Cisco live!

What You Make Possible

TOMORROW starts here.

CISCO

Cisco live!

Converged Access Architecture Overview

BRKARC-2665

TOMORROW starts here.

CISCO

Converged Access Architecture Overview

Diving into the "One Network"

BRKARC-2665 – **Session Overview and Objectives**

Come to this session to learn what Converged Access is – how it operates – and the benefits it provides.

Attendees at this session will gain a **greater understanding** of the design and operation of the exciting new Converged Access solution, be able to **understand** how it fits into the broader Cisco wired and wireless portfolio from both a product and a design perspective, and **recognise** the relevant benefits for their own network environments.

In addition to introducing the terminology and platforms that make up the Cisco Converged Access system, an in-depth review of the new Converged Access solution is provided, including coverage of client association, various roaming modes, high availability, and Quality of Service capabilities.

Converged Access Architecture Overview

Diving into the "One Network"

Your Instructor Today … **Dave Zacks**

I am a **Technical Solutions Architect**, and have been with Cisco for 13+ years.

I work primarily with large, high-performance Enterprise network architectures, designs, and systems. I have over 20 years of experience with designing, implementing, and supporting highly available network systems and solutions that have included many diverse network technologies and capabilities, using multiple different topologies.

I maintain a CCIE certification in Routing and Switching (ten years and counting!)

Dave Zacks
Technical Solutions Architect
dzacks@cisco.com
CCIE 8887

Cisco Converged Access

What I'm Going to Cover …

- System Architecture
- Roaming
- High Availability

Foundational Elements
for the Converged Access Solution

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – **Towards One Policy, One Management, One Network**

Converged Access – **Platform Overviews**

Wired and Wireless – **Deployment Options**

And a "double-click" deeper …

Existing Wireless Deployment – **Architecture Refresher**

The Converged Access Deployment in Detail –

- Components of the Deployment – **Terminology and Building Blocks**
- Converged Access Deployment – **Traffic Flows and Roaming**
- Converged Access Deployment – **High Availability**
- Converged Access Deployment – **Quality of Service**

Summary

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a “double-click” deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

Evolving User Workspace
Megatrends

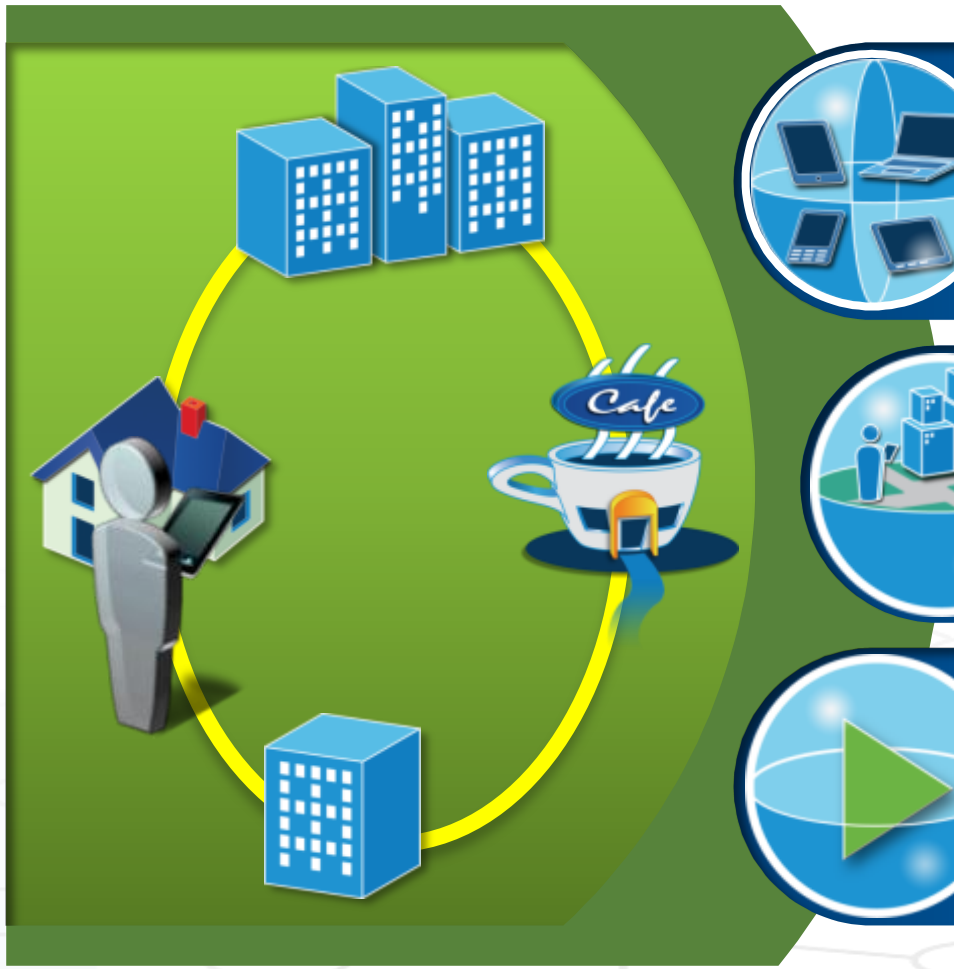

BYOD
- Secure access
- Customised experience
- Guest access

Mobility
- Seamless roaming
- Optimal client performance
- Cloud access/VXI

Video
- Multicast streaming
- Video conferencing
- Reliable performance

IT Requirement

Deliver an Uncompromised User Experience on Any Workspace

Wireless Standards

Past, Present, and Future

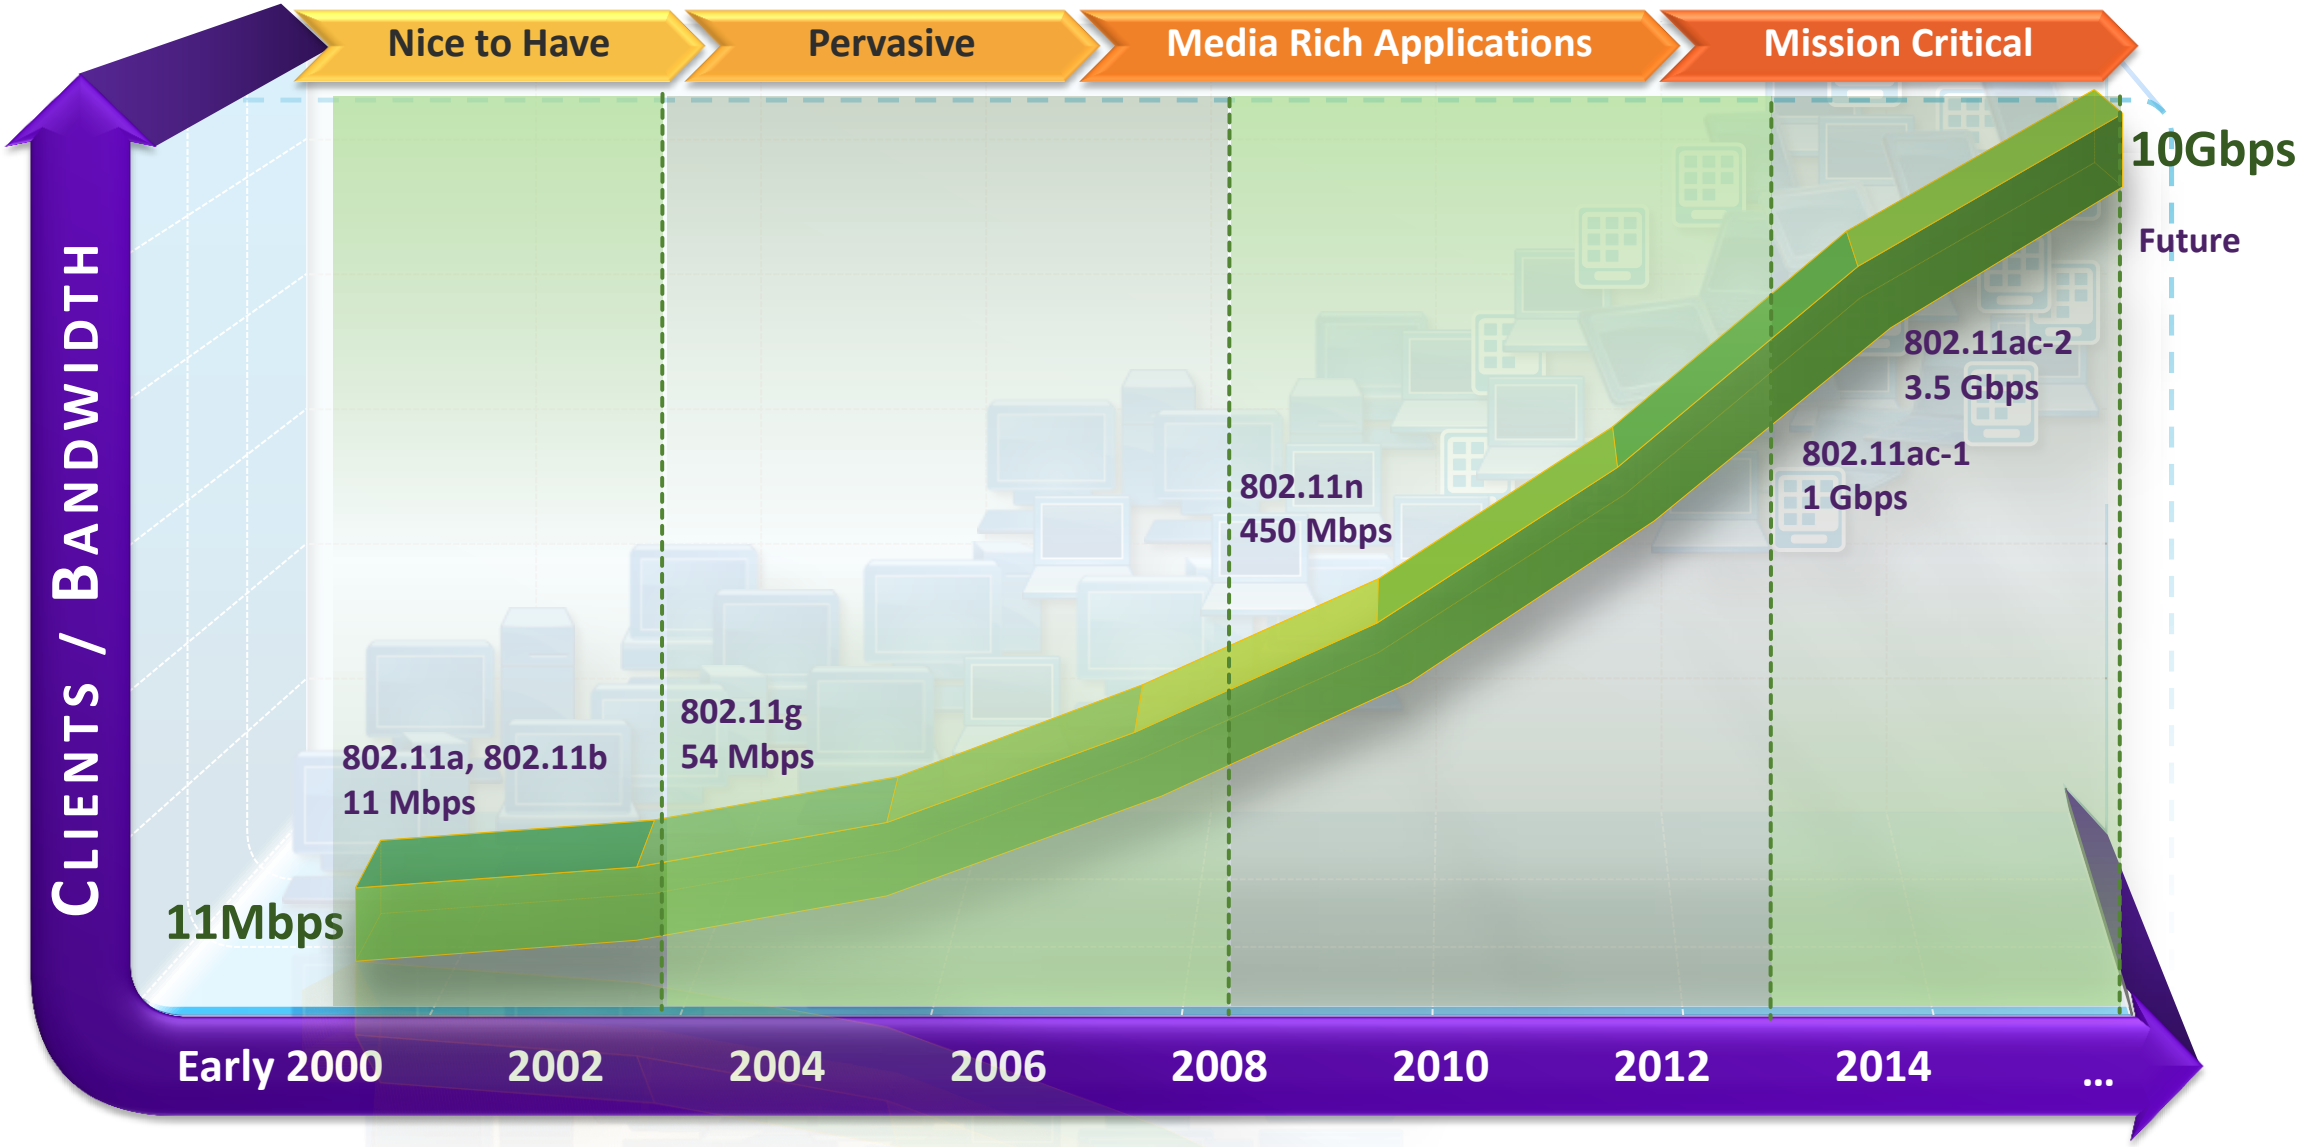

Converged Access

Uncompromised User Experience in Any Workspace

One Network, with Converged Access

A New Deployment Option for Wired / Wireless

IOS Based WLAN Controller

- Consistent IOS and ASIC as Catalyst 3850
- Required to scale beyond 250 AP or 16K client domains

Converged Access Mode

- Integrated wireless controller
- Distributed wired/wireless data plane (CAPWAP termination on switch)

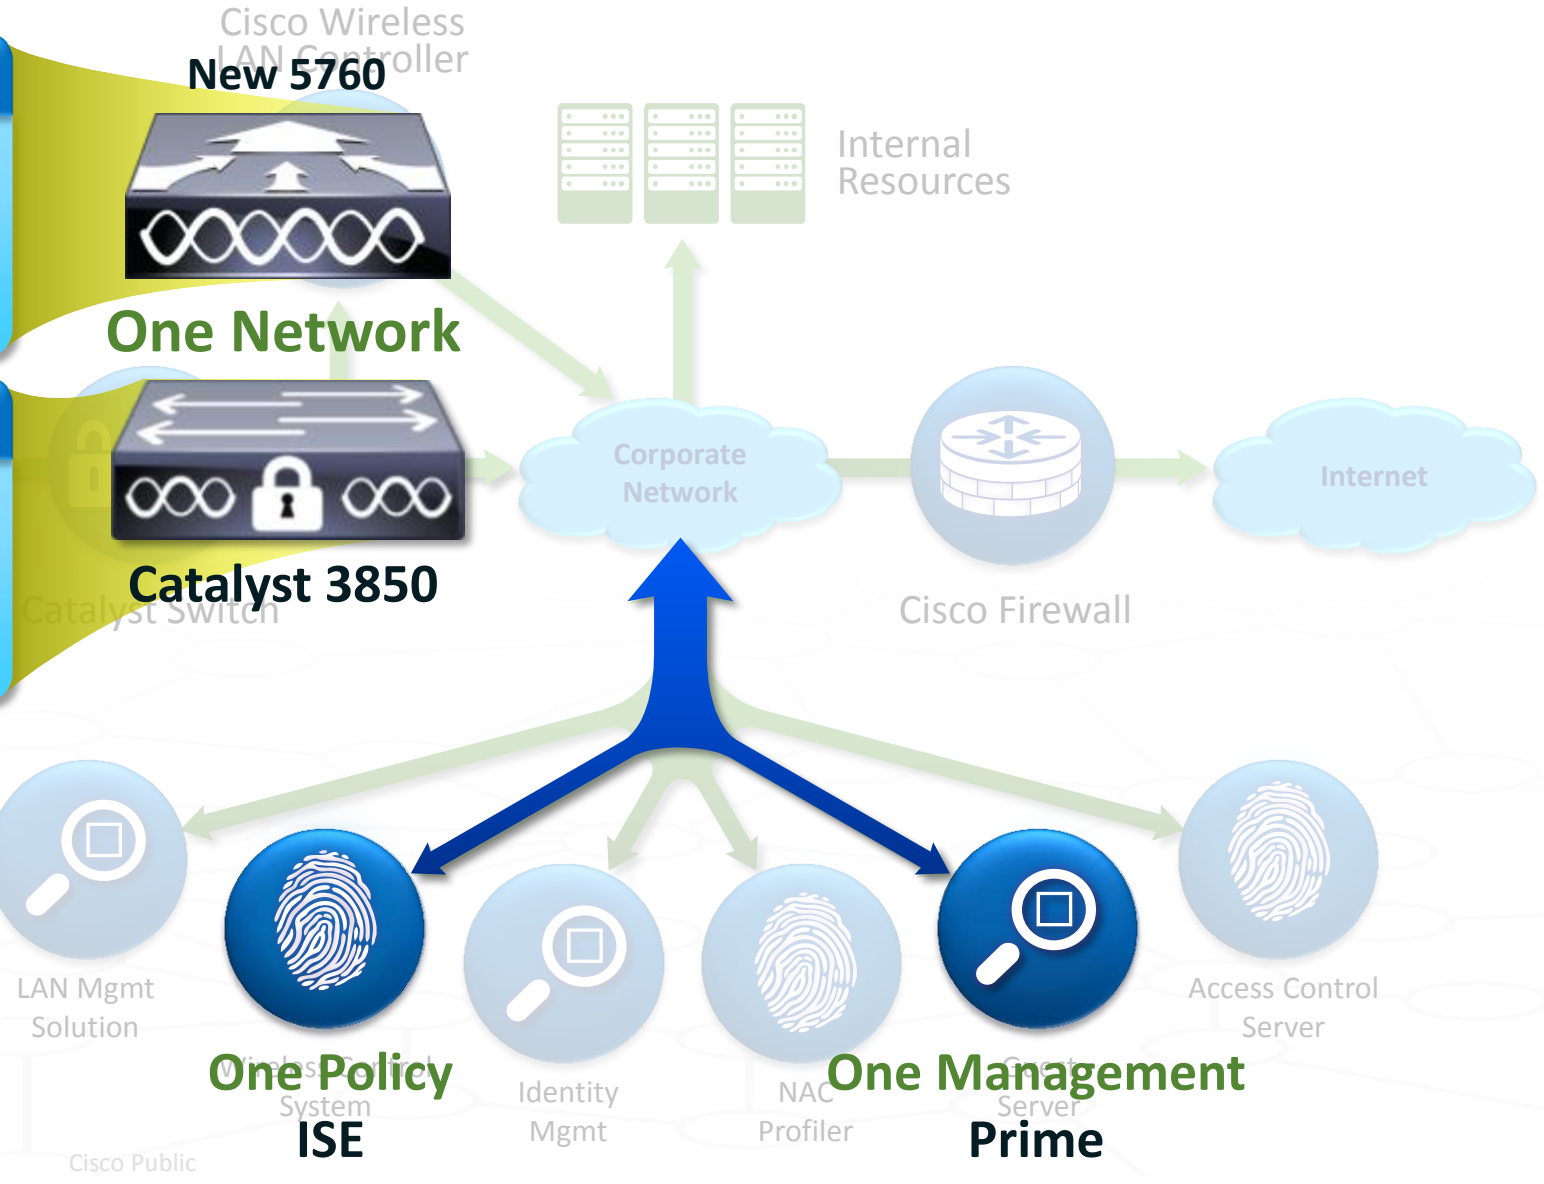

Converged Wired / Wireless Access

Benefits … Overview

Single platform for wired and wireless

Common IOS, same administration point, one release

Network wide **visibility** for faster troubleshooting

Wired and wireless traffic visible at every hop

Consistent security and Quality of Service **control**

Hierarchical bandwidth management and distributed policy enforcement

Maximum **resiliency** with fast stateful recovery

Layered network high availability design with stateful switchover

Scale with distributed wired and wireless data plane

480G stack bandwidth; 40G wireless / switch; efficient multicast

Unified Access - One Policy | One Management | One Network

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a “double-click” deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

Converged Access Components
Complete Overview

One Policy
with Identity Services Engine (ISE)

- BYOD policy management
- Device profiling and posture
- Guest access portal

One Management
with Cisco Prime 2.0

- Full wired and wireless management
- User/device centric view
- Intuitive troubleshooting workflows

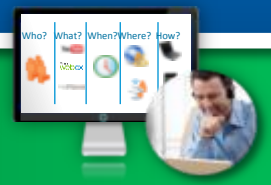

ISE

Cisco Prime

Catalyst 3850

5760 Wireless Controller

Catalyst 3850

- Industry's first fully integrated wired and wireless switch
- Wireless: 480G stack, 50 APs, 2K clients, 40G
- Flexible NetFlow, Granular QoS

5760 Wireless Controller

- Consistent IOS with Catalyst 3850
- 60G, 1K APs, 12K Clients, N+1 Redundancy
- Flexible Netflow, Granular QOS

Best-in-Class Performance, Security, and Resiliency

Catalyst 3850

Single Platform for Wired and Wireless

20+ Years of IOS Richness – Now on Wireless

Wireless

Features:

- Centralised deployment
- L2/L3 Fast Roaming
- Clean Air
- Video Stream
- Radio Resource Management (RRM)
- Wireless Security
- Radio performance
- 802.11ac Ready

Benefits

- Built on **UADP** – Cisco's Innovative Flexparser ASIC technology
- Eliminates operational complexity
- Single Operating System for wired and wireless

Wired

Features:

- Stacking, StackPower
- Advanced Identity
- Visibility and Control
- Flexible Netflow
- Granular QoS
- High Availability
- EEM, scripting
- IOS-XE Modular OS

Catalyst 3850
Platform Overview

Wireless CAPWAP Termination

Up to 50 APs/2000 clients per stack, and 40G per switch

480 Gbps Stacking Bandwidth

Up to 2000 Clients per Stack

FRU Fans, Power Supplies - HA

Full POE+

Stackpower

Granular QoS/Flexible NetFlow

Multi-Core CPU

40 Gbps Uplink Bandwidth (Modular)

Line Rate on All Ports

Built on Cisco's Innovative "UADP" ASIC

Catalyst 3850
Wireless Capabilities

- CAPWAP termination and DTLS in Hardware
- 40G wireless capacity/switch
  - Capacity increases with members
- 50 APs and 2000 clients/switch stack
- Wireless switch peer group support for faster roaming: latency sensitive applications
- Supports IPv4 and IPv6 client mobility

- APs must be directly connected to Catalyst 3850

Best-in-Class Wired Switch – with Integrated Wireless Mobility functionality

WLC 5760

Platform Overview

Centralised, or Converged Access Deployment Modes

Up to 1000 Access Points

First IOS-Based Wireless LAN Controller

Up to 12,000 Concurrent Clients

802.11ac Optimised

6x 1/10G SFP+ uplinks with LAG

Granular QoS

Flexible NetFlow

60 Gbps Wireless Bandwidth

FRU Fans

FRU Power Supplies

Built on Cisco's Innovative "UADP" ASIC

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a “double-click” deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

One Network

Wireless Deployment Mode Options, Overview

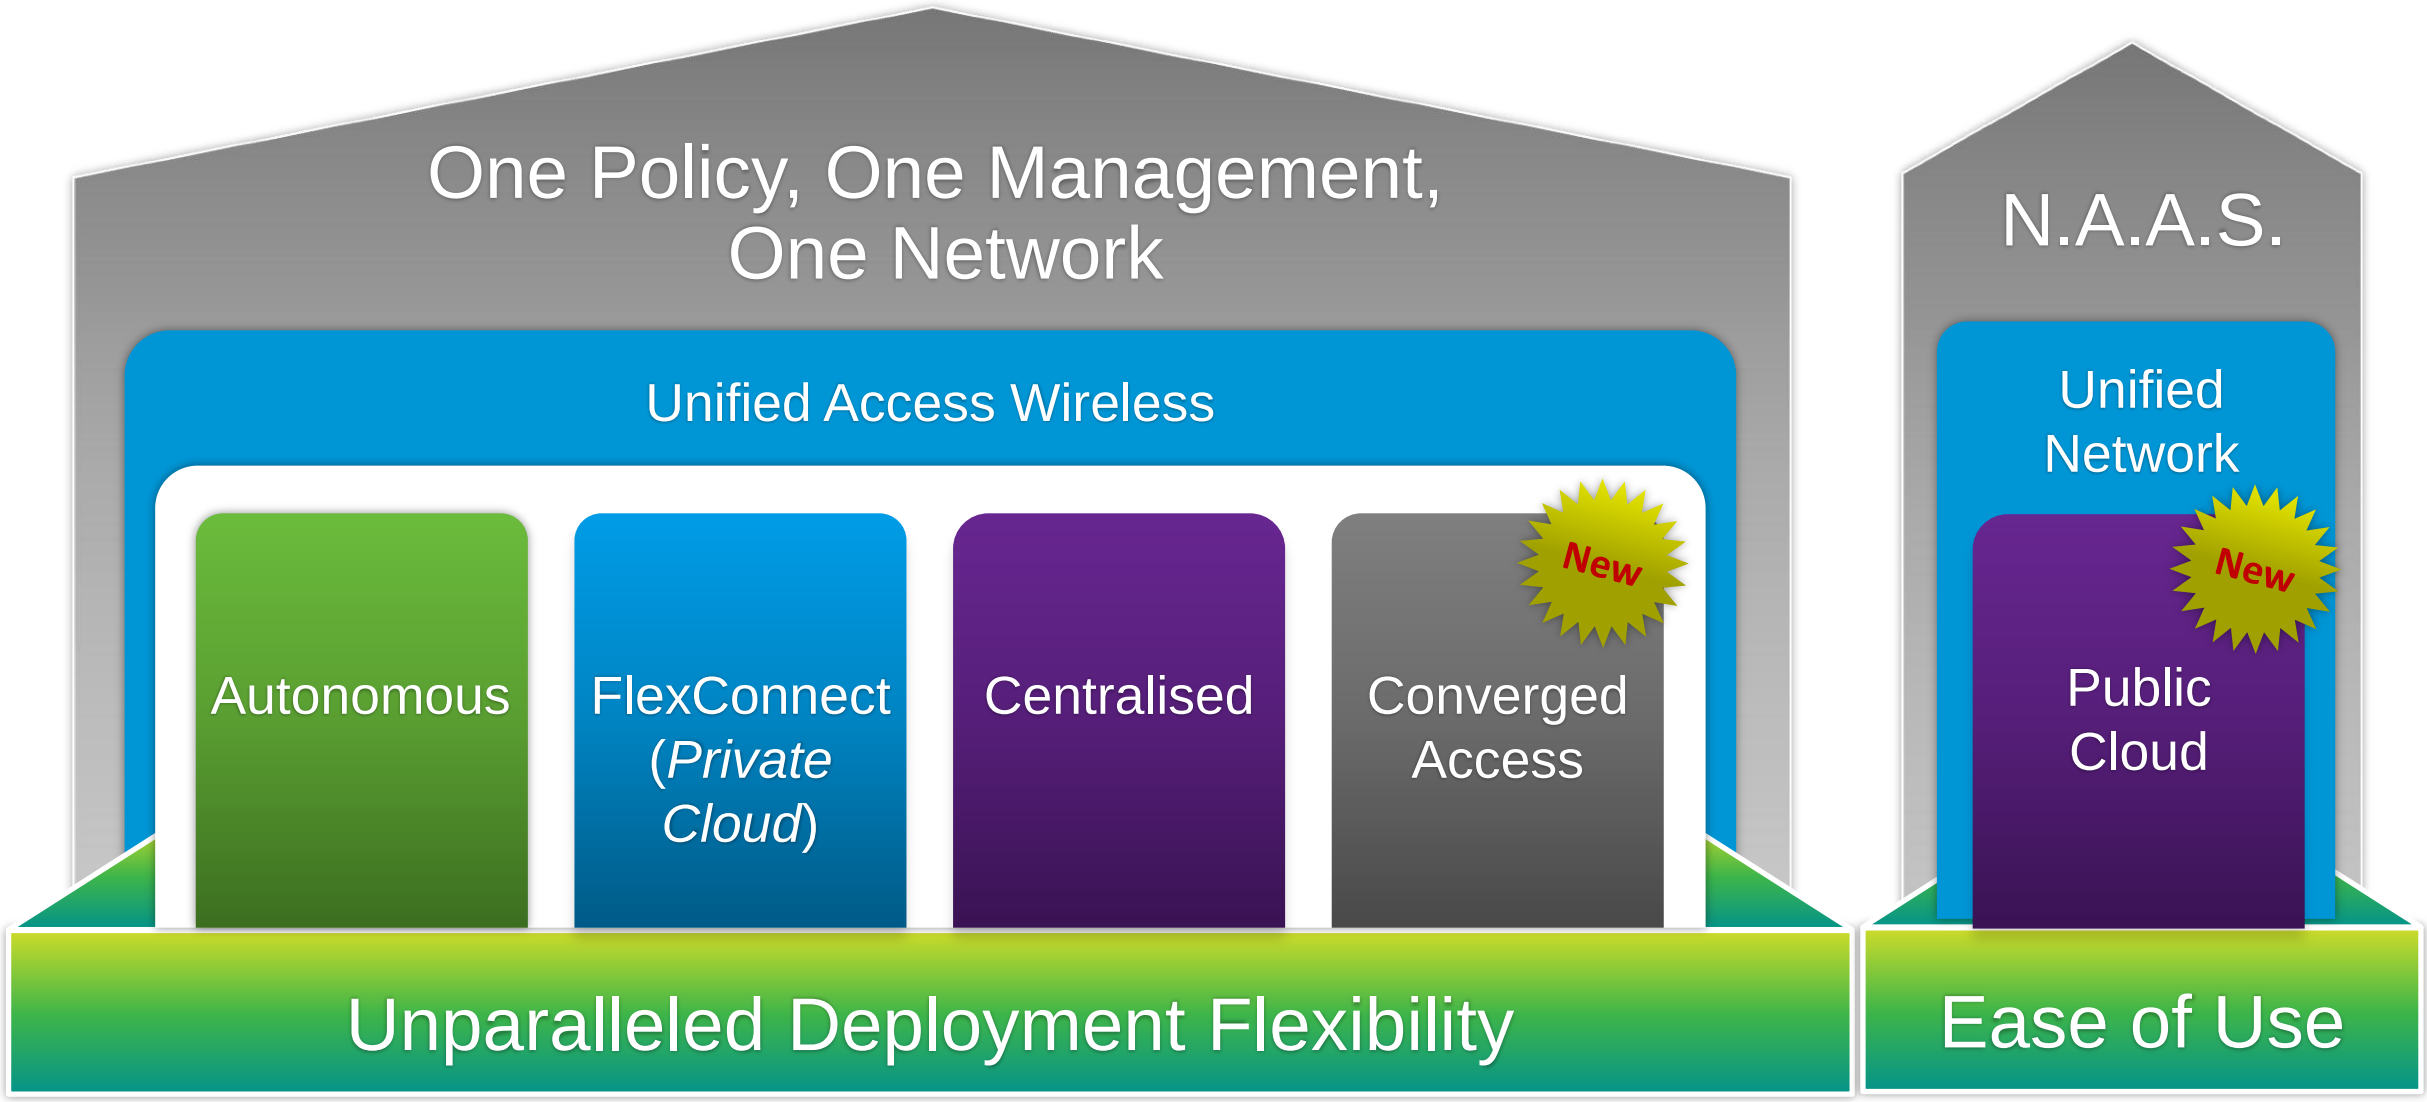

One Network

Wireless Deployment Mode Options … Detail

	Autonomous	FlexConnect	Centralised	Converged Access
			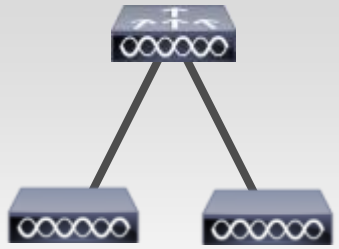	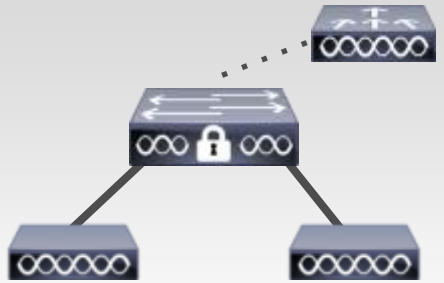
	Standalone APs	Traffic Distributed at AP	Traffic Centralised at Controller	Traffic Distributed at Switch

Target Positioning	Small Wireless Network	Branch	Campus	Branch and Campus
Purchase Decision	Wireless only	Wireless only	Wireless only	Wired and Wireless
Benefits	• Simple and cost-effective for small networks	• Highly scalable for large number of remote branches • Simple wireless operations with DC hosted controller	• Simplified operations with centralised control for Wireless • Wireless Traffic visibility at the controller	• Wired and Wireless common operations • One Enforcement Point • One OS (IOS) • Traffic visibility at every network layer • Performance optimised for 11ac
Key Considerations	• Limited RRM, no Rogue detection	• L2 roaming only • WAN BW and latency requirements	• System throughput	• Catalyst 3850 in the access layer

Converged Access Deployment
Use Cases

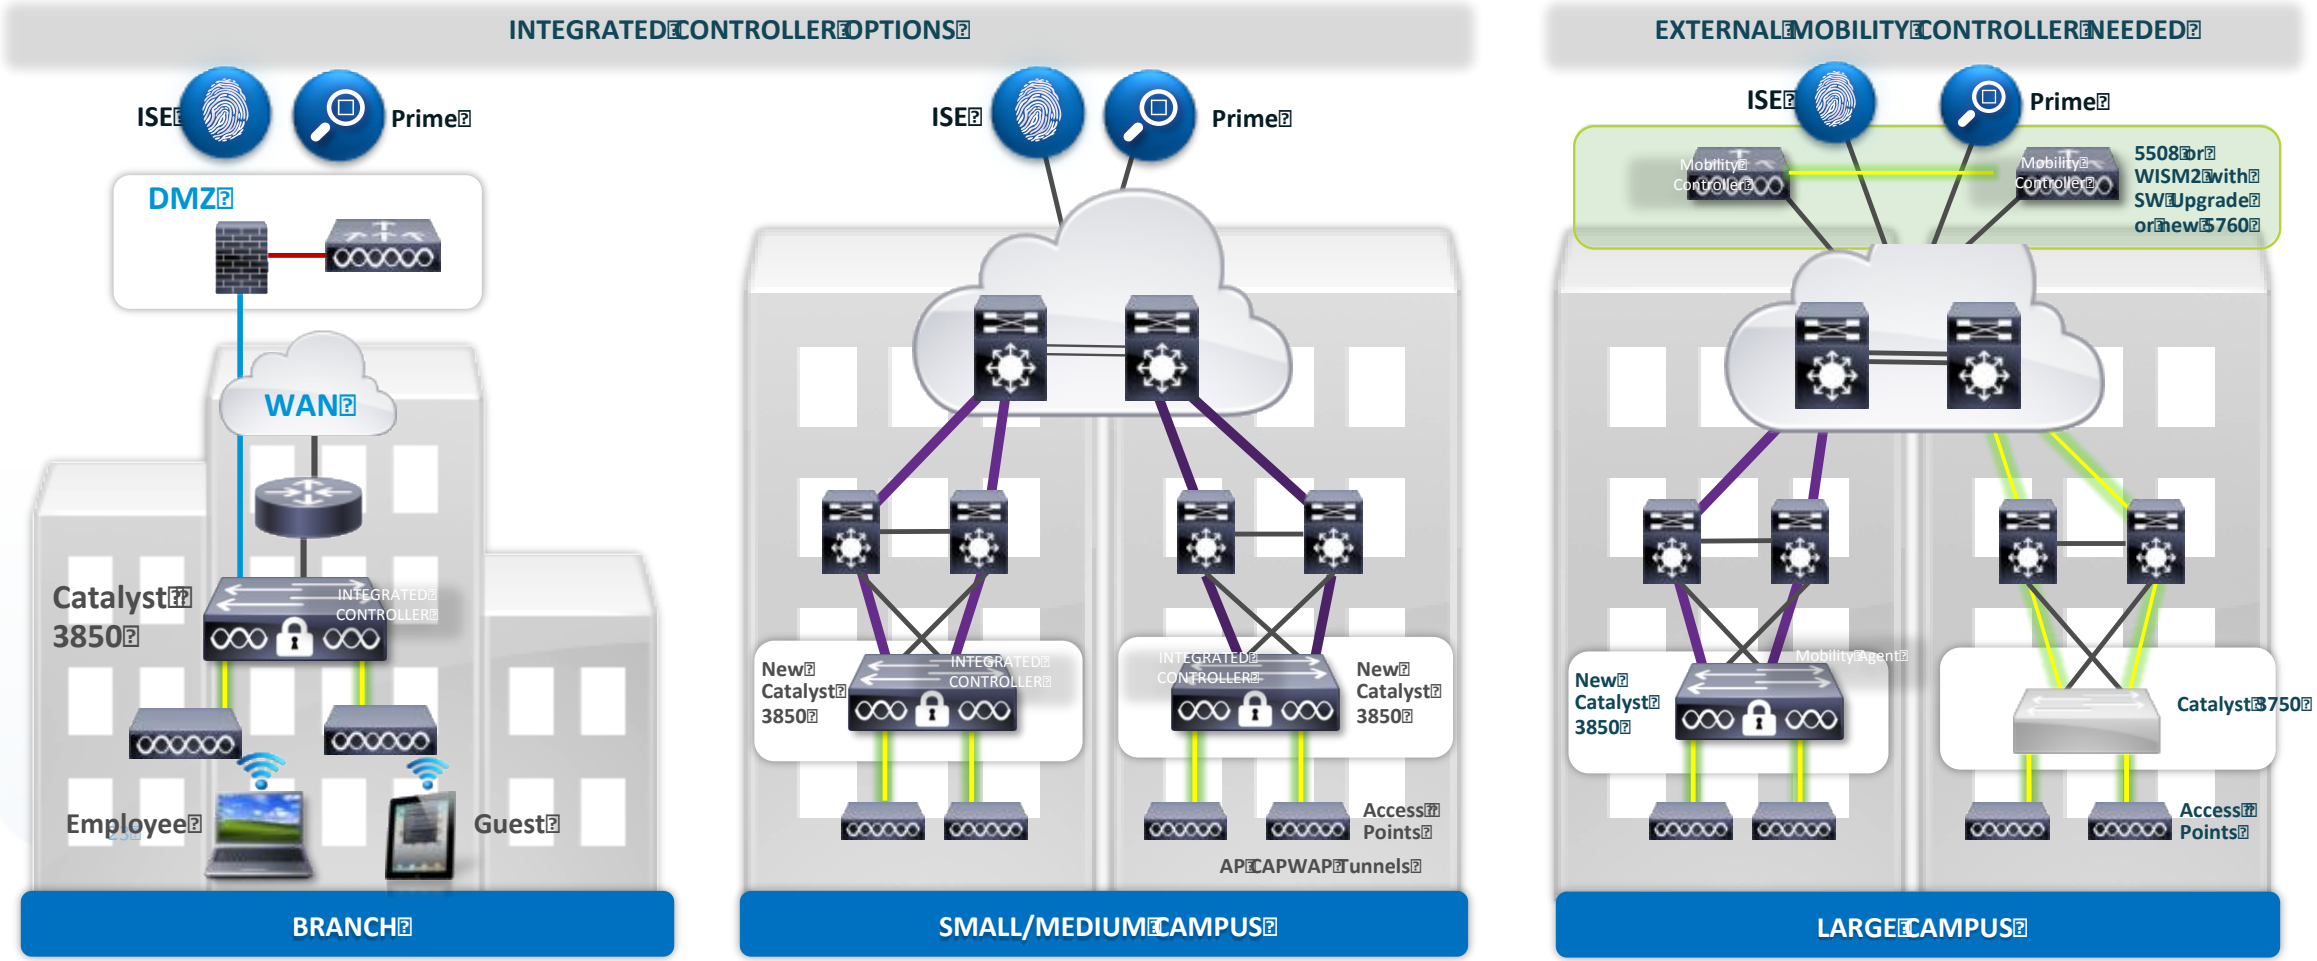

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a "double-click" deeper …

Existing Wireless Deployment – **Architecture Refresher**

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

Cisco Converged Access

Network Requirements Driving Wireless Evolution …

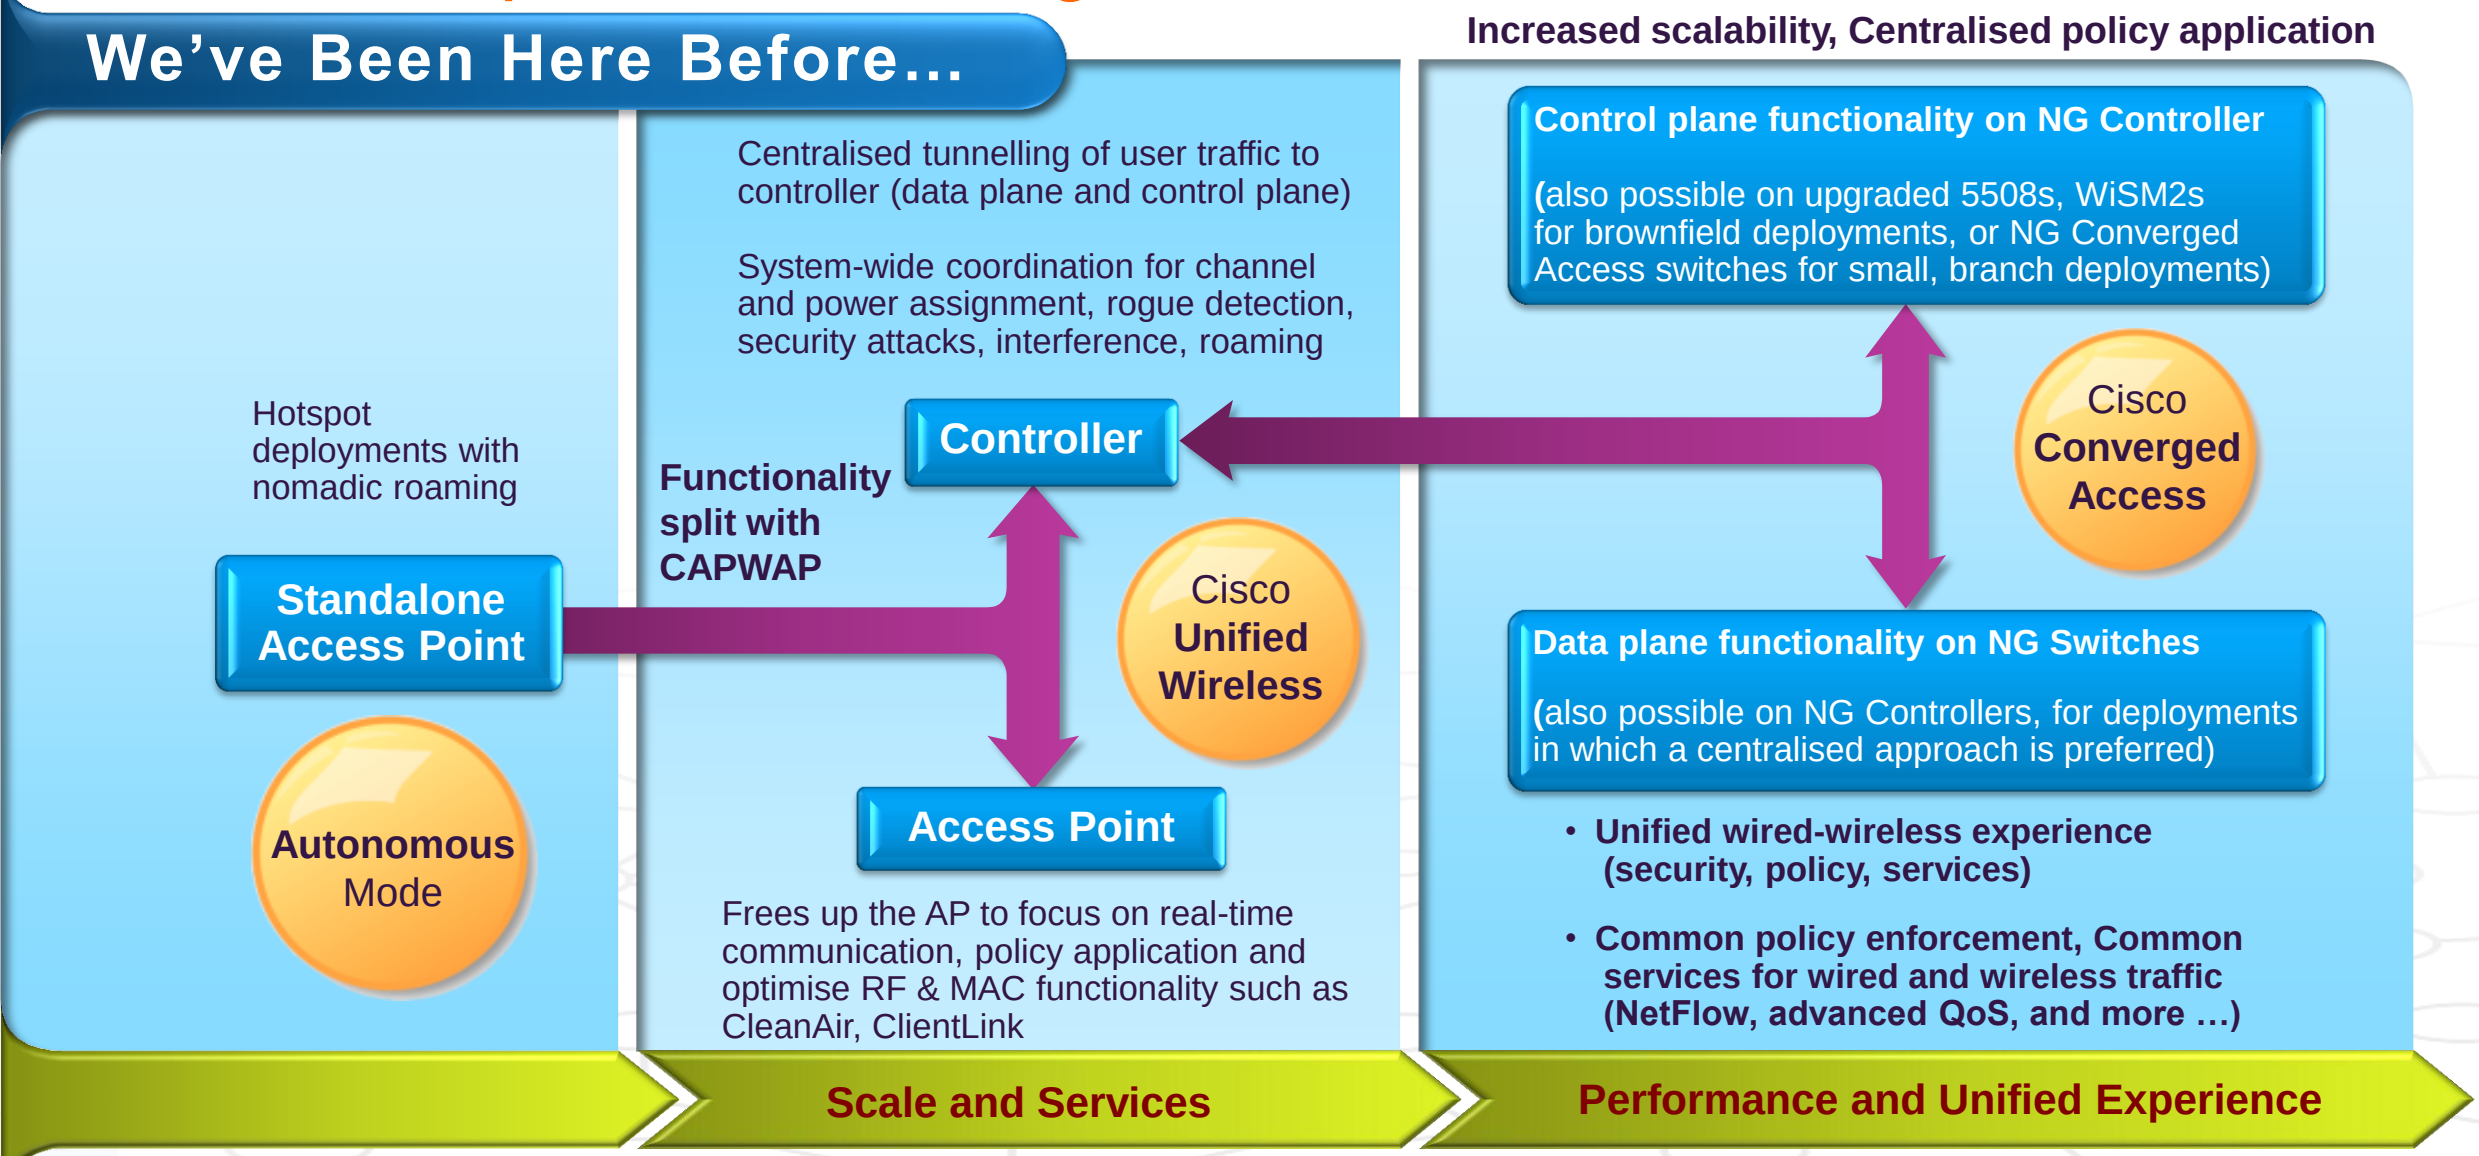

Architecture Constructs
CUWN Tunnel Types

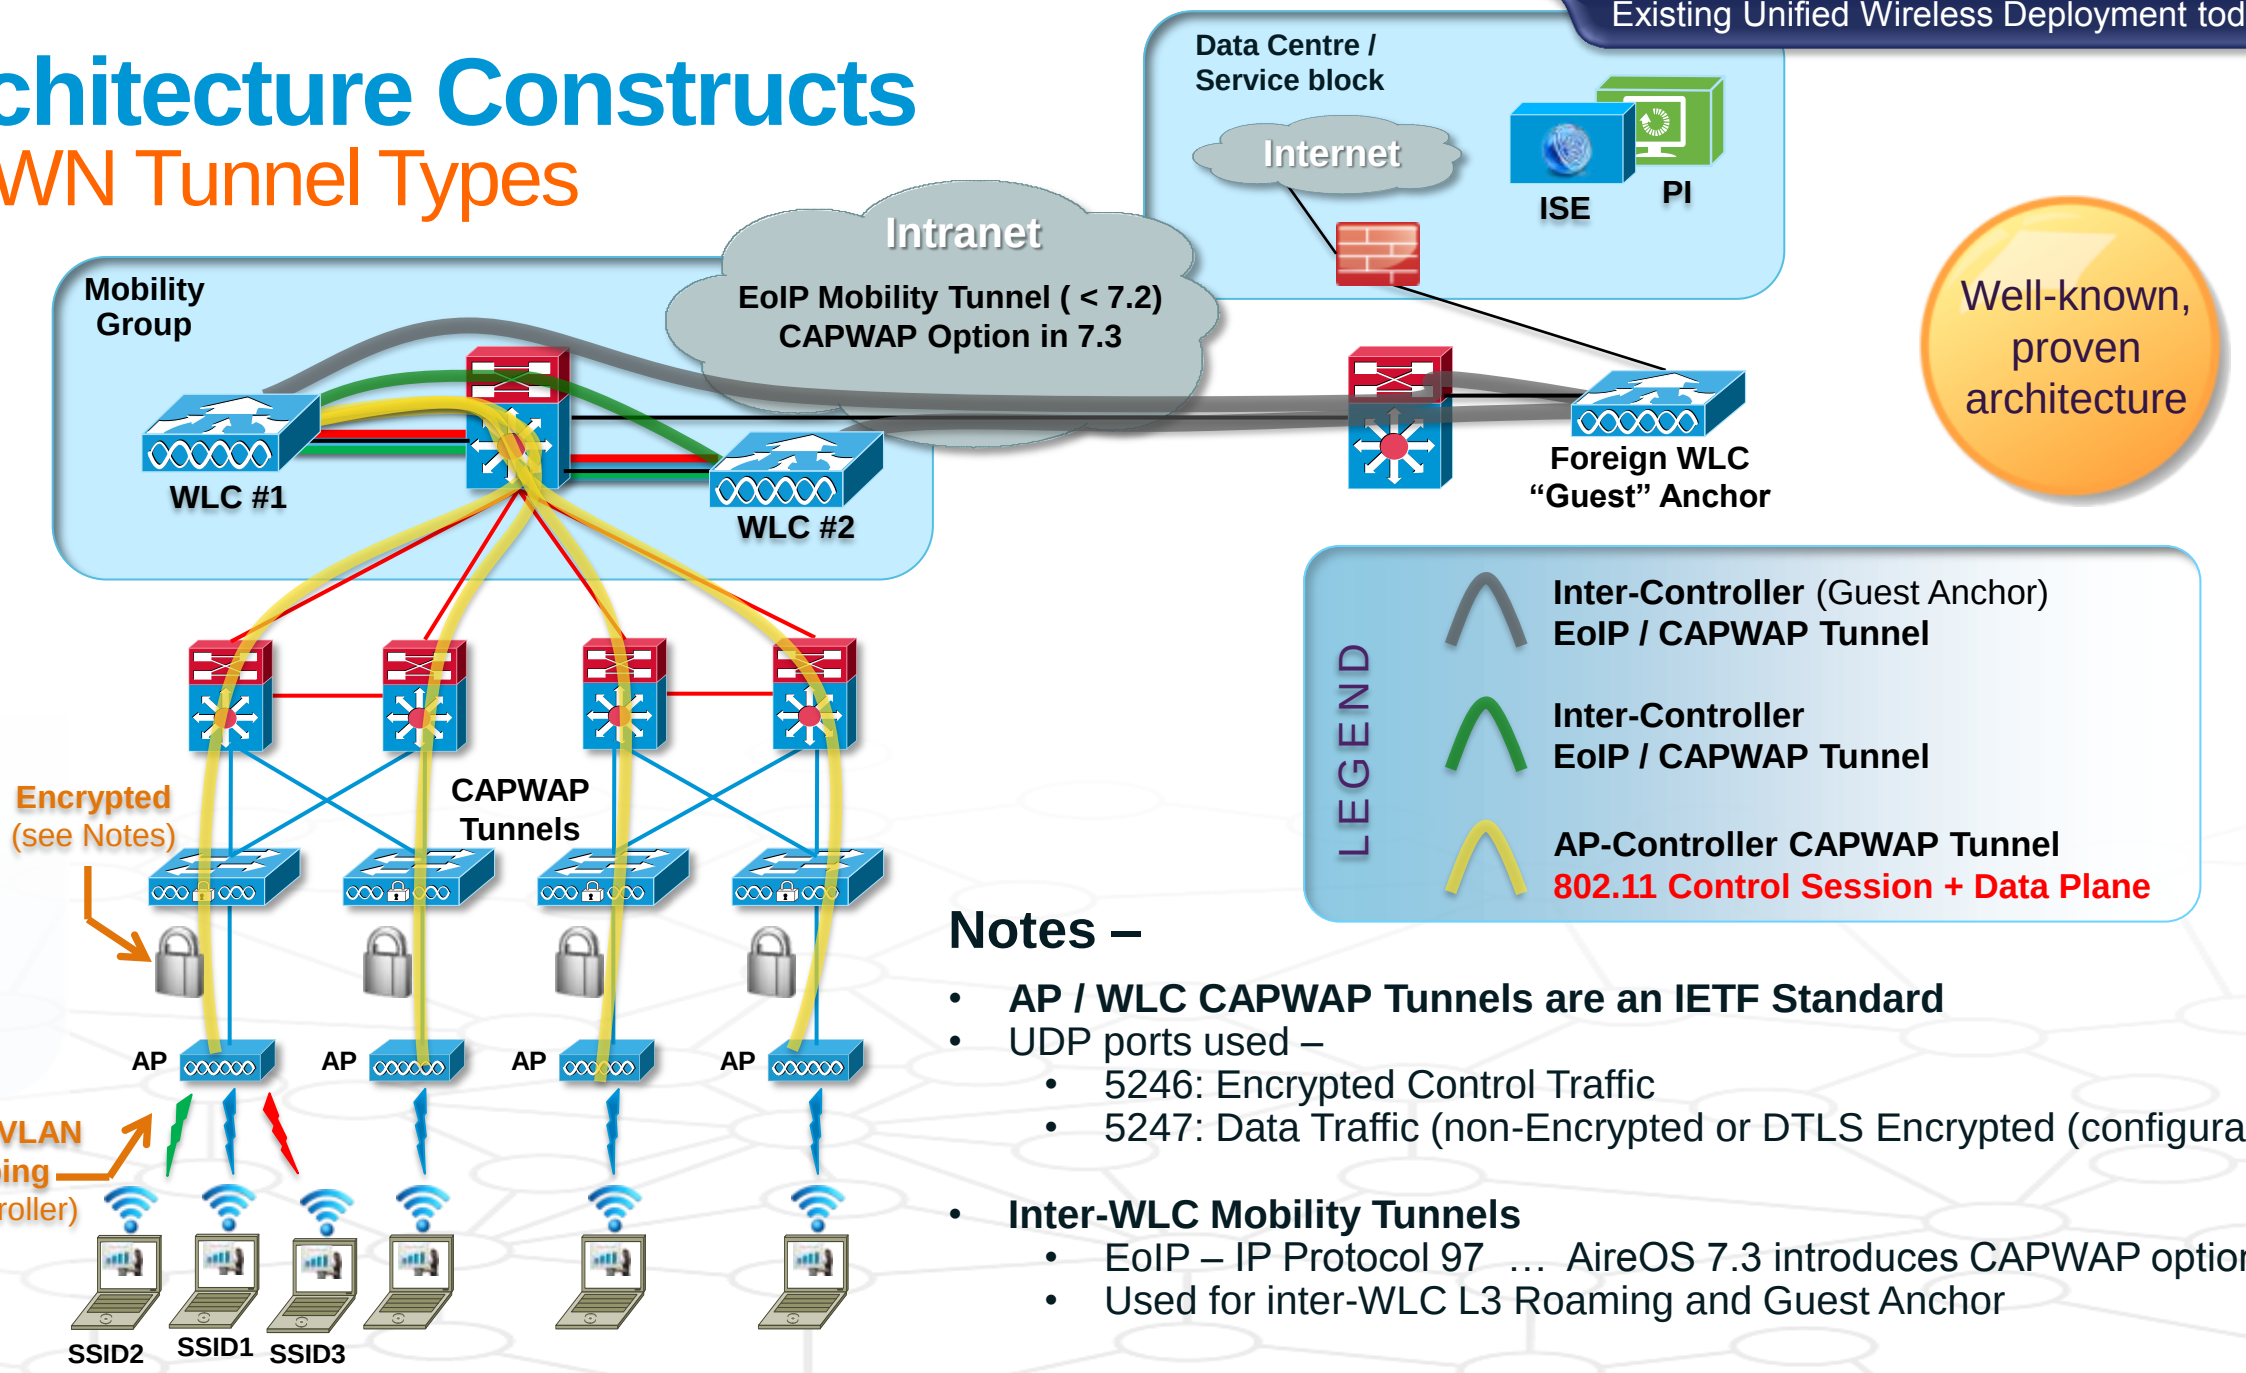

Well-known, proven architecture

Notes –

- **AP / WLC CAPWAP Tunnels are an IETF Standard**
- UDP ports used –
  - 5246: Encrypted Control Traffic
  - 5247: Data Traffic (non-Encrypted or DTLS Encrypted (configurable)
- **Inter-WLC Mobility Tunnels**
  - EoIP – IP Protocol 97 … AireOS 7.3 introduces CAPWAP option
  - Used for inter-WLC L3 Roaming and Guest Anchor

Architecture Constructs –
CUWN Control Functions

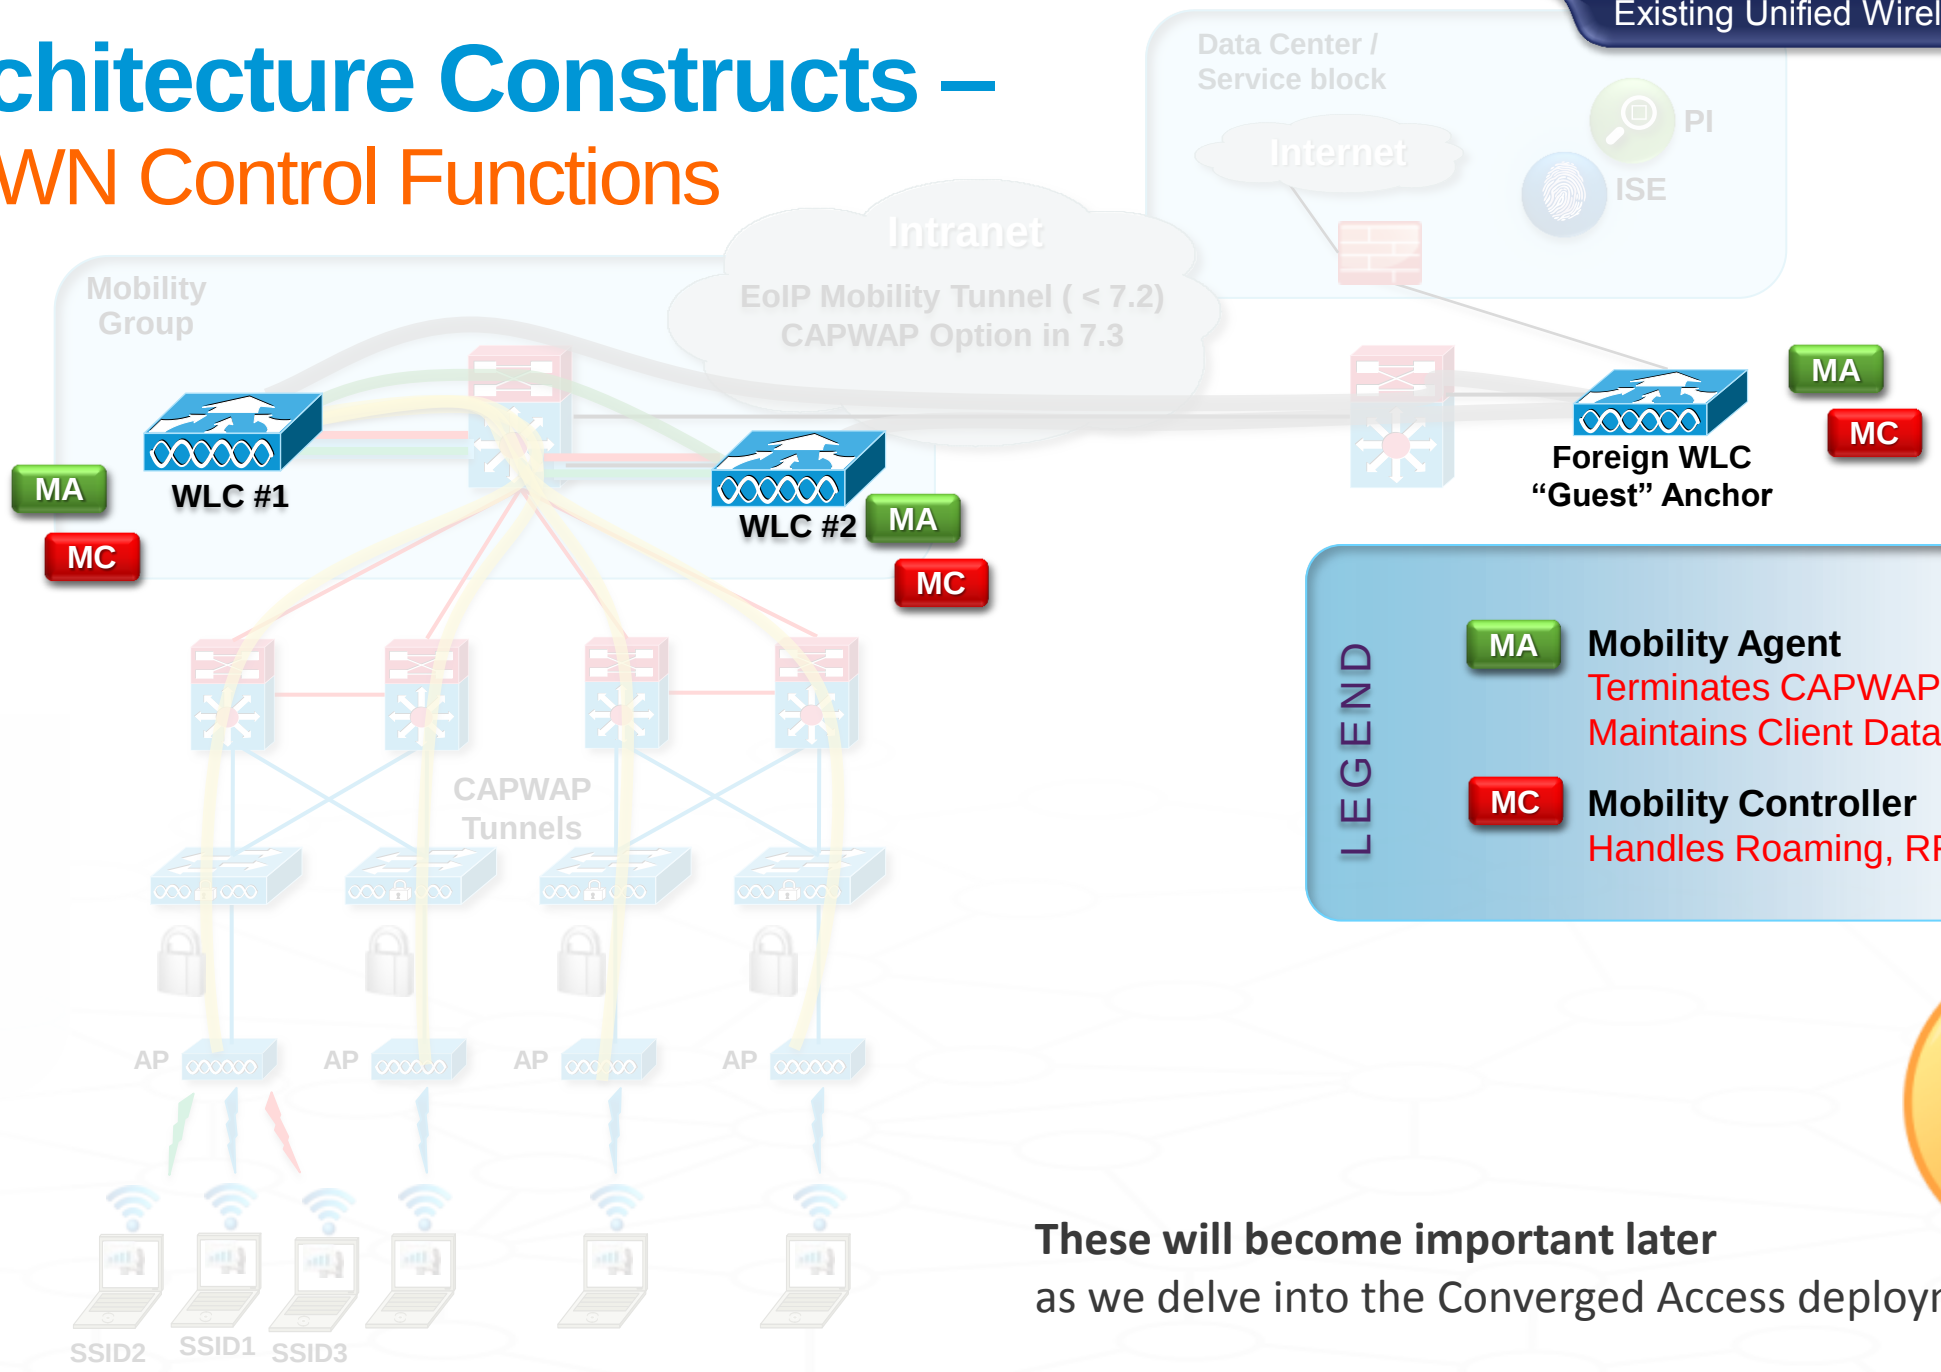

LEGEND

MA **Mobility Agent**
Terminates CAPWAP Tunnels,
Maintains Client Database

MC **Mobility Controller**
Handles Roaming, RRM, WIPS, etc.

