

TOMORROW starts here.



Cisco *live!*

Converged Access Architecture, Design and Deployment

BRKARC-2665

Dave Zacks

Distinguished Systems Engineer

Converged Access – Architecture, Design, and Deployment

BRKARC-2665 – Session Overview and Objectives

Cisco is bringing together the best of wired and wireless networking into “One Network” with Converged Access.

This session introduces the Converged Access solution, including the next-generation Catalyst 3850 switch, and how you can employ this exciting new platform within your network – discussing various design considerations and placement within various network deployments.

You will learn how the Converged Access architecture is designed and operates, how roaming works seamlessly, as well as many of the Multicast, NetFlow, QoS, and Security features inherent within the solution.

This session is targeted to Network Managers, Architects and Administrators.

Converged Access – Architecture, Design, and Deployment

Your Instructor Today ... **Dave Zacks**

I am a **Distinguished Systems Engineer**, and have been with Cisco for 14+ 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 have been involved with the Converged Access solution within Cisco for 3+ years.

Quick note – Throughout this presentation, the term “3x50” is used to refer to **both the Catalyst 3850 and 3650 platforms** together (saves space rather than spelling out 3850 / 3650 every time). **Only these two platforms are inferred when the term “3x50” is used in this presentation.**



Dave Zacks

Distinguished Systems Engineer

dzacks@cisco.com

Agenda – BRKARC-2665

► Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –

- IP Addressing

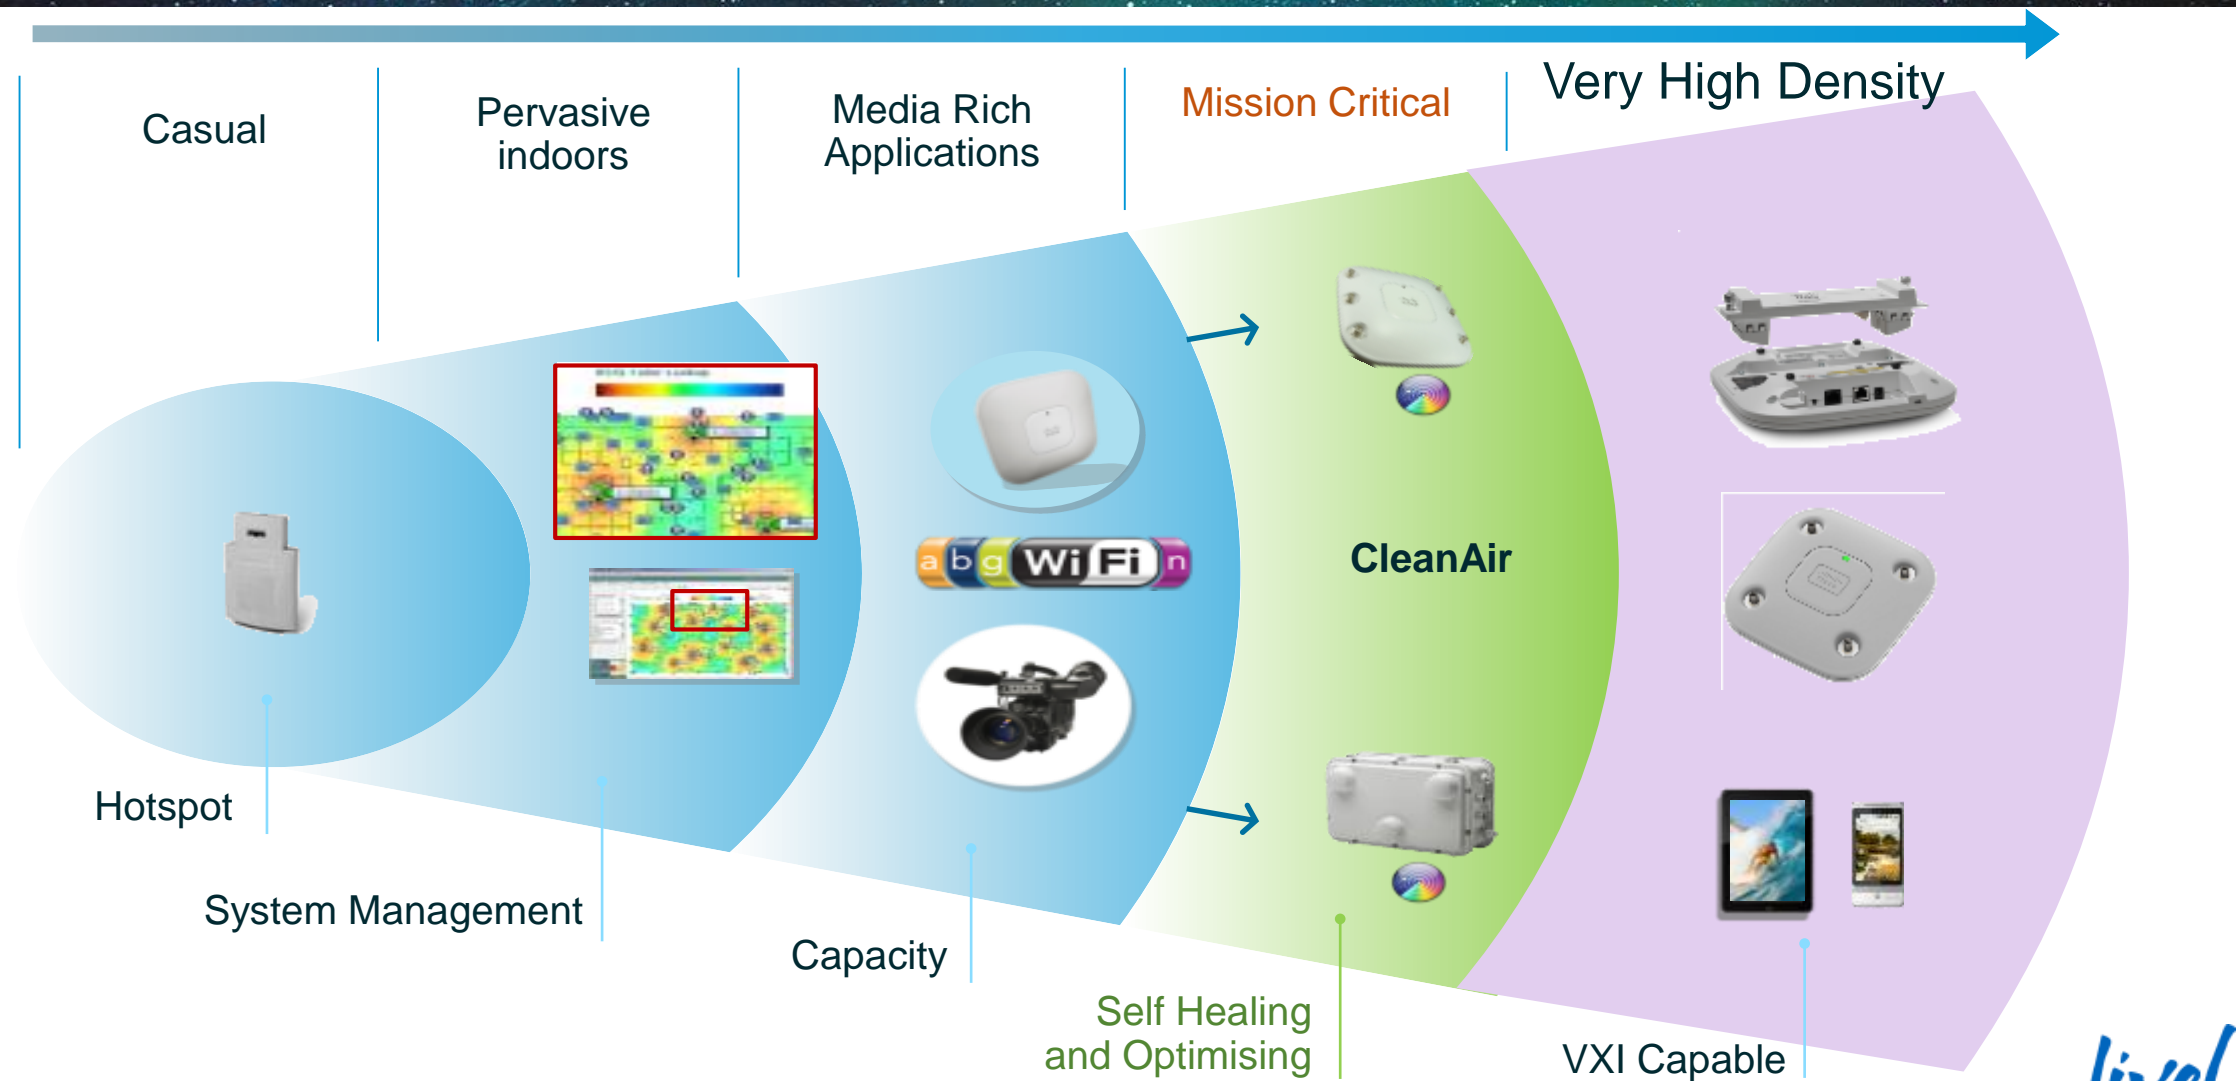
- Design Options

- Deployment Examples


Summary

Enterprise Wireless Evolution

From Best-Effort to Mission-Critical and Very High Density



Wireless Standards – Past, Present, and Future



How Many Mobile Data Devices Do You Think You Will Carry Everywhere in 2016?

Think about it, and choose the best answer

1

3

5

7



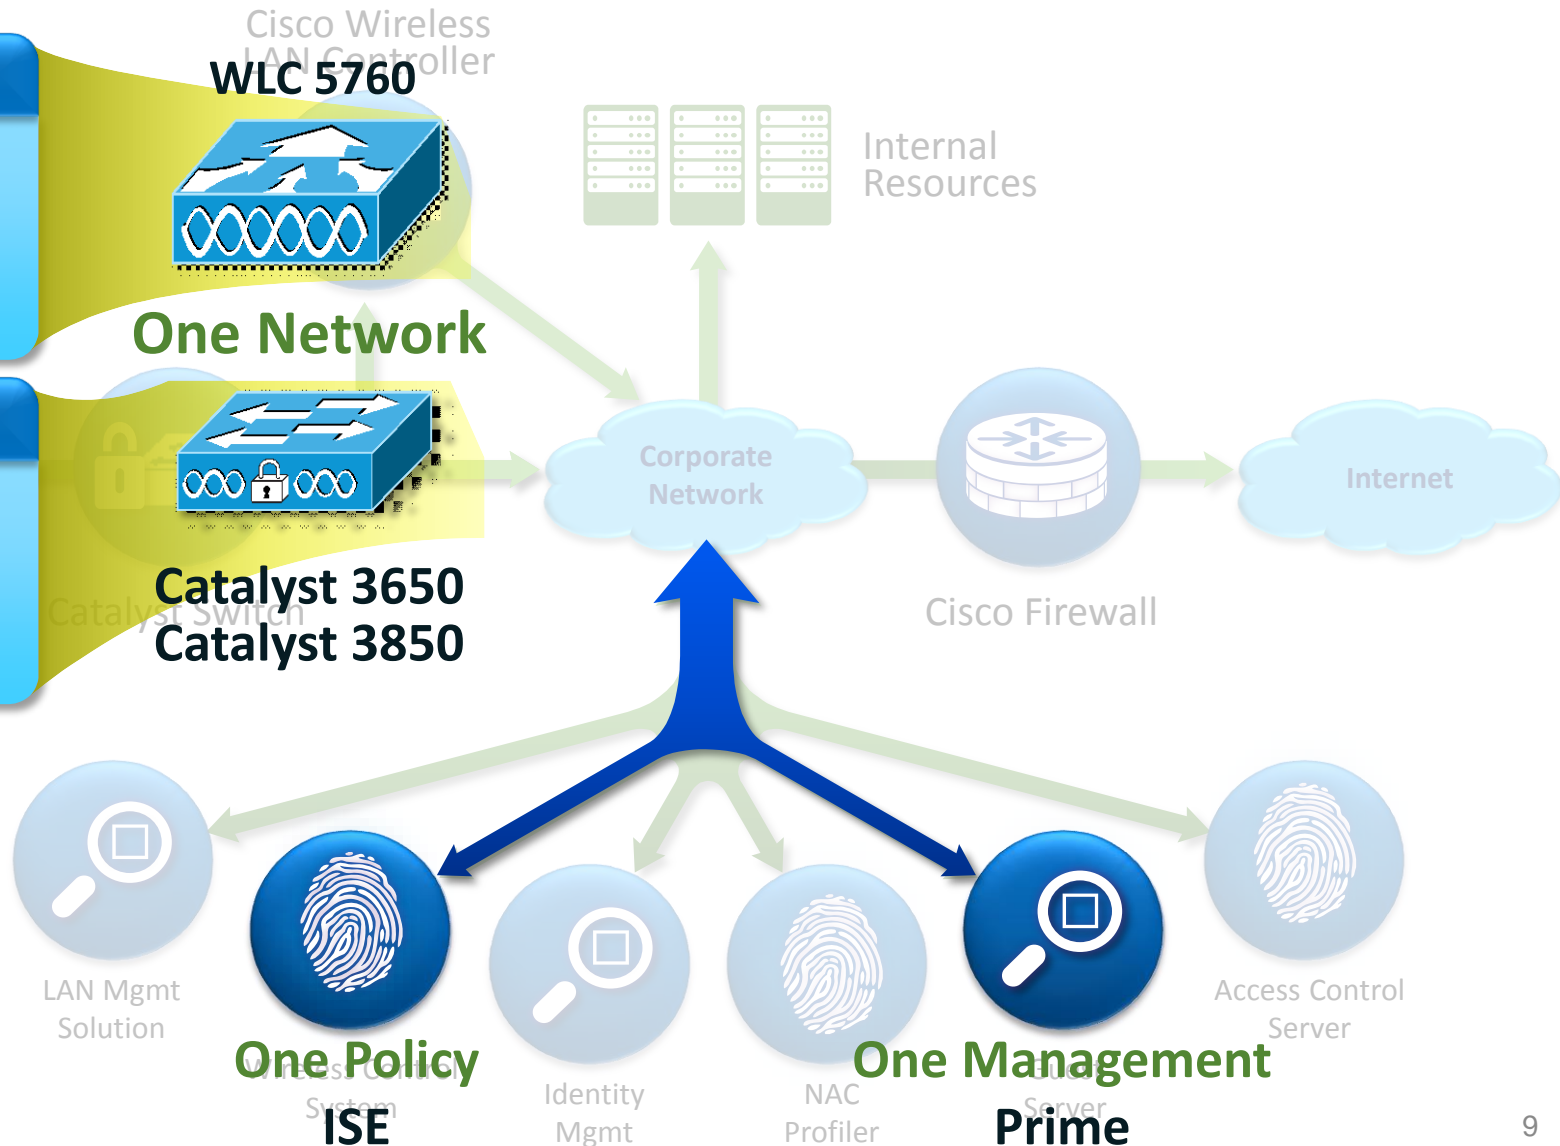
One Network, with Converged Access – A New Deployment Mode for Wired / Wireless

IOS Based WLAN Controller

- Consistent IOS and ASIC w/ Catalyst 3x50
- Required to scale beyond 200/250 AP or 8,000 / 16,000 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

Large stack bandwidth; 40G wireless / switch; efficient multicast; 802.11ac optimised

Unified Access - One Policy | One Management | One Network

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

► Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –

- IP Addressing

- Design Options

- Deployment Examples

Summary

Catalyst 3850 and 3650 Switches – 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



WIRED



Features:

