



*TOMORROW  
starts here.*

Cisco *live!*



# Advanced - ASA Clustering Deep Dive

BRKSEC-3032

Andrew Ossipov

Technical Marketing Engineer

#clmel

Cisco *live!*

# Your Speaker

Andrew Ossipov

[aeo@cisco.com](mailto:aeo@cisco.com)

Technical Marketing Engineer

8 years in Cisco TAC

18+ years in Networking



# Agenda

- Clustering Overview
- Unit Roles and Functions
- Packet Flow
- Control and Data Interfaces
- Configuring Clustering
- Advanced Deployment Scenarios
- Closing Remarks



A long-exposure photograph of a city street at night. The foreground is dominated by vibrant, multi-colored light trails from moving vehicles, creating a sense of motion and energy. In the background, a modern pedestrian bridge spans across the street, illuminated with blue lights. Tall buildings with lit windows and colorful architectural lighting (including red and blue) form the city skyline. The overall scene is a dynamic and visually rich urban environment.

# 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
  - Centralised 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
  - Centralised 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 8 ASA5580/5585-X in **ASA 9.0** and **9.1**; up to 16 ASA5585-X in **ASA 9.2(1)+**
  - Up to 2 ASA5500-X in **ASA 9.1(4)+**
  - SSP 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
- Limited switch chassis support for control and data interfaces
  - Catalyst 6500 with Sup32, Sup720, or Sup720-1GE and Nexus 7000 in **ASA 9.0+**
  - Catalyst 3750-X and Nexus 5000 in **ASA 9.1(4)+**
  - Catalyst 6500 and 6800 with Sup2T in **ASA 9.1(5)+**
  - Nexus 9300 and 9500 in **ASA 9.2(1)+**
  - Nexus 6000 in **ASA 9.3(2)+**



# Unsupported Features

- Remote Access VPN
  - SSL VPN, Clientless SSL VPN, and IPsec
- DHCP Functionality
  - DHCP client, DHCPD server, DHCP Proxy
  - DHCP Relay until **ASA 9.4(1)** in April 2015
- Advanced Application Inspection and Redirection
  - SIP until **ASA 9.4(1)** in April 2015
  - CTIQBE, WAAS, MGCP, MMP, RTSP, Scansafe, Skinny, H.323, GTP engines
  - Botnet Traffic Filter and WCCP
  - Next-Generation Firewall Services with CX
- Unified Communication Security
  - Phone Proxy, Intercompany Media Engine, and other TLS Proxy derivatives

# 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

# Centralised Features

- Not all features are distributed, some are **centralised**
  - Control and management connections
  - 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

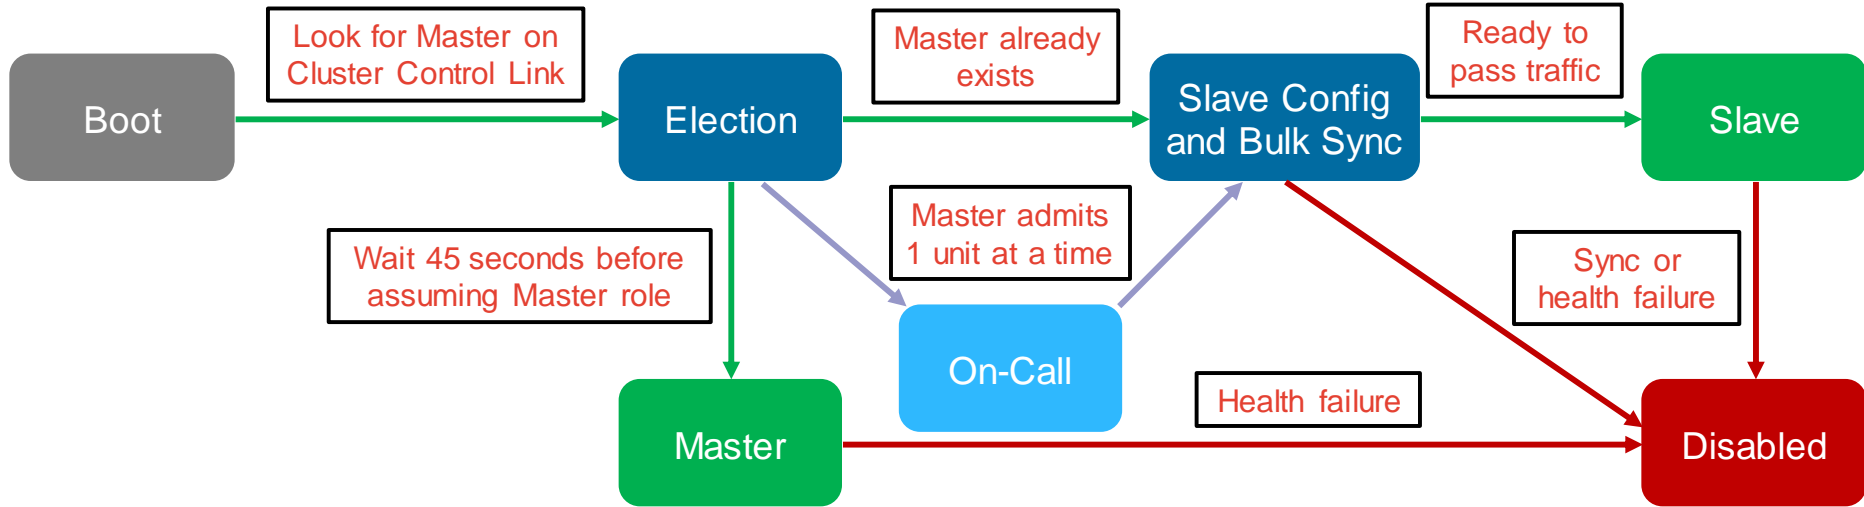


# Unit Roles and Functions

# Master and Slaves

- One cluster member is elected as the **Master**; other 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 centralised 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 centralised connections must re-establish upon Master failure
  - Disable or reload Master to transition the role; **do not** use **cluster master** command

# State Transition



```
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
```

Cisco live!