Additional details on controller functionality

These will become important later
as we delve into the Converged Access deployment …

Architecture Constructs

Point of Presence (PoP), Point of Attachment (PoA)

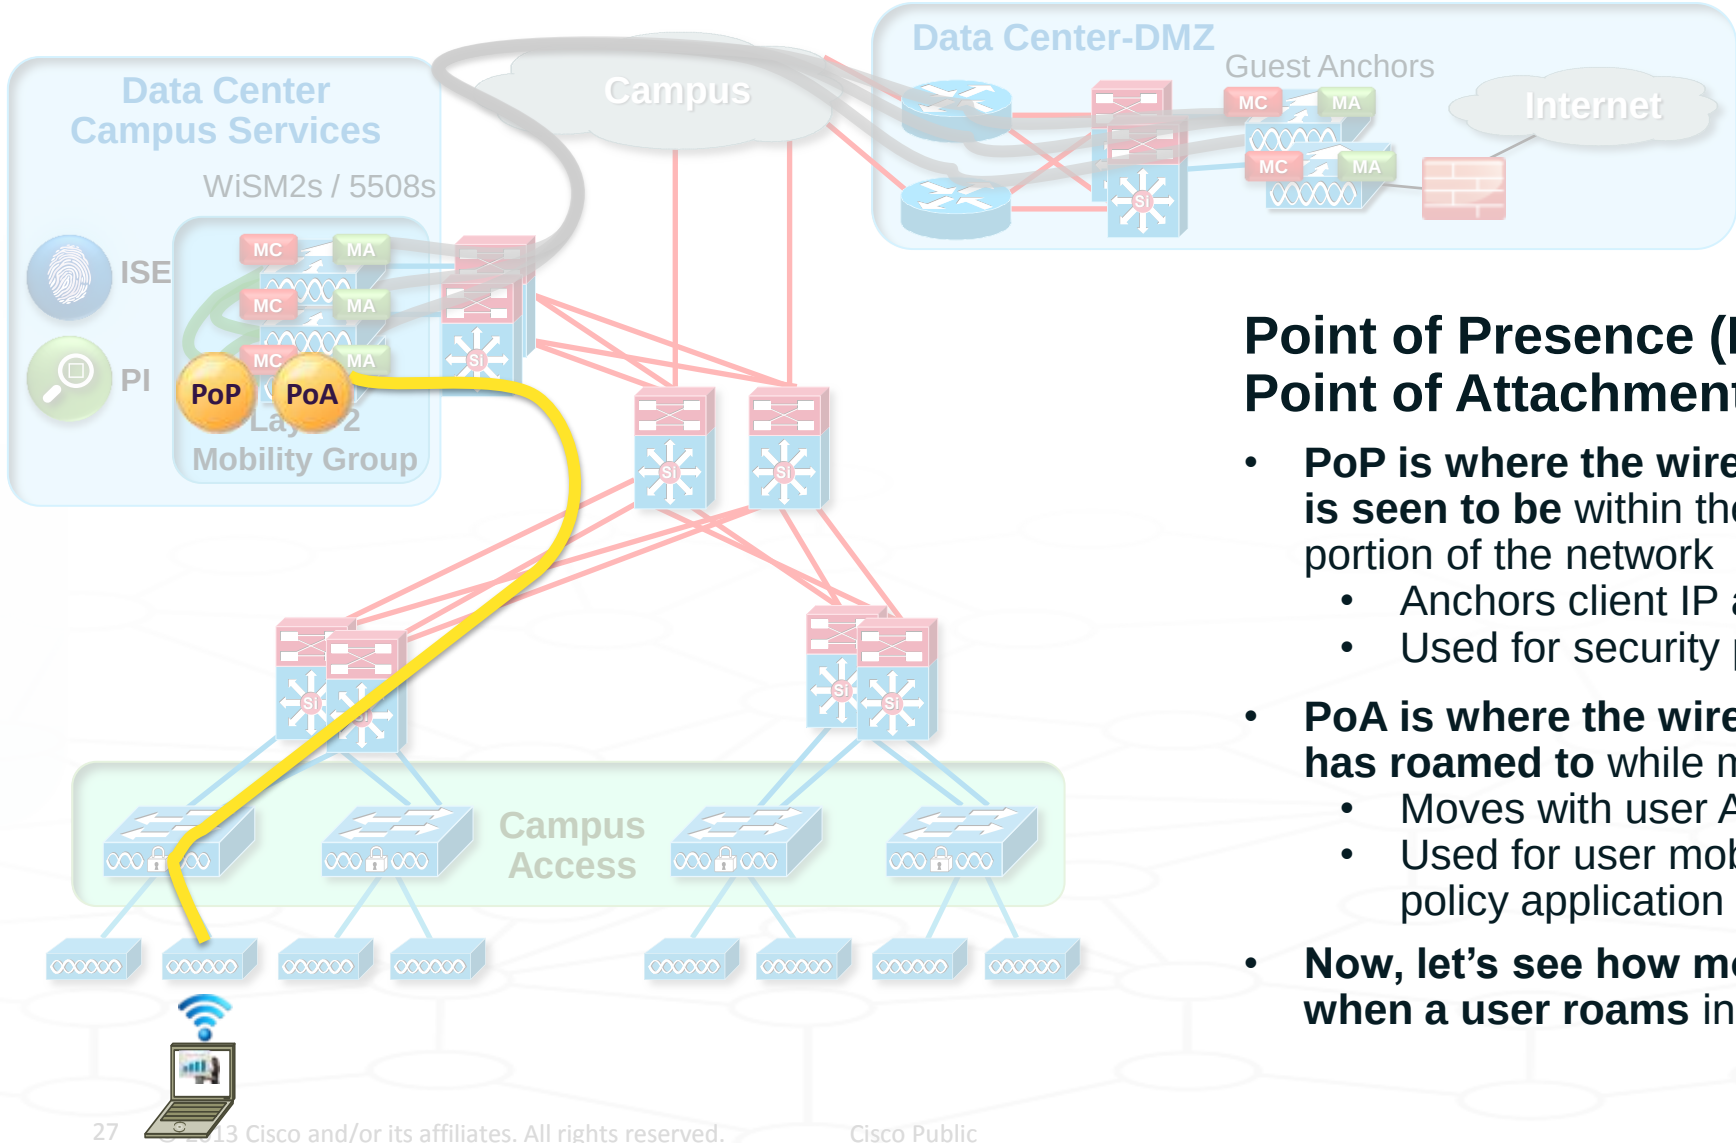

Point of Presence (PoP) vs. Point of Attachment (PoA) –

- **PoP is where the wireless user is seen to be** within the wired portion of the network
  - Anchors client IP address
  - Used for security policy application
- **PoA is where the wireless user has roamed to** while mobile
  - Moves with user AP connectivity
  - Used for user mobility and QoS policy application
- **Now, let's see how mobility works when a user roams** in this deployment model …

Architecture Constructs

Layer 2 Roaming (Campus Deployment)

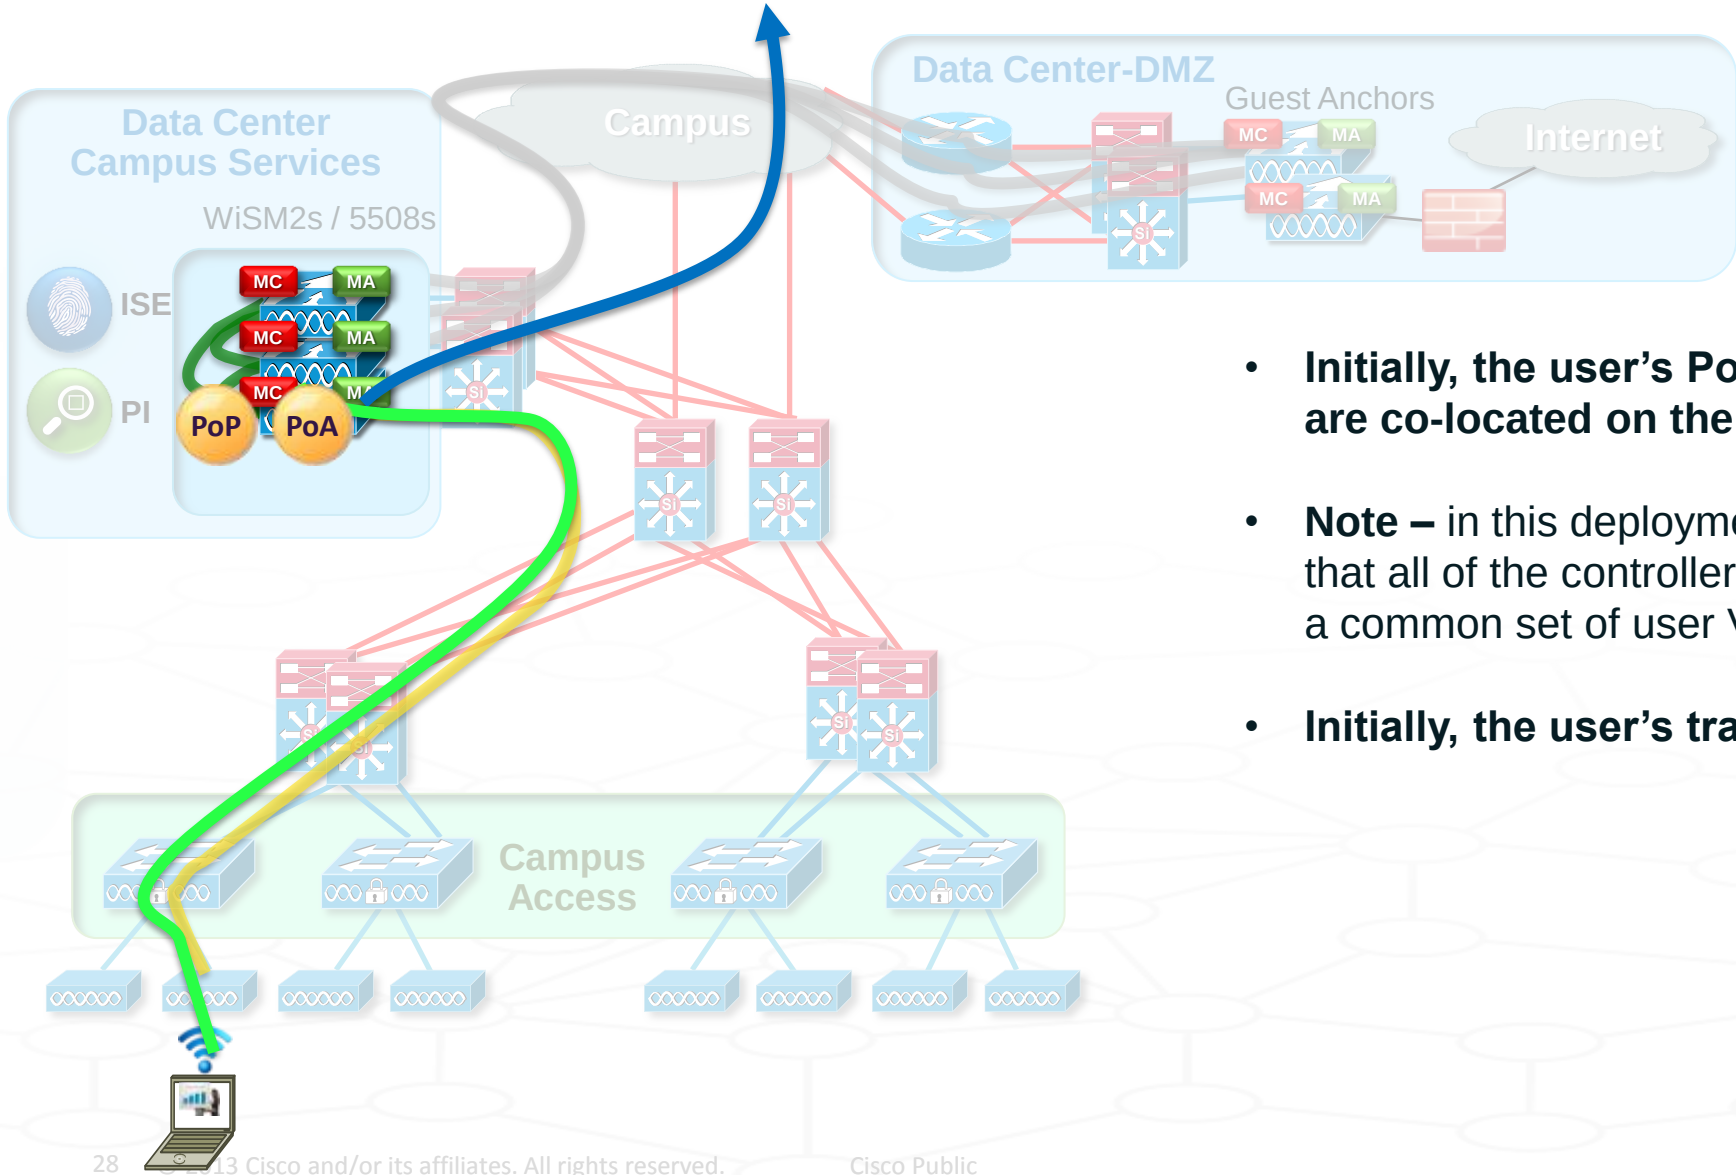

- **Initially, the user's PoP and PoA are co-located on the same controller**
- **Note –** in this deployment model, it is assumed that all of the controllers within the DC share a common set of user VLANs at Layer 2
- **Initially, the user's traffic flow is as shown …**

Architecture Constructs

Layer 2 Roaming (Campus Deployment)

Move of the user's entire Mobility Context

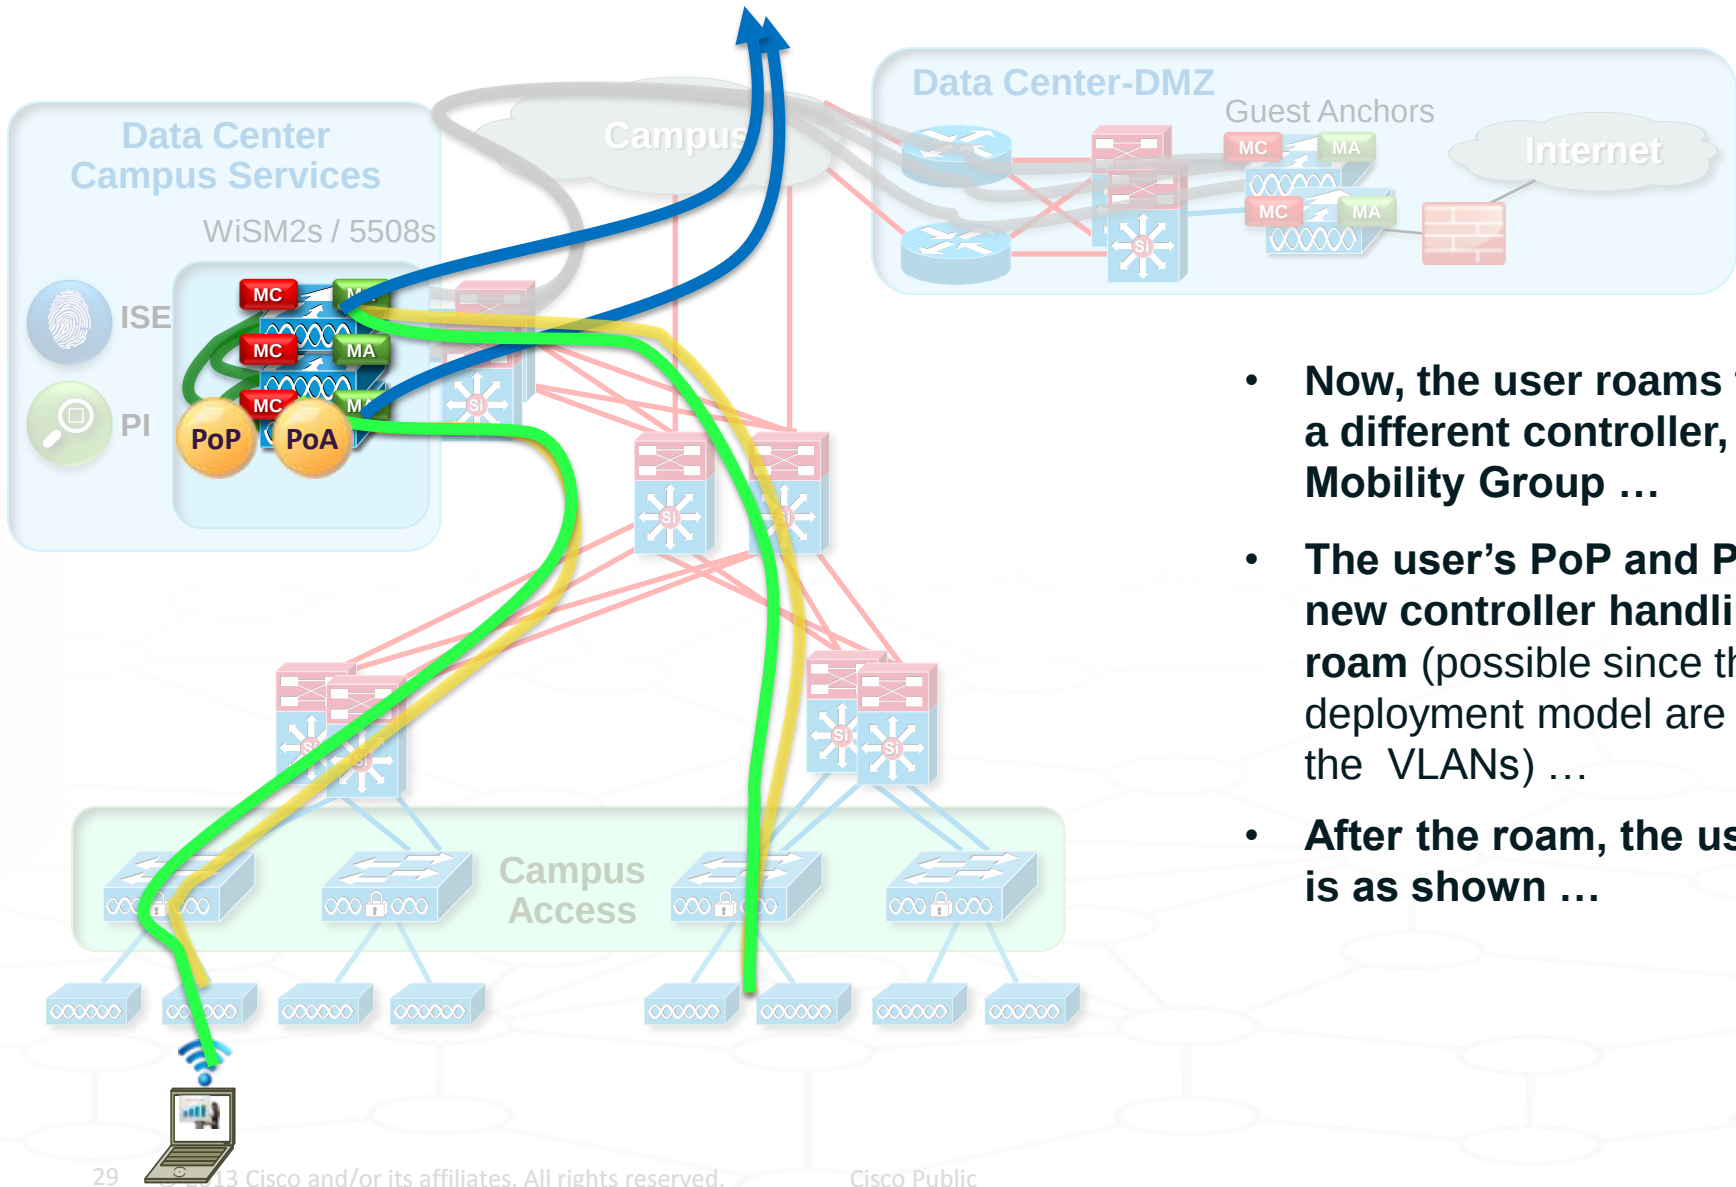

- **Now, the user roams to an AP handled by a different controller, within the same Mobility Group …**
- **The user's PoP and PoA both move to the new controller handling that user after the roam** (possible since the controllers in this deployment model are all L2-adjacent within the VLANs) …
- **After the roam, the user's traffic flow is as shown …**

Architecture Constructs

Layer 3 Roaming (Campus Deployment)

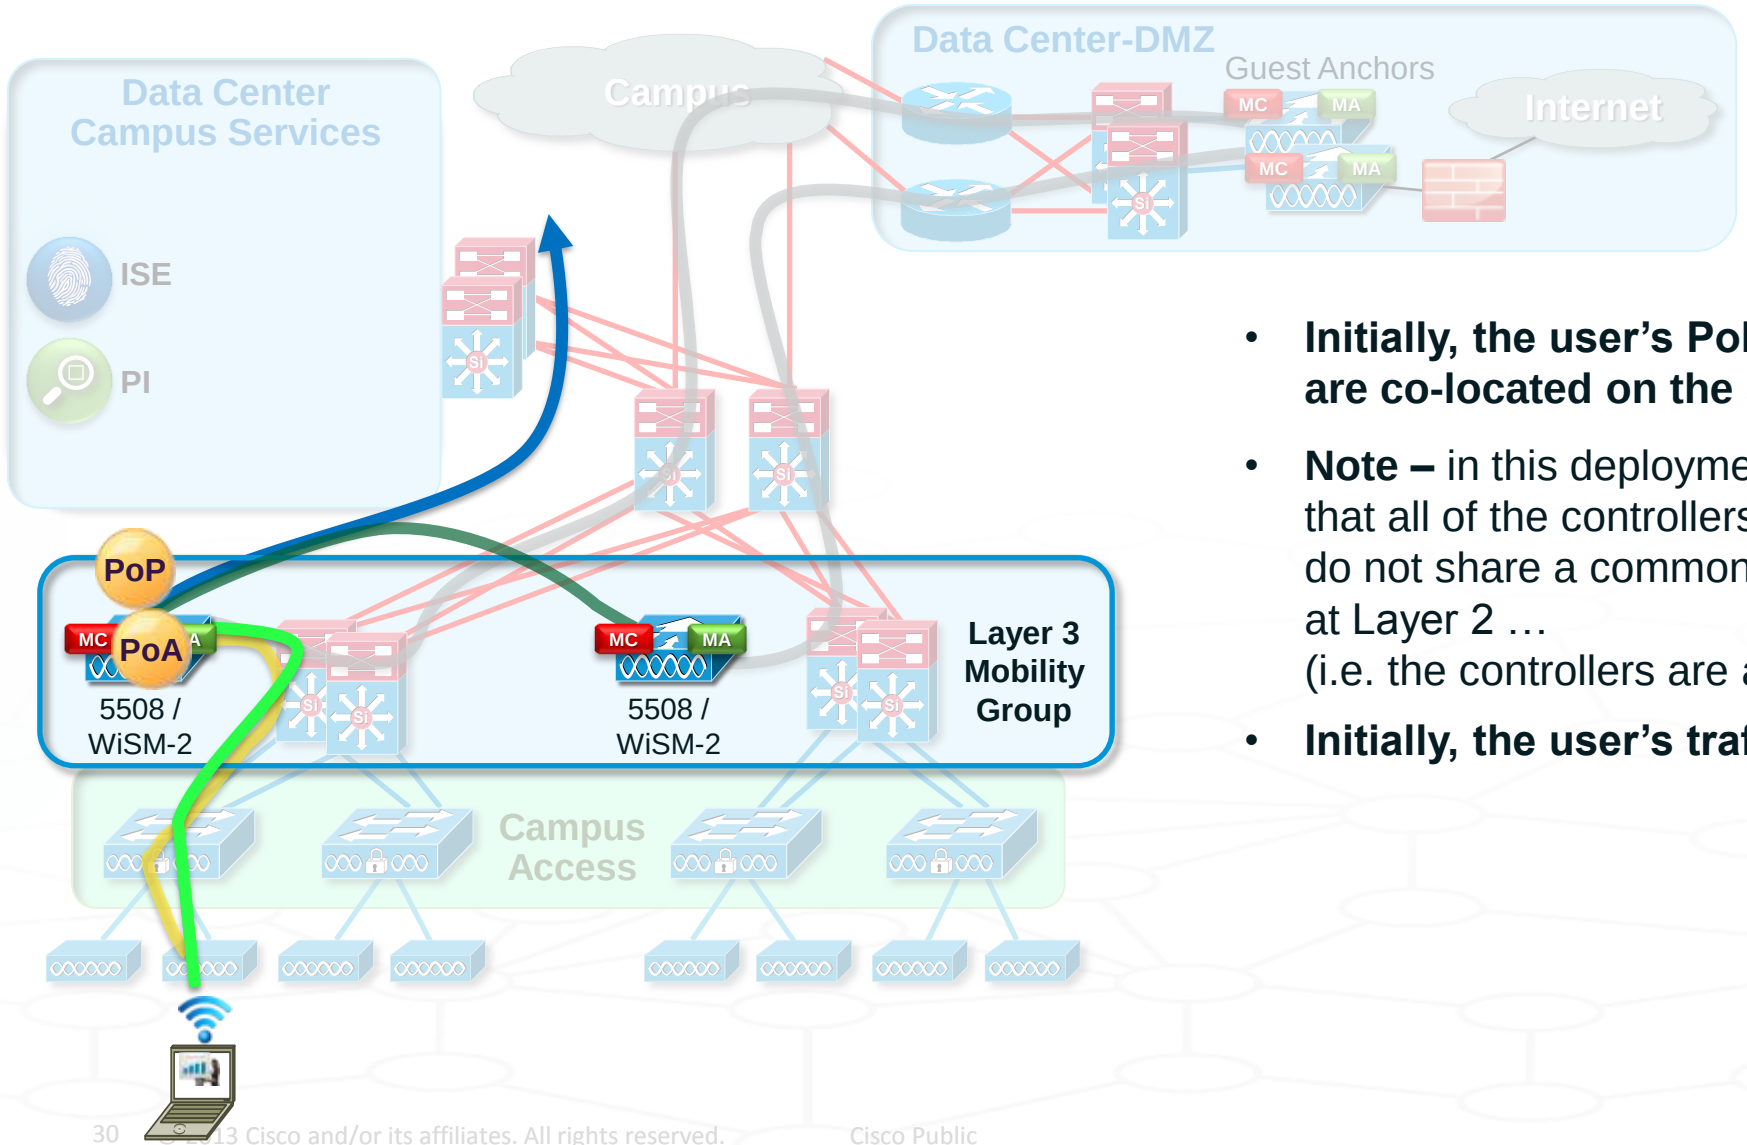

- **Initially, the user's PoP and PoA are co-located on the same controller**
- **Note –** in this deployment model, it is assumed that all of the controllers across the Campus do not share a common set of user VLANs at Layer 2 …
  (i.e. the controllers are all L3-separated)
- **Initially, the user's traffic flow is as shown …**

Architecture Constructs

Layer 3 Roaming (Campus Deployment)

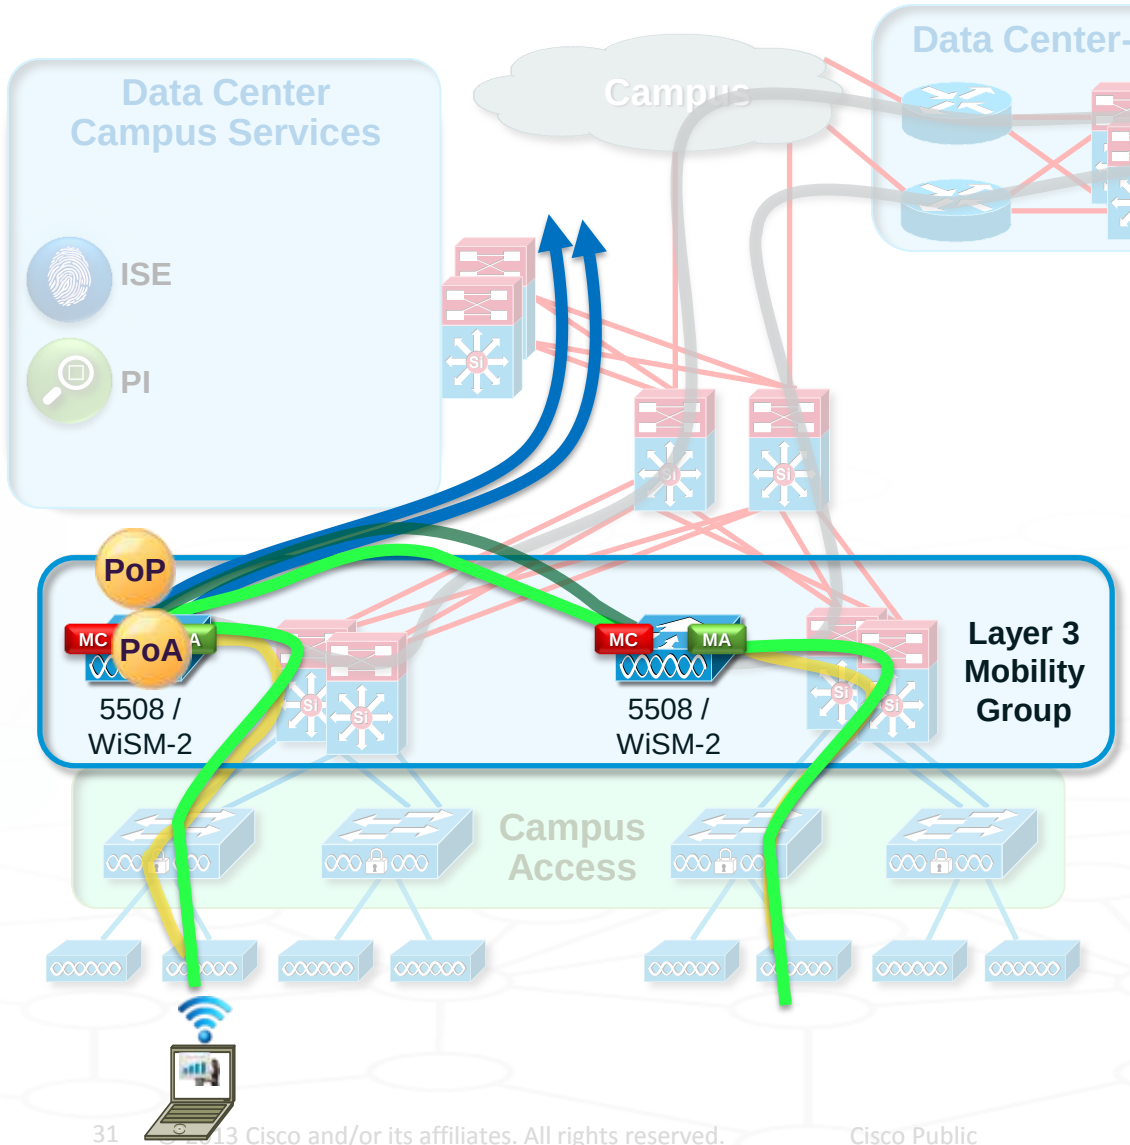

Symmetric Mobility Tunnelling

- **Now, the user roams to an AP handled by a different controller, within the same Mobility Group …**
- **The user's PoA moves to the new controller handling that user after the roam – but the user's PoP stays fixed on the original controller that the user associated to**
- **This is done to ensure that the user retains the same IP address across an L3 boundary roam –** and also to ensure continuity of policy application during roaming
- **After the roam, the user's traffic flow is as shown …**

Architecture Constructs
Roaming with Mobility Anchors

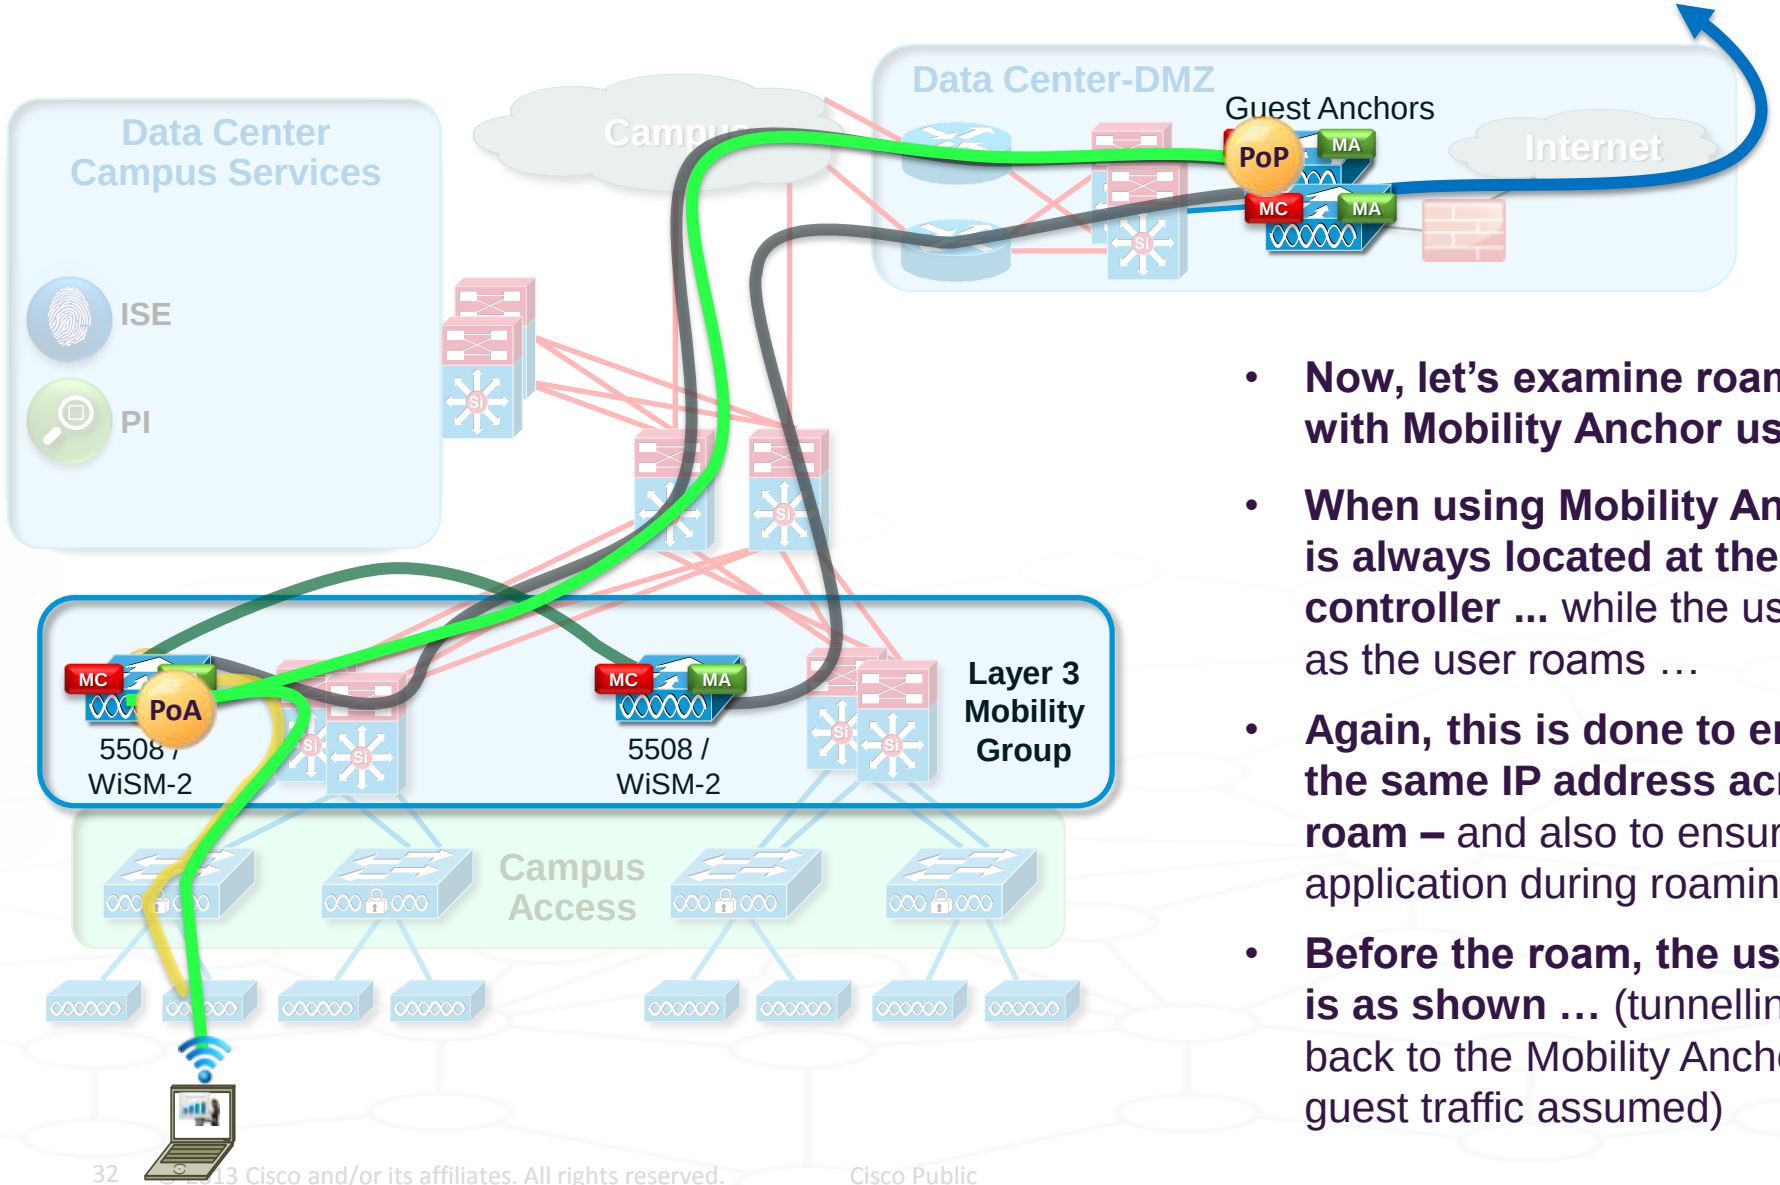

- **Now, let's examine roaming with Mobility Anchor use …**
- **When using Mobility Anchors, the user's PoP is always located at the Mobility Anchor controller ...** while the user's PoA moves as the user roams …
- **Again, this is done to ensure that the user retains the same IP address across an L3 boundary roam –** and also to ensure continuity of policy application during roaming
- **Before the roam, the user's traffic flow is as shown …** (tunnelling of user traffic back to the Mobility Anchor – guest traffic assumed)

Architecture Constructs
Roaming with Mobility Anchors

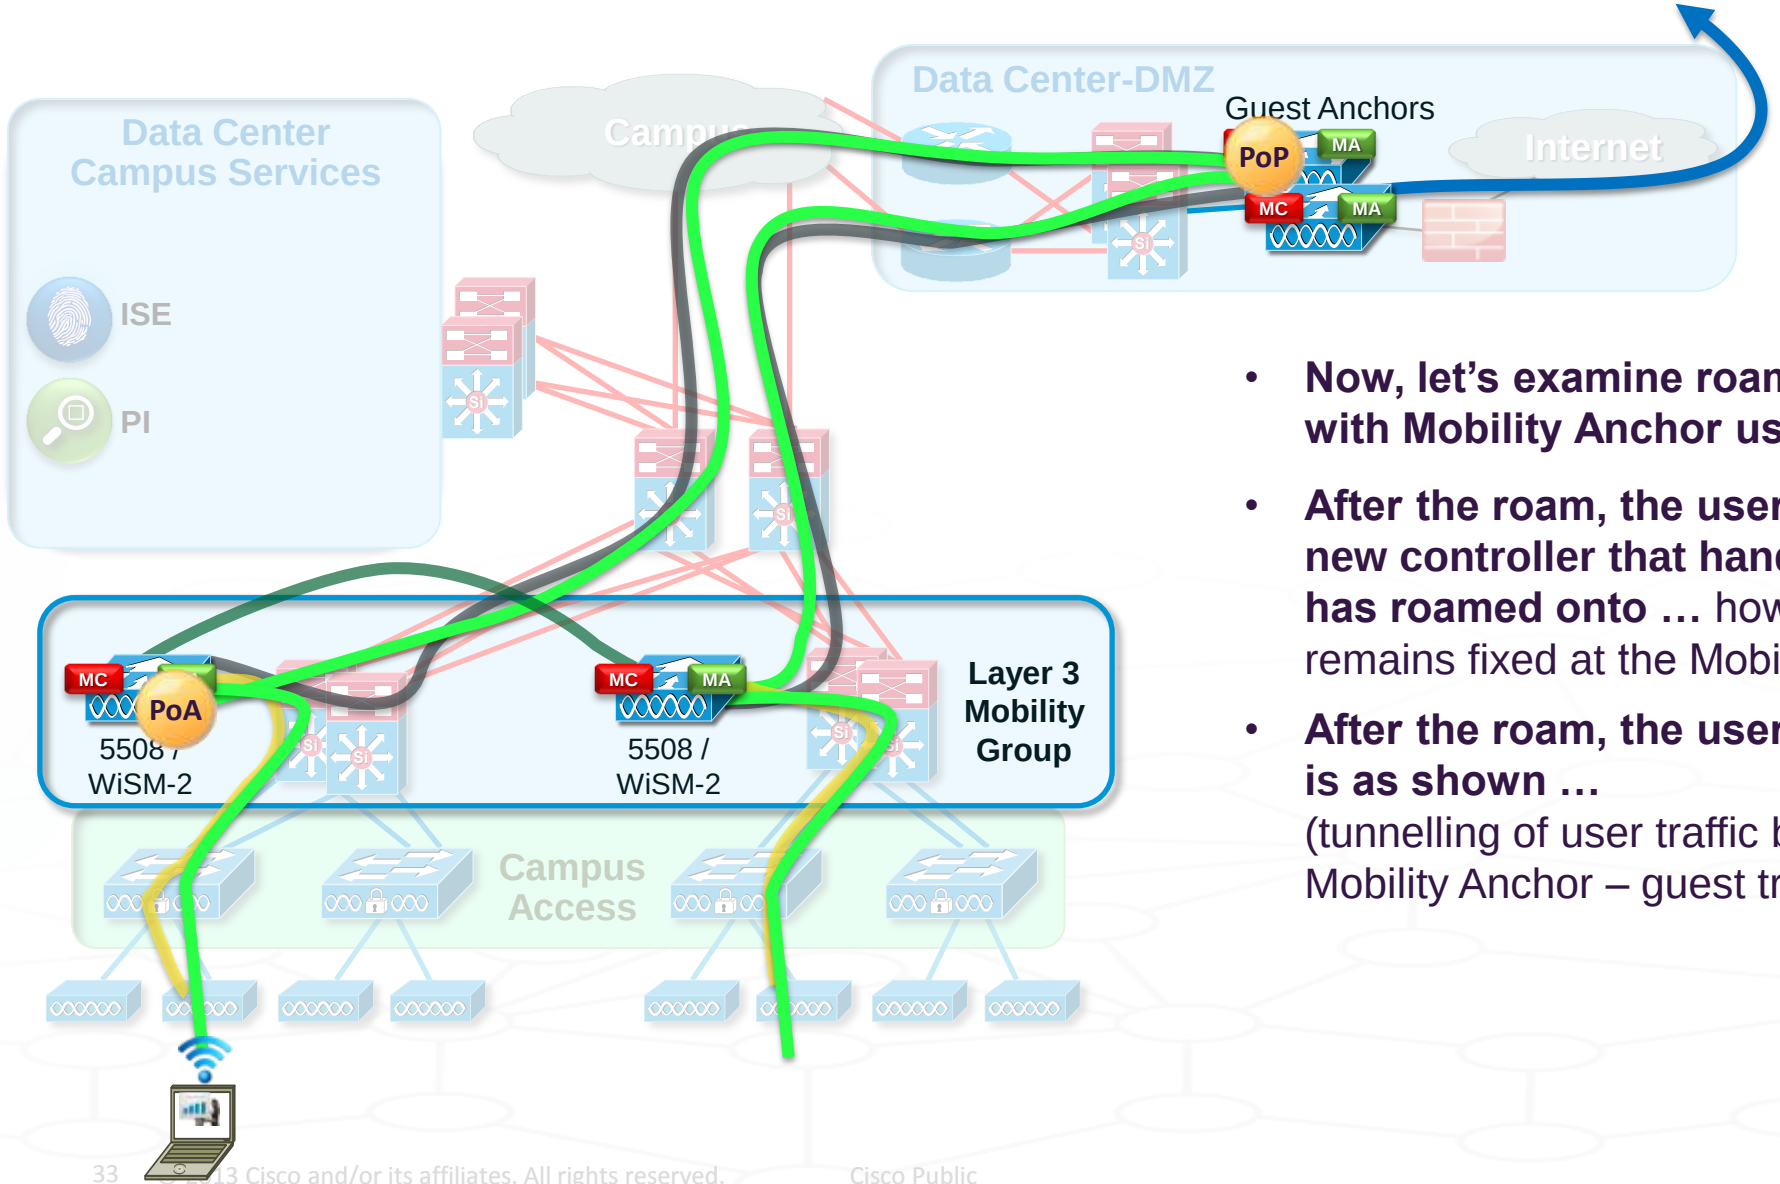

- **Now, let's examine roaming with Mobility Anchor use …**
- **After the roam, the user's PoA moves to the new controller that handles the AP the user has roamed onto …** however, the user's PoP remains fixed at the Mobility Anchor controller …
- **After the roam, the user's traffic flow is as shown …**
  (tunnelling of user traffic back to the Mobility Anchor – guest traffic assumed)

Unified Wireless
Traffic Flow

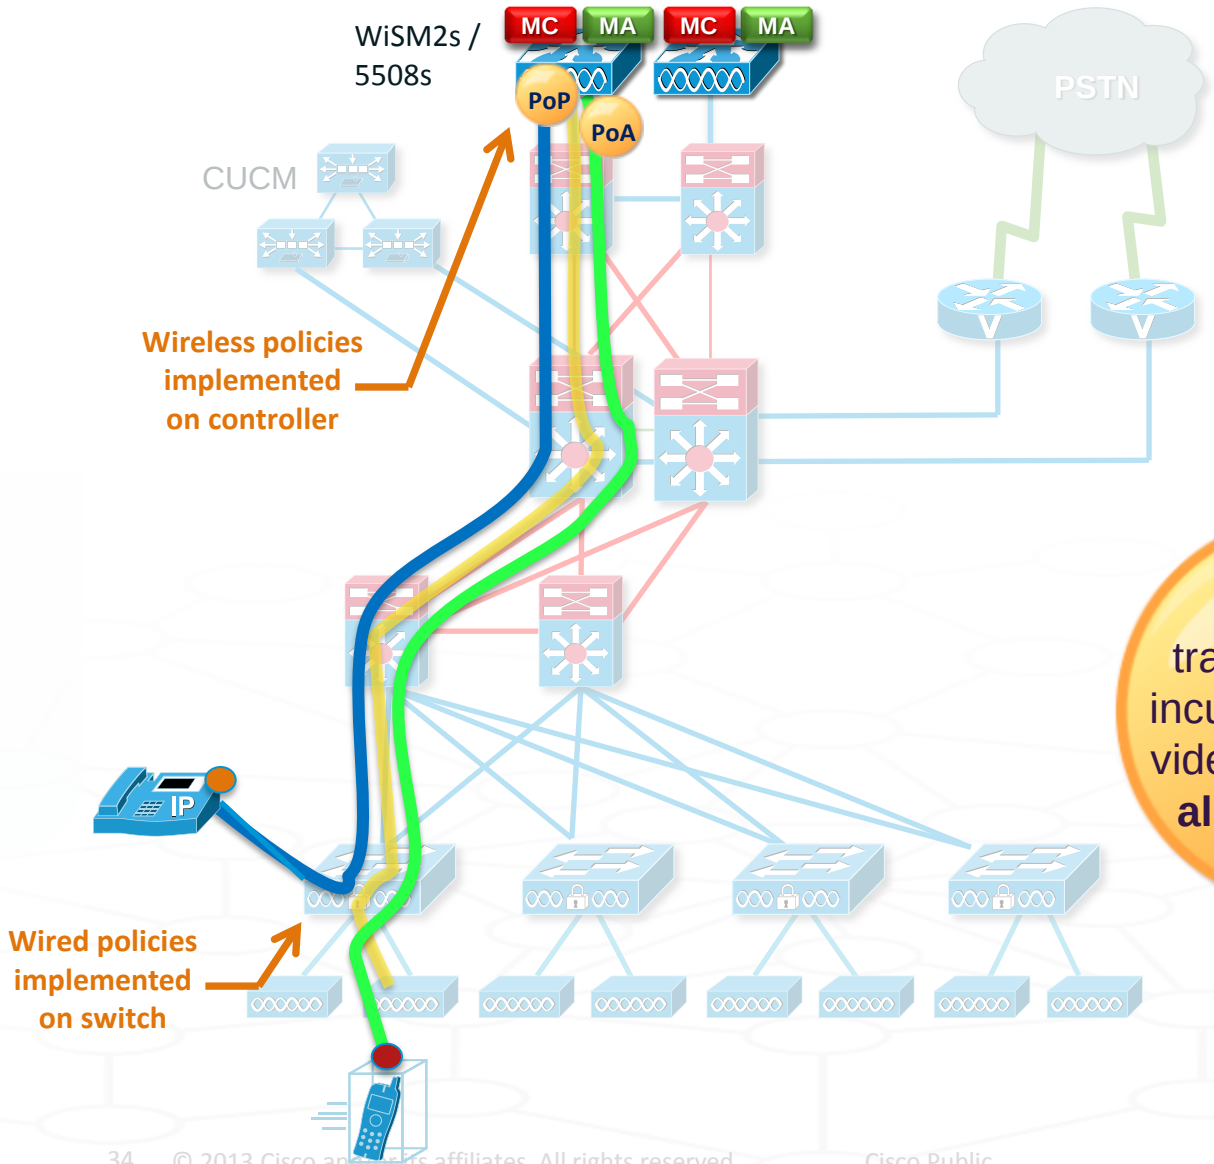

Separate policies and services for wired and wireless users

The same traffic paths are incurred for voice, video, data, etc. – **all centralised**

Traffic Flows, Unified Wireless –

- In this example, a VoIP user is on today's CUWN network, and is making a call from a wireless handset to a wired handset …
- **We can see that all of the user's traffic needs to be hairpinned back through the centralised controller,** in both directions …

In this example, a total of **9 hops** are incurred for each direction of the traffic path (including the controllers – Layer 3 roaming might add more hops) …

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a "double-click" deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – **Terminology and Building Blocks**
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

Converged Access
Deployment Overview

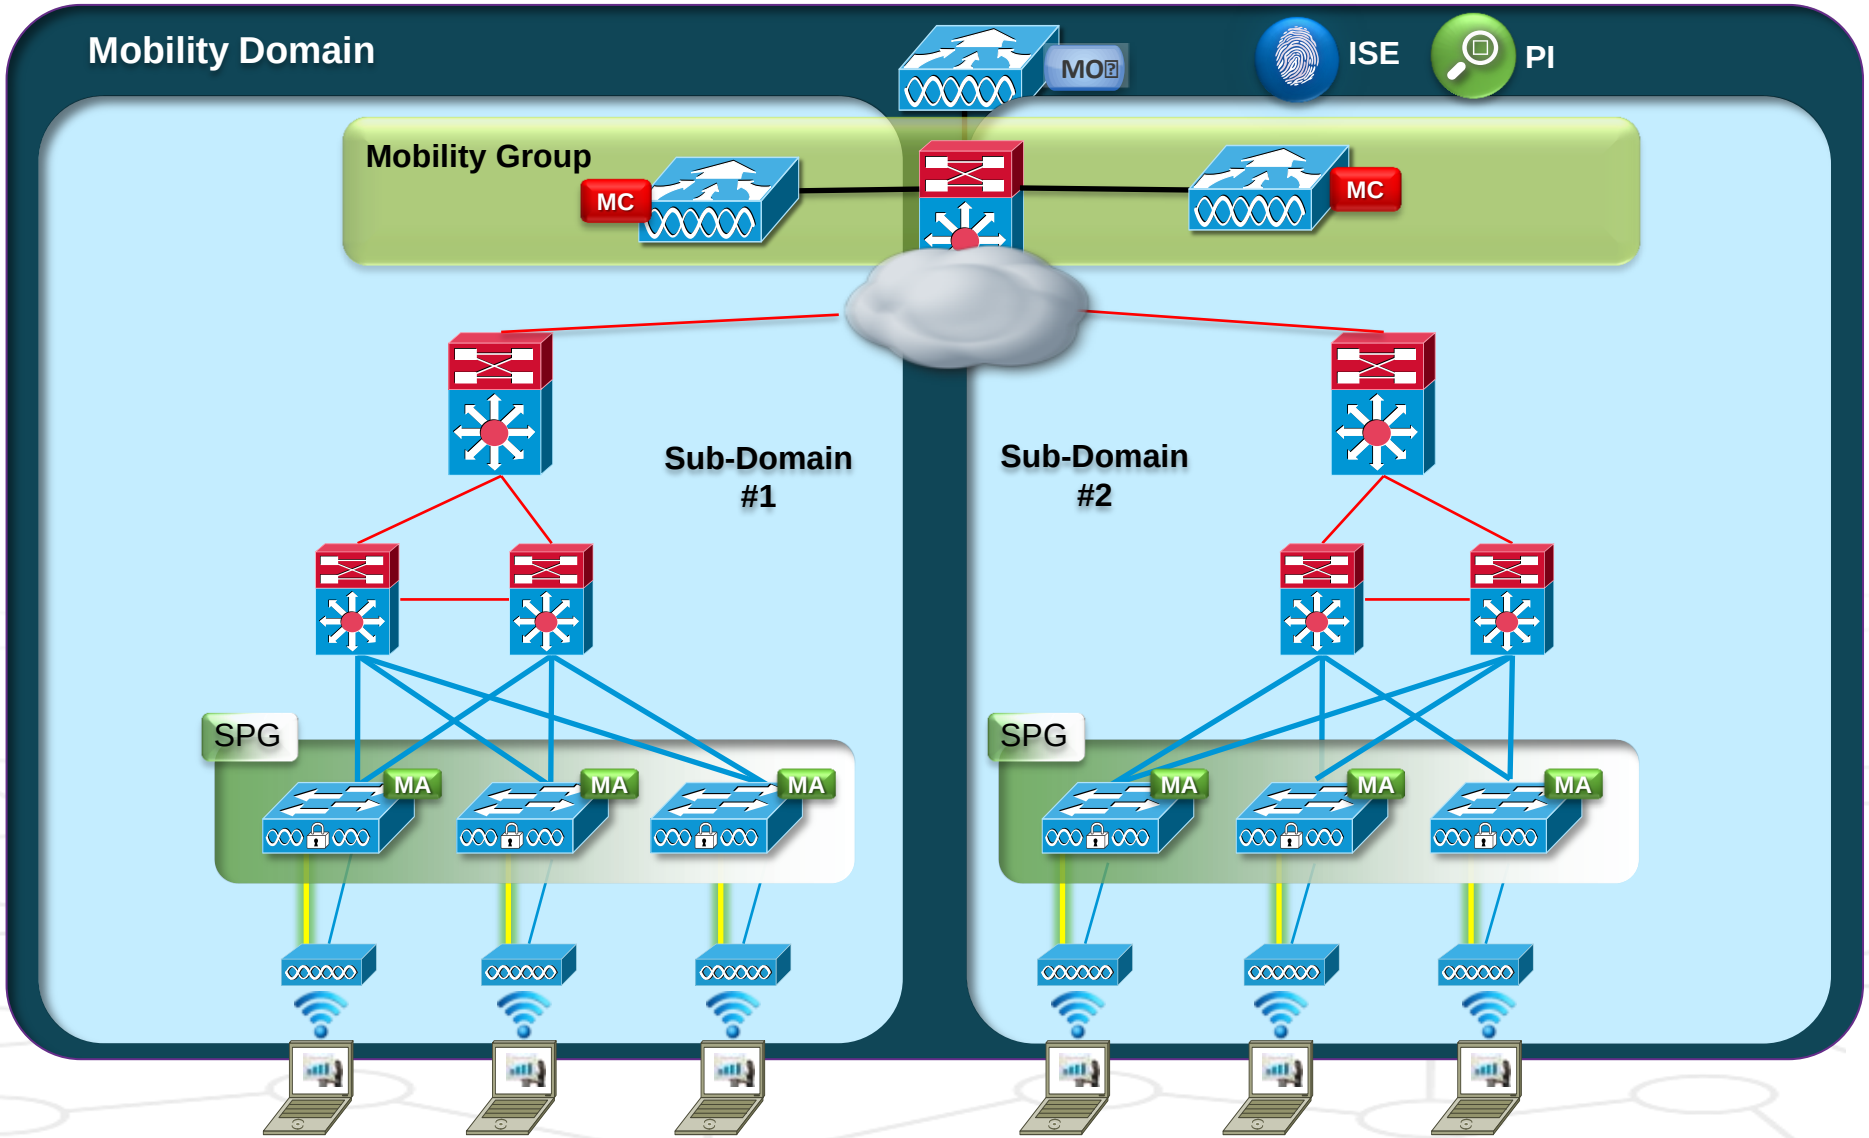

Converged Access

Components – Physical vs. Logical Entities

Physical Entities –

- **Mobility Agent (MA)** – Terminates CAPWAP tunnel from AP
- **Mobility Controller (MC)** – Manages mobility within and across Sub-Domains
- **Mobility Oracle (MO)** – Superset of MC, allows for Scalable Mobility Management within a Domain

Logical Entities –

- **Mobility Groups –** Grouping of Mobility Controllers (MCs) to enable Fast Roaming, Radio Frequency Management, etc.
- **Mobility Domain** – Grouping of MCs to support seamless roaming
- **Switch Peer Group (SPG) –** Localises traffic for roams within Distribution Block

MA, MC, Mobility Group functionality all exist in today's controllers (4400, 5500, WiSM2)

Converged Access

Physical Entities – Mobility Agents (MAs)