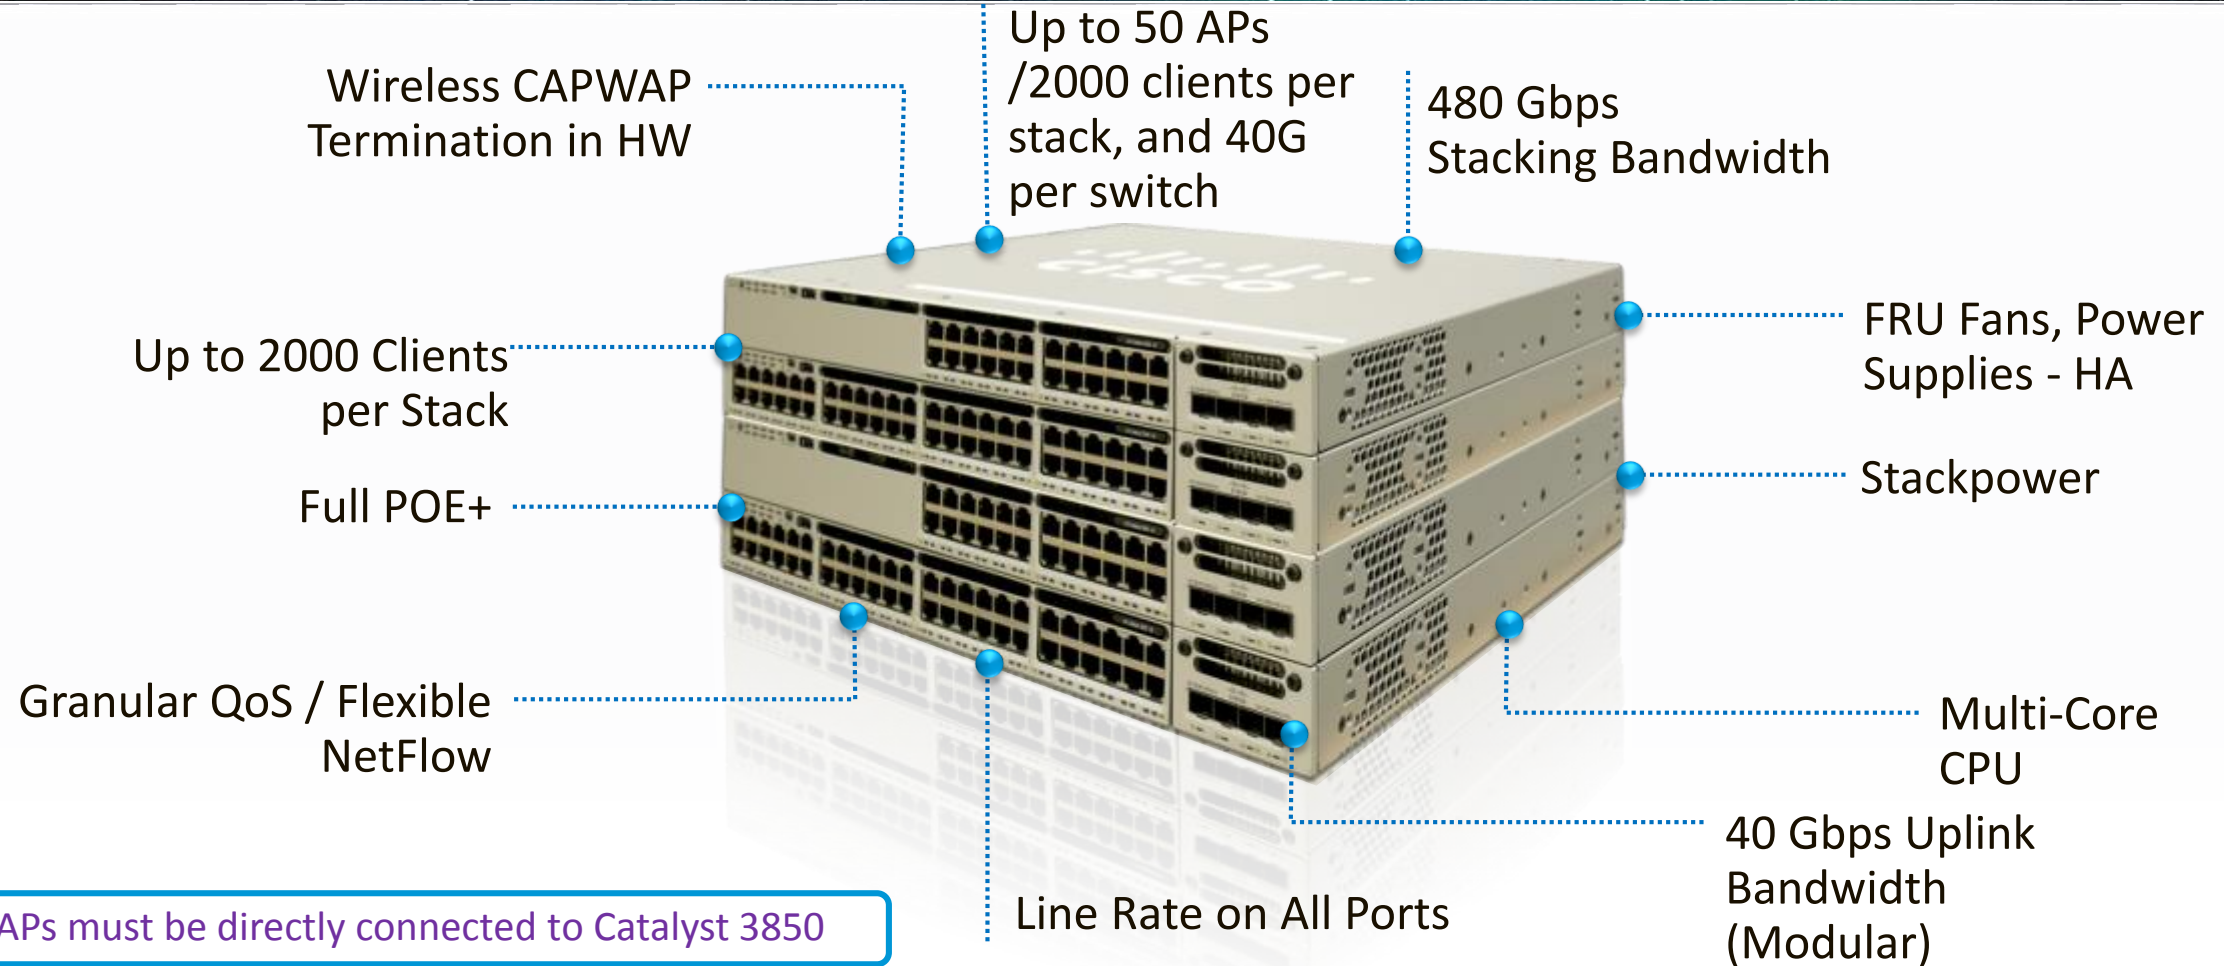
- Stacking and HA
- SGT & Advanced Identity
- Visibility and Control
- Flexible NetFlow
- Granular QoS
- Smart Operations
- EEM, scripting
- IOS-XE Modular OS



Benefits

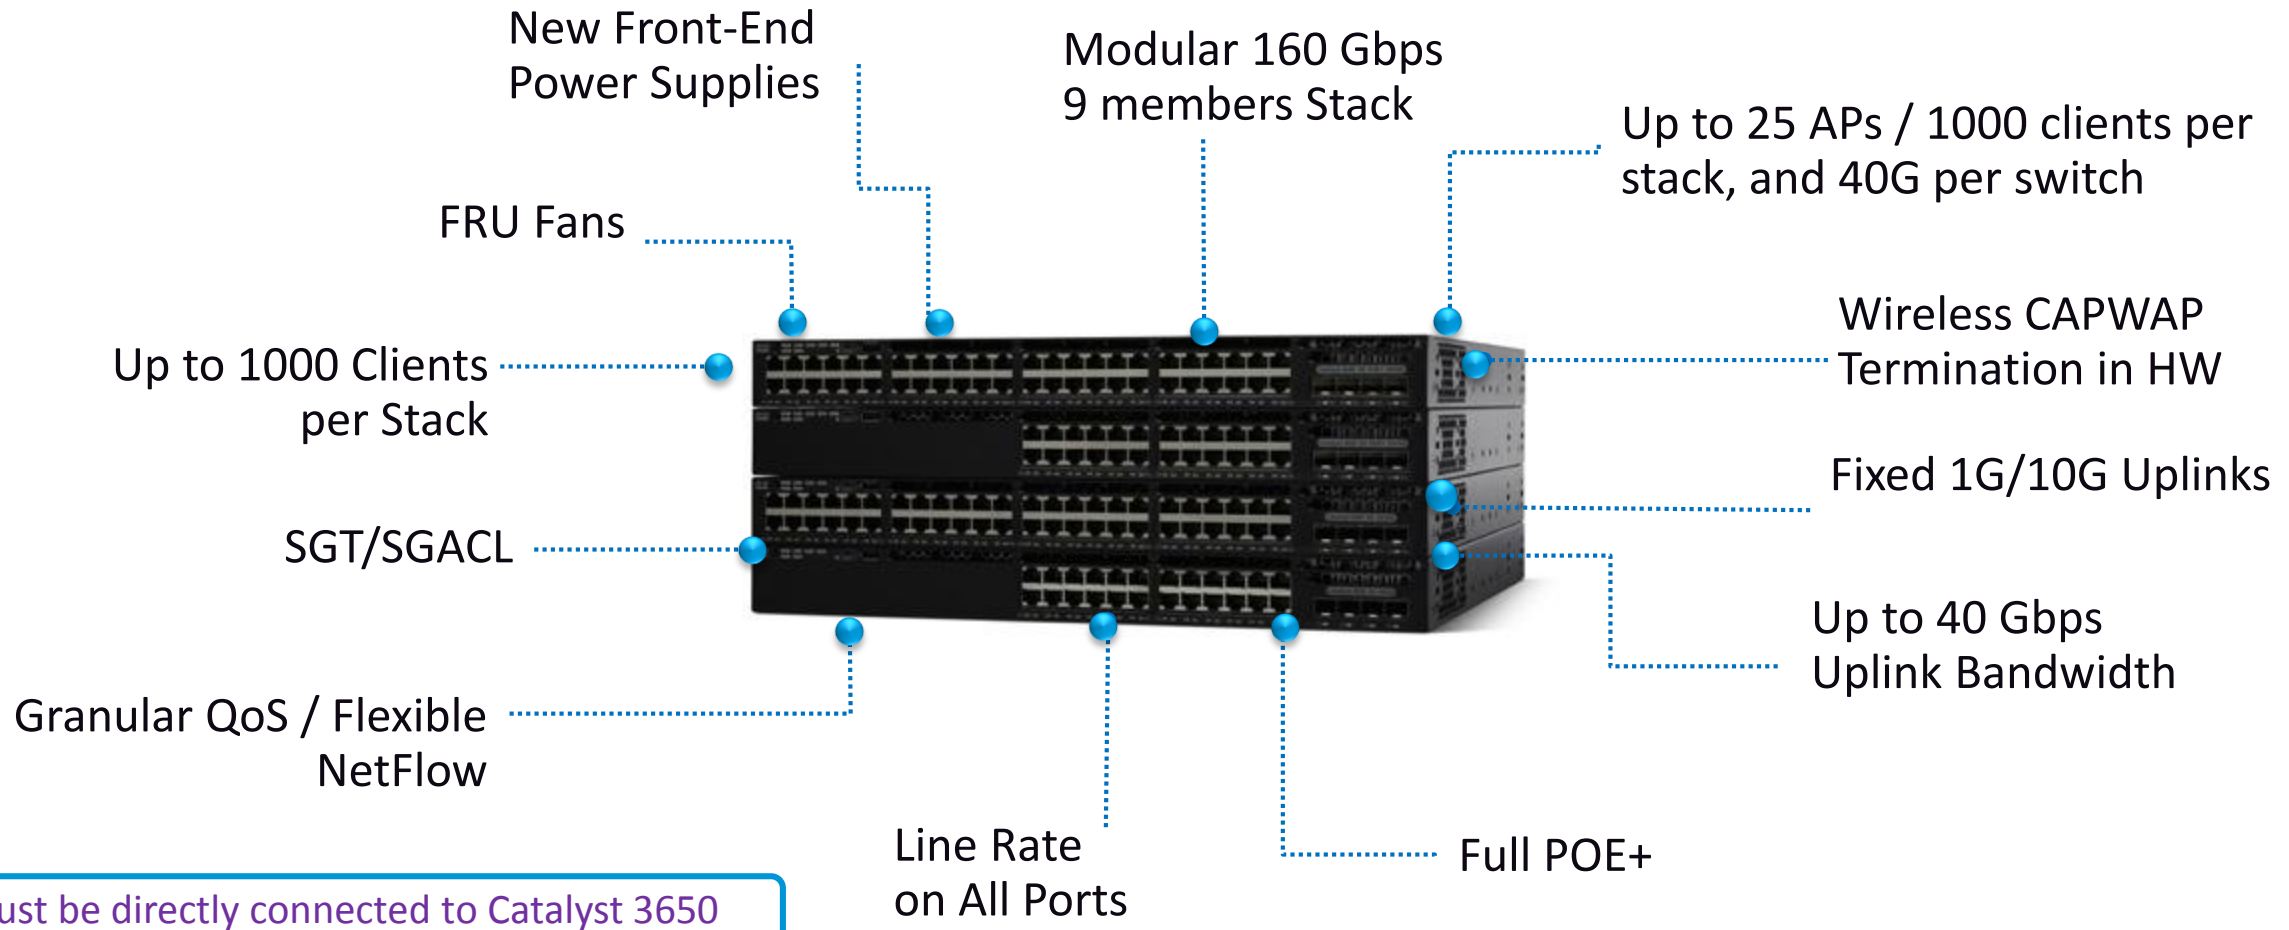
- Built on **UADP** – Cisco's Innovative Flexparser ASIC technology
- Eliminates operational complexity
- Single 'modern' Operating System for wired and wireless

