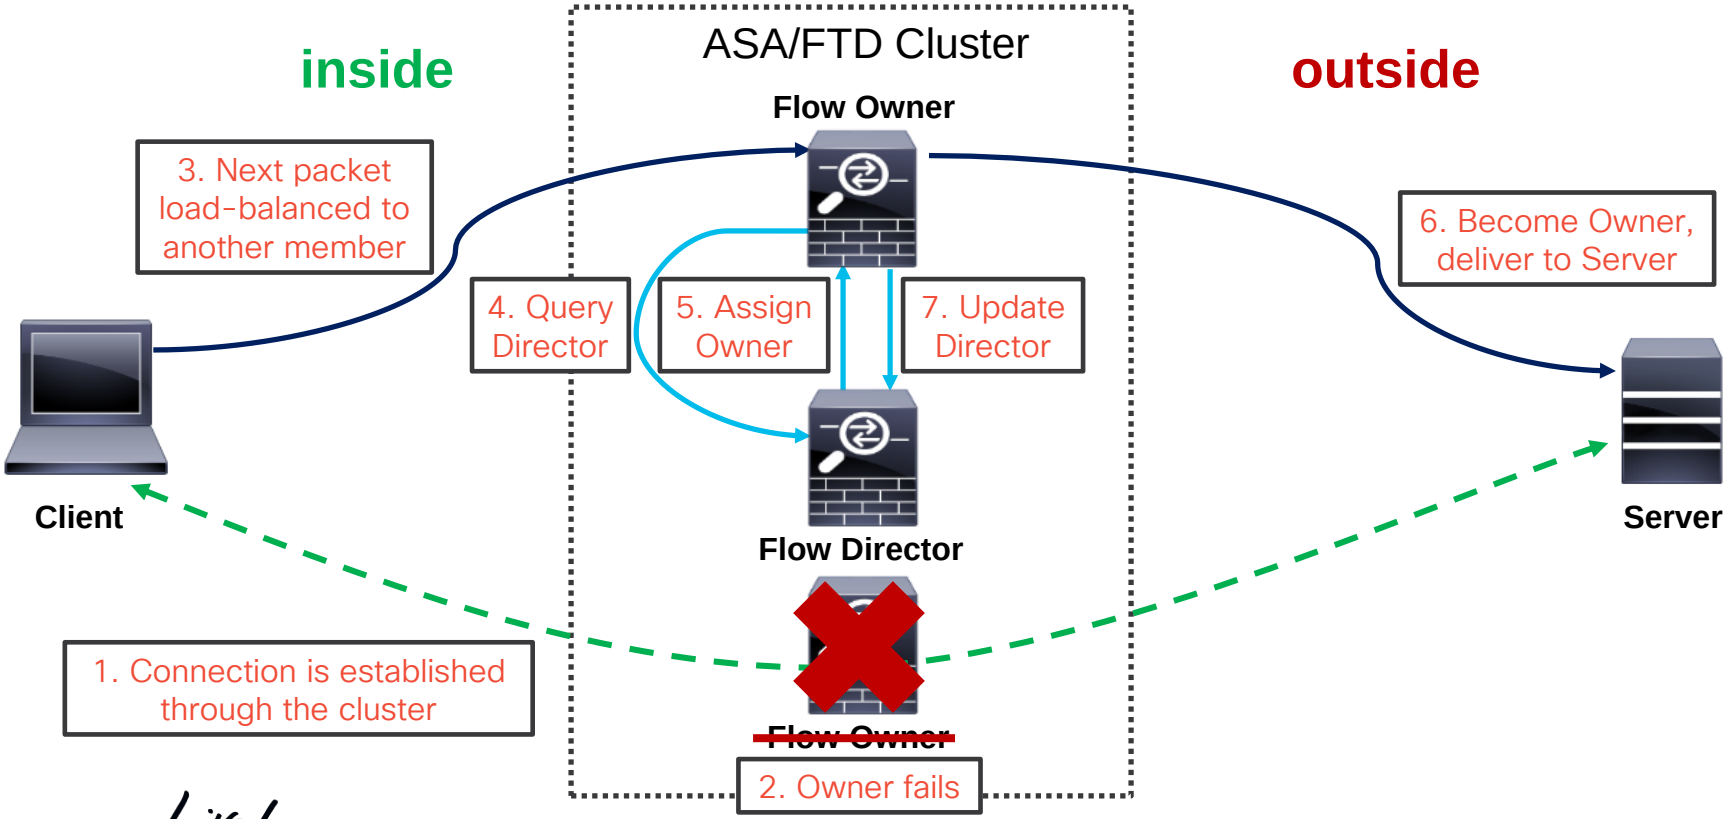

Basic Application Inspection

- 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 on ASA, FTP, GTP
- Semi Distributed
  - Control connections are processed independently by all units
  - Data connections are redirected to the associated control connections' Owners
  - Examples: SIP, SCTP, M3UA

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
- **Per-Session Xlate** feature allows immediate reuse of the mapped port
  - Enabled by default for all TCP and DNS connections

```
ftd# 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
- Master creates one-to-one Dynamic NAT xlates and replicates 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
- 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

Distributed PAT in Clustering

- Today PAT pool is uniformly distributed to all cluster members at IP level

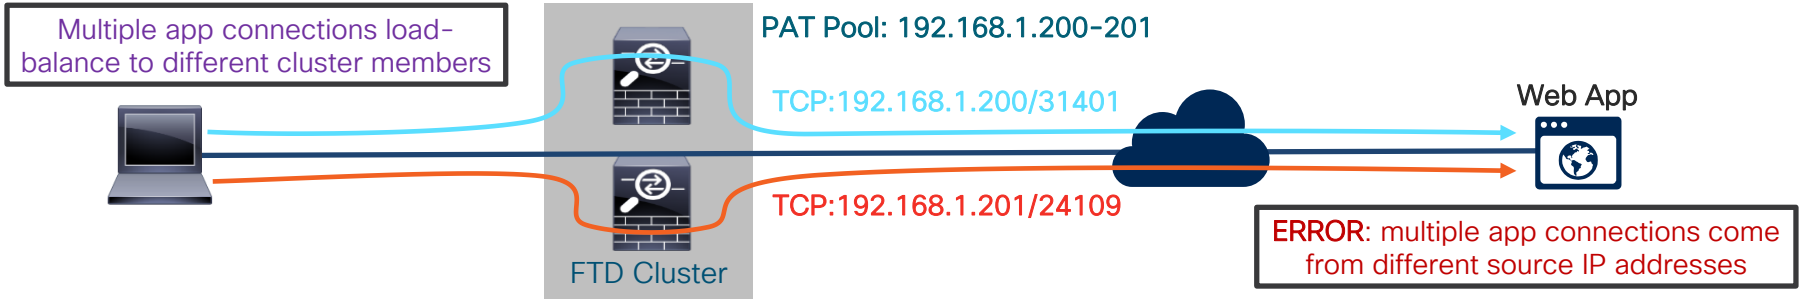

- FTD 6.7 and ASA 9.15 will distribute each PAT pool IP at port block level

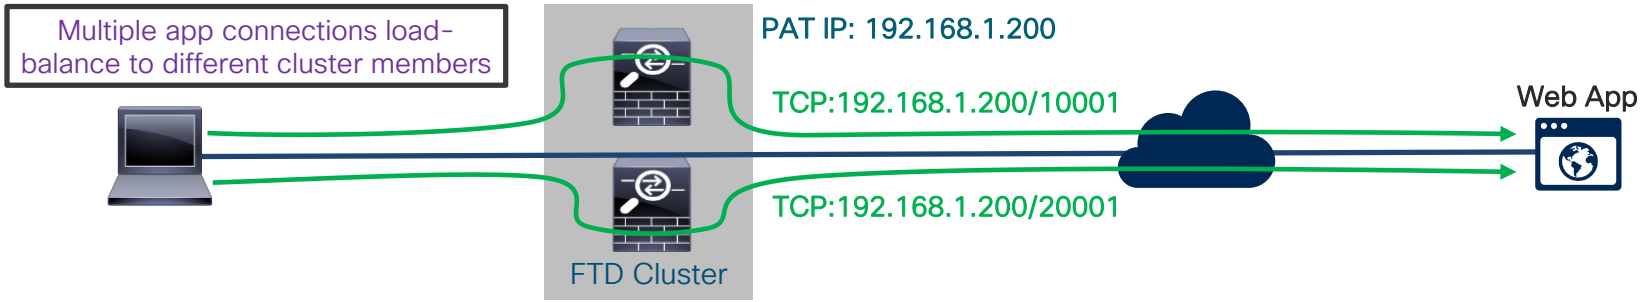

Site-to-Site (S2S) IKEv2 VPN in Distributed Mode

- Supported on Firepower 9300 with ASA 9.9(1) and Carrier license only

```
asa(cfg-cluster)# vpn-mode distributed backup flat
```

- Tunnel establishment (IKEv2) is done through per-session VPN Director
- VPN Session Owner handles IPsec and clear text traffic for a single tunnel
  - Backup Owner assures uninterrupted forwarding on failure
  - Optional Remote Chassis Backup protects against full chassis failure

```
asa(cfg-cluster)# vpn-mode distributed backup remote-chassis
```

- Scalability is constrained by multiple factors
  - Concurrent S2S VPN tunnels scale at ~45% of aggregated capacity
  - Throughput impact from cleartext traffic redirection to VPN session owner
  - Runtime manual tunnel redistribution with cluster redistribute vpn-sessiondb

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 of one-way latency
- CCL loss forces the member out of the cluster
  - No direct back-to-back connections

CCL Best Practices

- Use a per-unit LACP Etherchannel for link redundancy and aggregation
  - Bandwidth should match maximum forwarding capacity of each member
  - 40Gbps of data traffic on Firepower 4140 AVC+IPS → 4x10GE CCL
  - Dual-connect to different physical switches in vPC/VSS
- Set MTU 100 bytes above largest data interface MTU
  - Avoids fragmentation of redirected traffic due to extra trailer
  - Minimum supported value is 1400 bytes
- Ensure that CCL switches do not verify L4 checksums
- Enable Spanning Tree Portfast and align MTU on the switch side

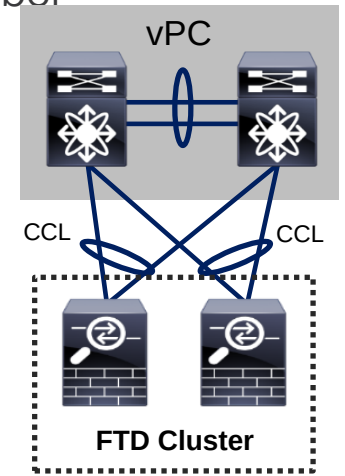

Data Interface Modes

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

```
asa5585(config)# interface Port-Channel1
asa5585(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 unique IP on each data interface in **Individual "L3"** mode
  - Available only on ASA5500-X and ASA5585-X appliances running ASA image
  - Use **Nexus ITD** or **PBR** or dynamic routing protocols to load-balance traffic
  - Virtual IPs are owned by Master, interface IPs are assigned from configured pools

```
asa5585(config)# ip local pool INSIDE 192.168.1.2-192.168.1.17
asa5585(config-if)# interface Port-Channel1
asa5585(config-if)# ip address 192.168.1.1 255.255.255.0 cluster-pool INSIDE
```

Spanned Etherchannel Interface Mode

- Transparent and routed mode on ASA/FTD; NGIPS interfaces in FTD
- 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

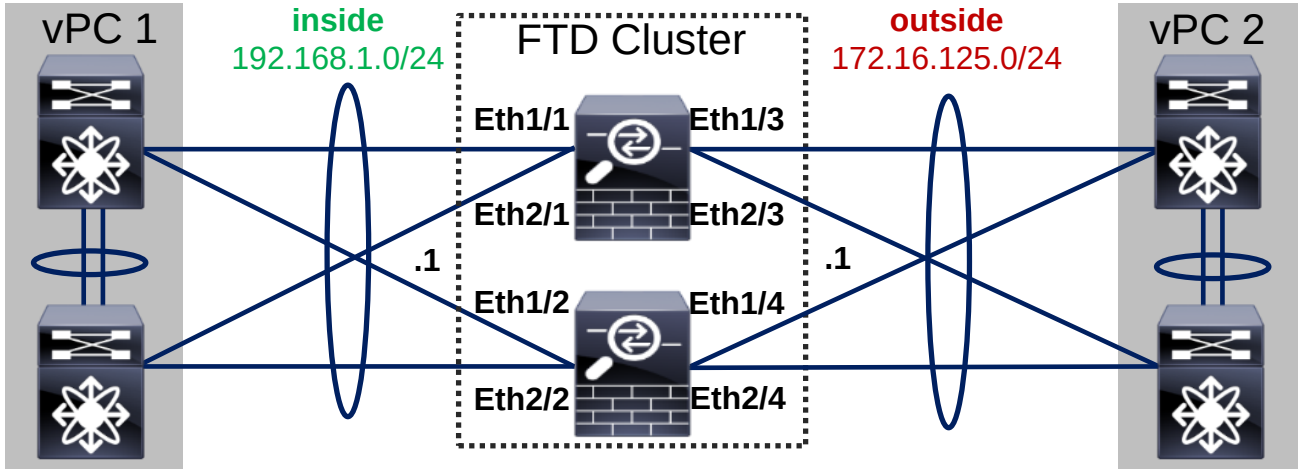

Clustering LACP (cLACP)

- Spanned Etherchannel is preferred for data interfaces on ASA appliances
  - Up to 32 (16 per unit) active total links with global static port priorities

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

  - Disable LACP Graceful Convergence and Adaptive Hash on adjacent NX-OS
- Supervisor bundles data and CCL interfaces on Firepower 4100 and 9300
  - Spanned Etherchannel only with up to 32 active total (up to 16 per chassis) links
  - Disable only Adaptive Hash on adjacent NX-OS
- Always configure virtual MAC for each data Etherchannel to avoid instability
- cLACP assumes a 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 ASA 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 admission
- Per-unit Etherchannel support up to 16 members

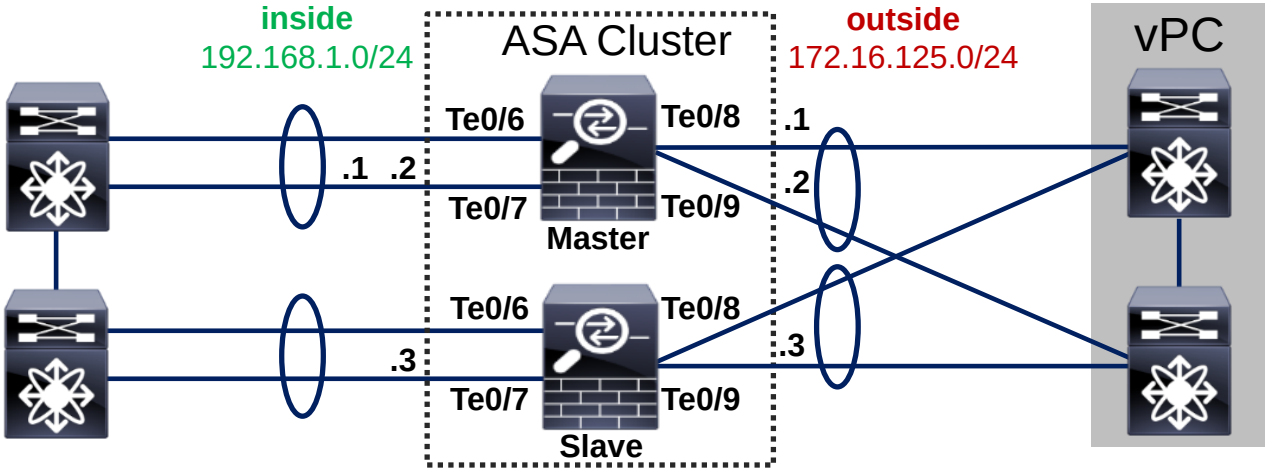

Traffic Load Balancing in Individual Mode

- 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, BGP-4 (IPv4 and IPv6), PIM
  - 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
  - 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**) avoid blackholing

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

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

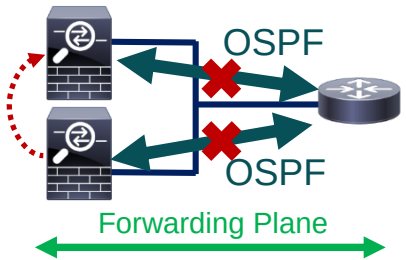

2. Router re-establishes adjacency with Master while retaining the stale routes; these routes are refreshed when the adjacency reestablishes.

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

NSF and GR Configuration on ASA

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

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

(Optional) Disable NSF if any adjacent device is

(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.

Faster Dynamic Routing Convergence on ASA

- Reduce protocol timers on all adjacent segments to improve 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

Check conn and packet distribution

```
ftd# 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
ftd# 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
```

Avoid too much forwarding

Cluster Management

- Dedicated management interface is required on FTD and preferred on ASA
  - SNMP typically requires per-unit IP, syslog/NSEL can share IP on a data interface
  - management-only allows MAC/IP pools in Spanned Etherchannel mode on ASA
- A regular data interface can be used for managing an ASA cluster in-band
  - Connecting to cluster data interface IP always reaches the master
- Use cluster exec for non-configuration commands on some/all members

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

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

Health Monitoring

- A unit shuts down all data interfaces and disables clustering on CCL failure
- Each member generates keepalives on CCL every 1 second by default
  - Master removes 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
  - Rejoin attempted 3 times (after 5, 10, 20 minutes), then the unit disables clustering

```
a/master# cluster group sjfw
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 2 1
a/master(cfg-cluster)# health-check data-interface auto-rejoin 10 2 1
a/master(cfg-cluster)# health-check system auto-rejoin 5 2 1
```

Keepalive is always 1/3 of the configured holdtime

Exempt non-critical interfaces from monitoring

Configurable re-join attempts, interval, and interval multiplier

Re-join on internal system failures in ASA 9.9(2) and FTD 6.2.3

Configuring Clustering

Preparation Checklist for ASA Appliances

- 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 entire cluster

Setting Interface Mode on ASA Appliances

- 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 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

Management Access to ASA Appliances

- 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 Wizard

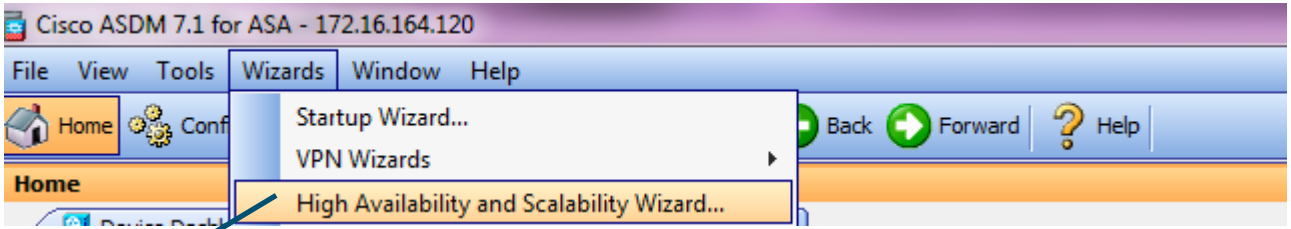

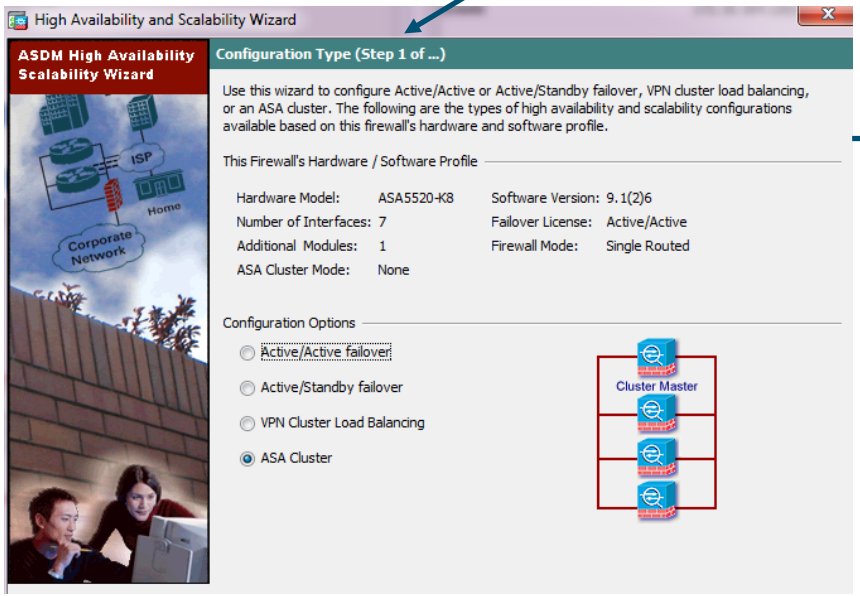

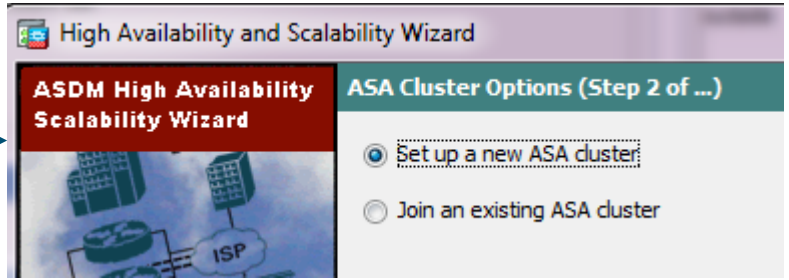

Fully configure Master in 4 easy steps, then have ASDM add Slaves one by one over basic HTTPS management connection.