# 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 are re-hosted only on unit departure

# 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 Randomisation
- **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
```

- When Flow Director and Owner are the same, another unit has Backup Stub Flow

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



# Flow Forwarder

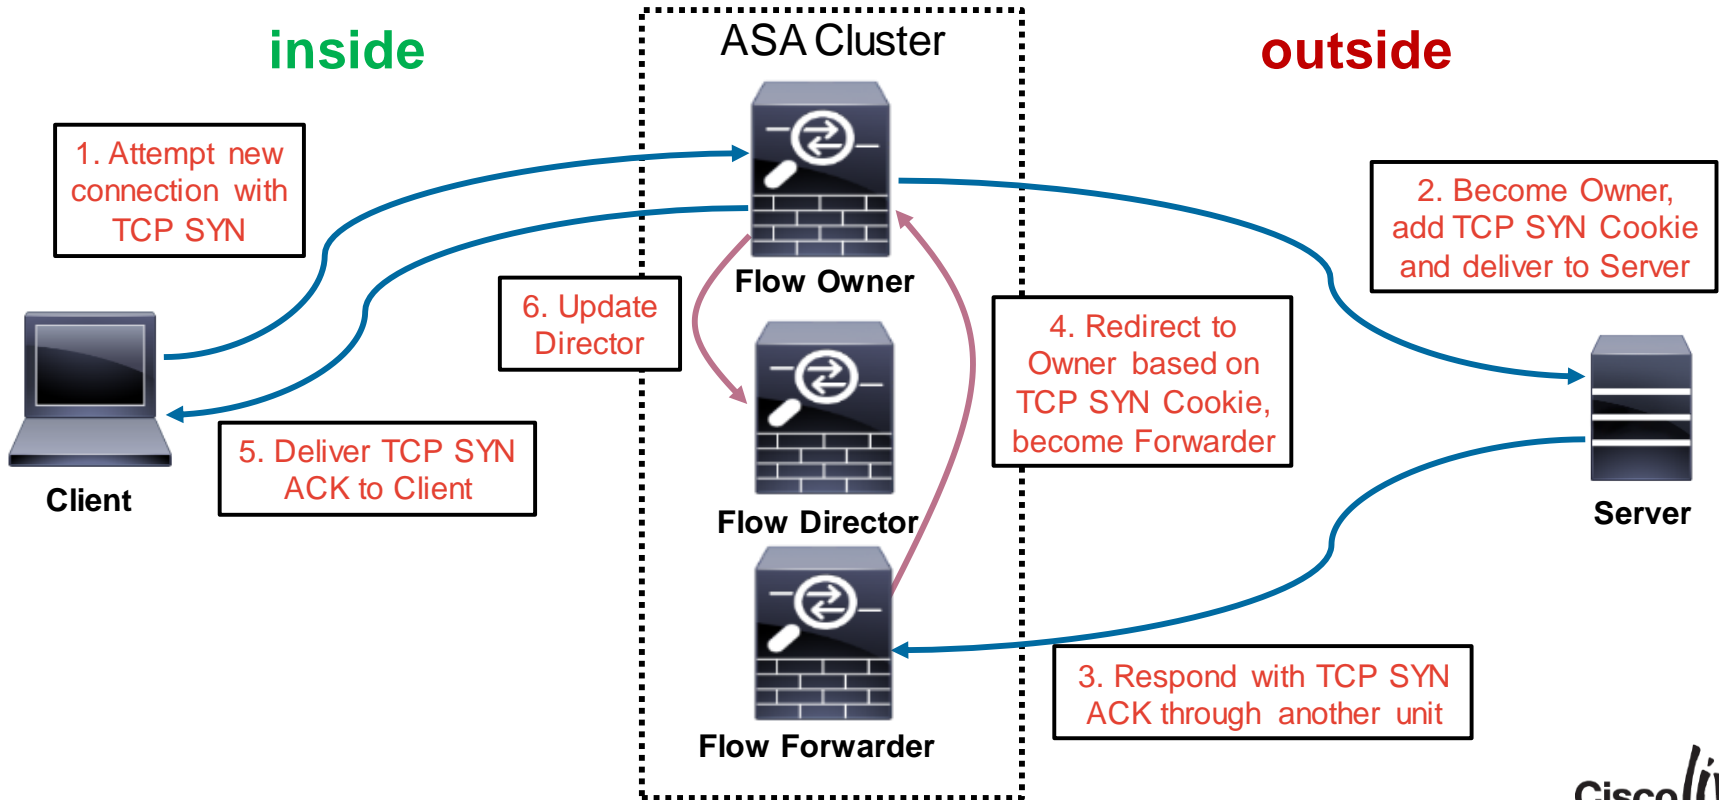
- External stateless load-balancing mechanism does not guarantee symmetry
  - Only TCP SYN packets can reliably indicate that the connection is new
- Cluster member receiving a non-TCP-SYN packet must query Flow Director
  - No existing connection → Drop if TCP, become Flow Owner if UDP
  - Existing 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 centralised 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)
```