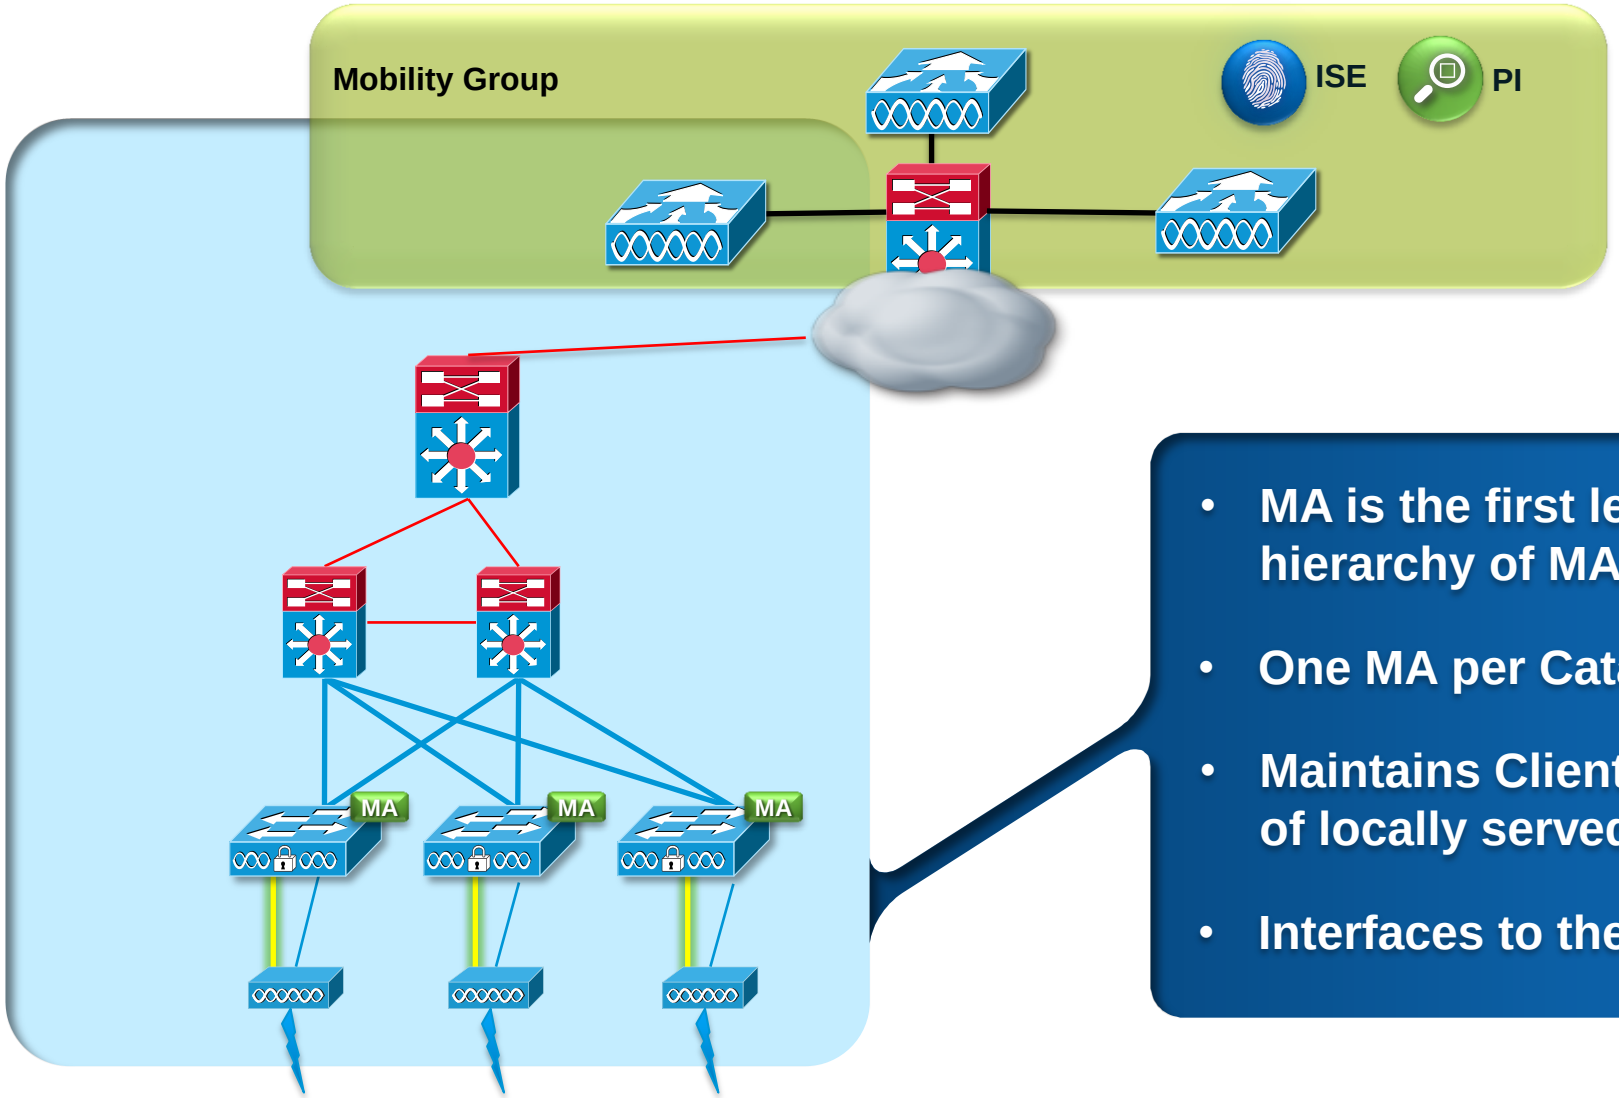

- MA is the first level in the hierarchy of MA / MC / MO
- One MA per Catalyst 3850 Stack
- Maintains Client DB of locally served clients
- Interfaces to the Mobility Controller (MC)

Converged Access

Physical Entities – Mobility Controllers (MCs)

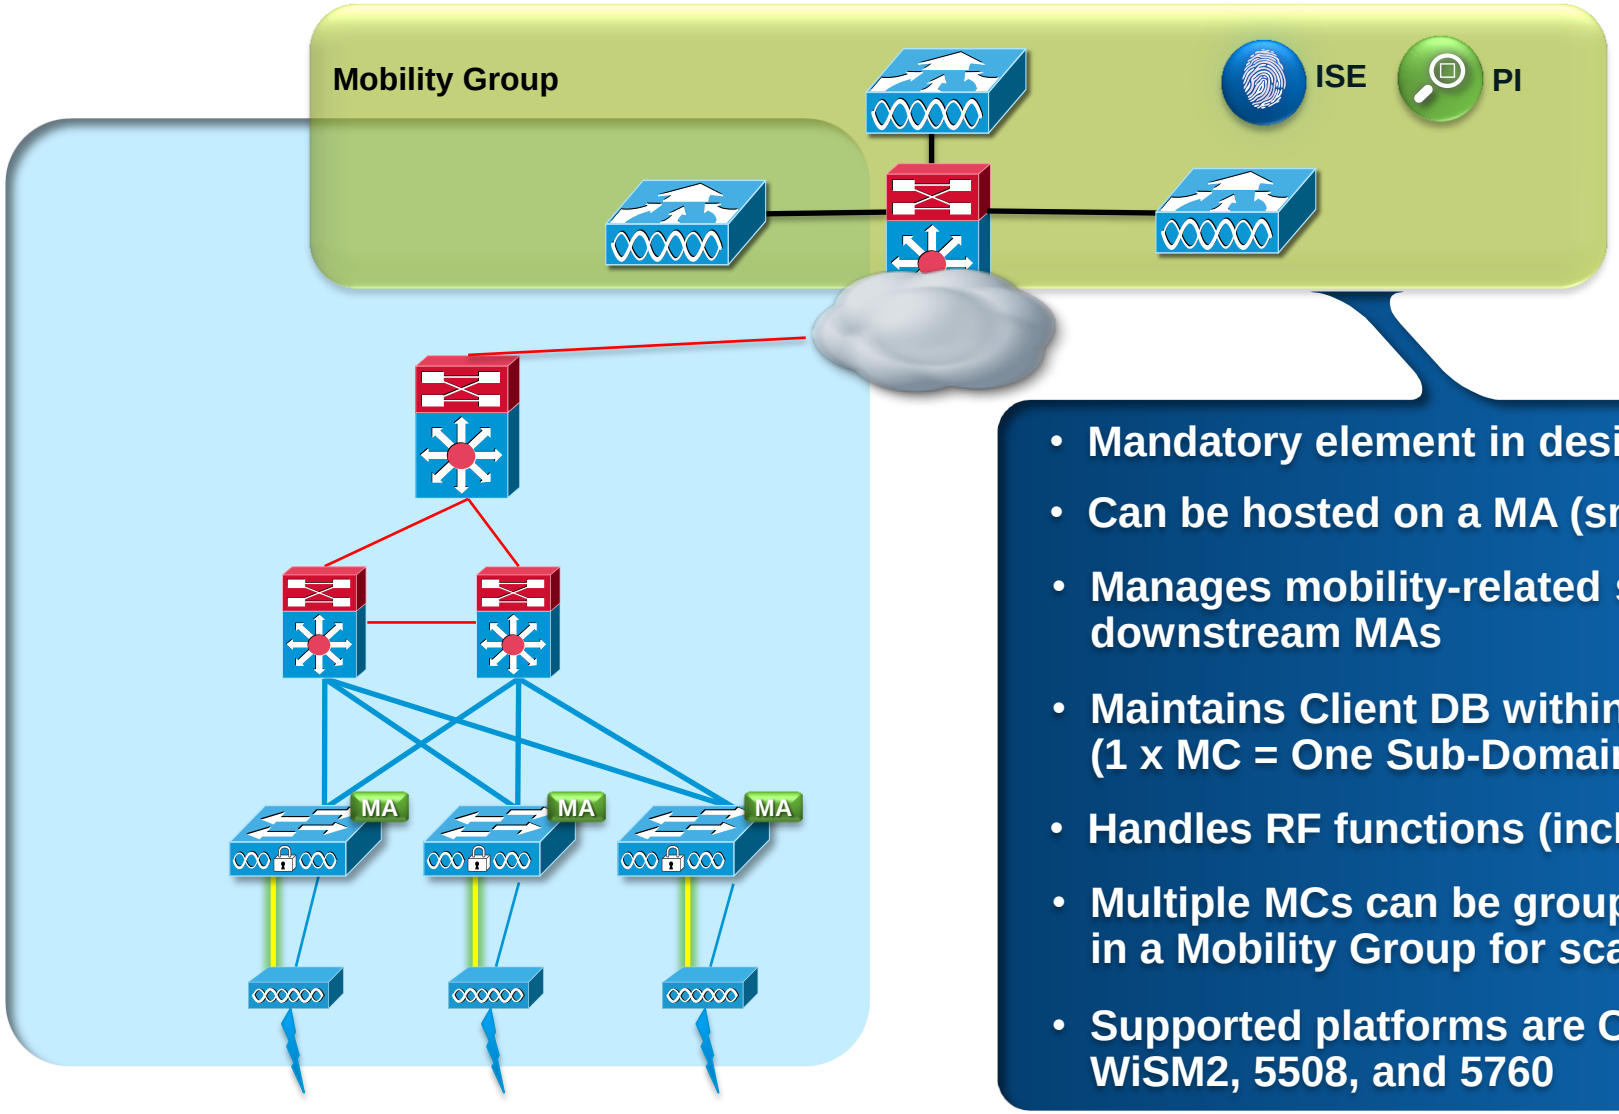

- Mandatory element in design
- Can be hosted on a MA (smaller deployments)
- Manages mobility-related state of the downstream MAs
- Maintains Client DB within a Sub-Domain (1 x MC = One Sub-Domain)
- Handles RF functions (including RRM)
- Multiple MCs can be grouped together in a Mobility Group for scalability
- Supported platforms are Catalyst 3850, WiSM2, 5508, and 5760

Converged Access

Physical Entities – Catalyst 3850 Switch Stack

Best-in-Class Wired Switch – with Integrated Wireless Mobility functionality

MA

- Can act as a **Mobility Agent** (MA)
  for terminating CAPWAP tunnels for locally connected APs …

MC

- as well as a **Mobility Controller** (MC)
  for other Mobility Agent (MA) switches, in small deployments

  - MA/MC functionality works on a Stack of Catalyst 3850 Switches
  - MA/MC functionality runs on Stack Master
  - Stack Standby synchronises some information (useful for intra-stack HA)

Converged Access

Logical Entities – Switch Peer Groups

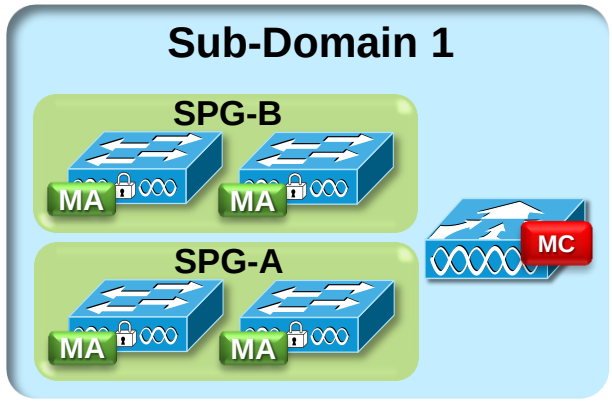

- **Made up of multiple Catalyst 3850 switches as Mobility Agents** (MAs), **plus an MC** (on controller as shown)
- **Handles roaming across SPG** (L2 / L3)
- **MAs within an SPG are fully-meshed** (auto-created at SPG formation)
- **Fast Roaming within an SPG**
- **Multiple SPGs under the control of a single MC form a Sub-Domain**

SPGs are a logical construct, not a physical one …

SPGs can be formed across Layer 2 or Layer 3 boundaries

SPGs are designed to constrain roaming traffic to a smaller area, and optimise roaming capabilities and performance

Current thinking on best practices dictates that **SPGs will likely be built around buildings, around floors within a building, or other areas that users are likely to roam most within**

Roamed traffic within an SPG moves directly between the MAs in that SPG (CAPWAP full mesh)

Roamed traffic between SPGs moves via the MC(s) servicing those SPGs

Hierarchical architecture is optimised for scalability and roaming

Converged Access

Logical Entities – Switch Peer Groups and Mobility Group

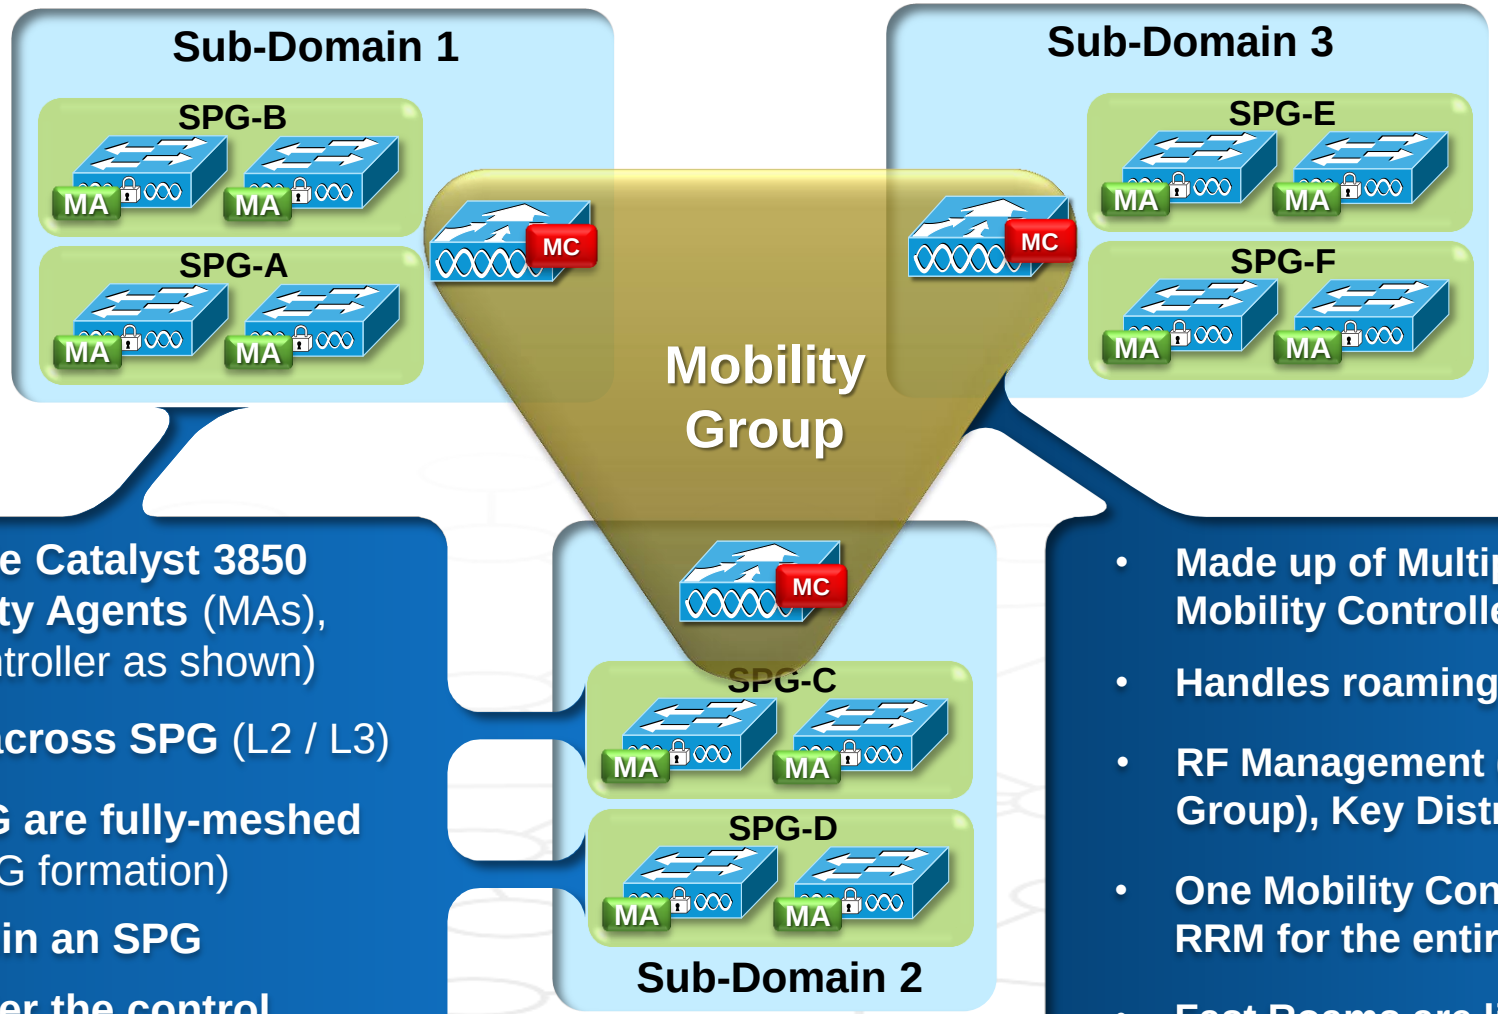

- **Made up of multiple Catalyst 3850 switches as Mobility Agents** (MAs), **plus an MC** (on controller as shown)
- **Handles roaming across SPG** (L2 / L3)
- **MAs within an SPG are fully-meshed** (auto-created at SPG formation)
- **Fast Roaming within an SPG**
- **Multiple SPGs under the control of a single MC form a Sub-Domain**

- **Made up of Multiple Mobility Controllers (MCs)**
- **Handles roaming across MG** (L2 / L3)
- **RF Management (RRM, handled by RF Group), Key Distribution for Fast Roaming**
- **One Mobility Controller** (MC) **manages RRM for the entire RF Group**
- **Fast Roams are limited to Mobility Group member MCs**

Converged Access

Scalability Considerations

As with any solution – there are scalability constraints to be aware of …

- **These are summarised below, for quick reference**
- **Full details on scalability – for both CUWN as well as Converged Access deployments – is located in the Reference section at the end of this slide deck**

Scalability	3850 as MC	5760	5508	WiSM2
Max number of MCs in a Mobility Domain	8	72	72	72
Max number of MCs in a Mobility Group	8	24	24	24
Max number of MAs in a Sub-domain (per MC)	16	350	350	350
Max number of SPGs in a Mobility Sub-Domain (per MC)	8	24	24	24
Max number of MAs in a SPG	16	64	64	64
Max number of WLANs	64	512	512	512

Converged Access

Physical Entities – Mobility Oracle (MO)

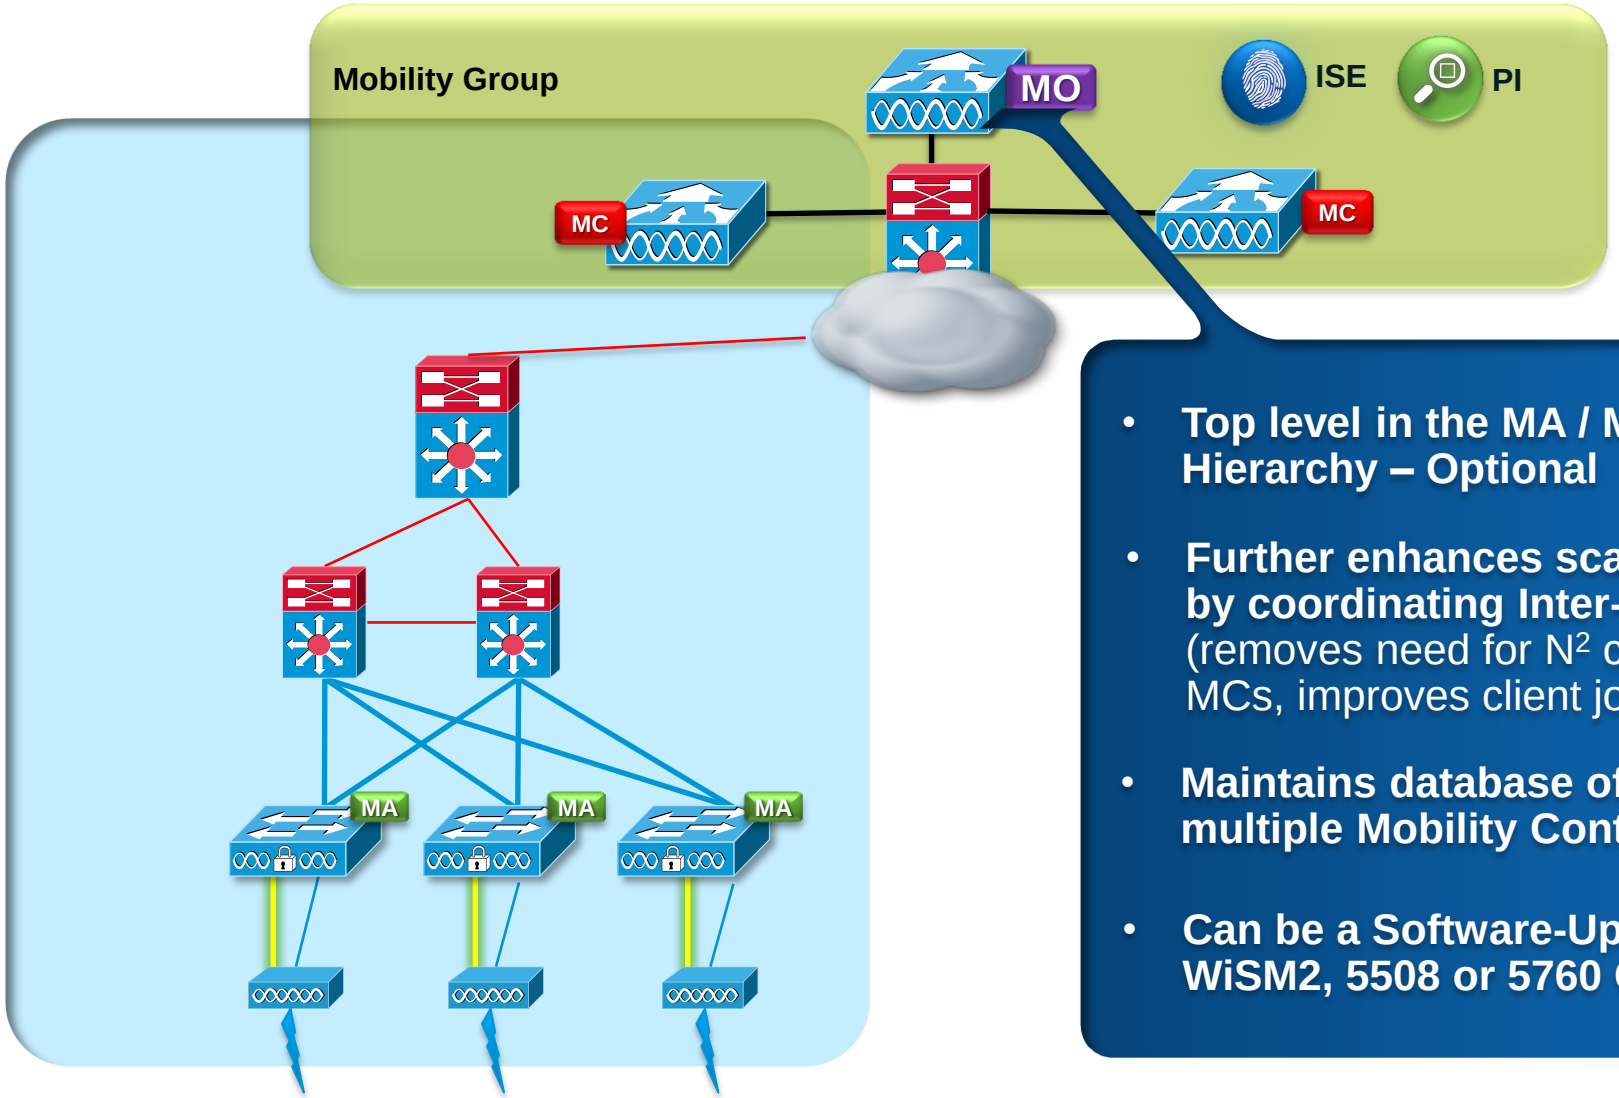

- **Top level in the MA / MC / MO Hierarchy – Optional**
- **Further enhances scalability and performance by coordinating Inter-MC roams** (removes need for N^2 communications between MCs, improves client join performance)
- **Maintains database of clients across multiple Mobility Controllers (MCs)**
- **Can be a Software-Upgraded WiSM2, 5508 or 5760 Controller**

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a "double-click" deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – **Traffic Flows and Roaming**
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

Converged Access

Roaming – Point of Presence (PoP), Point of Attachment (PoA)

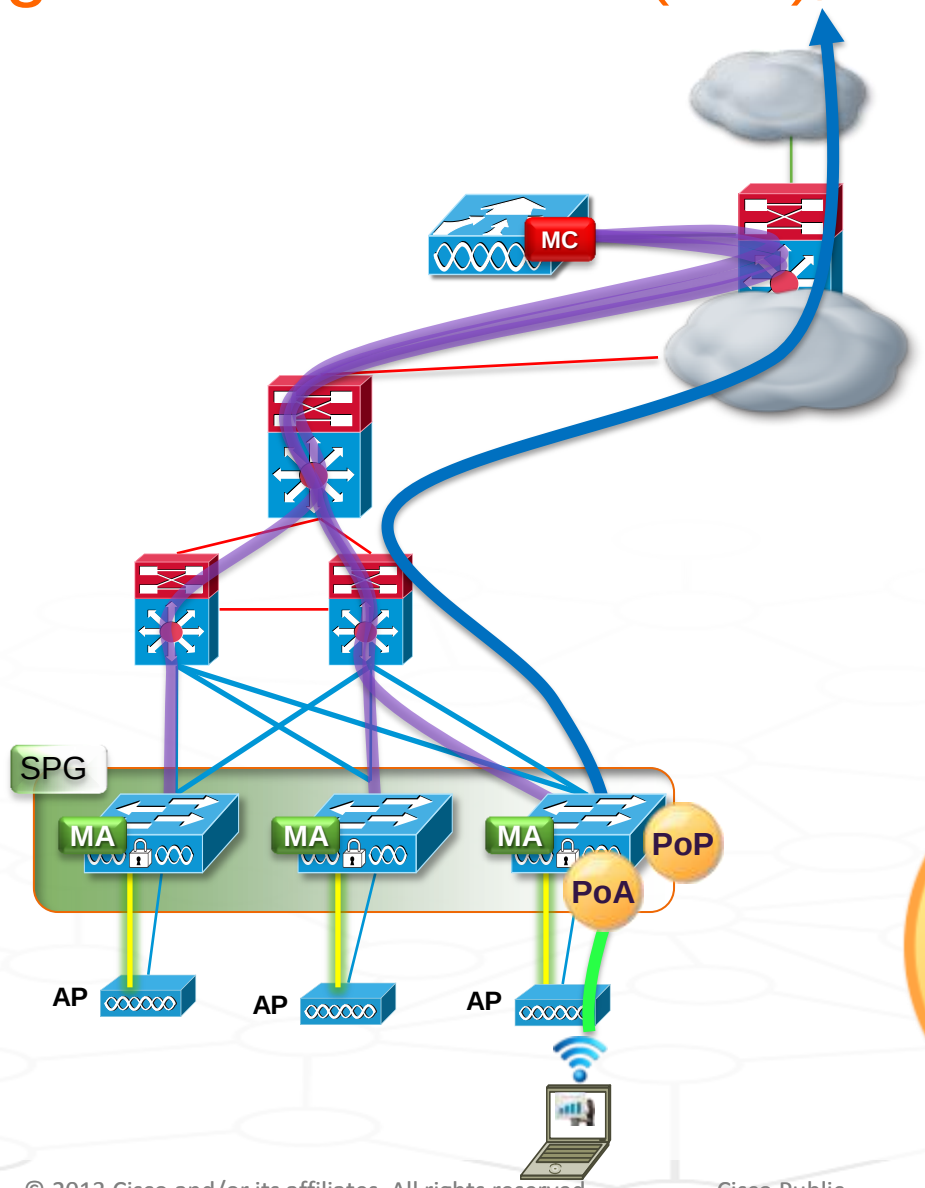

If users associate and remain stationary, this is their traffic flow

Point of Presence (PoP) vs. Point of Attachment (PoA) –

- **PoP is where the wireless user is seen to be** within the wired portion of the network
- **PoA is where the wireless user has roamed to** while mobile
- **Before a user roams, PoP and PoA are in the same place**

Note – for the purposes of illustrating roaming, we are showing the purple connections herein that indicate the connections between the MAs and their corresponding MC for the Switch Peer Group (or Groups) involved on each slide … notice that, in this example, **the traffic does NOT flow through the MC …**

Converged Access

Traffic Flow and Roaming – Branch, Single Catalyst 3850 Stack

Notice how the 3850 switch stack shown is an MC (as well as an MA) – in a branch such as this with 50 APs or less, no discrete controller is necessarily required …

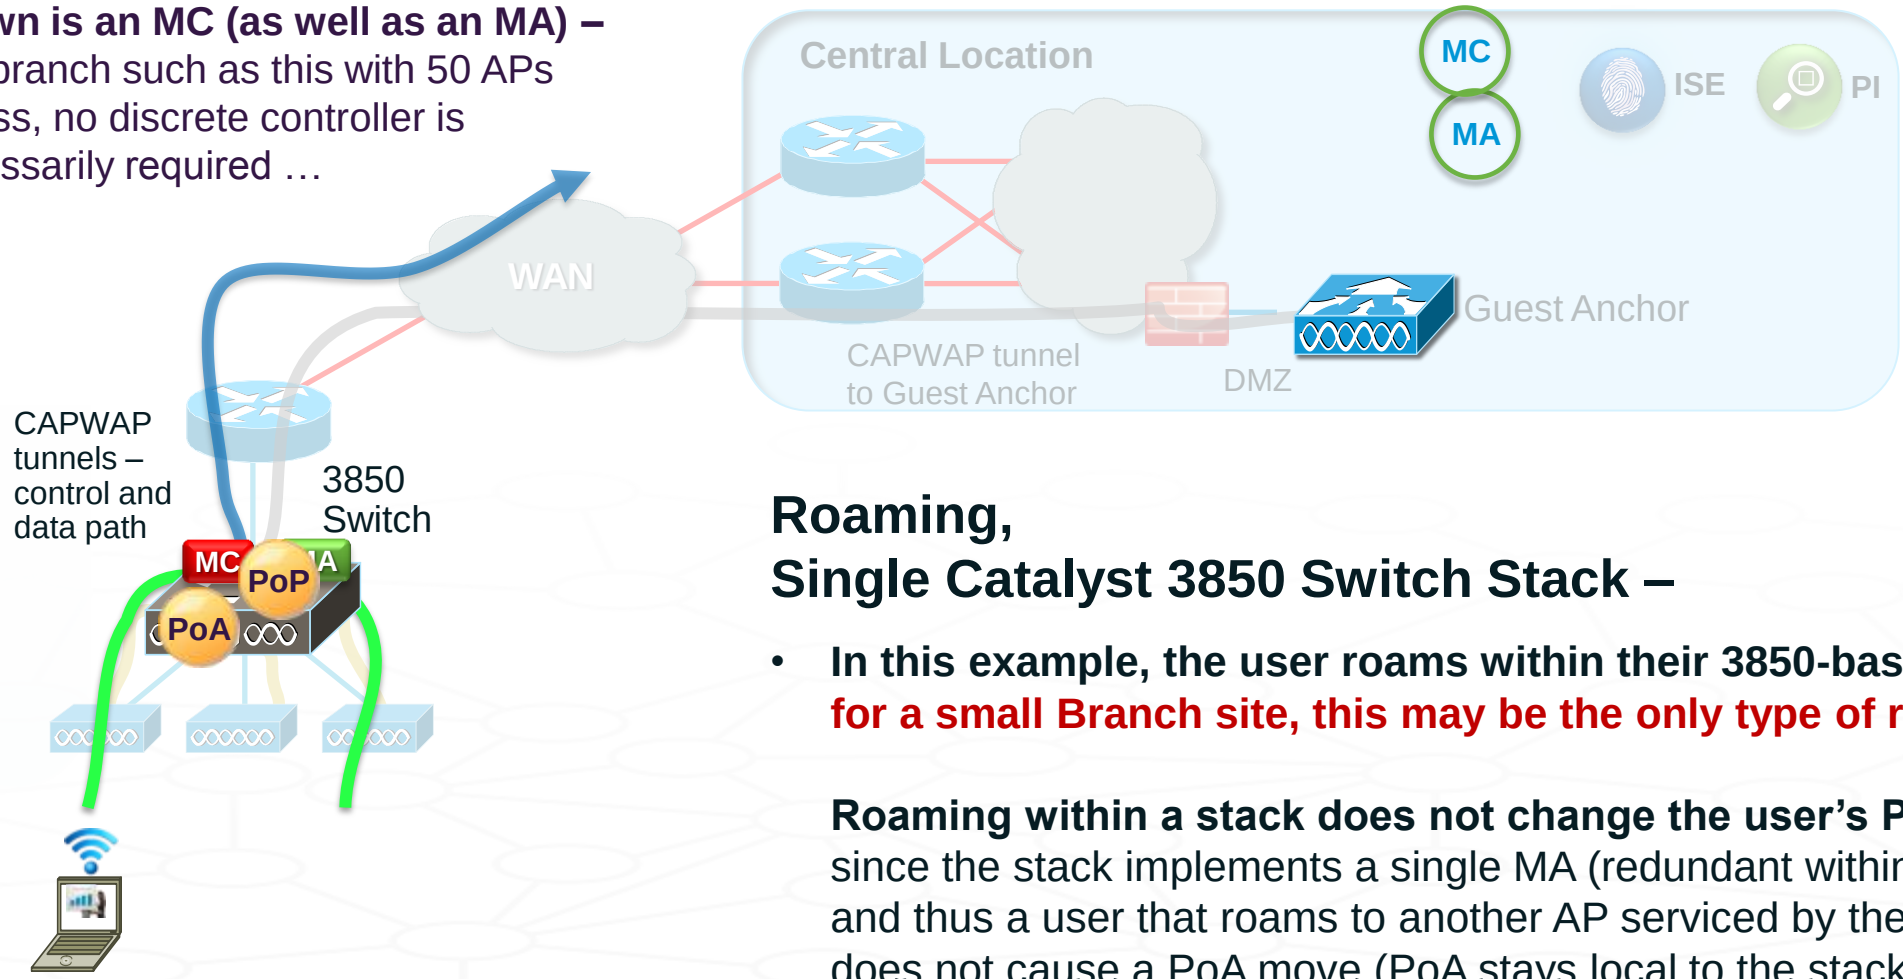

Roaming across Stack (small branch)

Roaming,
Single Catalyst 3850 Switch Stack –

- **In this example, the user roams within their 3850-based switch stack – for a small Branch site, this may be the only type of roam**

Roaming within a stack does not change the user's PoP or PoA – since the stack implements a single MA (redundant within the stack), and thus a user that roams to another AP serviced by the same stack does not cause a PoA move (PoA stays local to the stack)

Converged Access

Traffic Flow and Roaming – Branch, L2 / L3 Roam (within SPG)

Roaming across Stacks (larger branch)

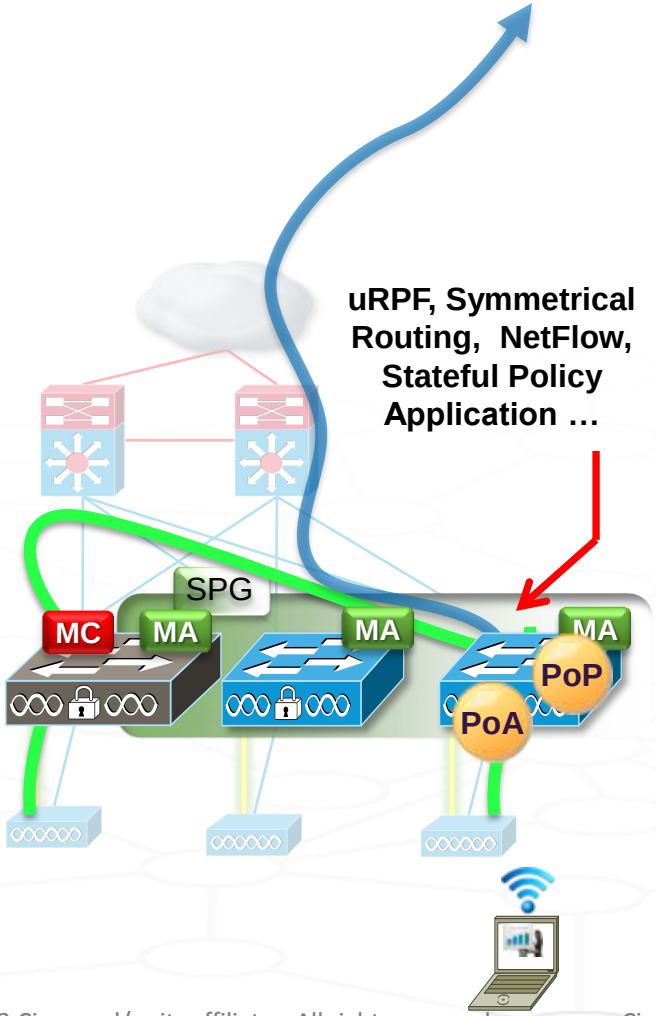

Roaming, Within a Switch Peer Group (Branch) –

- Now, let's examine a roam at a larger branch, with multiple 3850-based switch stacks joined together via a distribution layer
- **In this example, the larger Branch site consists of a single Switch Peer Group – and the user roams within that SPG – again, at a larger Branch such as this, this may be the only type of roam**

The user may or may not have roamed across an L3 boundary (depends on wired setup) – **however, users are always* taken back to their PoP** for policy application

Again, notice how the 3850 switch stack on the left is an MC (as well as an MA) in this picture – in a larger branch such as this with 50 APs or less, no discrete controller is necessarily required …

** Adjustable via setting, may be useful for L2 roams (detailed on slides in following section of this slide deck)*

Converged A

Traffic Flow and

```
L09-3850s-3# show wireless client summary
Number of Local Clients : 1

MAC Address    AP Name                                WLAN State                 Protocol
-------------------------------------------------------------------------------------------
001e.65b7.7d1a L09-AP1142-1                           2    UP                    11n(5)

L09-3850s-3# show wcdb database all
        Total Number of Wireless Clients = 1
                Local Clients                = 1

Mac Address    VlanId IP Address      Src If              Auth      Mob
-------------- ------ --------------- ------------------- --------- -------
001e.65b7.7d1a 2003   10.101.203.1    0x009350C0000000E4  RUN       LOCAL

L09-3850s-3# show wireless client mac 001e.65b7.7d1a detail

Client MAC Address : 001e.65b7.7d1a
Mobility State : Local
```

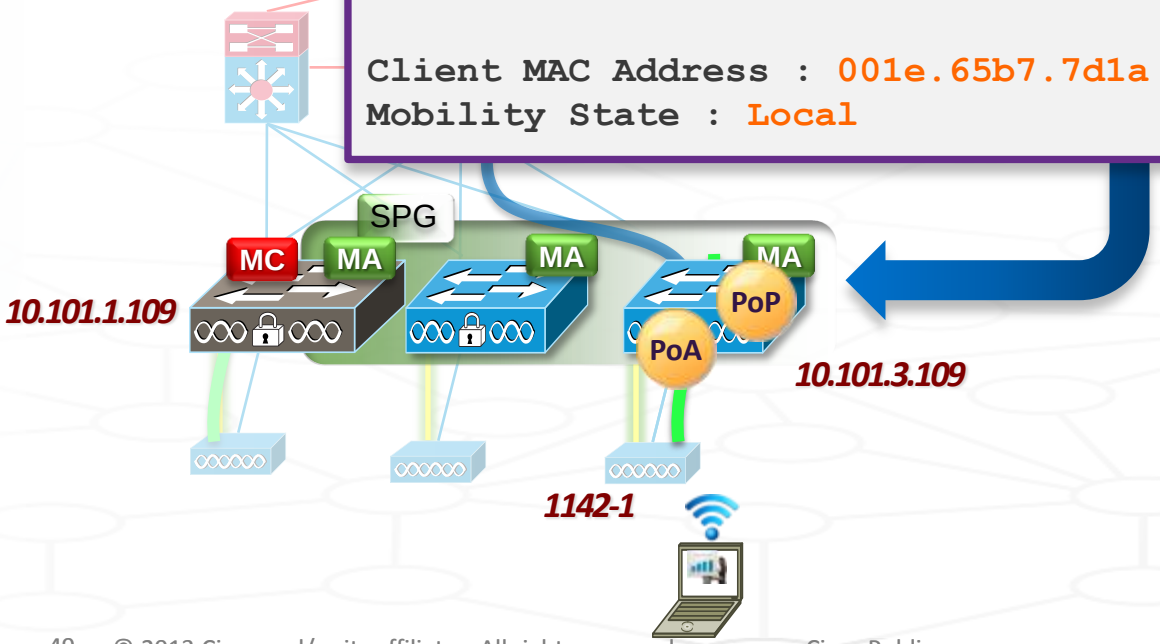

Converged A

Traffic Flow and

```
L09-3850s-3# show capwap summary

Name   APName                           Type PhyPortIf Mode      McastIf
------ -------------------------------- ---- --------- --------- -------
Ca1    -                                mob  -         unicast   -
Ca2    -                                mob  -         unicast   -
Ca5    L09-AP1142-1                     data Gi1/0/7   multicast Ca4
```

```
Name   SrcIP           SrcPort DestIP          DstPort DtlsEn MTU
------ --------------- ------- --------------- ------- ------ -----
Ca1    10.101.3.109    16667   10.101.1.109    16667   No     1464
Ca2    10.101.3.109    16667   10.101.2.109    16667   No     1464
Ca5    10.101.3.109    5247    10.101.3.98     31901   No     1449
```

```
L09-3850s-3# show ip dhcp snooping binding
MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
00:1E:65:b7:7d:1a   10.101.203.1     10617       dhcp-snooping  2003  Capwap5
```

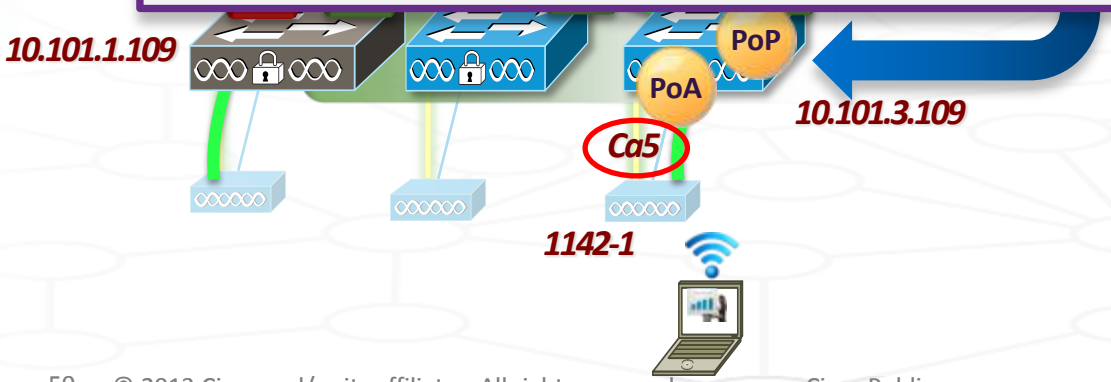


```
L09-3850s-3# show mac address dynamic | inc Ca
2003        001e.65b7.7d1a    DYNAMIC     Ca5
```

```
L09-3850-1# show wireless client summary

MAC Address    AP Name                                WLAN State                Protocol
-------------------------------------------------------------------------------------
001e.65b7.7d1a 3600-1                                 2    UP                   11n(5)


L09-3850-1# show wcdb database all
        Total Number of Wireless Clients = 1
               Foreign Clients           = 1

Mac Address    VlanId IP Address       Src If              Auth      Mob
-------------- ------ ---------------- ------------------- --------- -------
001e.65b7.7d1a 2001   10.101.203.1     0x00C55A40000000A6 RUN       FOREIGN
```

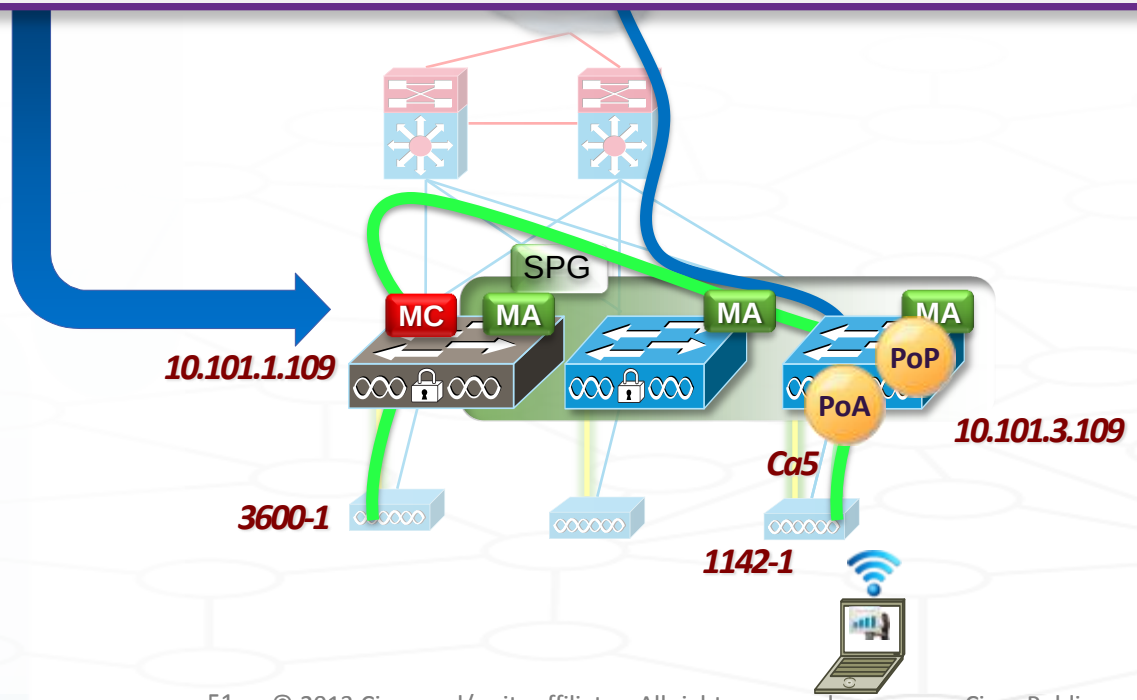

Converged Access

Traffic Flow and Roaming – Branch, L2 / L3 Roam (within SPG)

```
L09-3850-1# show wireless client mac 001e.65b7.7d1a detail
...
Client MAC Address : 001e.65b7.7d1a
Mobility State : Foreign
Mobility Anchor IP Address : 10.101.3.109
```

```
L09-3850-1# show mac address dynamic | inc Ca
4095    001e.65b7.7d1a     DYNAMIC     Ca3
```

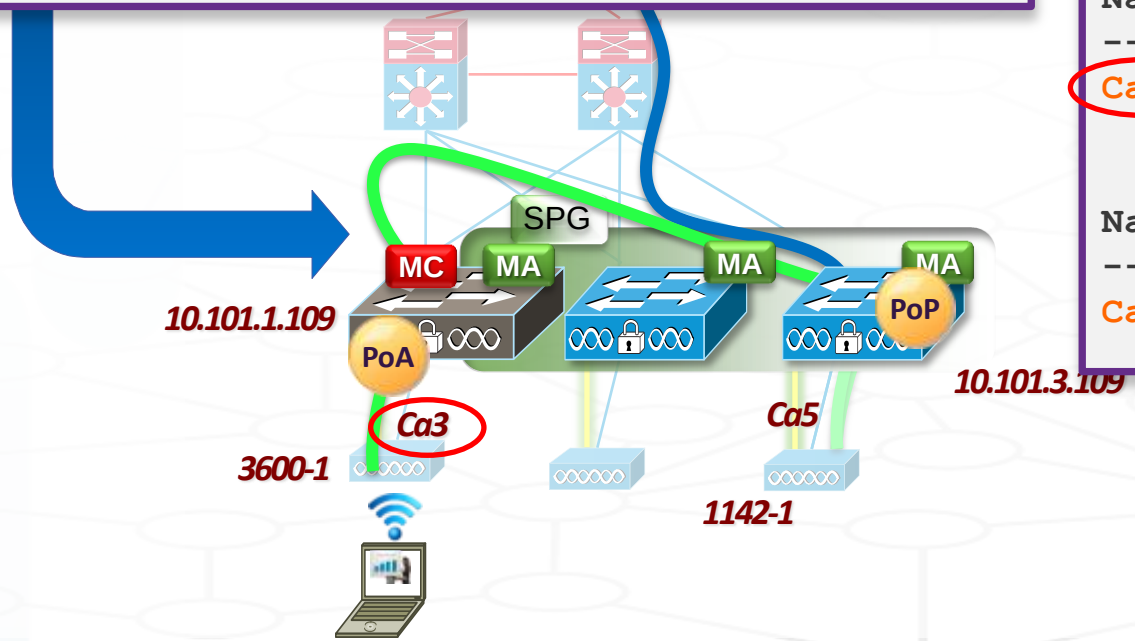


```
L09-3850-1# show capwap summary

Name    APName                 Type PhyPortIf Mode       McastIf
------ -------------------- ---- --------- --------- -------
Ca3     3600-1                 data Gi1/0/9   multicast Ca1


Name    SrcIP         SrcPort DestIP        DstPort DtlsEn MTU
----- ------------ ------- ------------ ------- ------ -----
Ca3     10.101.1.109 5247    10.101.1.98   16370   No     1449
```

Conve
Traffic F

```
L09-3850s-3# show wireless client summary
Number of Local Clients : 1

MAC Address    AP Name                                WLAN State                 Protocol
-------------------------------------------------------------------------------------
001e.65b7.7d1a 10.101.1.109                           2    UP                    Mobile
```

```
L09-3850s-3# show wcdb database all

Mac Address    VlanId IP Address      Src If              Auth      Mob
-------------- ------ --------------- ------------------- --------- -------
001e.65b7.7d1a  2003 10.101.203.1     0x00B72D4000000002 RUN       ANCHOR
```

```
L09-3850s-3# show wireless client mac 001e.65b7.7d1a detail
...
Client MAC Address : 001e.65b7.7d1a
Mobility State : Anchor
Mobility Foreign IP Address : 10.101.1.109
```

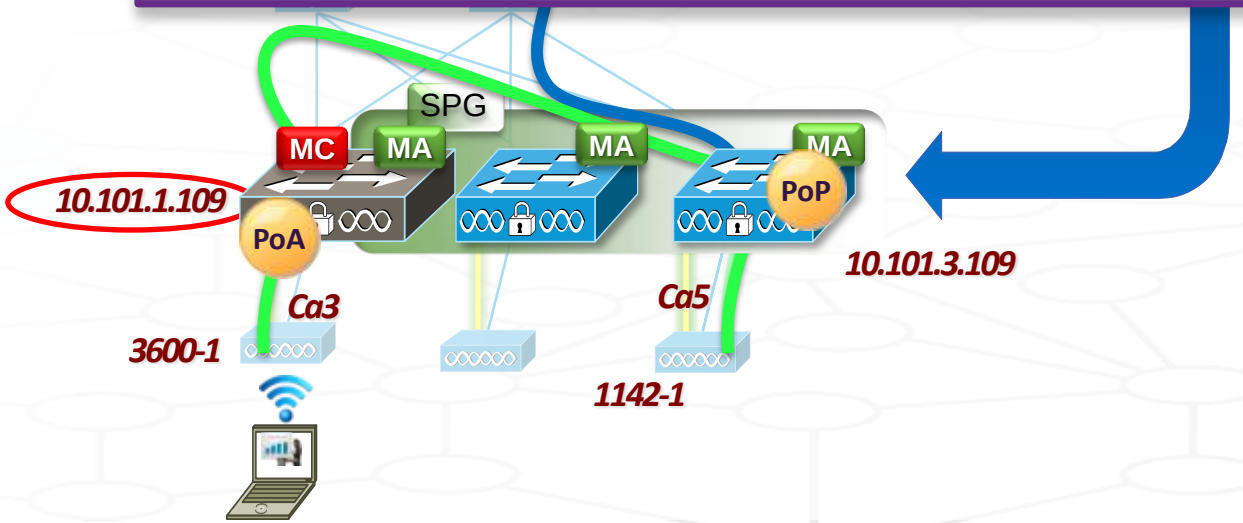

Converg
Traffic Flow

```
L09-3850s-3# show ip dhcp snooping binding
MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
00:1E:65:b7:7d:1a   10.101.203.1     10720       dhcp-snooping   2003  Capwap1
```

```
L09-3850s-3# show capwap summary

Name   APName                            Type PhyPortIf Mode       McastIf
------ -------------------------------- ---- --------- ---------- -------
Ca1    -                                mob  -         unicast    -

Name   SrcIP           SrcPort DestIP          DstPort DtlsEn MTU
------ --------------- ------- --------------- ------- ------ -----
Ca1    10.101.3.109    16667   10.101.1.109    16667   No     1464
```

```
L09-3850s-3# show mac address dynamic | inc Ca
2003    001e.65b7.7d1a    DYNAMIC     Ca1
```

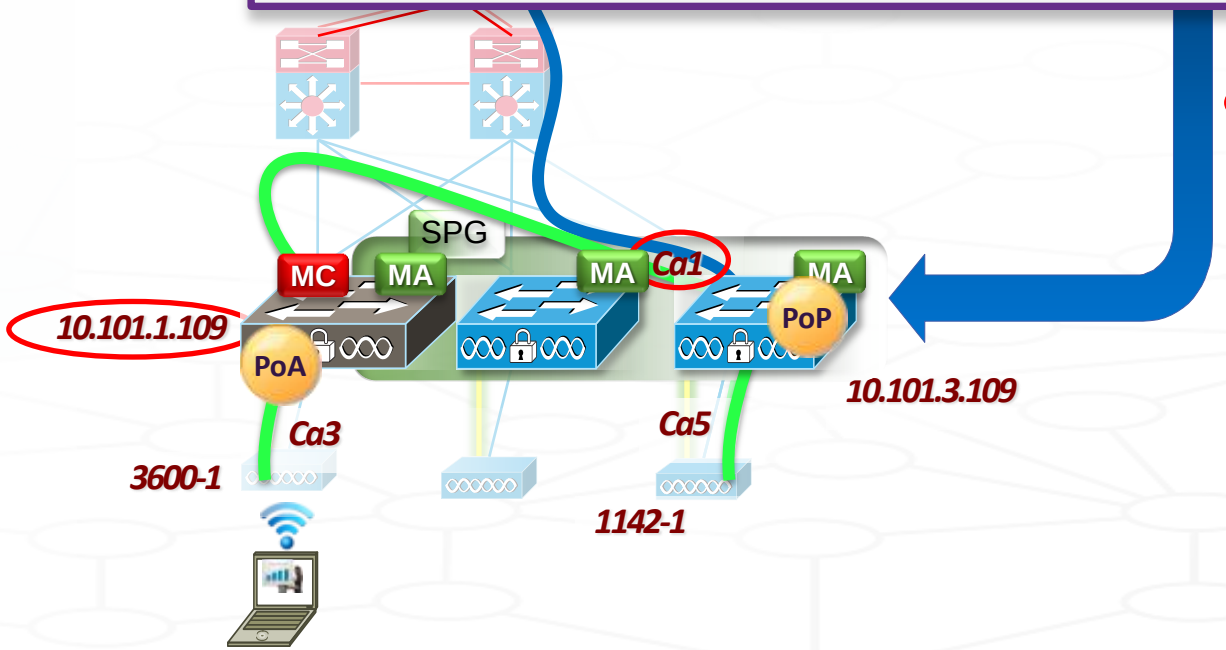

Overall observation –

This looks exactly the same as a Layer 3 inter-controller roam in CUWN …
because it is exactly the same process –
Just distributed, rather than centralised …

Converged Access

Traffic Flow and Roaming – Campus L2 / L3 Roam (within SPG)

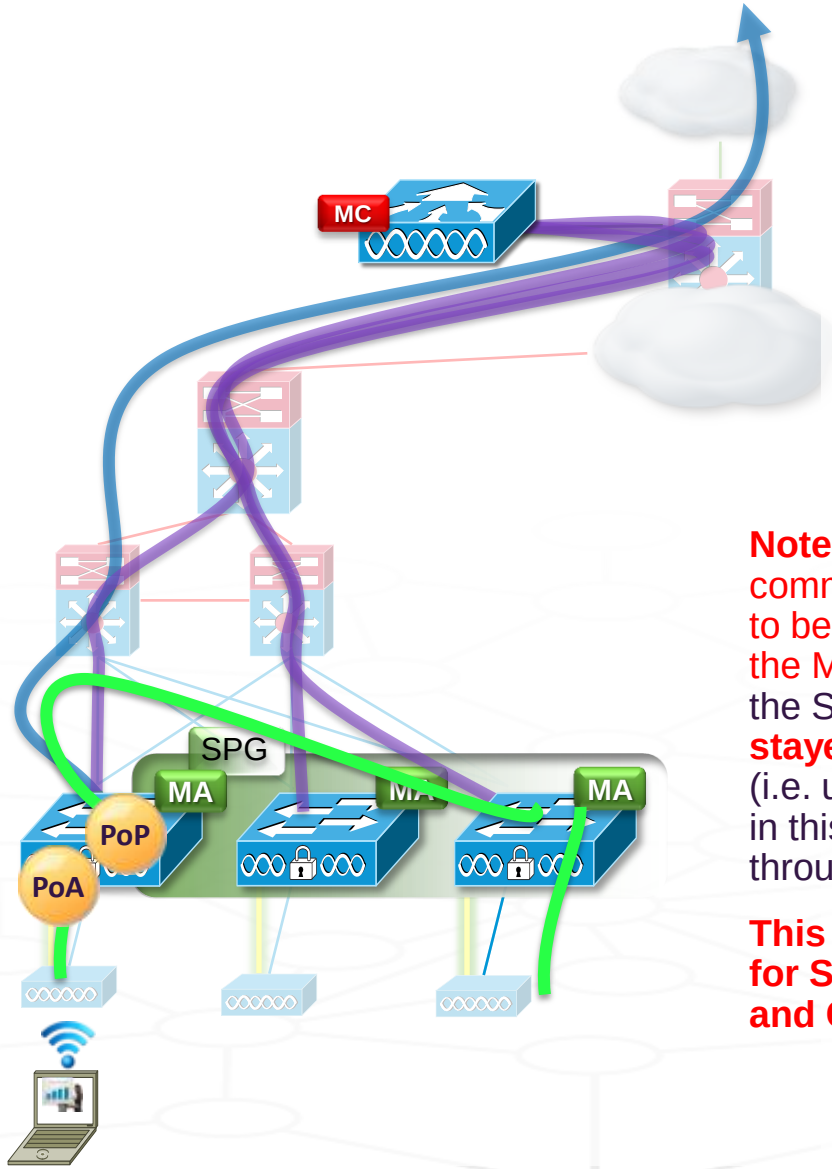

Roaming within an SPG

(L3 behaviour and default L2 behaviour)

Note – the traffic in this most common type of roam did **not** have to be transported back to, or via, the MC (controller) servicing the Switch Peer Group – **traffic stayed local to the SPG only** (i.e. under the distribution layer in this example – not back through the core).

This is an important consideration for Switch Peer Group, traffic flow, and Controller scalability.