... or use good old CLI ;-)

ASA CLI: 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
```

ASA CLI: Cluster Group

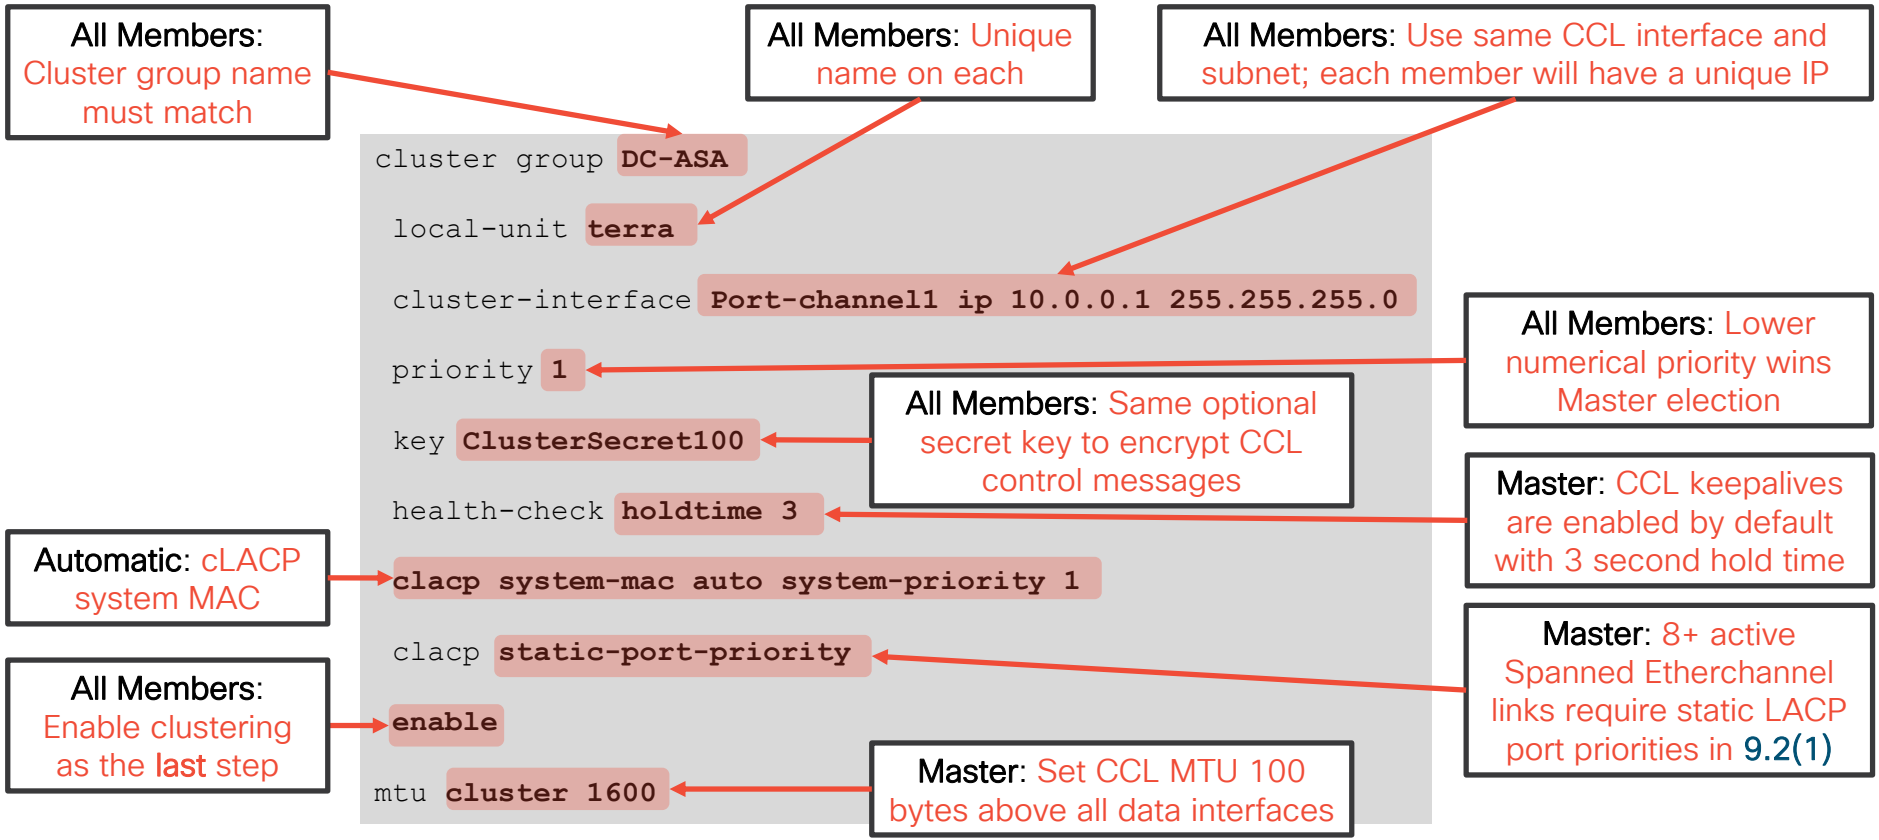

ASA CLI: Data Interfaces on Master

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

ASA CLI: 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
```