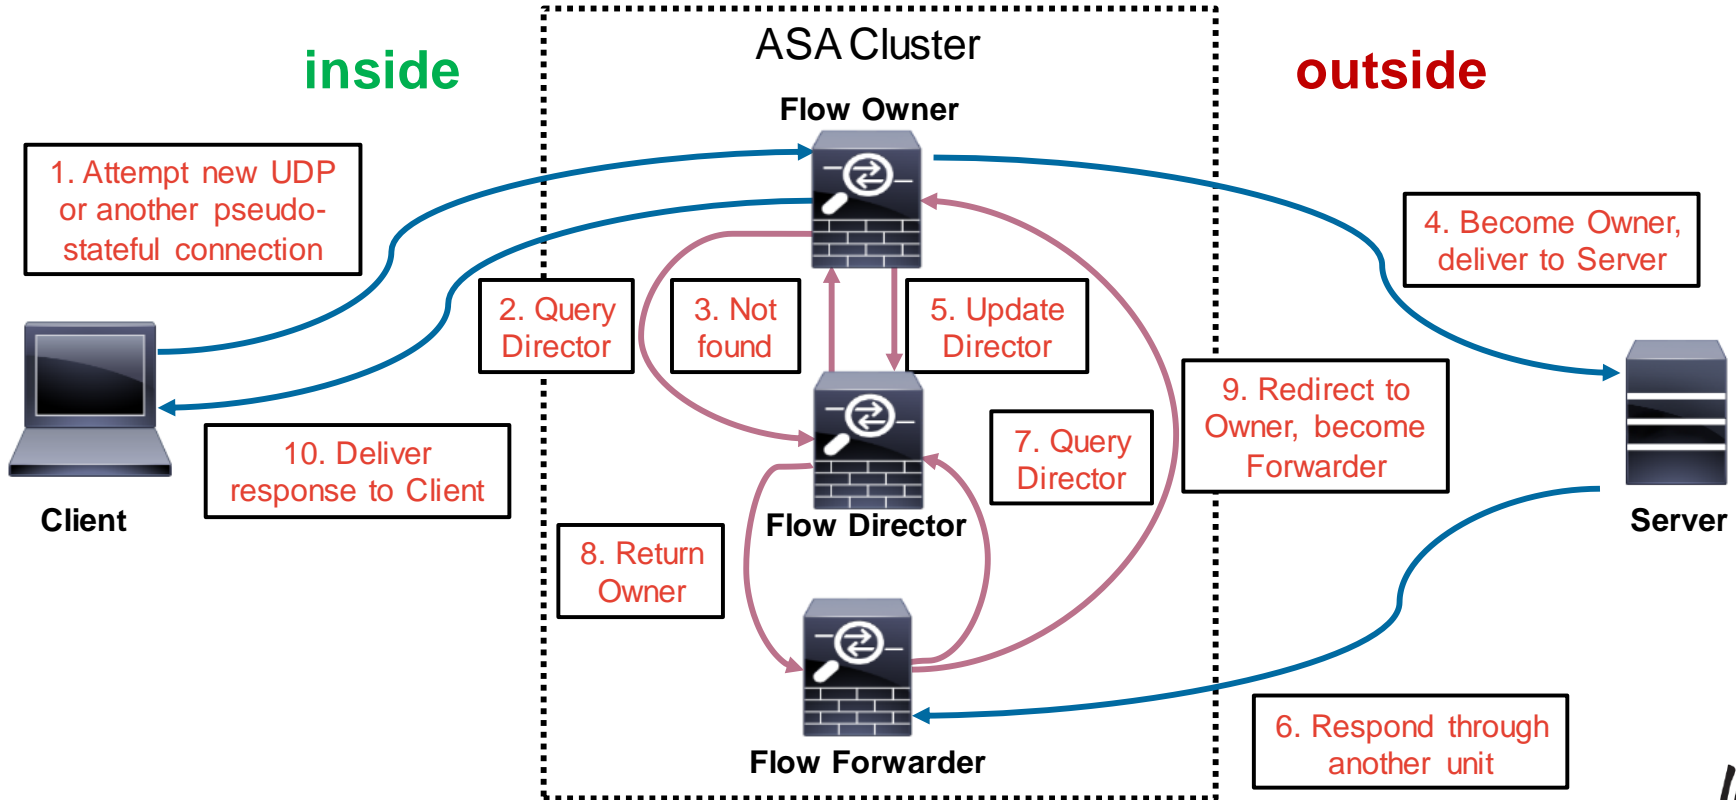


# Packet Flow

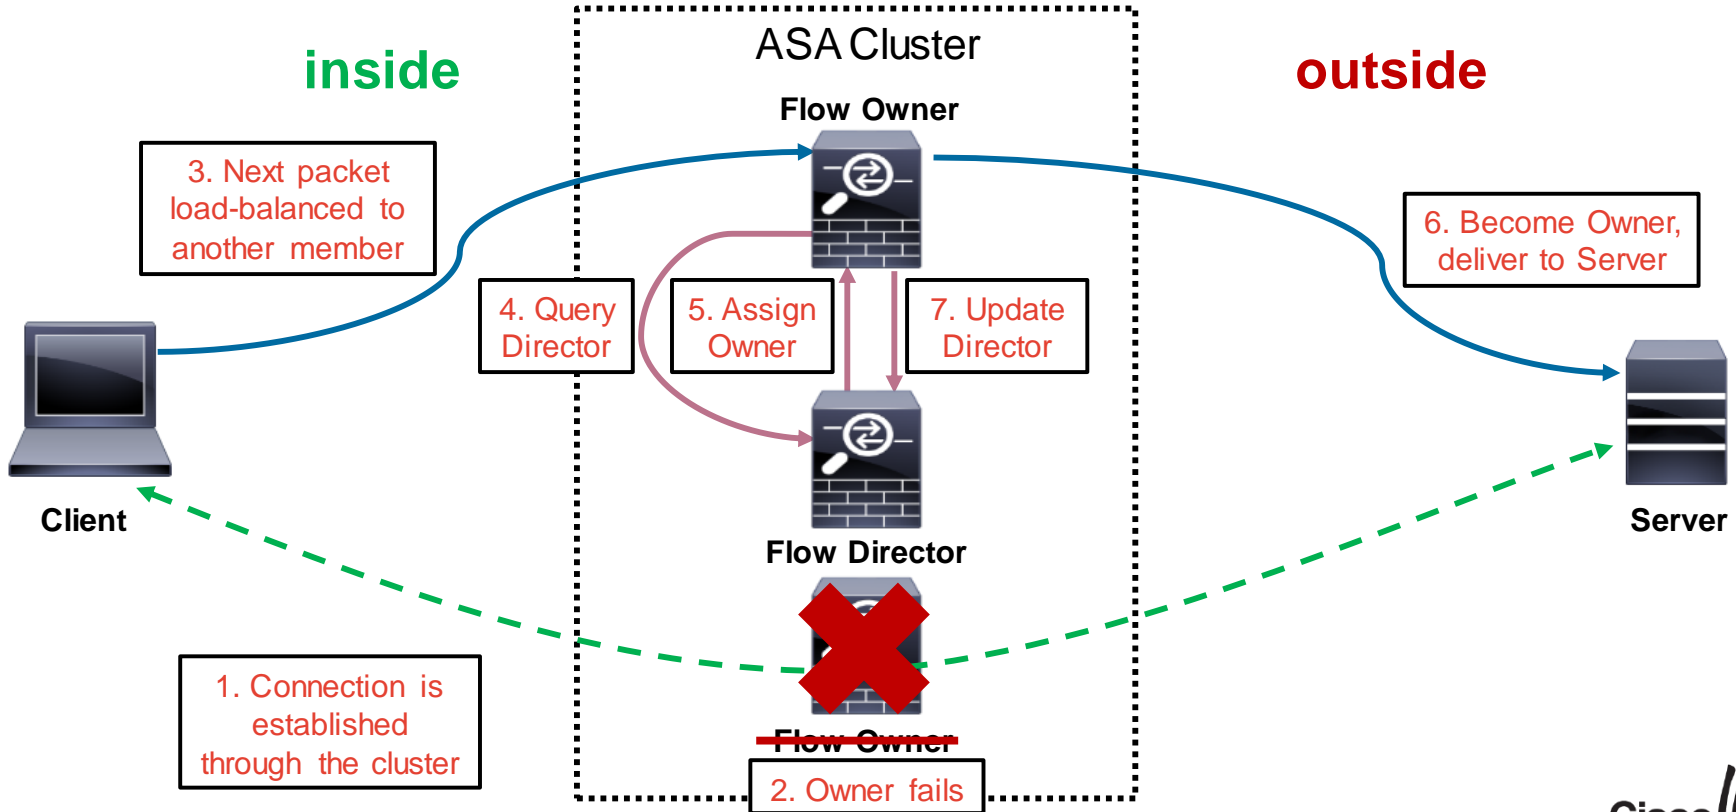
# New TCP Connection



# New UDP-Like Connection



# Owner Failure



# Application Inspection

- Centralised
  - 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
- Semi Distributed with **ASA 9.4(1)+** in April 2015
  - Control connections are processed independently by all units
  - Data connections are redirected to the associated control connections' Owners
  - Example: SIP

# 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 connection failures
  - Per-session xlates are local to the Owner with an Xlate backup
  - Some connections require non-per-session xlates which are centralised 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
  - Static, one-to-one dynamic, and non-per-session NAT **does not** scale in clustering

Cisco *live!*






# Control and Data Interfaces

# Cluster Control Link (CCL)

- Carries all data and control communication between cluster members
  - Master discovery, initial negotiation, keepalives, interface status updates
  - Configuration synchronisation from Master to Slaves
  - Centralised resource allocation (such as PAT/NAT, pinholes)
  - Flow Director updates and Owner queries
  - Centralised and asymmetric traffic redirection from Forwarders to Owners
- Must use same dedicated interfaces on each member
  - Separate physical interface(s), no sharing or VLAN subinterfaces
  - 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

# 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 MTU 100 bytes above largest data interface 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 of all members bundle into a single Etherchannel


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

- Peer switch sees the cluster as a single logical entity
  - 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 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-Channell
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)

- Recommended way to bundle **data** interfaces into a Spanned Etherchannel
  - Up to 8 active and 8 standby links in **9.0/9.1** with dynamic port priorities in vPC/VSS

```
asa(config)# interface Port-Channel 1
asa(config-if)# port-channel span-cluster vss-load-balance
asa(config-if)# 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)+**

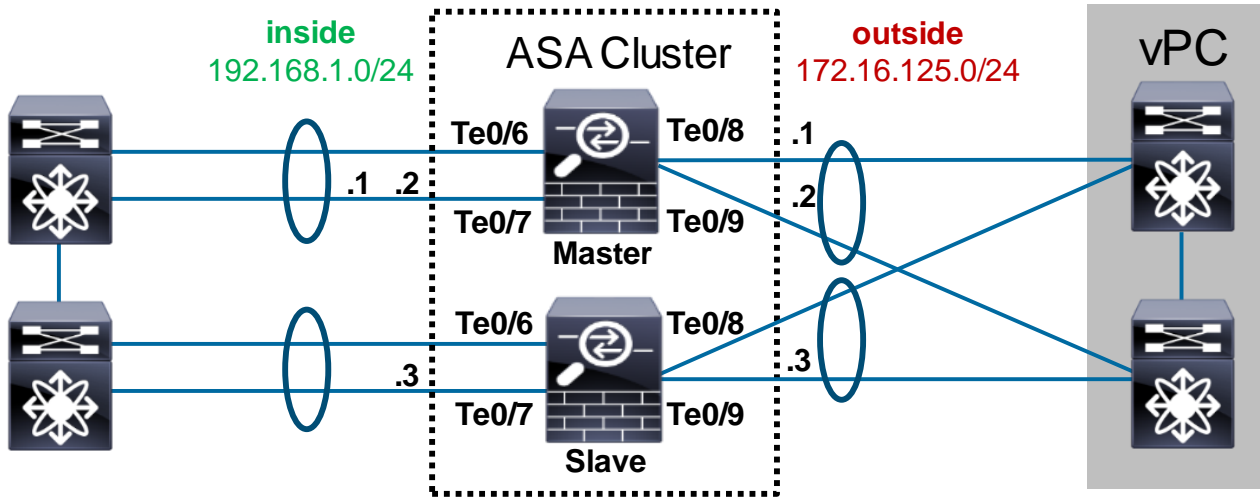
```
asa(config)# cluster group DC_ASA
asa(cfg-cluster)# clacp static-port-priority
```

- Use static LACP port priorities to avoid problems with unsupported switches
  - Always configure virtual MAC addresses for each Etherchannel to avoid instability
  - **Disable** LACP Graceful Convergence and Adaptive Hash on adjacent NX-OS
- cLACP **assumes** each Spanned Etherchannel connects to one logical switch
    - LACP actor IDs between member ports are not strictly enforced, allowing creativity

Cisco *live!*

# Individual Interface Mode

- **Routed** firewalls only
- 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)** is very static by definition
  - Use static route maps on adjacent routers to fan flows across all cluster members
  - 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
- 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 synchronised 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 **centralised** 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.

# Dynamic Routing Convergence Optimisation

- 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 neighbour 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 neighbour

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

CiscoLive!

# 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
```