Catalyst 3850 – Platform Overview



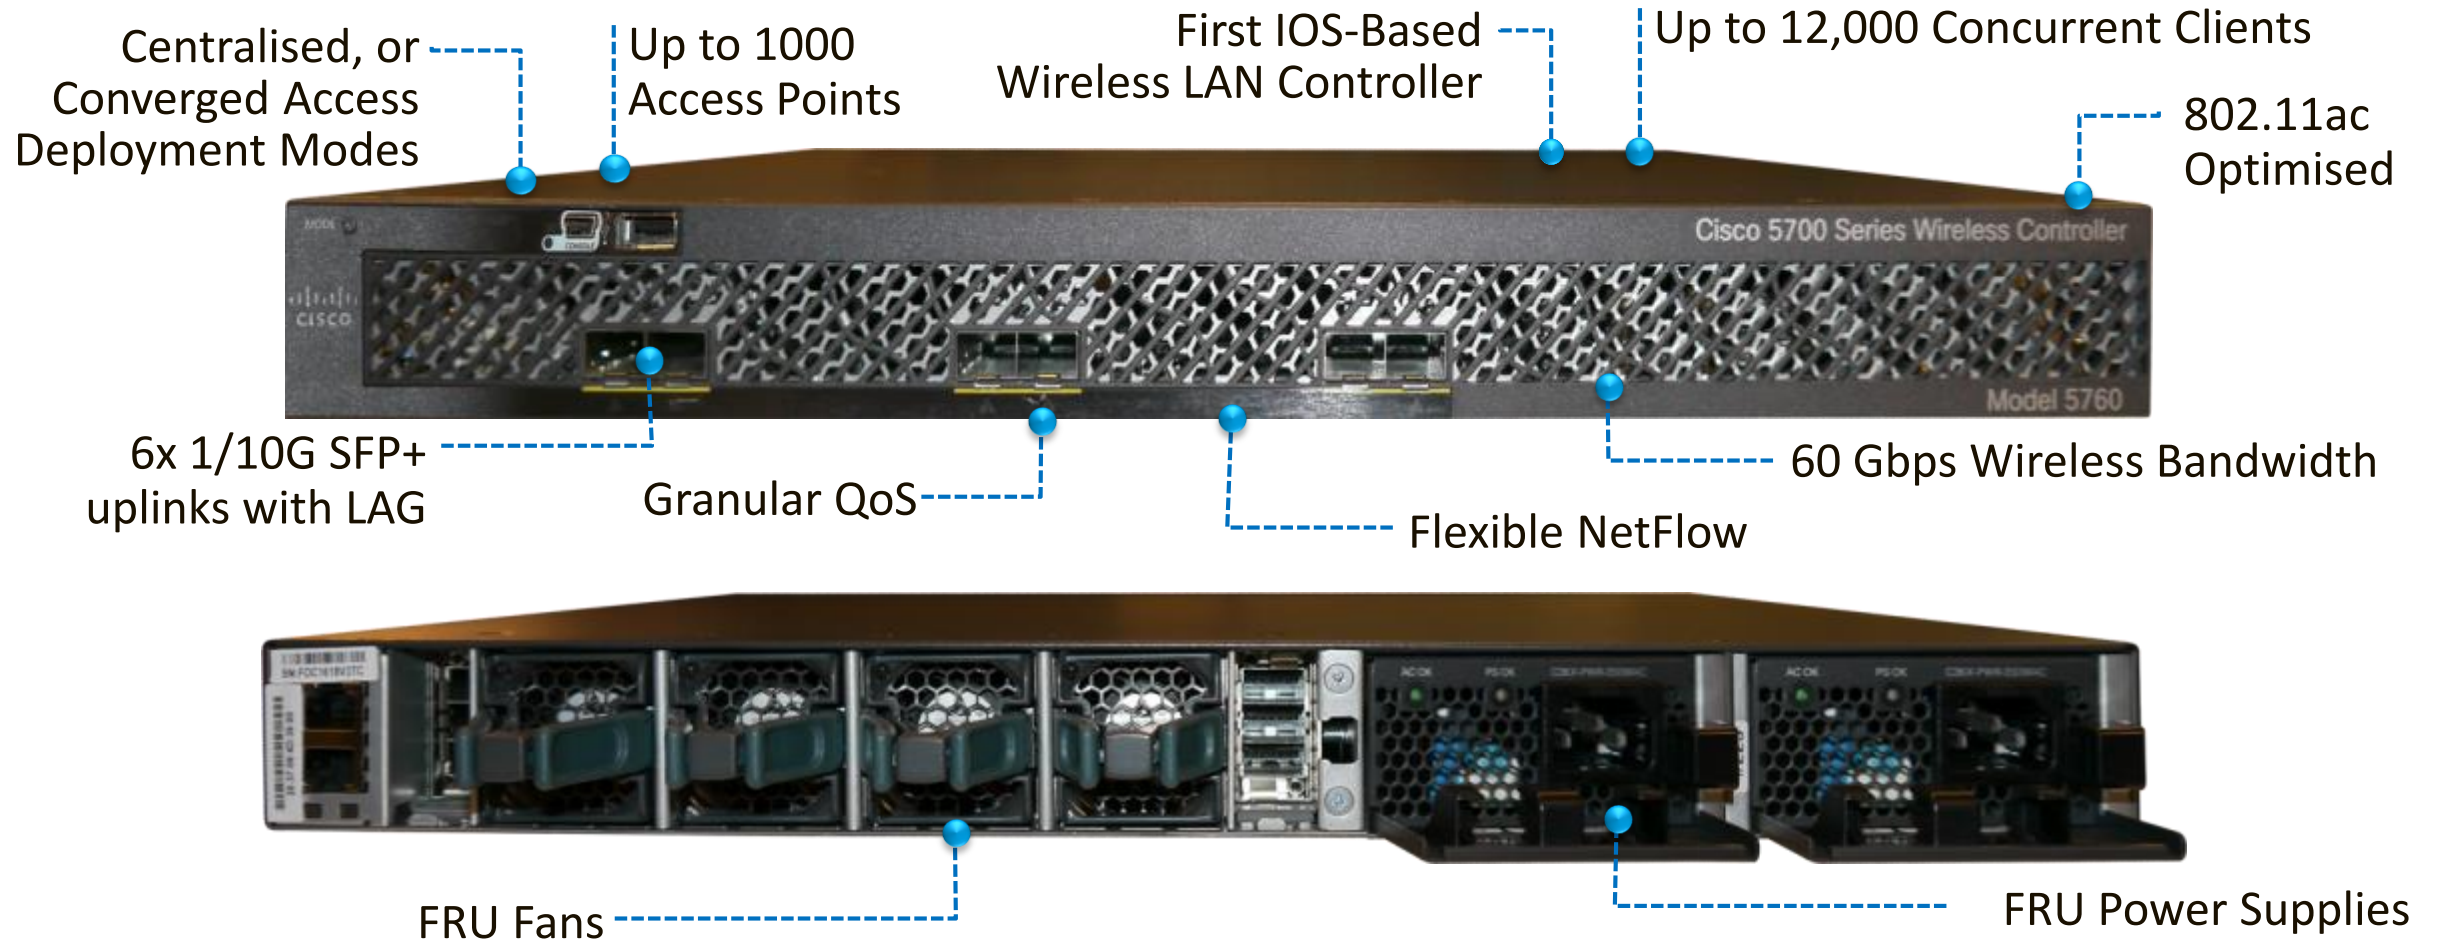
Built on Cisco's Innovative "UADP" ASIC

New Catalyst 3650 Switch – Platform Overview



Built on Cisco's Innovative "UADP" ASIC

Wireless LAN Controller (WLC) 5760 – Platform Overview



Built on Cisco's Innovative "UADP" ASIC

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network
Converged Access – Platform Overviews

► Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks
- Traffic Flows and Roaming
- High Availability
- Quality of Service
- Security
- Multicast
- NetFlow

Converged Access Design and Deployment –

- IP Addressing
- Design Options
- Deployment Examples

Summary

Converged Access Architecture – What We're Going to Cover



Corner Stones

CA System Architecture

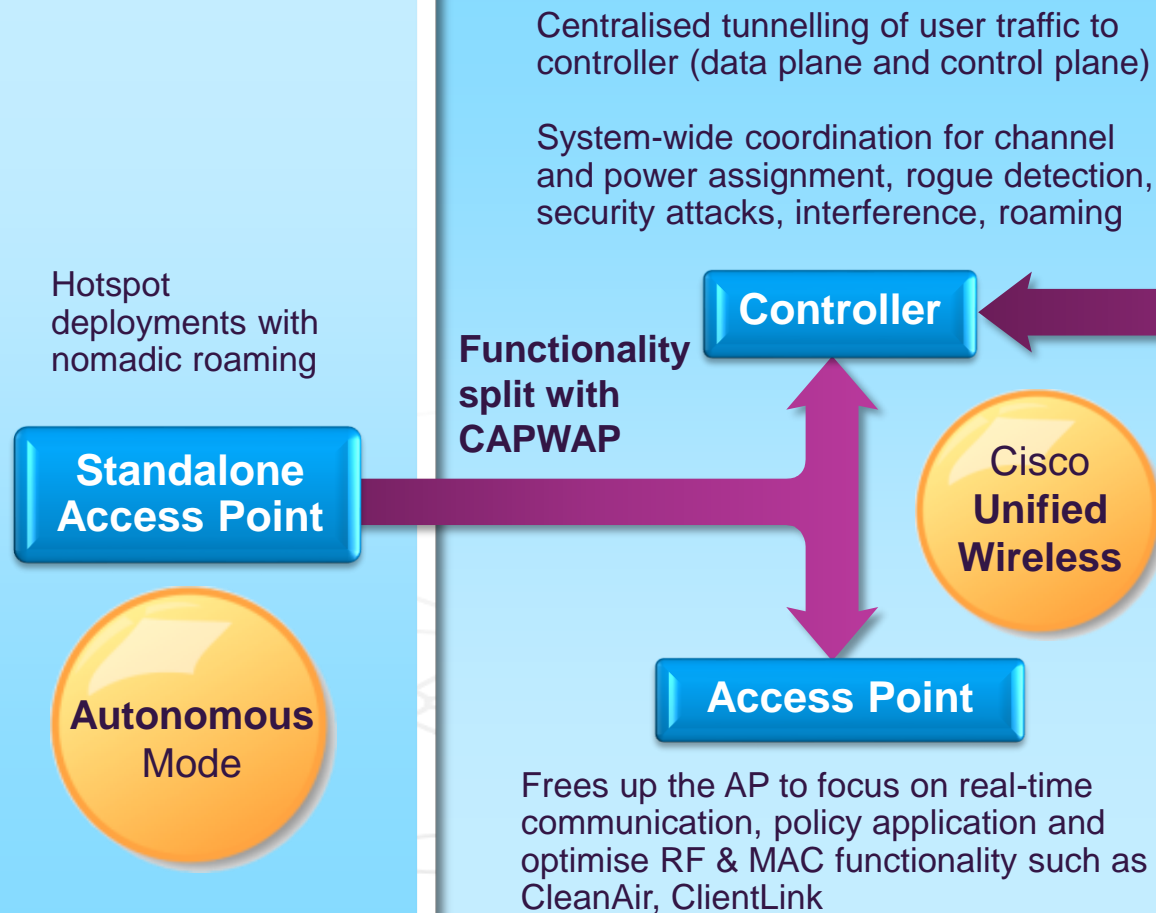
**Roaming, HA, QoS,
Security, NetFlow, Mcast**

Deployment and Design

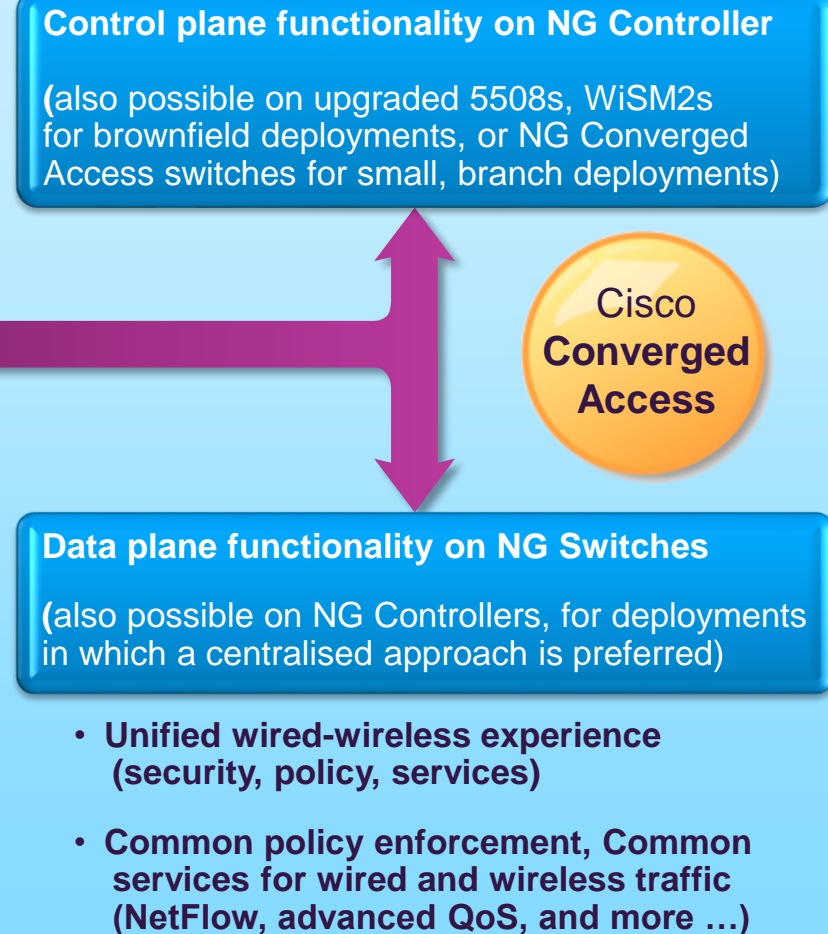
**Foundational Elements
for the Converged Access Solution**

Converged Access – Network Requirements Driving Wireless Evolution

We've Been Here Before...



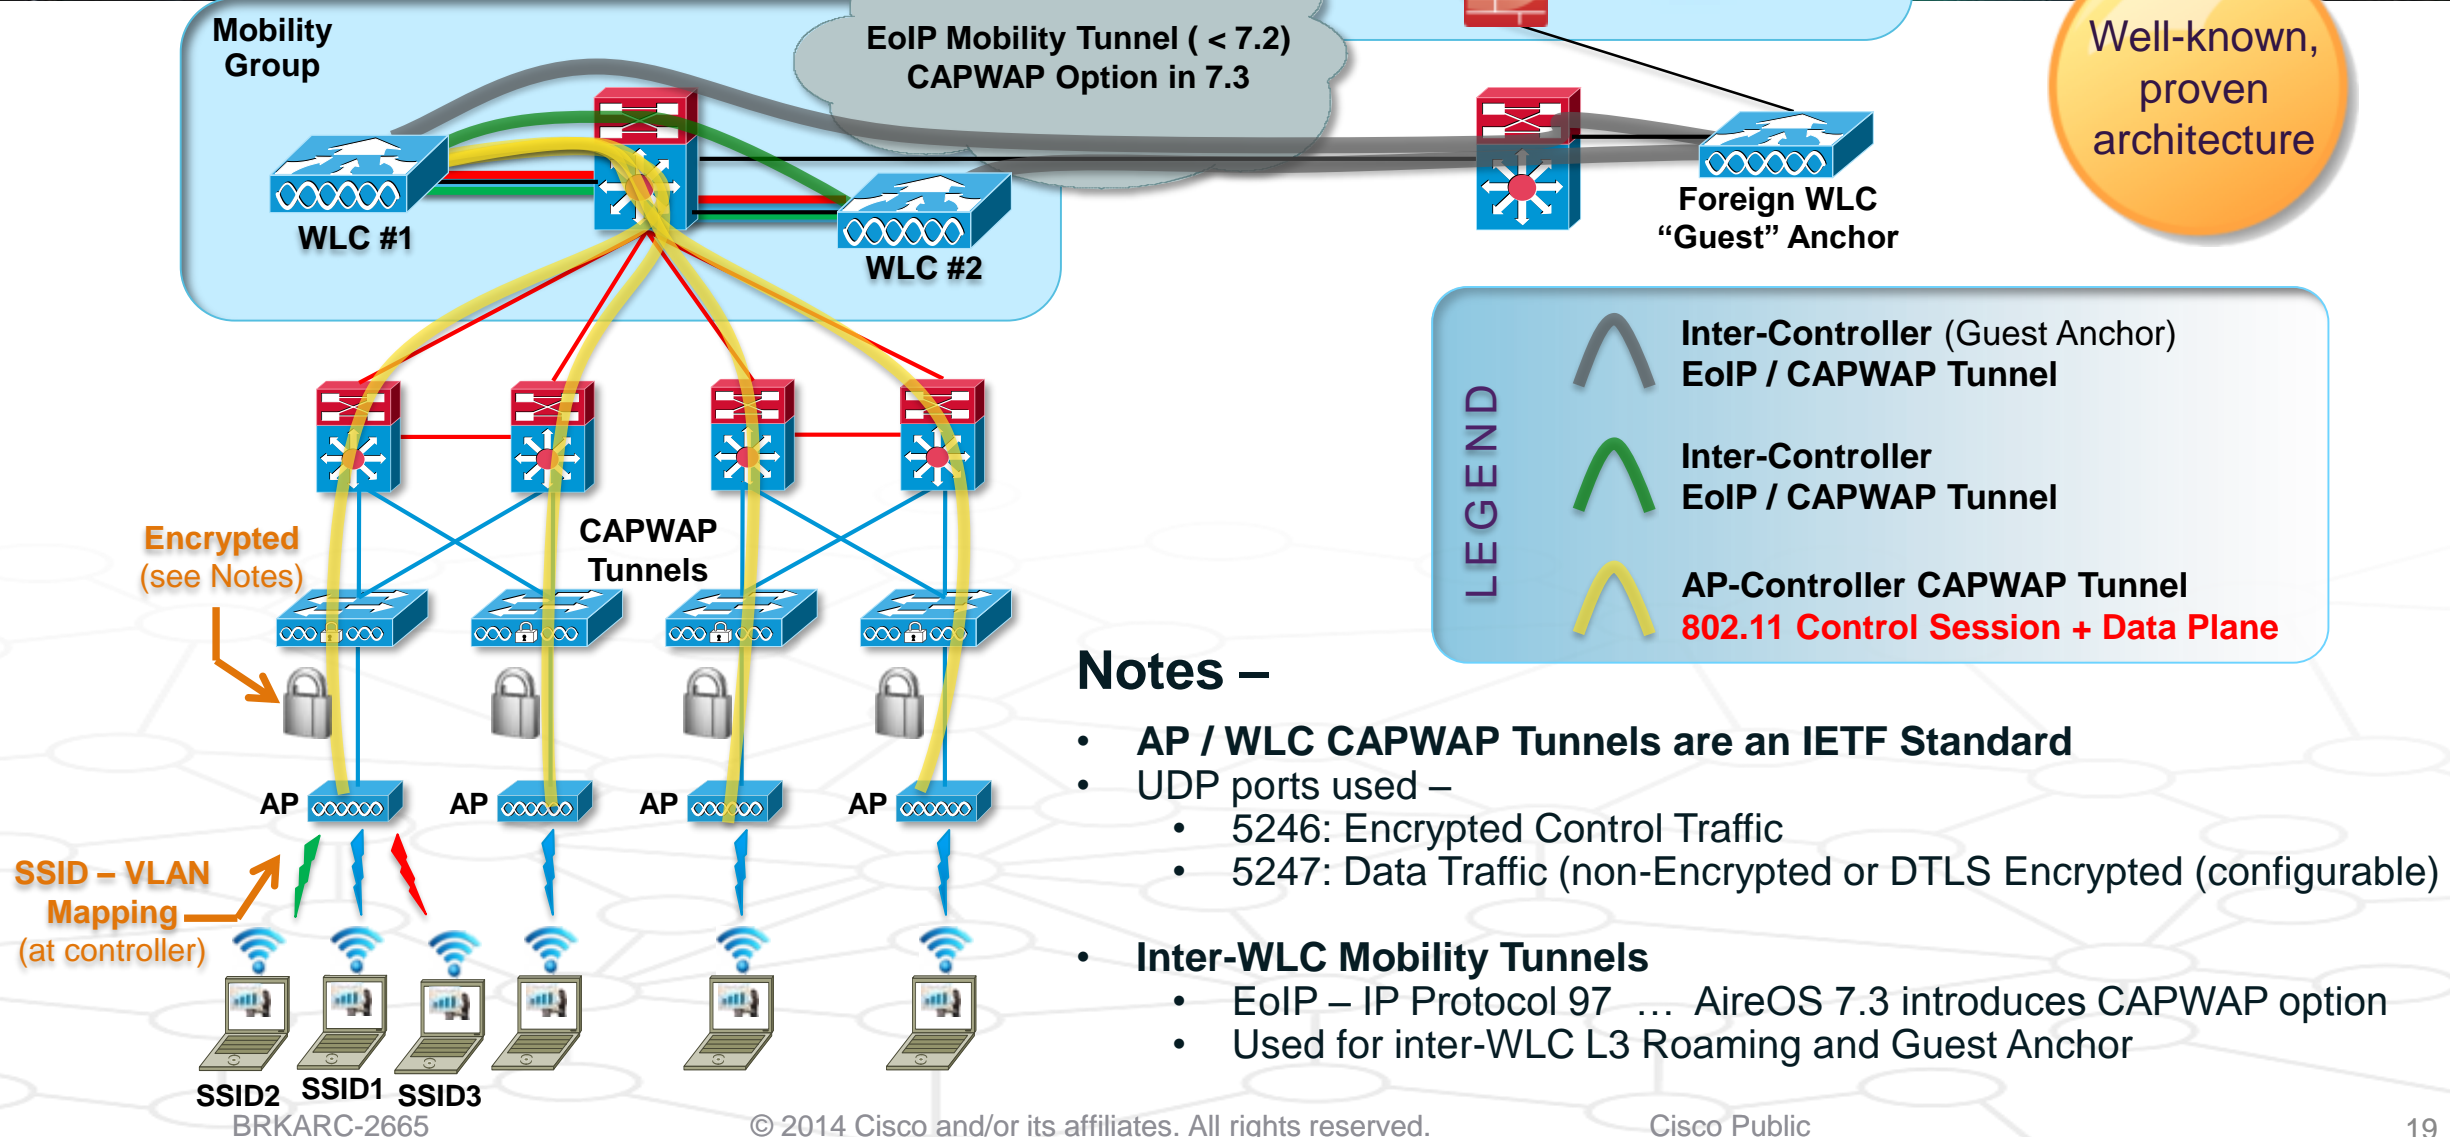
Increased scalability, Centralised policy application