ASA: 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

Clustering on Firepower 4100 and 9300

- Only Spanned Etherchannel interface mode is supported
- Supervisor pushes cluster configuration during logical device deployment
  - Site ID for inter-site clustering is optional
  - Firewall context mode and TLS/SSL ciphers are replicated in ASA
- Remote flow backup for N+1 chassis fault tolerance on Firepower 9300
- Module- and chassis-level overflow warning 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.
```

Preparation Checklist for Firepower Appliances

- Set up and cable identical Firepower chassis and modules for the cluster
- Ensure over-the-network Supervisor management access
- Bring up Firepower Management Center for FTD
- Generate device token and enable Smart Licensing on chassis and/or FMC
- Delete all pre-existing logical devices from the chassis
- Download application images for FTD or ASA
- Designate a dedicated application management interface (one per chassis)
- Designate one or more physical interfaces per chassis for CCL
- Assign an isolated subnet for CCL on a separate switch or VDC

Sample FTD Cluster Deployment Topology

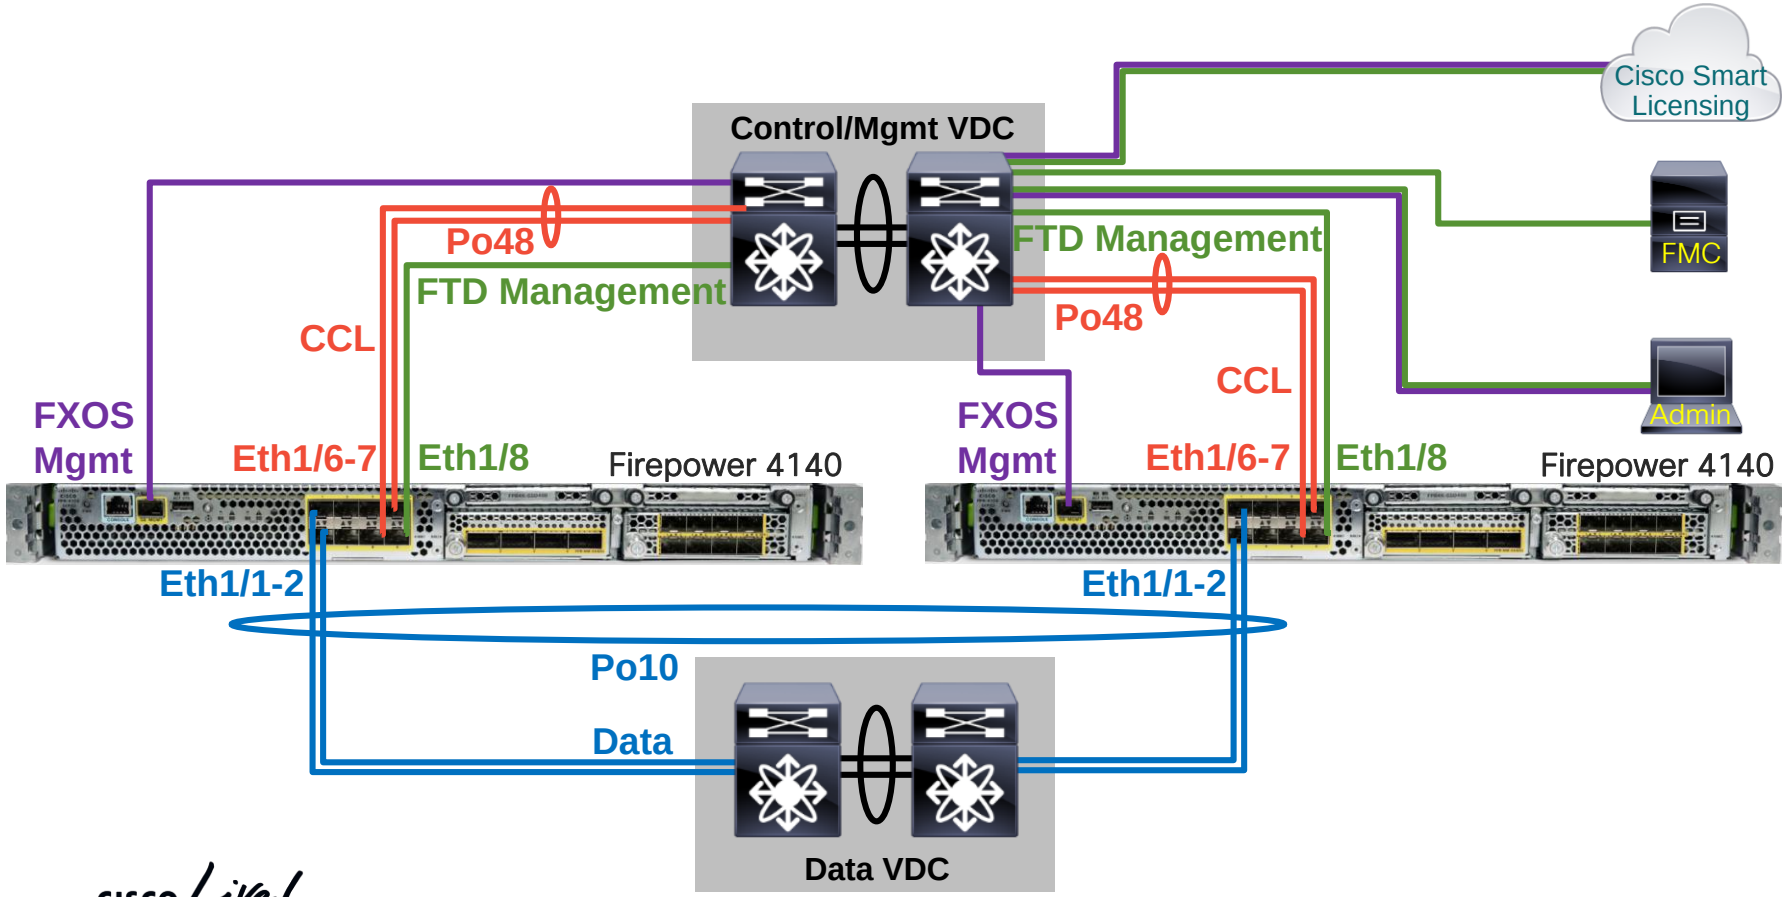

Chassis Manager: Interface Configuration

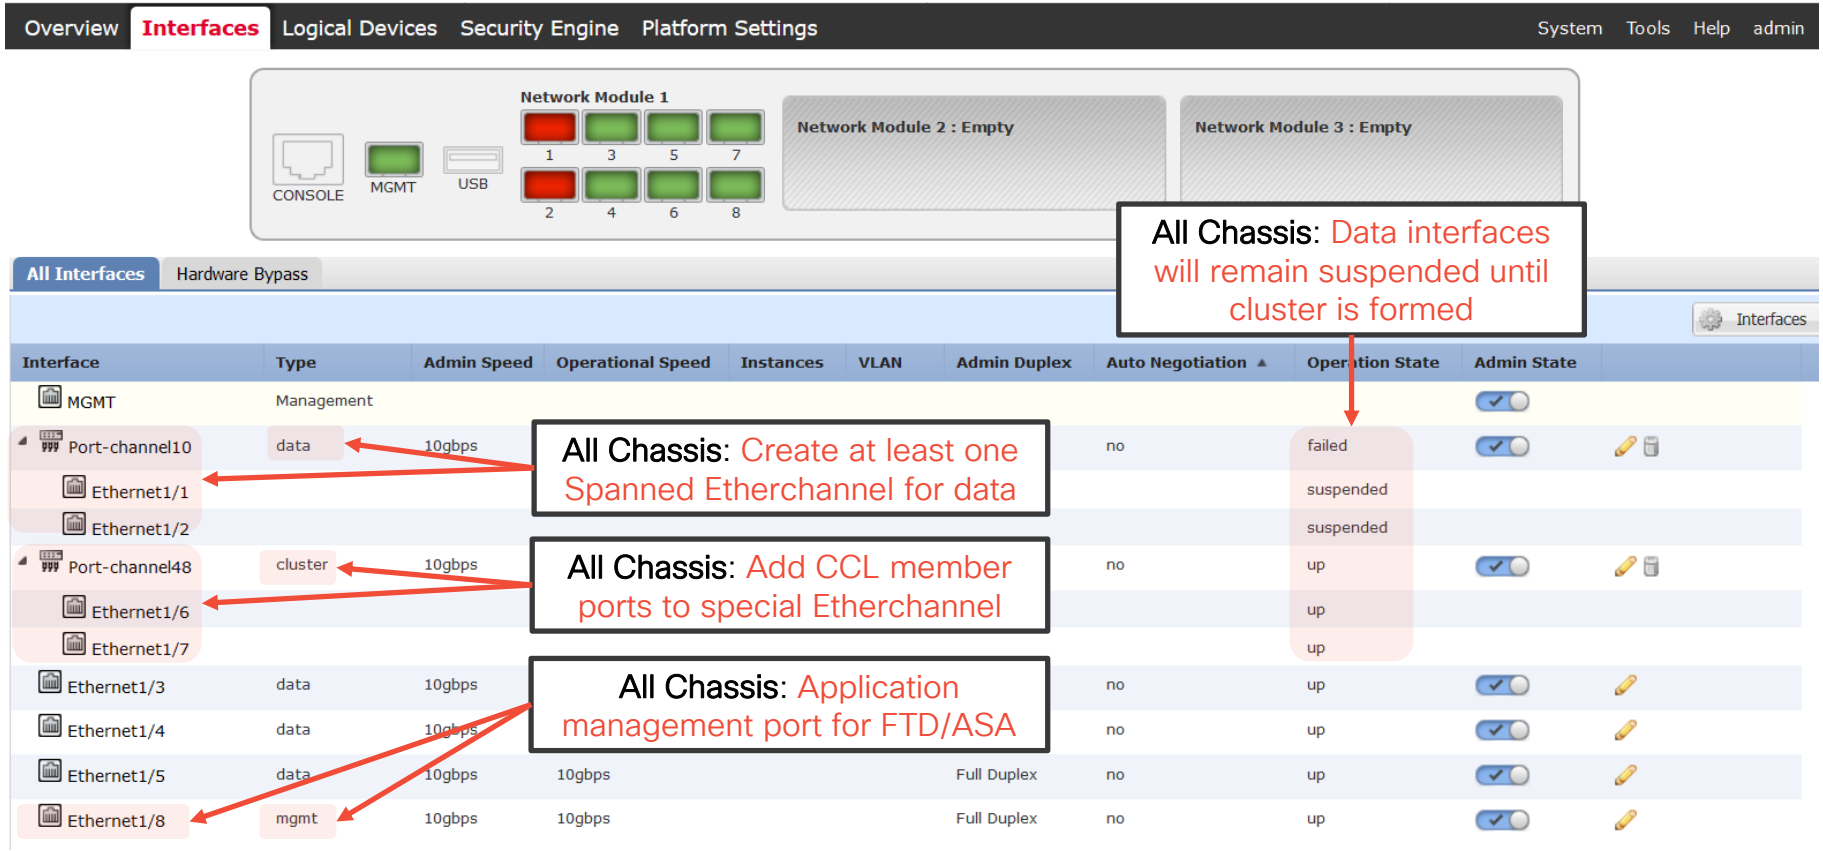

Chassis Manager: Add Logical Device

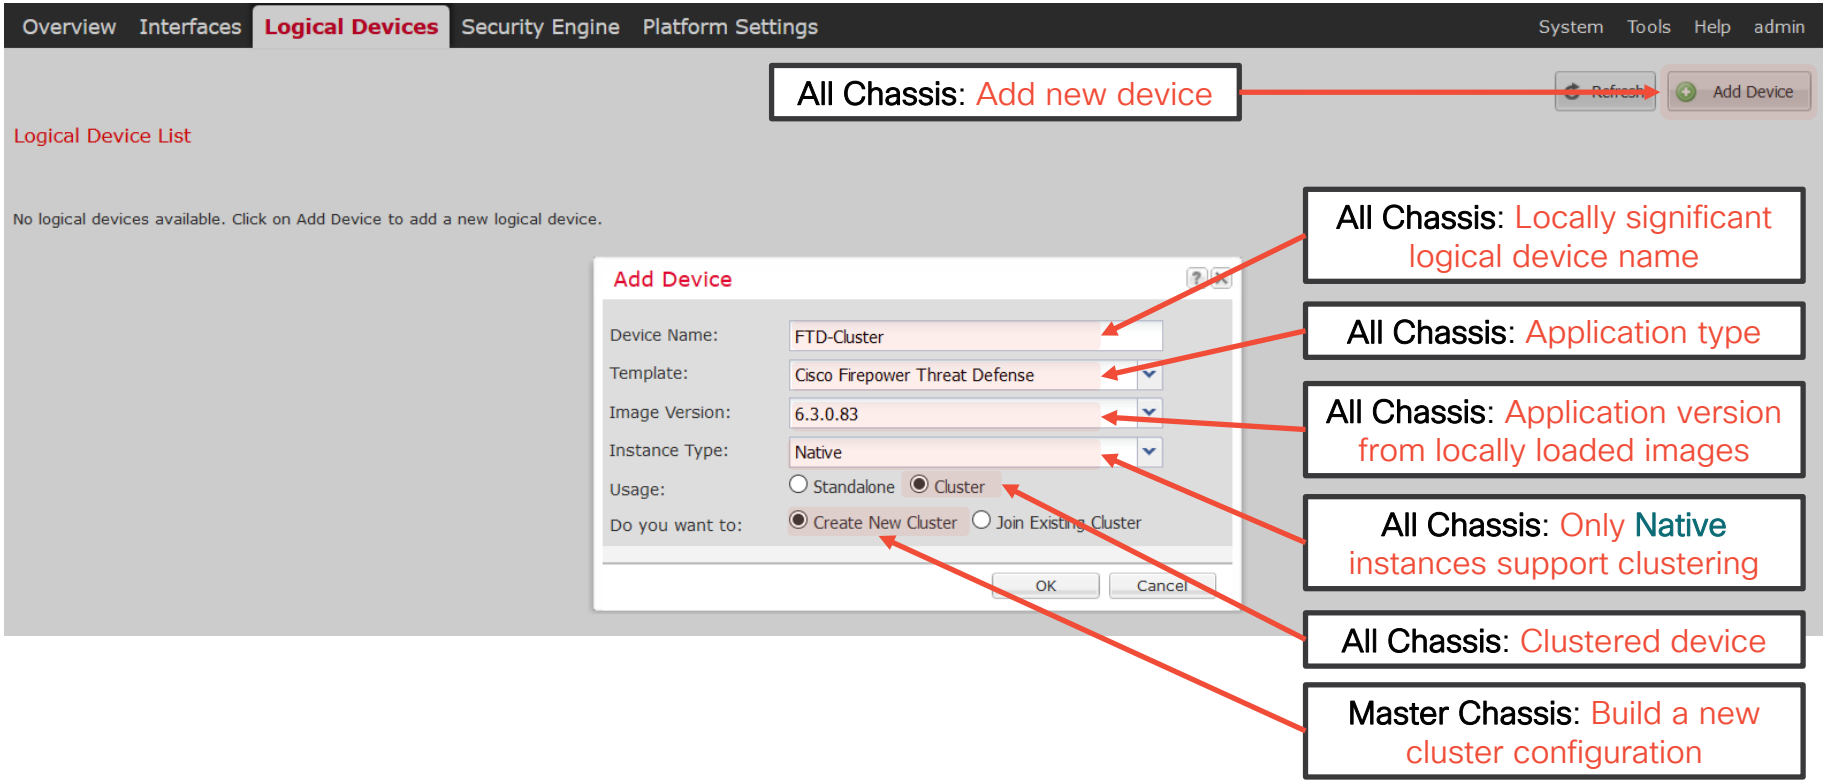

Chassis Manager: FTD Interface Assignment

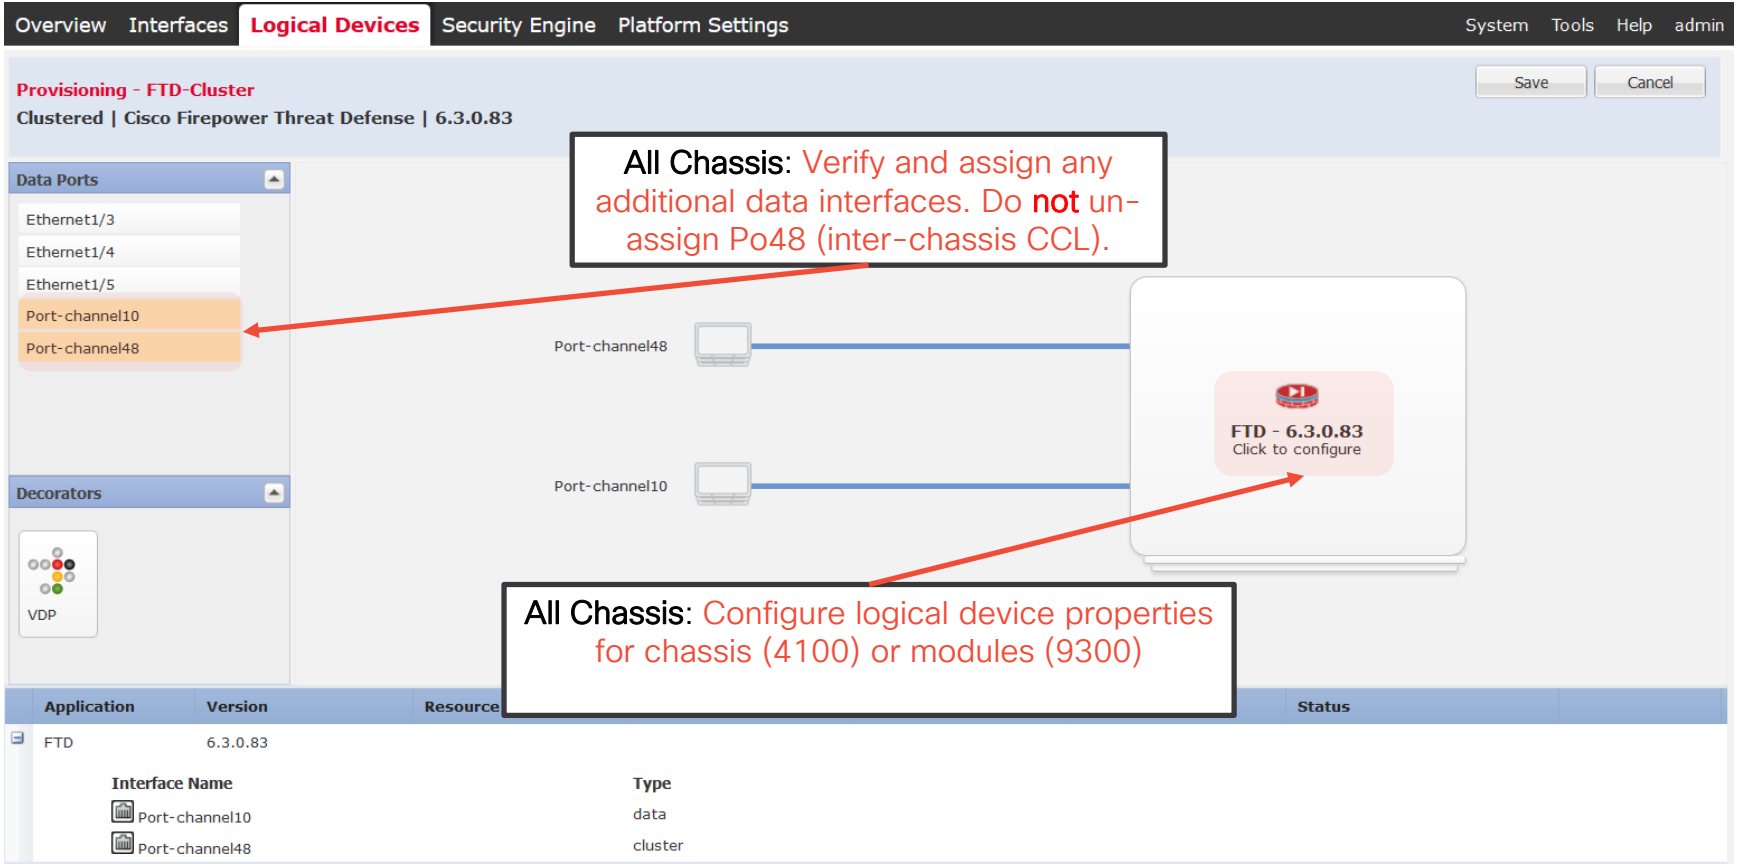