# 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
- 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
- Each unit monitors the health of its interfaces only locally
  - 500ms reaction time on standalone interface failures, 9 seconds for LACP
- Disable health check during changes and use selective interface monitoring

```
asa/master# cluster group sjfw
asa/master(cfg-cluster)# no health-check
asa/master(cfg-cluster)# health-check holdtime 1
asa/master(cfg-cluster)# no health-check monitor-interface Management0/0
```

Keepalive is always 1/3 of the configured holdtime

ASA 9.4(1) feature in April 2015



# Configuring Clustering

# 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
  - A warning prompt will indicate conflicts and available options
  - 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

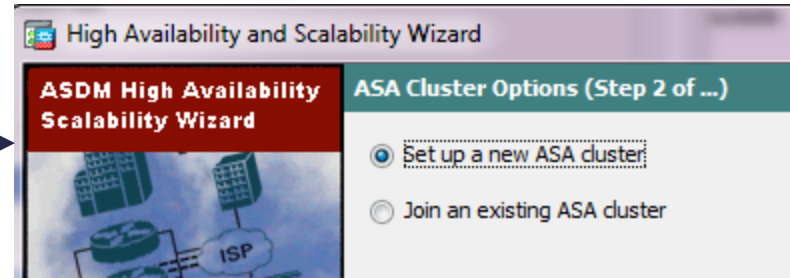
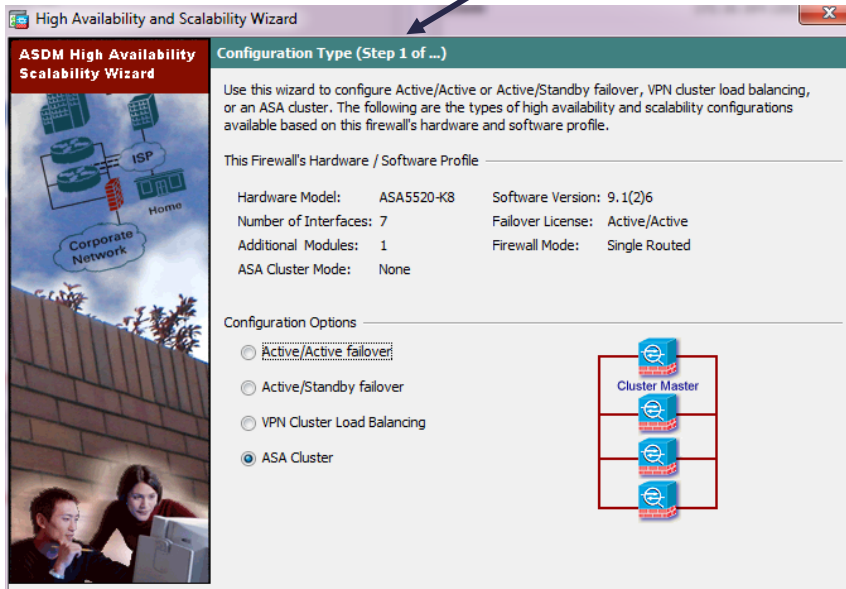

```
ip local pool CLUSTER_MANAGEMENT 172.16.162.243-172.16.162.250
!
interface Management0/0
  description management interface
  management-only
  nameif mgmt
  security-level 0
  ip address 172.16.162.242 255.255.255.224 cluster-pool CLUSTER_MANAGEMENT
!
route mgmt 0.0.0.0 0.0.0.0 172.16.162.225 1
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



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

**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

```
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
```

**Automatic: cLACP**  
system MAC

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

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

**All Members: Lower**  
numerical priority wins  
Master election

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

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

**Master: Set CCL MTU**  
100 bytes above all data  
interfaces

# 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
```

Spanned Etherchannel bundles ports across entire cluster

Single virtual IP for all members

Up to 32 ports with cLACP in 9.2(1)