Roaming, Within an SPG (Campus) –

- Now, let's examine a few more types of user roams
- **In this example, the user roams within their Switch Peer Group –** since SPGs are typically formed around floors or other geographically-close areas, **this is the most likely and most common type of roam**

The user may or may not have roamed across an L3 boundary (depends on wired setup) – **however, users are always* taken back to their PoP** for policy application

Converged Access
Traffic Flow

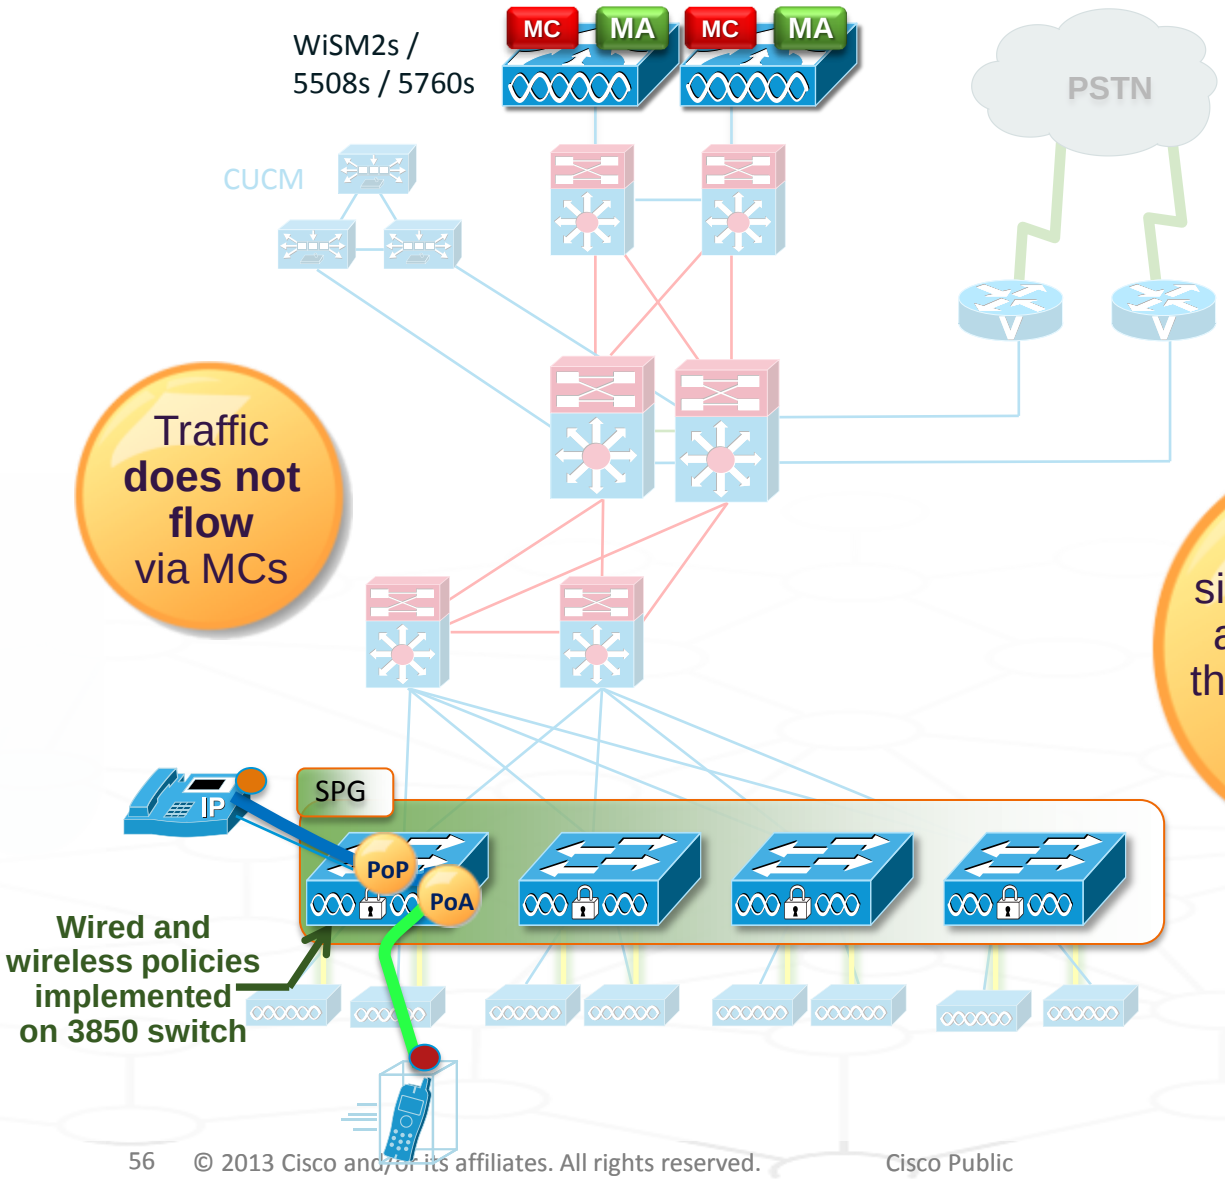

Converged policies and services for wired and wireless users

Traffic Flows, Comparison (Converged Access) –

- Now, our VoIP user is on a Cisco Converged Access network, and is again making a call from a wireless handset to a wired handset …
- **We can see that all of the user's traffic is localised to their Peer Group, below the distribution layer,** in both directions …

  In this example, a total of **1 hop** is incurred for each direction of the traffic path (assuming no roaming) … two additional hops may be incurred for routing …

Converged Access

Traffic Flow – with Intra-SPG Roam

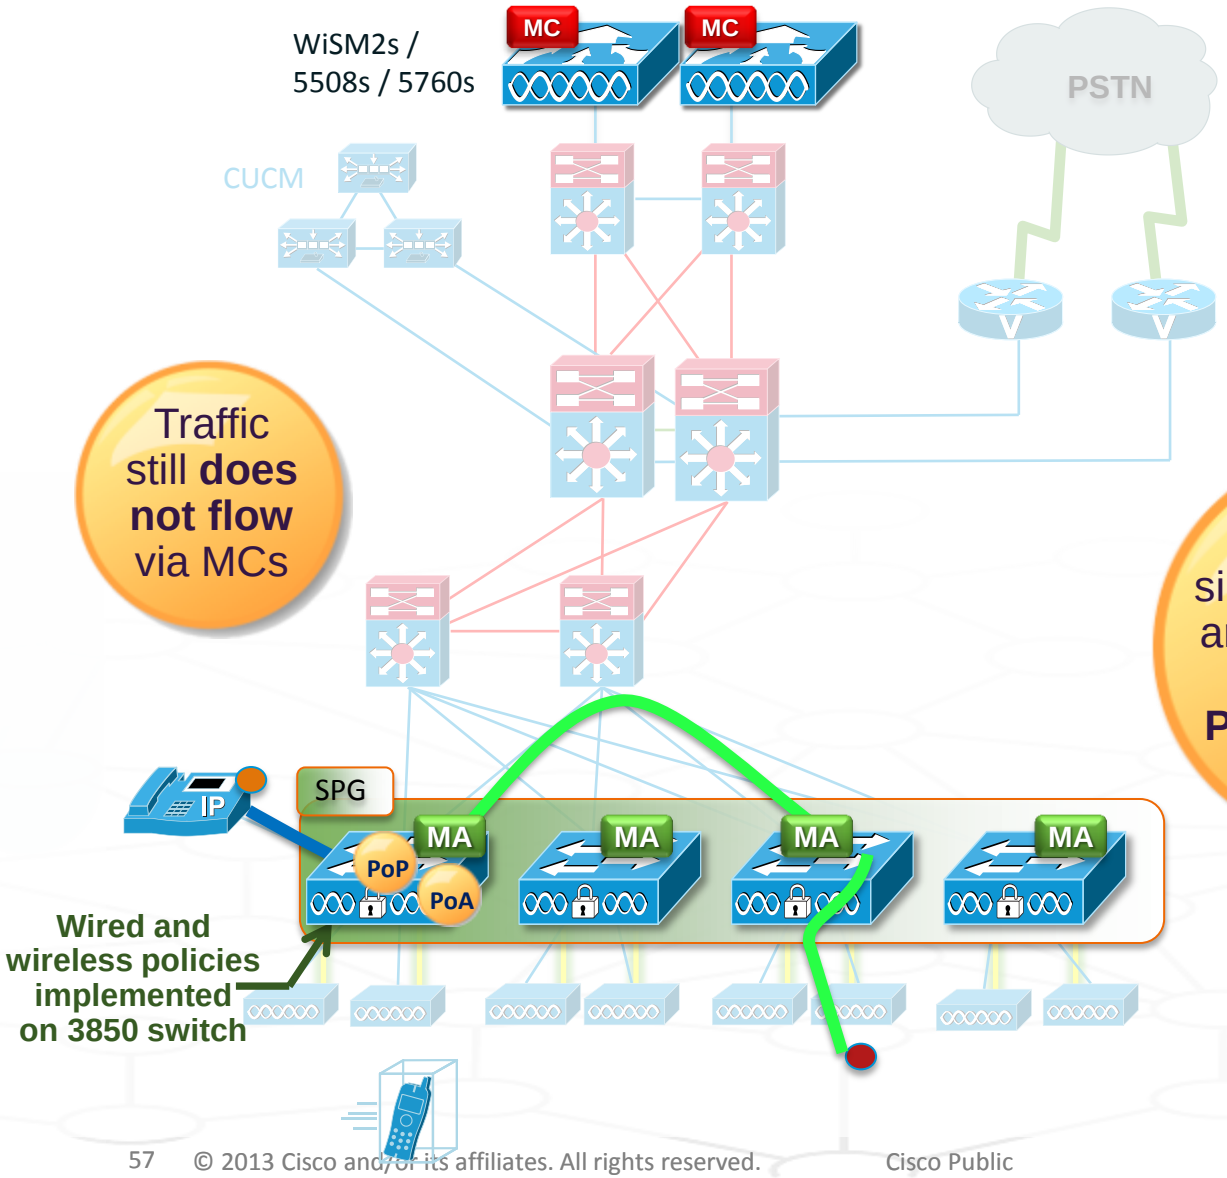

Converged policies and services for wired and wireless users

Traffic Flows, Comparison (Converged Access) –

- Now, our VoIP user on the Cisco Converged Access network roams, while a call is in progress between the wireless and wired handsets …
- **We can see that all of the user's traffic is still localised to their Switch Peer Group, below the distribution layer,** in both directions …

  In this example, a total of **3 hops** is incurred for each direction of the traffic path (assuming intra-SPG roaming) … two additional hops may be incurred for routing …

Converged Access

Traffic Flow and Roaming – Campus, L2 / L3 Roam (across Switch Peer Groups)

Roaming, Across SPGs (Campus) –

- Now, let's examine a few more types of user roams
- **In this example, the user roams across Switch Peer Groups –** since SPGs are typically formed around floors or other geographically-close areas, **this type of roam is possible, but less likely than roaming within an SPG**

Typically, this type of roam will take place across an L3 boundary (depends on wired setup) – **however, users are always* taken back to their PoP** for policy application

Converged Acce

Traffic Flow and Roar

```
L09-5760-1# show wcdb database all
        Total Number of Wireless Clients = 1
                MTE Clients                = 1

Mac Address     VlanId IP Address       Src If                Auth      Mob
--------------  ------ ---------------  --------------------  --------  -------
001e.65b7.7d1a       0 0.0.0.0          0x008FF88000000007 RUN       MTE
```

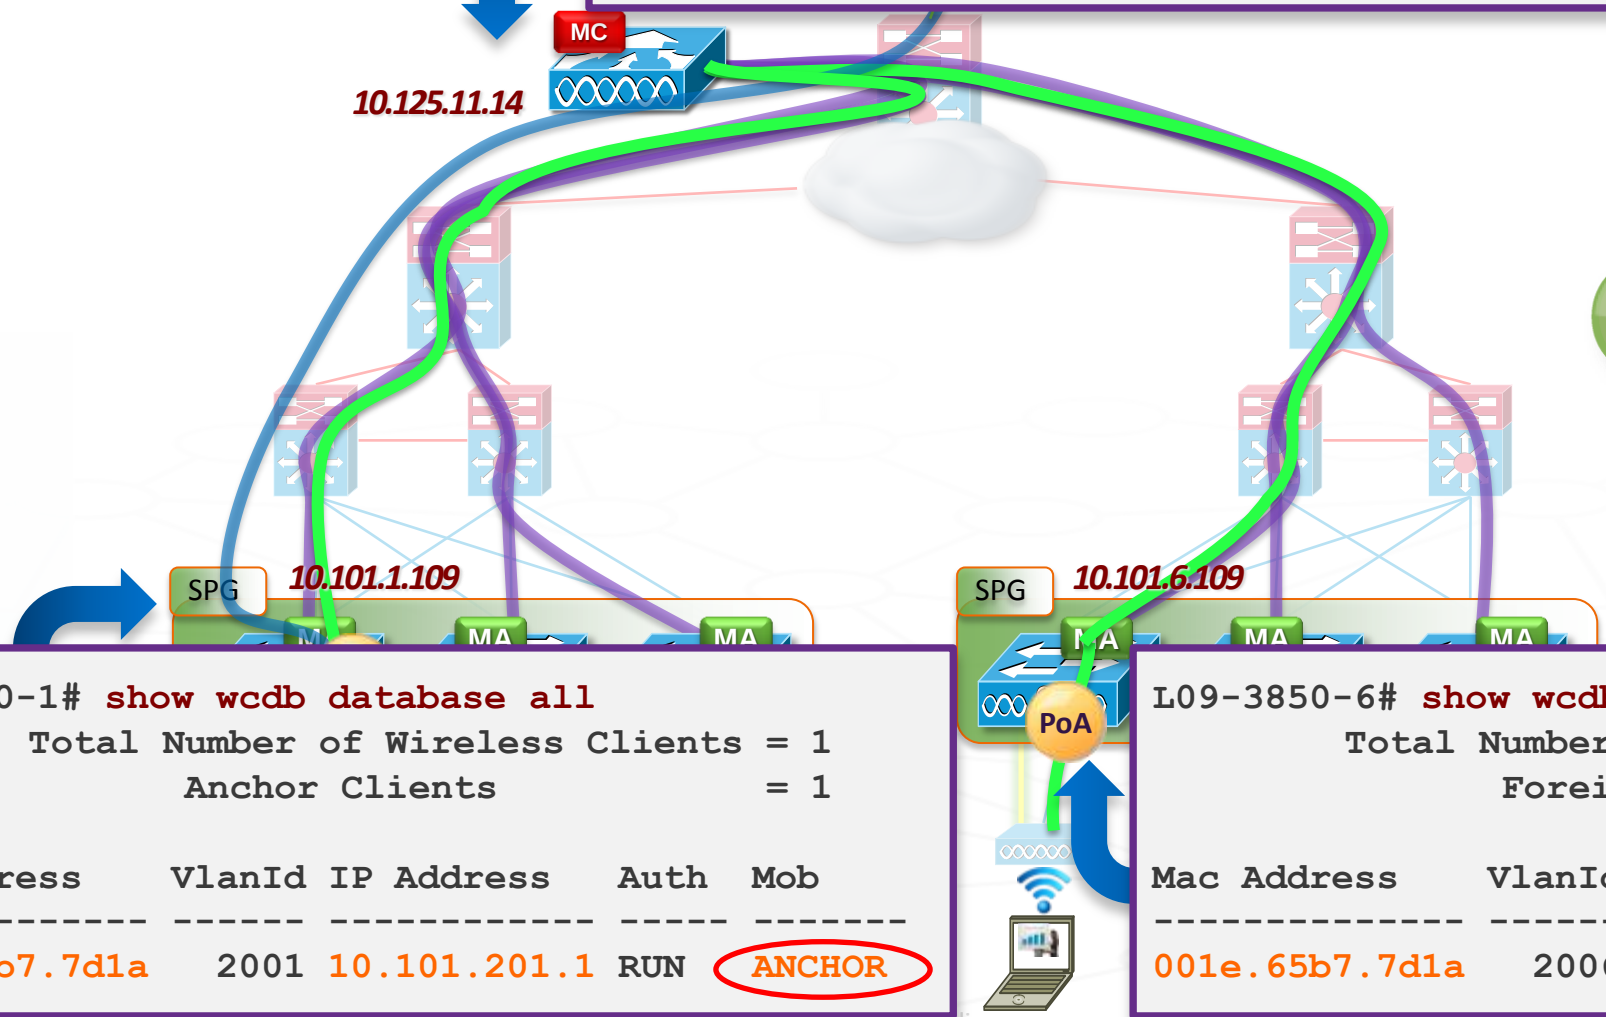

Control Plane view

```
L09-3850-1# show wcdb database all
        Total Number of Wireless Clients = 1
                Anchor Clients           = 1

Mac Address     VlanId IP Address    Auth  Mob
--------------  ------ ------------  ----- -------
001e.65b7.7d1a    2001 10.101.201.1  RUN   ANCHOR
```

```
L09-3850-6# show wcdb database all
        Total Number of Wireless Clients = 1
                Foreign Clients          = 1

Mac Address     VlanId IP Address    Auth  Mob
--------------  ------ ------------  ----- -------
001e.65b7.7d1a    2006 10.101.201.1  RUN   FOREIGN
```

Converged Acces

Traffic Flow and Roam ch Peer Groups)

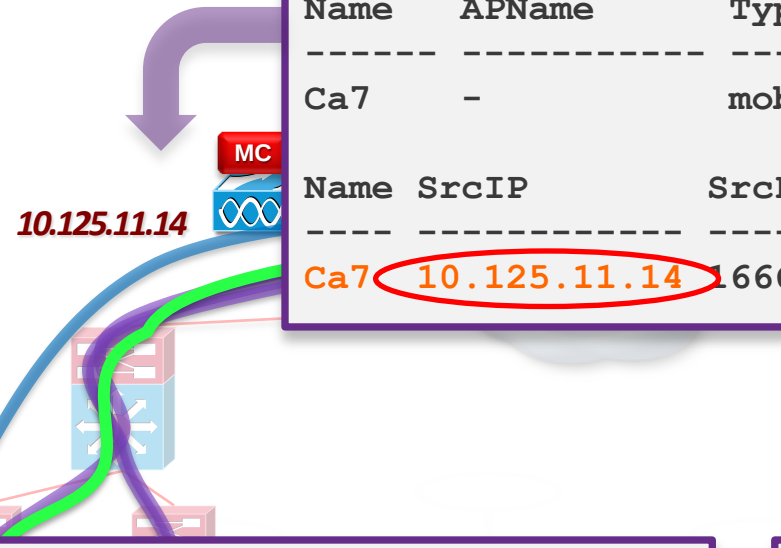


```
L09-5760-1# show mac address dynamic | inc Ca
4095      001e.65b7.7d1a    DYNAMIC     Ca7
```

```
L09-5760-1# show capwap summary

Name    APName          Type PhyPortIf Mode       McastIf
------ ----------- ---- --------- --------- -------
Ca7     -               mob  -          unicast    -

Name SrcIP          SrcPort DestIP         DstPort MTU
---- ------------ ------- ------------ ------- ----
Ca7  10.125.11.14 16667   10.101.6.109 16667   1464
```

Data Plane view

```
L09-3850-1# show mac address dynamic | inc Ca
2001      001e.65b7.7d1a    DYNAMIC     Ca1
```

```
L09-3850-1# show capwap summary
Name    APName          Type PhyPortIf Mode       McastIf
------ ----------- ---- --------- --------- ----
Ca1     -               mob  -          unicast    -

Name SrcIP          SrcPort DestIP         DstPort MTU
---- ------------ ------- ------------ ------- ----
Ca1  10.101.1.109 16667   10.125.11.14 16667   1464
```

```
L09-3850s-6# show mac address dynamic | inc Ca
4095      001e.65b7.7d1a    DYNAMIC     Ca3
```

```
L09-3850s-6# show capwap summary
Name    APName          Type PhyPortIf Mode       McastIf
------ ----------- ---- --------- --------- ----
Ca3     L09-AP1042-2    data Gi2/0/7    multicast Ca1

Name  SrcIP          SrcPort DestIP       DstPort MTU
----- ------------ ------- ----------- ------- ----
Ca3    10.101.6.109 5247    10.101.6.97 65321   1449
```

Converged Access

Traffic Flow and Roaming – Campus, L2 / L3 Roam (across Switch Peer Groups)

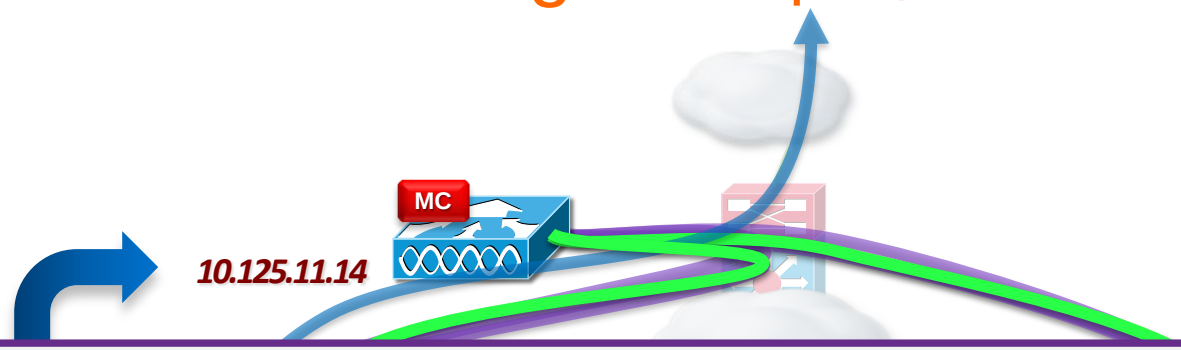

Overall view – across the entire Sub-Domain controlled by the MC

```
L09-5760-1# show wireless mobility controller client summary
Number of Clients : 5

State is the Sub-Domain state of the client.
* indicates IP of the associated Sub-domain
Associated Time in hours:minutes:seconds

MAC Address     State       Anchor IP          Associated IP    Associated Time
-------------------------------------------------------------------------------
001e.65b7.7d1a  Local      10.101.1.109       10.101.6.109     00:04:36
b817.c2f0.61b2  Local       0.0.0.0           10.101.7.109     00:21:07
74e1.b65a.a8f3  Local      10.101.3.109       10.101.1.109     00:03:27
cc08.e028.6fdd  Local       0.0.0.0           10.101.1.109     00:04:57
a467.06e2.813d  Local       0.0.0.0           10.101.3.109     00:02:56
```

Roamed client, Switch 1 to Switch 6 (inter-SPG)

Stationary client, Switch 7

Roamed client, Switch 3 to Switch 1 (intra-SPG)

Stationary client, Switch 1

Stationary client, Switch 3

Converged Access

Traffic Flow and Roaming – Campus, L2 / L3 Roam (across SPGs and MCs)

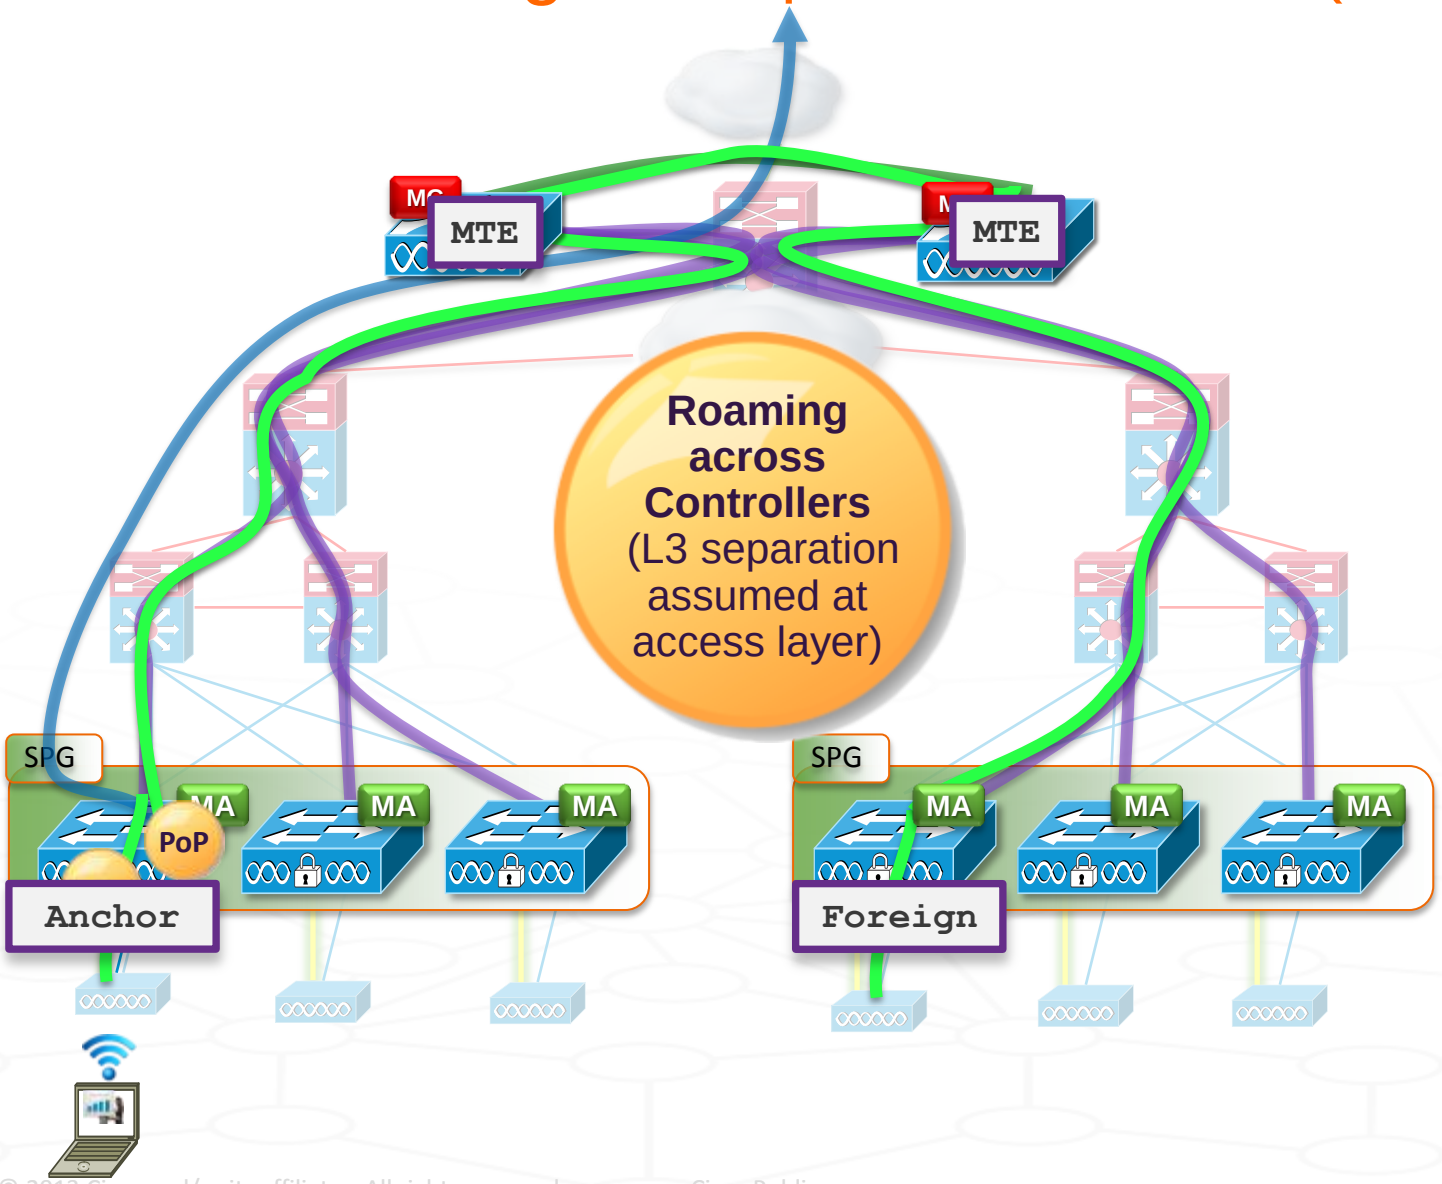

Roaming, Across SPGs and MCs (Campus) –

- Now, let's examine a few more types of user roams
- **In this example, the user roams across Switch Peer Groups and Controllers –** (within the same Mobility Group) … again, **this type of roam is possible, but less likely than intra-SPG roaming**

  Typically, this type of roam will take place across an L3 boundary (depends on wired setup) – **however, users are always* taken back to their PoP** for policy application

Converged Access

Traffic Flow and Roaming – Guest / Mobility Anchor

Quick note –

Catalyst 3850 switches can be used as MCs (with appropriate licensing) – however, they **cannot** be used as the destination for Guest Anchor tunnels (Guest Anchors must be discrete controllers)

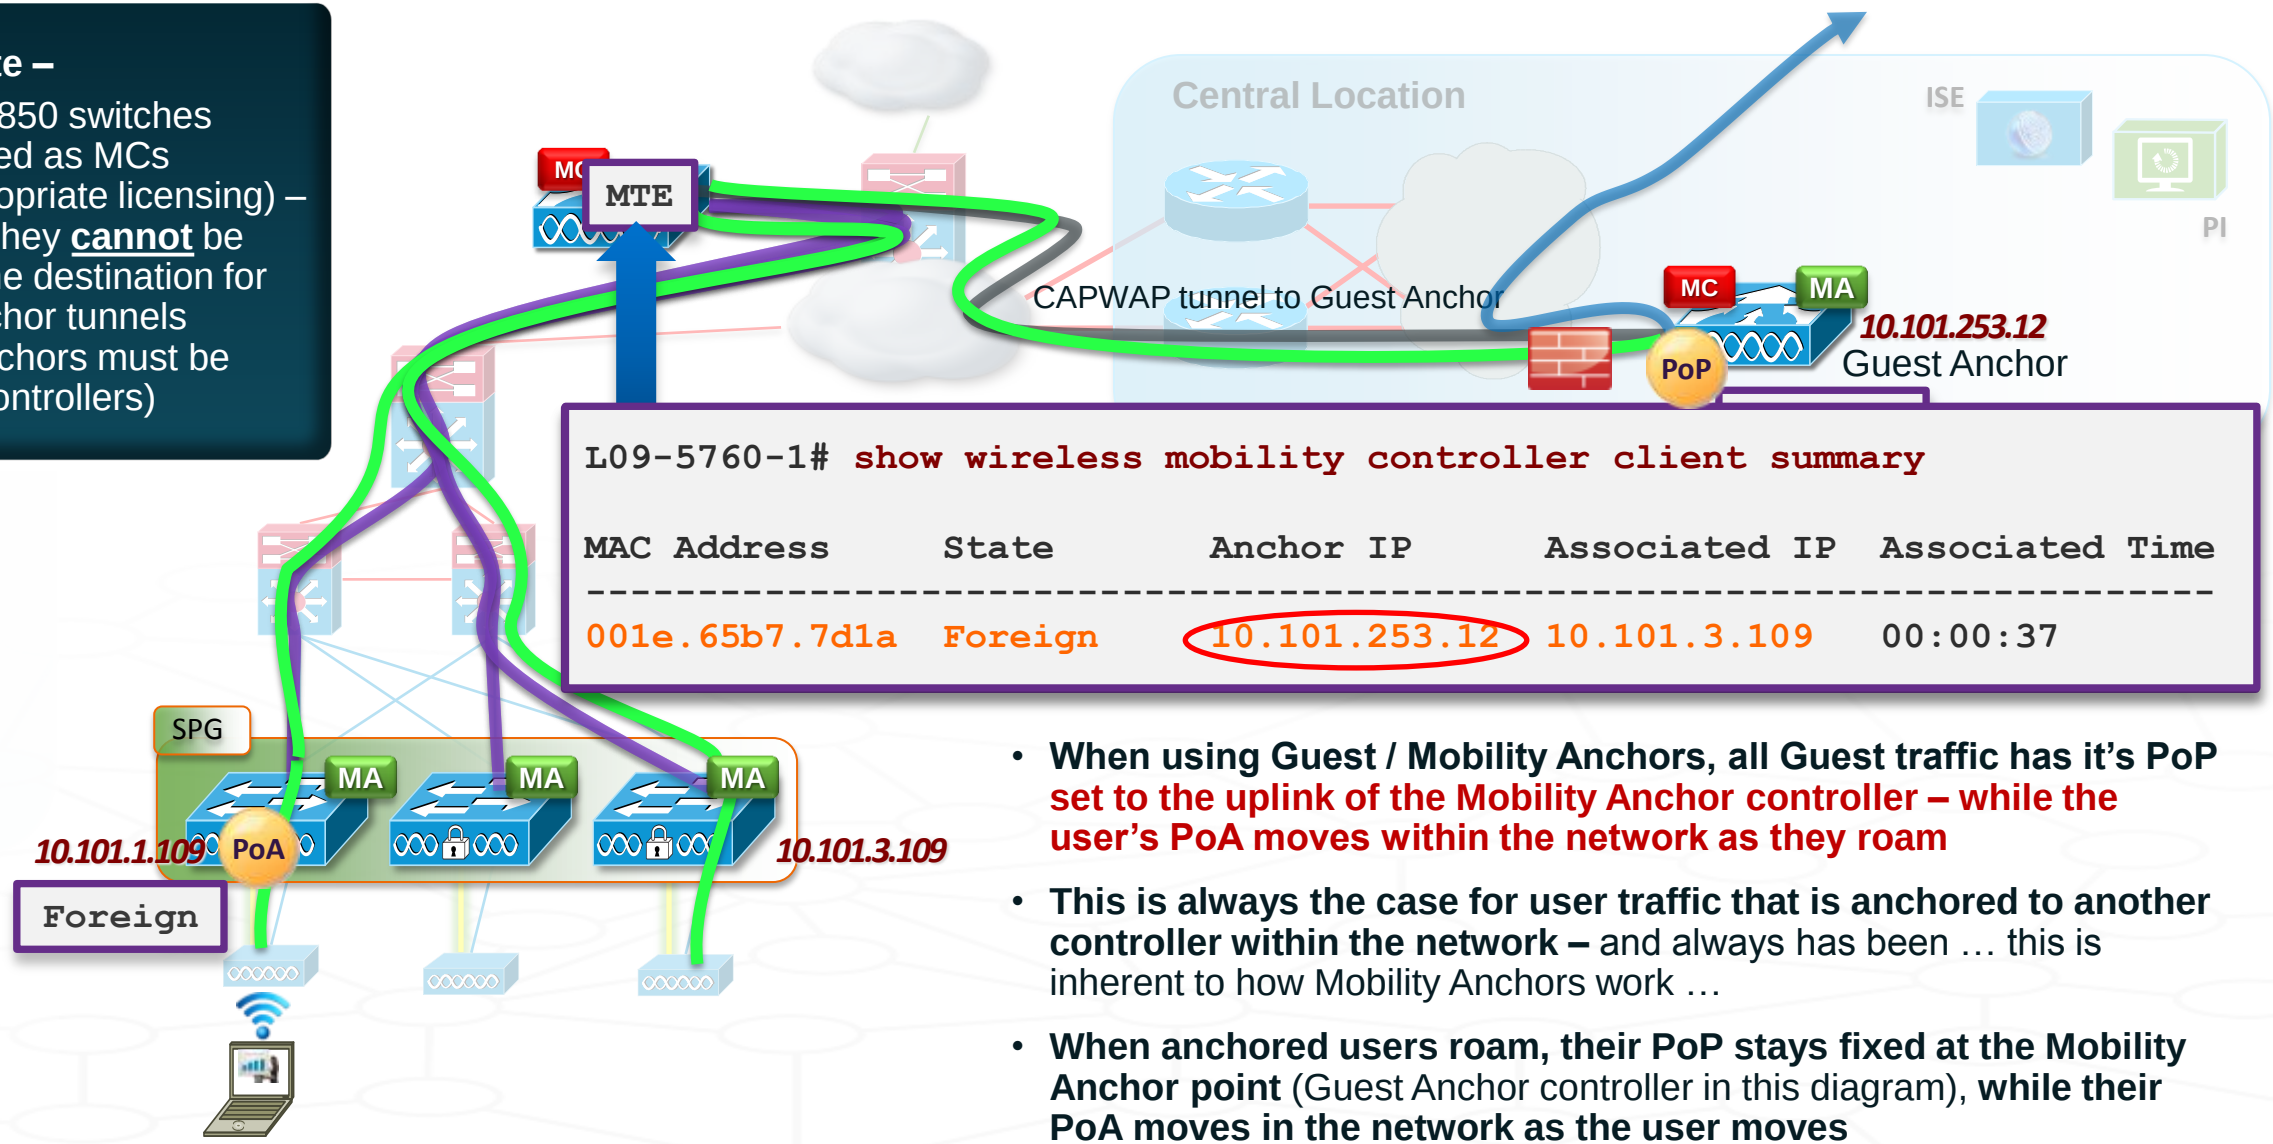


```
L09-5760-1# show wireless mobility controller client summary

MAC Address     State       Anchor IP       Associated IP  Associated Time
-------------------------------------------------------------------------
001e.65b7.7d1a  Foreign     10.101.253.12   10.101.3.109   00:00:37
```

- **When using Guest / Mobility Anchors, all Guest traffic has it's PoP set to the uplink of the Mobility Anchor controller – while the user's PoA moves within the network as they roam**
- **This is always the case for user traffic that is anchored to another controller within the network –** and always has been … this is inherent to how Mobility Anchors work …
- **When anchored users roam, their PoP stays fixed at the Mobility Anchor point** (Guest Anchor controller in this diagram), **while their PoA moves in the network as the user moves**

Converged Access

Traffic Flow and Roaming – L2 Roam (adjustable via setting)

```
L09-5760-1# show wireless mobility controller client summary

MAC Address      State        Anchor IP           Associated IP
-------------------------------------------------------------
001e.65b7.7d1a   Local        0.0.0.0             10.101.1.109
```

```
wlan UA-D3 4 UA-D3
 client vlan 3000
 no mobility anchor sticky
 security dot1x authenticat
 session-timeout 86400
 no shutdown
```

Move of th… user's enti… Mobility Context

Roaming across network

(L2 extension below distribution layer in this example)

```
…0-1# show wireless mobility controller client summary

…ress      State        Anchor IP           Associated IP
-------------------------------------------------------------
….65b7.7d1a  Local        0.0.0.0             10.101.3.109
```

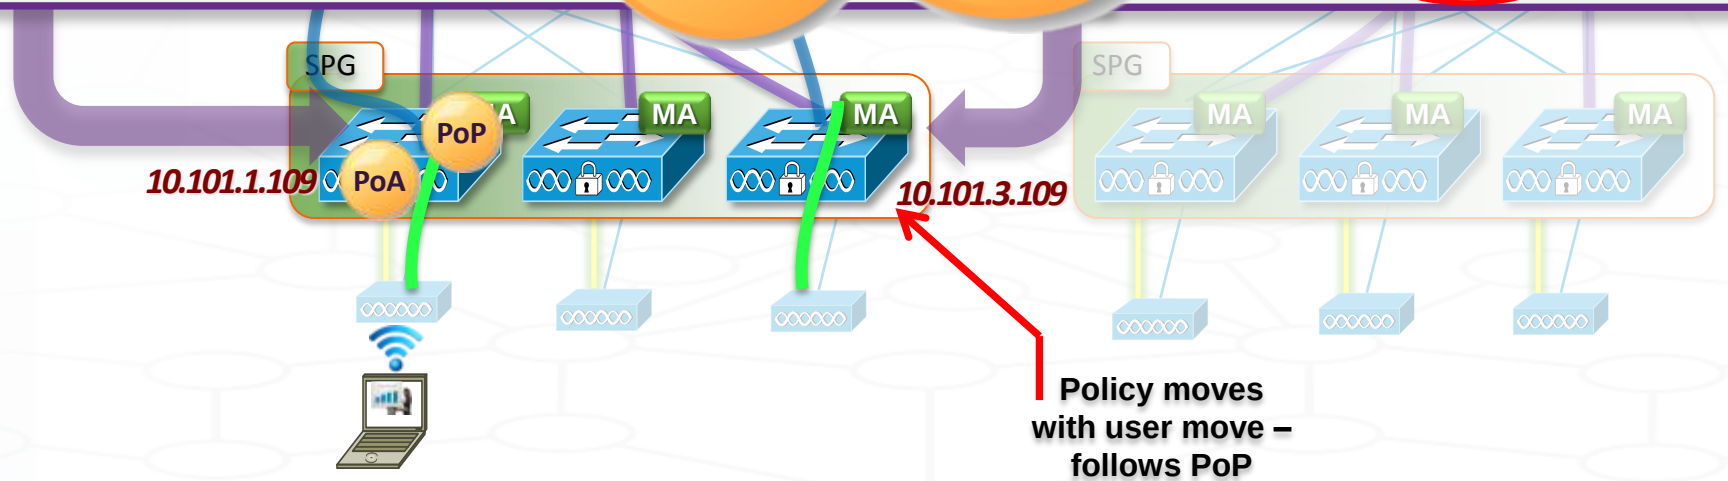

Roaming, Within an SPG (Layer 2) –

- Now, let's examine a few more types of user roams
- **In this example, the user roams within a Switch Peer Group …**

This type of L2 VLAN extension below the distribution layer may or may not be present in a given Enterprise network deployment – however, if this setup is present, an available setting allows for L2 roaming (move of both PoP and PoA)

Converged Access

Traffic Flow and Roaming – L2 Roam (impact of policy moves)

As Noted –

- **When a user roams in a L2 environment, an optional setting allows for both the user's PoA and PoP to move**.
- The benefits that accrue to a PoP move for an L2 user roam are **reduced end-to-end latency** for the user (less traffic hops), as well as a **reduction of state held within the network** (as the user needs to be kept track of only at the roamed-to switch).
- The drawback to a PoP move for an L2 user roam are likely **increased roam times**, as user policy may be retrieved from the AAA server, and applied at the roamed-to switch. The combination of these two elements may introduce a level of non-deterministic behaviour into the roam times if this option is used.
- **Default Behaviour –**
  - **L2 Roams Disabled** – by default, all roams (whether across an L3 boundary or not) carry the user's traffic from their roamed-to switch (where the user's PoA has moved to), back to the original switch the user associated through (where the user's PoP remains). In this case, **the user's policy application point remains fixed**, and roam times are more **deterministic**.
  - However, if desired, **this behaviour can be modified via a setting to allow for an L2 roam** – assuming the network topology involved allows for the appropriate Layer 2 extension across the network.

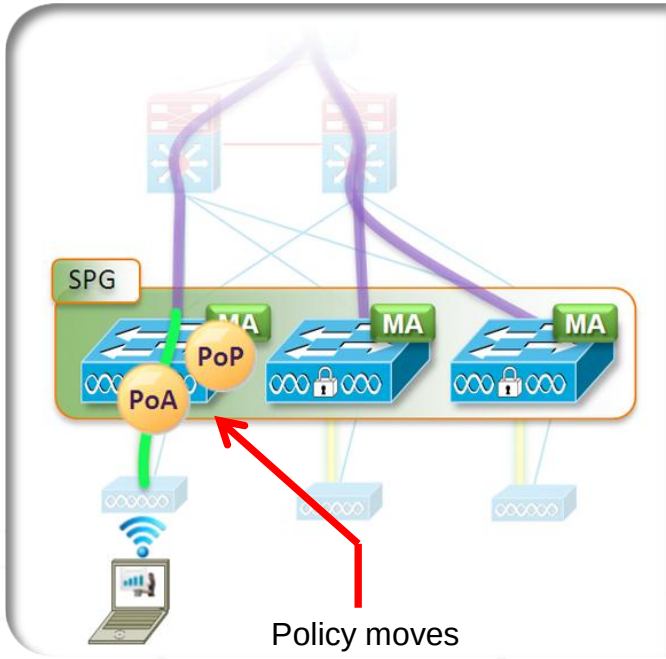

Converged Access

Catalyst 3850-based MCs – Functionality

As we saw previously, we can also optionally use a Catalyst 3850 switch as an MC + co-located MA for a Switch Peer Group … let's explore this in more detail –

- **Single Catalyst 3850 MC** supported per Switch Peer Group …
- which can have **up to 16 x MAs (stacks) per 3850-based MC**
- Single Catalyst 3850 MC can handle up to **50 APs and 2,000 clients total … therefore, up to 50 APs and 2,000 clients in a Catalyst 3850-based Switch Peer Group**
- MC handles inter-SPG roaming, RRM, Guest Access, etc.
- More scalable MC capability can be provided by 5760 / WiSM2

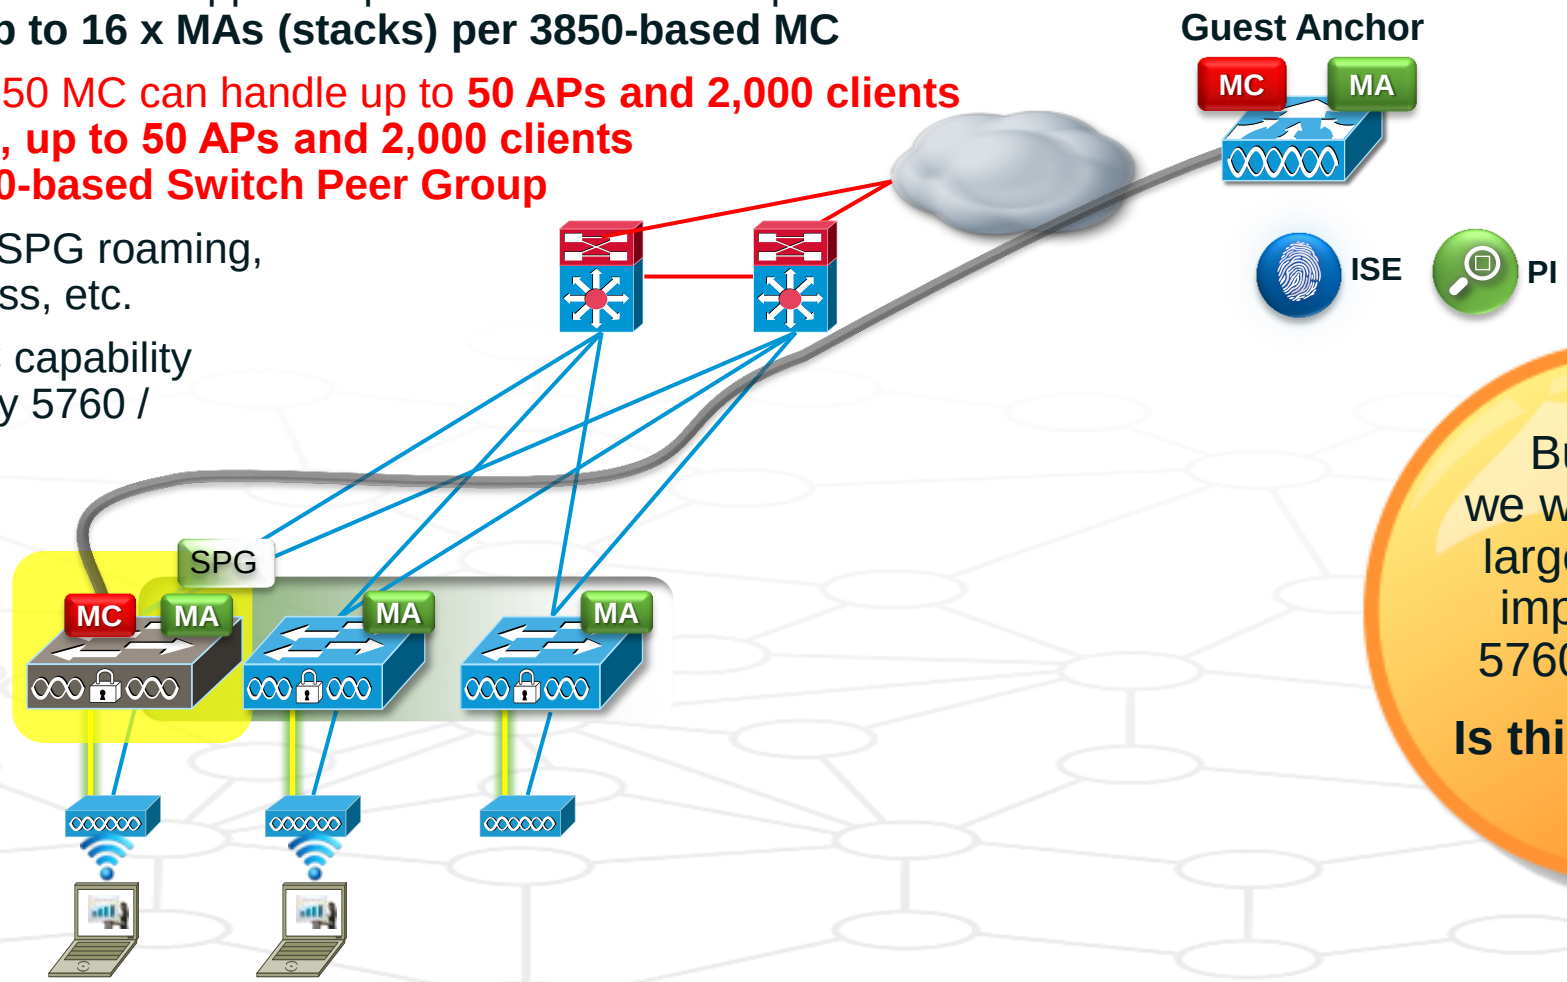

But what if we want to scale larger, **without** implementing 5760 / WiSM2?

Is this possible?

Converged Access

Catalyst 3850-based MCs – Scaling

Switch Peer Group / Mobility Group Scaling with Catalyst 3850 –

- **Up to 8 x Catalyst 3850 MCs** can be formed into a Mobility Group
- **Up to 250 APs total and 16,000 clients supported (maximum)** across a Mobility Group made up solely of Catalyst 3850 switches
  - **Licensing is per MC** – not pooled across MCs
  - **RRM, etc. is coordinated across the MCs** in the same Mobility Group
- **Guest tunnelling is per MC** – to Guest Anchor controller

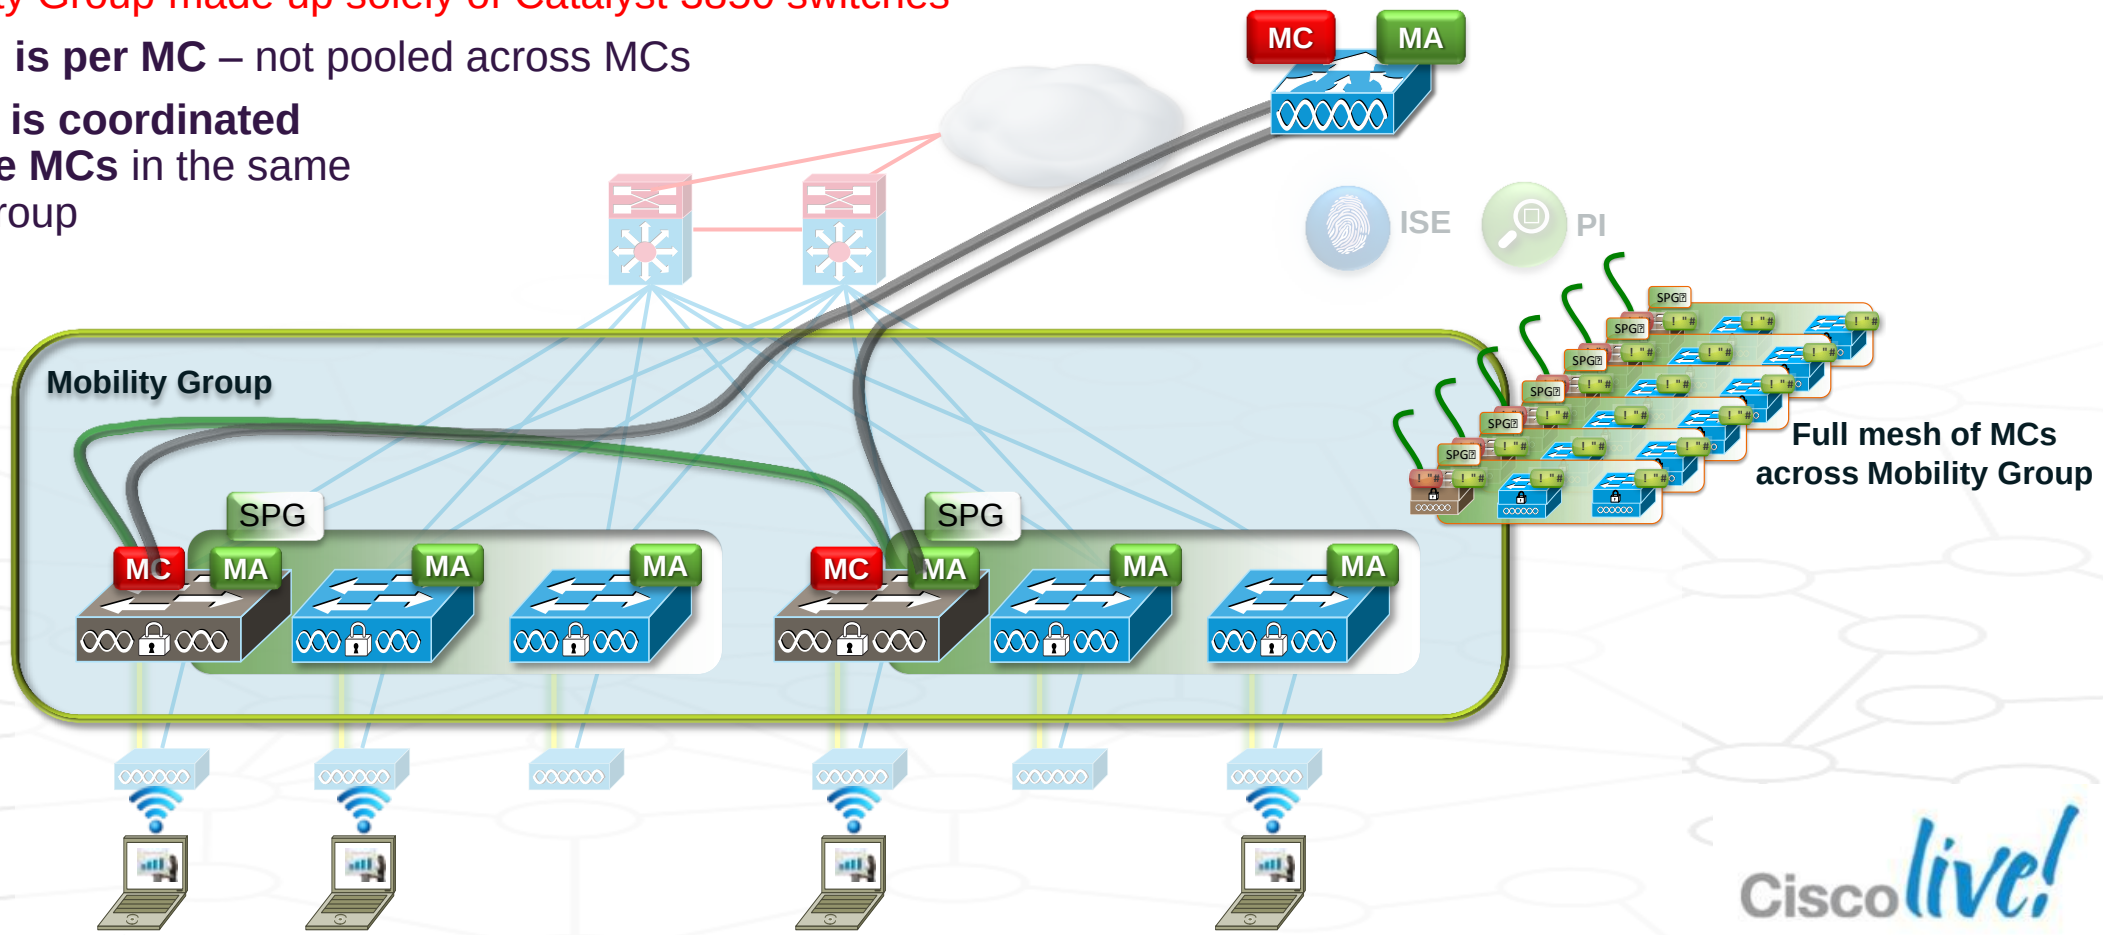

Converged Access

Catalyst 3850-based MCs – Roaming

There are multiple roaming scenarios with Catalyst 3850-based MCs –

- **These replicate the traffic flow expectations seen elsewhere with Converged Access**
- **Traffic within an SPG flows directly between MAs – traffic between SPGs flows via MCs**
  - **Which, in this case,** are Catalyst 3850 switches operating as MCs
  - **Catalyst 3850-based MC deployments are likely to be common in branches and even possibly smaller Campuses**
  - **Larger deployments are likely to use discrete controllers** (5760, 5508, WiSM2s) as MCs, for **scalability and simplicity**
  - **Rather than detail every roaming case here, these are summarised below – Full details are given in the Reference section at the end of this slide deck …**

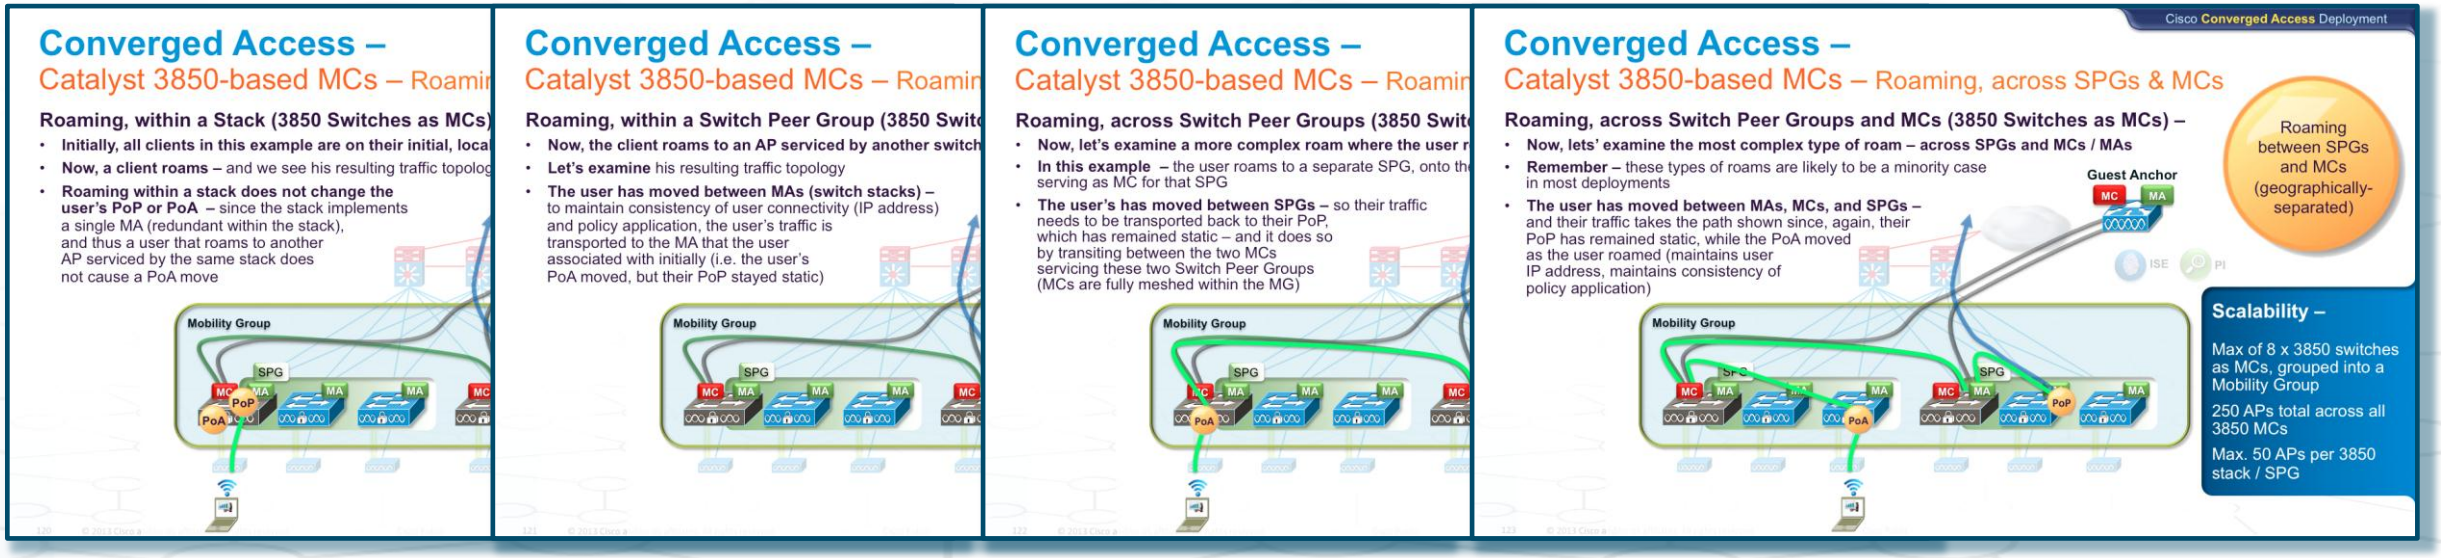

Converged Access

Catalyst 3850-based MCs – When to Use

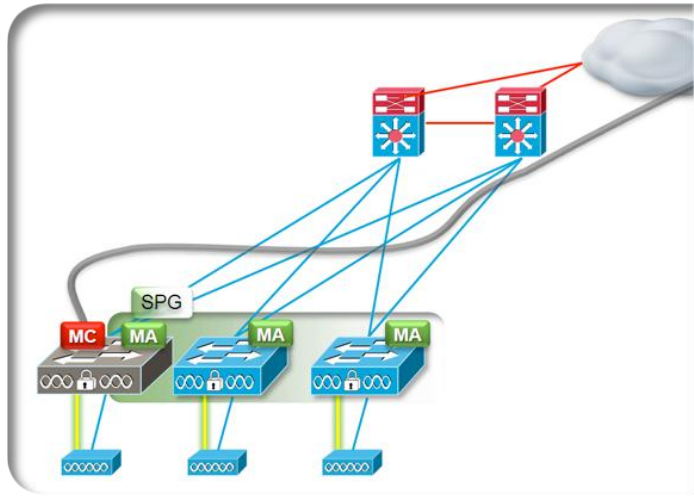

Considerations –

- **Many larger designs (such as most Campuses) will likely utilise a discrete controller, or group of controllers, as MCs.** Combined with Catalyst 3850 switches as MAs, this likely provides the most scalable design option for a larger network build.
- **However, if using 3850 switches as MCs for smaller builds – and with the scaling limits detailed on the previous slide in mind – we need to determine where to best use this capability.**
- **Pros –**
  - **CapEx cost savings** – via the elimination of a discrete-controller-as-MC in some designs (typically, smaller use cases and deployments) … cost also needs to take into consideration licensing on the Catalyst 3850 switches.
- **Cons –**
  - **OpEx complexity** – due to some additional complexity that comes into roaming situations when using multiple 3850 switch-based MCs (as detailed in the preceding slide). While not insurmountable, this does need to be factored in as part of the decision process.

Roaming details provided on Reference slides

Conclusion –

In smaller designs (such as branches), the use of Catalyst 3850 switches as MCs is likely workable. In mid-sized designs, this may also be workable, but does lead to some additional roaming considerations (as detailed on the following slides). In large campus deployments, the use of controllers as MCs is more likely, due to economies of scale.