Chassis Manager: FTD Cluster Bootstrap

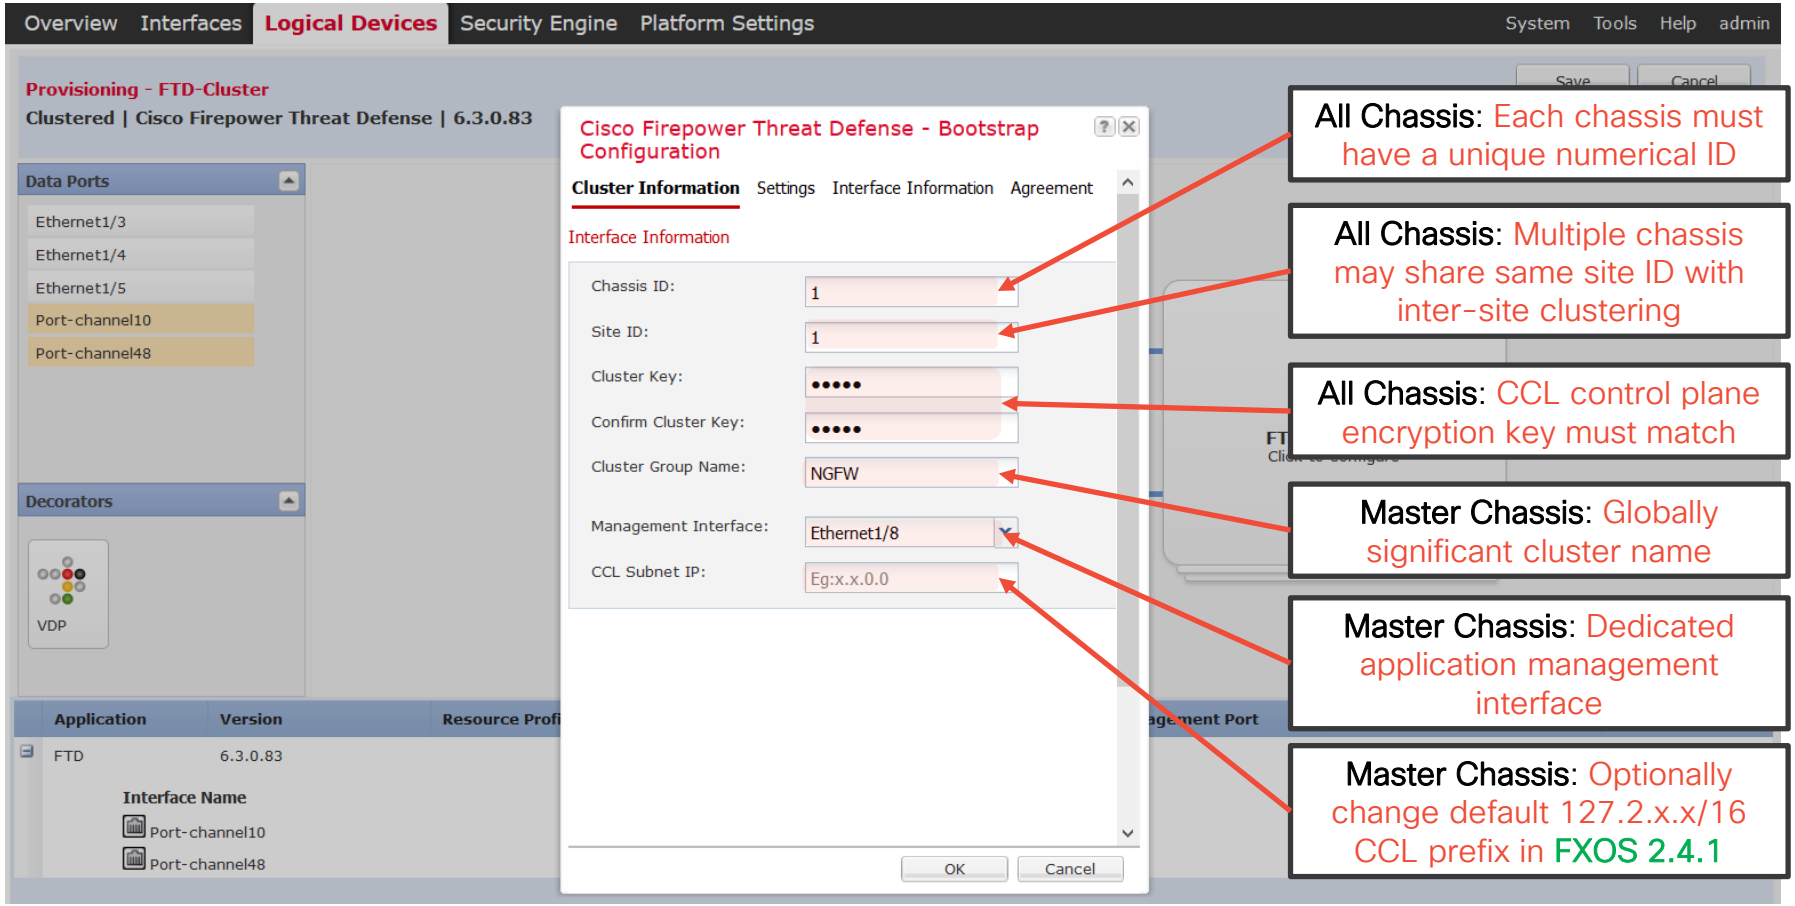

Chassis Manager: FTD Device Settings

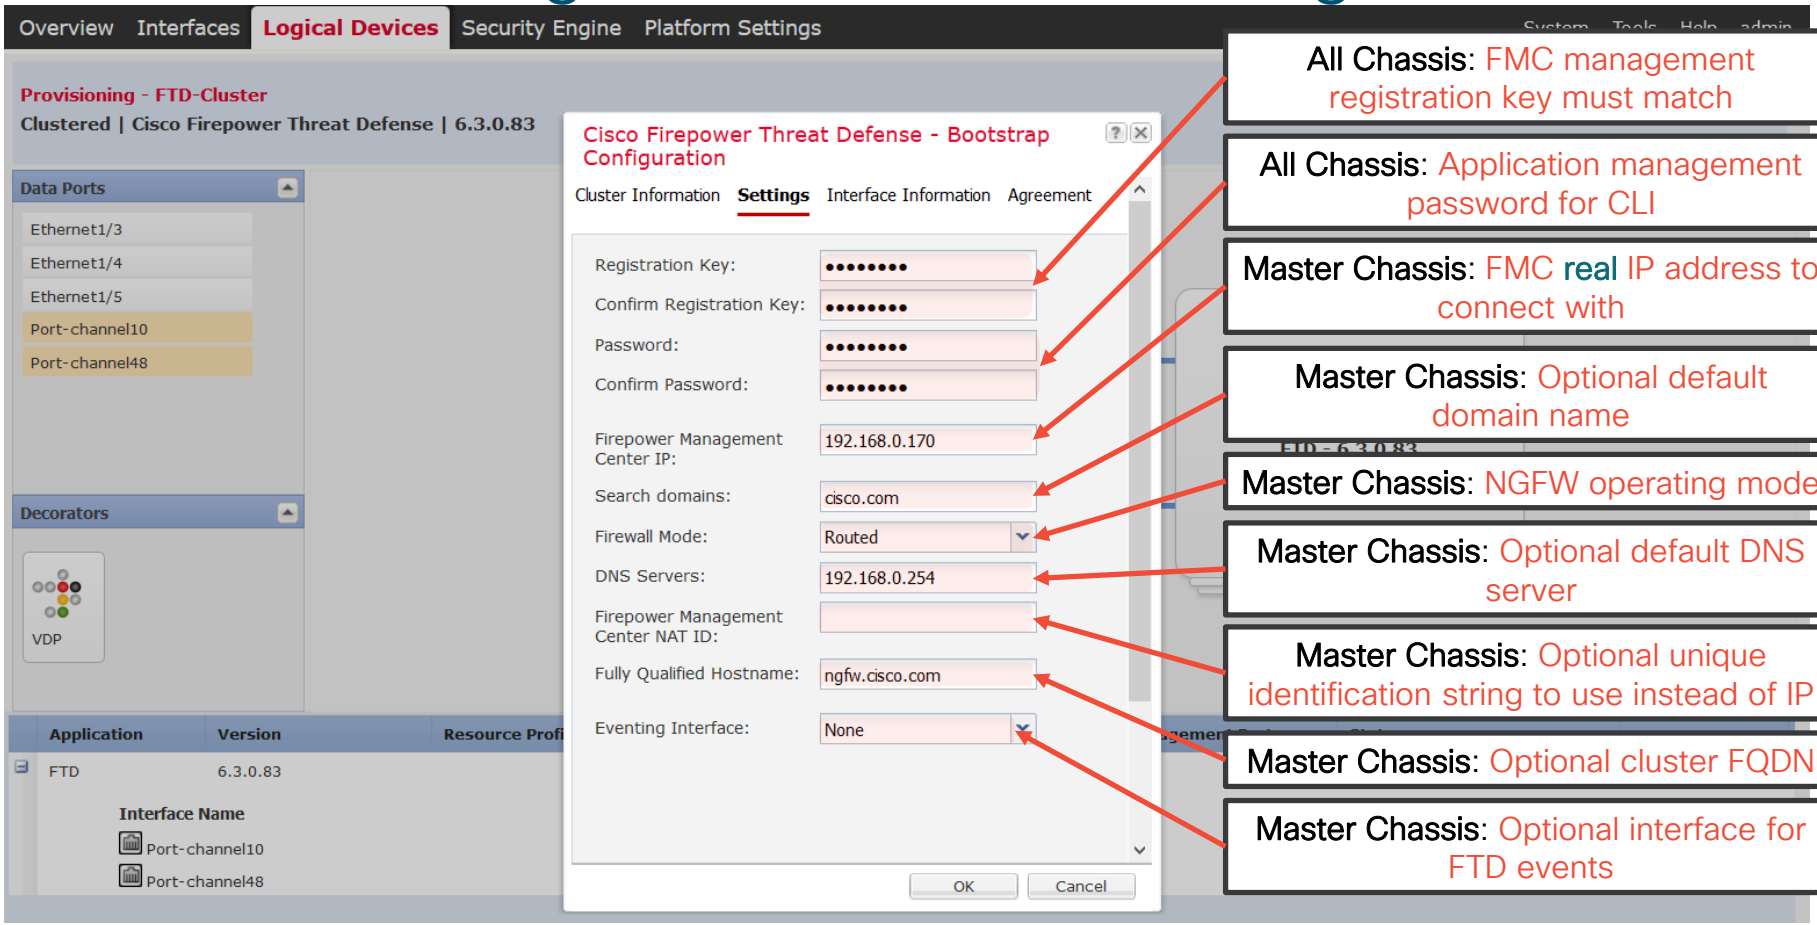

Chassis Manager: FTD Management Interface

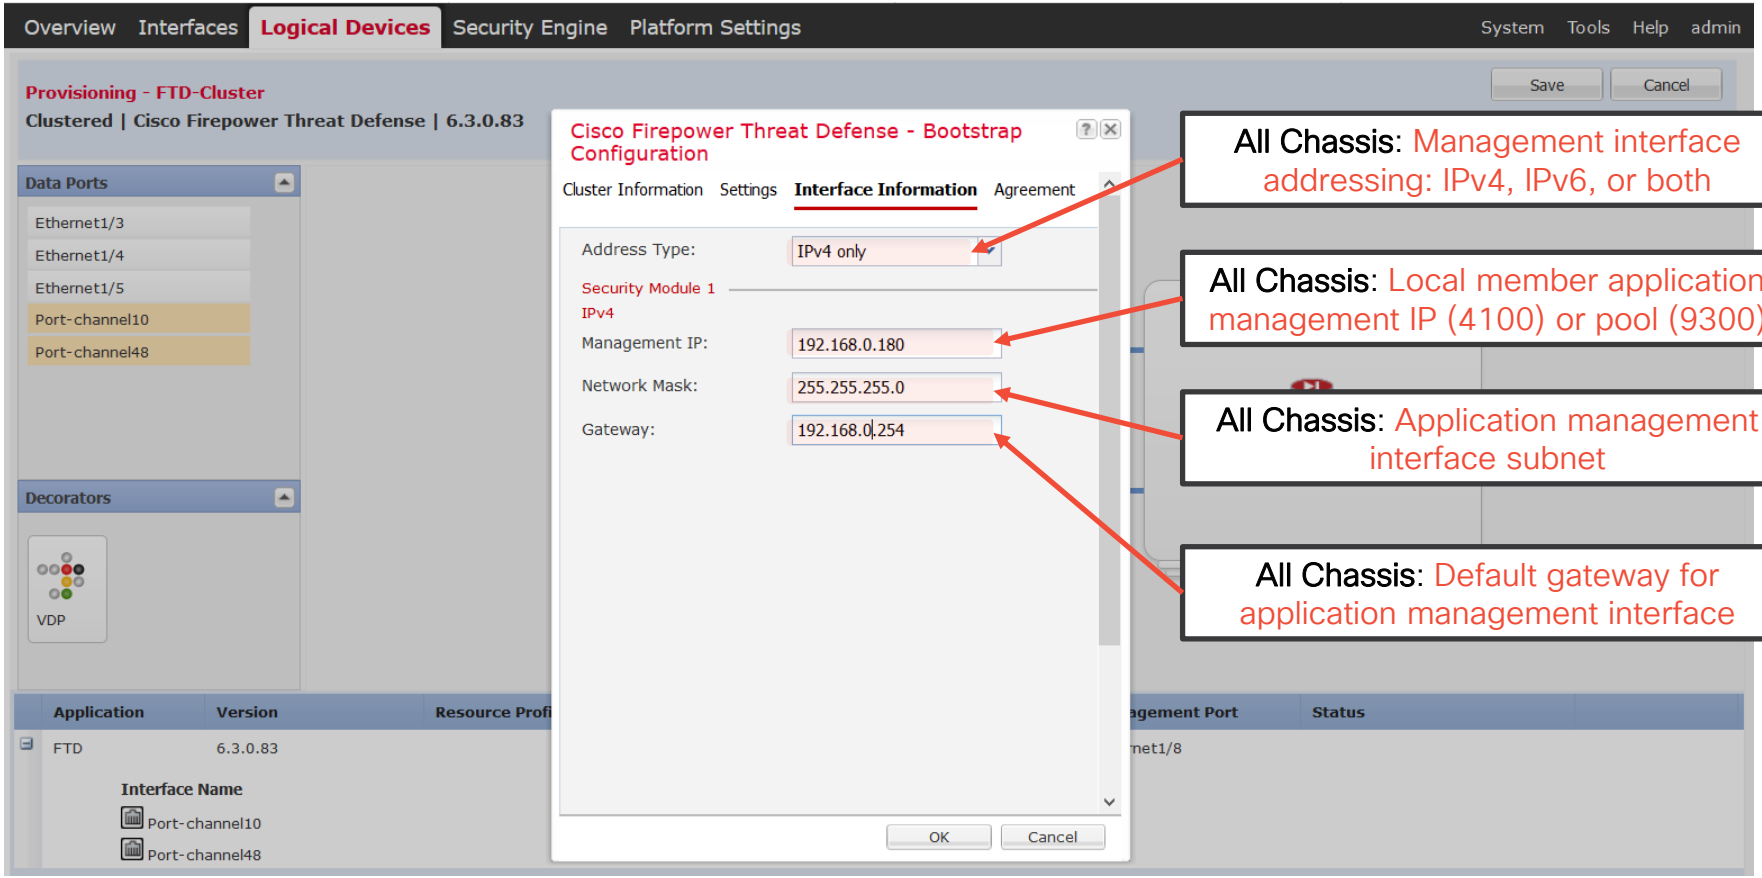

Chassis Manager: FTD Device Installation

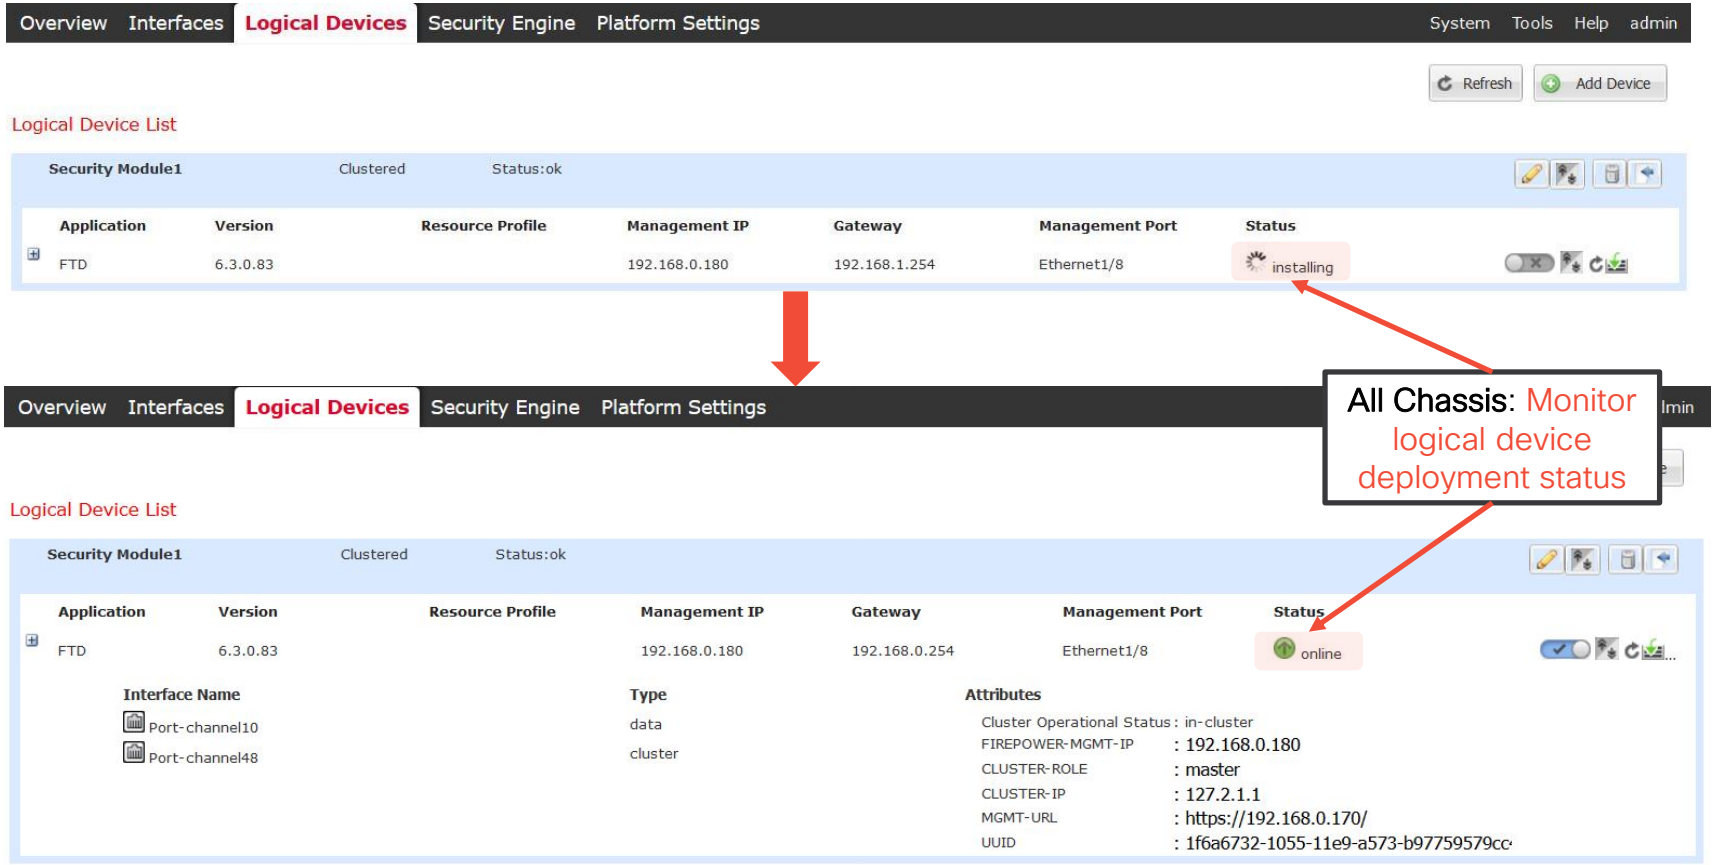

Chassis Manager: Export Cluster Configuration

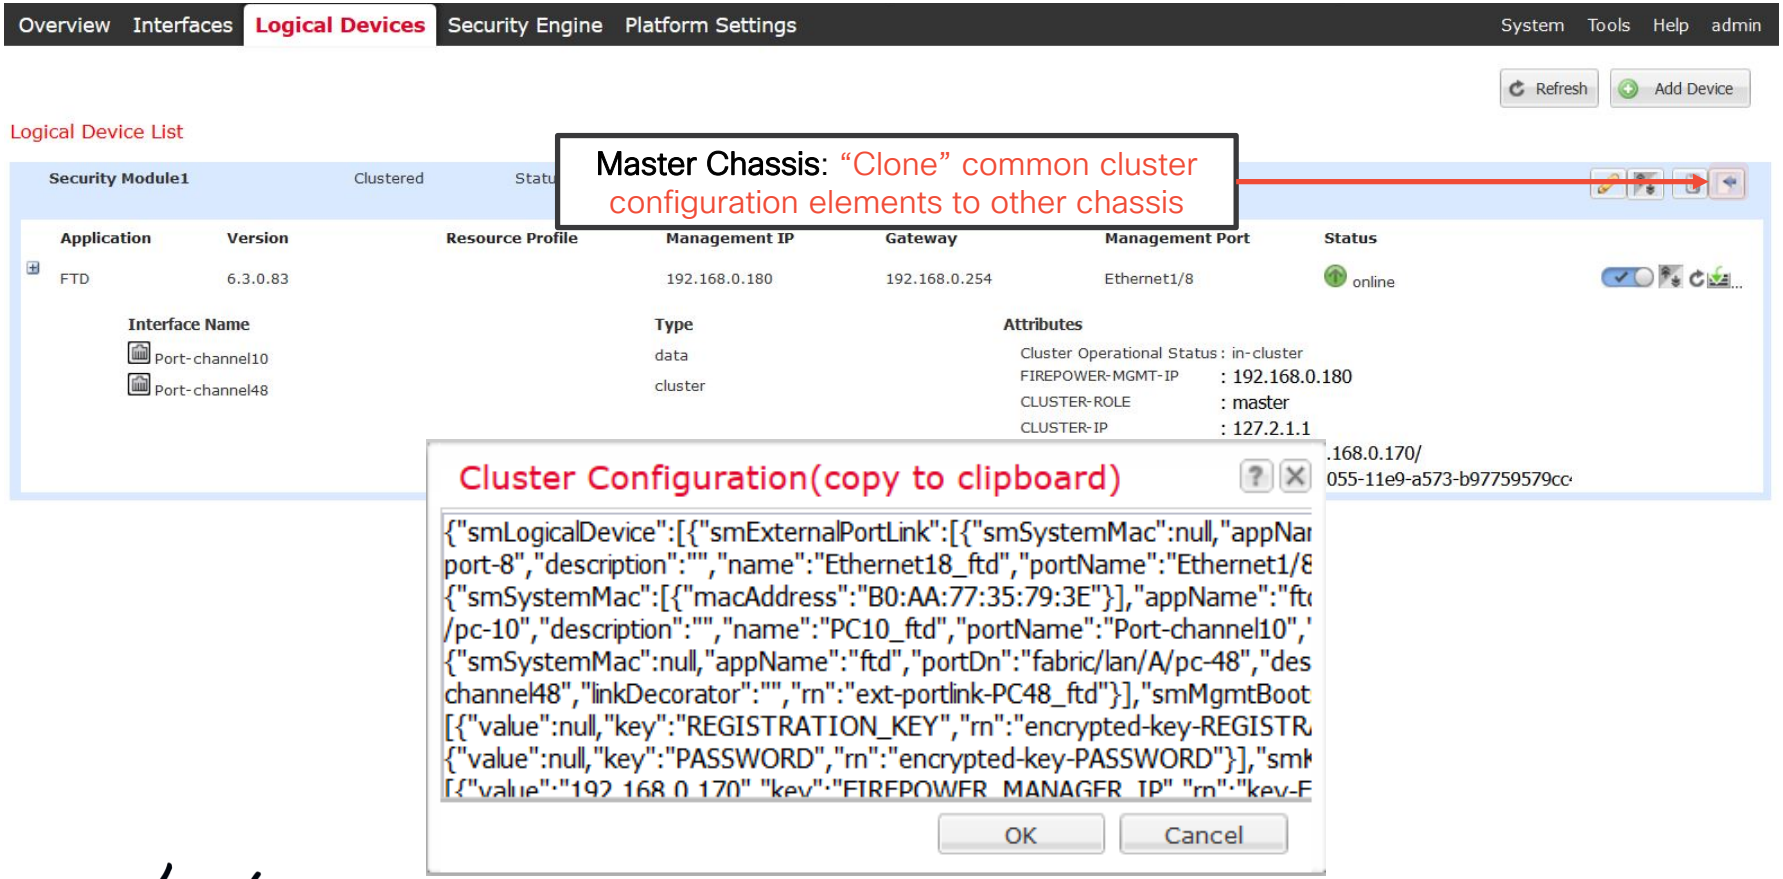

Chassis Manager: Adding Chassis to Cluster

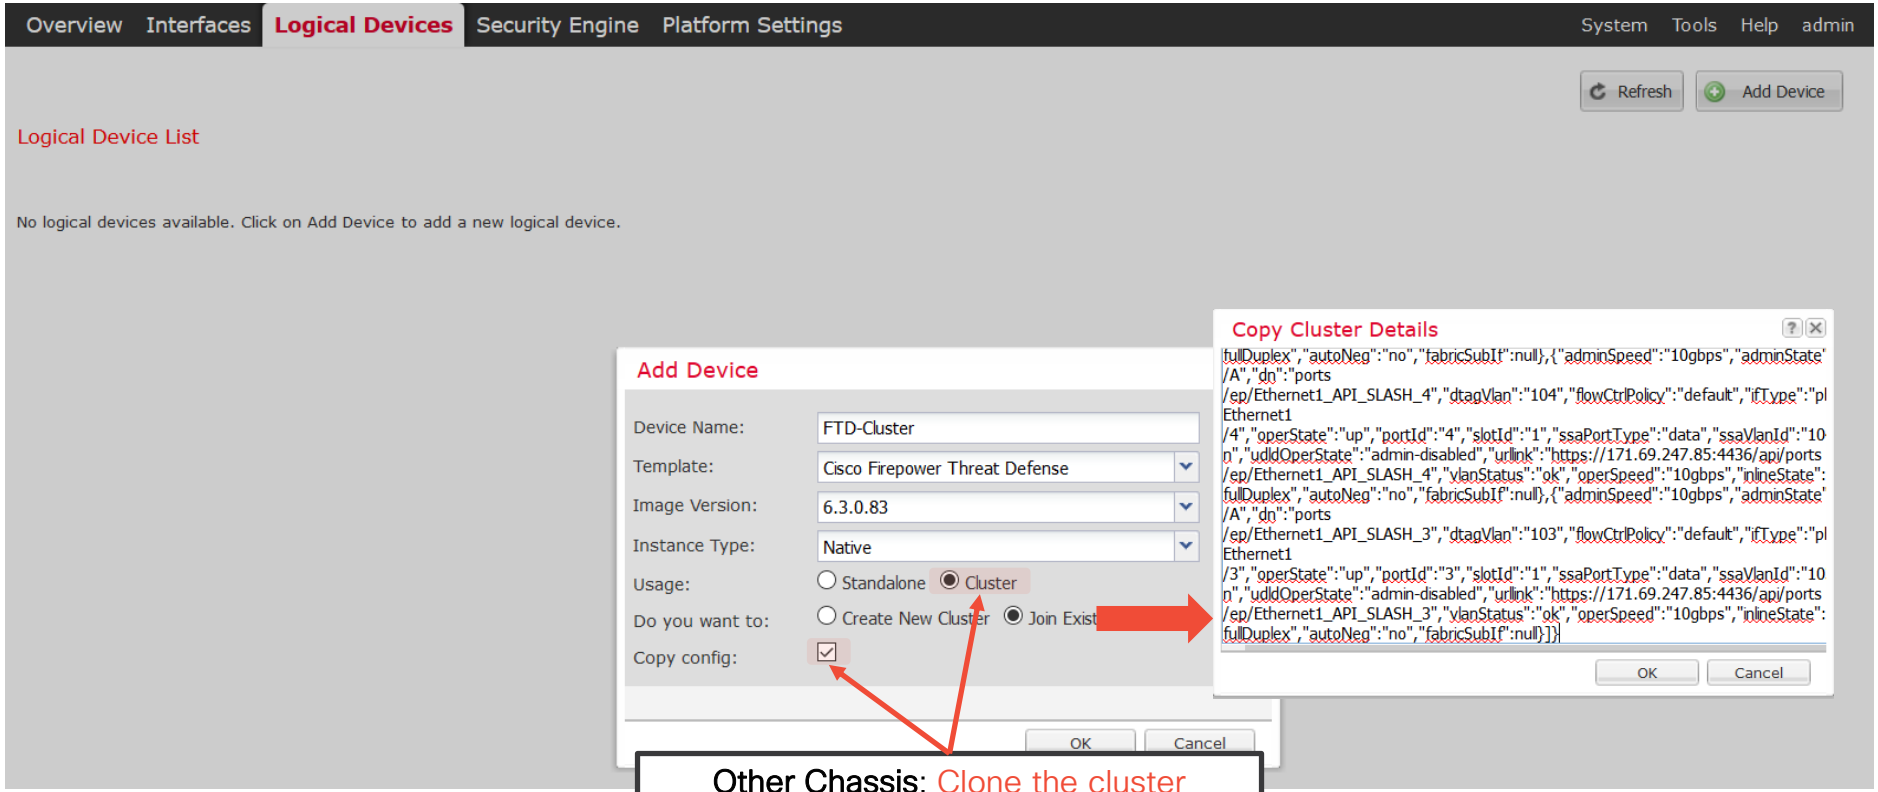

Before FMC 6.3: Add Individual Cluster Members

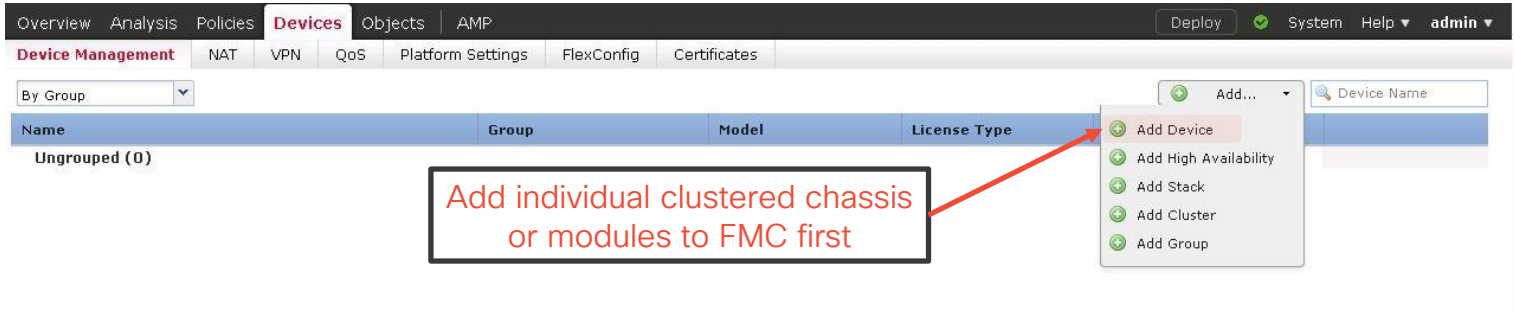

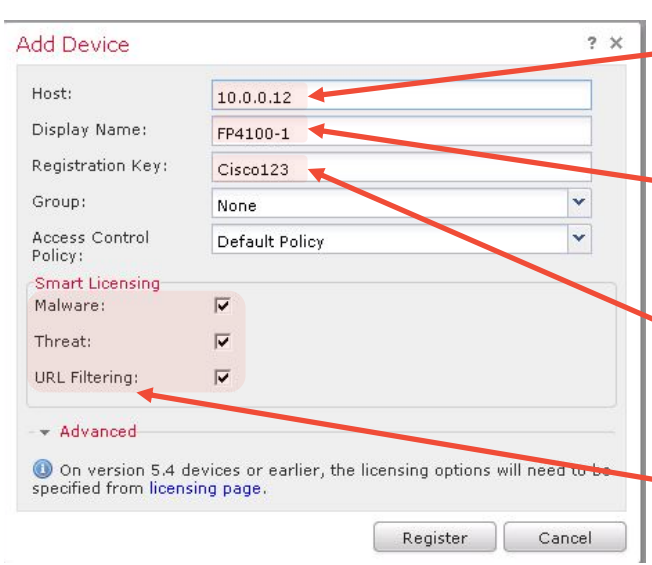

All Chassis/Modules: FTD application management IP

All Chassis/Modules: Unique display name in FMC

All Chassis: FMC registration key **must** match logical device configuration

All Chassis/Modules: Feature licenses must match across entire cluster

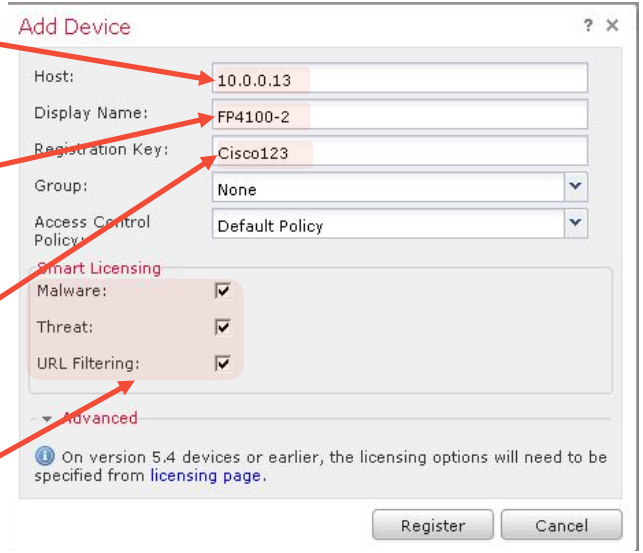

Before FMC 6.3: Add Cluster

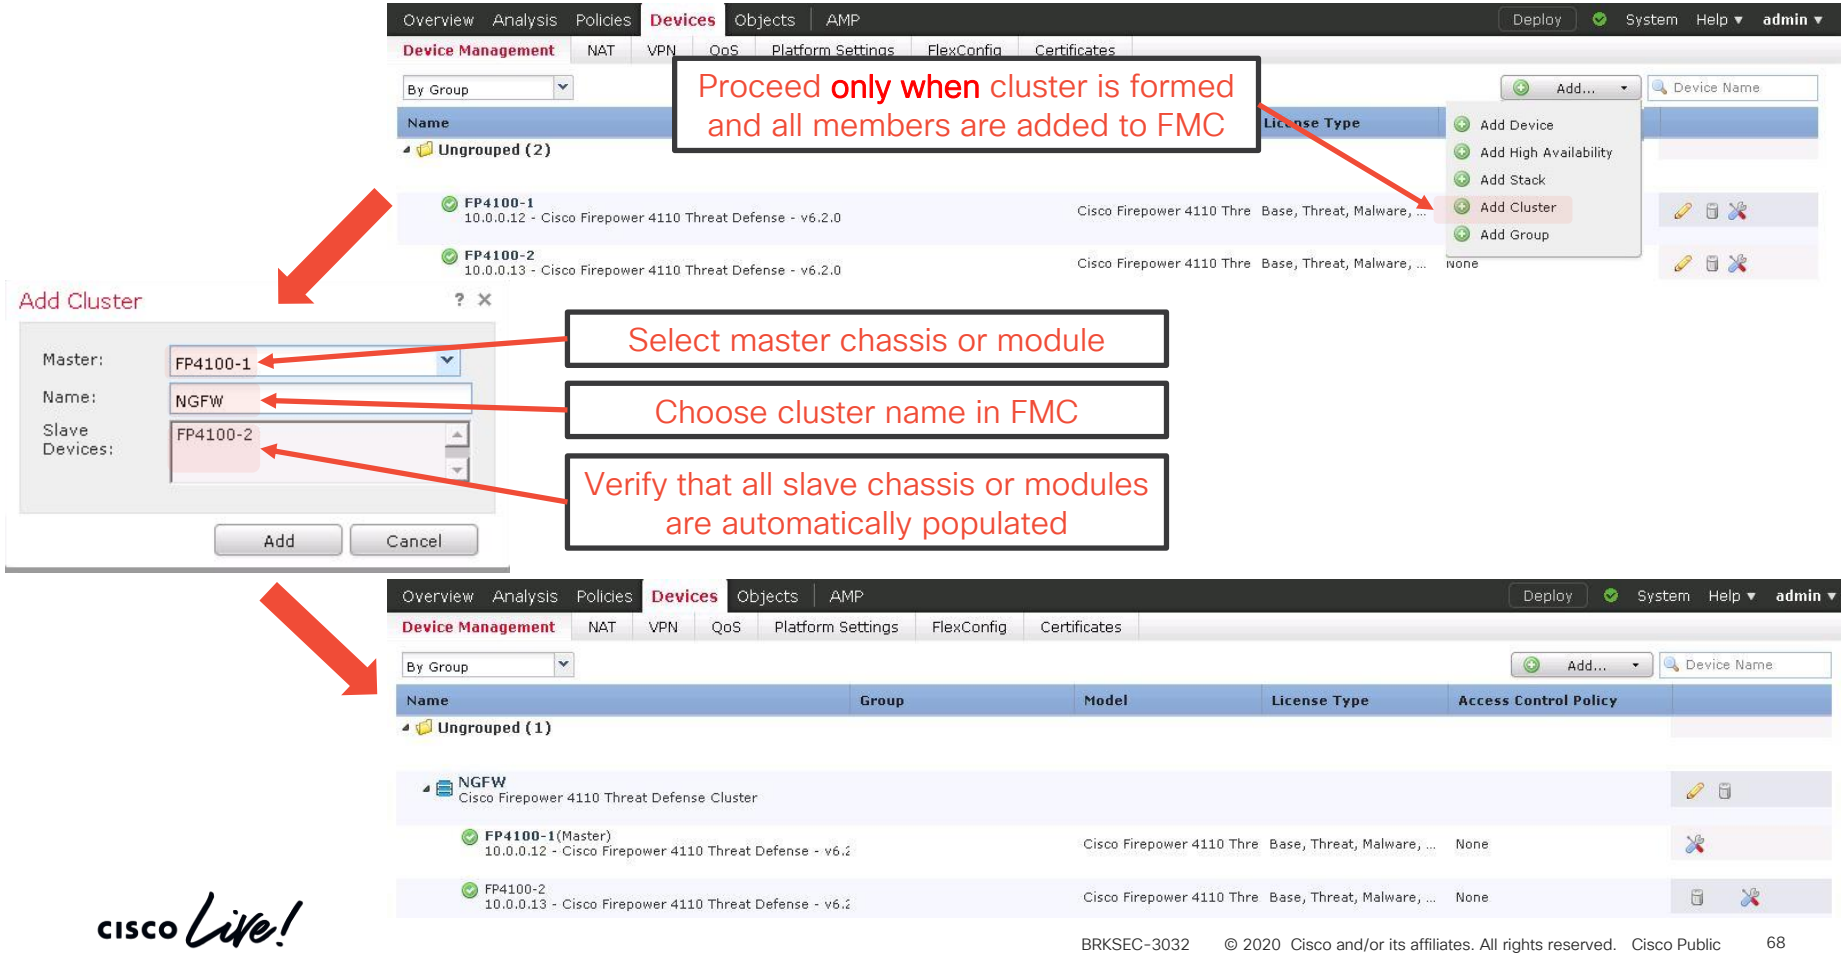

FMC 6.3: Add Entire Cluster

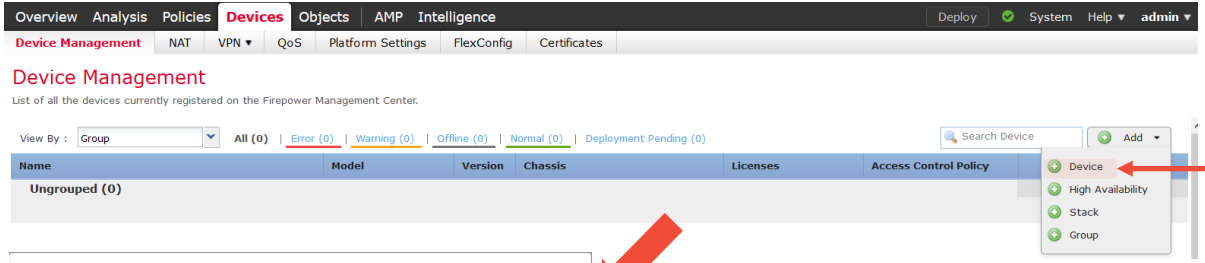

Add **master** unit **only**, other cluster members are discovered automatically

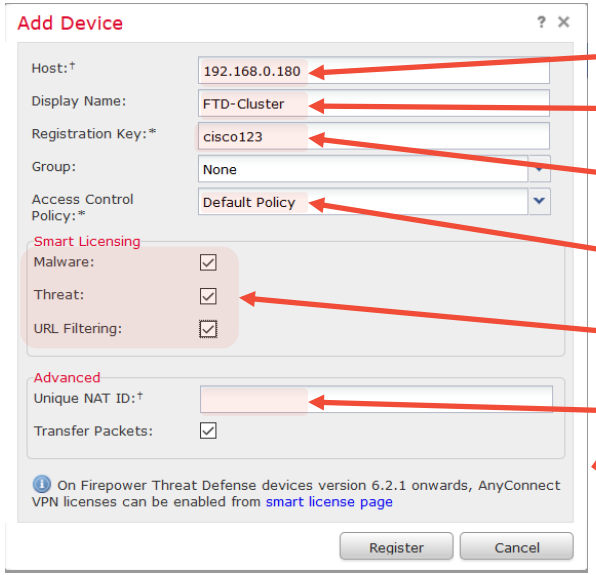

FTD application **real** management IP

Master unit unique display name in FMC

FMC registration key must match logical device configuration across all members

Must assign a default **Main Access Control Policy**

Feature licenses must match across entire cluster