Virtual MAC is required for Etherchannel stability

## 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

Virtual IP is owned by Master for management only

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

Cisco *live!*

# 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-channel11 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
```

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 utilisation 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 identify reporting members for connection events

Cisco *live!*



# Advanced Deployment Scenarios

# Inter Data Centre (DC) Clustering

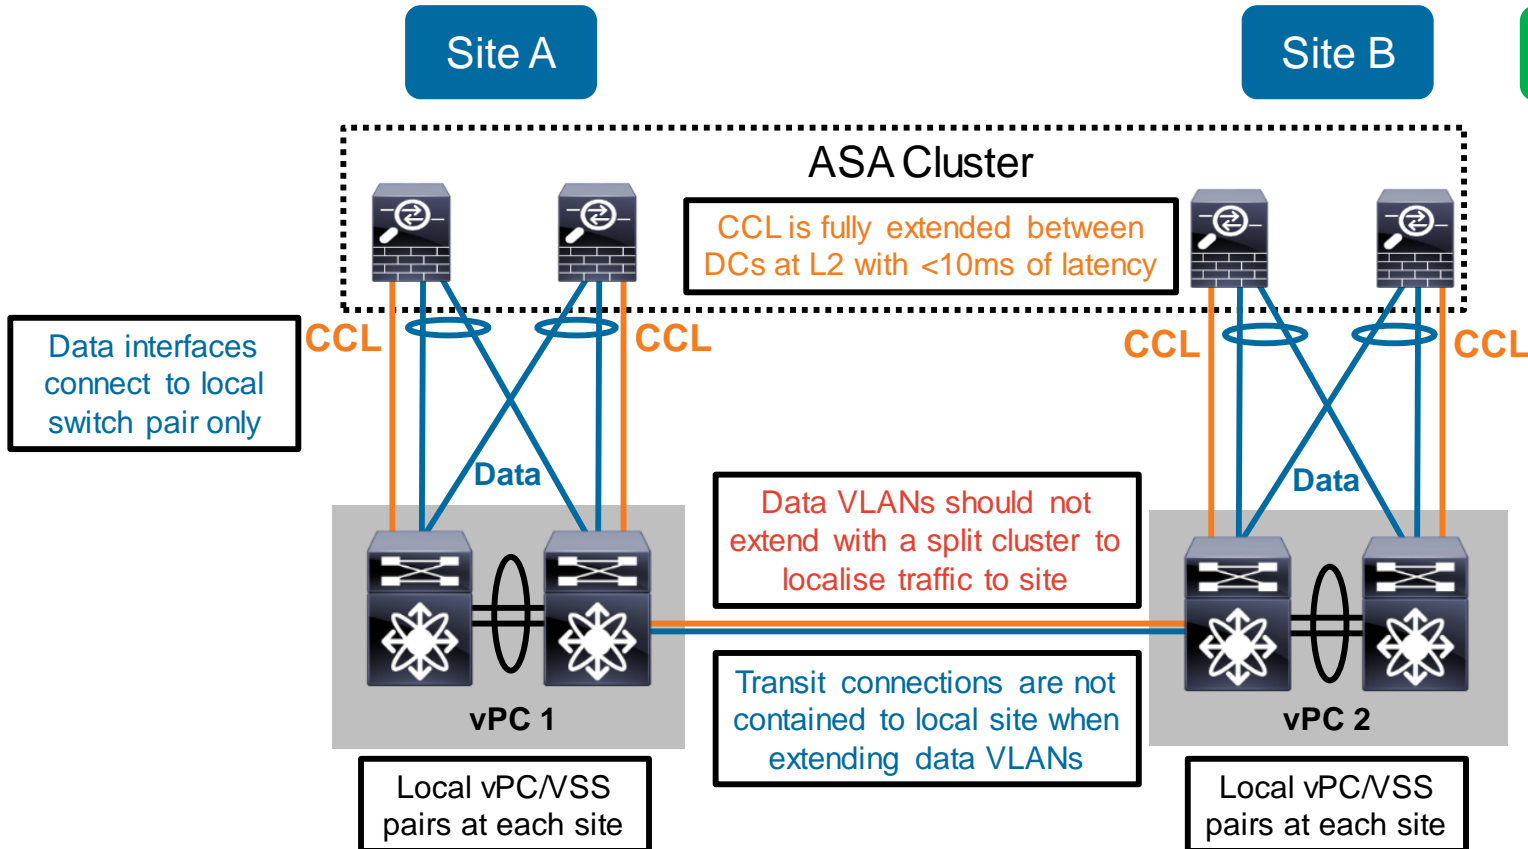
- Clustering **assumes rather than requires** data interface adjacency at Layer 2
- Geographically separated clusters supported in **ASA 9.1(4)+**
  - “Dark Media” CCL with up to 10ms of one-way latency