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a "double-click" deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – **High Availability**
- Converged Access Deployment – Quality of Service

Summary

High Availability

State Held within the Network – for Local Users and Roamed Users

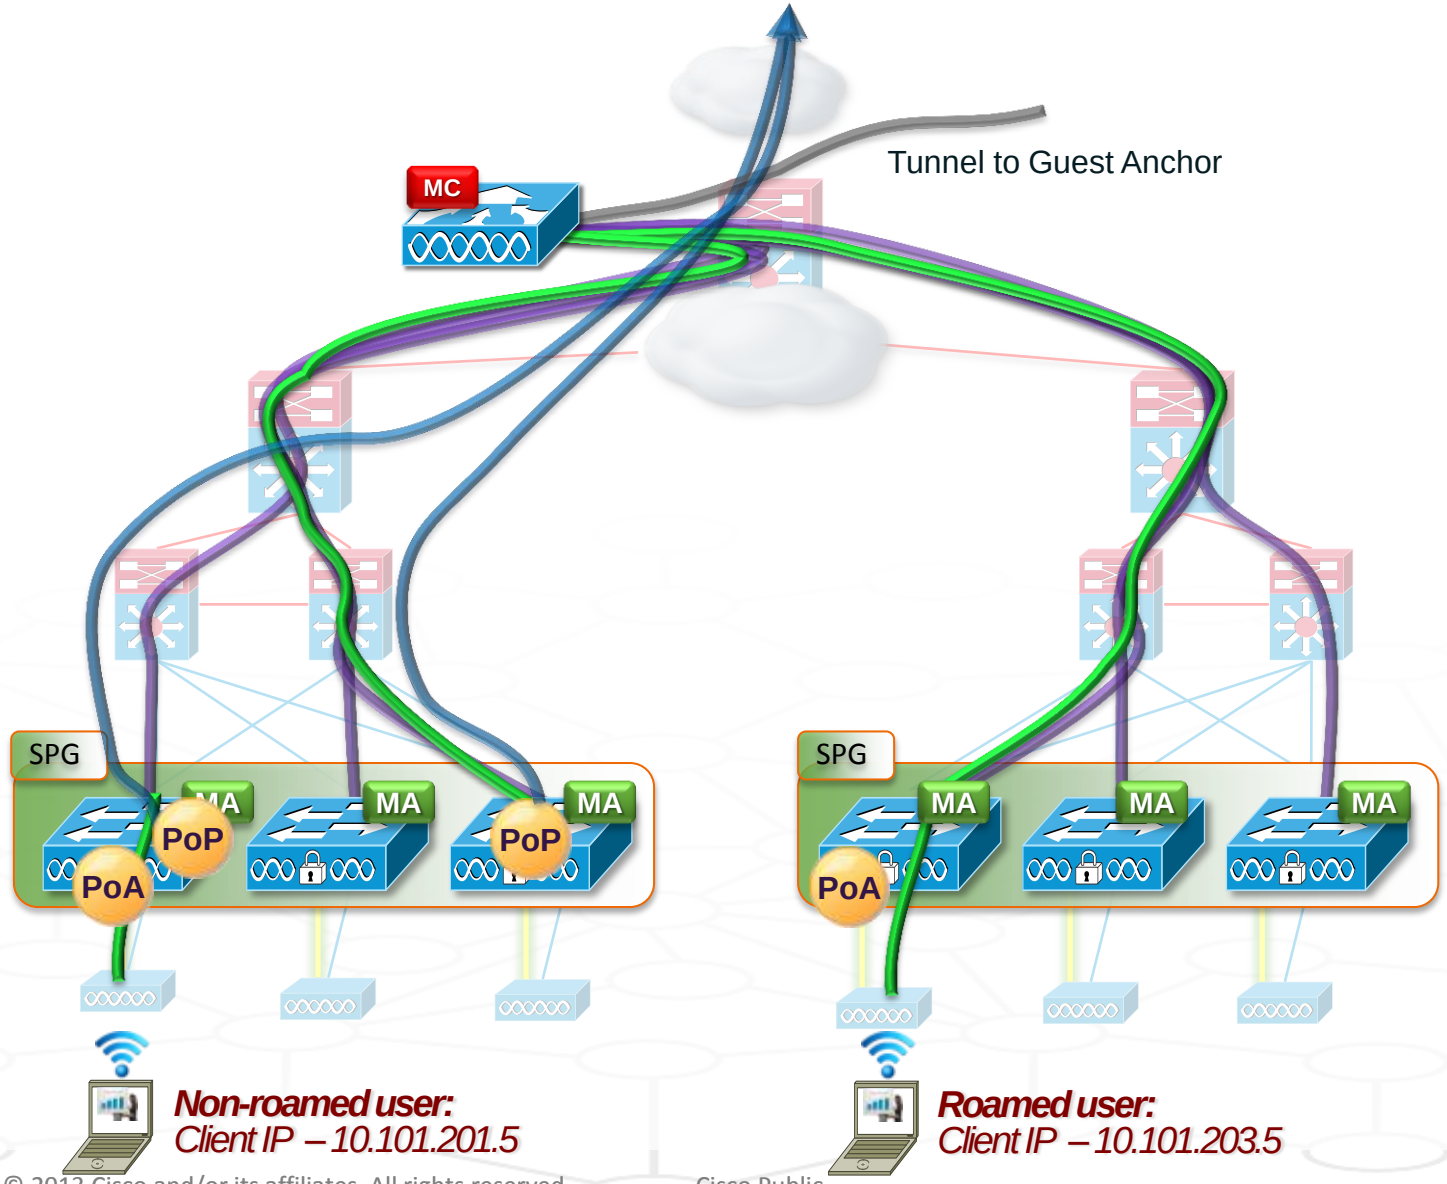

Roamed and Local users, High Availability Considerations –

- State for users is held within the network (on MAs and MCs) – in this case, we are using a discrete controller (5760, 5508, or WiSM2) as an MC
- **In this example as shown,** we have two users – one local (non-roaming), and the other roamed across SPGs (same MC) …
- **Note that in this case, the roamed user's client IP address is associated with the IP address pool on the right-hand switch in the left-side SPG** (where the user originally associated) …

High Availability

State Held within the Network – Local Users and Roamed Users … MC View

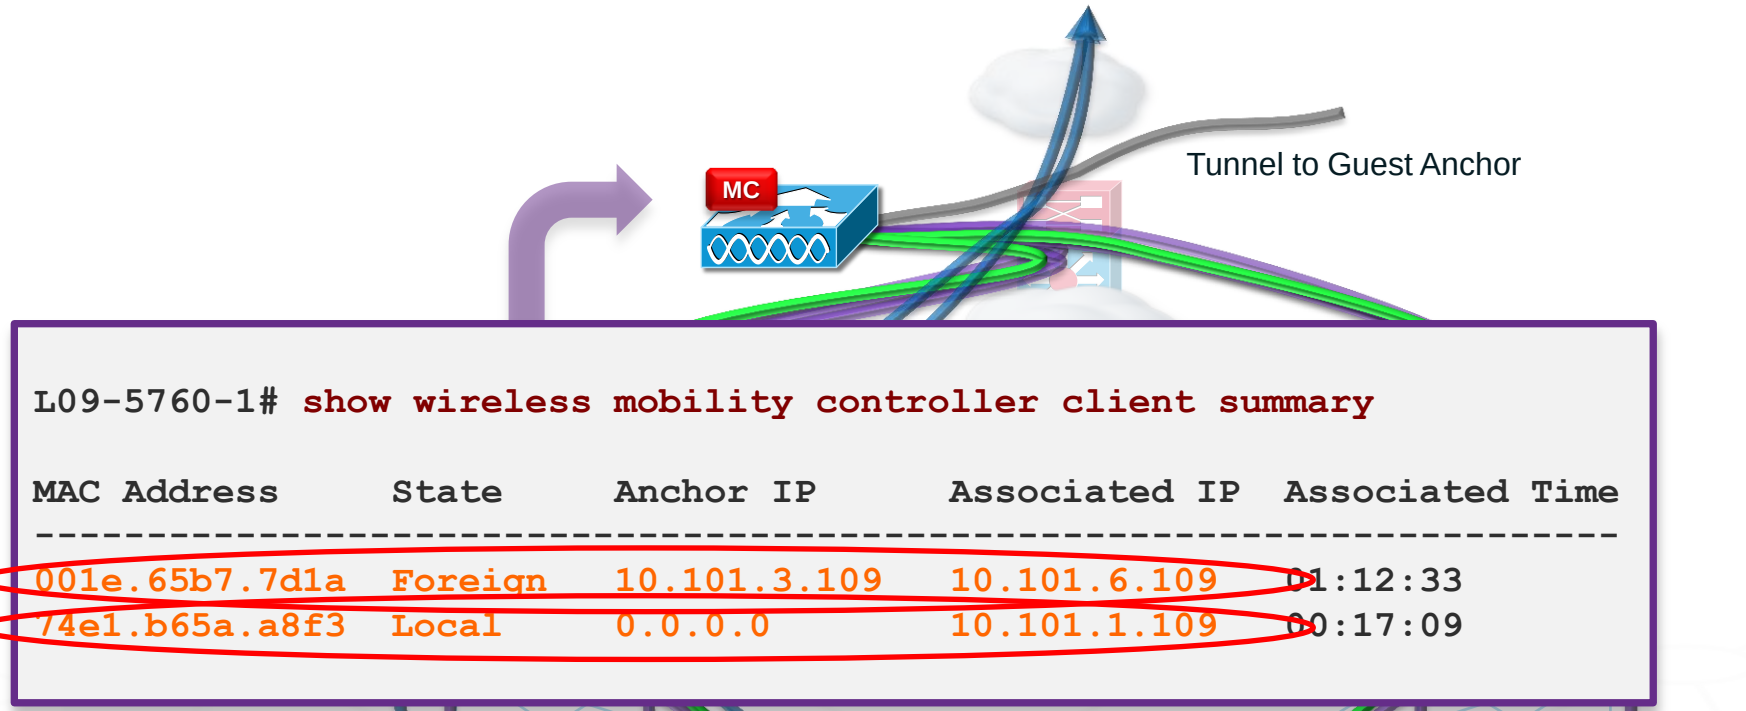

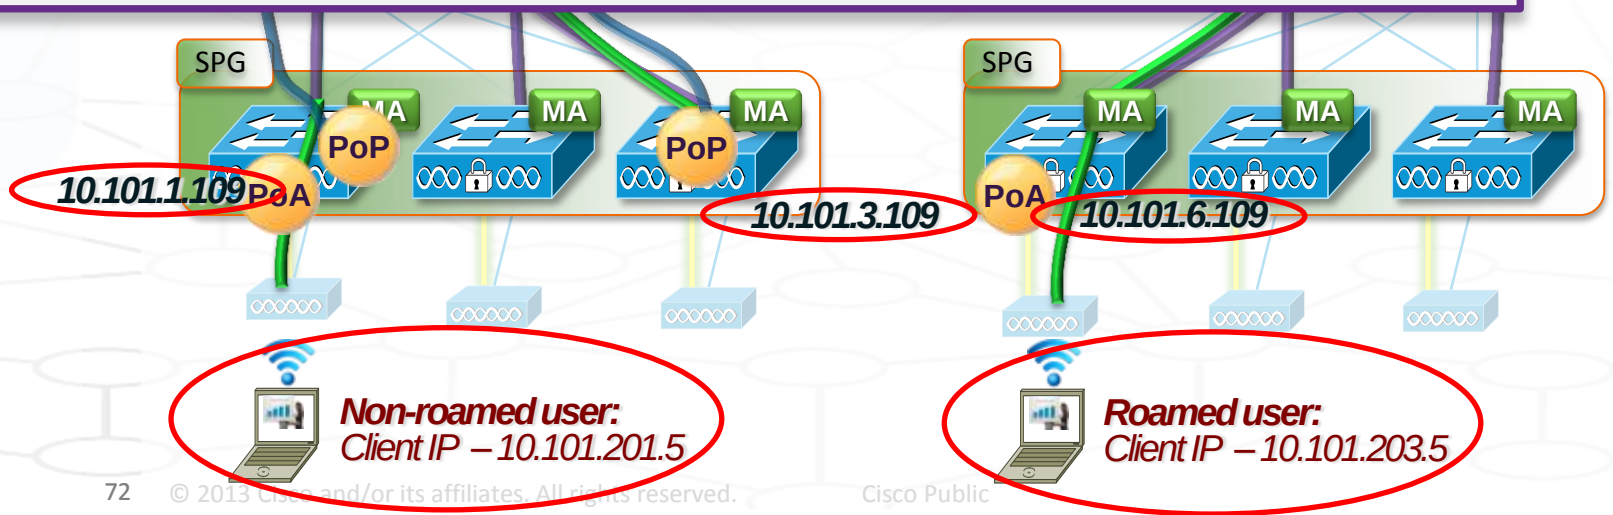

Roamed and Local users, High Availability Considerations –

- **Here, we can see the state of the network for the roamed and non-roamed clients**, as reflected at the MC itself (shows a snapshot of the traffic flows within the Mobility Sub-Domain the MC controls) …

High Availability

State Held within the Network – Local Users and Roamed Users … MA Views

```
L09-3850-1# show wcdb database all

Mac Address     VlanId IP Address       Auth Mob
--------------- ------ ---------------- ---- -----
001e.65b7.7d1a    2001 10.101.201.5     RUN  LOCAL
```

Tunnel to Guest Anchor

```
L09-3850-3# show wcdb database all

Mac Address     VlanId IP Address       Auth Mob
--------------- ------ ---------------- ---- ------
74e1.b65a.a8f3 2003   10.101.203.5      RUN  ANCHOR
```

```
L09-3850-6# show wcdb database all

Mac Address     VlanId IP Address          Auth Mob
--------------- ------ ------------------- -------
74e1.b65a.a8f3 2006   10.101.203.5         RUN  FOREIGN
```

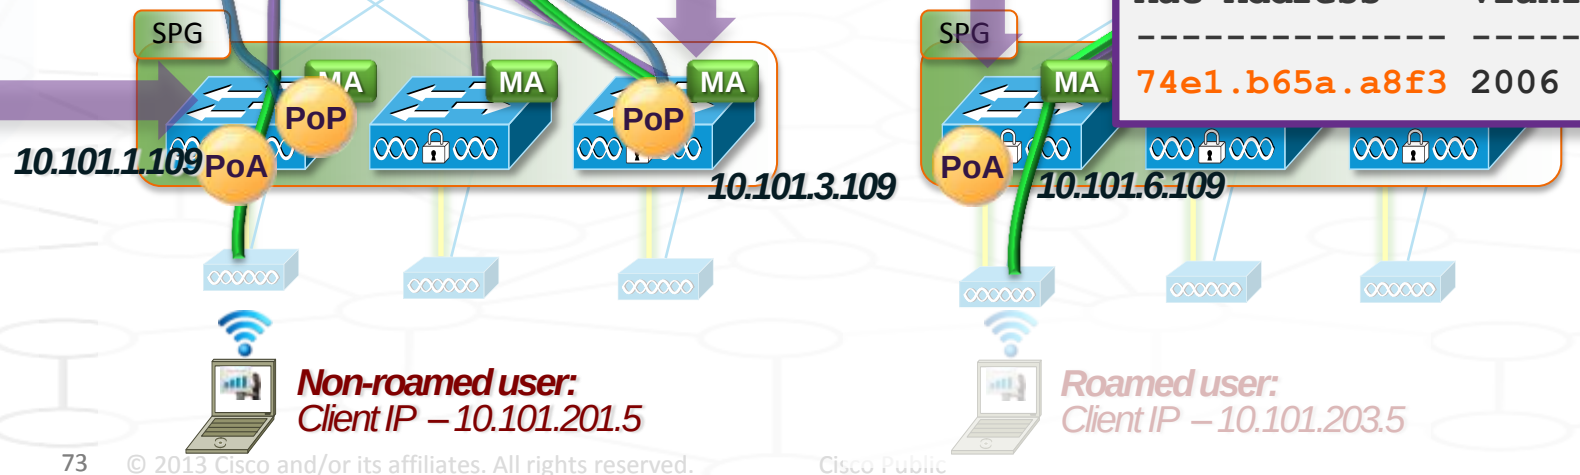

Roamed and Local users, High Availability Considerations –

- **And here, we see the state of the network**, as reflected at the MAs involved …

High Availability

MC Failure – and the Effect on the MC's Sub-Domain and Anchor Connections

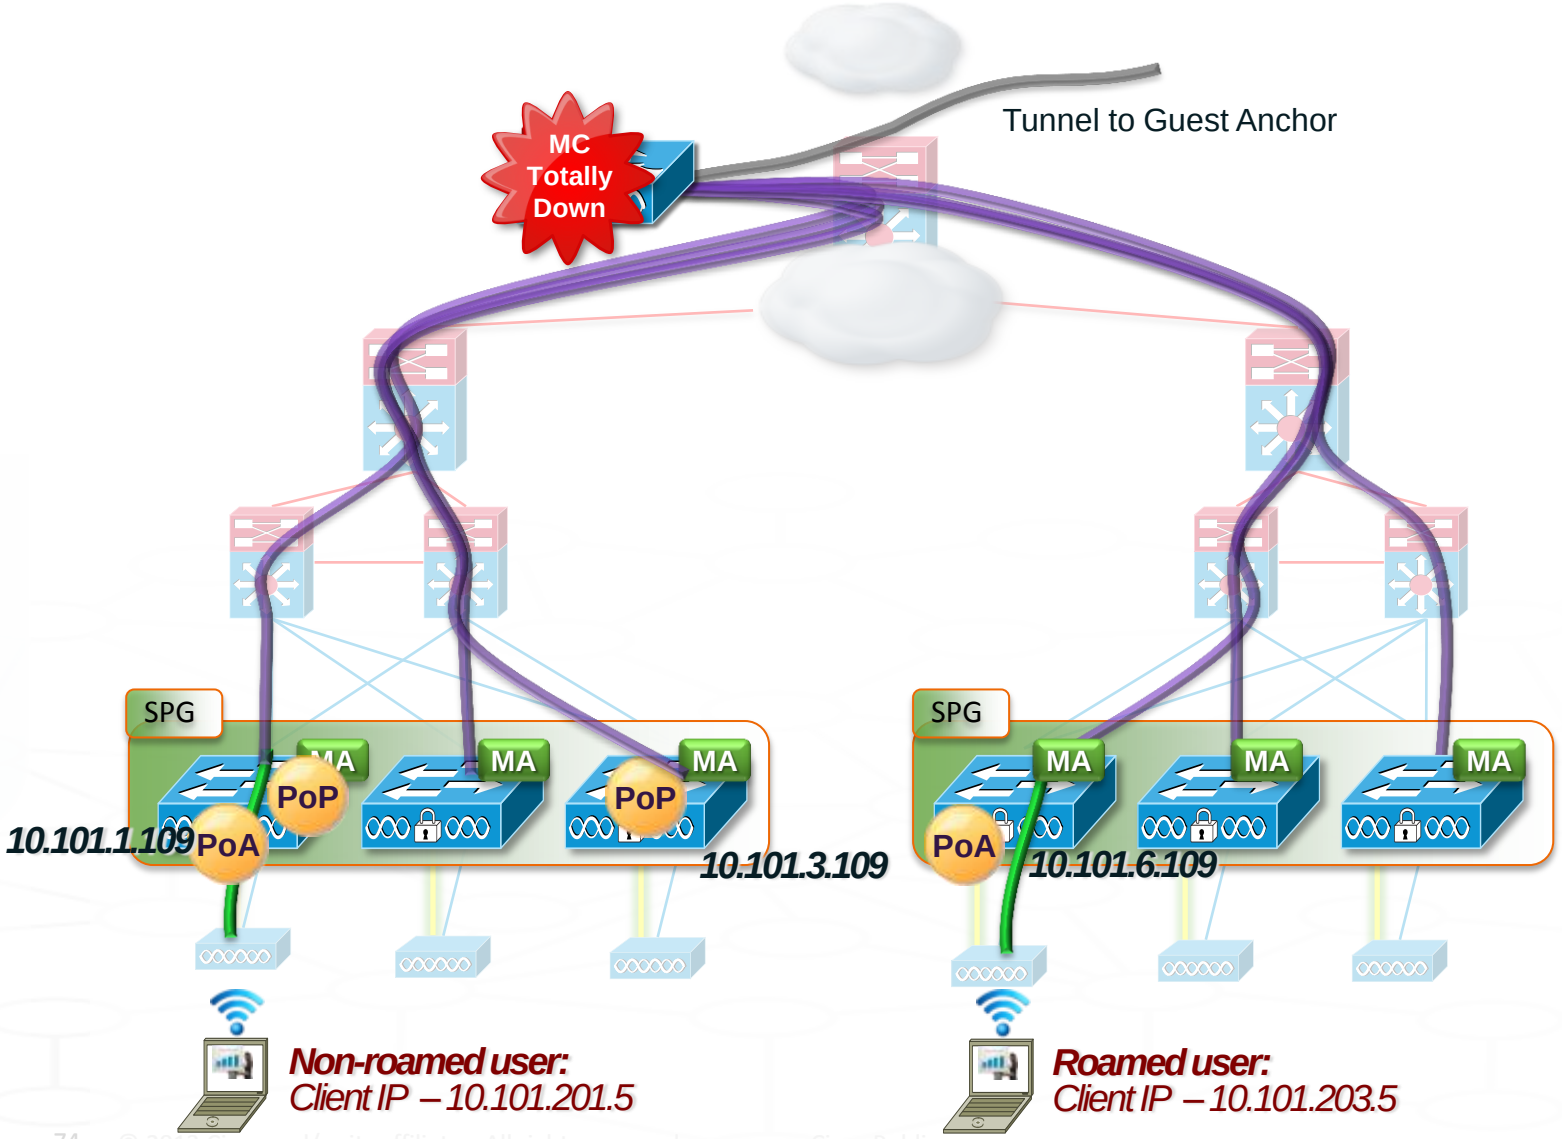

Roamed and Local users, High Availability Considerations –

- **Now, the MC fails (power down in this case)** … let's examine the effects …
- **When the MC for a given Sub-Domain goes down, all of the tunnels serviced by that MC go down** – this includes all MA-MC tunnels (purple tunnels as shown on this diagram), as well as any MC-Guest Anchor tunnel (if present – grey tunnel as shown on this diagram)

  Note that all of the tunnel connections between switches within the SPGs themselves stay up – as these are pre-formed at SPG creation, and once up, do not depend on the MC to stay up …

High Availability

MC Failure – Effect on Local (Non-Roamed) Clients

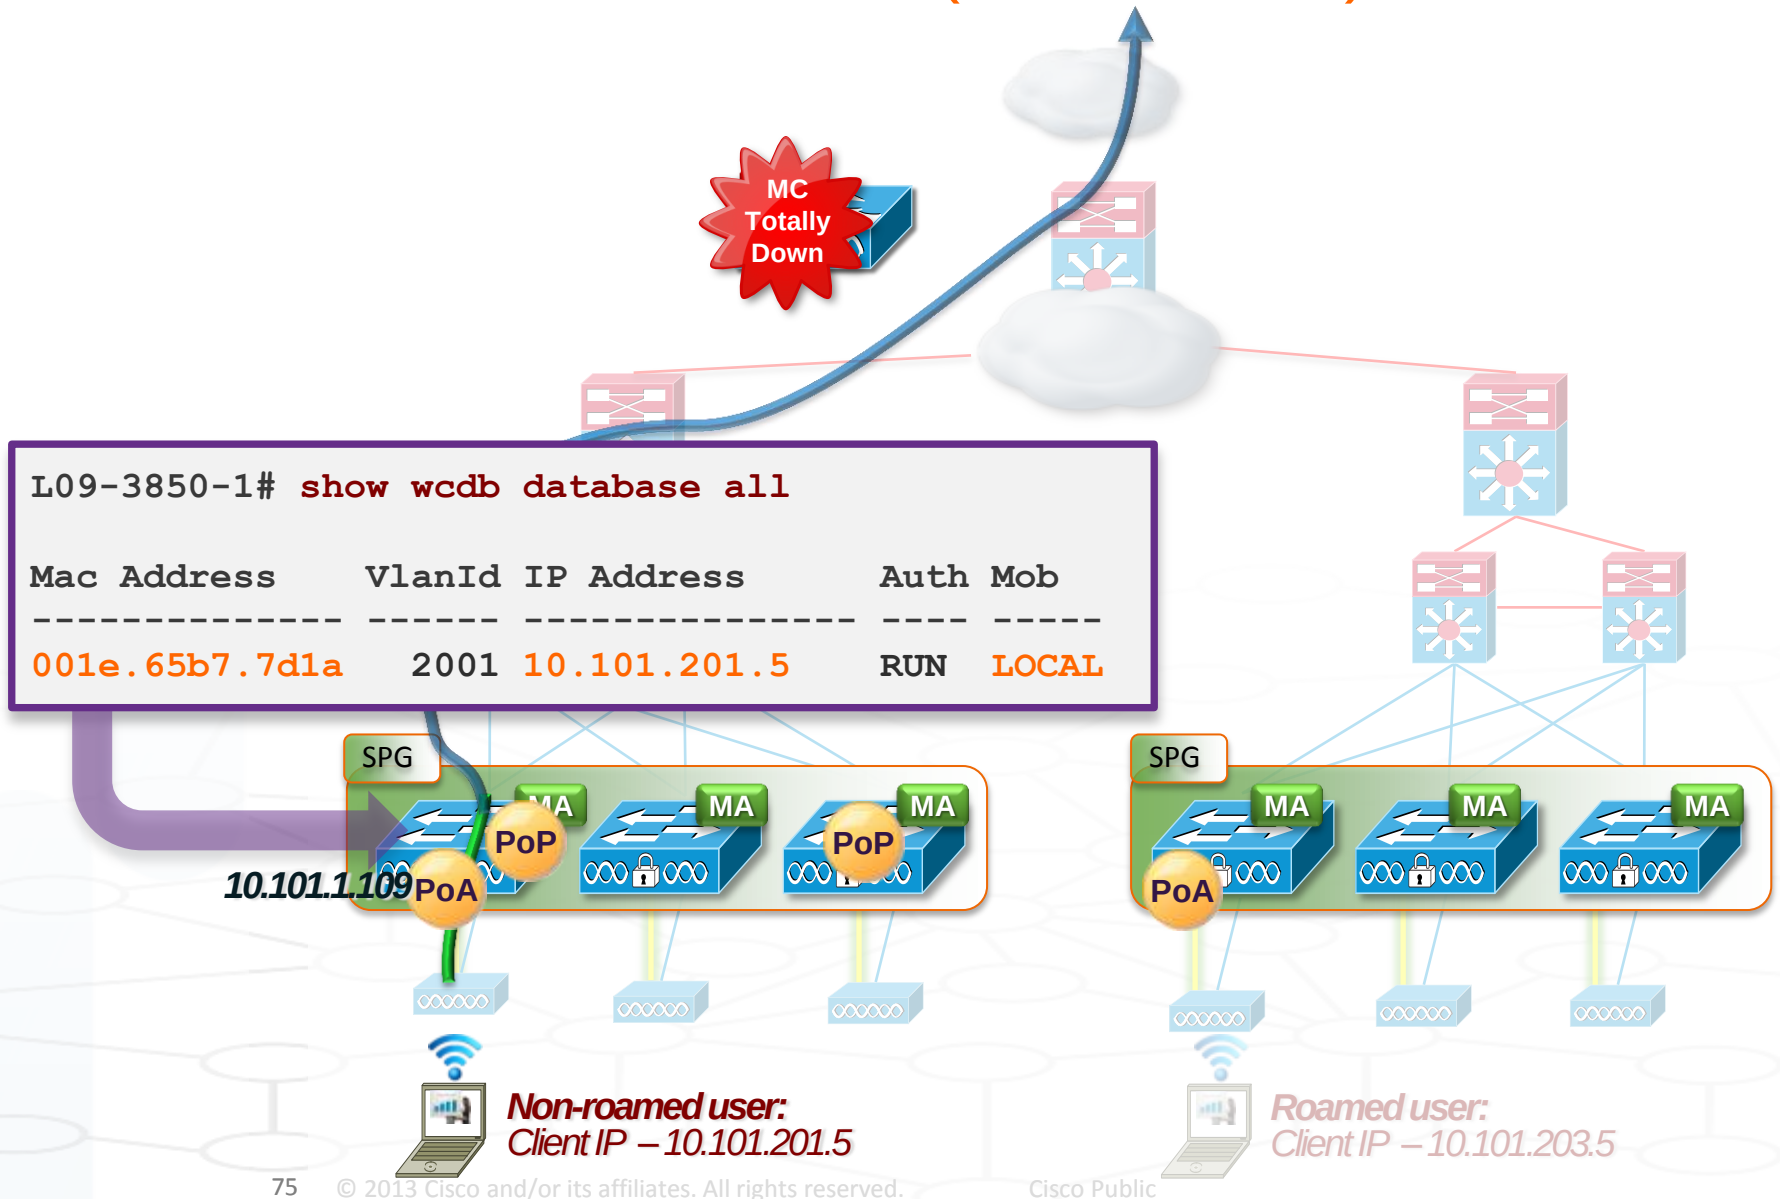

Roamed and Local users, High Availability Considerations –

- **For a local (non-roamed) user, the effect of an MC failure is not that severe …**
- **The local user still continues to operate,** as their traffic flow is terminated locally at their MA switch …
- **However, the user may be missing some services (Guest Access, RRM, Fast Roaming, etc) for the duration of the MC failure …** as these functions depend on the MC servicing the SPG(s) …

  … and as well, **inter-SPG roaming will be affected**, as shown on the following slides

High Availability

MC Failure – Effect on Previously-Roamed Clients

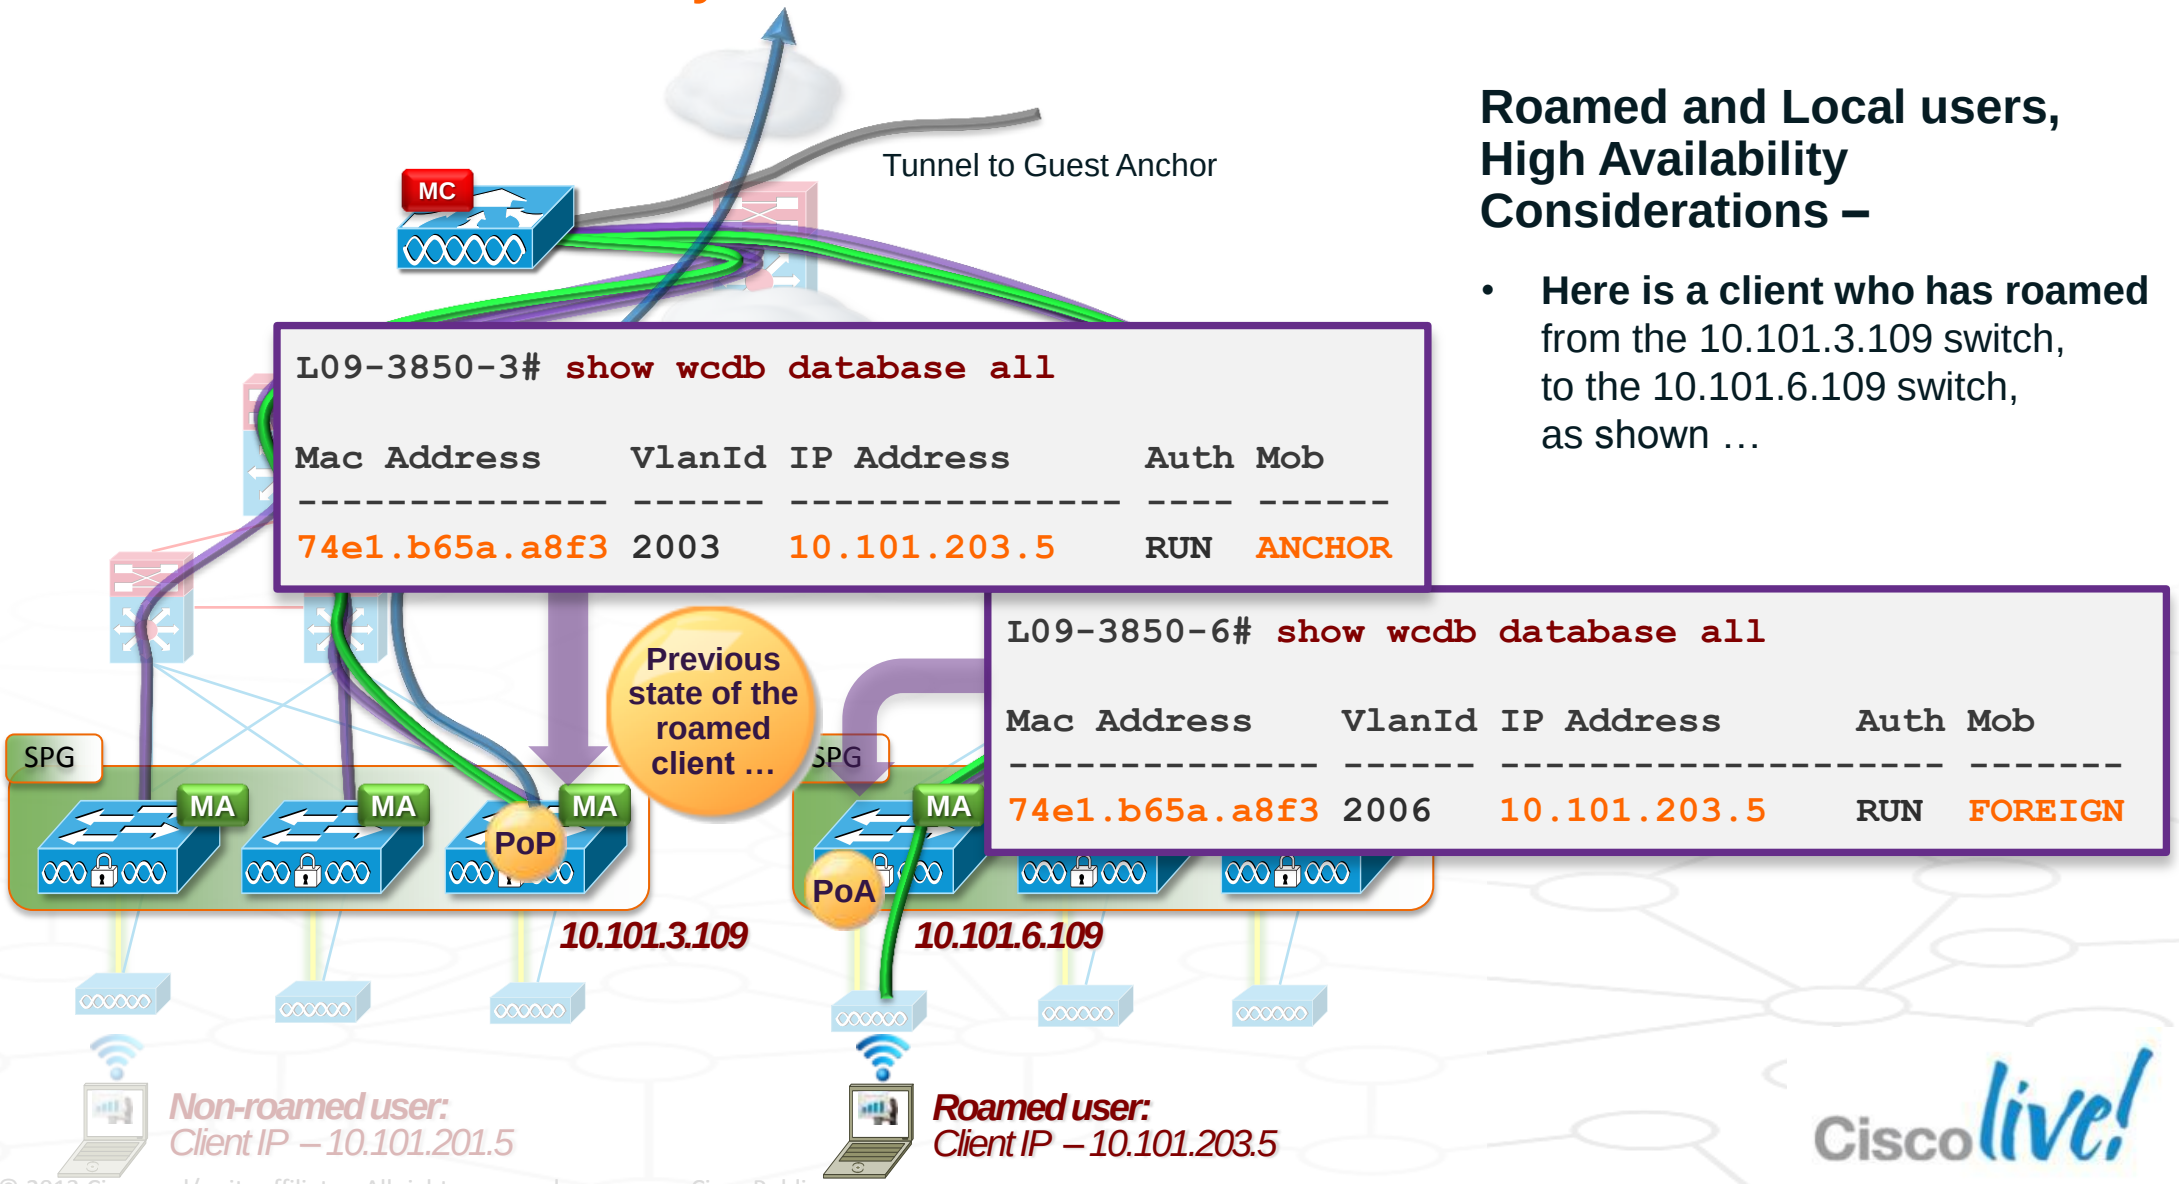

Roamed and Local users, High Availability Considerations –

- **Here is a client who has roamed** from the 10.101.3.109 switch, to the 10.101.6.109 switch, as shown …

High Availability

MC Failure – Effect on Previously-Roamed Clients

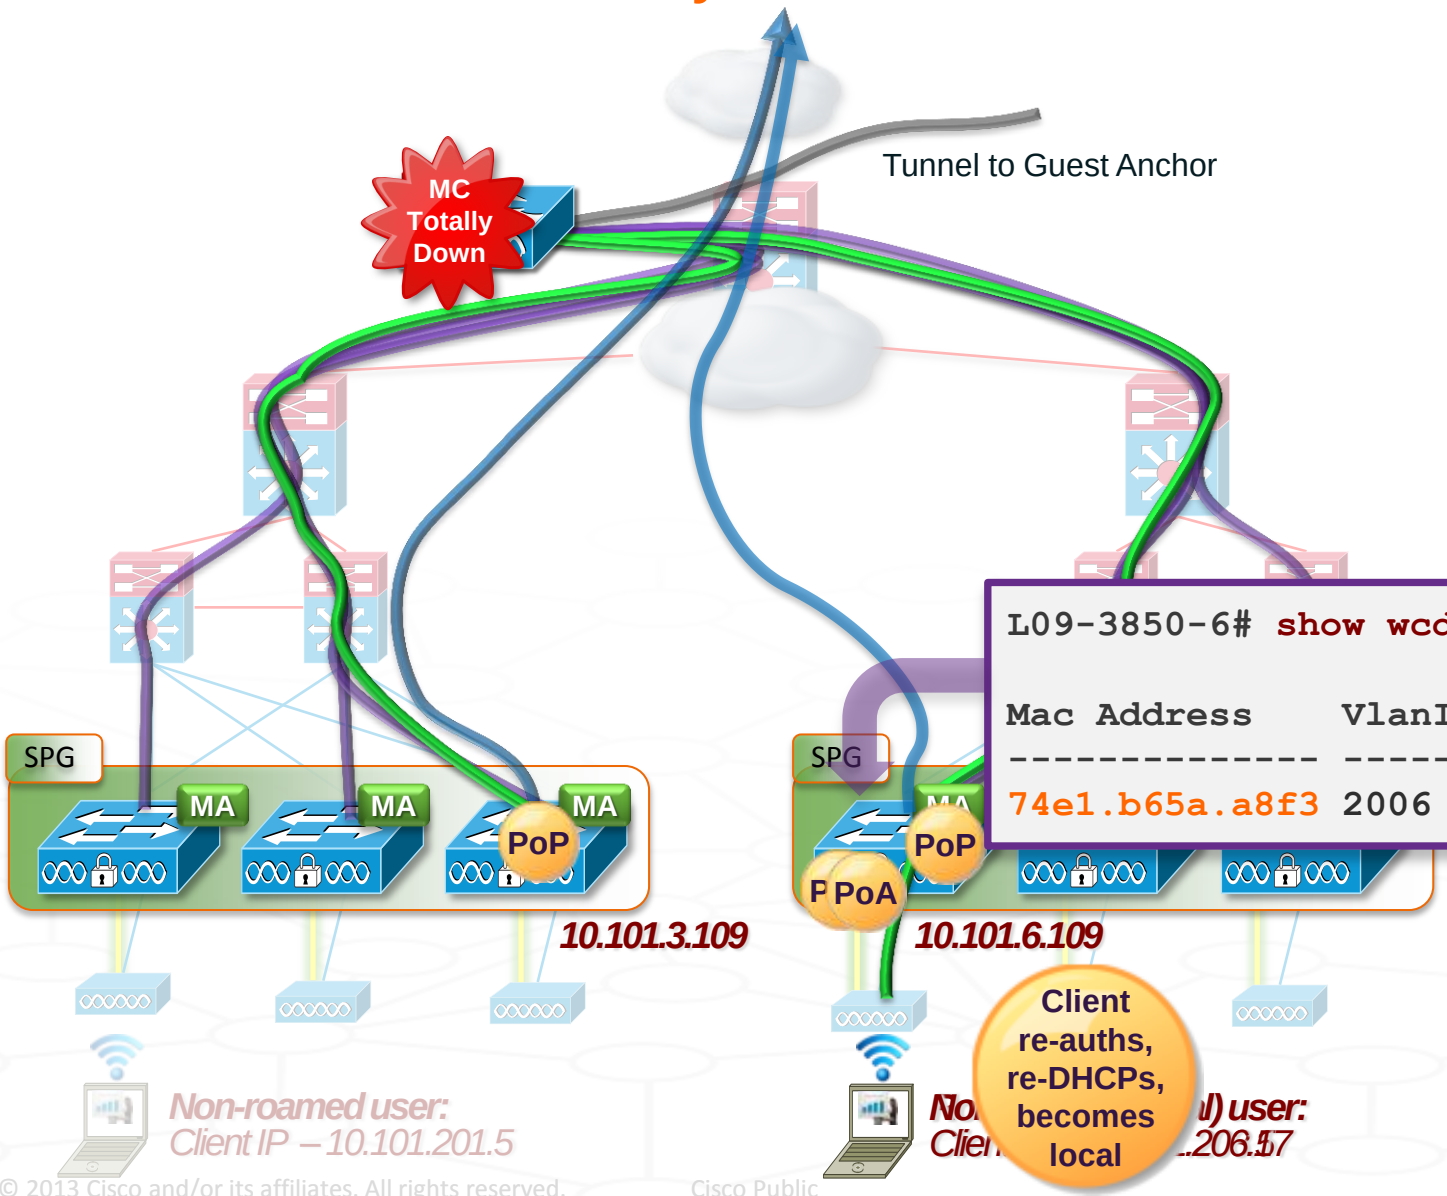


```
L09-3850-6# show wcdb database all

Mac Address     VlanId IP Address          Auth Mob
--------------- ------ ------------------- -------
74e1.b65a.a8f3 2006   10.101.206.17       RUN  LOCAL
```

Roamed and Local users, High Availability Considerations –

- **For a previously-roamed client, the loss of the MC also results in the roamed client "becoming local" on their roamed-to switch**, as shown …

High Availability

MC Failure – Effect on Intra-SPG Client Roams after MC Down

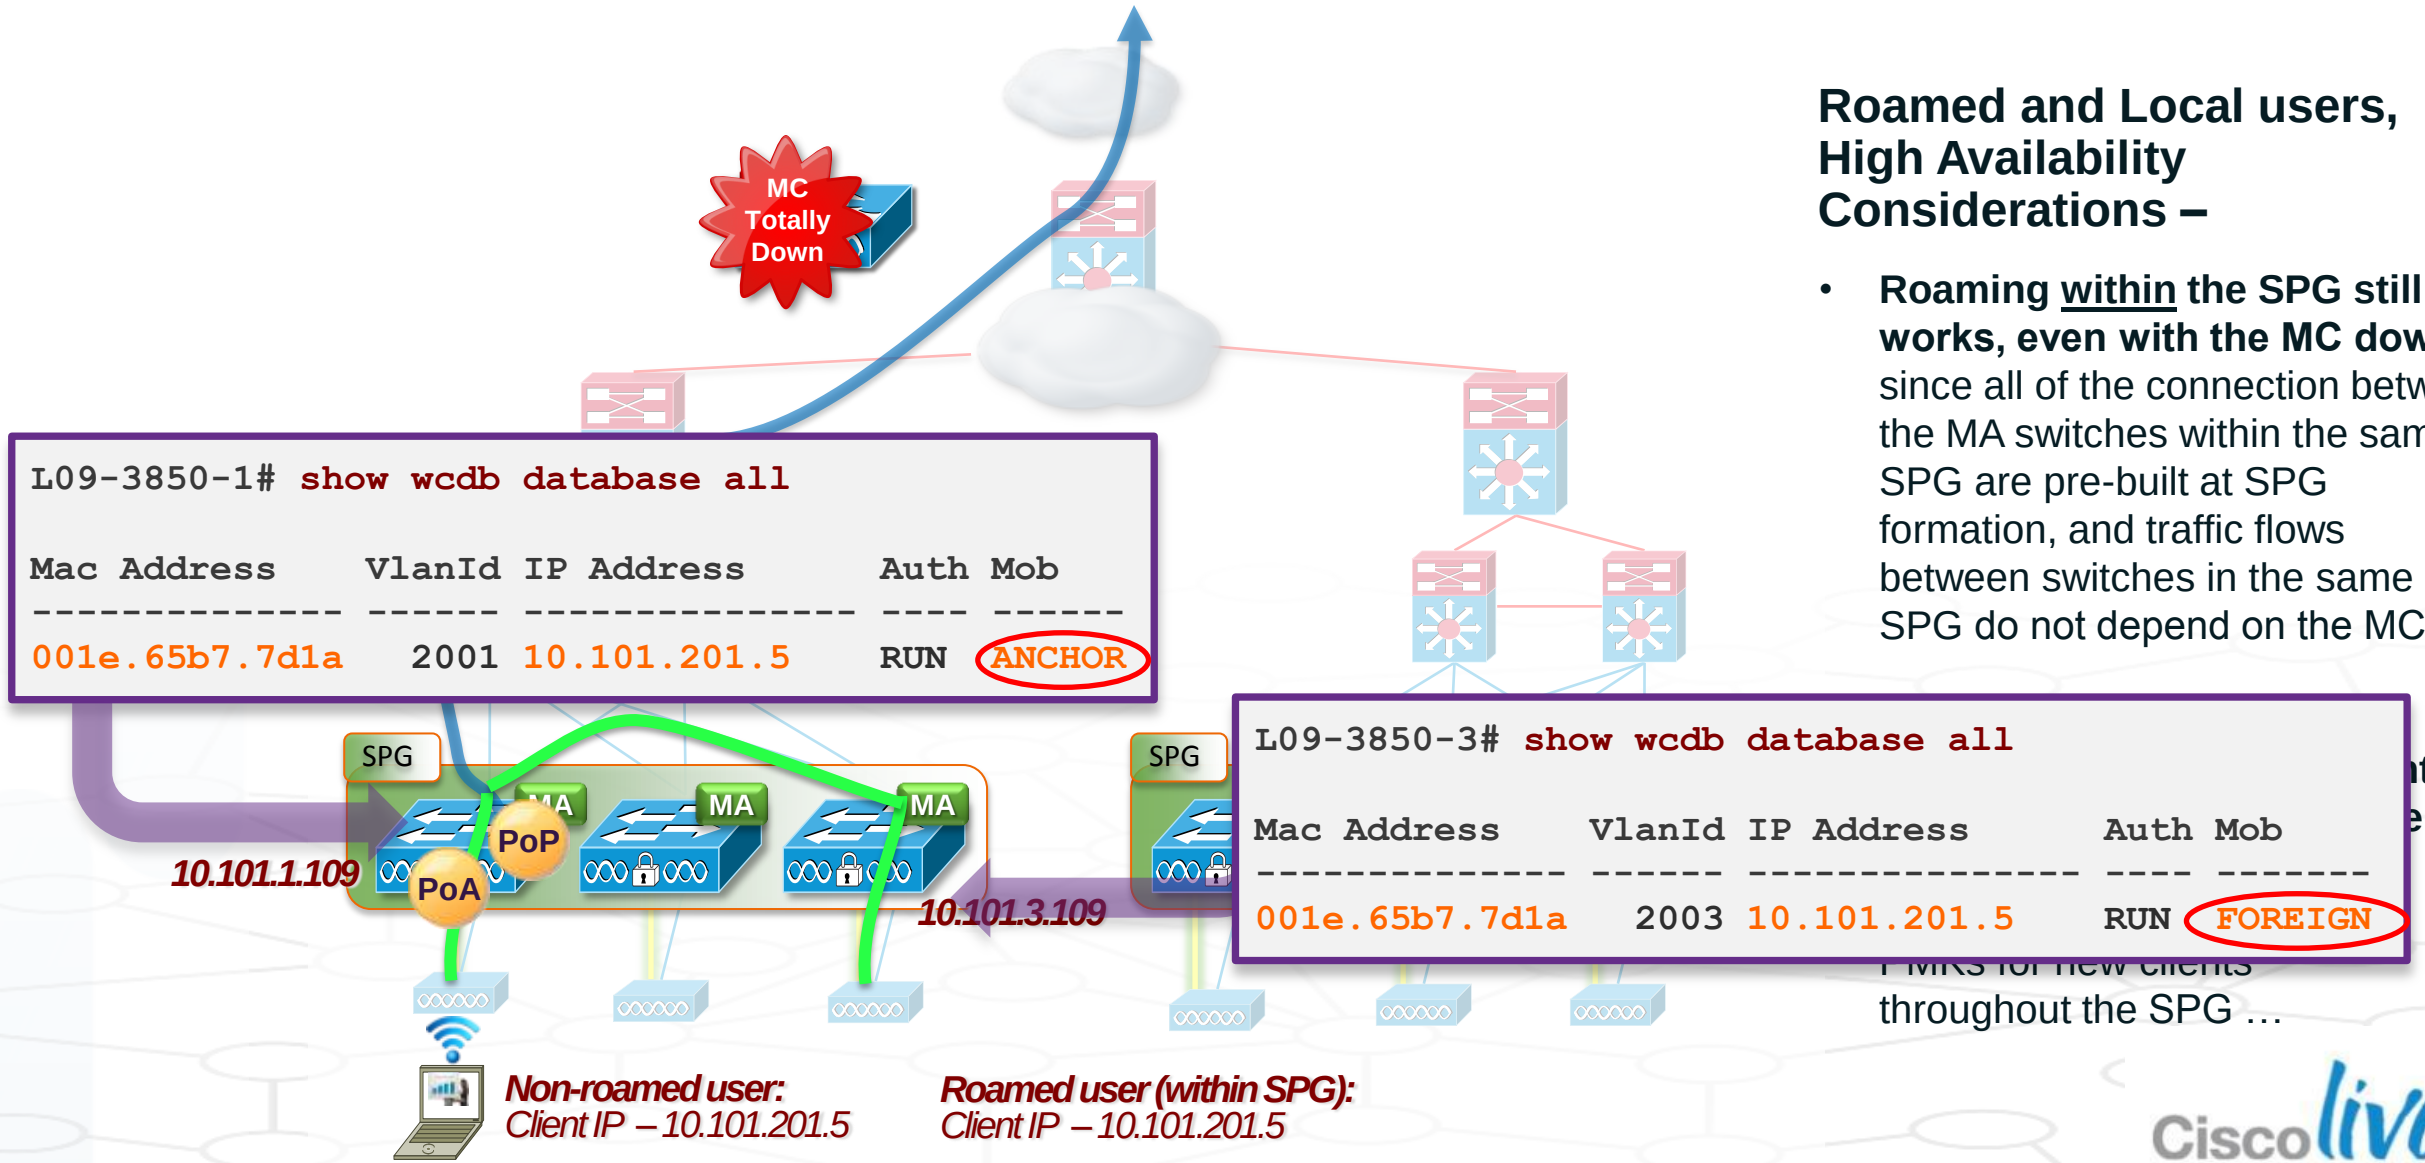

Roamed and Local users, High Availability Considerations –

- **Roaming within the SPG still works, even with the MC down …** since all of the connection between the MA switches within the same SPG are pre-built at SPG formation, and traffic flows between switches in the same SPG do not depend on the MC …

throughout the SPG …

High Availability

MC Failure – Effect on Inter-SPG Client Roams after MC Down

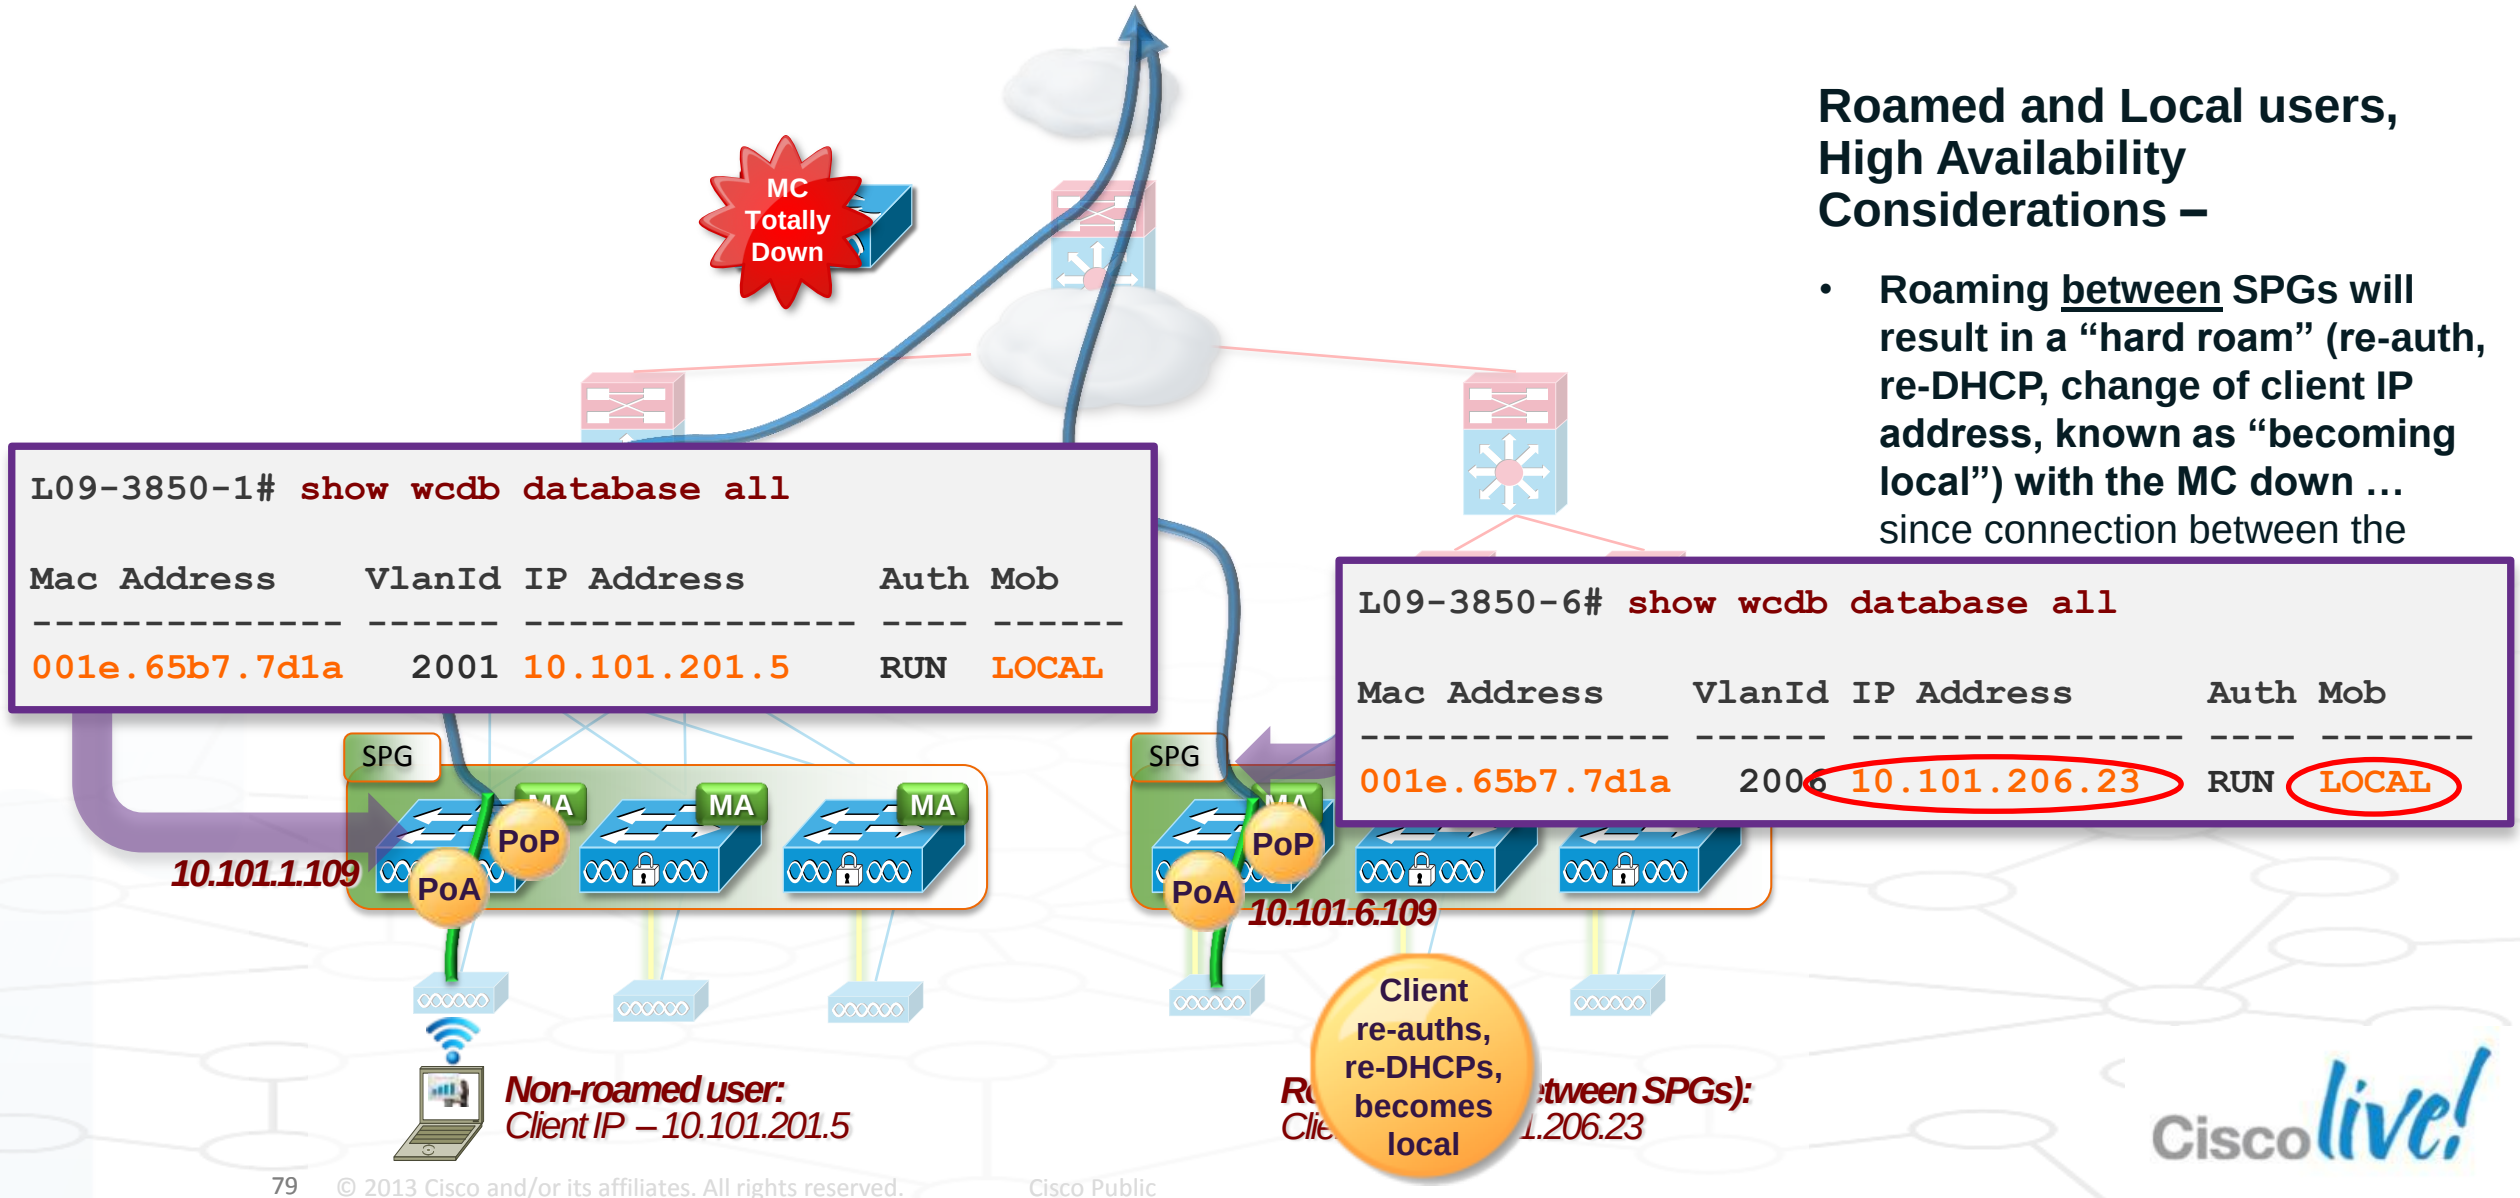

Roamed and Local users, High Availability Considerations –

- **Roaming between SPGs will result in a "hard roam" (re-auth, re-DHCP, change of client IP address, known as "becoming local") with the MC down …** since connection between the