Optional matching identification string to use instead of IP

FMC: Change CCL MTU Settings

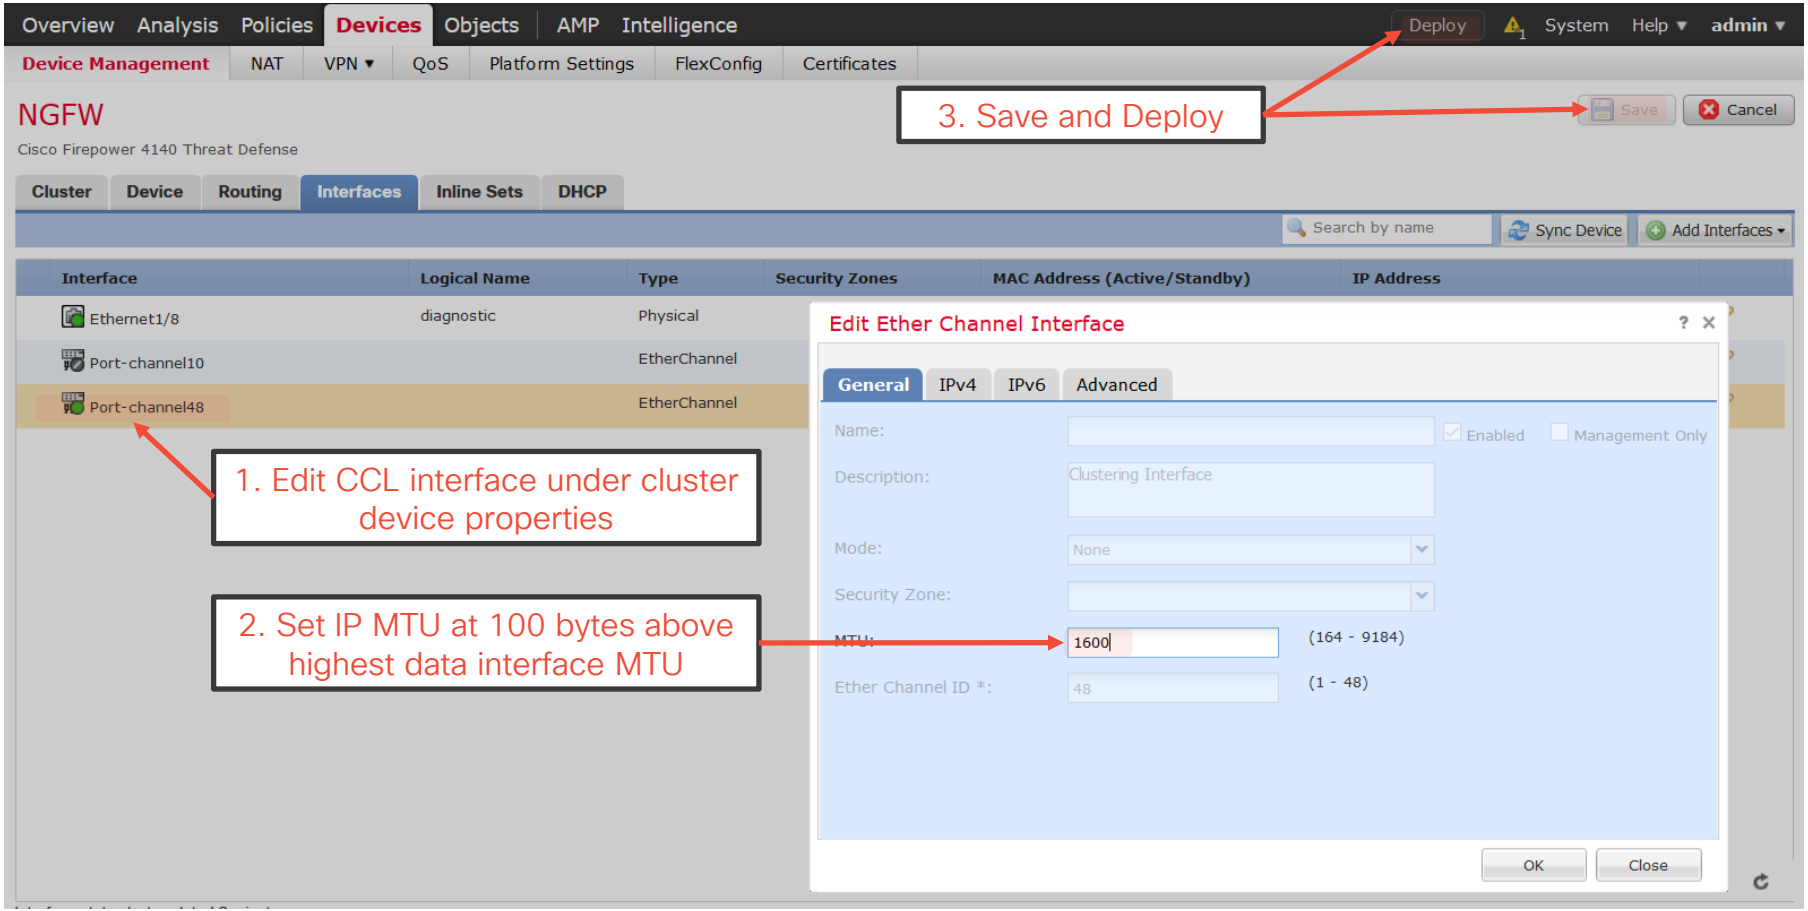

Monitoring and Troubleshooting Clustering

- **show cluster** command group displays aggregated statistics
  - **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 syslog messages 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

Multi-Site Clustering

Inter Data Center (DC) Clustering

- Clustering assumes, but not 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 and no packet loss
  - Routed firewall in Individual interface mode only
- ASA 9.2(1) extends inter-DC clustering to Spanned Etherchannel mode
  - Transparent firewall only
  - Routed firewall support presented design challenges
- ASA 9.5(1) adds inter-DC Spanned Etherchannel clustering in routed mode
- FTD 6.2 adds NGFW inter-site clustering through FlexConfig only
- ACI 3.2 Anycast Services for routed ASA and FTD clusters with Multi-Pod

Split or Single Individual Mode Cluster

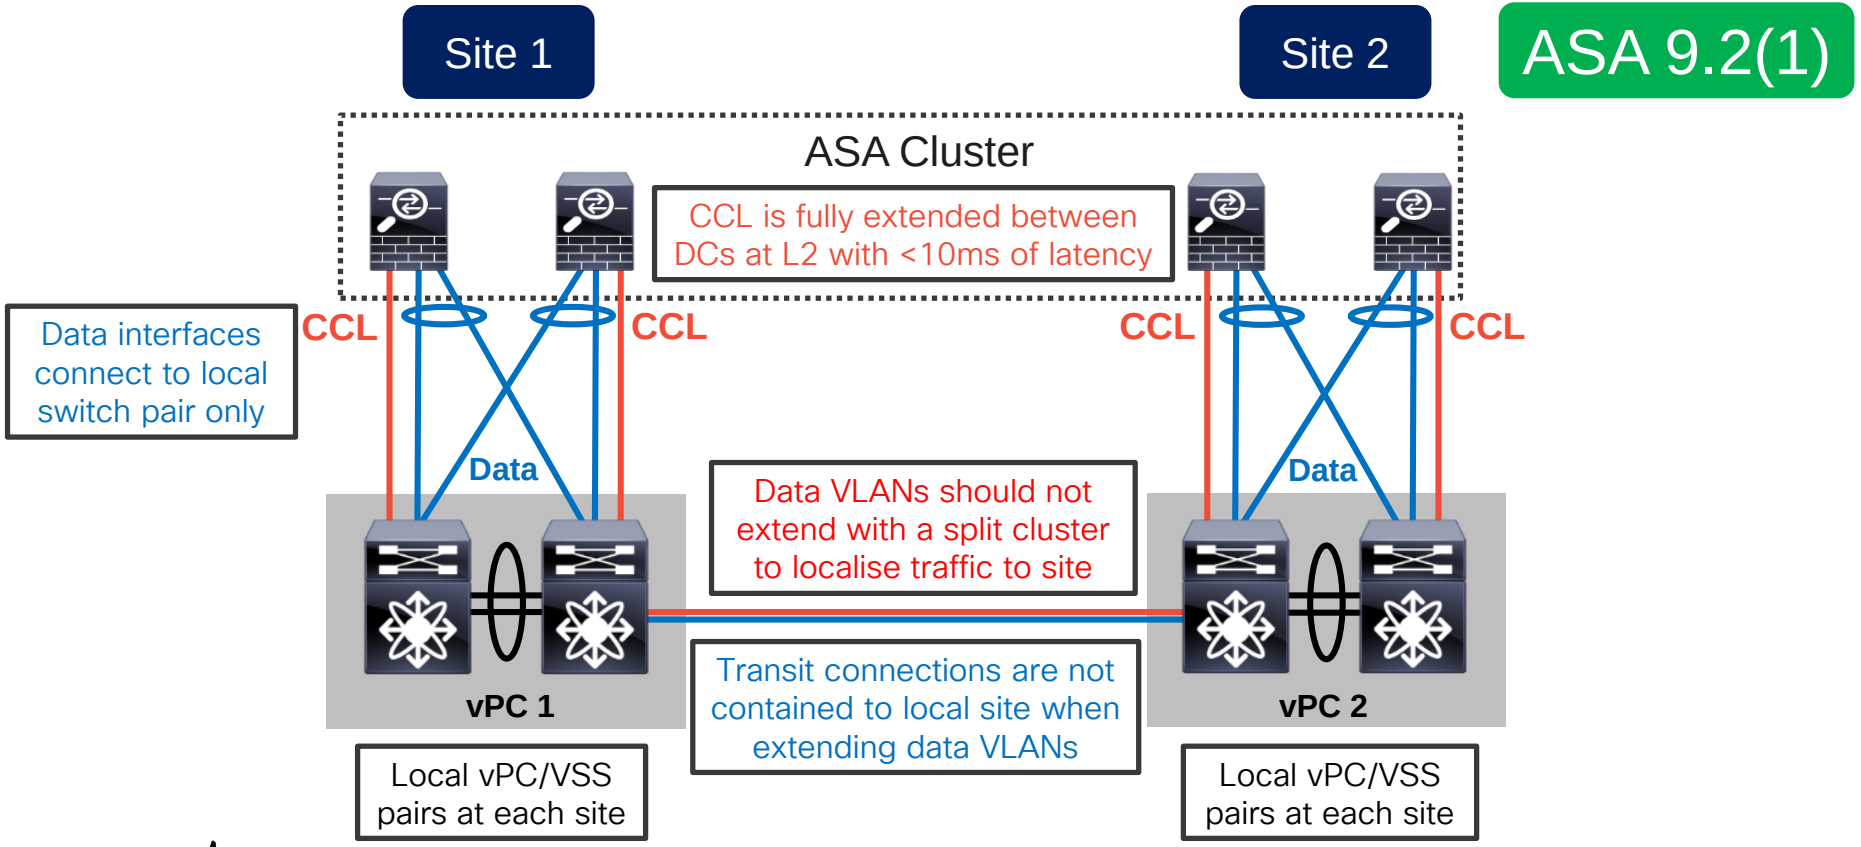

Extended Spanned Etherchannel Cluster

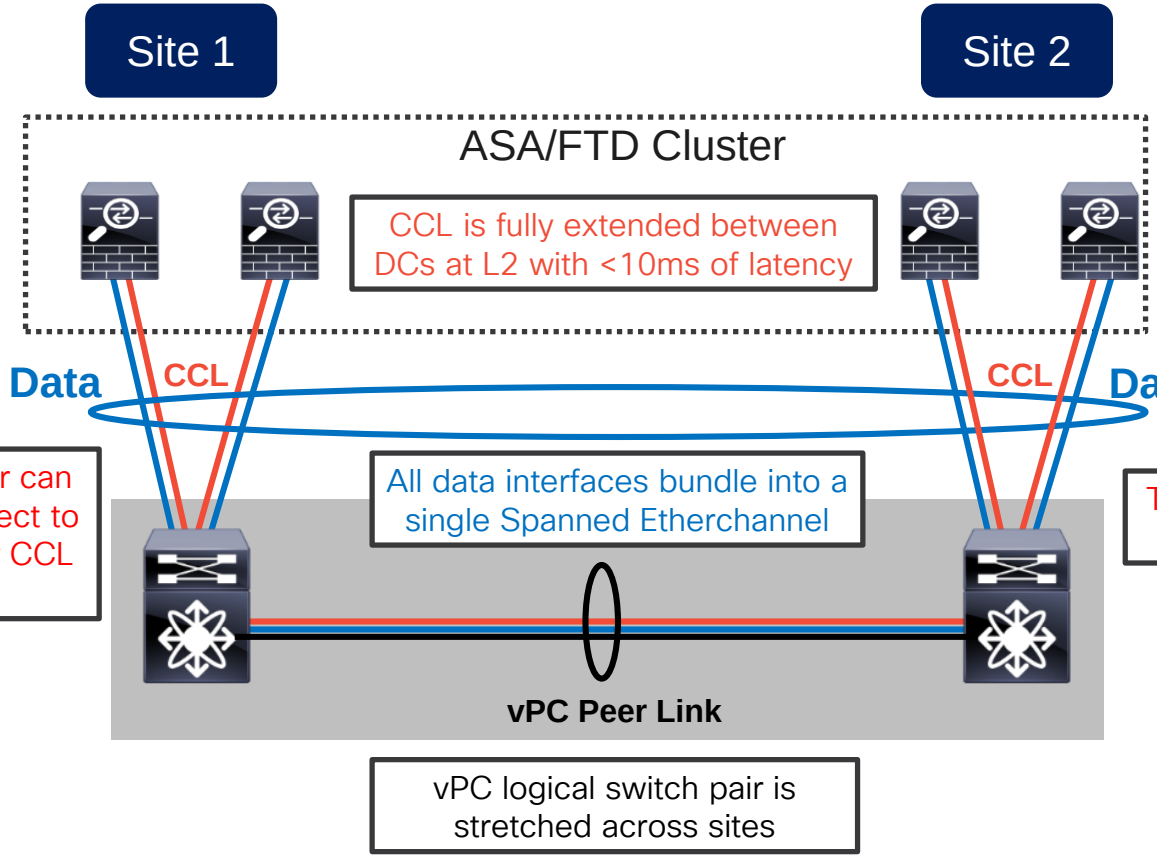

Split Spanned Etherchannel Cluster

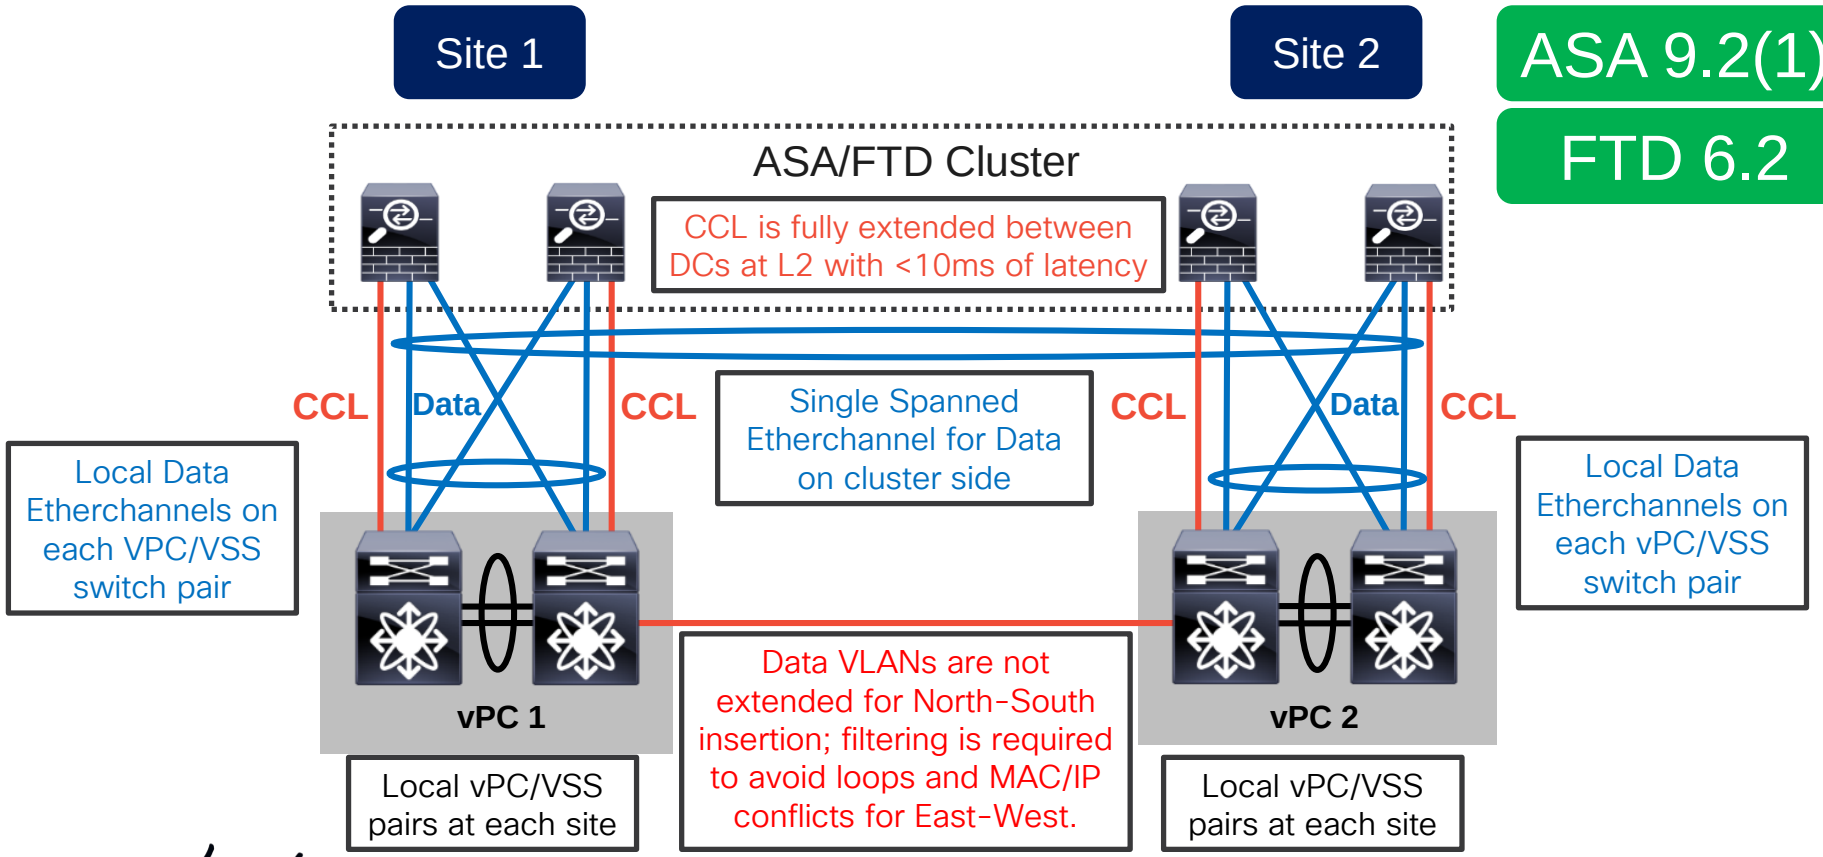

North-South (NS) Inter DC Cluster

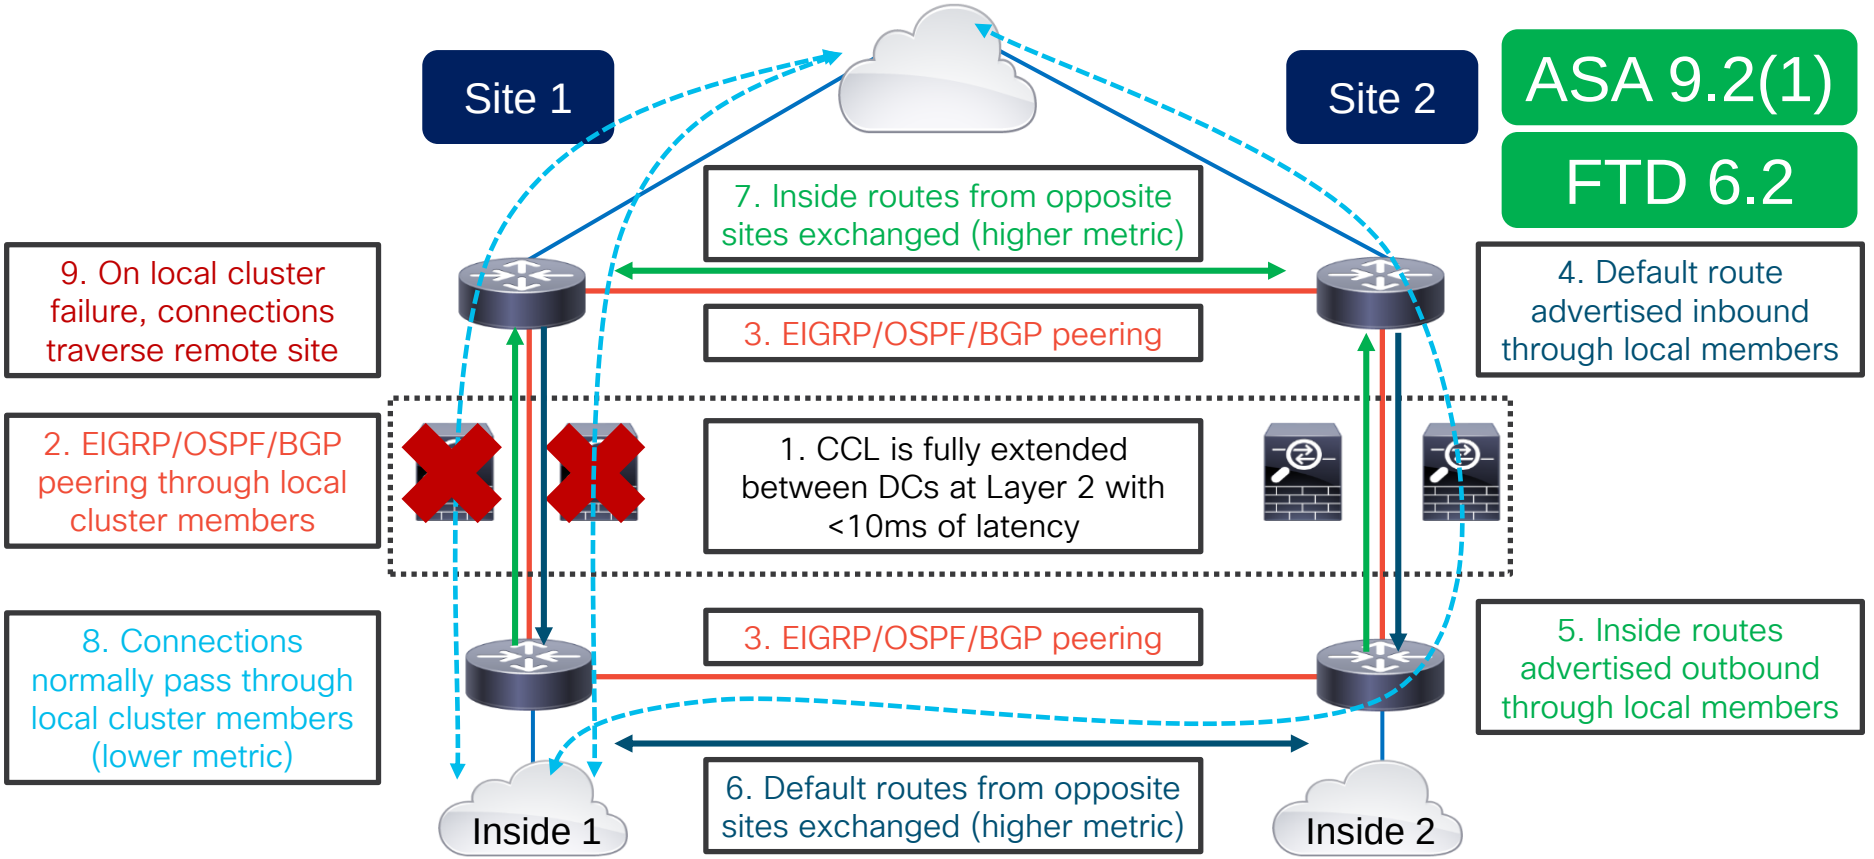

Example: NS Split Spanned Etherchannel Cluster

- 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

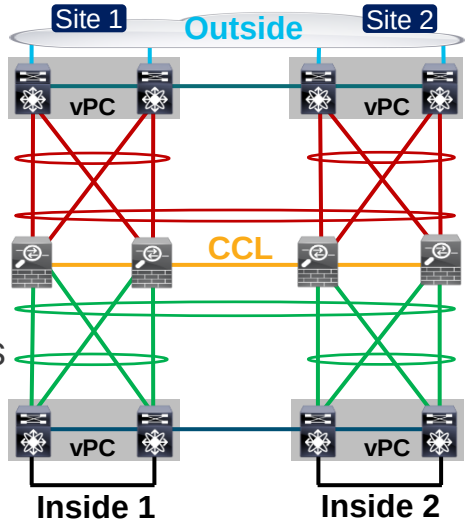

NS Split Spanned Cluster Configuration