  - No tolerance for packet re-ordering or loss
  - Routed firewall in Individual interface mode **only**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

# Split or Single Individual Mode Cluster in Inter DC

Site A

Site B

ASA 9.1(4)



# Extended Spanned Etherchannel Cluster in Inter DC

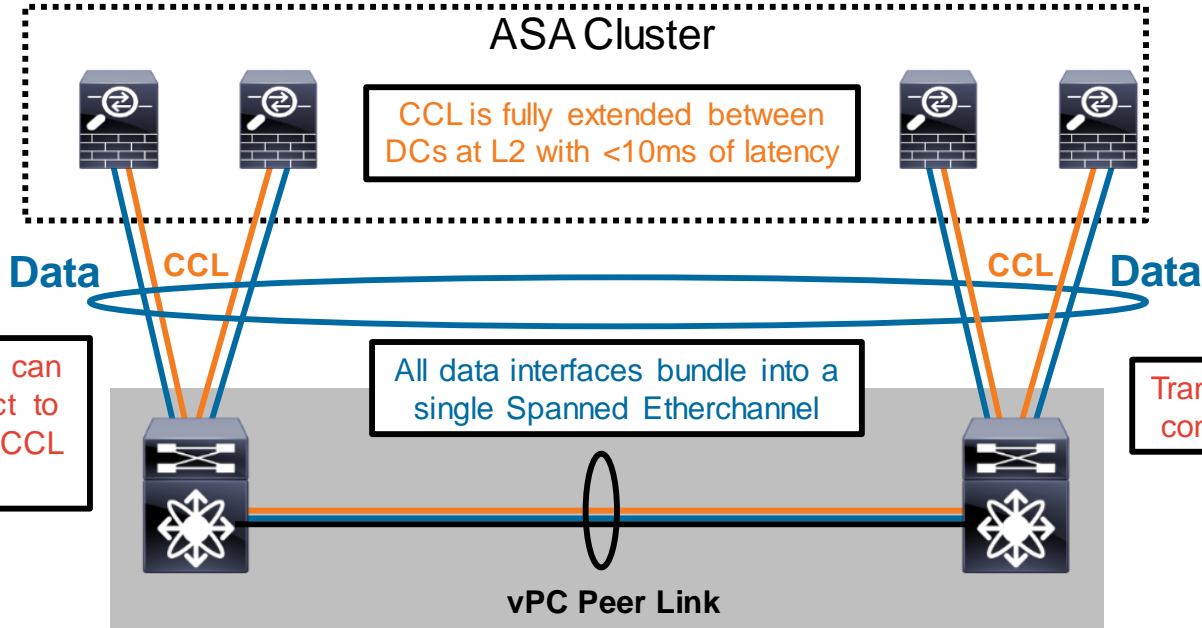
Site A

Site B

ASA 9.2(1)

ASA Cluster

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

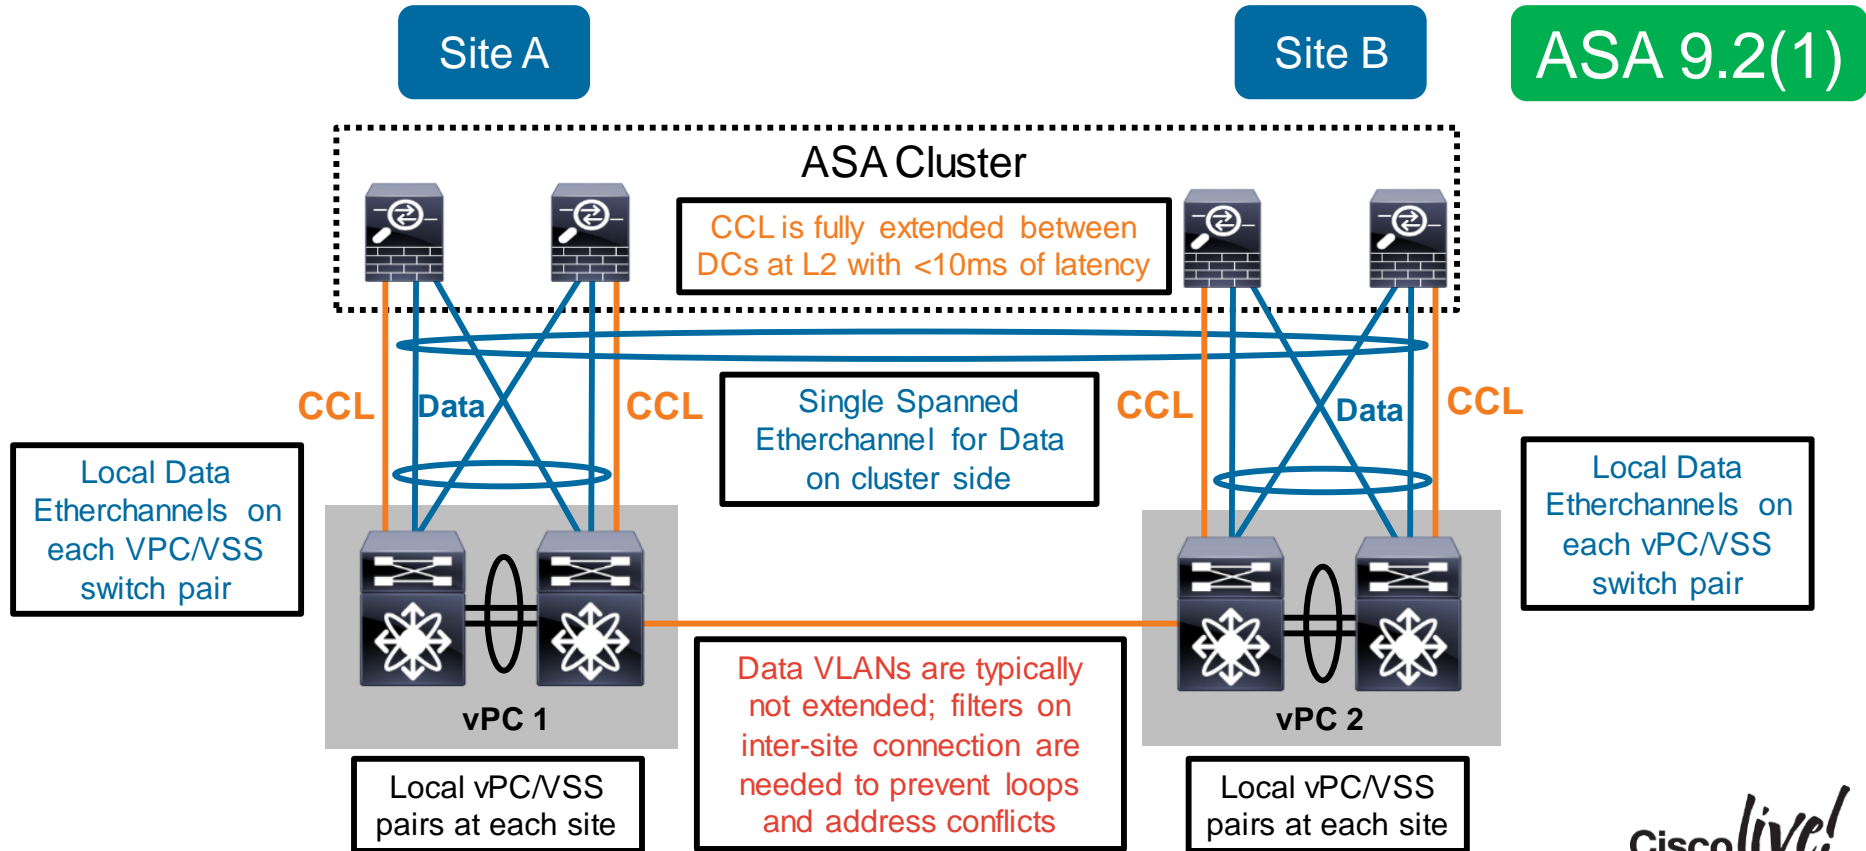


Each cluster member can single- or dual-connect to the VSS/vPC pair for CCL and Data

All data interfaces bundle into a single Spanned Etherchannel

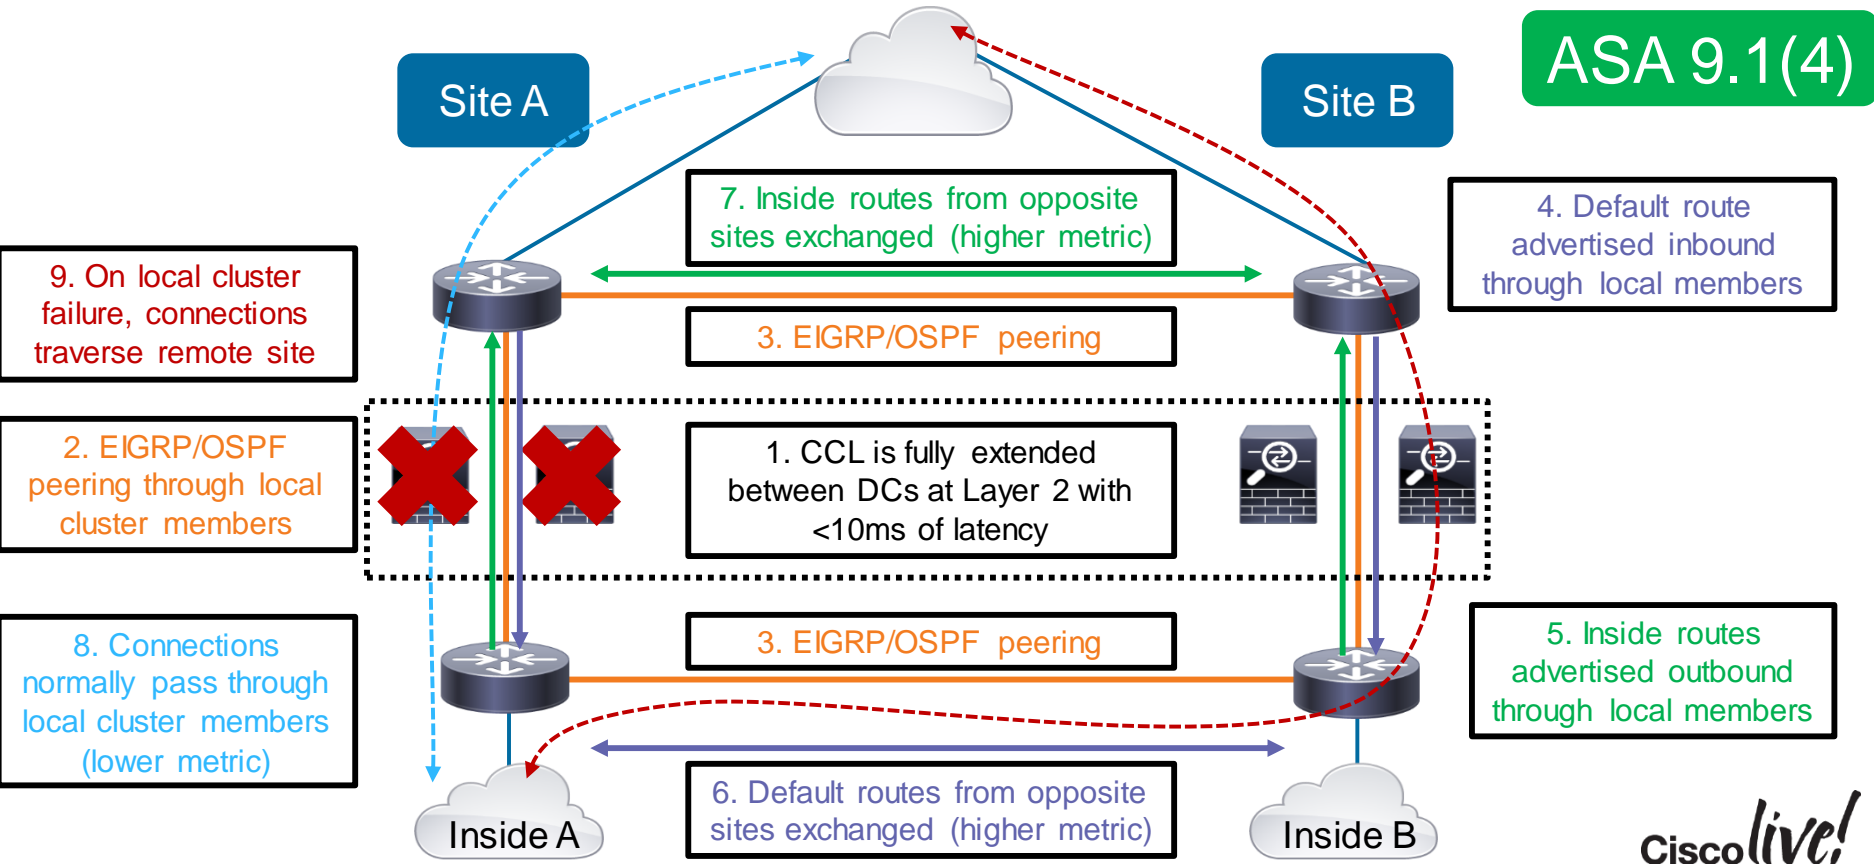
Transit connections are not contained to the local site

# Split Spanned Etherchannel Cluster in Inter DC



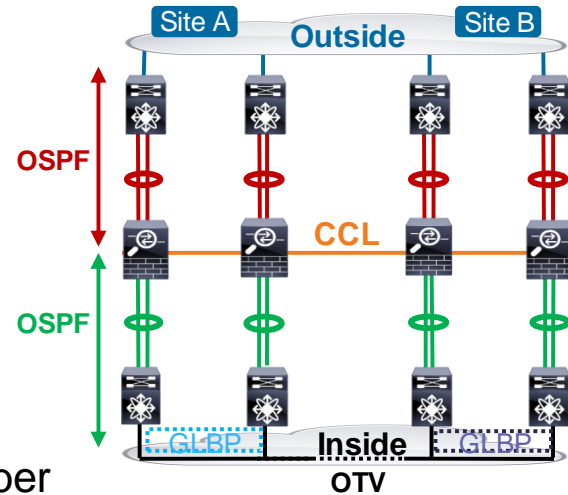