High Availability

MC Redundancy with 1:1 AireOS 7.3 HA – 5508 and WiSM2

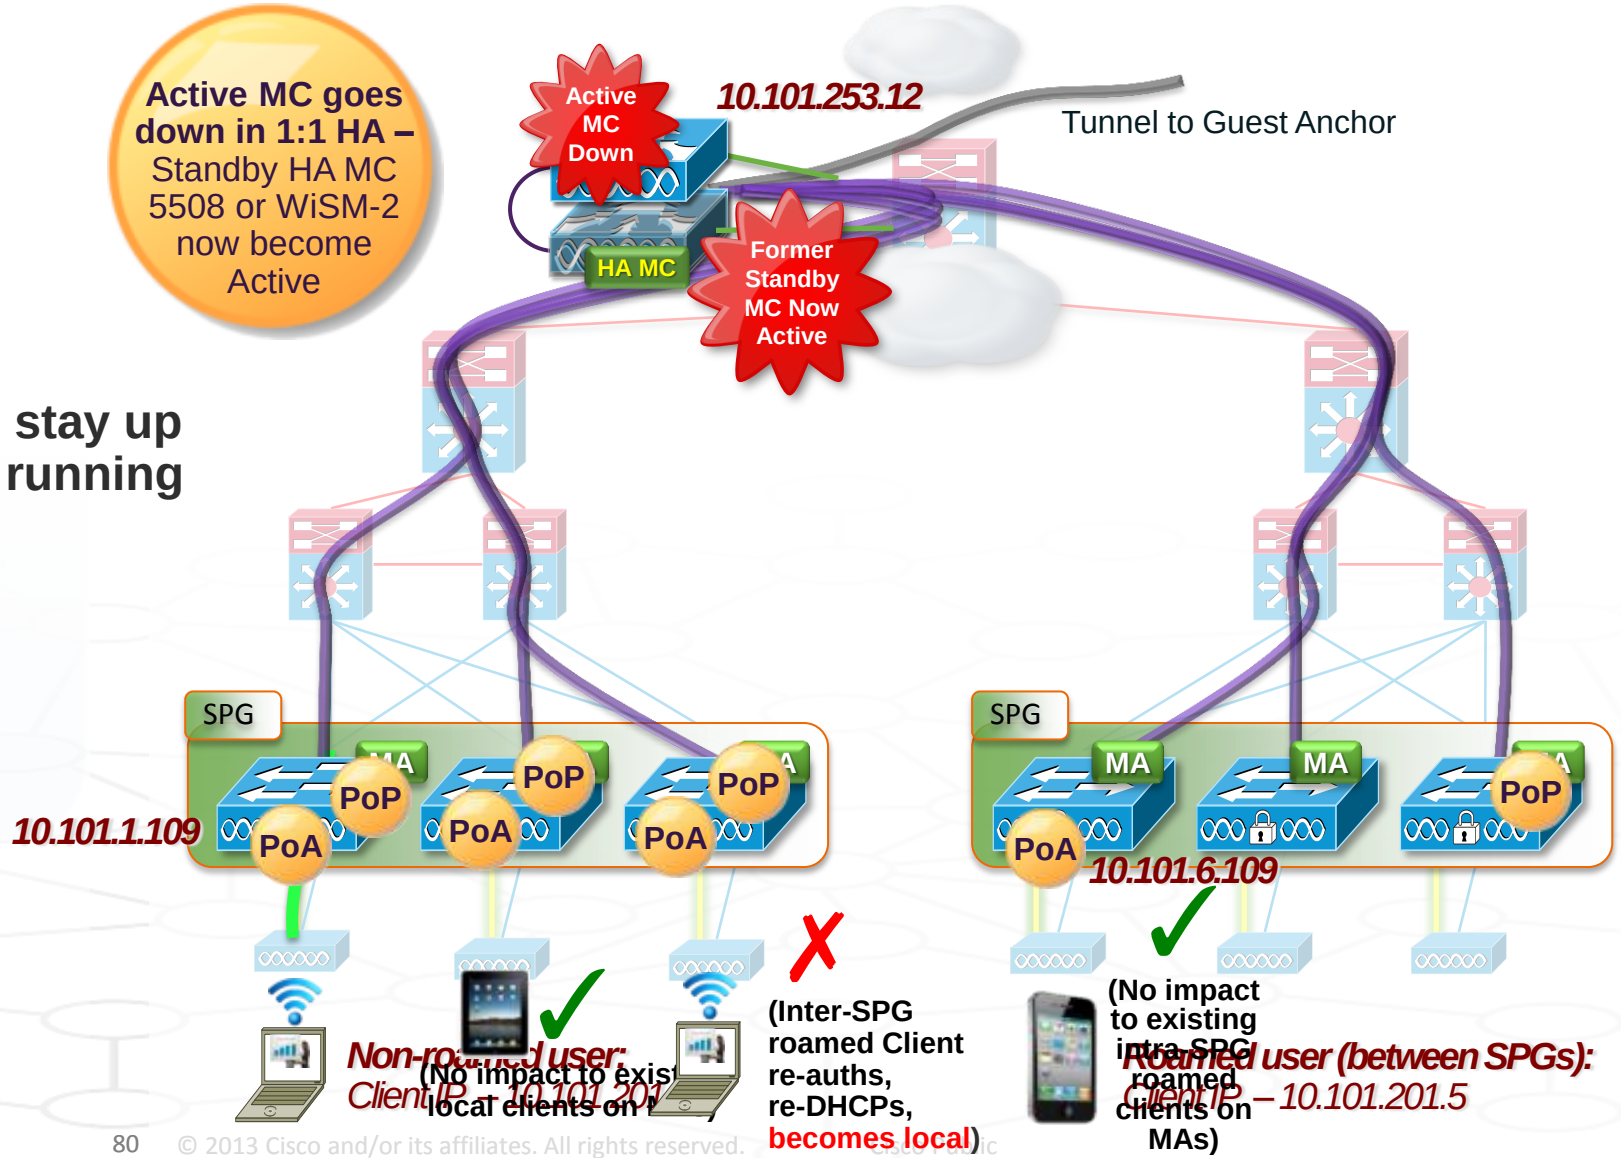

Roamed and Local users, High Availability Considerations –

- **Local users on their MAs** have no impact following a HA MC failover event
- **Intra-SPG roamed users** also have no impact following the MC HA failover
- **All previously-roamed clients (inter-SPG) will result in a "hard roam" after MC failover** (re-auth, re-DHCP, change of client IP address, known as "becoming local")
- **Any new intra-SPG or inter-SPG roaming happening after MC HA failover** from local MA clients will be handled normally

High Availability

MC Redundancy with 1:1 AireOS 7.3 HA – 5508 and WiSM2

All Existing AireOS '7.3 HA' features Maintained –

- Box-to-Box High Availability = 1:1
- One WLC in Active state and Second WLC in Hot Standby State – Standby monitors the health of the Active WLC
- Configuration on Active is synched to Standby WLC via Redundant Port
- Both the WLCs share the same set of configurations, including the IP address of the Management Interface
- AP SSO - APs CAPWAP state (only APs which are in RUN state) is also synched … this state is retained for APs hosted directly by the WLC 5508 or WiSM-2
- APs do not go in Discovery state when Active WLC fails
- Downtime between failover reduced to 5 - 996 msec in case of Box failover and up to 3 seconds in case of Network Issues
- 1:1 HA Licensing Model – Supported on WLC 5500 and WiSM-2 ONLY

High Availability

HA Connectivity on WLC 5508

How It Works –

- 5508 WLC have dedicated Redundancy Port which is used to synchronise the configuration from the Active to the Standby WLC
- Keepalives are sent on RP port from Standby to Active WLC every 100 msec (default timer) to check the health of the Active WLC
- ICMP packets are also sent every one second from each WLC to check reachability to gateway using the Redundant Management interface

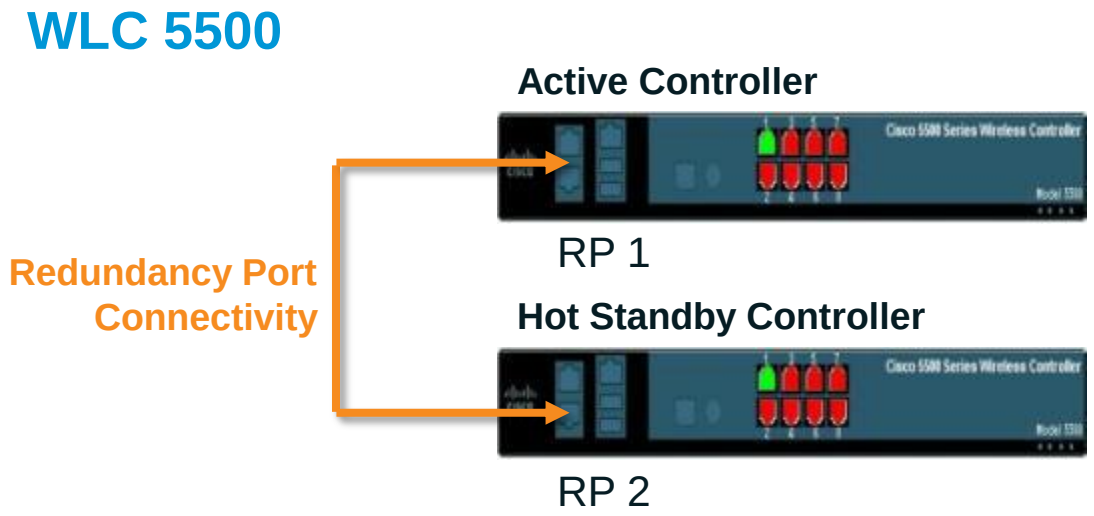

High Availability

HA Connectivity on WiSM2, Multi-Chassis

How It Works –

- When using WiSM-2s, HA can be deployed in single chassis OR can also be deployed between multiple chassis using VSS
- WiSM-2 has a dedicated Redundancy VLAN which is used to sync the configuration from the Active to the Standby WLC
- A redundancy VLAN should be a non-routable VLAN, meaning a Layer 3 interface should not be created for this VLAN
- Keepalives are sent on the Redundancy VLAN … in this case VLAN 169
- Using this VLAN, keepalives are sent from the Standby to the Active WLC every 100 msec (default timer) to check the Active WLC's health

Multi Chassis Connectivity

WISM2 configuration on Cat6k

wism service-vlan 192 (service port Vlan)
wism redundancy-vlan 169(redundancy port Vlan)
wism module 6 controller 1 allowed-vlan 24-38(data vlan)

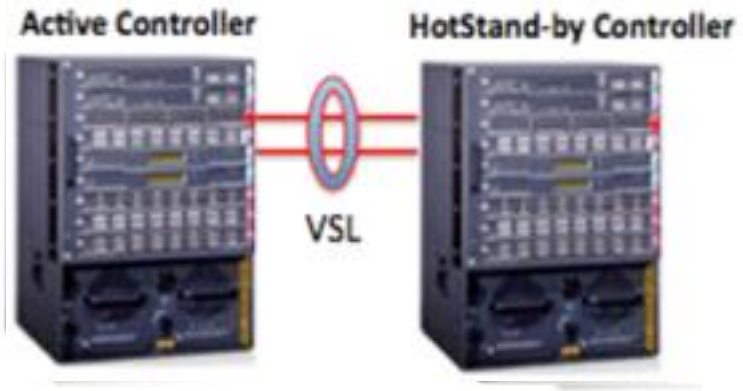

Single Chassis Connectivity

Slot 8: Active WiSM-2
Slot 9: Hot Standby WiSM-2

High Availability

Catalyst 3850-based MCs – Fault Tolerance in Stack

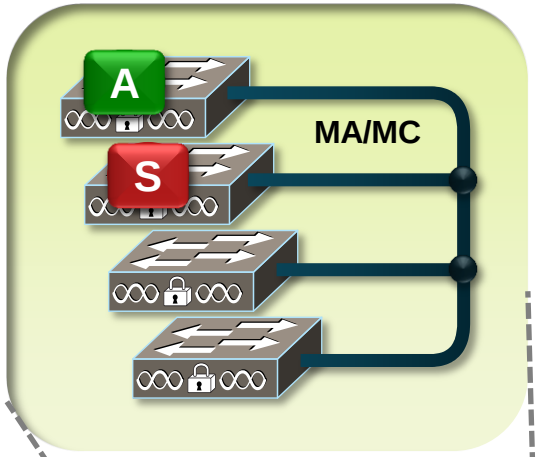

Examining state within the stack (for MC) –

- **Let's now examine the state maintained by the MC within a stack,** and see what redundancy we provide for this …

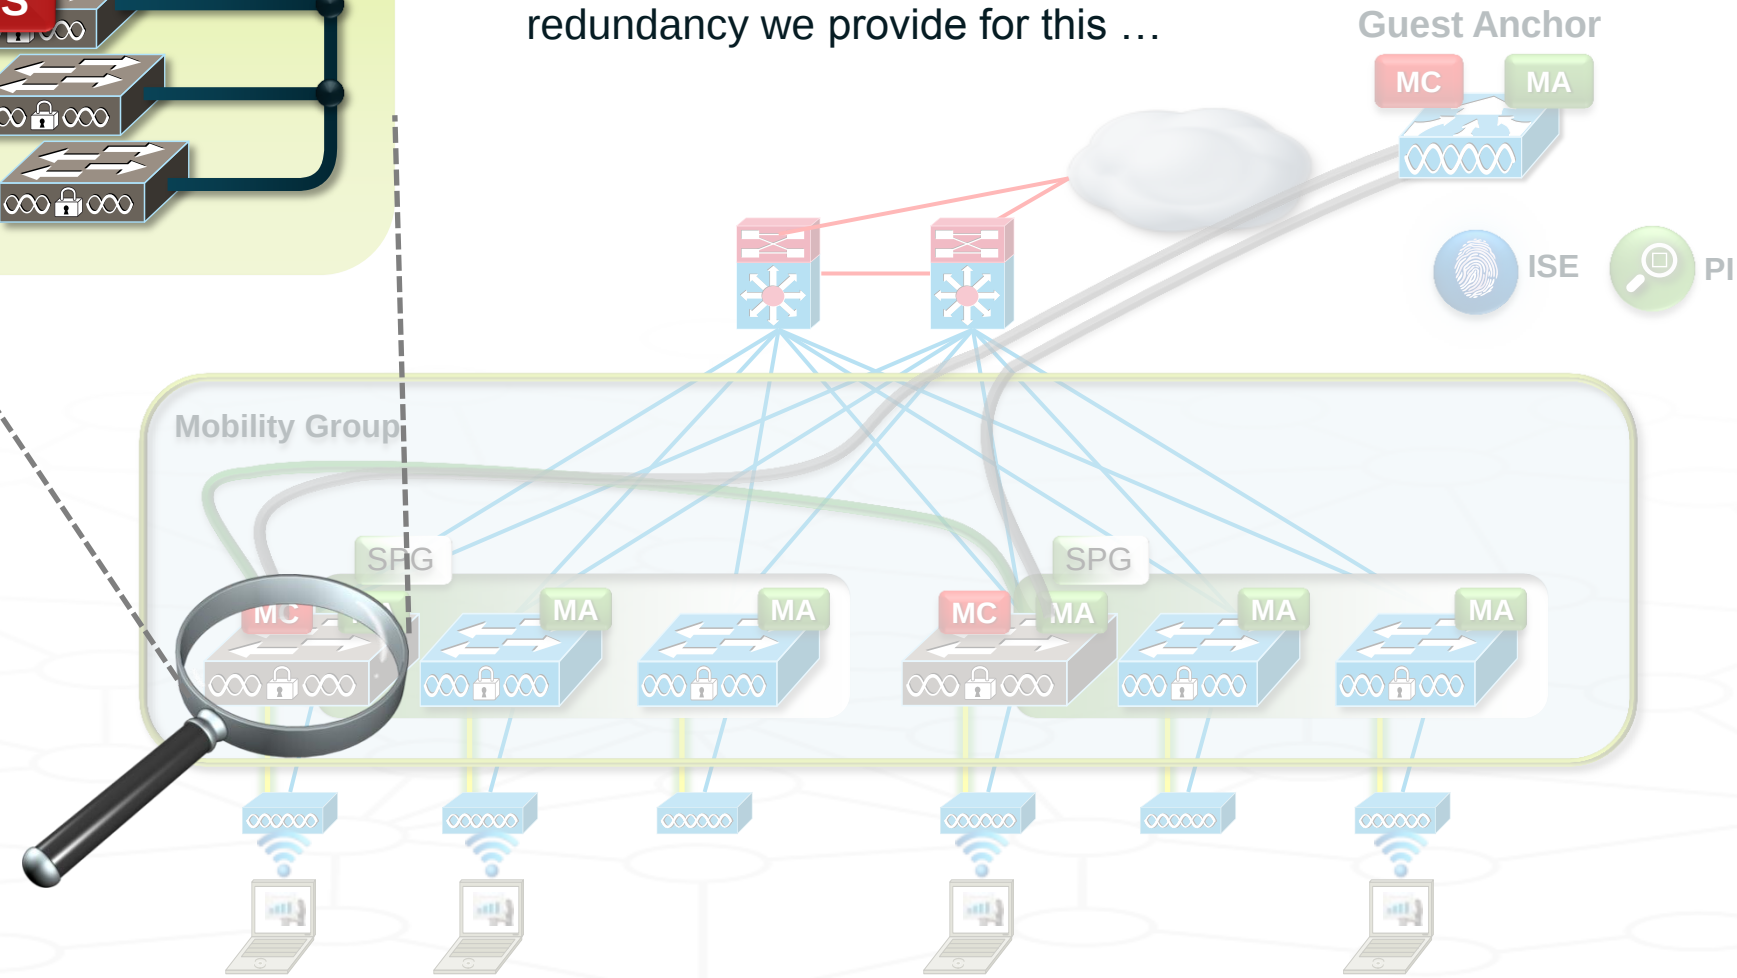

High Availability

Catalyst 3850-based MCs – Tunnel SSO

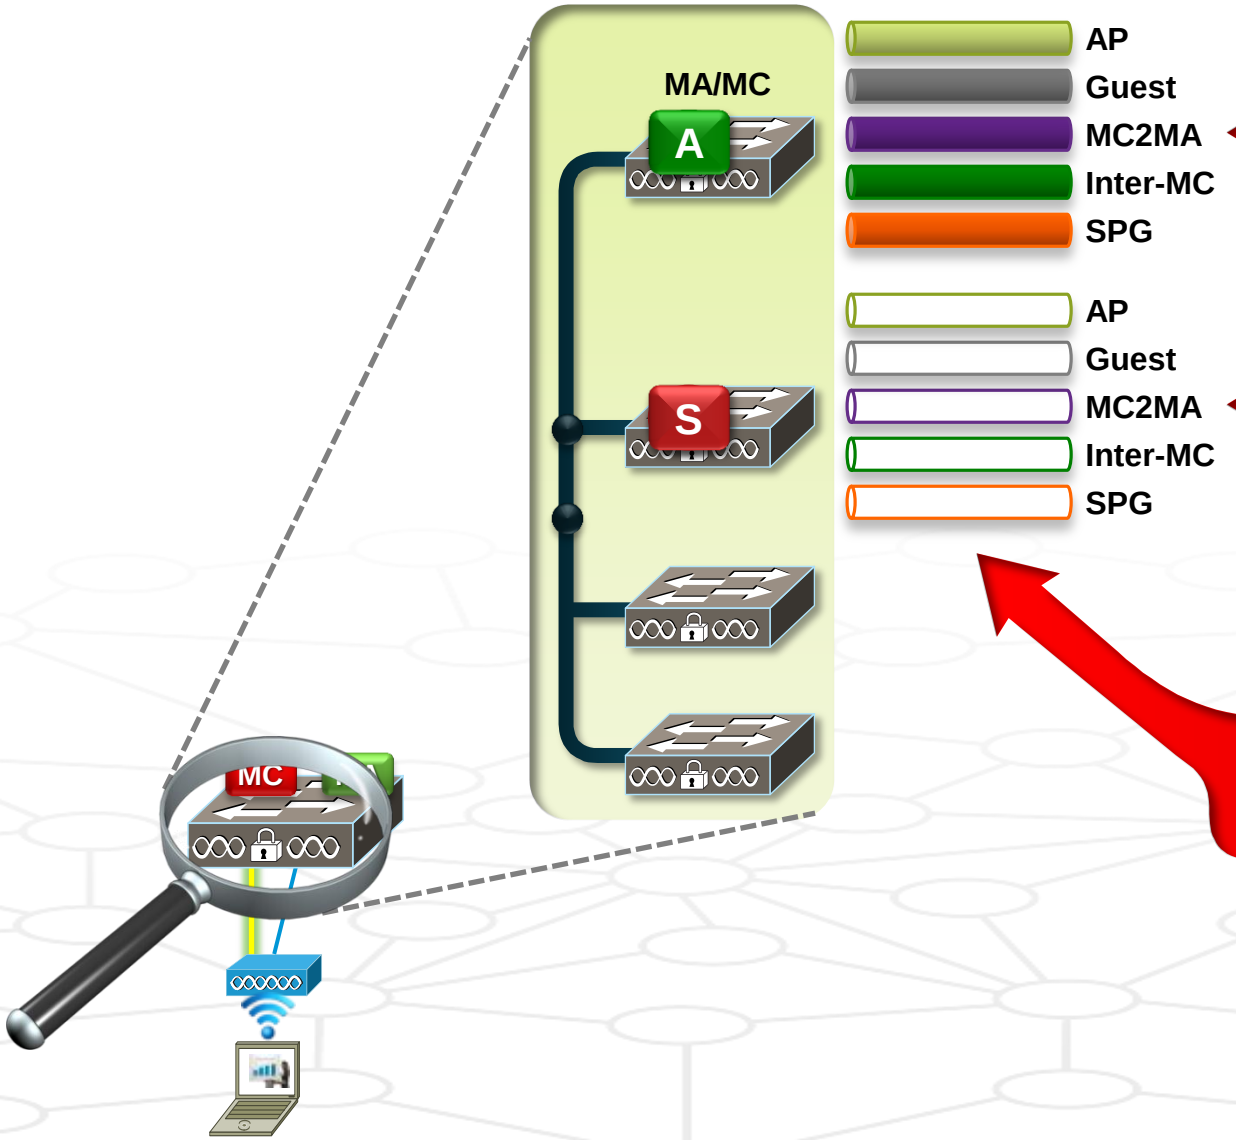

Tunnel State is synced between Active and Standby Member in stack

Tunnel States are inactive on Standby Member

High Availability

Catalyst 3850-based MCs – Fault Tolerance in Stack

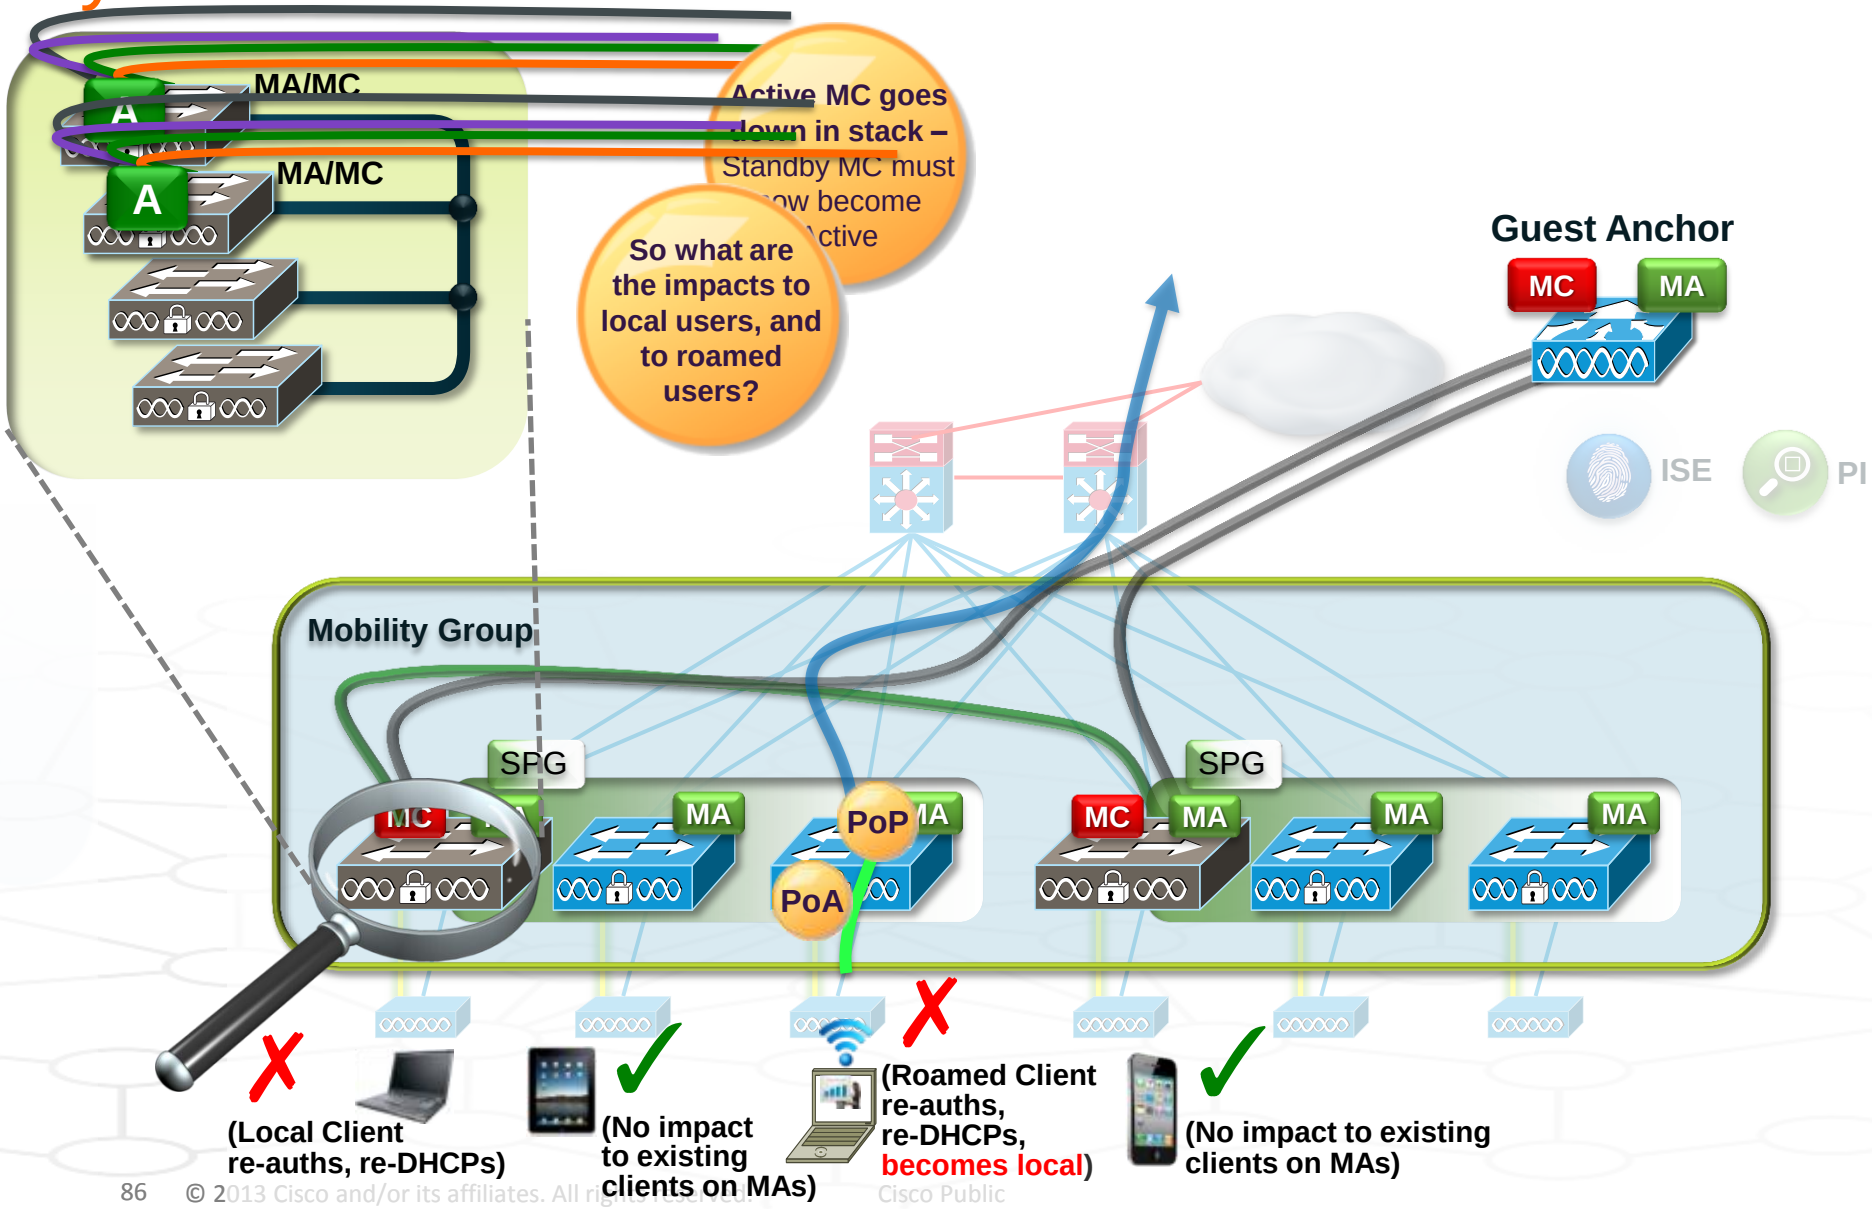

High Availability

Catalyst 3850-based MCs – Fault Tolerance across Stacks

Switch Peer Group Fault Tolerance with Catalyst 3850 –

- **If an Catalyst 3850-based stack, operating as an MC, completely goes down** in a Switch Peer Group –
  - **Roaming within a Switch Peer Group** still works seamlessly
  - **Roaming between Switch Peer Groups** does not work (re-DHCP)
  - **PMKs (via PKC) will not be distributed if the MC is down –** so no Fast Roaming for new clients until the MC is restored

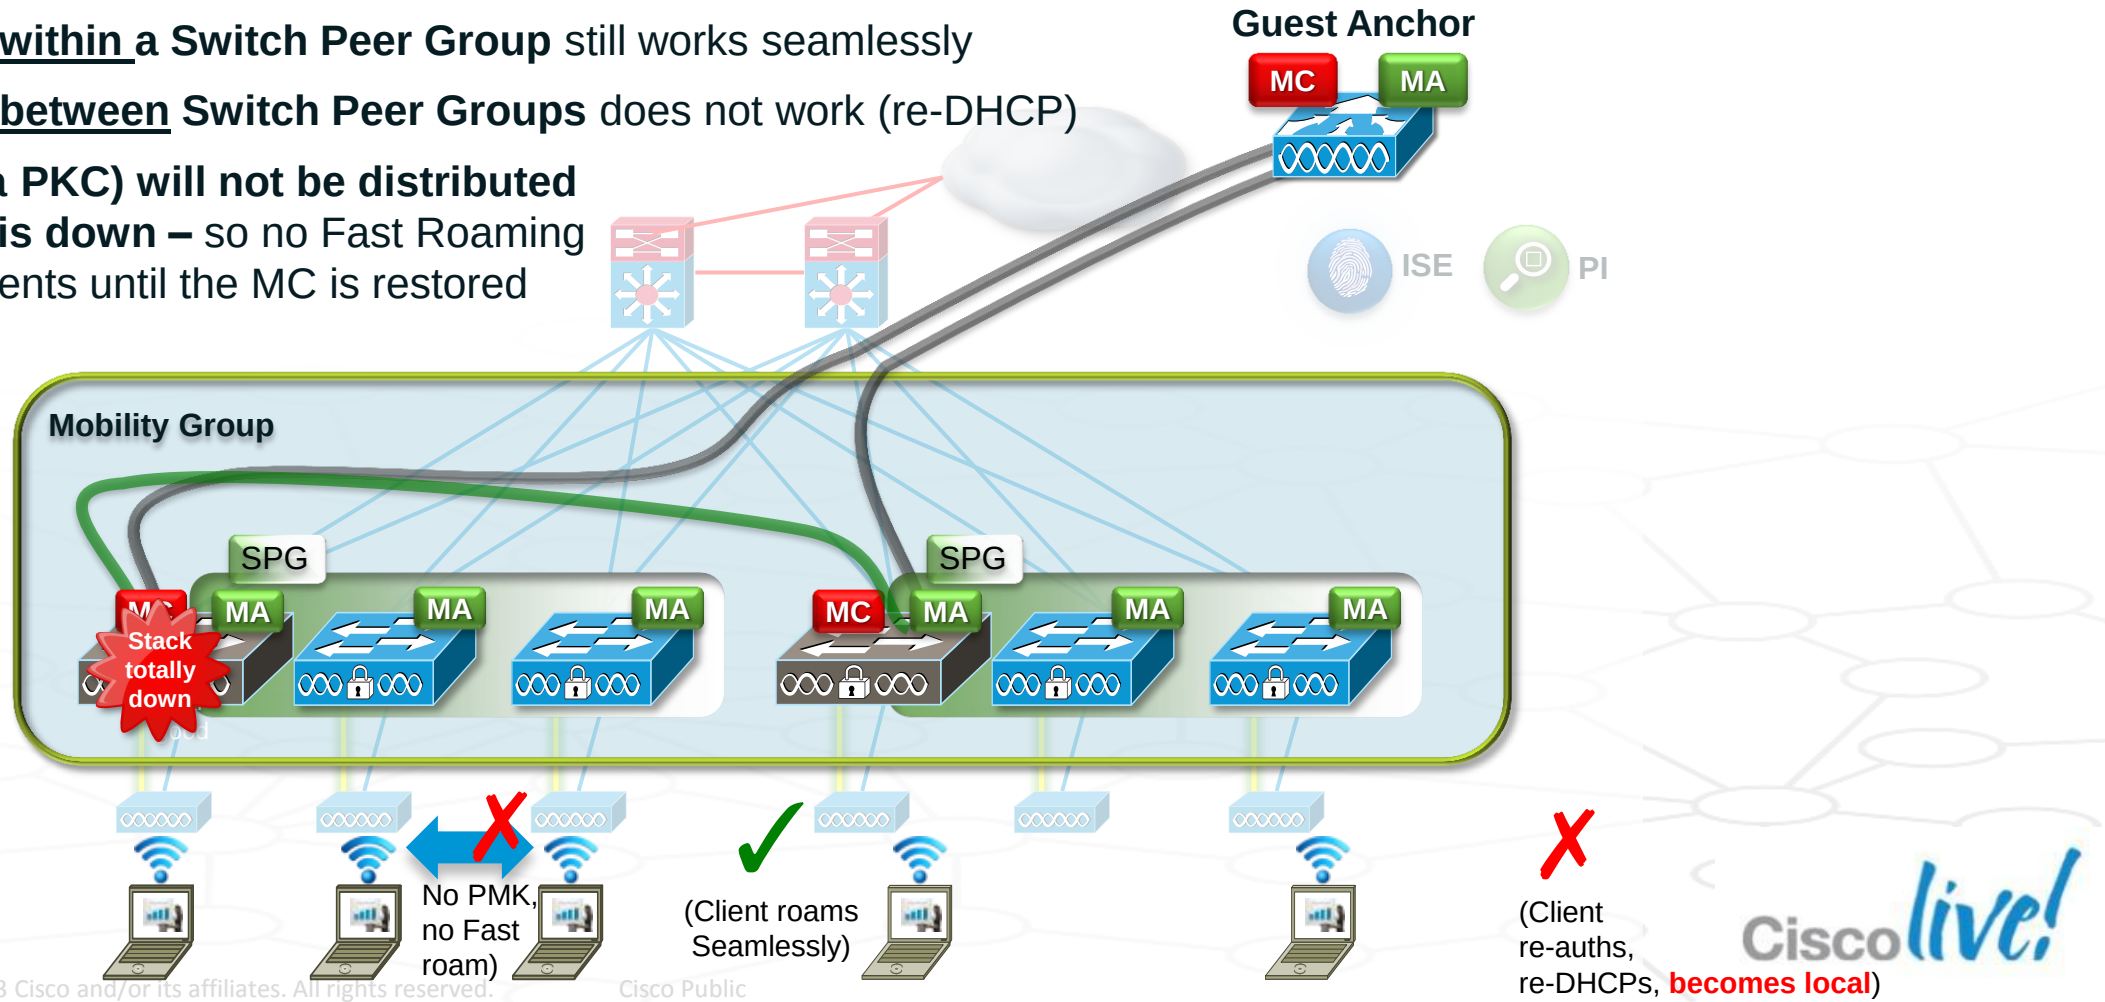

High Availability

Catalyst 3850-based MCs – Fault Tolerance across Stacks

Switch Peer Group Fault Tolerance with Catalyst 3850 –

- **If an Catalyst 3850-based MC is completely down** in a Switch Peer Group –
  - **When the MC is down**, RRM, Guest Access, (guest tunnelling) and other MC-based functions do not operate within the affected Switch Peer Group – other Switch Peer Groups are unaffected, however

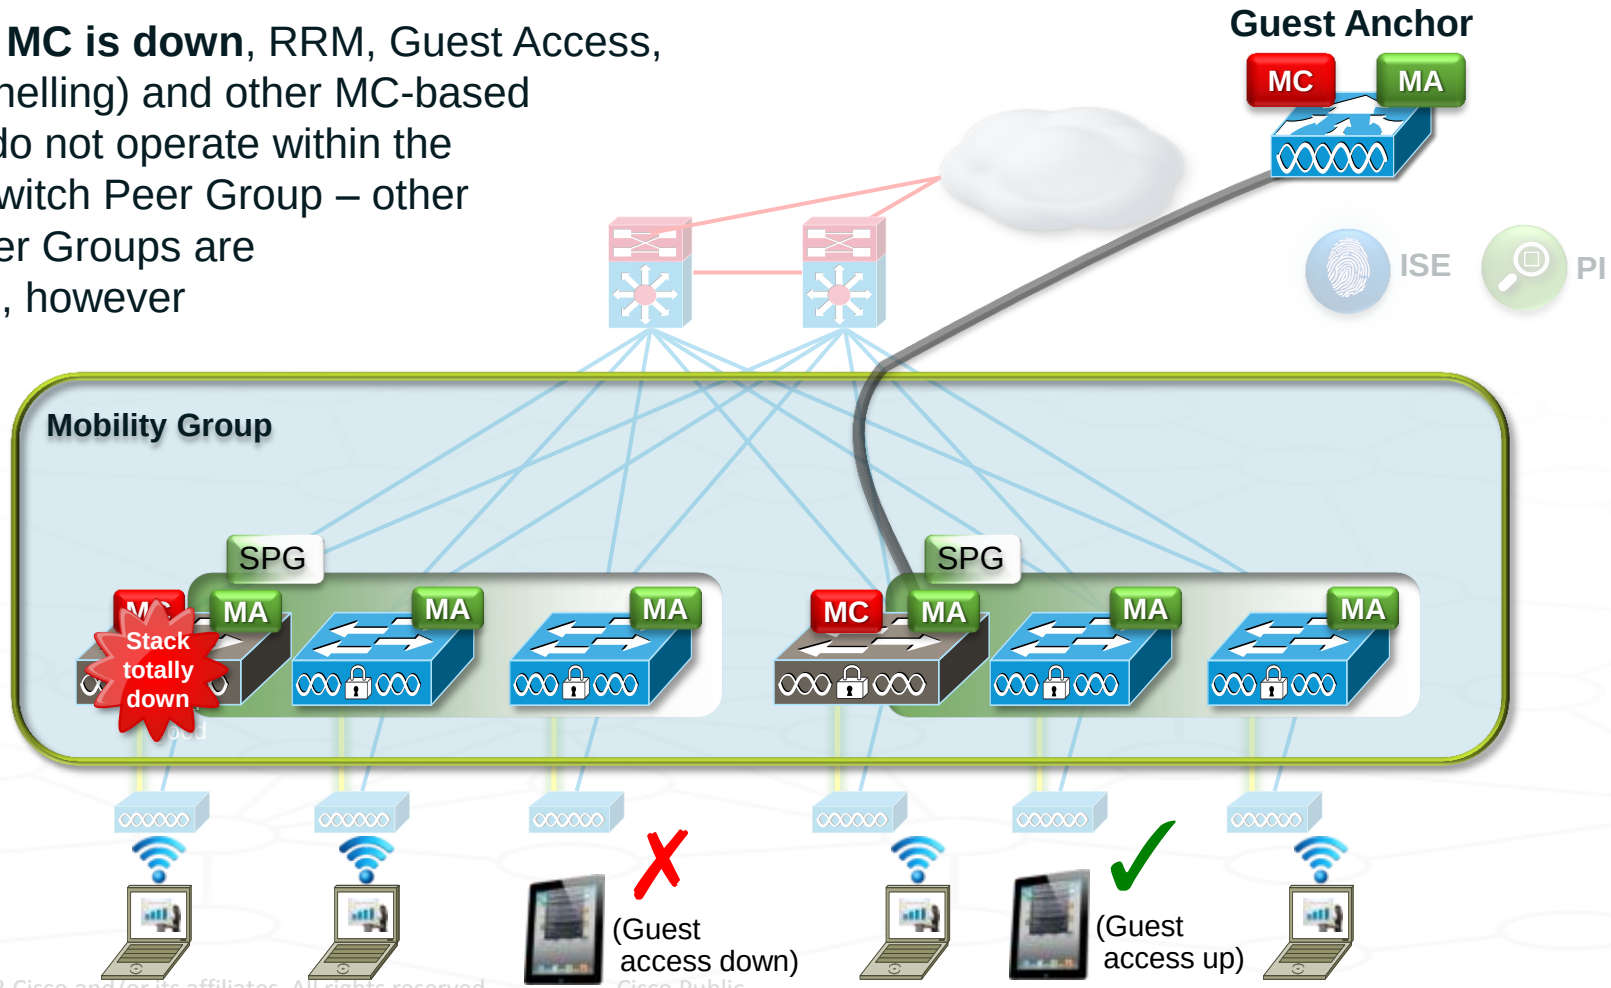

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a "double-click" deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – **Quality of Service**

Summary

What is QoS Made of?

- Classification
- Marking/Mutation
- Shaping/Policing
- Queuing
- Bandwidth Allocation
- Trust

What is QoS Made of?

- Classification
- Marking/Mutation
- Shaping/Policing
- Queuing
- Bandwidth Allocation
- Trust

DSCP 46

CUWN Architecture

Overview – Challenges of QoS

Current Mobility Architecture

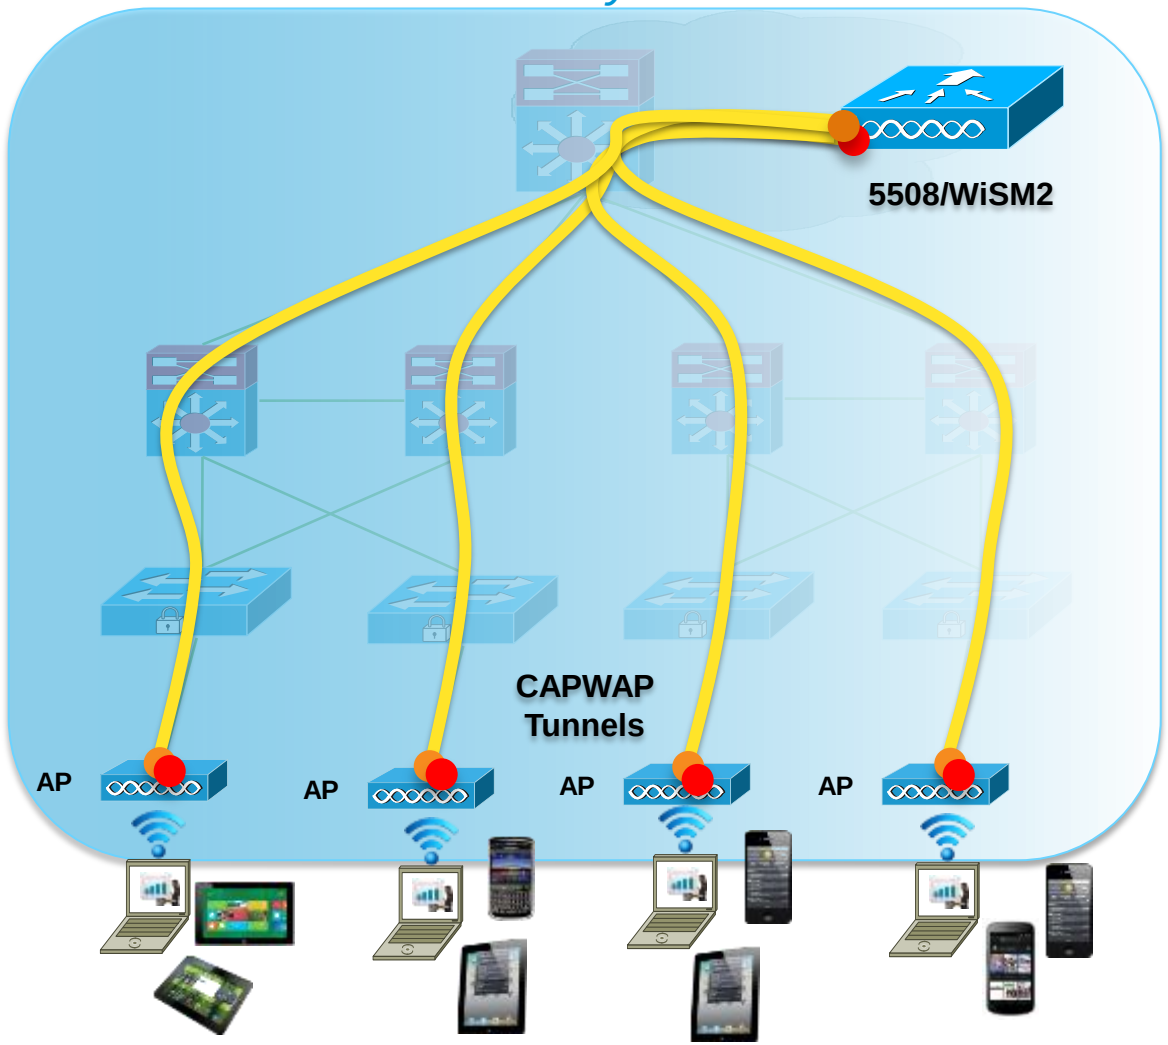

Challenges –

Overlay model with multiple points of policy application*

Limited **visibility** into applications

Lack of **granular classification**

Software based **QoS**

*** Overlay model applies to CUWN local mode and FlexConnect centralised mode**

Marking Policing

Existing QoS Deployments

How We Overlay QoS Policies Today

Current QoS Architecture

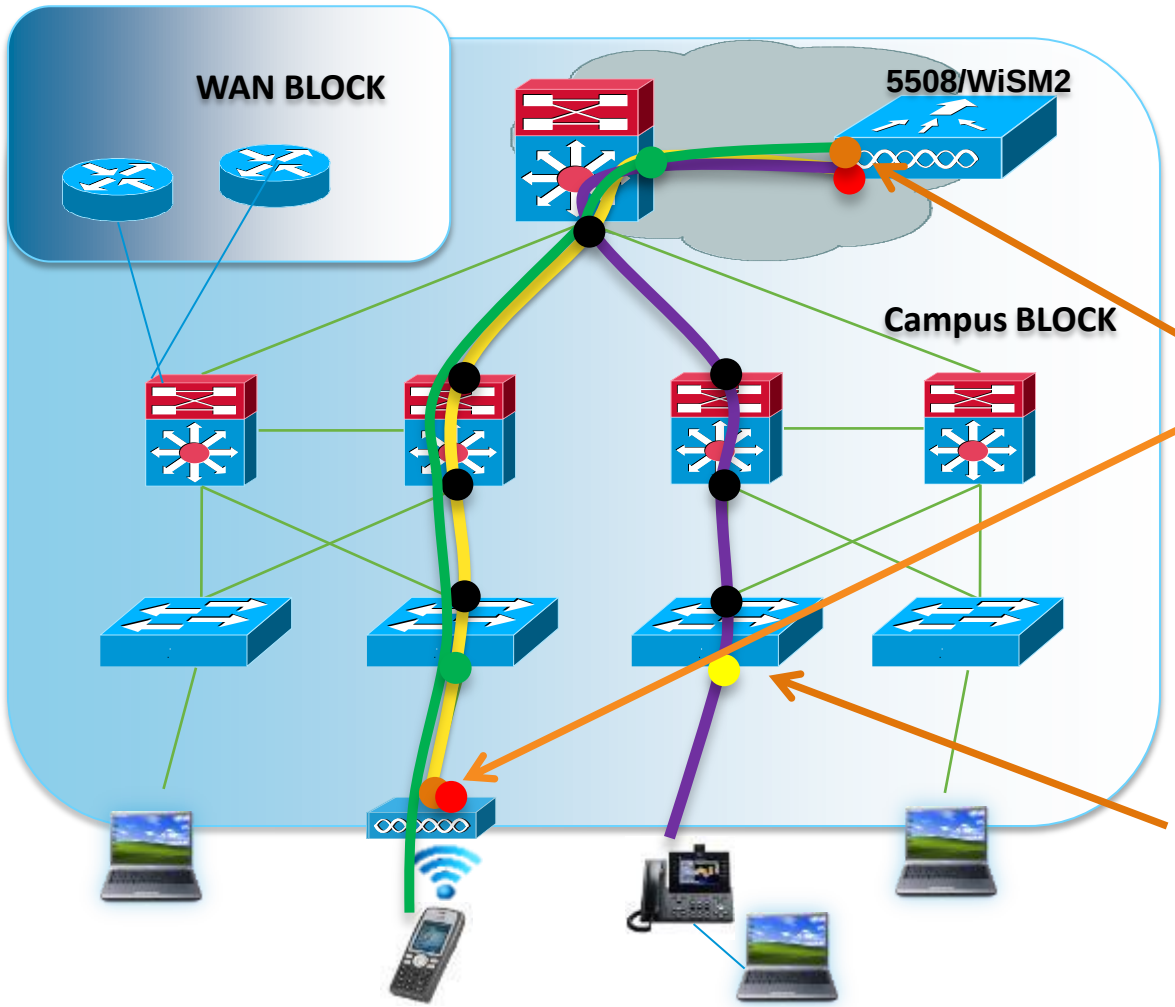

Distributed Management Configuration and Deployment

Separate policies and services for wired and wireless users

Wireless policies implemented on controller pushed to AP

Wired policies implemented on switch

Marking Policing Queuing

QoS – What's New with Converged Access

Wired (Cat 3850)

- Modular QoS based CLI (MQC)
  - Alignment with 4500E series (Sup6, Sup7)
  - Class-based Queueing, Policing, Shaping, Marking
- More Queues
  - Up to 2P6Q3T queuing capabilities
  - Standard 3750 provides 1P3Q3T
  - Not limited to 2 queue-sets
  - Flexible MQC Provisioning abstracts queuing hardware

Wireless(Cat 3850 & CT 5760)

- Granular QoS control at the wireless edge
  - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- Enhanced Bandwidth Management
  - Approximate Fair Drop (AFD) Bandwidth Management ensures fairness at Client, SSID and Radio levels for NRT traffic
- Wireless Specific Interface Control
  - Policing capabilities Per-SSID, Per-Client upstream*** and downstream
  - AAA support for dynamic Client based QoS and Security policies
- Per SSID Bandwidth Management

*** **NOT** available on CT 5760 at FCS

QoS – What's New with Converged Access

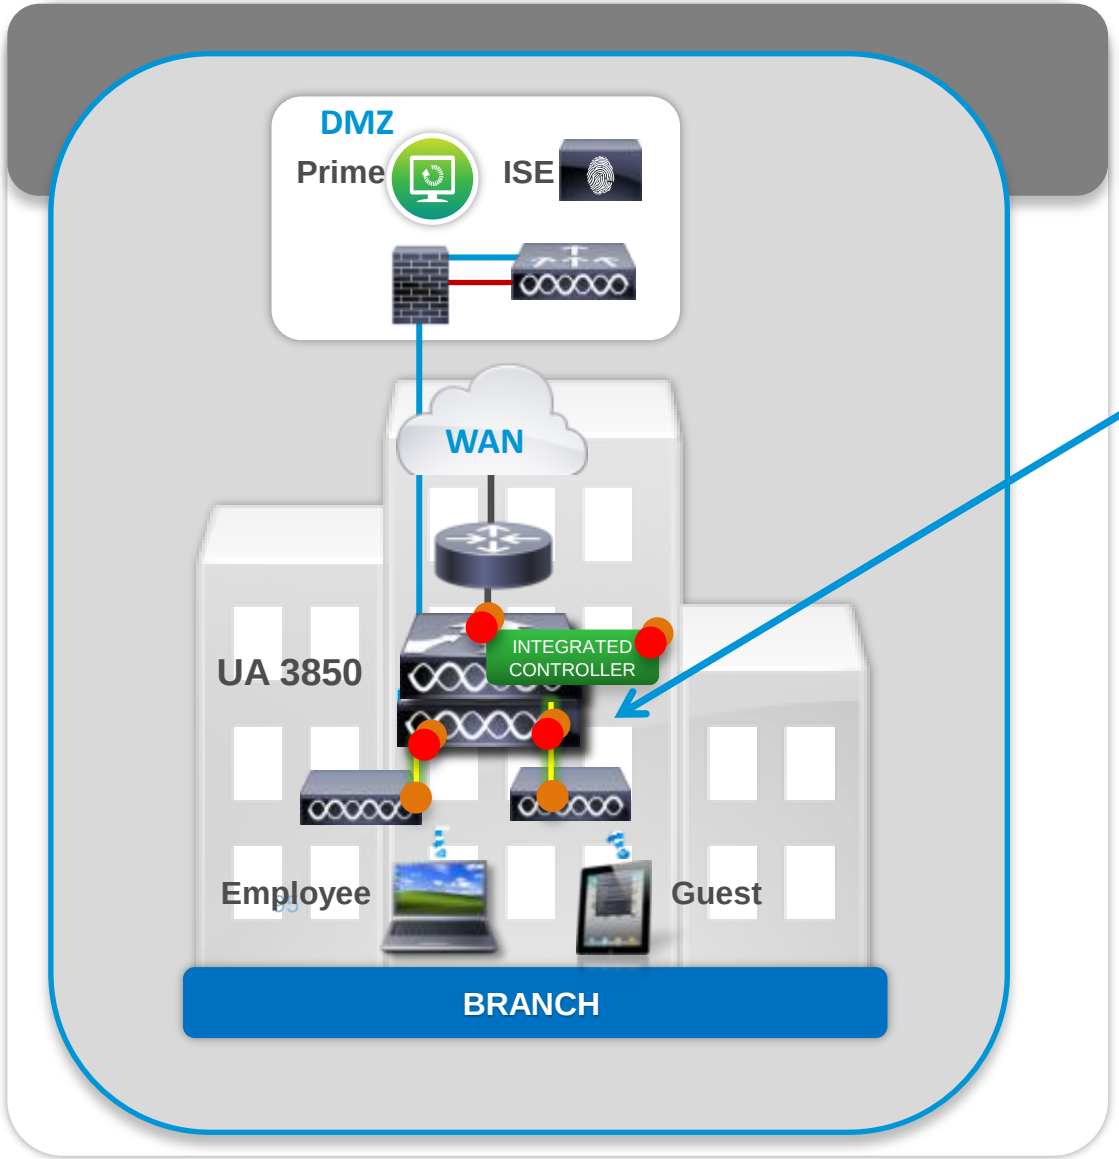

Wireless(Cat 3850 & CT 5760)

- **Granular QoS control at the wireless edge**
  - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- Enhanced Bandwidth Management
  - Approximate Fair Drop (AFD) Bandwidth Management ensures fairness at Client, SSID and Radio levels for NRT traffic
- Wireless Specific Interface Control
  - Policing capabilities Per-SSID, Per-Client upstream*** and downstream
  - AAA support for dynamic Client based QoS and Security policies
- Per SSID Bandwidth Management

*** **NOT** available on CT 5760 at FCS

QoS – What's New with Converged Access

With the CT 5760 or CAT 3850
Usage based fair allocation without configuration

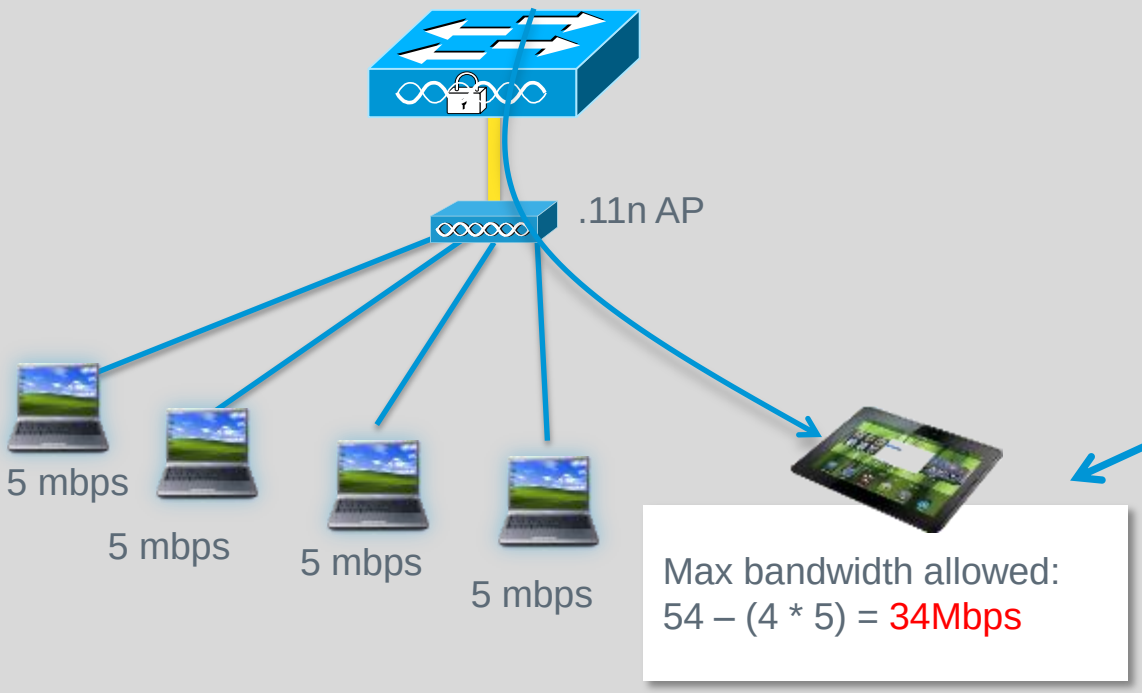

Wireless(Cat 3850 & CT 5760)

- Granular QoS control at the wireless edge
  - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- **Enhanced Bandwidth Management**
  - Approximate Fair Drop (AFD) Bandwidth Management ensures fairness at Client, SSID and Radio levels for NRT traffic
- Wireless Specific Interface Control
  - Policing capabilities Per-SSID, Per-Client upstream*** and downstream
  - AAA support for dynamic Client based QoS and Security policies
- Per SSID Bandwidth Management

*** **NOT** available on CT 5760 at FCS

QoS – What's New with Converged Access

With the 3850
Bidirectional policing at the edge per- user , per-SSID and in Hardware

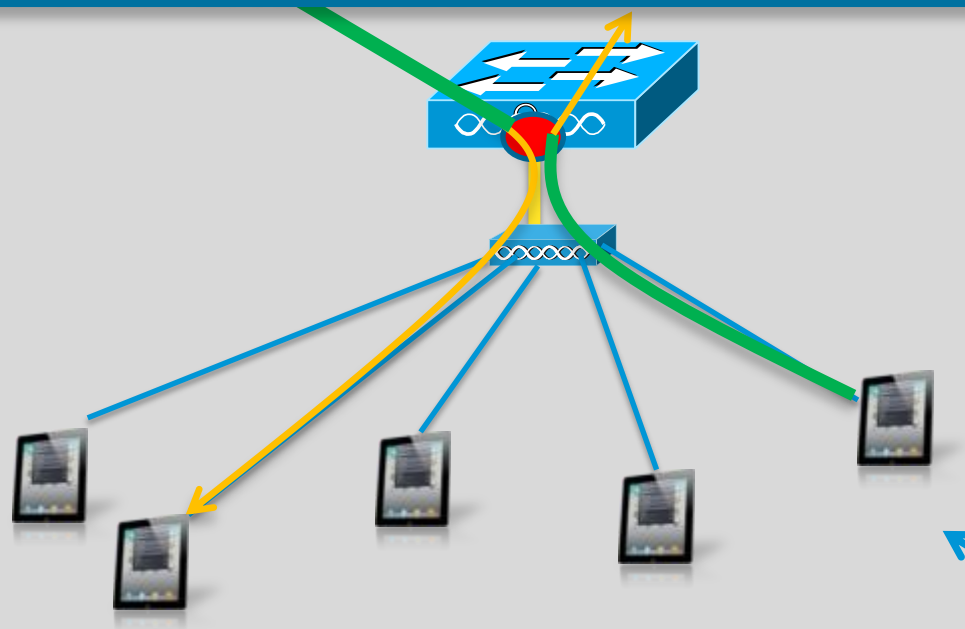

- SSID: BYOD
- QoS policy on 3850 used to police each client bidirectionally
- Policy can be sent via AAA to provide specific per-client policy
- Allocate Bandwidth or police/shape SSID as a whole

Wireless(Cat 3850 & CT 5760)

- Granular QoS control at the wireless edge
  - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- Enhanced Bandwidth Management
  - Approximate Fair Drop (AFD) Bandwidth Management ensures fairness at Client, SSID and Radio levels for NRT traffic
- Wireless Specific Interface Control
  - **Policing capabilities** Per-SSID, Per-Client upstream*** and downstream
  - AAA support for dynamic Client based QoS and Security policies
- Per SSID Bandwidth Management

*** **NOT** available on CT 5760 at FCS

QoS – What's New with Converged Access

With the CT 5760 or CAT 3850
Deterministic bandwidth is allocated per SSID

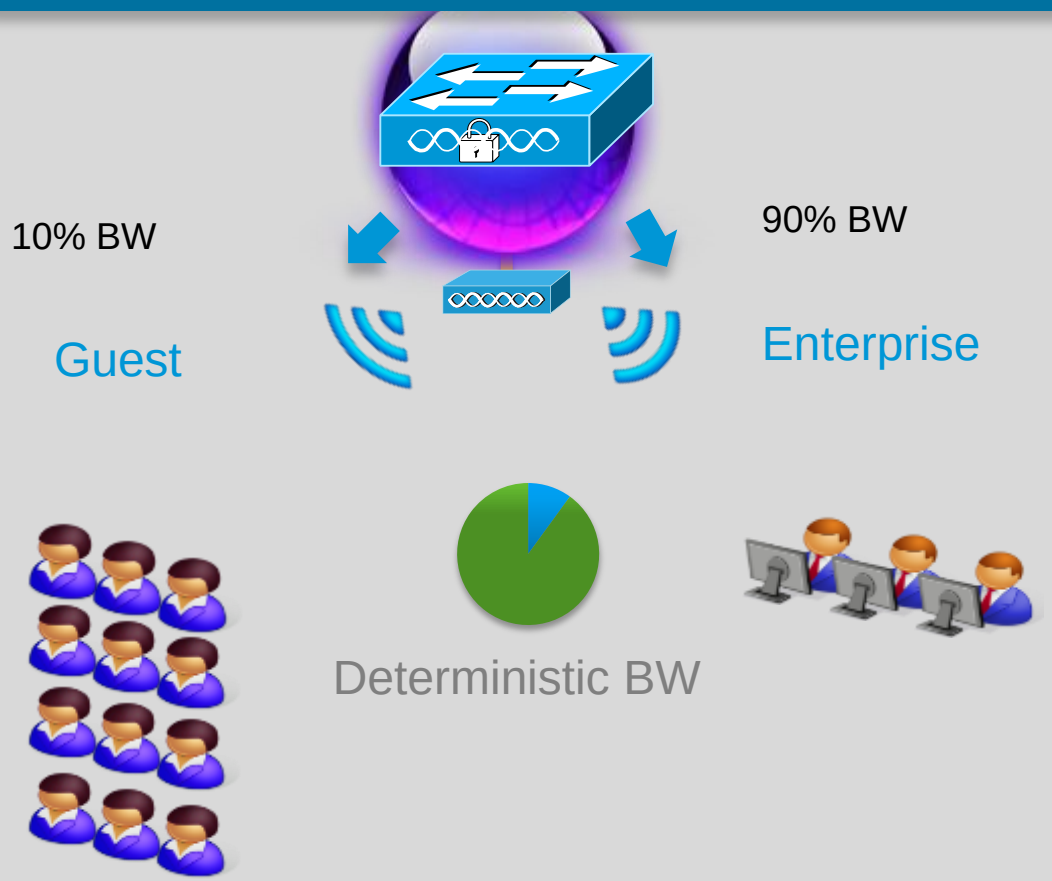

Wireless(Cat 3850 & CT 5760)

- Granular QoS control at the wireless edge
  - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- Enhanced Bandwidth Management
  - Approximate Fair Drop (AFD) Bandwidth Management ensures fairness at Client, SSID and Radio levels for NRT traffic
- Wireless Specific Interface Control
  - Policing capabilities Per-SSID, Per-Client upstream*** and downstream
  - AAA support for dynamic Client based QoS and Security policies
- **Per SSID Bandwidth Management**

*** **NOT** available on CT 5760 at FCS

QoS – What's New with Converged Access

Wired (Cat 3850)