```
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-Channel10.10
 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
```

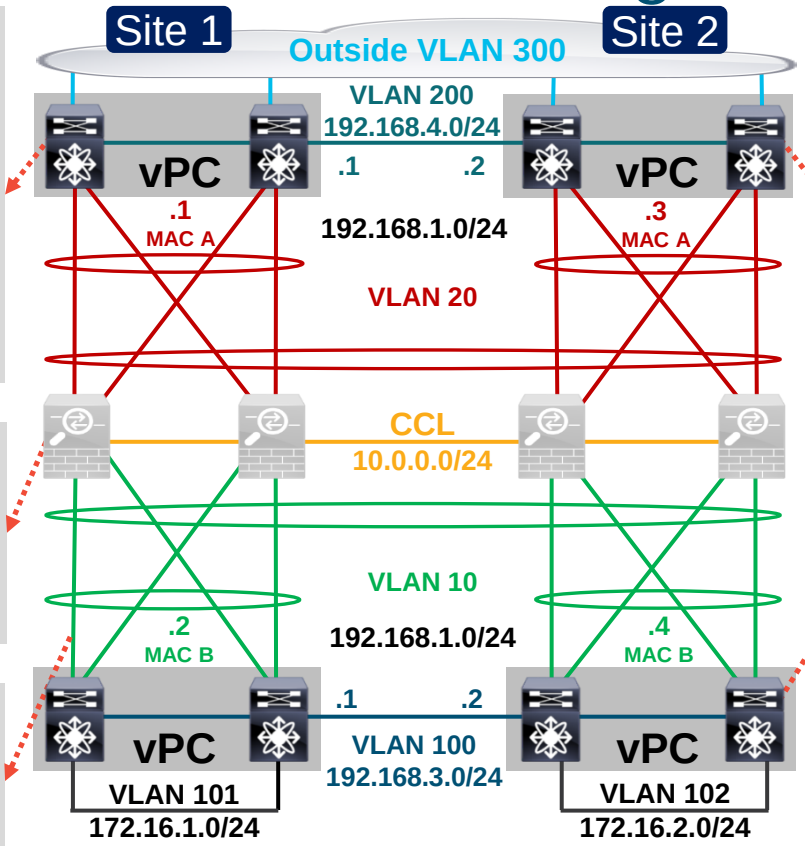