Scale and Services

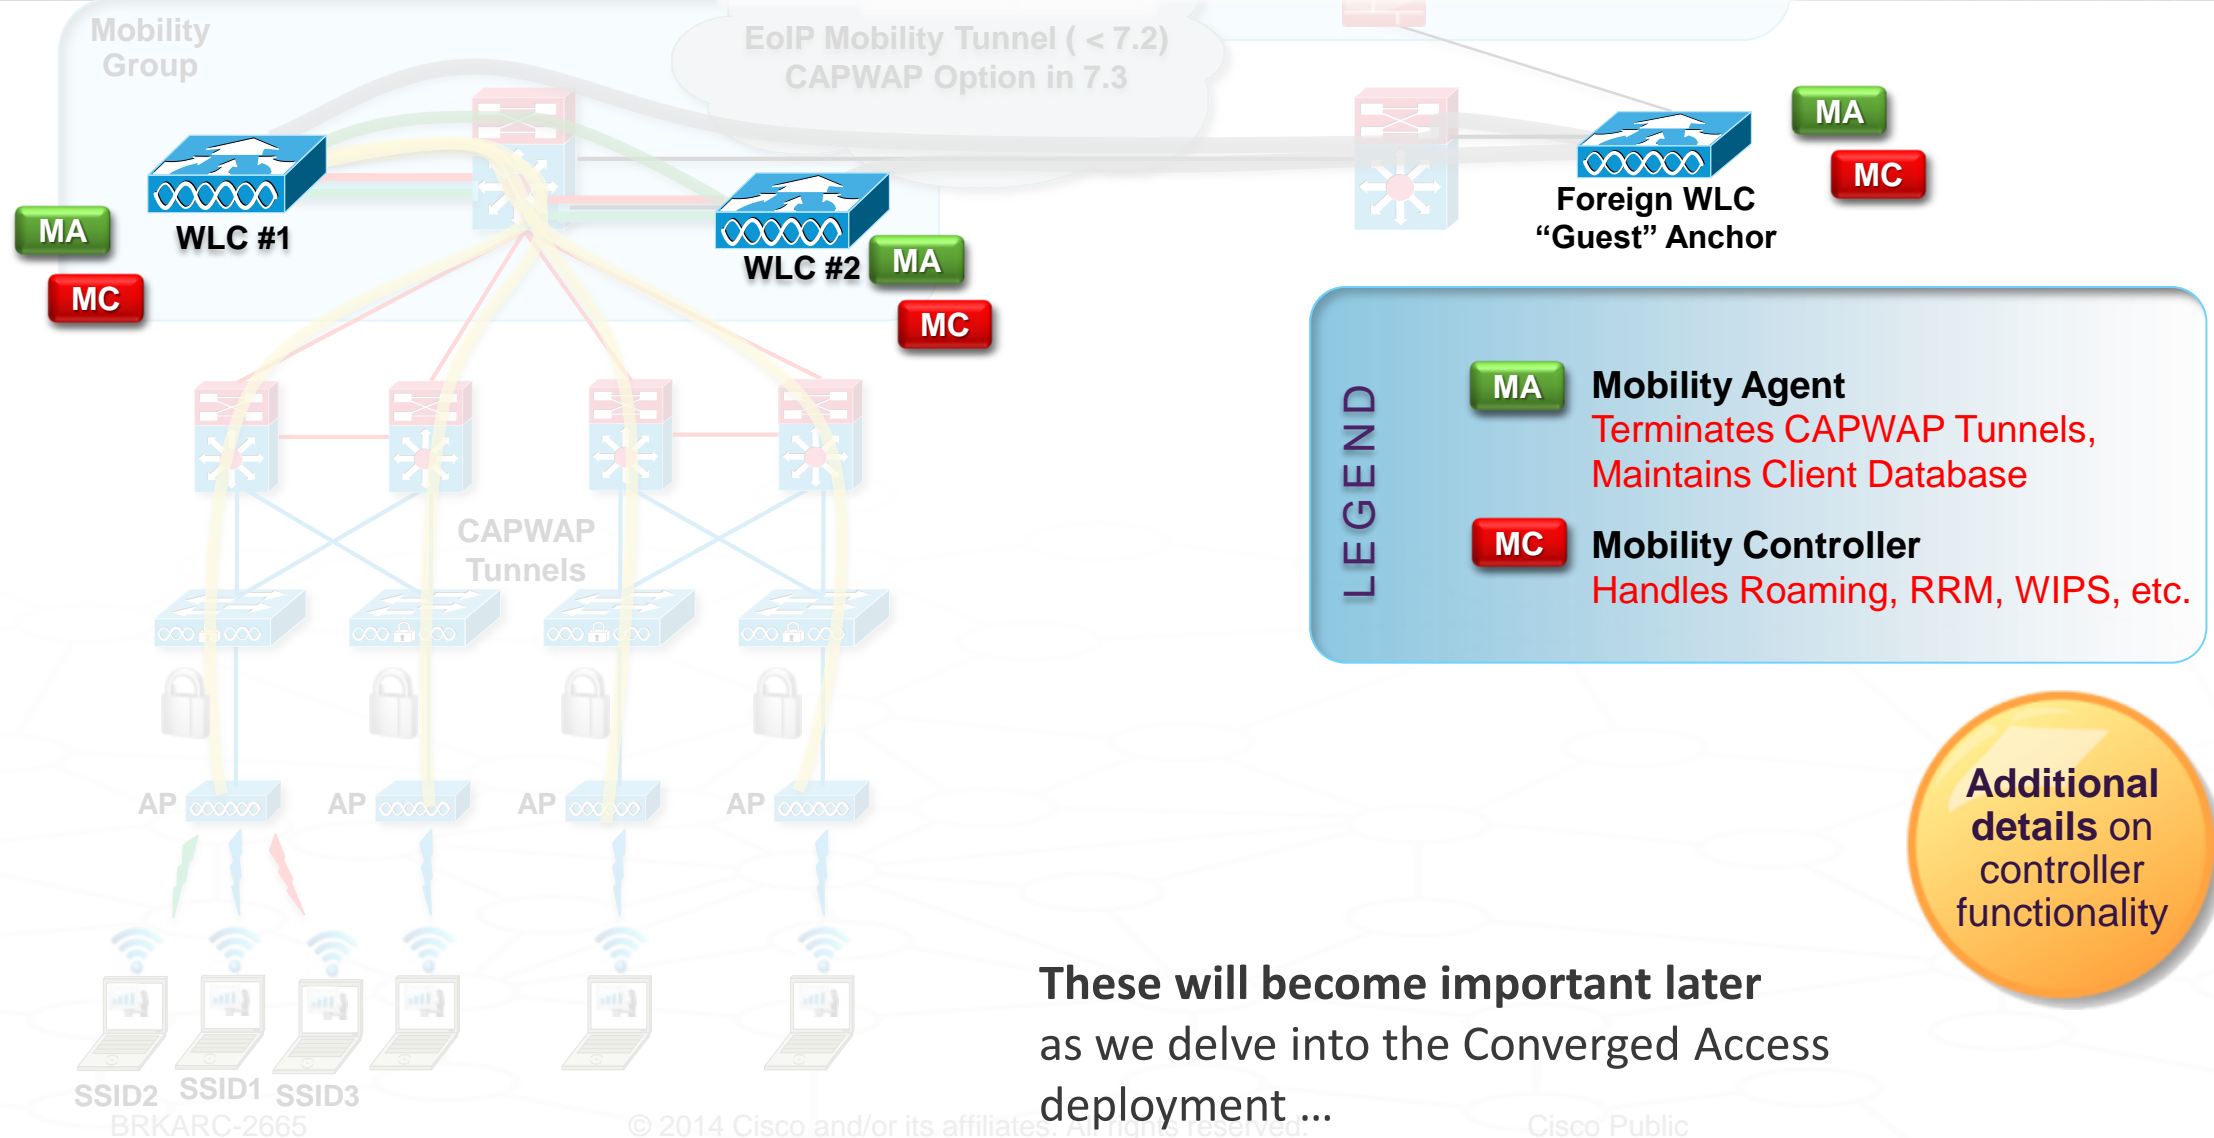
Performance and Unified Experience

Architecture Constructs – CUWN Tunnel Types



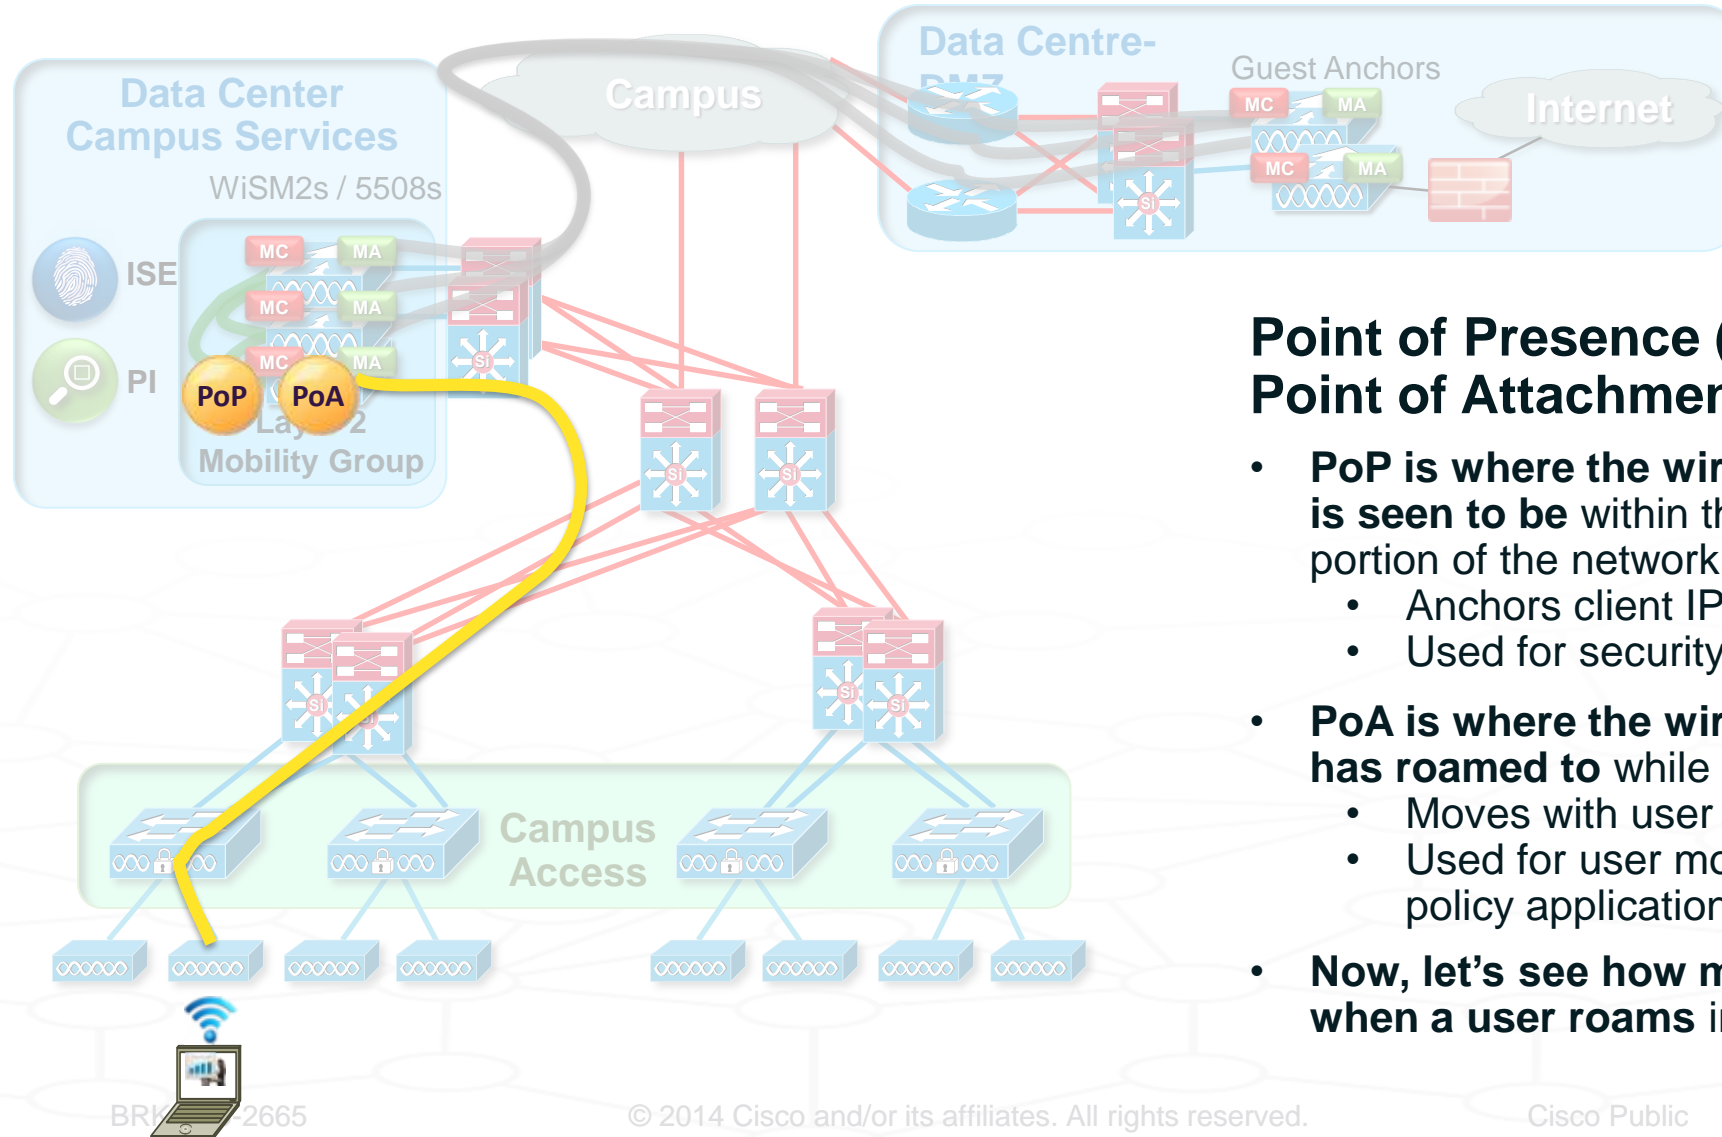
Architecture Constructs – CUWN Control Functions

Existing Unified Wireless Deployment today ...