- **Modular QoS based CLI (MQC)**
  - Alignment with 4500E series (Sup6, Sup7)
  - Class-based Queueing, Policing, Shaping, Marking
- More Queues
  - Up to 2P6Q3T queuing capabilities
  - Standard 3750 provides 1P3Q3T
  - Not limited to 2 queue-sets
  - Flexible MQC Provisioning abstracts queuing hardware

```
Policy-map PER-PORT-POLICING
 Class VOIP
   set dscp ef
    police 128000 conform-action transmit exceed-action drop
 Class VIDEO
   set dscp CS4
    police 384000 conform-action transmit exceed-action drop
 Class SIGNALING
   set dscp cs3
    police 32000 conform-action transmit exceed-action drop
 Class TRANSACTIONAL-DATA
   set dscp af21
 Class class-default
   set dscp default
```

Wireless(Cat 3850 & CT 5760)

- Granular QoS control at the wireless edge
  - ers to per-
  - gement
  - ures levels
  - Control
  - er-Client
  - upstream *** and downstream
  - AAA support for dynamic Client based QoS and Security policies
- Per SSID bandwidth allocation

*** **NOT** available on CT 5760 at FCS

QoS Touch Points

Port, Radio, SSID, Client – What Features Apply at Each Level, Downstream

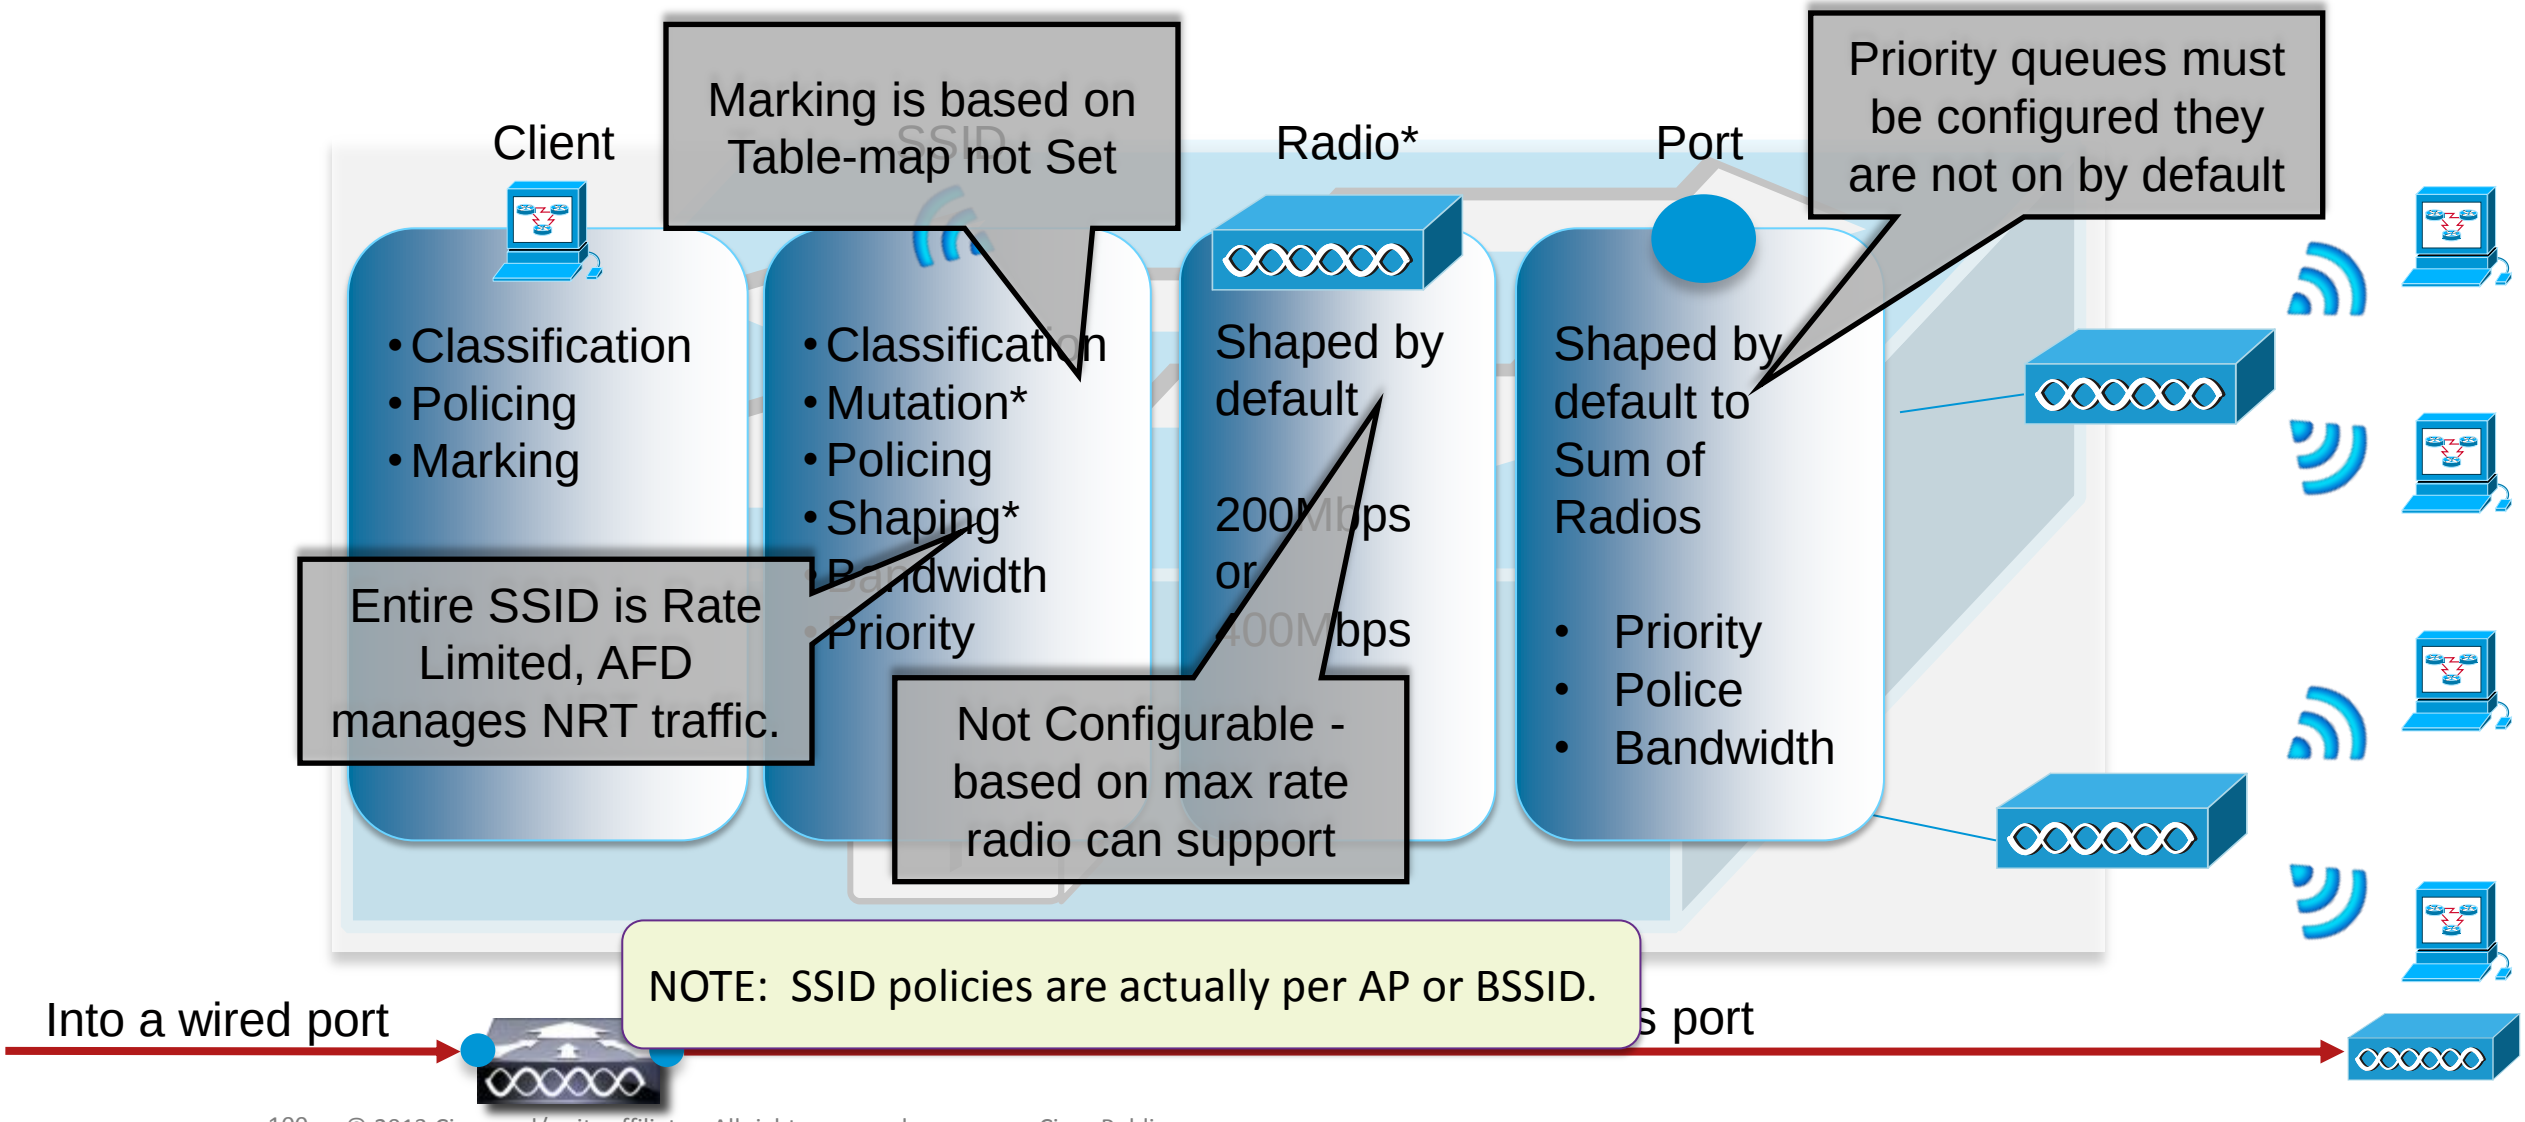

The Catalyst 3850 QoS Toolbox
Wired to Wireless

Conceptual View

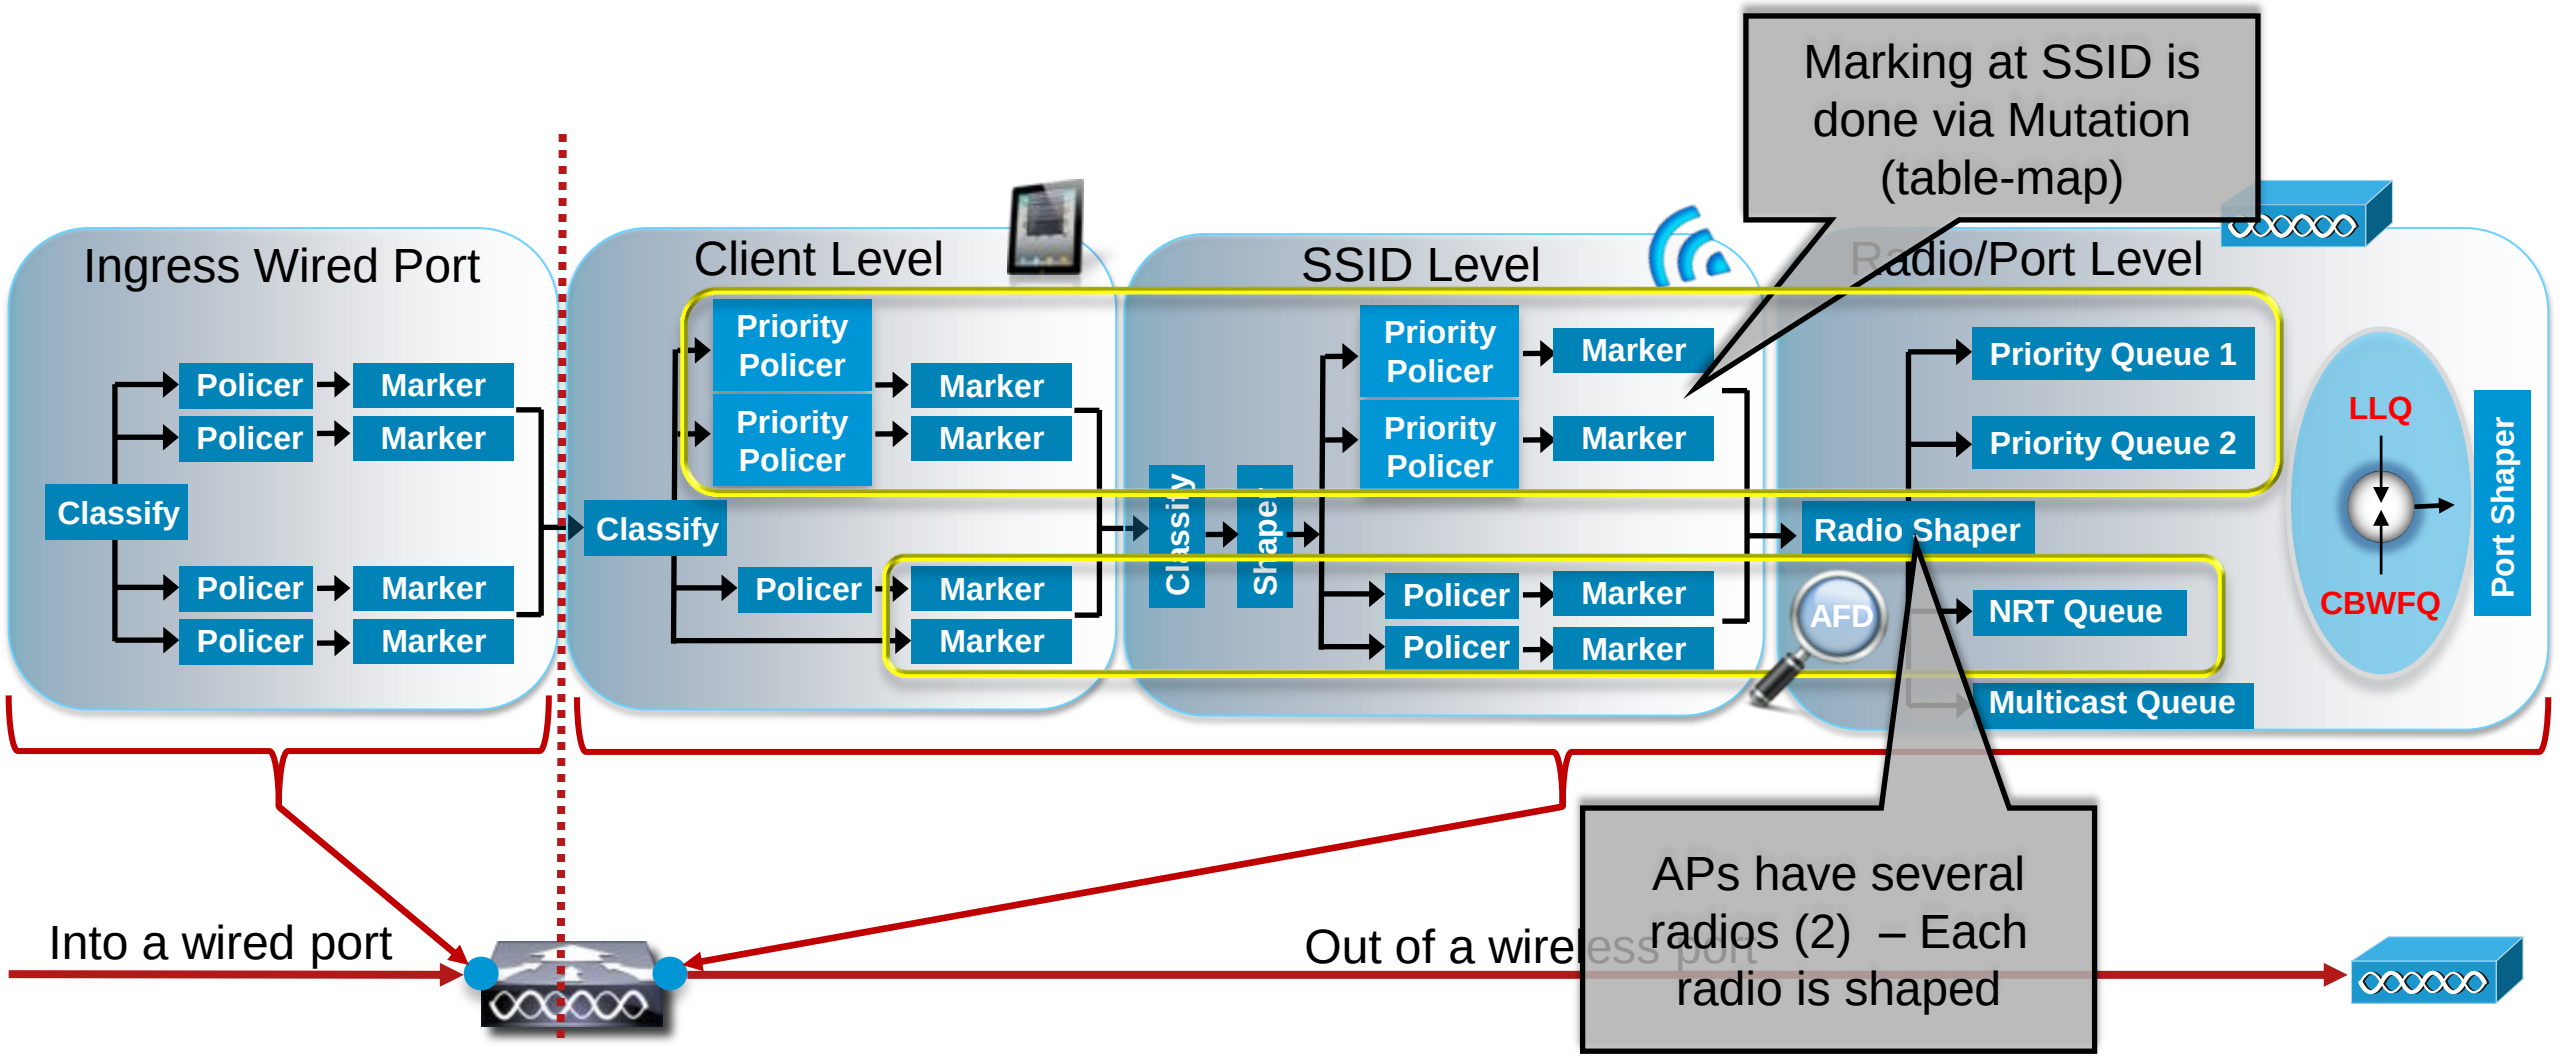

QoS Touch Points

Port, Radio, SSID, Client – What Features Apply at Each Level, Upstream

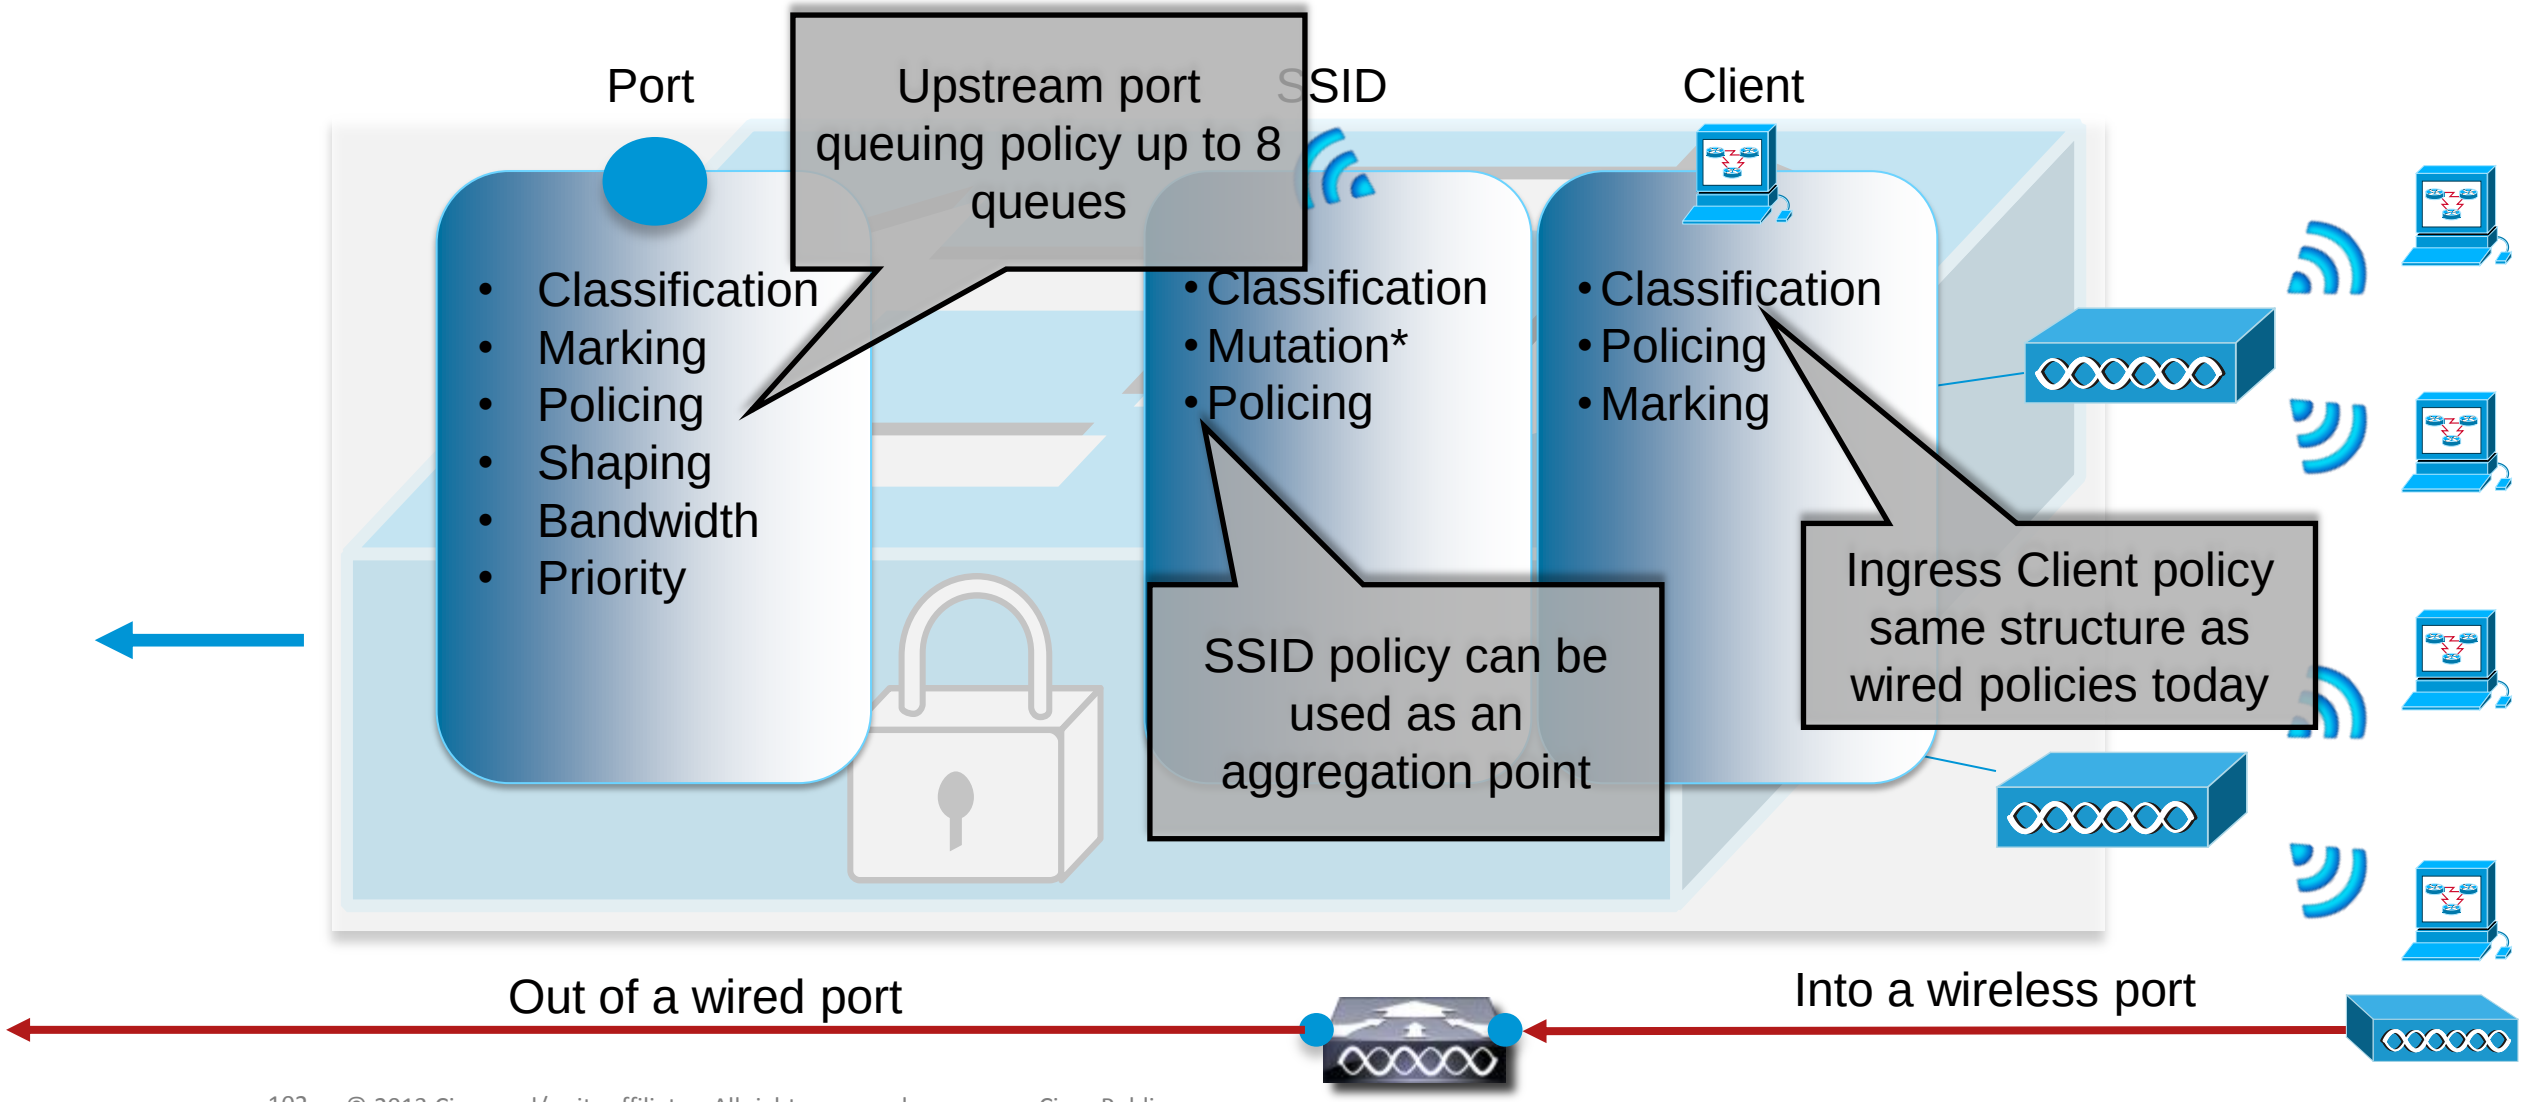

The Catalyst 3850 QoS Toolbox

Wireless to Wired

Conceptual View

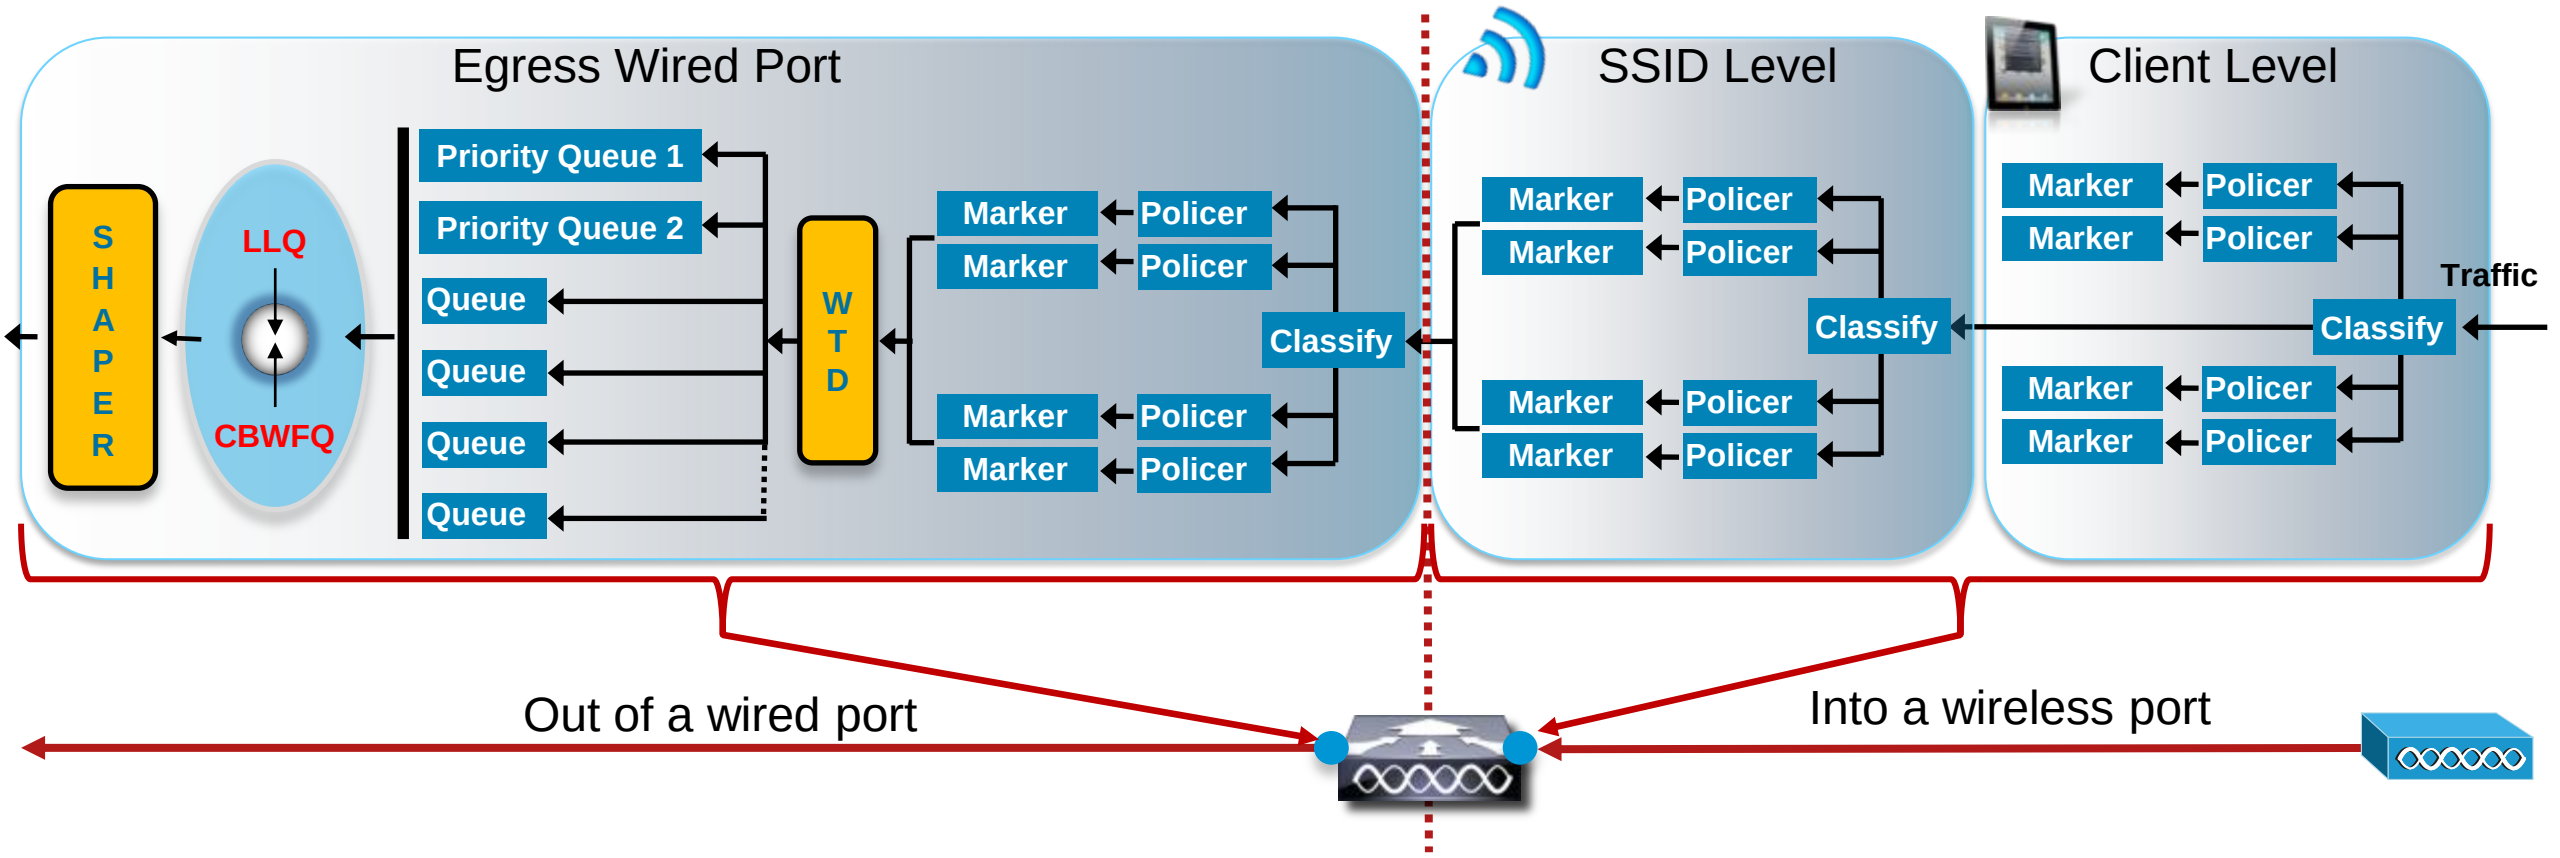

QoS
Default Behaviour

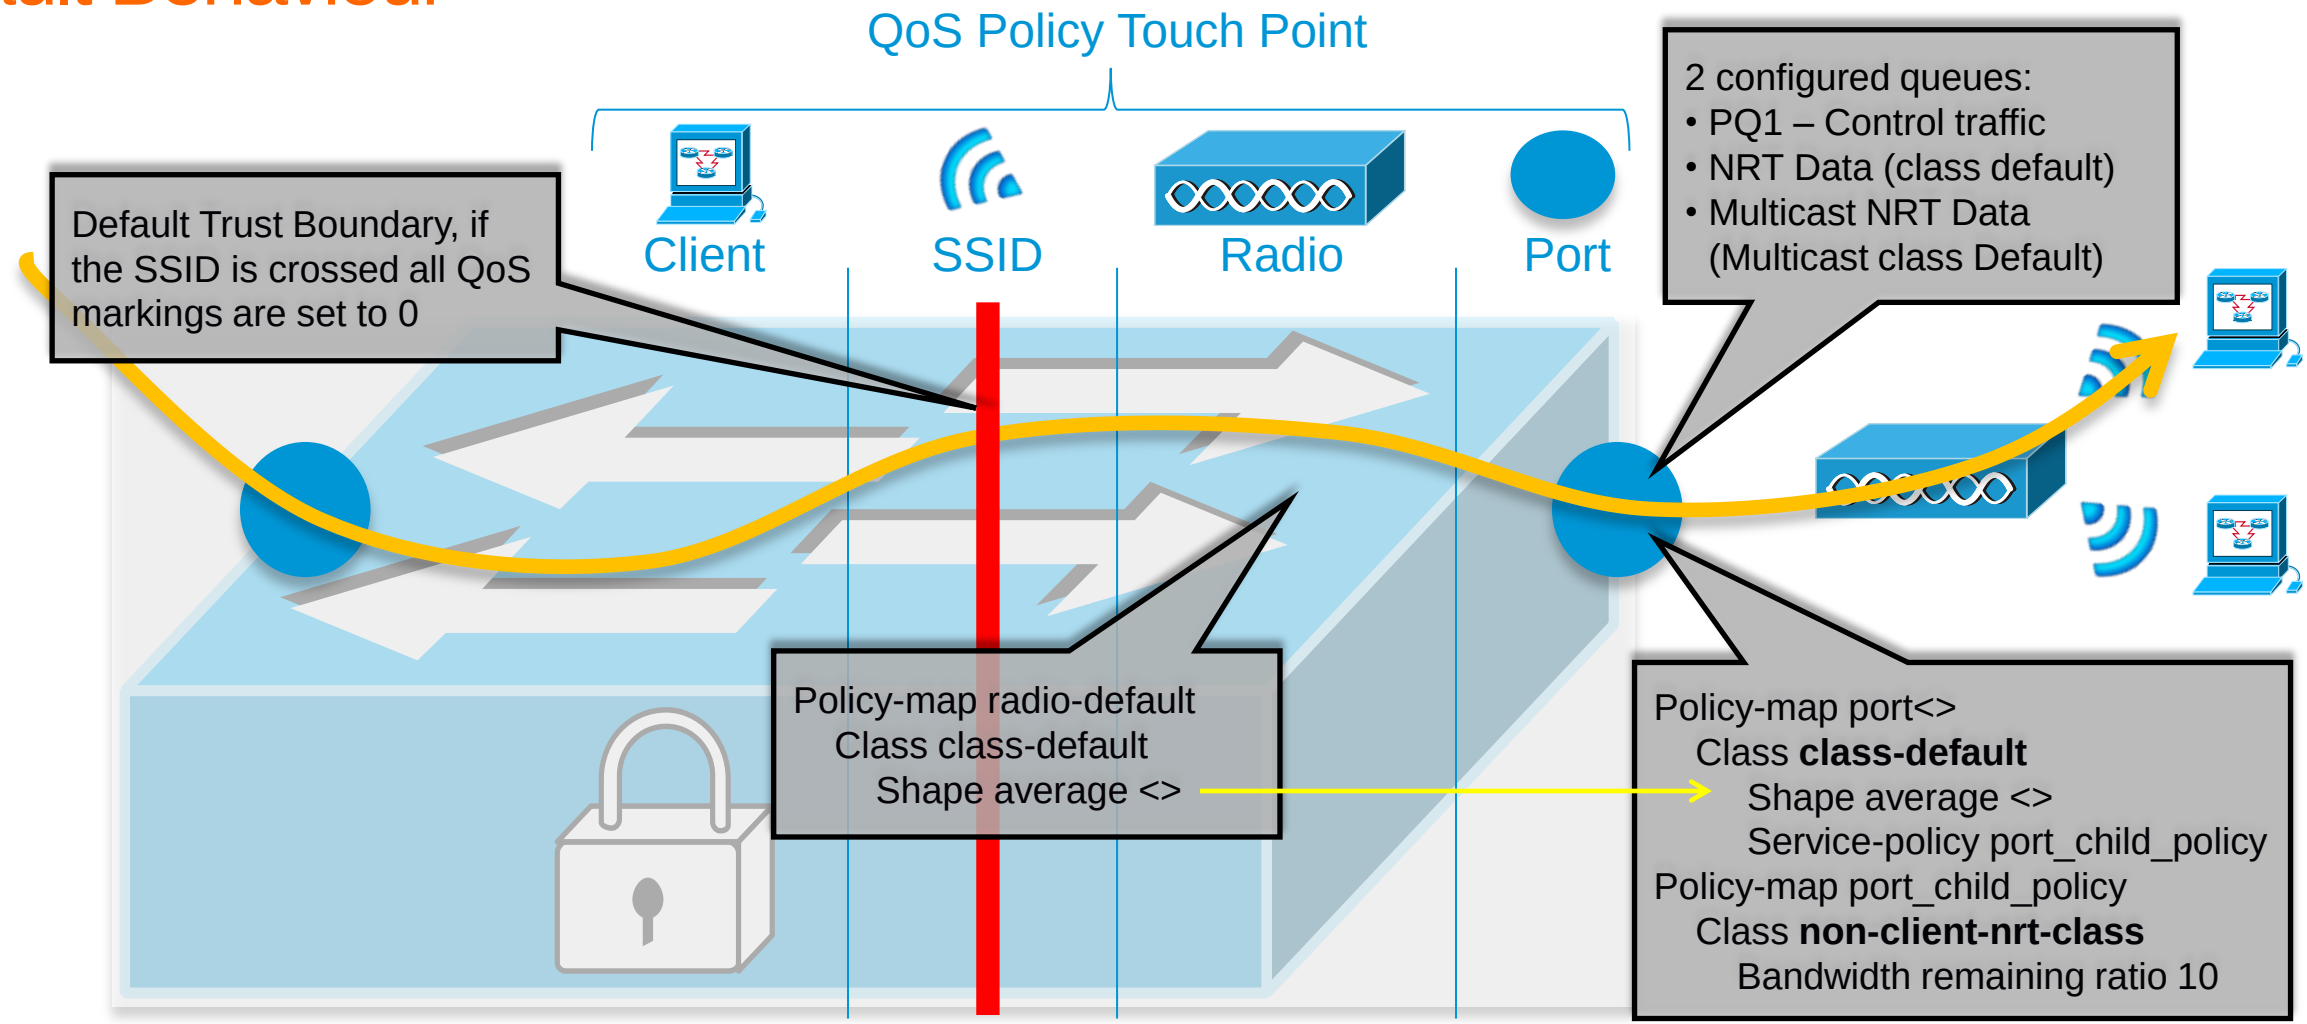

Marking with Table-Maps

And the end of "trust" – Table-Map Example

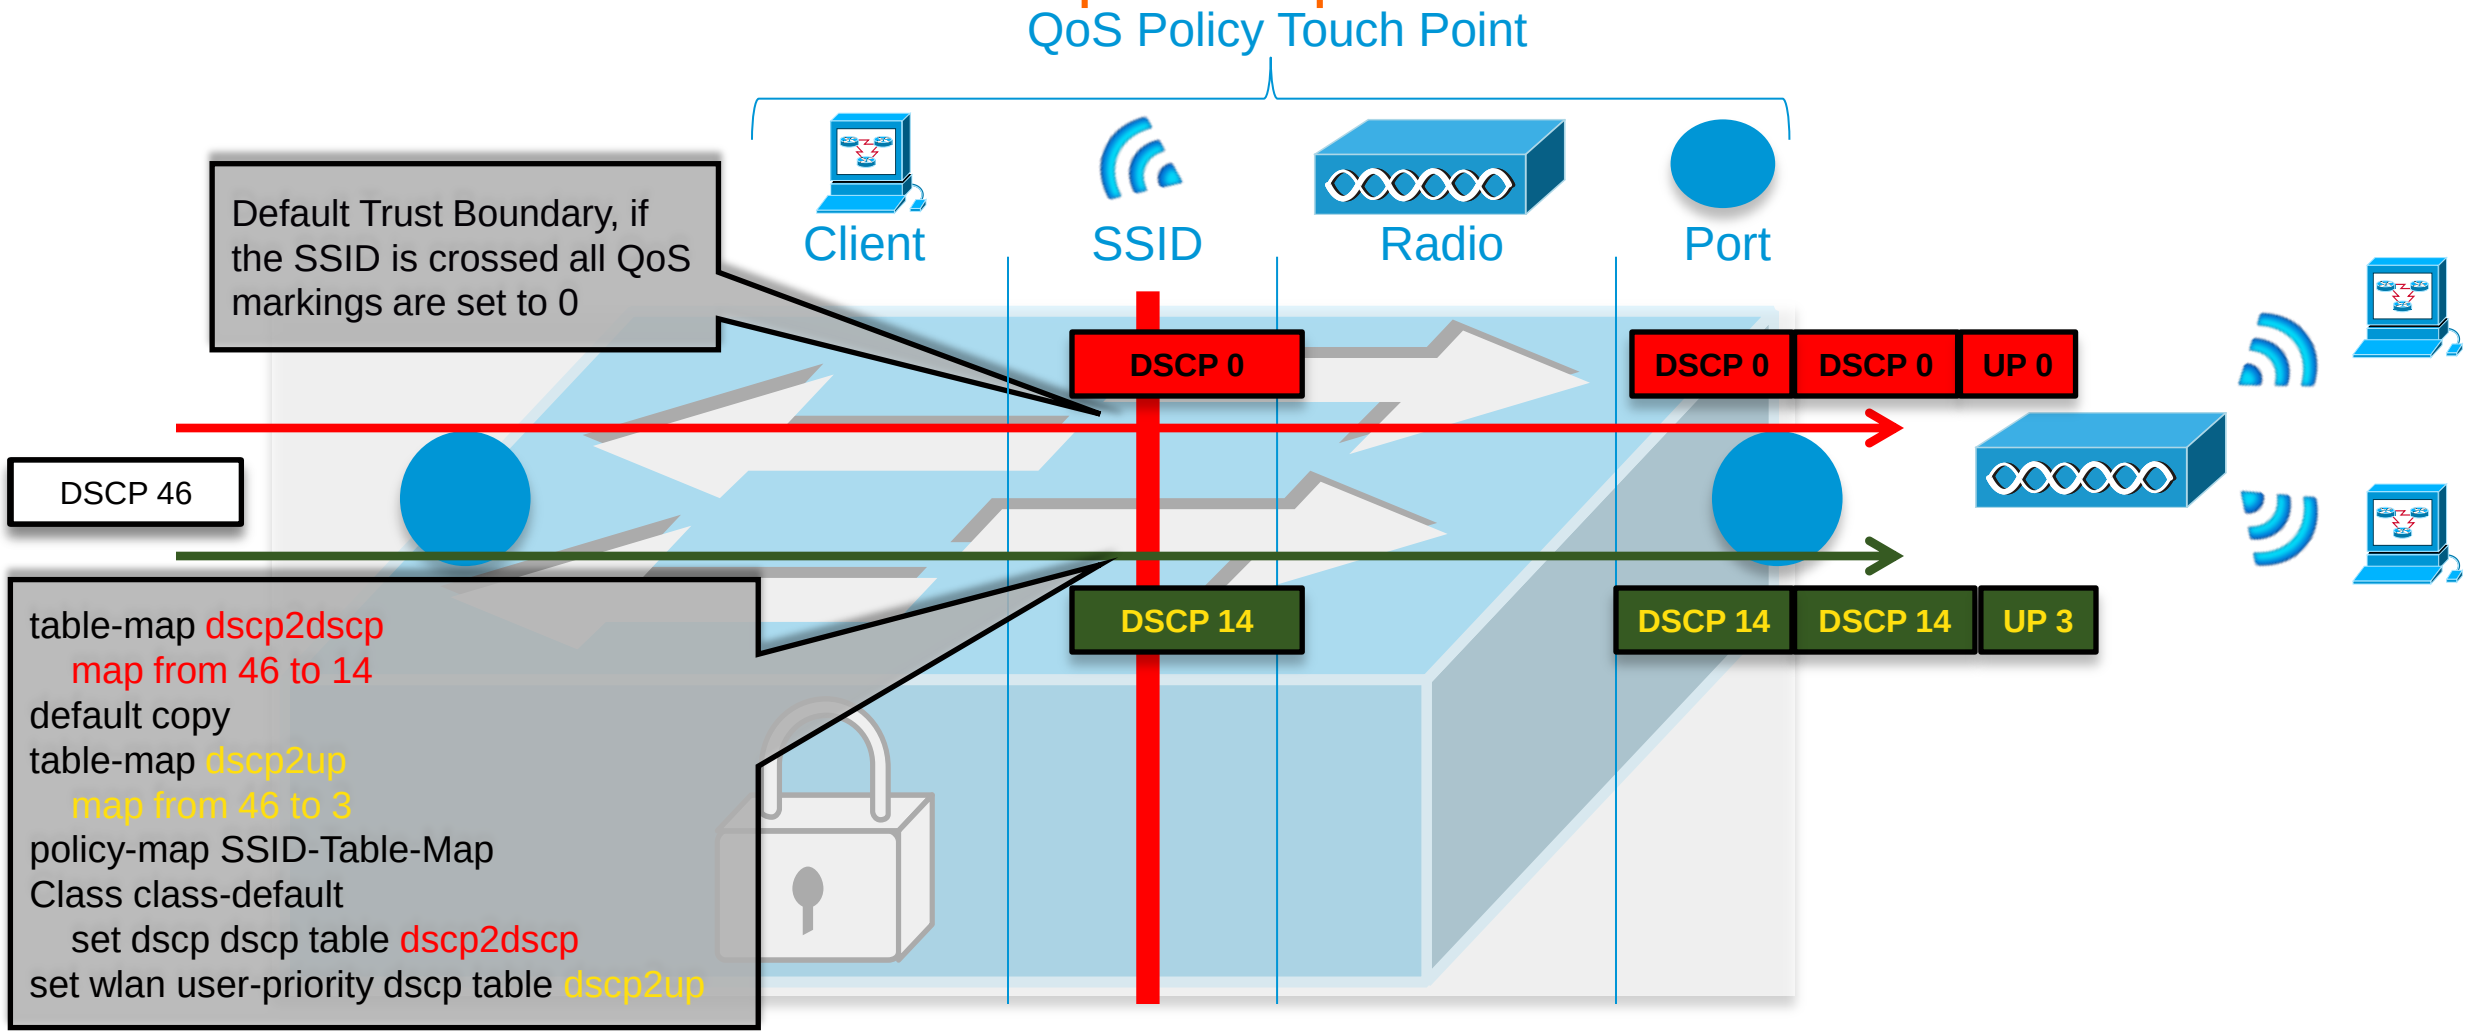

Approximate Fair Drop and Wireless Queuing

Agenda BRKARC-2665 … Converged Access Architecture Overview

Evolution – Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Wired and Wireless – Deployment Options

And a “double-click” deeper …

Existing Wireless Deployment – Architecture Refresher

The Converged Access Deployment in Detail –

- Components of the Deployment – Terminology and Building Blocks
- Converged Access Deployment – Traffic Flows and Roaming
- Converged Access Deployment – High Availability
- Converged Access Deployment – Quality of Service

Summary

Bringing Together Wired and Wireless

How Are We Addressing This Shift?

Controller

Control plane functionality on NG Controller

(also possible on upgraded 5508s, WiSM2s for brownfield deployments, or NG Converged Access switches for small, branch deployments)

Next-Generation WLAN Controller (5760)

Data plane functionality on NG Switches

(also possible on NG Controllers, for deployments in which a centralised approach is preferred)

Next-Generation Switches (Catalyst 3850s)

Enabled by Cisco's strength in Silicon and Systems … **UADP ASIC**

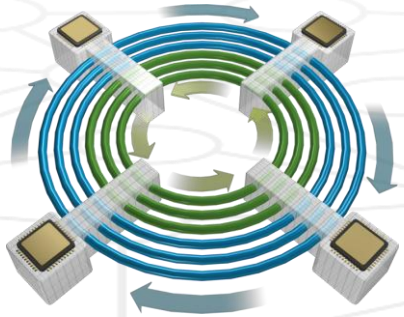

An Evolutionary Advance to Cisco's Wired + Wireless Portfolio, to address device and bandwidth scale, and services demands ….

Bringing Together Wired and Wireless

With a Next-Generation Deployment and Solution

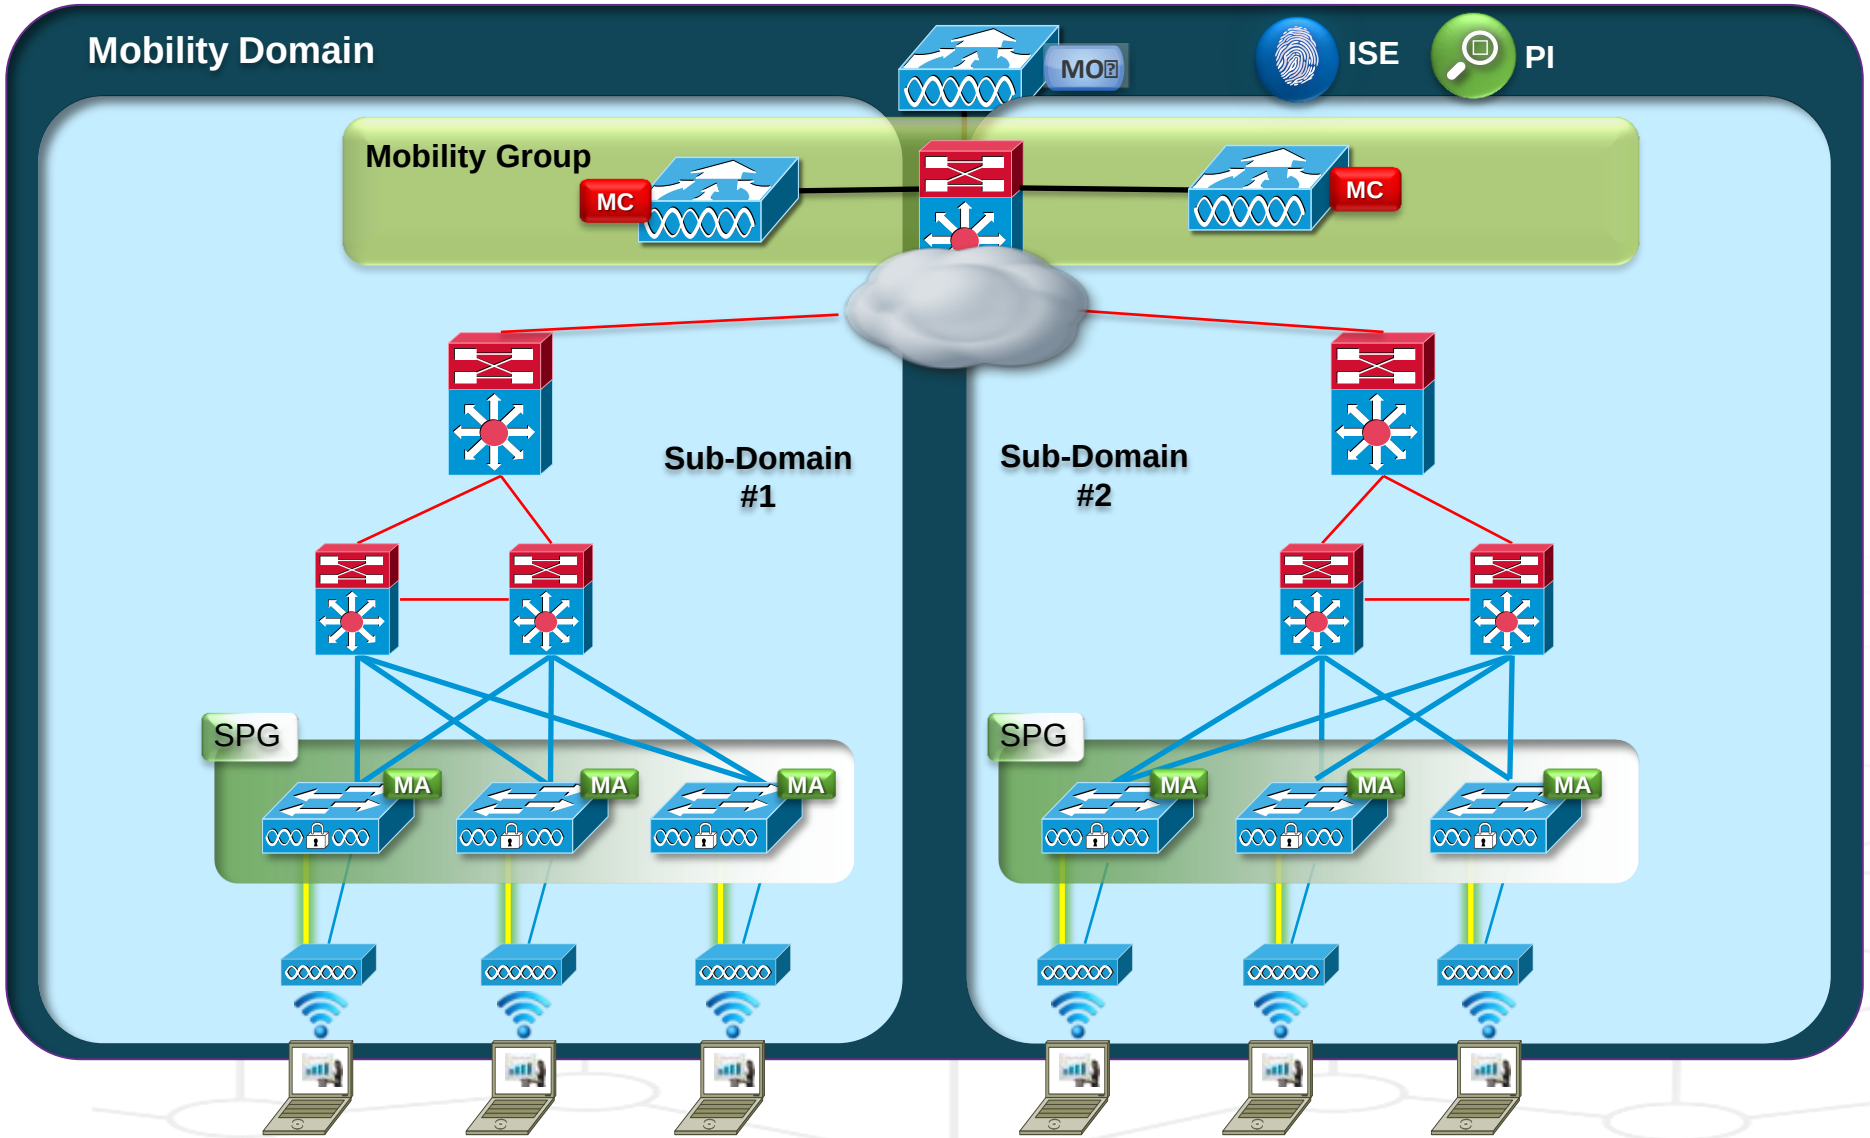

Cisco **Converged Access** Deployment

An Evolutionary Advance to Cisco's Wired + Wireless Portfolio, to address device and bandwidth scale, and services demands ….

Converged Access

Additional Areas of Interest – Reference Material

Additional topics exist, which time precludes us covering here …

However, these are detailed in the Reference Slides which accompany this presentation …

Scalability Details – for both CUWN and Converged Access deployments

Catalyst 3850-based MCs – Examination of Roaming details, and additional design options

Lobby Issue and Solution – Examination of issues with Building entrances / common lobbies, and their impact on client distribution, DHCP scope usage, etc. in Converged Access deployments

Please refer to these slides for additional information on these topics, and feel free to reach out to the presenter with any questions that you may have.

Converged Access

Tell Me How I Did!

Did I Achieve My Objective?

Do You Have a Better Understanding …

of **what Converged Access is** …

of **how Converged Access works** …

and do you now have the basic "Building Blocks" for Converged Access?

Don't Forget

to fill out your evaluations!

Converged Access

For Further Information …

Check out **BRKARC-2666 –**
Converged Access, Campus and Branch Design Guidance!

In the BRKARC-2666 session …

we will **review key elements** from this session …

and **build on this** with information on …

Catalyst 3850 Platform Details,
QoS Implementation, Security,
IP addressing, and Design Options!

Building Out the Complete Solution!

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo 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 www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm

Cisco live! 365

Don't forget to activate your Cisco Live 365 account for access to all session material, communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.
www.ciscoliveaustralia.com/portal/login.ww

Cisco live!