# North-South Inter DC Clustering

ASA 9.1(4)



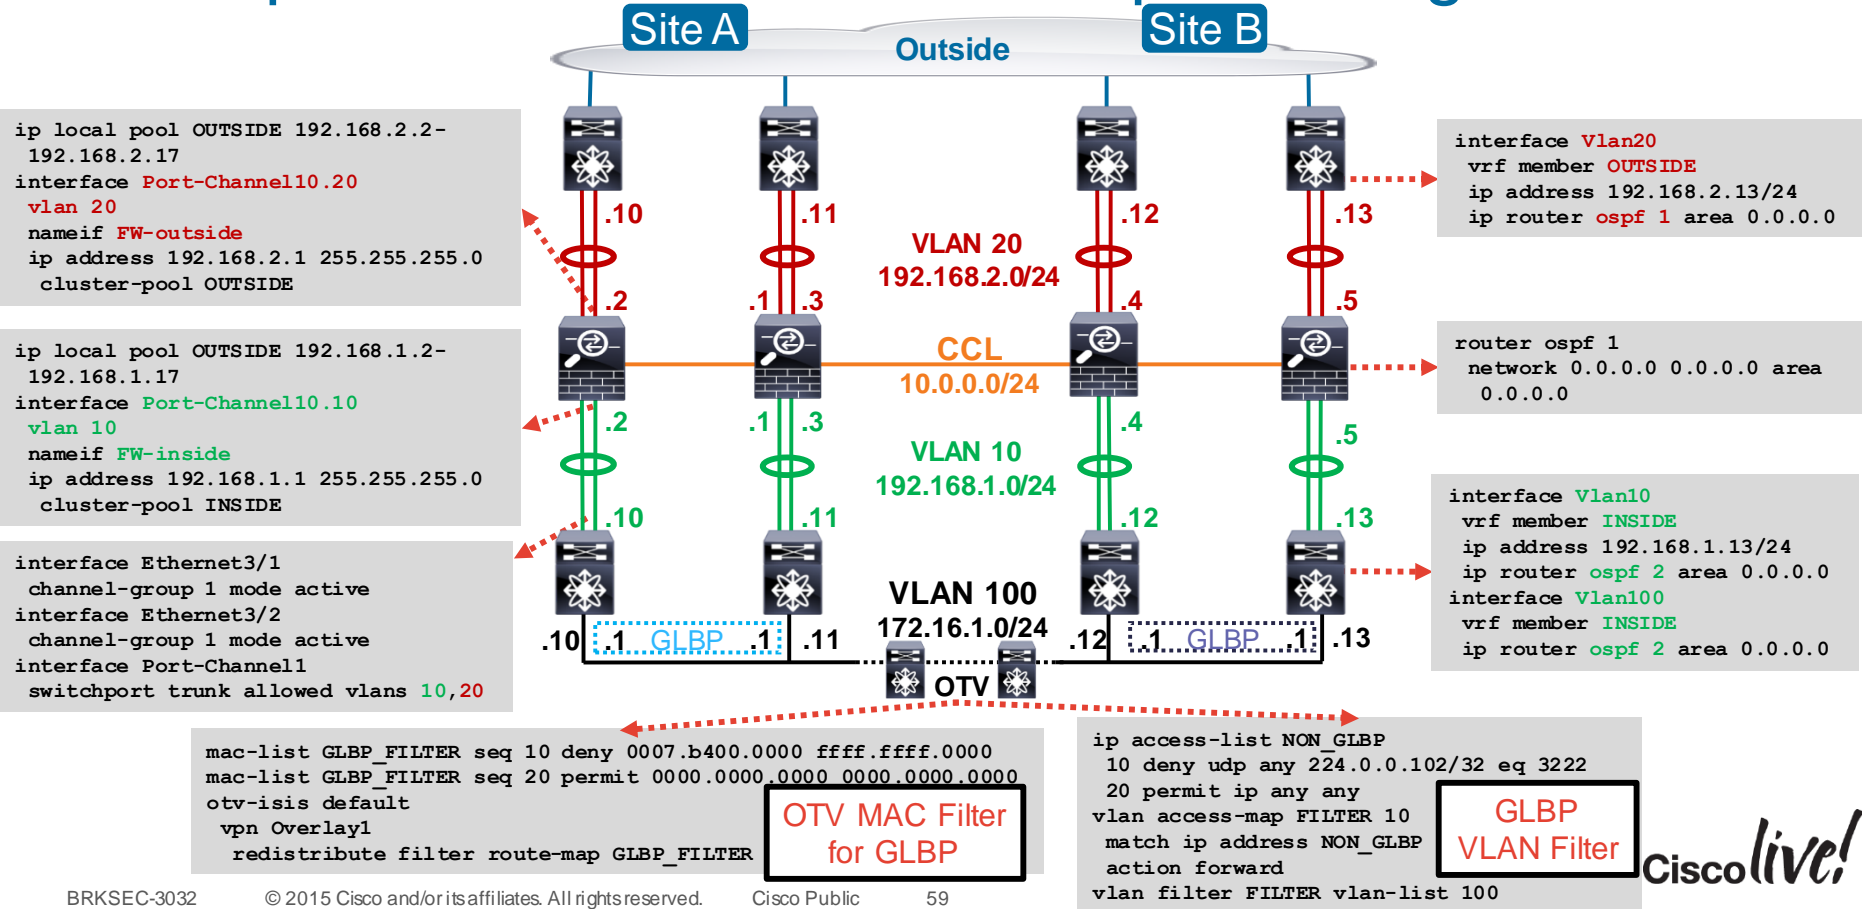
# Example: N-S Split Individual Mode Cluster

- 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 localised 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

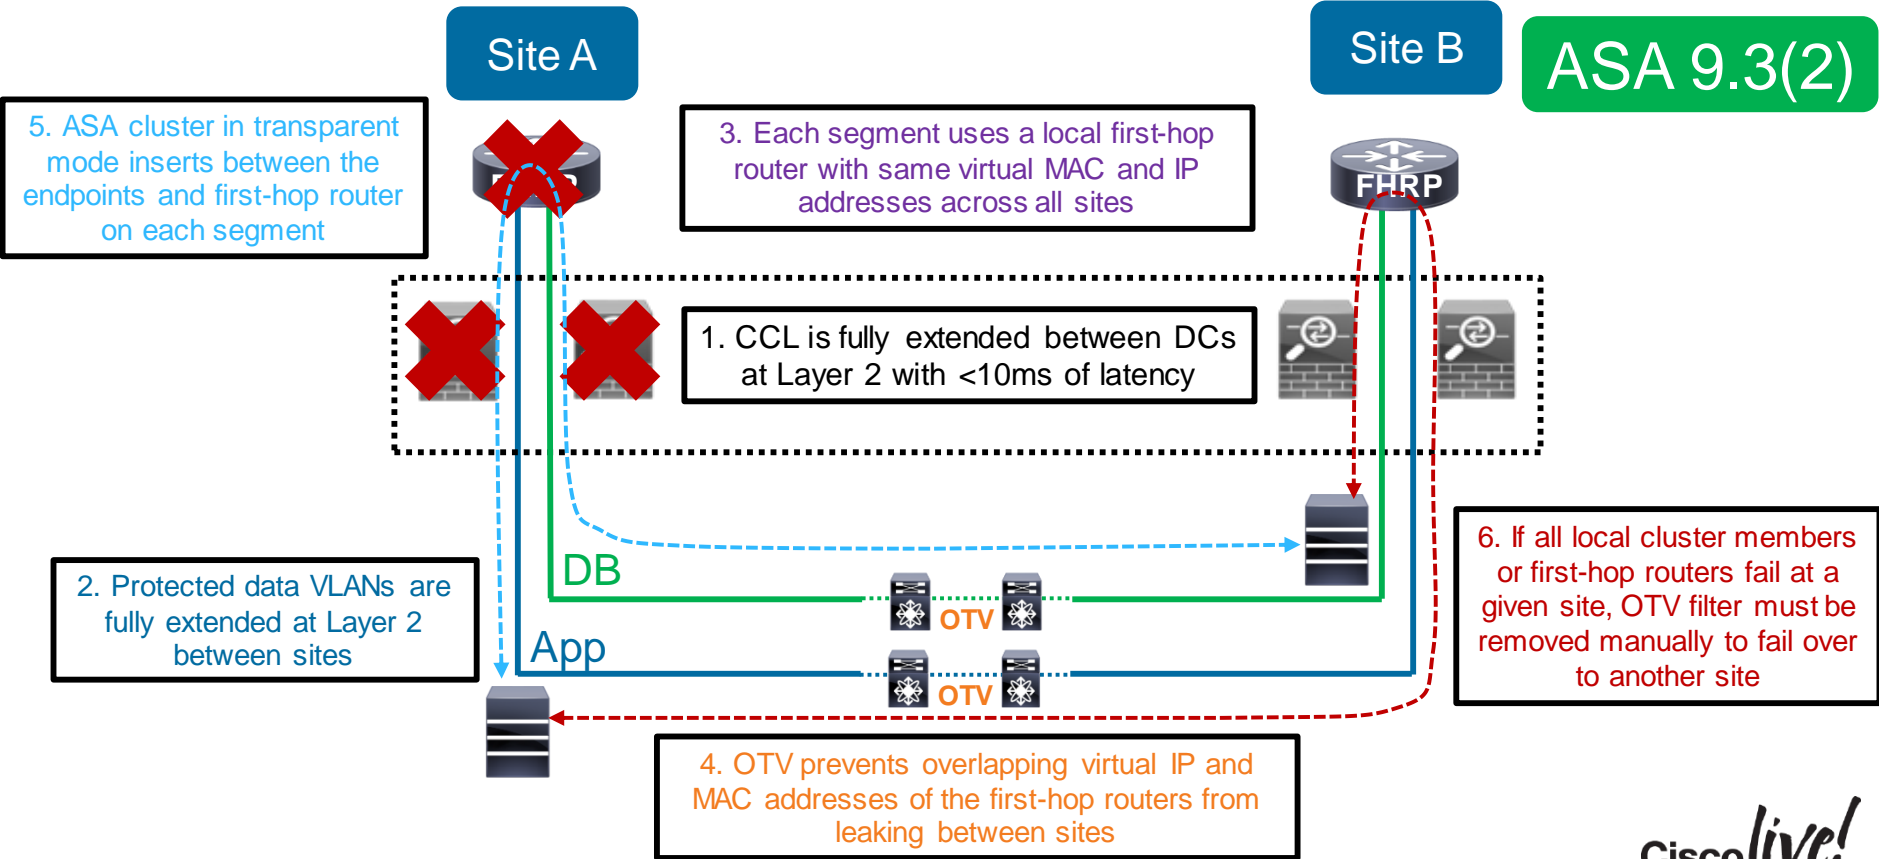




# N-S Split Individual Cluster Sample Configuration




# East-West Inter DC Clustering

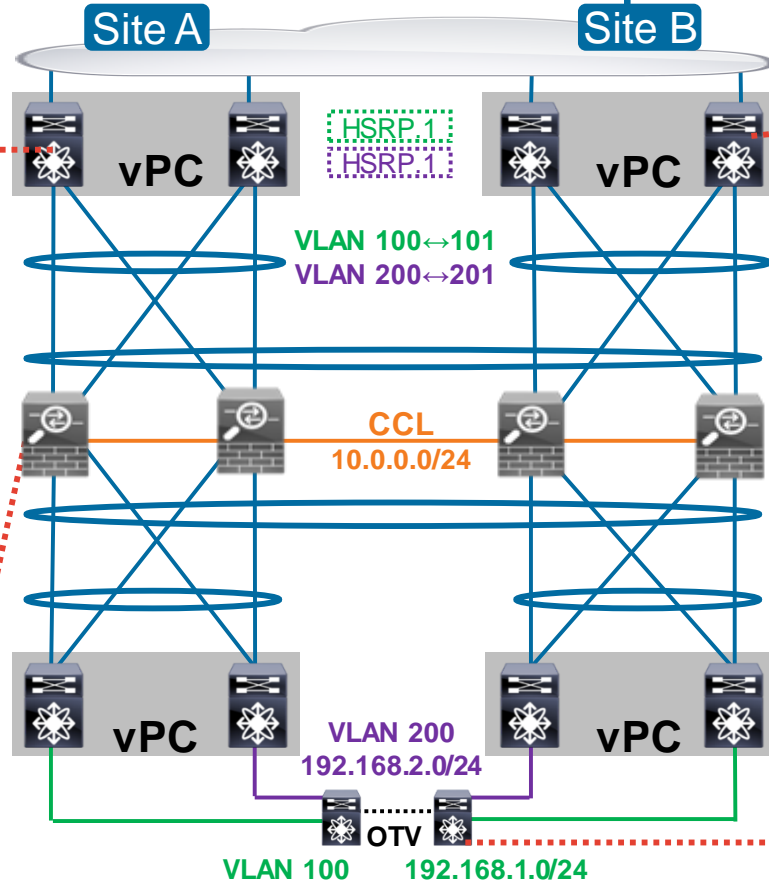


# Example: E-W Split Spanned Etherchannel Cluster

- 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 localised 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
- Traffic symmetry is achievable without NAT
  - Fully localised processing between same-site applications
  - First-hop routers may run LISP for North-South site symmetry



# E-W Split Spanned Cluster Sample 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
```