Architecture Constructs –

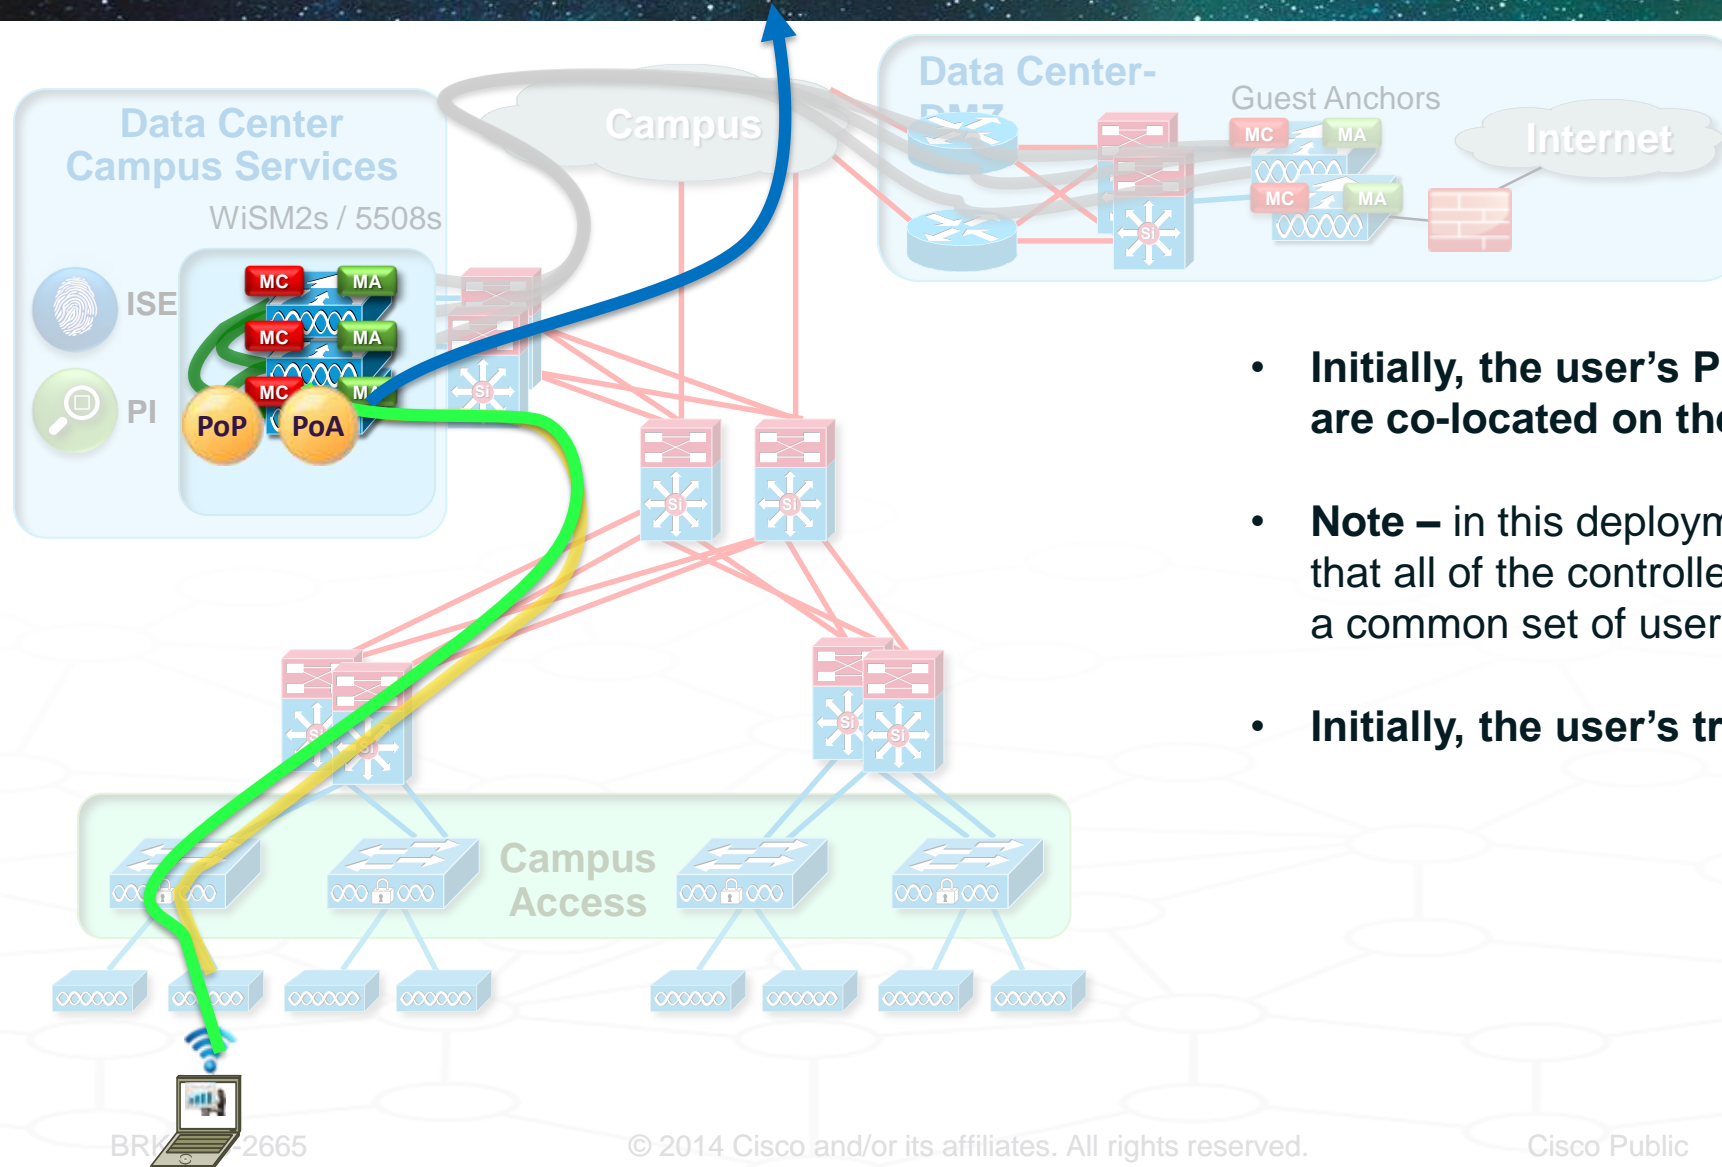
Point of Presence (PoP), Point of Attachment (PoA)



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 ...**

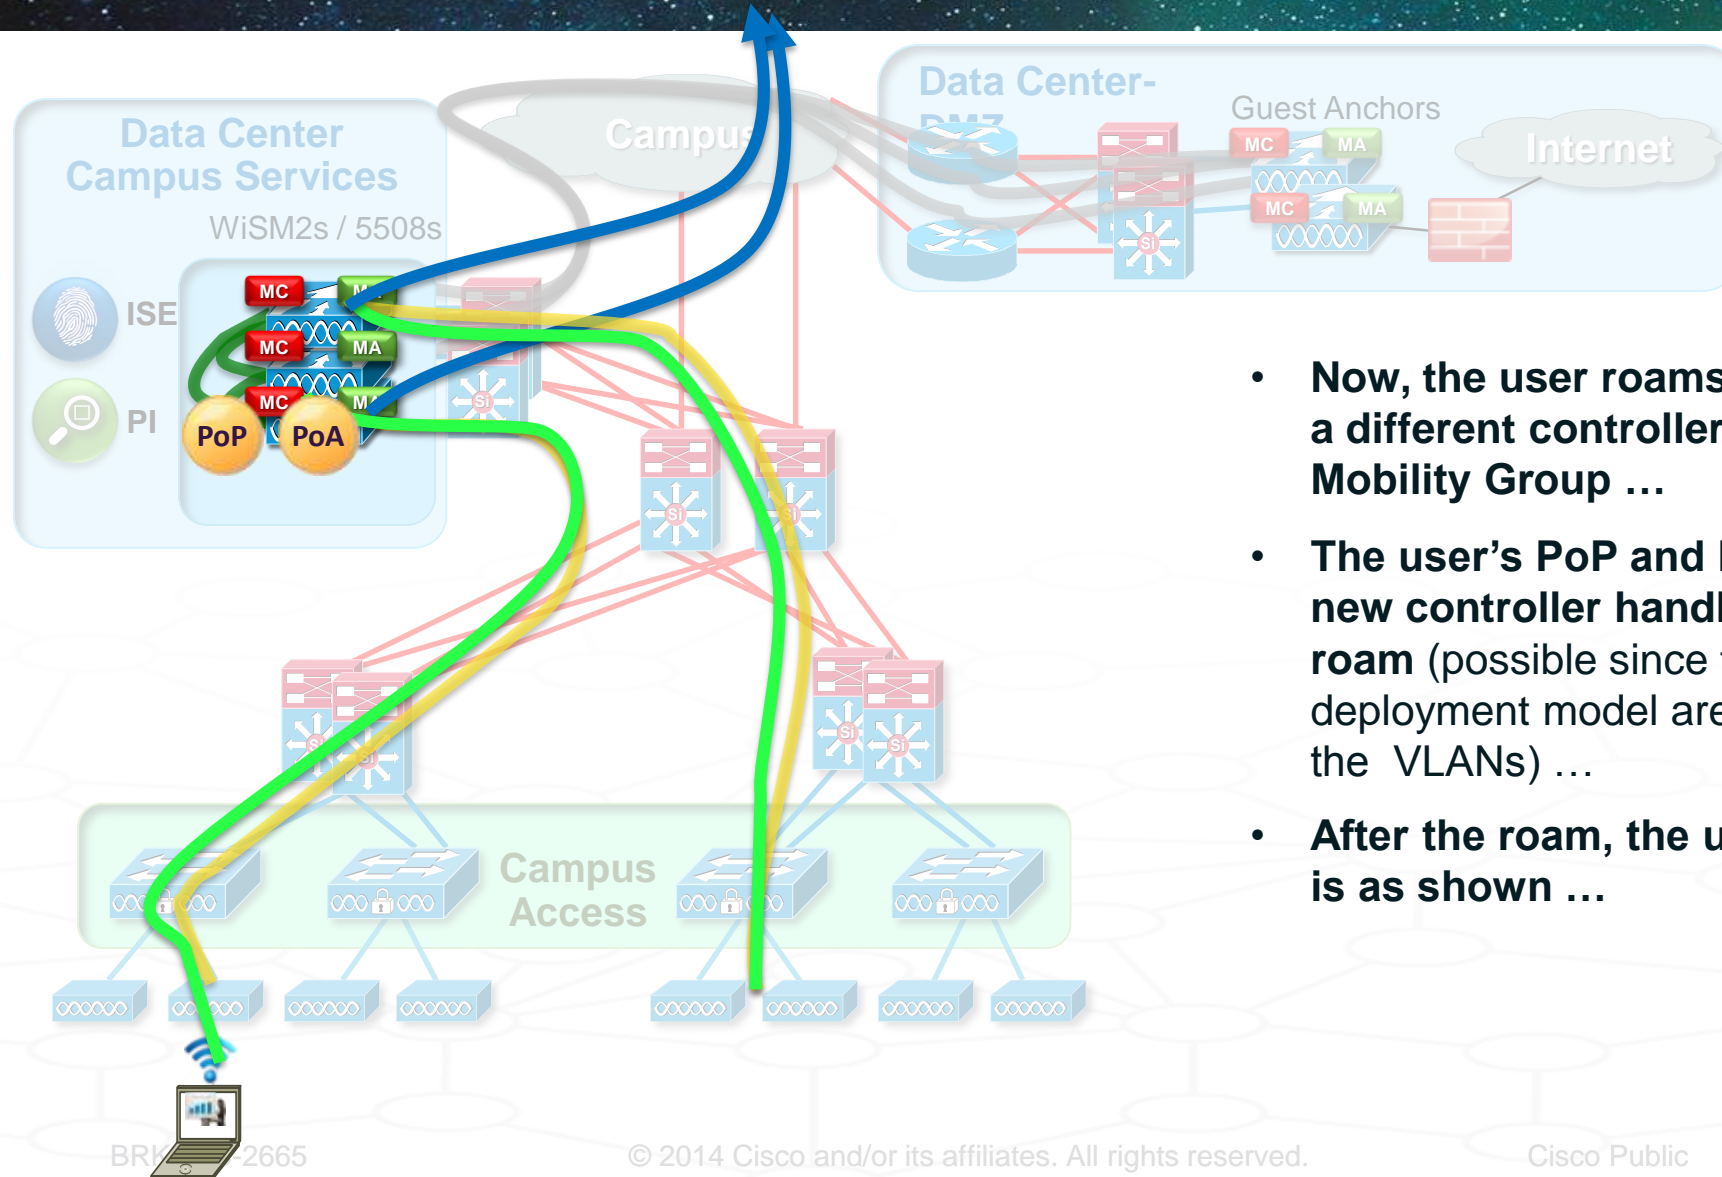
Architecture Constructs – Layer 2 Roaming (Campus Deployment)



- **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 ...**

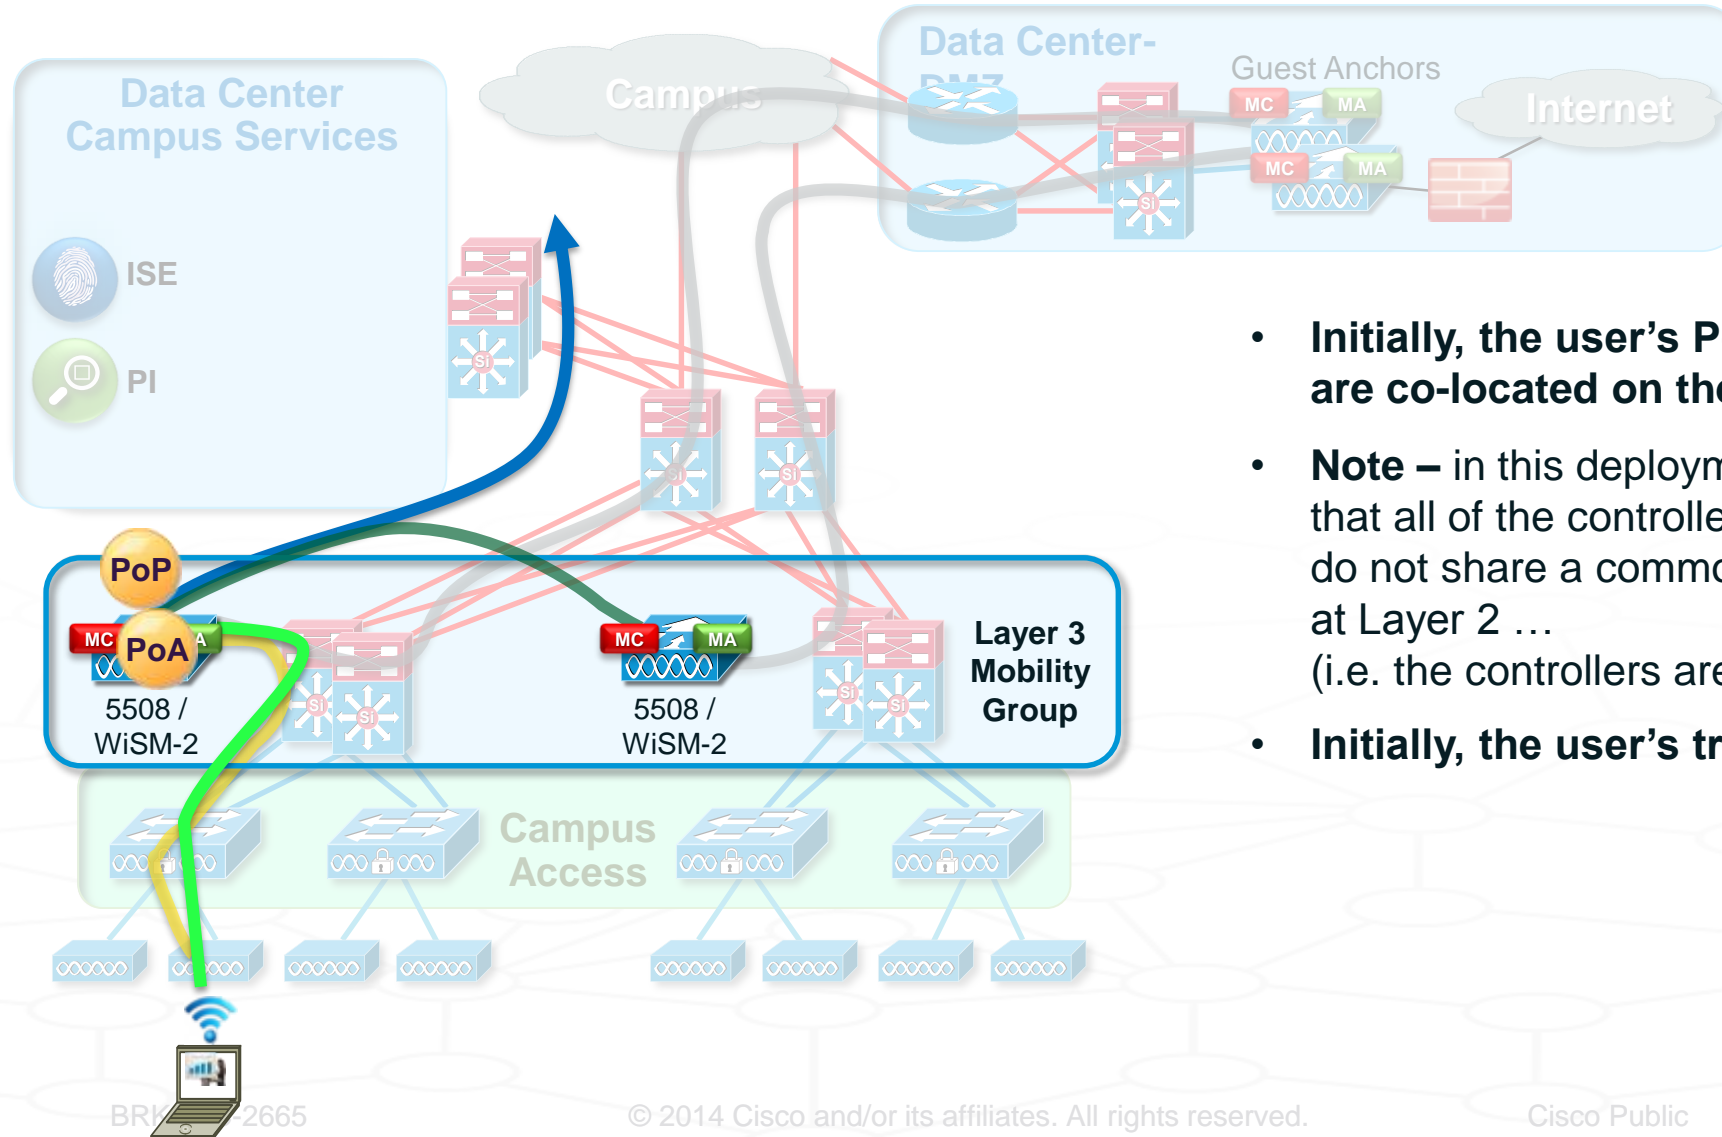
Architecture Constructs – Layer 2 Roaming (Campus Deployment)

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 ...