```
ip sla 1
 icmp-echo 192.168.1.4
ip sla schedule 1 life forever start-
 time now
track 1 ip sla 1 reachability
ip access-list PBR
  permit ip any 172.16.2.0 255.255.255.0
route-map PBR
 match ip address PBR
 set ip next-hop verify-availability
  192.168.1.4 track 1
 set ip next-hop 192.168.4.1
interface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

Example: NS 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

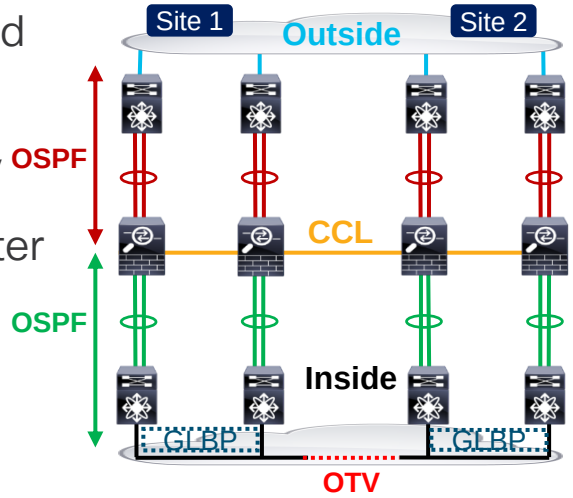

NS Split Individual Cluster Configuration

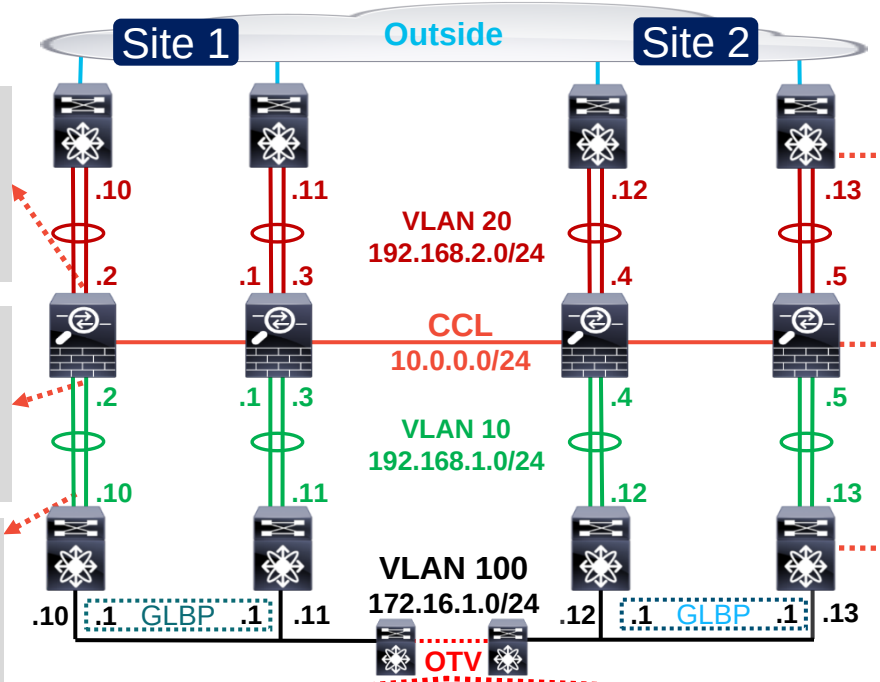


```
ip local pool OUTSIDE 192.168.2.2-
 192.168.2.17
interface Port-Channel10.20
 vlan 20
 nameif FW-outside
 ip address 192.168.2.1 255.255.255.0
  cluster-pool OUTSIDE
```

```
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

Locator/Identifier Separation Protocol (LISP)

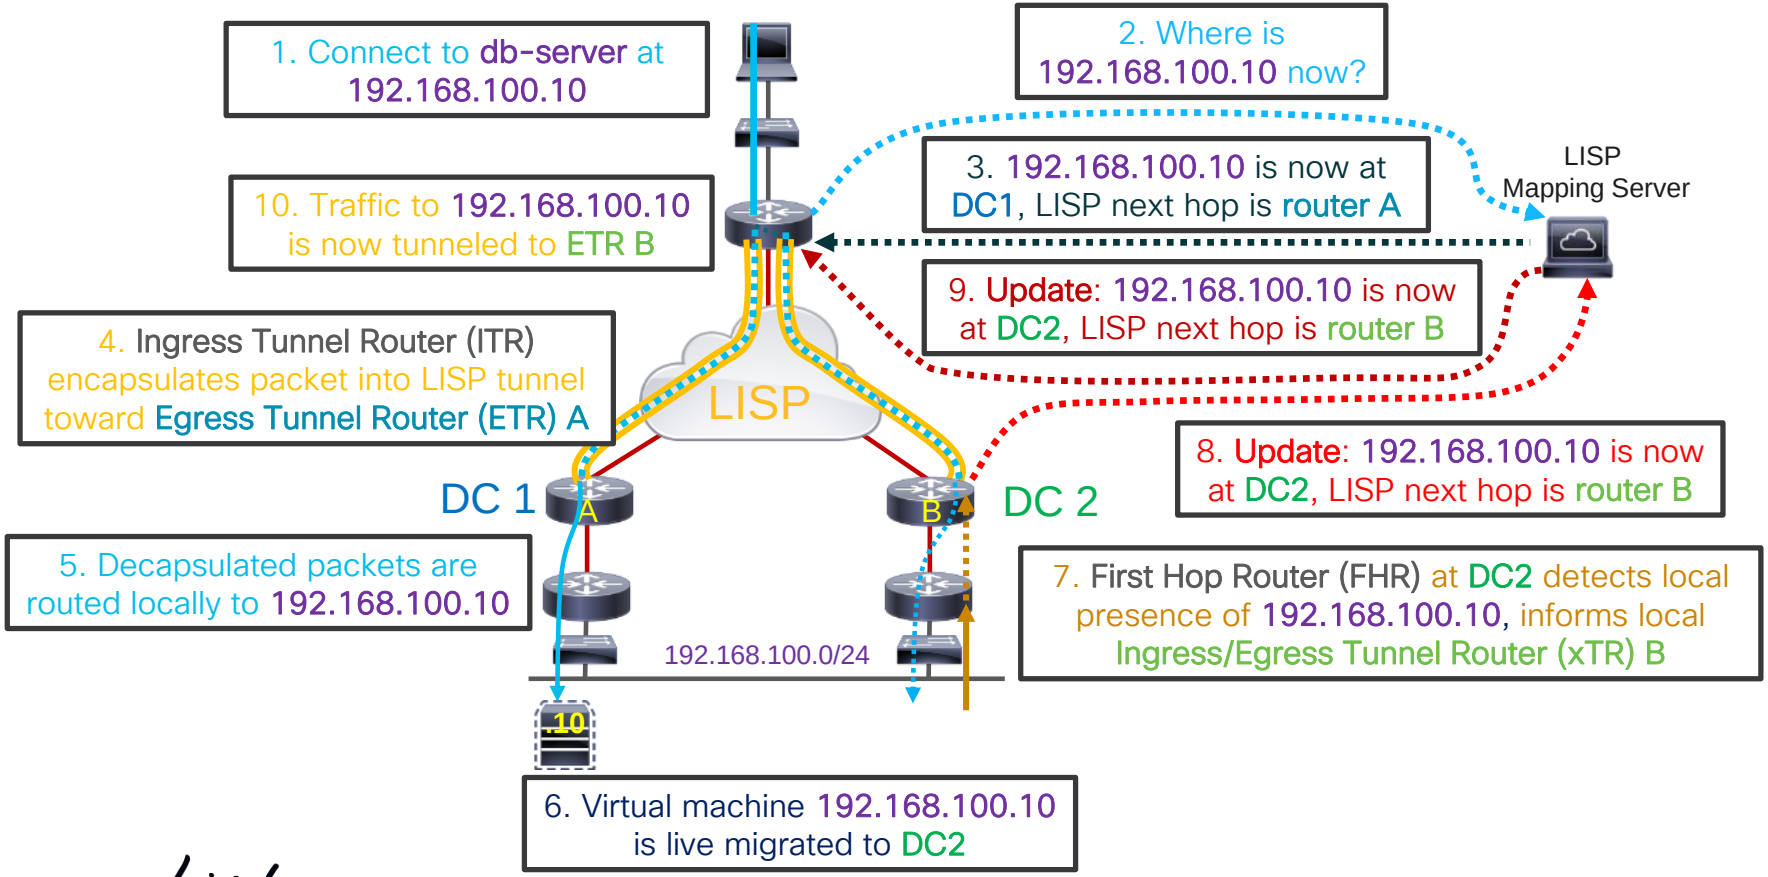

Dynamic Owner Reassignment with LISP

- Move flow ownership with Virtual Machines
  - 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
```

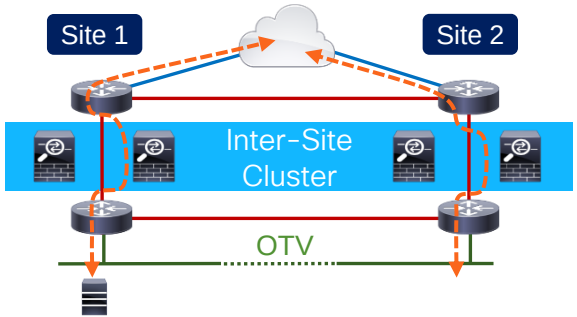

Select specific applications or flows that are eligible for owner reassignment

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 (EW) Inter DC Cluster

Not recommended due to OTV filtering complexity; use Routed East-West insertion instead.

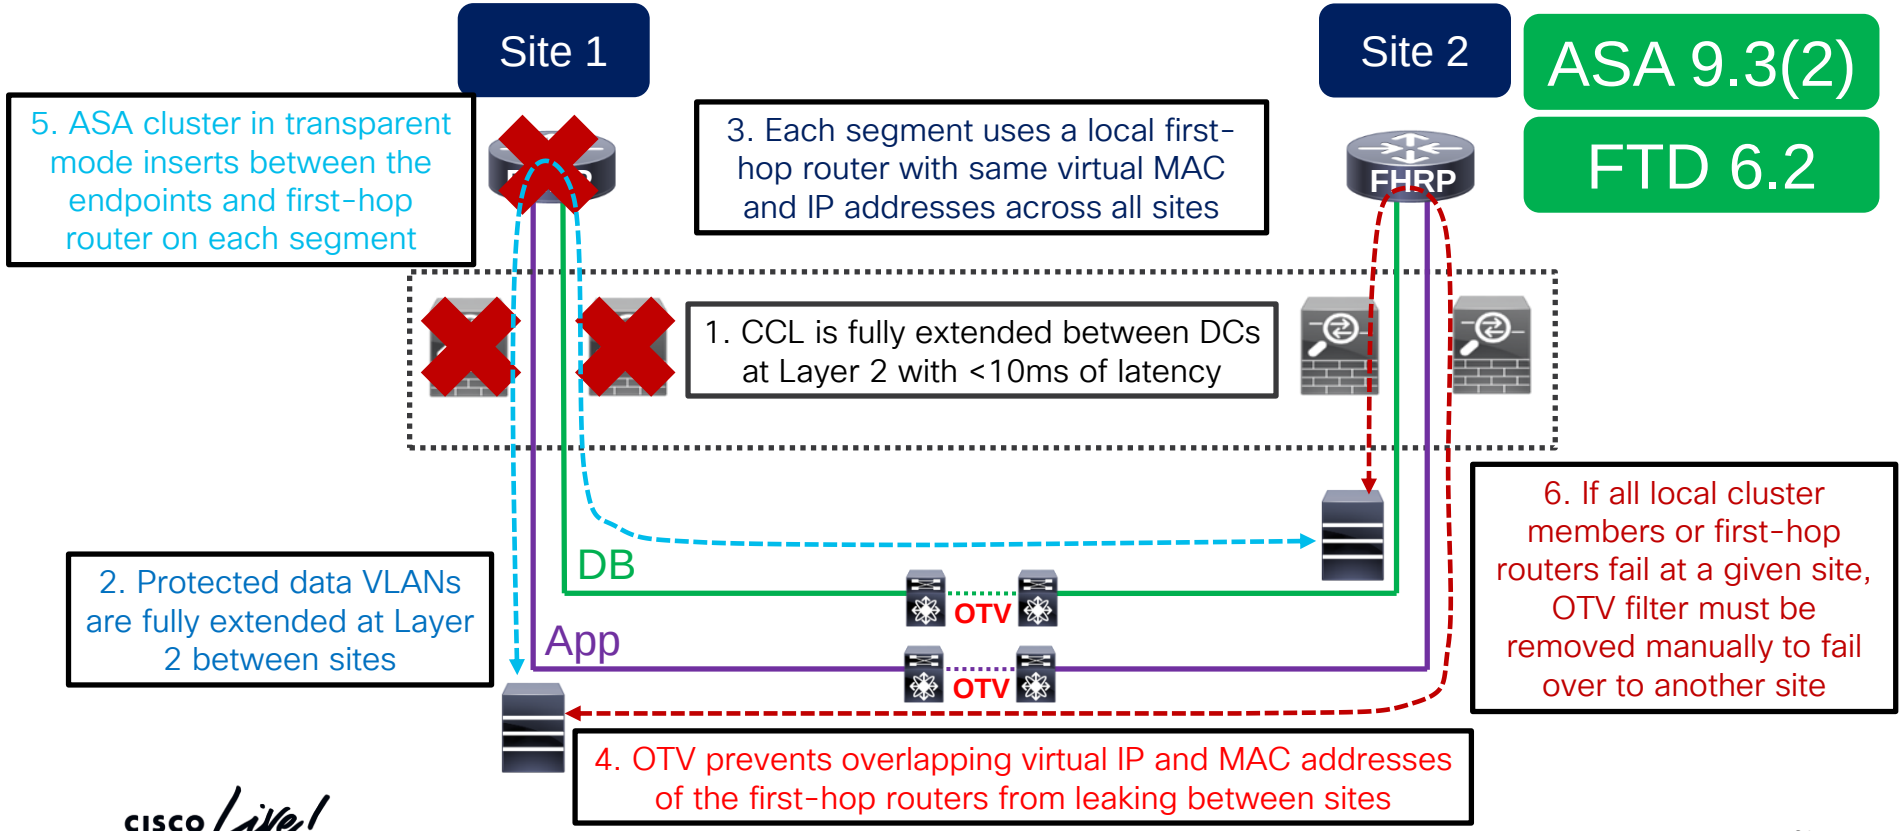

Example: EW Transparent Spanned Cluster

- A vPC pair of Nexus switches at each site
  - Transparent Split Spanned Etherchannel cluster in to separate internal segments
  - Separate Etherchannel to local cluster members per vPC pair
  - Passing firewall twice between segments is acceptable
- Internal VLANs are fully extended between sites with OTV
  - Each site uses localized HSRP as first hop router
  - HSRP traffic is blocked between sites
  - Upstream SVI/HSRP MAC statically bound to outside on cluster
  - Full Layer 2 reachability from each router to remote site
  - Must manually remove OTV filters on full upstream path failure
- Must implement LISP to avoid excessive flow redirection

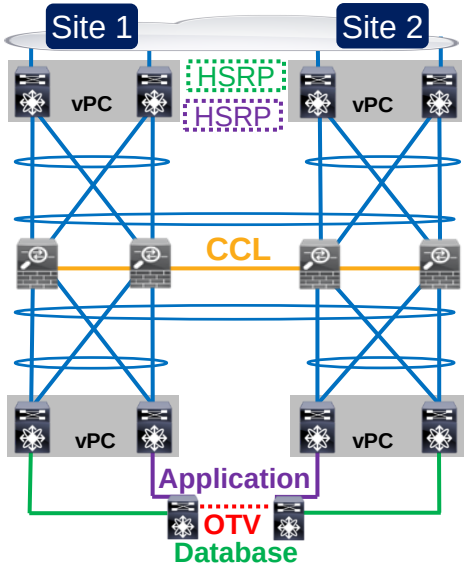

EW Transparent Spanned Cluster Configuration

```
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-Channel10.100
 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
mac-address-table static DB-outside A
mac-address-table static DB-outside B
mac-address-table static App-outside A
mac-address-table static App-outside B
```