REFERENCE MATERIAL

SCALABILITY

Scalability

CUWN – Using 5508 Controllers

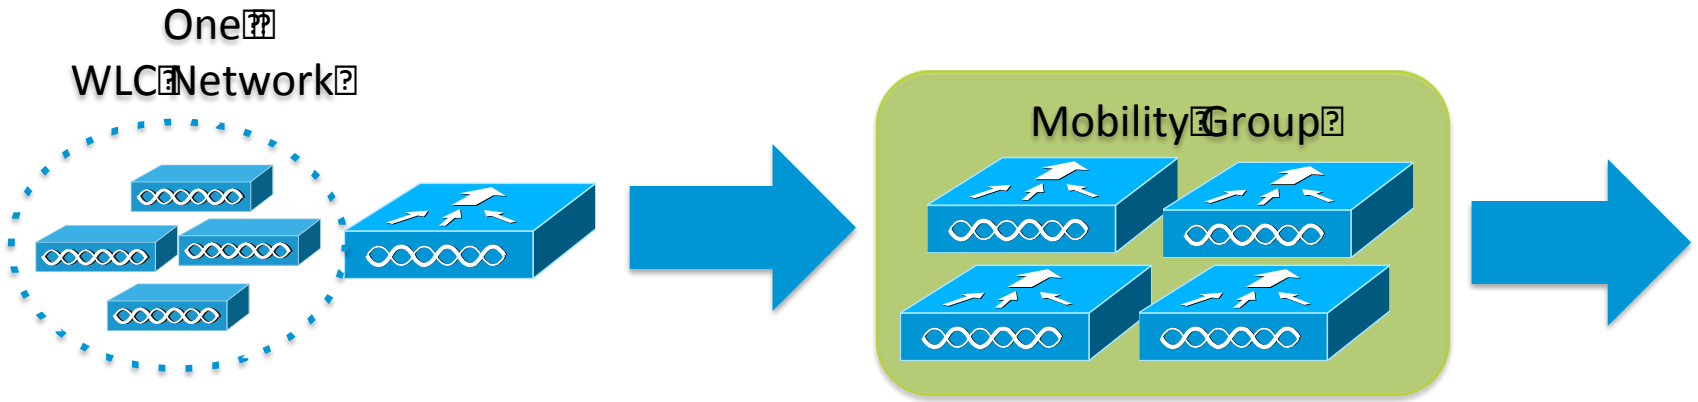

- Up to 500 APs
- Up 7K Clients
- Up to 8 GB 1/O for AP Traffic

- Up to 12K APs
- Up 168K Clients
- Up to 24 WLCs in a MG
- Up to 192 GB 1/O for AP Traffic

- Up to 36K APs
- Up to 504K Clients
- Up to 72 WLCs in a MD
- Up to 576 GB 1/O for AP Traffic

- CT5508 rel 7.3
- Max theoretical scalability numbers
- Without Considering FlexConnect

Scalability

CUWN – Using WiSM2 Controllers

One WLC Network

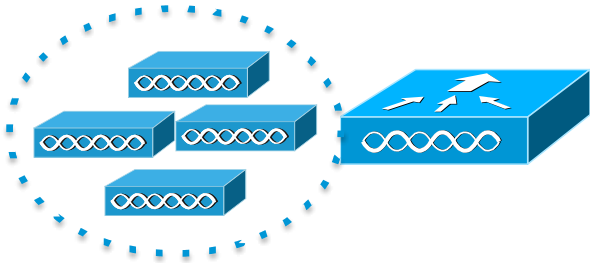

- Up to 1K APs
- Up 15K Clients
- Up to 20 GB I/O for AP Traffic

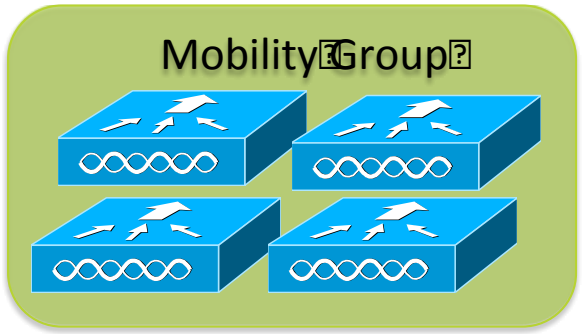

- Up to 24K APs
- Up 360K Clients
- Up to 24 WLCs in a MG
- Up to 480 GB I/O for AP Traffic

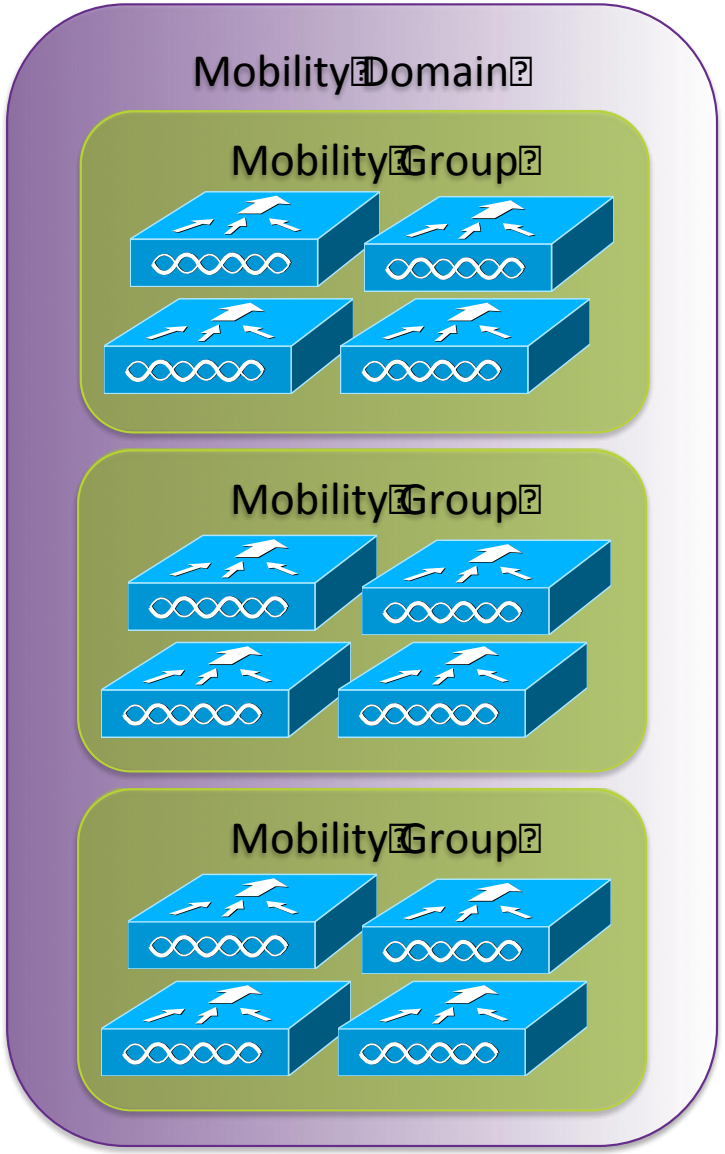

- Up to 72K APs
- Up to 1.08M Clients
- Up to 72 WLCs in a MD
- Up to 1.44TB I/O for AP Traffic

- WiSM-2 rel 7.3
- Max theoretical scalability numbers
- Without Considering FlexConnect

Scalability

Converged Access – Using Catalyst 3850 as MC

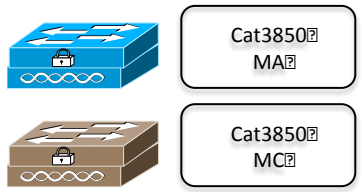

MA=Mobility Agent, MC=Mobility Controller
SPG=Switch Peer Group, SD=Sub-Domain

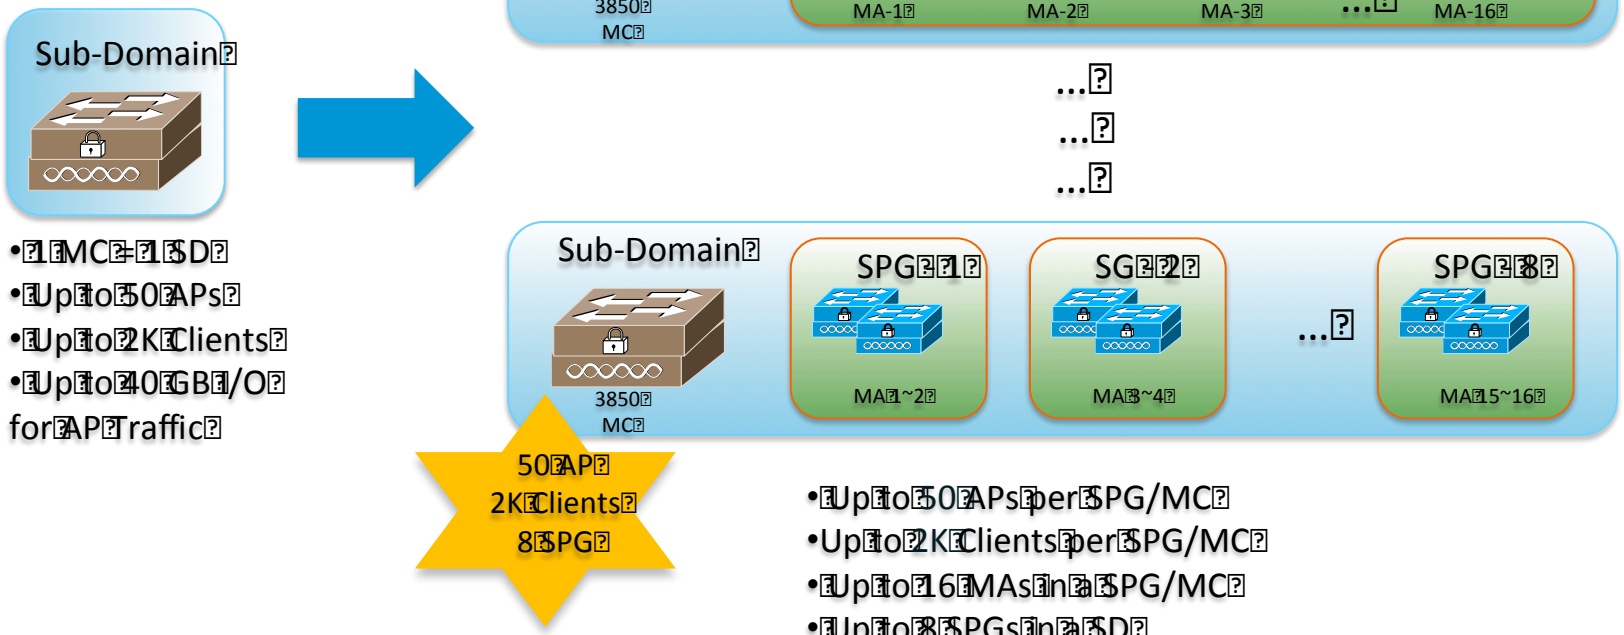

- 1 MC = 1 SD
- Up to 50 APs
- Up to 2K Clients
- Up to 40 GB I/O for AP Traffic

- Up to 50 APs per SPG/MC
- Up to 2K Clients per SPG/MC
- Up to 16 MAs in a SPG/MC
- Up to 8 SPGs in a SD
- Up to 50 GB I/O for AP Traffic

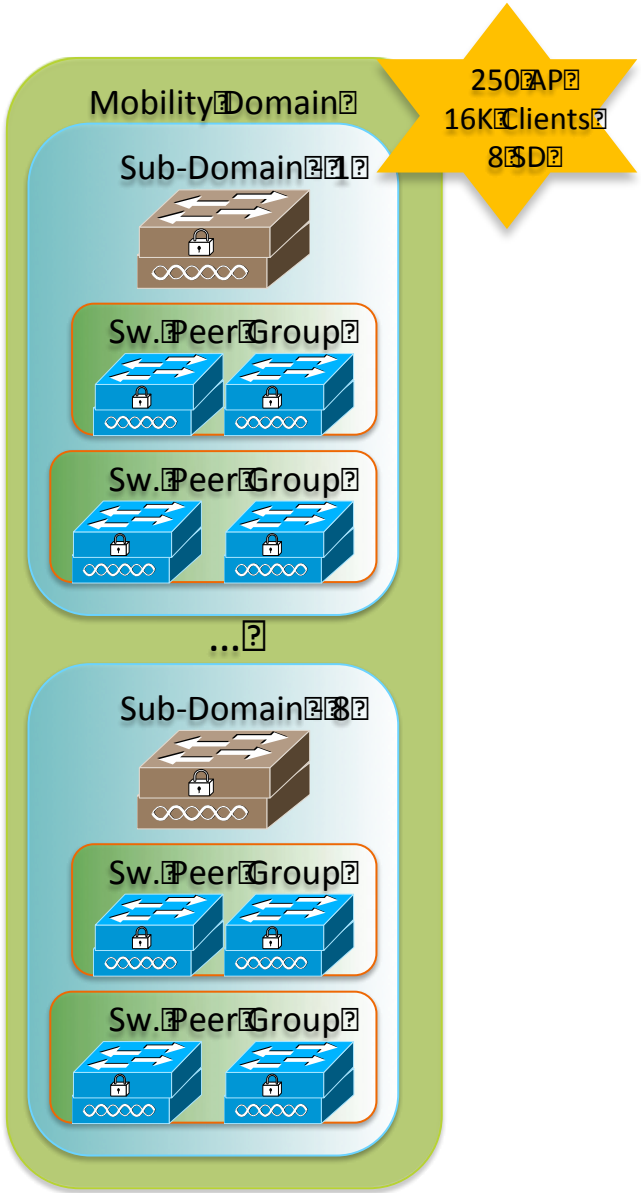

- Up to 250 APs per MD
- Up to 8 SDs per MD
- Up to 128 MAs per MD
- Up to 16K Clients per MD
- Up to 250 GB I/O for AP Traffic

Scalability

Converged Access – 5760 as MC, 3850s as MAs

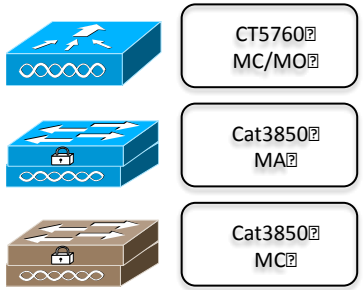

MA=Mobility Agent, MC=Mobility Controller, SPG=Switch Peer Group, SD=Sub-Domain, MG=Mobility Group

MC/MA on one Switch

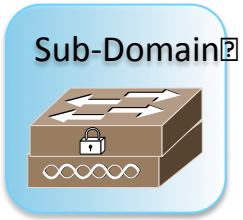

- 1 MC + 1 SD
- Up to 50 APs
- Up to 2K Clients
- Up to 40GB I/O for AP Traffic

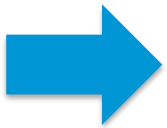

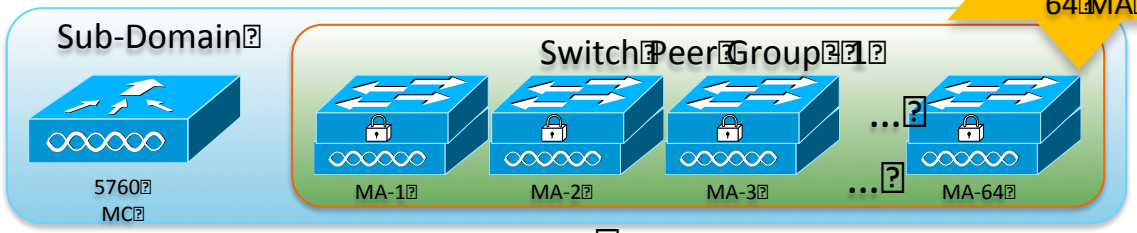

1K AP
12K Clients
64 MA

...

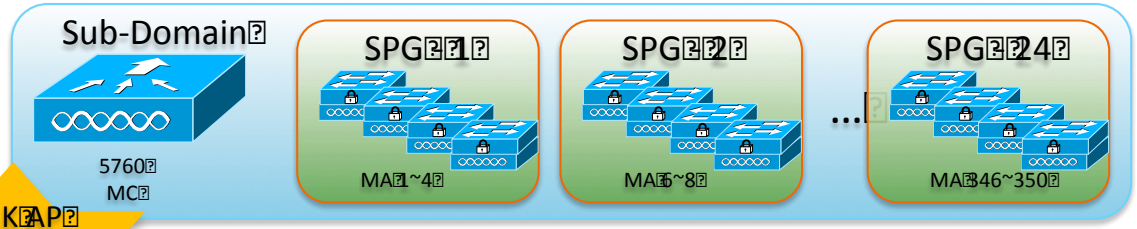

1K AP
12K Clients
24 SPG

- Up to 1K APs per SD/MC
- Up to 12K Clients per SD/MC
- Up to 64 MAs per SPG
- Up to 24 SPGs per SD/MC
- Up to 24 SD/MC per MG
- Up to 350 MAs per SD/MC
- Up to 1TB I/O for AP Traffic

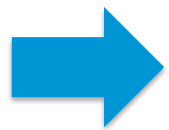

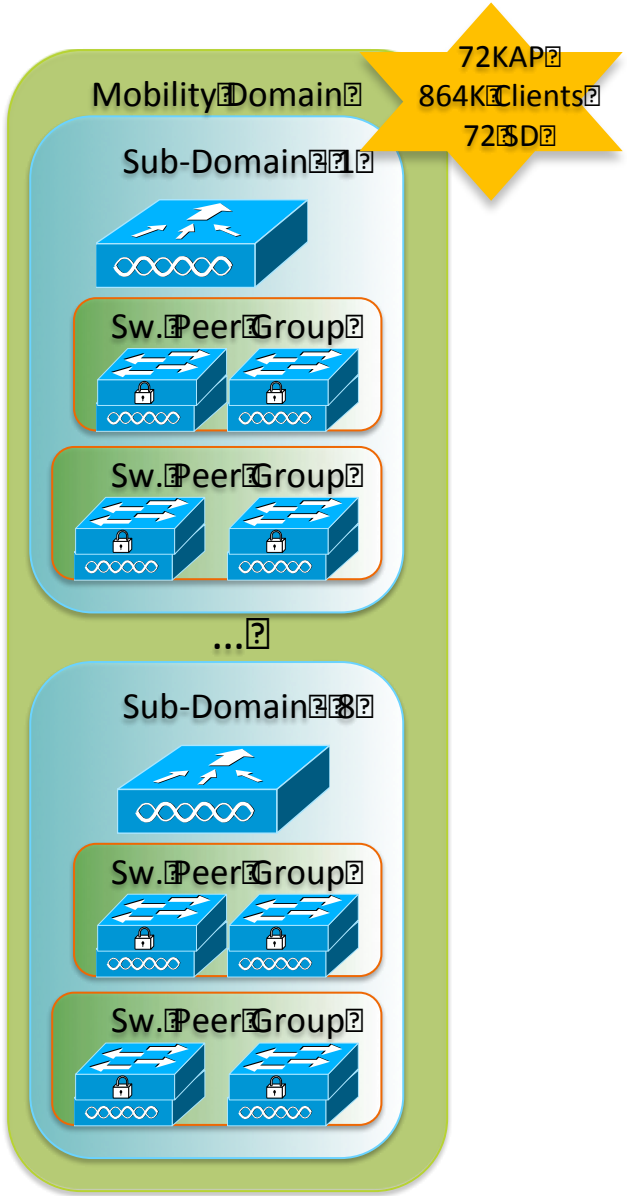

- Up to 72K APs per MD
- Up to 864K Clients per MD
- Up to 72 SD per MD
- Up to 25,200 MAs per MD
- Up to 72TB I/O for AP Traffic

Scalability

Converged Access – 5508 as MC, 3850s as MAs

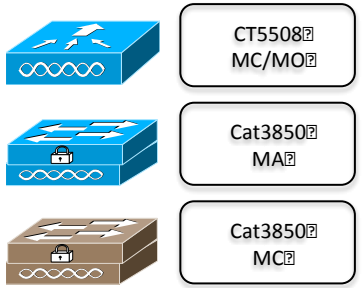

MA=Mobility Agent, MC=Mobility Controller, SPG=Switch Peer Group, SD=Sub-Domain, MG=Mobility Group

MC/MA on one Switch

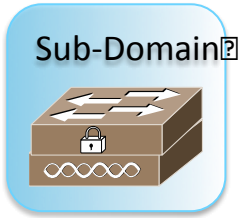

- 1 MC = 1 SD
- Up to 50 APs
- Up to 2K Clients
- Up to 40GB I/O for AP Traffic

500 AP
7K Clients
24 SPG

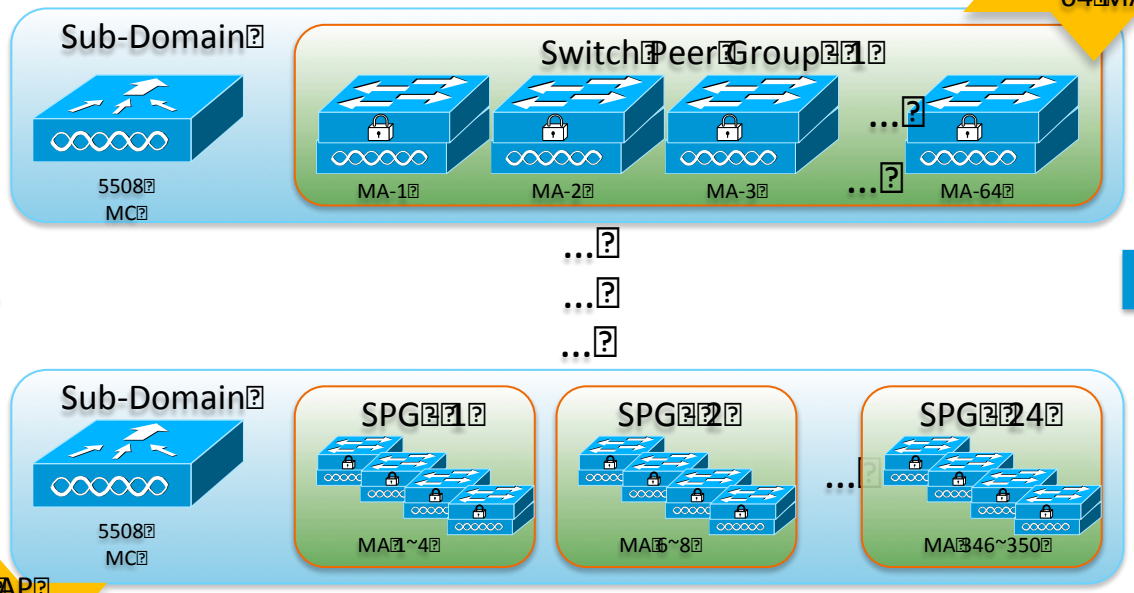

- Up to 500 APs per SD/MC
- Up to 7K Clients per SD/MC
- Up to 64 MAs per SPG
- Up to 24 SPGs per SD/MC
- Up to 24 SD/MC per MG
- Up to 350 MAs per SD/MC
- Up to 500GB I/O for AP Traffic

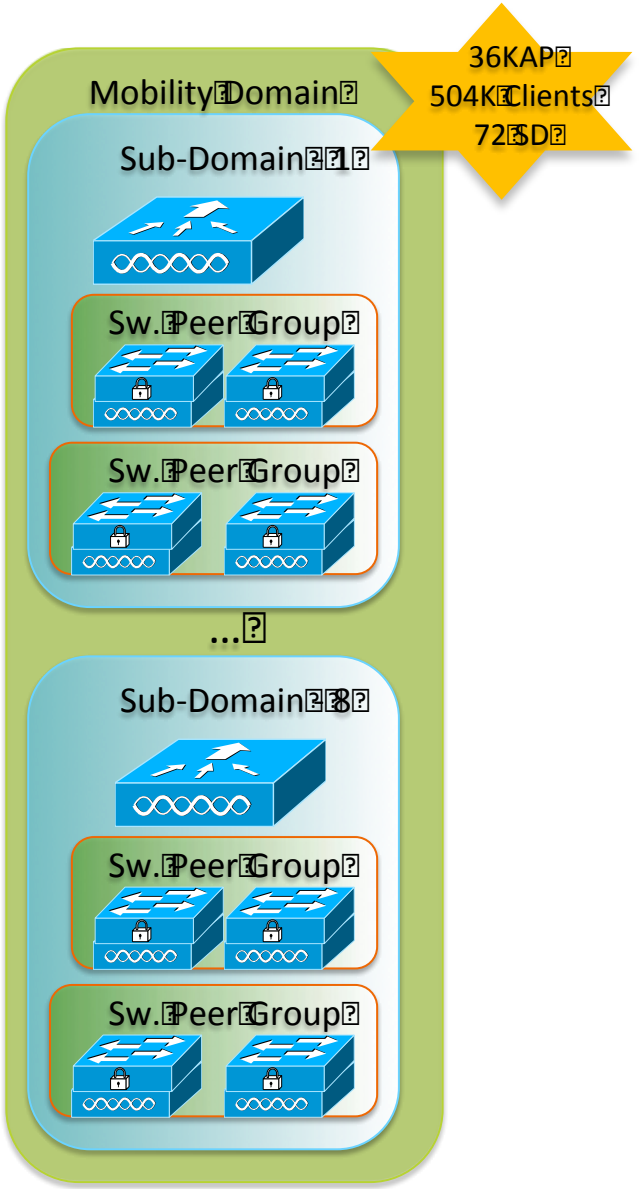

- Up to 36K APs per MD
- Up to 504K Clients per MD
- Up to 72 SD per MD
- Up to 25,200 MAs per MD
- Up to 36TB I/O for AP Traffic

Scalability

Converged Access – WiSM2 as MC, 3850s as MAs

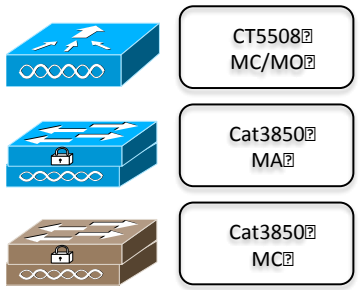

MA=Mobility Agent, MC=Mobility Controller, SPG=Switch Peer Group, SD=Sub-Domain, MG=Mobility Group

MC/MA on one Switch

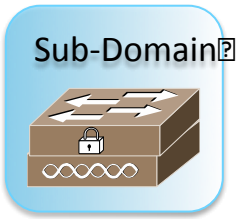

- 1 MC = 1 SD
- Up to 50 APs
- Up to 2K Clients
- Up to 40GB I/O for AP Traffic

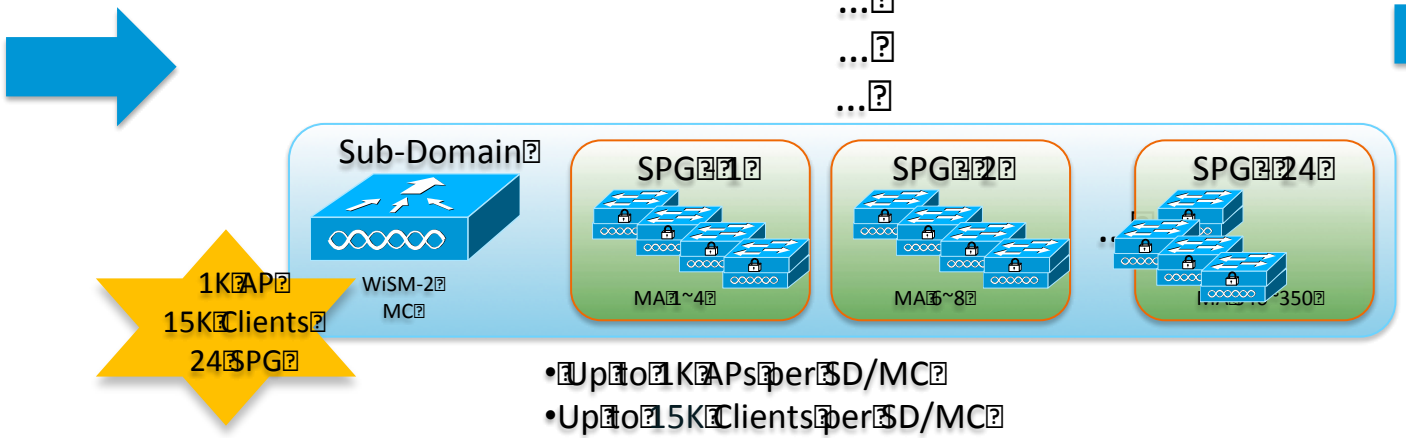

- Up to 1K APs per SD/MC
- Up to 15K Clients per SD/MC
- Up to 16 MAs per SPG
- Up to 24 SPGs per SD/MC
- Up to 24 SD/MC per MG
- Up to 350 MAs per SD/MC
- Up to 1TB I/O for AP Traffic

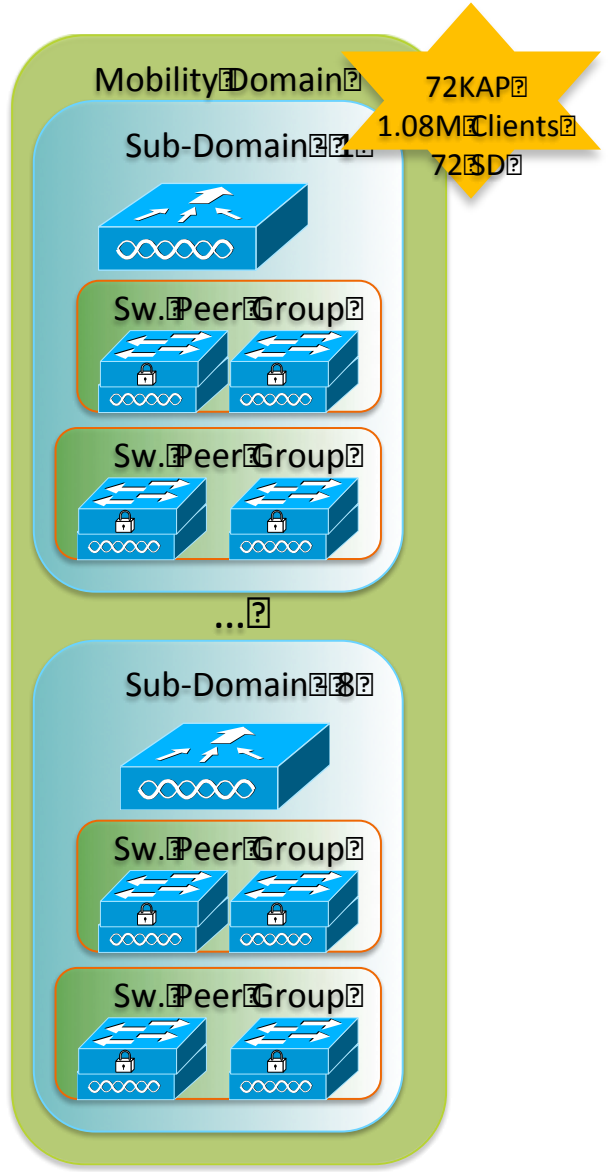

- Up to 72K APs per MD
- Up to 1.08M Clients per MD
- Up to 72 SD per MD
- Up to 25,200 MAs per MD
- Up to 72TB I/O for AP Traffic

REFERENCE MATERIAL

CATALYST 3850-BASED MCS – ROAMING DETAILS

Converged Access

Catalyst 3850-based MCs – Roaming, within a Stack

Roaming, within a Stack (3850 Switches as MCs) –

- **Initially, all clients in this example are on their initial, local Converged Access switches**
- **Now, a client roams –** and we see his resulting traffic topology
- **Roaming within a stack does not change the user's PoP or PoA –** since the stack implements a single MA (redundant within the stack), and thus a user that roams to another AP serviced by the same stack does not cause a PoA move

No change to user's PoP or PoA

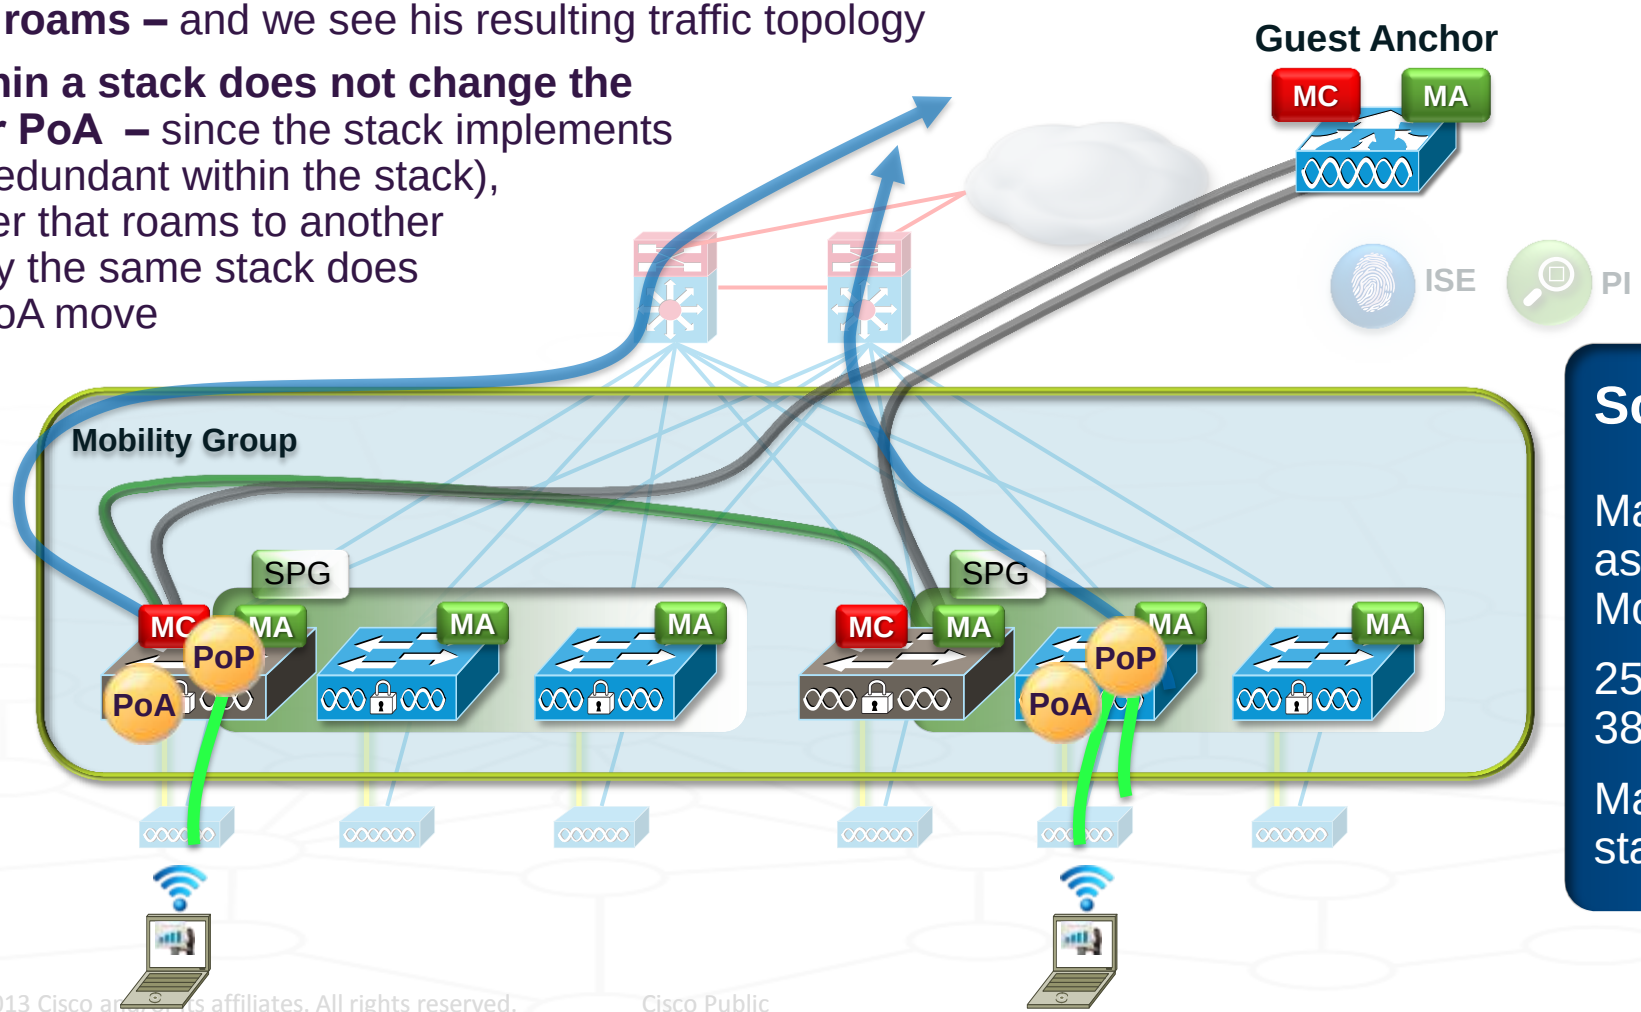

Scalability –

Max of 8 x 3850 switches as MCs, grouped into a Mobility Group

250 APs total across all 3850 MCs

Max. 50 APs per 3850 stack / SPG

Converged Access

Catalyst 3850-based MCs – Roaming, within an SPG

Roaming, within a Switch Peer Group (3850 Switches as MCs) –

- **Now, the client roams to an AP serviced by another switch stack (within the same SPG)**
- **Let's examine** his resulting traffic topology
- **The user has moved between MAs (switch stacks) –** to maintain consistency of user connectivity (IP address) and policy application, the user's traffic is transported to the MA that the user associated with initially (i.e. the user's PoA moved, but their PoP stayed static)

Most Common Roaming Case

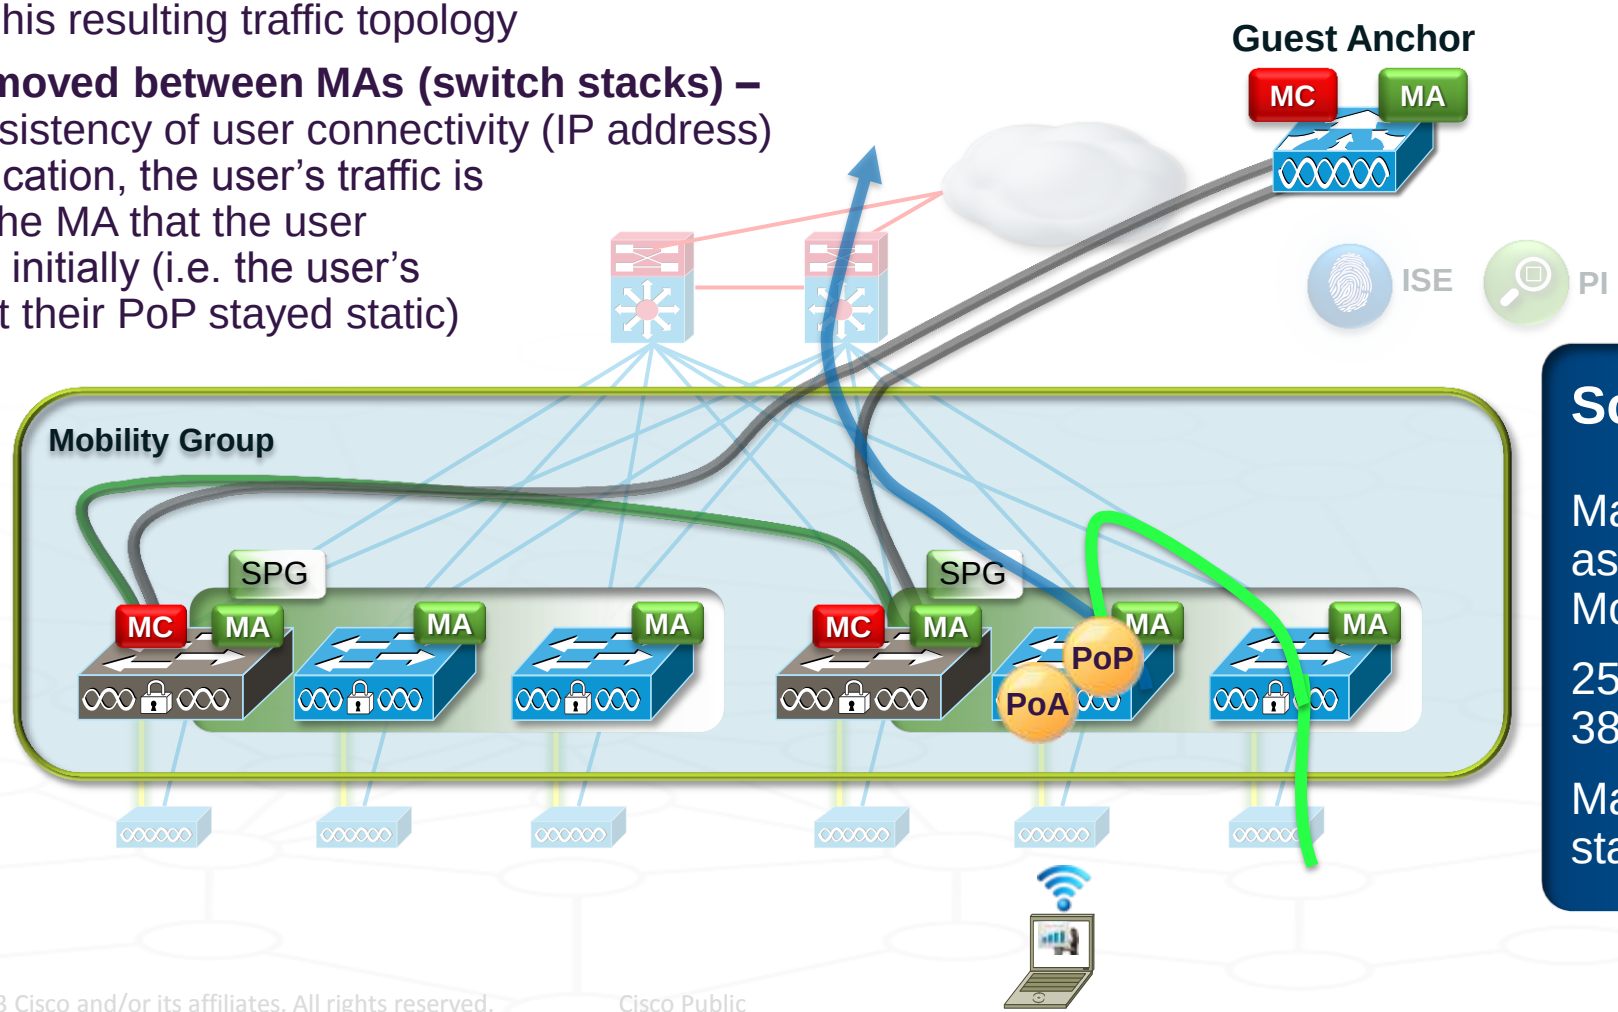

Scalability –

Max of 8 x 3850 switches as MCs, grouped into a Mobility Group

250 APs total across all 3850 MCs

Max. 50 APs per 3850 stack / SPG

Converged Access

Catalyst 3850-based MCs – Roaming, across SPGs

Roaming, across Switch Peer Groups (3850 Switches as MCs) –

- **Now, let's examine a more complex roam where the user roams across SPGs**
- **In this example –** the user roams to a separate SPG, onto the stack serving as MC for that SPG
- **The user's has moved between SPGs –** so their traffic needs to be transported back to their PoP, which has remained static – and it does so by transiting between the two MCs servicing these two Switch Peer Groups (MCs are fully meshed within the MG)

Roaming between SPGs (geographically-separated)

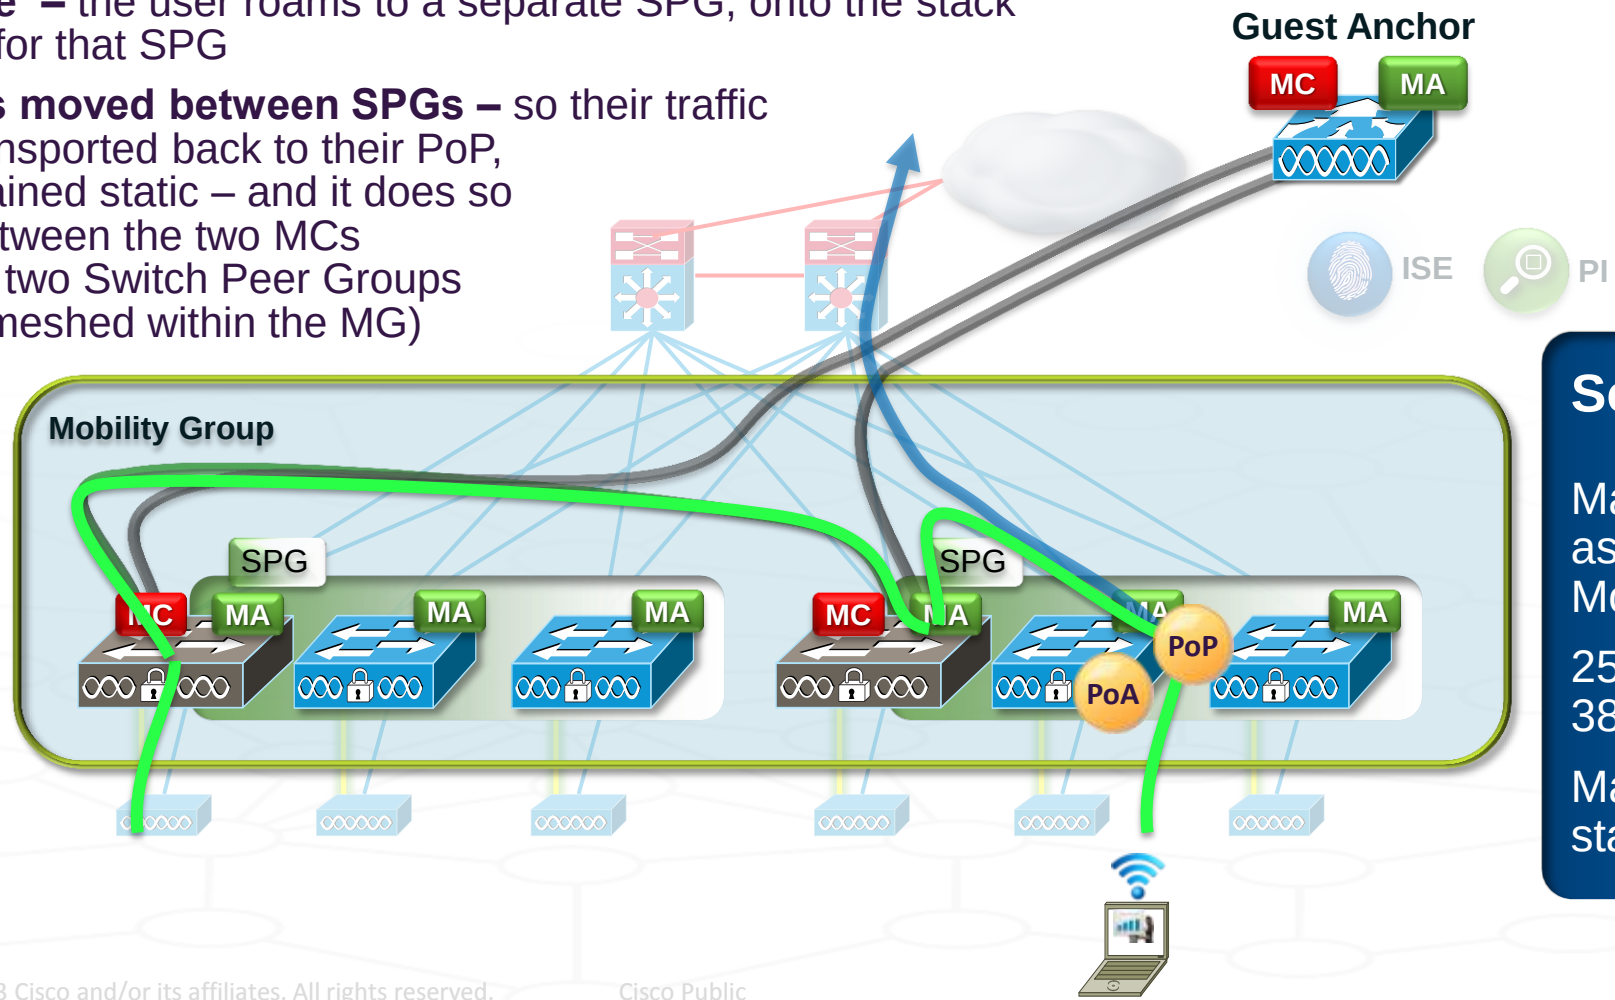

Scalability –

Max of 8 x 3850 switches as MCs, grouped into a Mobility Group

250 APs total across all 3850 MCs

Max. 50 APs per 3850 stack / SPG

Converged Access

Catalyst 3850-based MCs – Roaming, across SPGs & MCs

Roaming, across Switch Peer Groups and MCs (3850 Switches as MCs) –

- **Now, lets' examine the most complex type of roam – across SPGs and MCs / MAs**
- **Remember –** these types of roams are likely to be a minority case in most deployments
- **The user has moved between MAs, MCs, and SPGs –** and their traffic takes the path shown since, again, their PoP has remained static, while the PoA moved as the user roamed (maintains user IP address, maintains consistency of policy application)

Roaming between SPGs and MCs (geographically-separated)

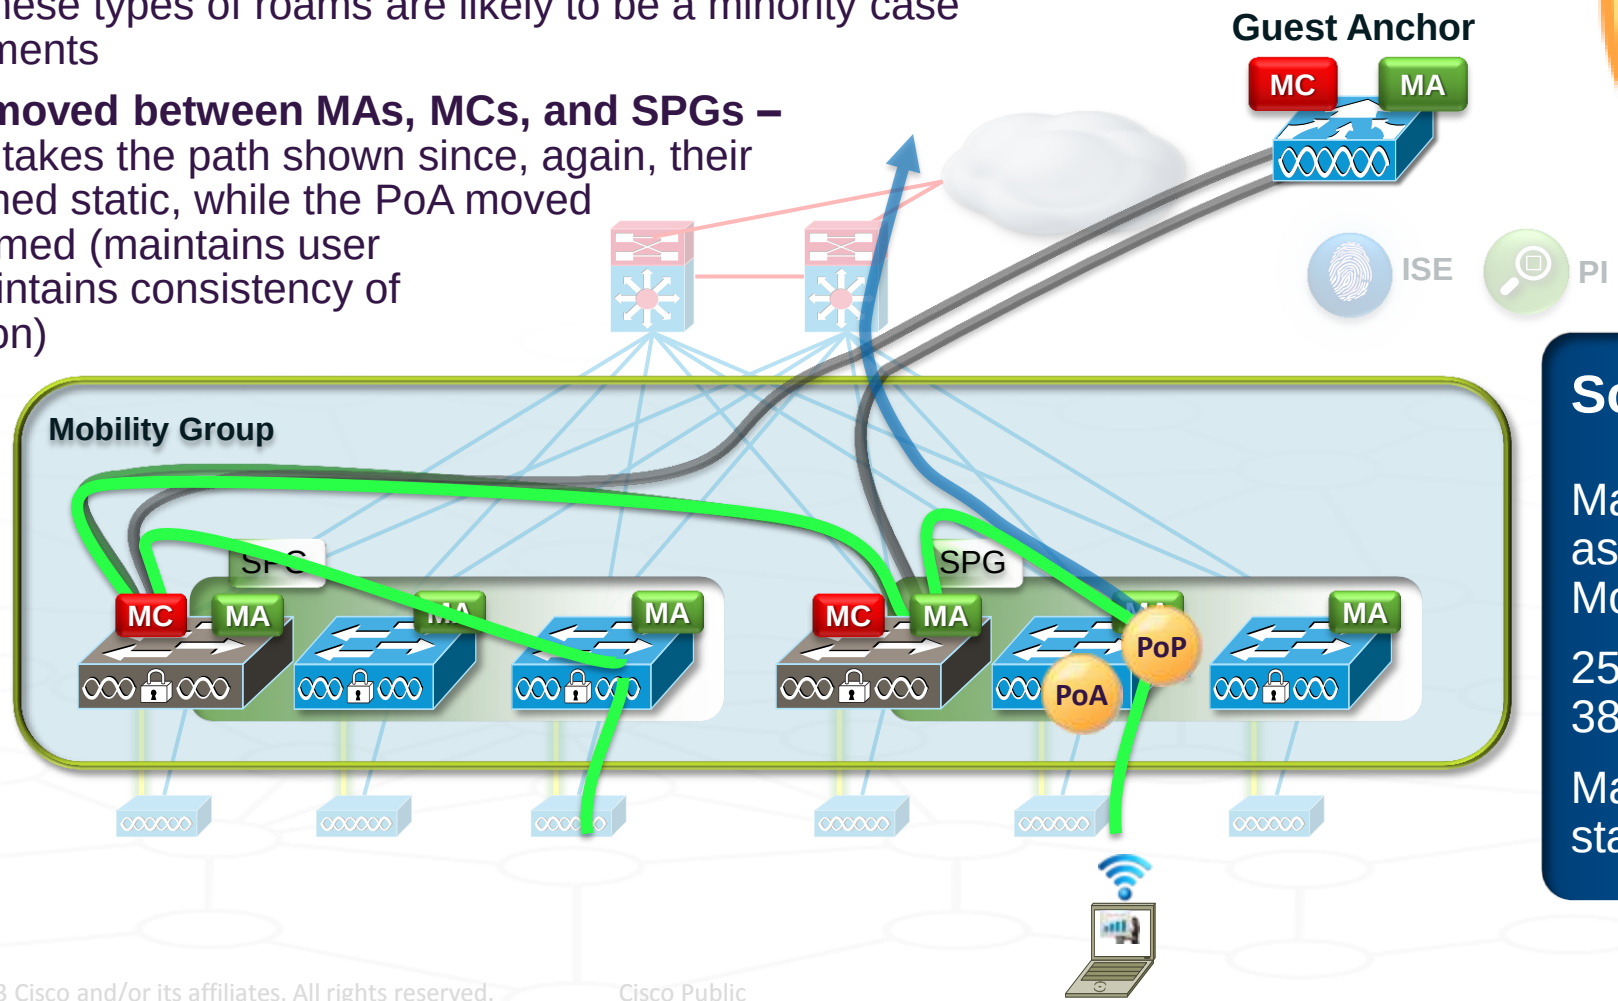

Scalability –

Max of 8 x 3850 switches as MCs, grouped into a Mobility Group

250 APs total across all 3850 MCs

Max. 50 APs per 3850 stack / SPG

REFERENCE MATERIAL

LOBBY ISSUE / SOLUTION

Converged Access

Common Building Access – The "Lobby Issue"

What happens when –

- **Everyone enters the building** via a common lobby
- **APs in that lobby are controlled by one Converged Access switch stack**
- All the users, and their traffic –
  - **Gets "pinned" to that switch** ... causing issues for traffic load, switch load, DHCP pool exhaustion, etc. …

Many users could end up "staying in the lobby" logically

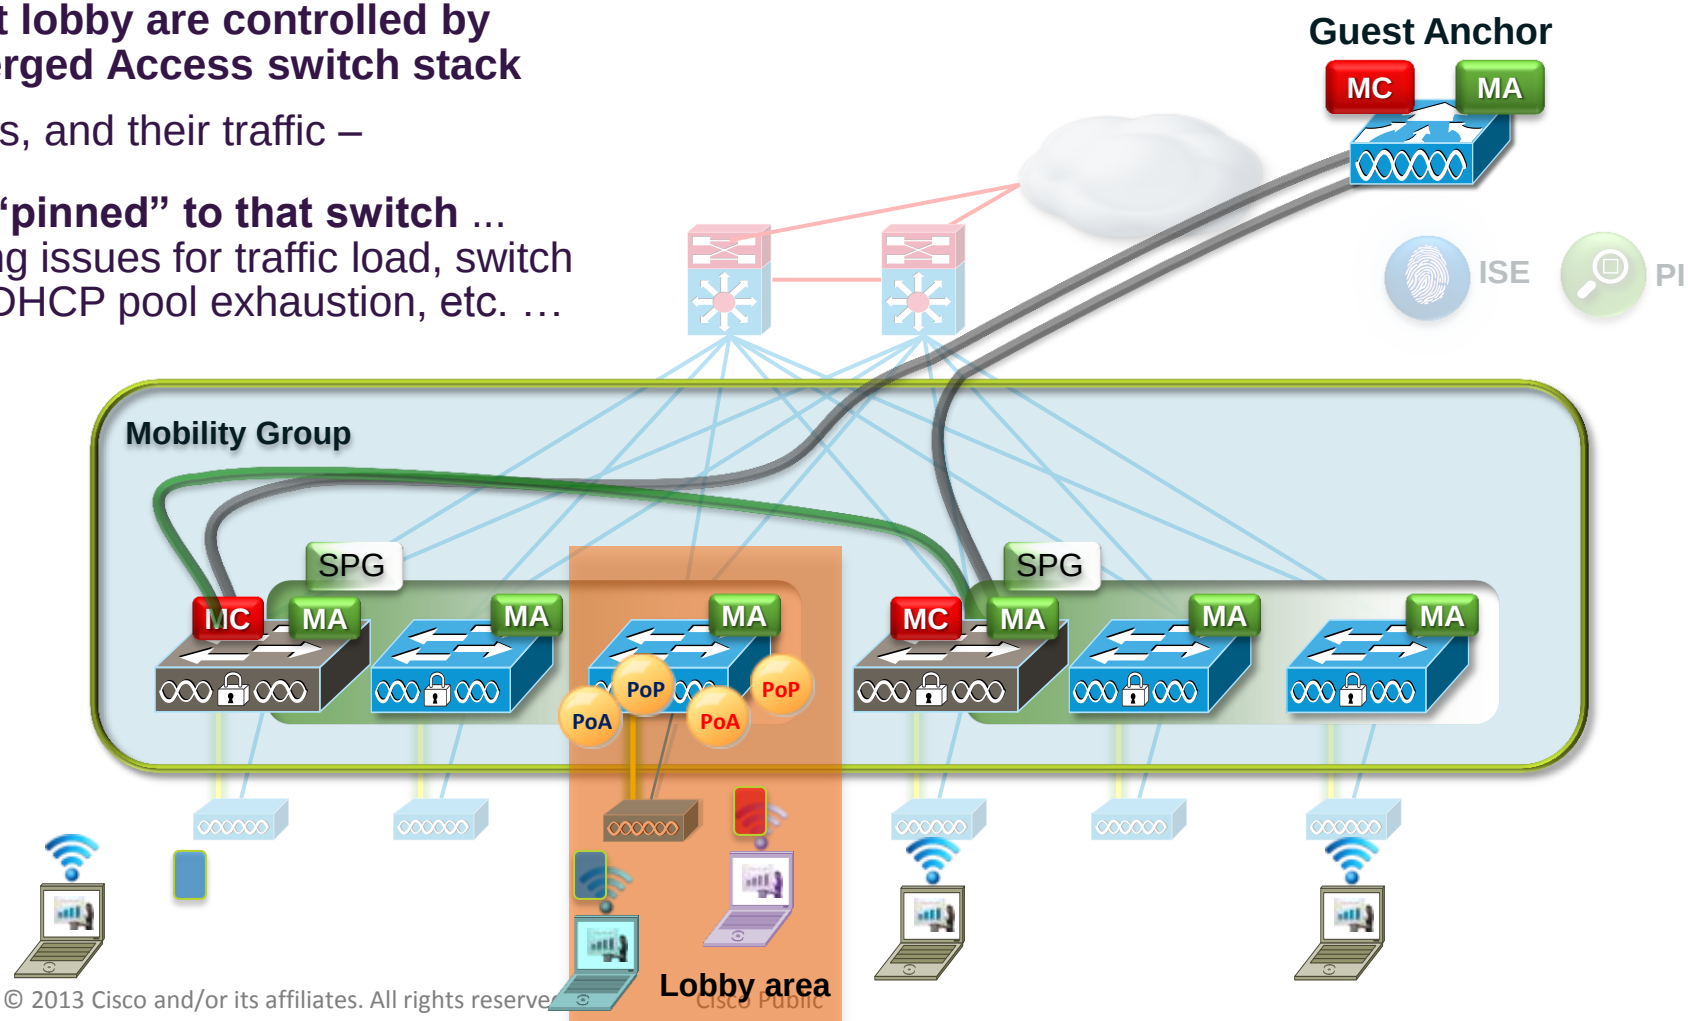

Converged Access

Common Building Access – The "Lobby Solution"

What can we do to address this issue?

- **User client association can be distributed** across Converged Access switches in the Switch Peer Group
- **User load info is constantly shared within the SPG –** with heartbeat (10s default, adjustable 1s-30s)
- **At a defined client load**, the lobby Catalyst 3850 switch distributes incoming client association requests to it's SPG members … the client is anchored based on reported client load
  - **Addresses** traffic load, switch load, DHCP pool exhaustion, etc.

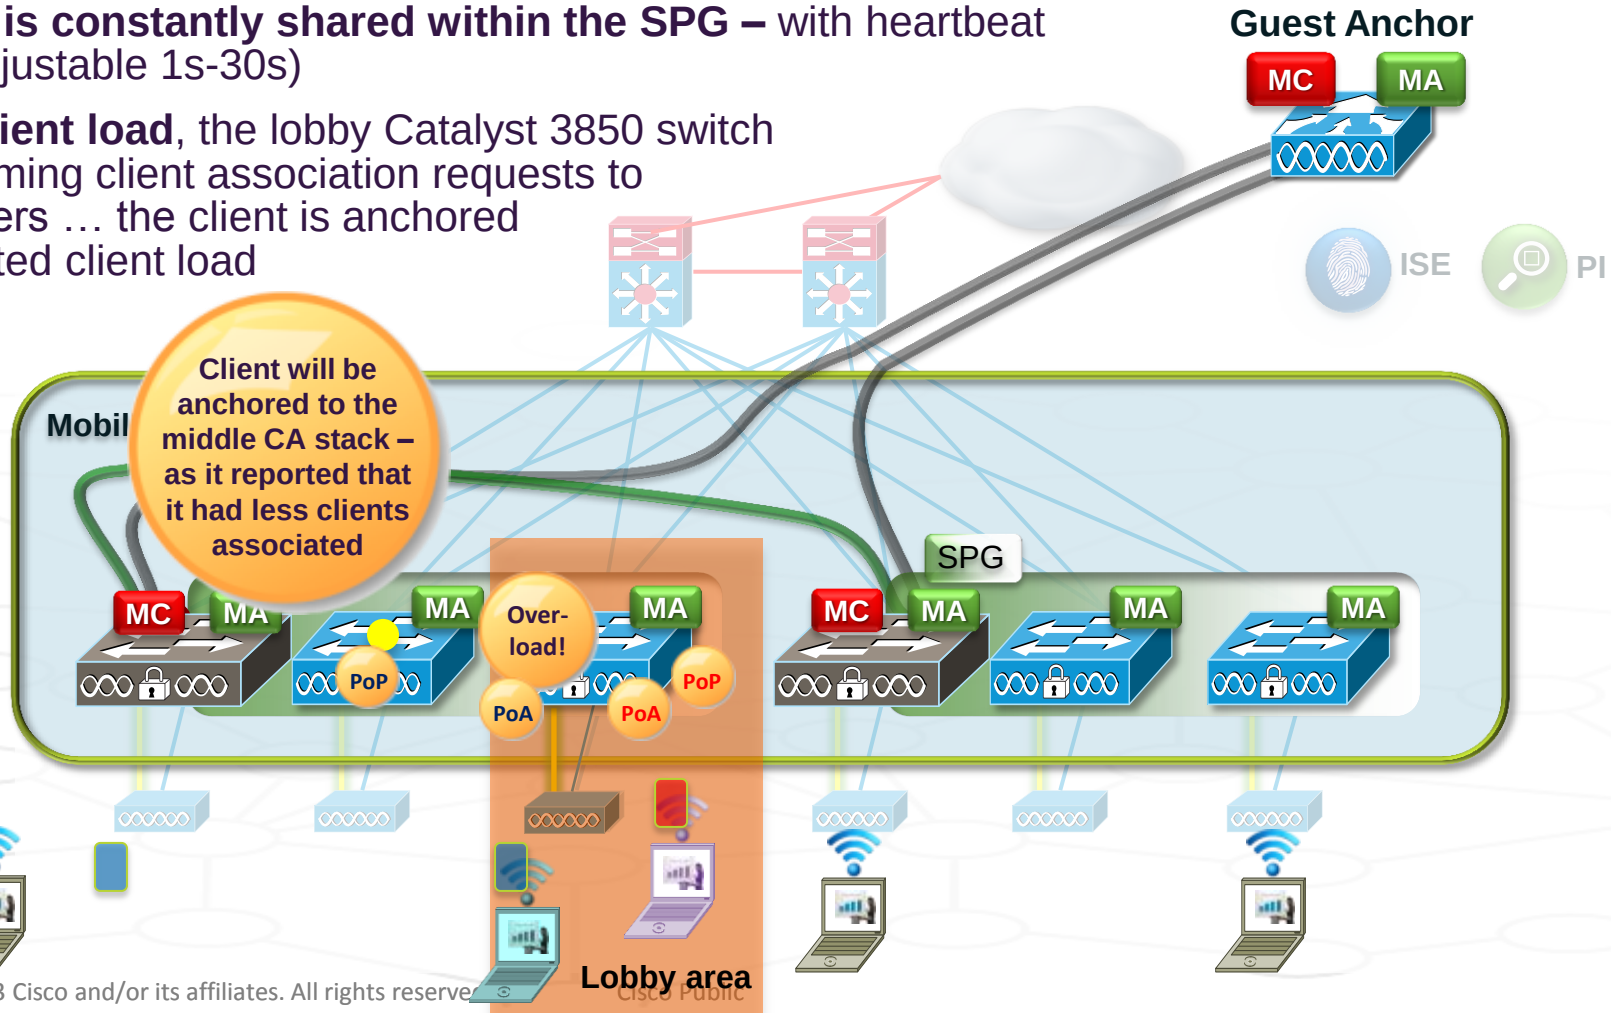

Converged Access

Common Building Access – The "Lobby Solution", Detail

- **What**: when configured, the client first PoA is load balanced across the switches in the SPG. When the client joins, the switch checks if its load is over a configurable threshold and send a message to anchor the client to least loaded switch in the SPG.
- **Why**: large number of clients could potentially attach to a single MA whose APs are situated close to the front door / lobby. This would result into congestion at that home switch, whereas other MAs would be under-utilised. This is even worse if the client's data path is anchored at the home switch.
- **How to configure it**: the feature is ON by DEFAULT and it's possible to change the threshold value. By default is 50% (of the max client allowed)

  To configure a different threshold use the following command on a per MA basis –

```
3850(config)# wireless mobility load-balance threshold ?

 <100-2000>  Threshold value for number of clients that can be anchored locally
```