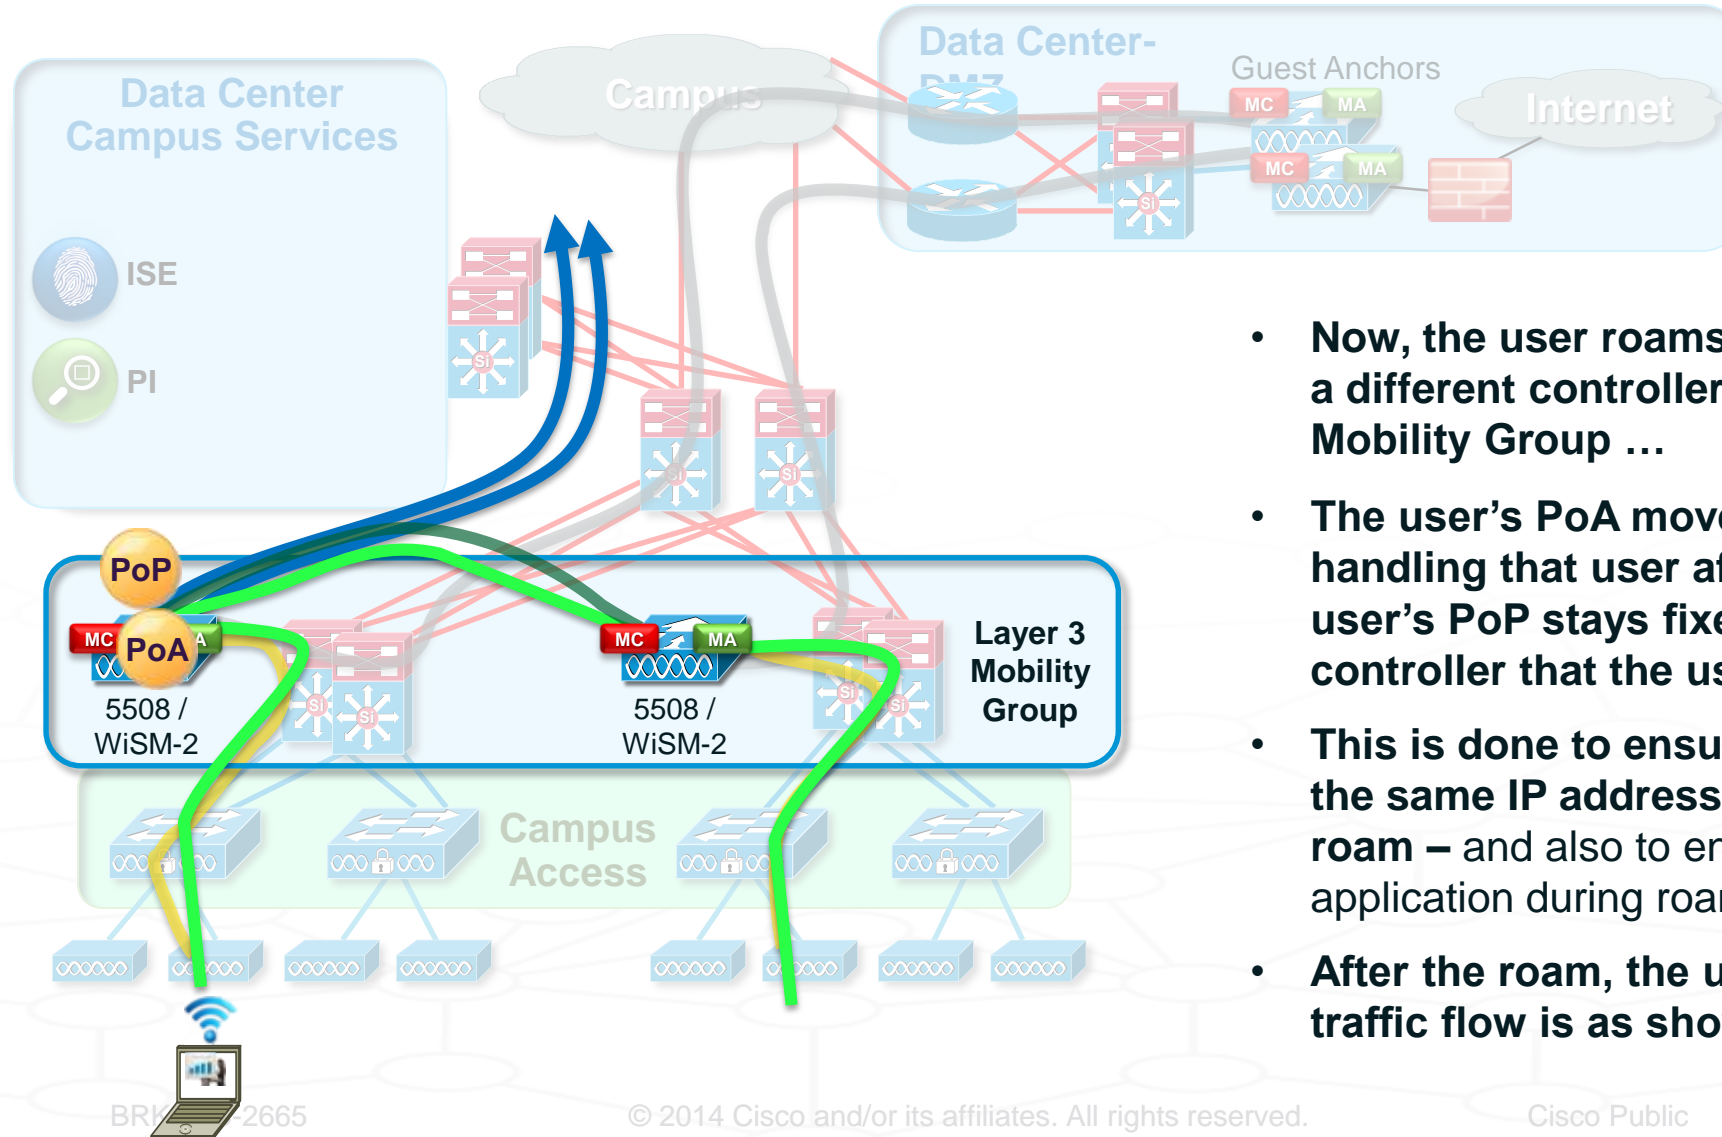
Architecture Constructs – Layer 3 Roaming (Campus Deployment)



- 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 ...

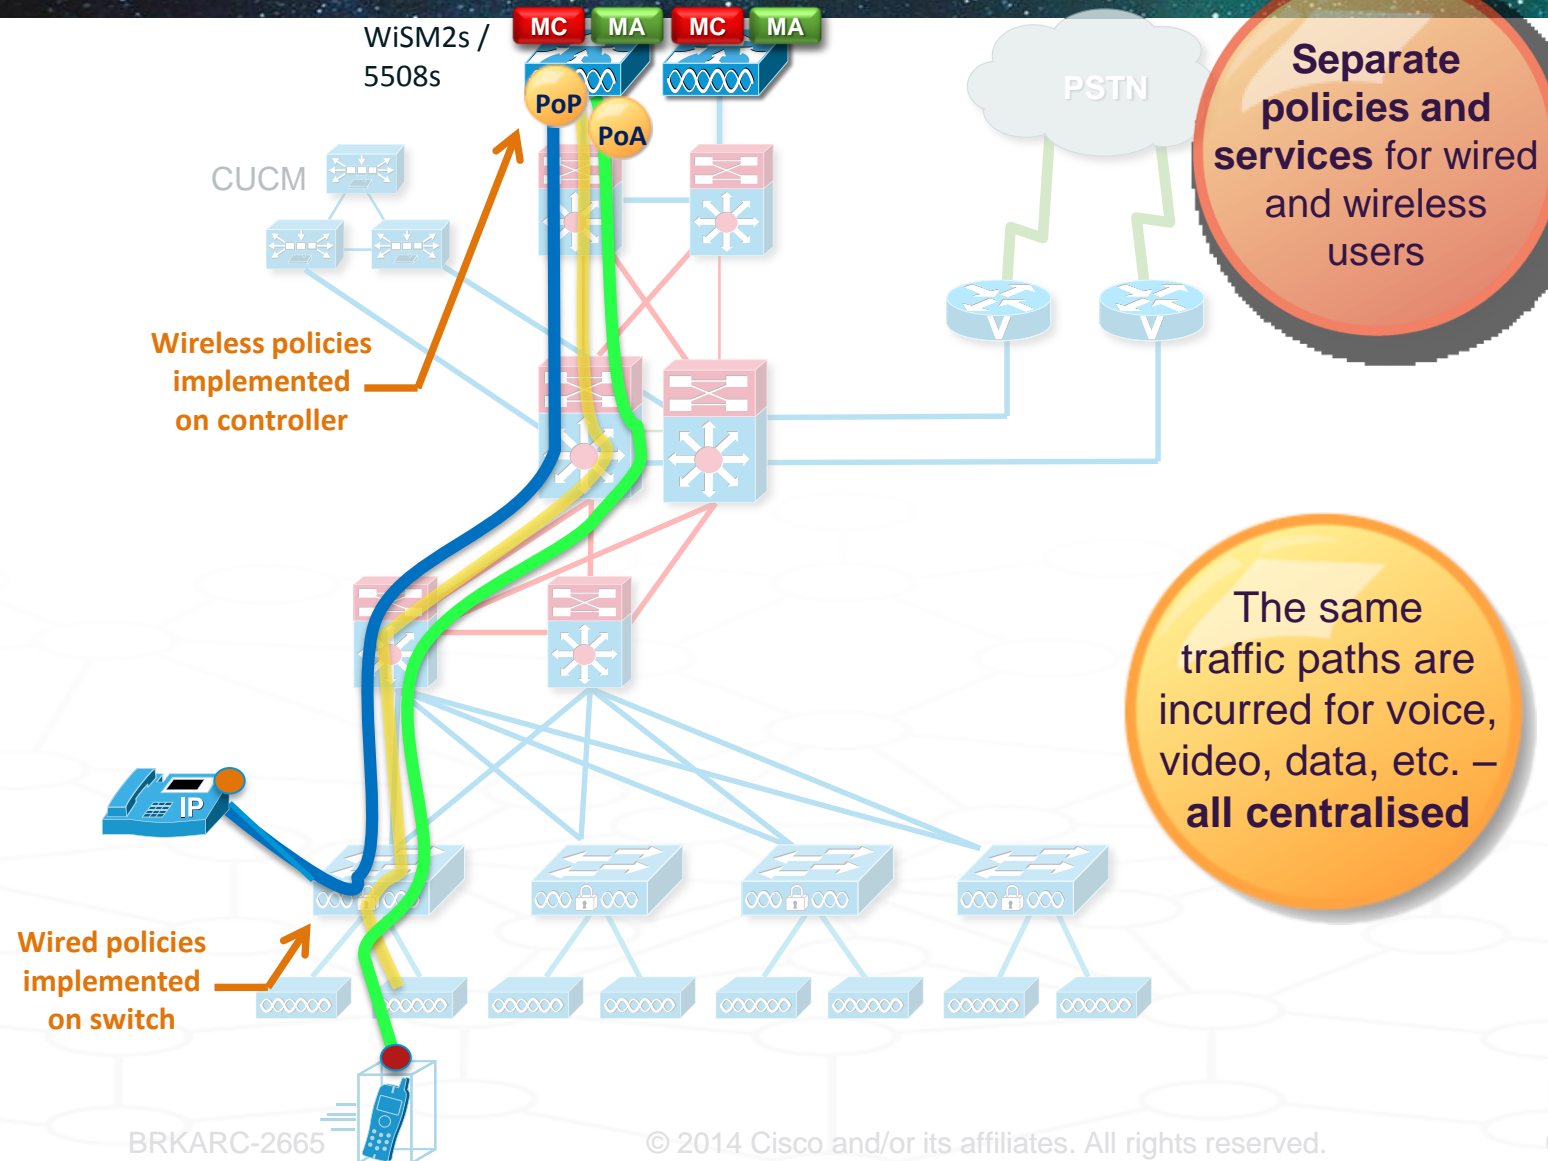
Architecture Constructs – Layer 3 Roaming (Campus Deployment)

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 – Traffic Flow



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 centralized 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