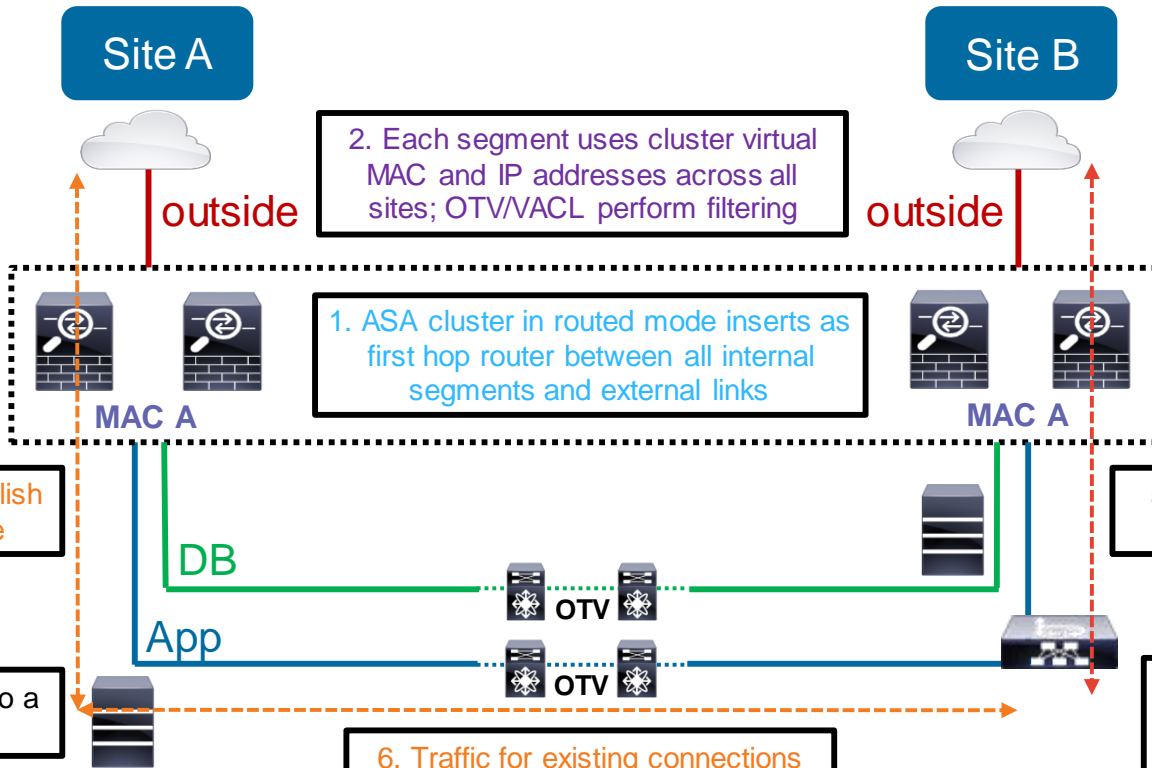
```
interface Port-Channel10
port-channel span-cluster
interface Port-Channel10.100
vlan 100
nameif DB-inside
bridge-group 1
interface Port-Channel10.101
vlan 101
nameif DB-outside
bridge-group 1
interface Port-Channel10.200
vlan 200
nameif App-inside
bridge-group 2
interface Port-Channel10.201
vlan 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.fff.ff00
mac-list HSRP_FILTER seq 20 deny
0000.0c9f.f000 ffff.fff.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
```

# Unsupported East-West Inter DC Scenario

Future



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

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

3. Connections establish locally at each site

5. New connections establish locally through new site

4. VM live-migrates to a different 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



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 Randomisation enabled for SYN Cookies
  - Minimise centralised features and NAT/PAT
  - Use Spanned Etherchannel mode for better load distribution
  - Match Etherchannel hashing algorithms on all connected switches

# The Challenges Come from Every Direction







## Cisco 2015 Annual Security Report

Now available:

[cisco.com/go/asr2015](http://cisco.com/go/asr2015)



Q & A

# Complete Your Online Session Evaluation

**Give us your feedback and receive a Cisco Live 2015 T-Shirt!**

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site  
<http://showcase.genie-connect.com/clmelbourne2015>
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm



**Learn online with Cisco Live!**

Visit us online after the conference for full access to session videos and presentations. [www.CiscoLiveAPAC.com](http://www.CiscoLiveAPAC.com)

**Cisco** *live!*



Thank you.

Cisco *live!*



**CISCO**