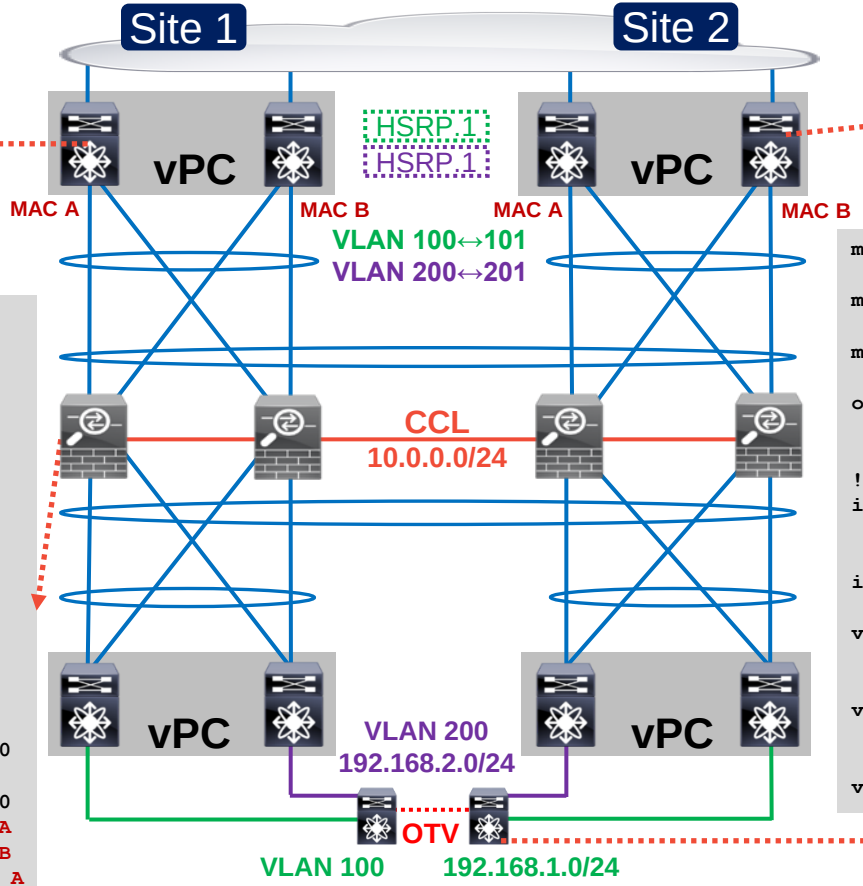


```
interface Vlan101
 ip address 192.168.1.3/24
 hsrp 10
  ip 192.168.1.1
interface Vlan201
 ip address 192.168.2.3/24
 hsrp 20
  ip 192.168.2.1
```

```
mac-list HSRP_MAC seq 10 deny
     0000.0c07.ac00 ffff.ffff.ff00
mac-list HSRP_MAC seq 20 deny
     0000.0c9f.f000 ffff.ffff.ff00
mac-list HSRP_MAC seq 30 permit
     0000.0000.0000 0000.0000.0000
otv-isis default
 vpn Overlay1
  redistribute filter route-map HSRP_MAC
!
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 HSRP_FILTER vlan-list 100, 200
```


Routed East-West (EW) Inter DC Cluster

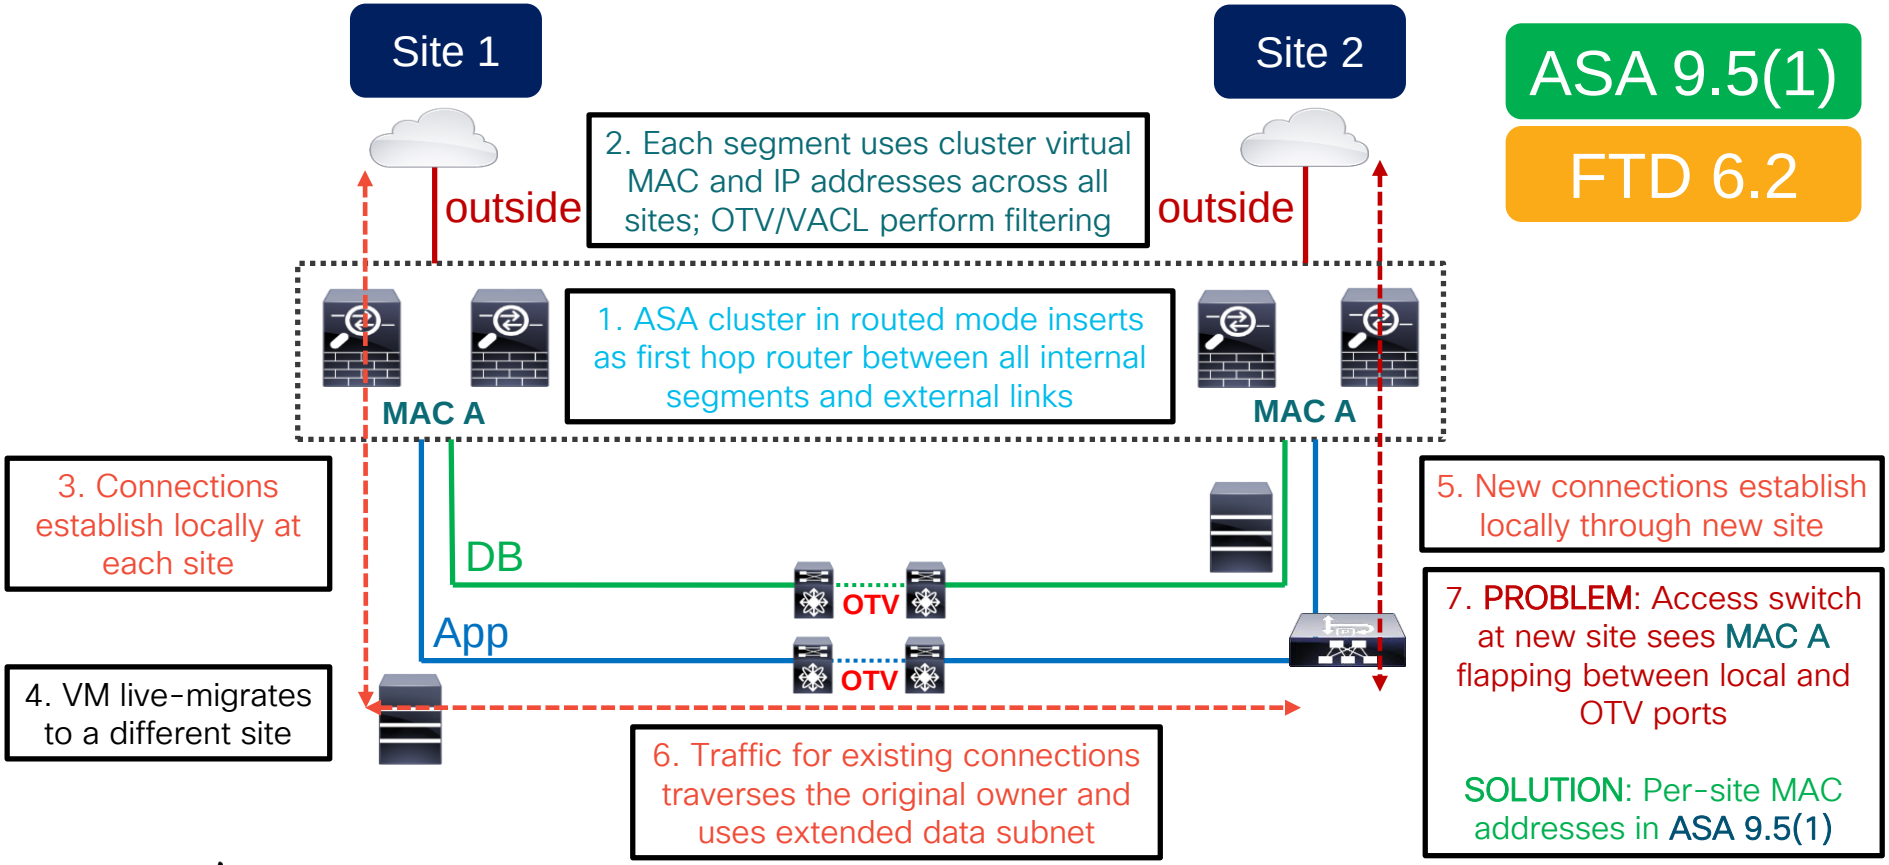

Per-Site MAC Addresses

- Routed Spanned Etherchannel cluster extends MAC addresses in 9.5(1)
  - Global interface MAC address is used to receive and source frames by default
  - Per-site MAC addresses can be 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, but possible with a shared outside segment
- Global MAC address localization is required by OTV or similar mechanisms

OTV Silent Host Problem

- OTV suppresses **unicast flooding** for unknown MAC addresses by default
  - Hosts that mostly generate local traffic quickly become unreachable across OTV
  - Recommended to set ARP timeout below MAC address table timeout
- **ASA 9.8(3)** and **FTD 6.2.2.2** replicate ARP replies to all sites
  - Refresh MAC table entries in OTV to **partially** combat the **Silent Host** problem
- Cluster global MAC becomes a silent host when per-site MAC is used
  - **ASA 9.12(1)** and **FTD 6.4** generate a periodic GARP for global MAC/IP

```
asa(cfg-cluster)# site-periodic-garp interval 280
```

One unit at each site generates a GARP at this frequency in seconds; default is 280

Example: EW 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 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
  - Must manually remove OTV filters on full upstream path failure
  - One silent host with a very long ARP timeout at site 1

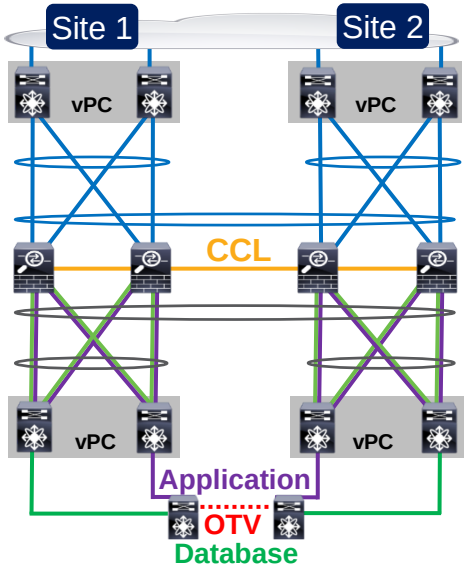

EW Routed Spanned Cluster Configuration

```
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
```

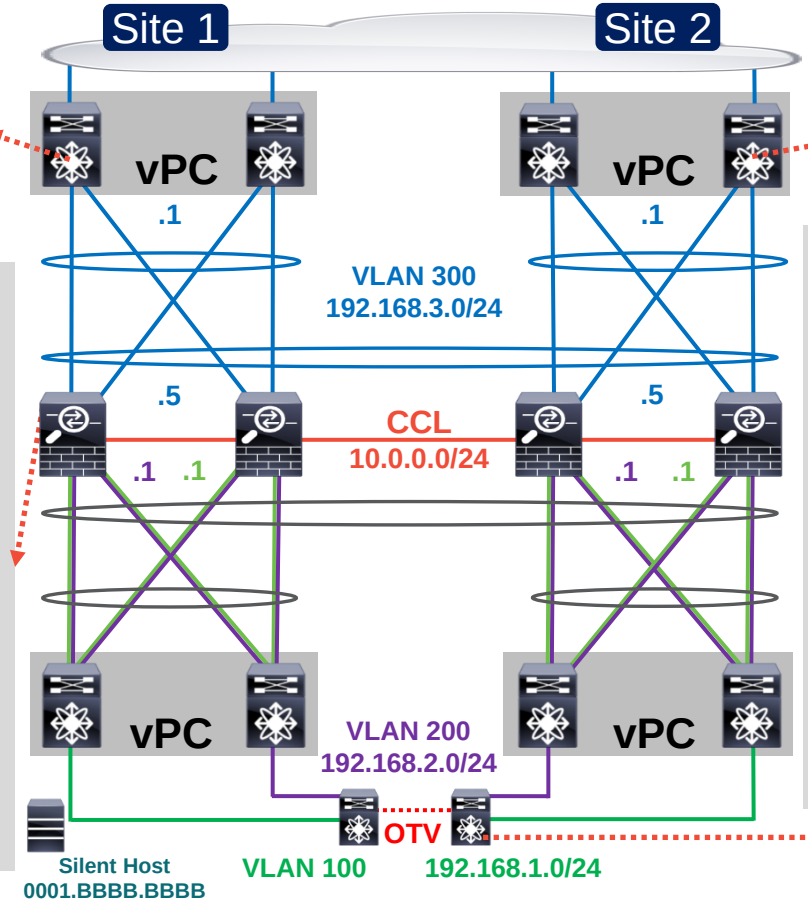


```
interface Vlan300
 ip address 192.168.3.3/24
hsrp 10
  ip 192.168.3.1
ip route 192.168.1.0/24 192.168.3.5
ip route 192.168.2.0/24 192.168.3.5
```

```
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 GMAC_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
```

EW Routed Cluster with Upstream BGP

```
route-map NHOP permit 10
 set ip next-hop 10.1.1.1
router bgp 65000
 address-family ipv4 unicast
  neighbor 10.1.1.2 route-map NHOP in
  neighbor 10.1.1.3 route-map NHOP in
```

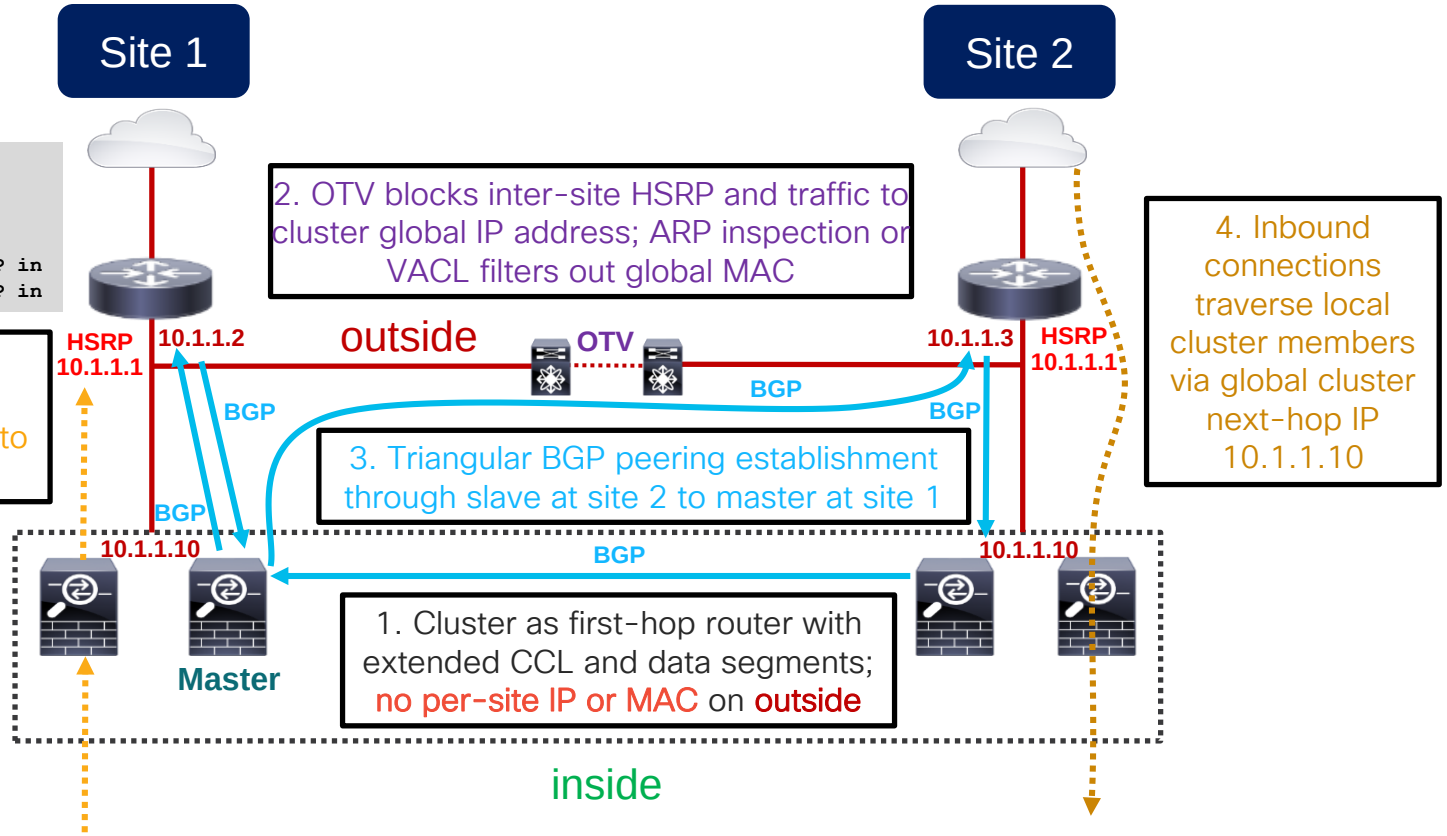

Inter DC Cluster with ACI Anycast Services

- Routed ASA or FTD as first-hop gateway or PBR node in ACI Multipod
  - Split Spanned Etherchannel insertion with each pod as a separate vPC
- Cluster global interface IP/MAC are configured as Anycast gateways
  - No need for per-site IP/MAC addresses or FTD FlexConfig
  - ACI always directs outgoing traffic to closest cluster member group in local pod
  - Automatic switchover to next closest cluster group with no manual filters on failure

ACI Anycast Service IP:
192.168.1.1, 10.1.1.1

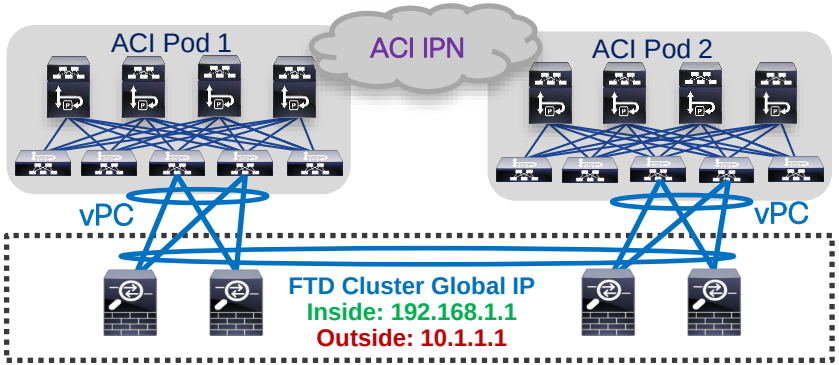

Workloads in Pod 2 direct traffic toward 192.168.1.1 and 10.1.1.1 through local cluster members only; switch to Pod 1 cluster members if no local ones are available

Director Localization and Site Redundancy

- Flow Director selection logic is not site-aware by default
  - A flow owned at one site may select Flow Director at a different site
  - Excessive inter-site traffic on CCL for director lookups is expensive
- Director Localization can be enabled to create two Directors

```
asa(cfg-cluster)# site-id 1
asa(cfg-cluster)# director-localization
```

  - Local Director is at the same site as Flow Owner, primary lookup path

```
TCP outside 85.2.2.123:22 inside 85.2.1.122:58772, idle 0:00:07, bytes 0, flags yl
```

  - Global Director is at a different site from Flow Owner, backup lookup path
  - Lookups for NAT/PAT, IP fragments, or SCTP inspected flows are not localized
- Site Redundancy adds a Director at remote site in ASA 9.9 and FTD 6.2.3

```
asa(cfg-cluster)# site-redundancy
```

Closing Remarks

Clustering Best Practices

- Use a validated switch or verify documented requirements
- Leverage LACP Etherchannel for CCL and dual-connect to VSS/vPC
  - Match the data forwarding capacity of each member
  - Set CCL MTU to 100 bytes above all data interfaces and no less than 1400 bytes
- 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
  - Use firewall-on-a-stick in Spanned Etherchannel mode for best load distribution
  - Minimize centralized features and NAT/PAT

Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

cisco Live!

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

Thank you

cisco Live!

You make **possible**