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

► Terminology and Building Blocks

Traffic Flows and Roaming

High Availability

Quality of Service

Security

Multicast

NetFlow

Converged Access Design and Deployment –

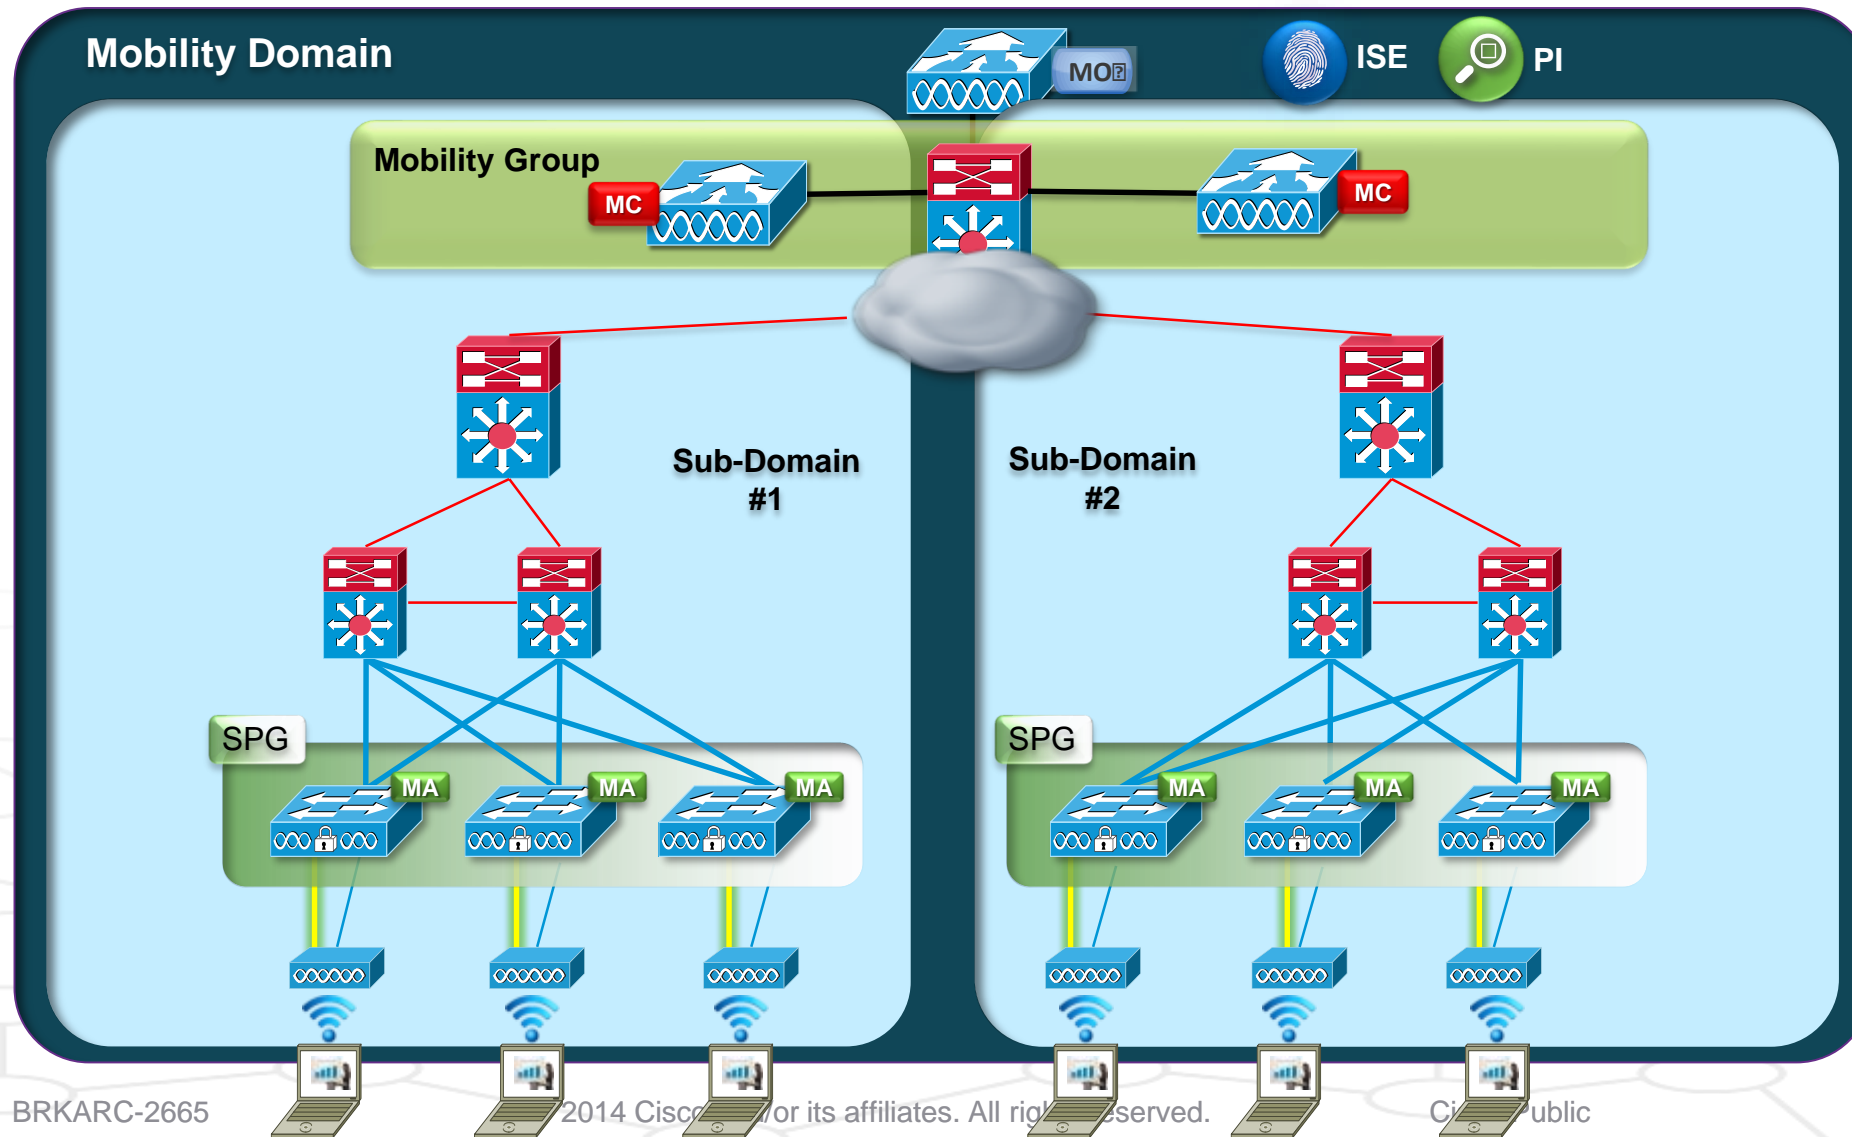
IP Addressing

Design Options

Deployment Examples

Summary

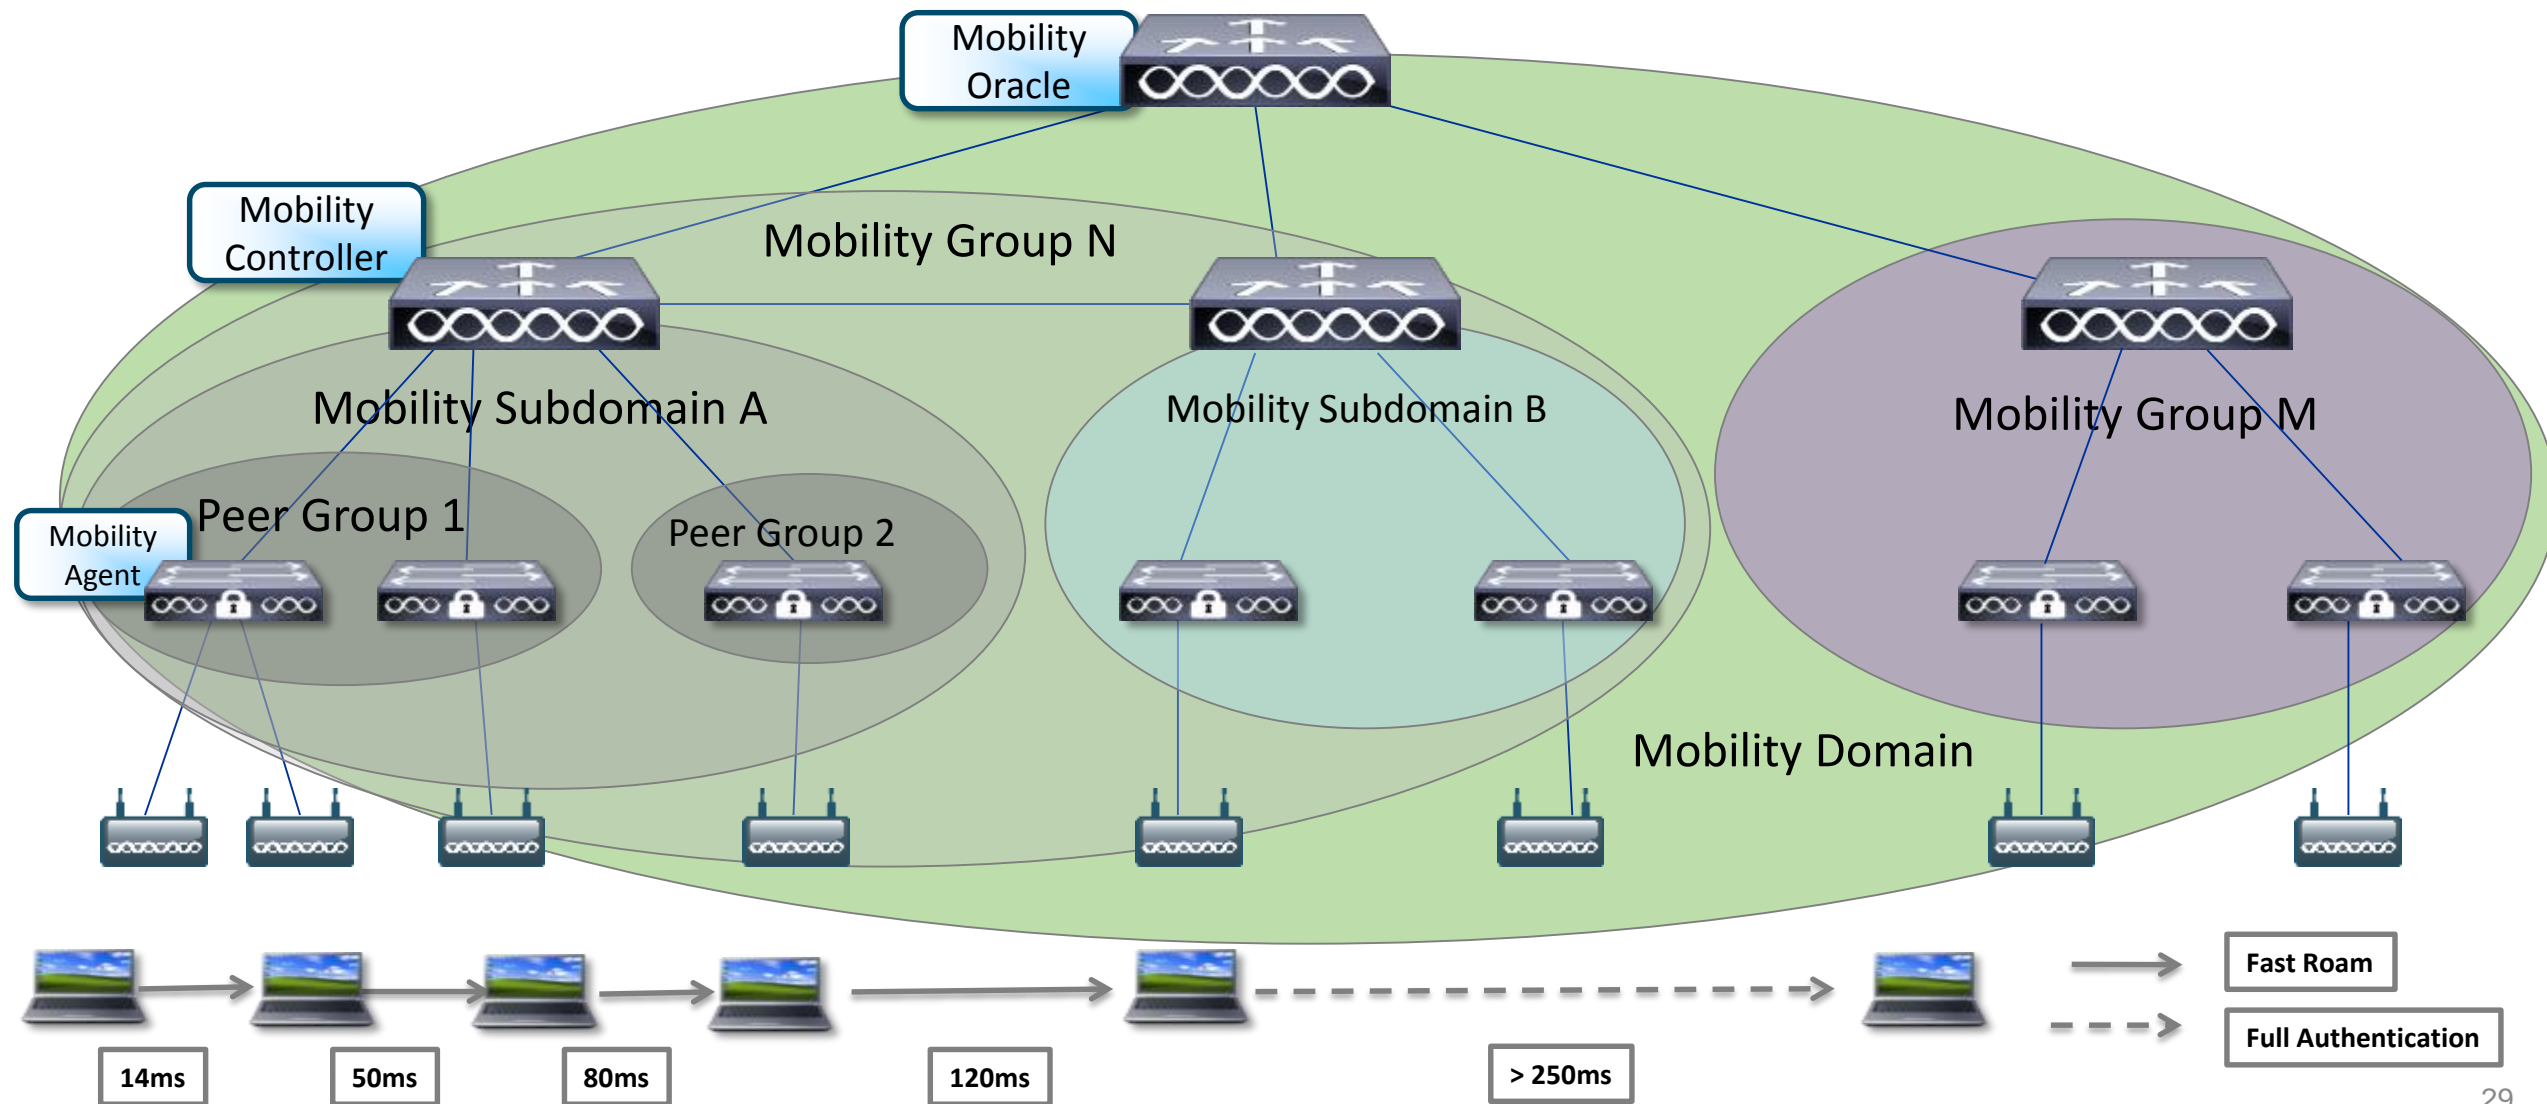
Converged Access – Deployment Overview



Converged Access – Mobility Architecture



For Your
Reference



Converged Access – Components – Physical and 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 a Distribution Block

Converged Access – Physical Entities – Catalyst 3850 / 3650 Switch Stack



Best-in-Class
Wired Switches –
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 / 3650 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 (SPGs)



- Made up of multiple Catalyst 3x50 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)

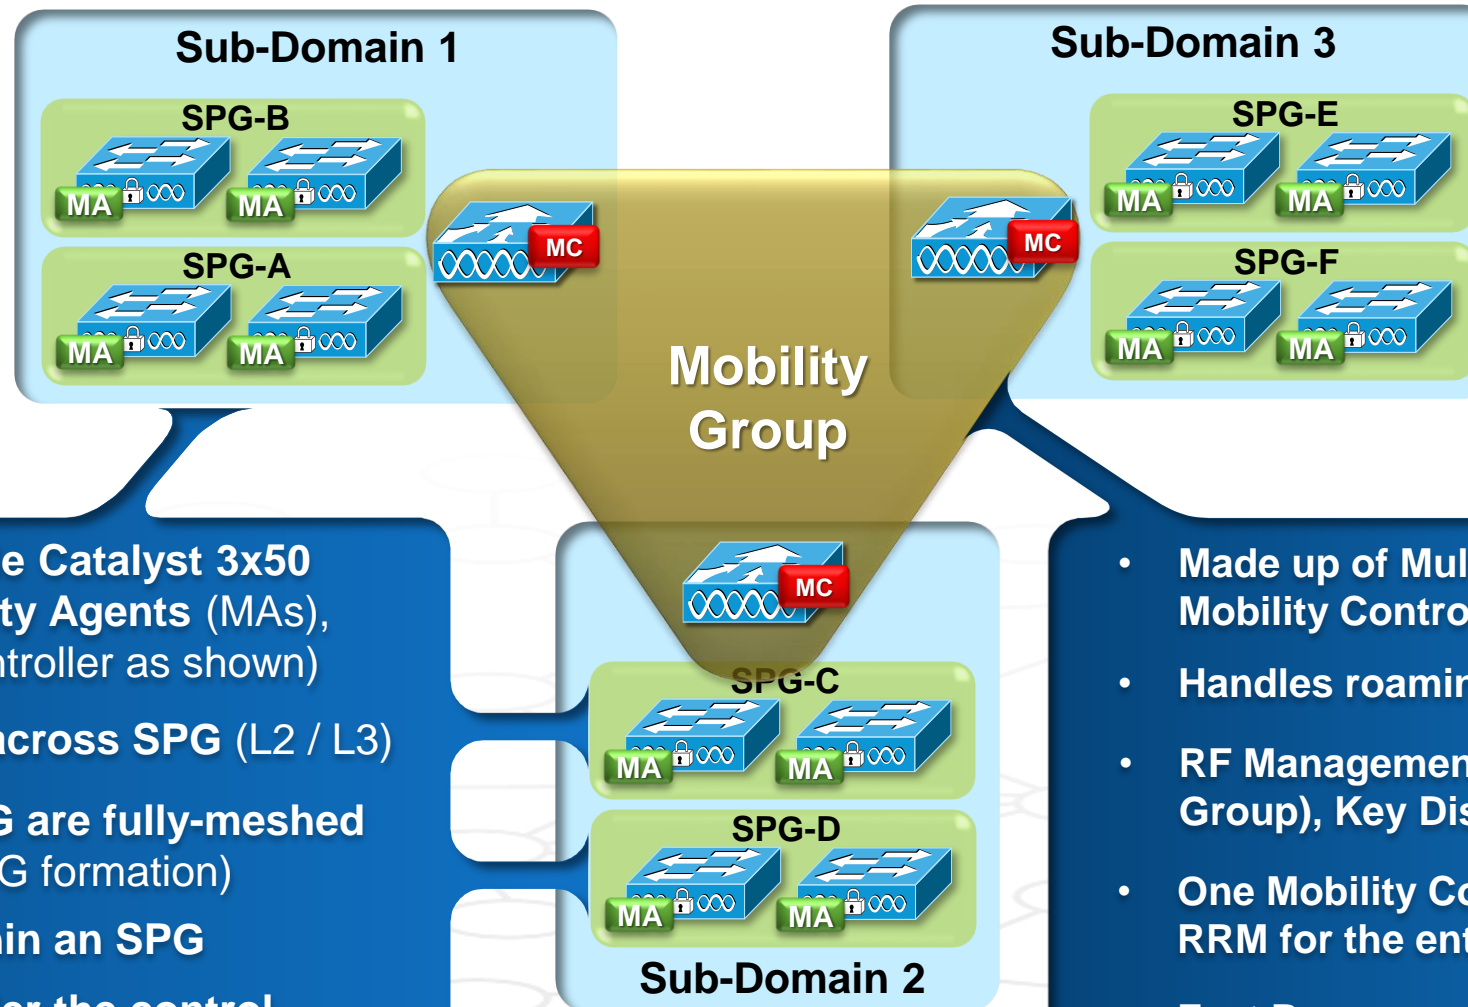
Roamed traffic between SPGs moves via the MC(s) servicing those SPGs



Hierarchical architecture is optimised for scalability and roaming

Converged Access –

Logical Entities – Switch Peer Groups and Mobility Group



- Made up of multiple Catalyst 3x50 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



For Your
Reference

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	3650 as MC (3.3.1SE)	3850 as MC (3.3.1SE)	WLC2504 (7.6)	WLC5760 (7.6)	WLC5508 (7.6)	WiSM2 (7.6)
Max APs Supported per MC	25	50	75	1000	500	1000
Max APs Supported in overall Mobility Domain	200	250	5400	72000	36000	72000
Max Clients Supported per MC	1000	2000	1000	12000	7000	15000
Max Clients Supported in overall Mobility Domain	8000	16000	72000	864000	504000	1.08M
Max number of MC in Mobility Domain	8	8	72	72	72	72
Max number of MC in Mobility Group	8	8	24	24	24	24
Max number of MAs in Sub-domain (per MC)	16	16	350	350	350	350
Max number of SPGs in Mobility Sub-Domain (per MC)	8	8	24	24	24	24
Max number of MAs in a SPG	16	16	64	64	64	64
Max number of WLANs	64	64	16	512	512	512

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

Terminology and Building Blocks

▶ **Traffic Flows and Roaming**

High Availability

Quality of Service

Security

Multicast

NetFlow

Converged Access Design and Deployment –

IP Addressing


Design Options

Deployment Examples

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 ...**

Very common roaming case

The diagram illustrates a network architecture. On the left, a 'Central Location' contains two blue circular routers connected to a grey cloud. Below the routers is the text 'CAPWAP tunnel to Guest Anchor'. To the right of the cloud is a red brick wall labeled 'DMZ'. Further right is a blue cube labeled 'Guest Anchor'. Above the Guest Anchor are three green circles containing the text 'MC', 'MA', and 'F'. To the right of these circles is a blue circle with a fingerprint icon labeled 'ISE', and a green circle with a magnifying glass icon labeled 'F'.

Roaming
across Stack
(small branch)

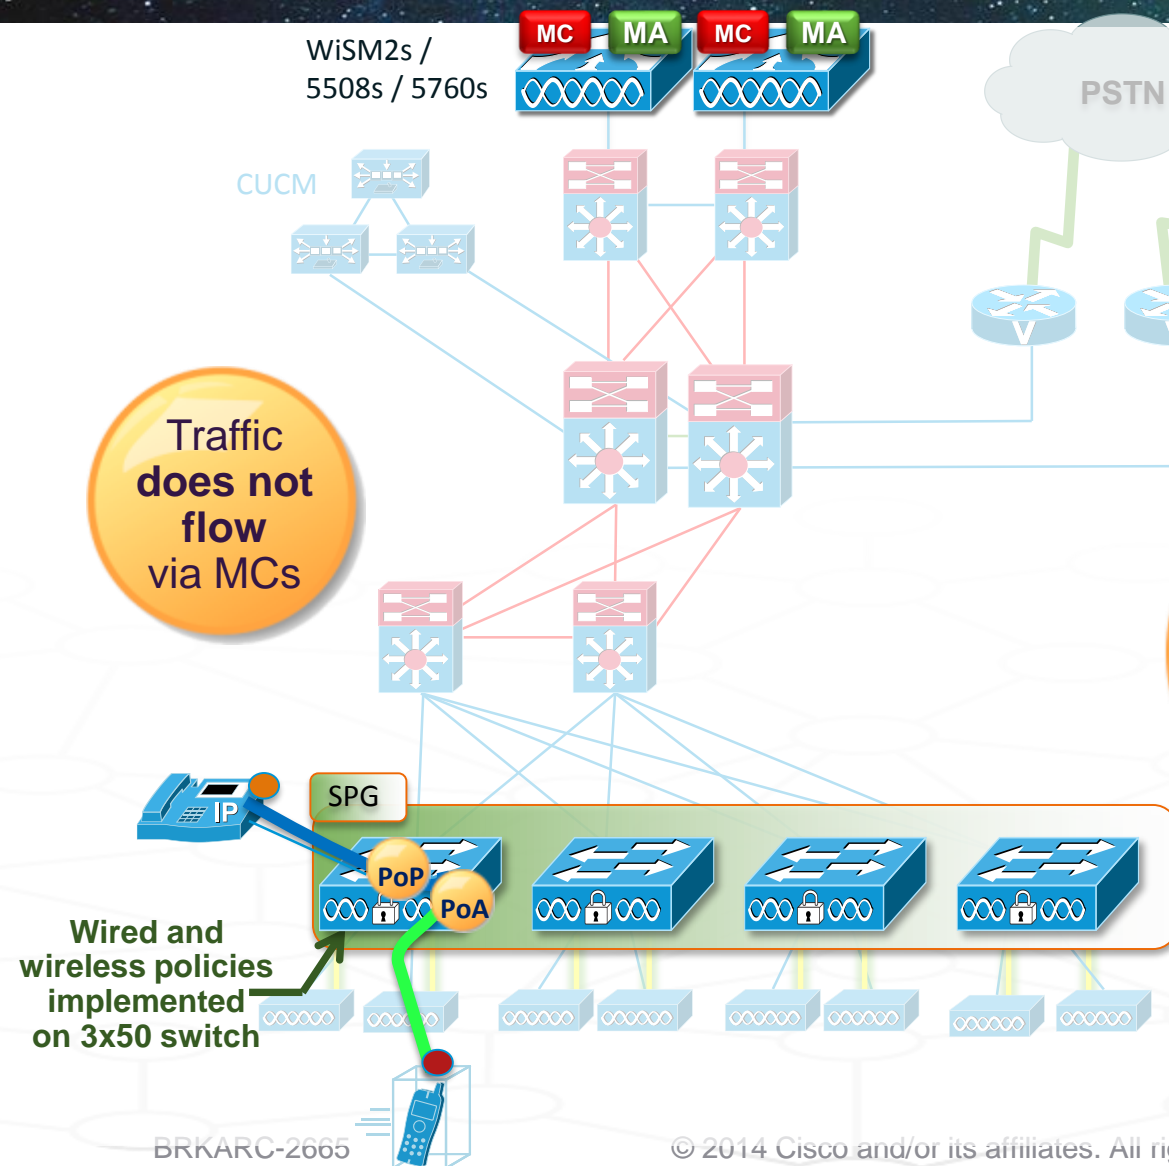


Roaming, Single Catalyst 3x50 Switch Stack –

- In this example, the user roams within their 3x50-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



Traffic Flows, Comparison (Converged Access) –

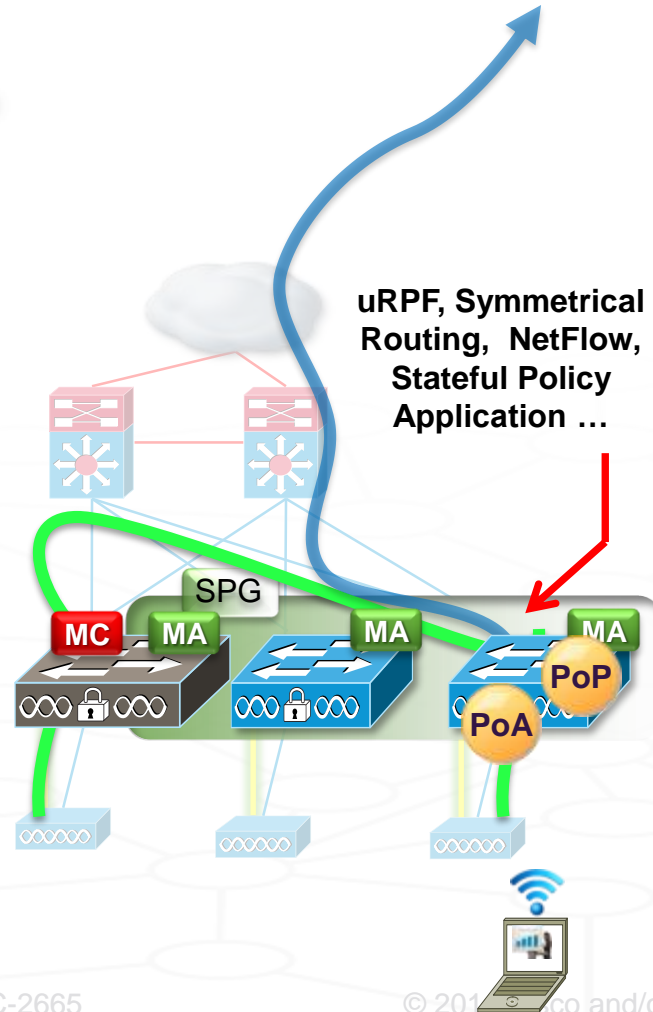
- Now, our VoIP user is on a Cisco Converged Access network, and is again making a call from a wireless handset to a wired handset ...
- We can see that all of the user's traffic is localised to their Peer Group, below the distribution layer, in both directions ...

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

Very
common
roaming
case

Roaming
across Stacks
(larger branch)



Roaming, Within a Switch Peer Group (Branch) –

- Now, let's examine a roam at a larger branch, with multiple 3x50-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 3x50 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)*

Number of Local Clients : 1

MAC Address	AP Name	WLAN State		Protocol
001e.65b7.7d1a	LD9-AP1142-1	2	UP	11n(5)

Total Number of Wireless Clients = 1

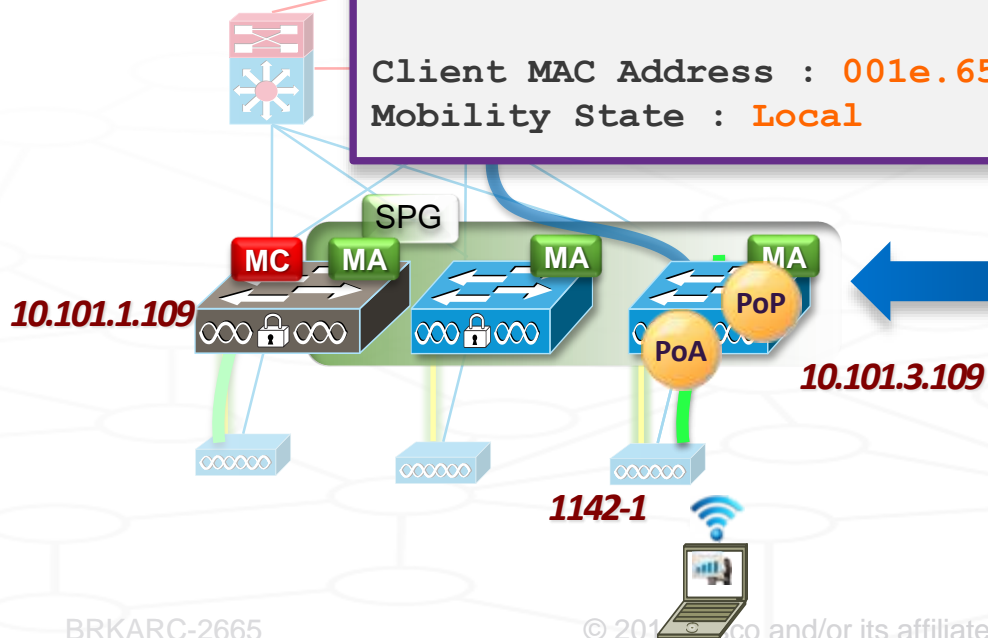
Local Clients = 1

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 Access

Traffic Flow and

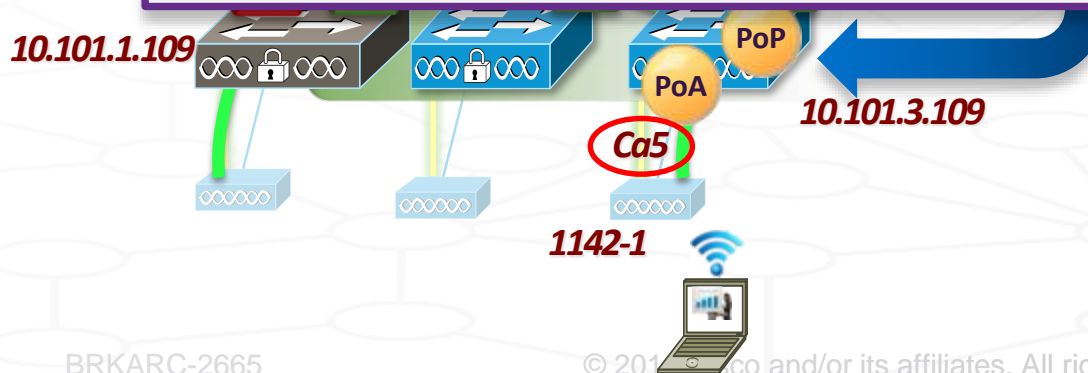
L09-3850s-3# **show capwap summary**

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

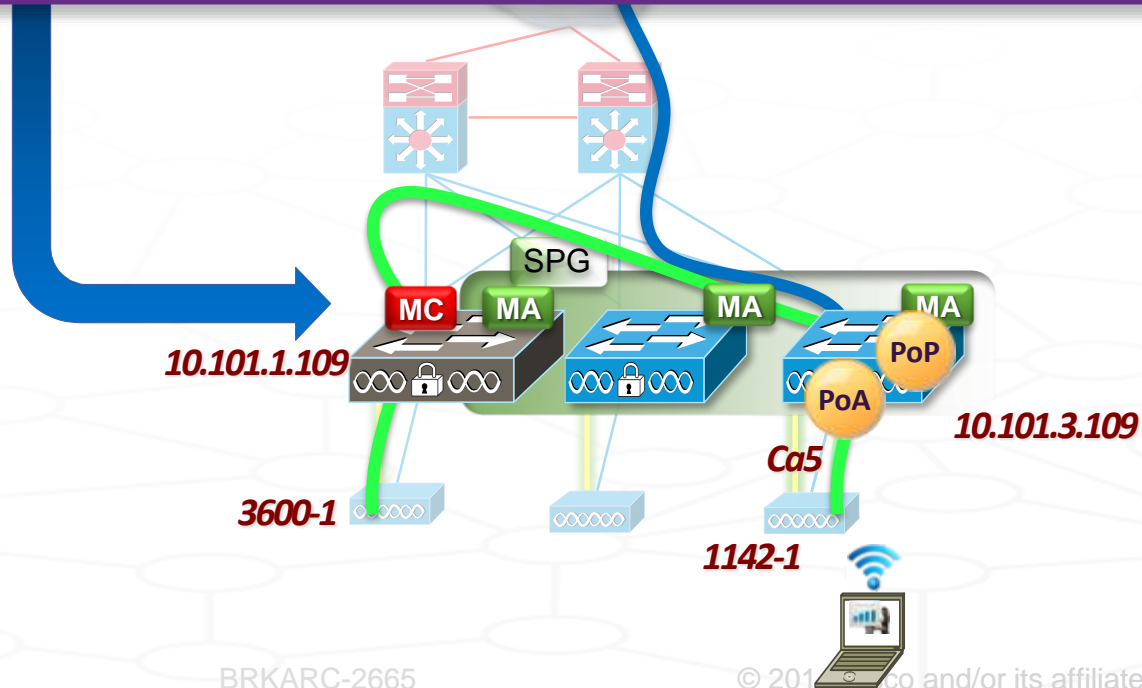
L09-3850s-3# **show ip dhcp snooping binding**

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

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




```
...
Client MAC Address : 001e.65b7.7d1a
Mobility State : Foreign
Mobility Anchor IP Address : 10.101.3.109
```

The diagram illustrates a network topology for a multi-homed PoP (PoA). The PoA is connected to two upstream providers, Ca3 and Ca5, via SPG (Service Provider Gateway) and MA (Multi-Access) components. The PoA is labeled with IP address 10.101.1.109 and interface 3600-1. The upstream providers are labeled with IP addresses 10.101.3.109 and interface 1142-1. A laptop is shown connected to the PoA. The diagram also shows a large blue arrow pointing towards the PoA, and a red circle highlighting the Ca3 label.

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

Converged Traffic Flow

```
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
```

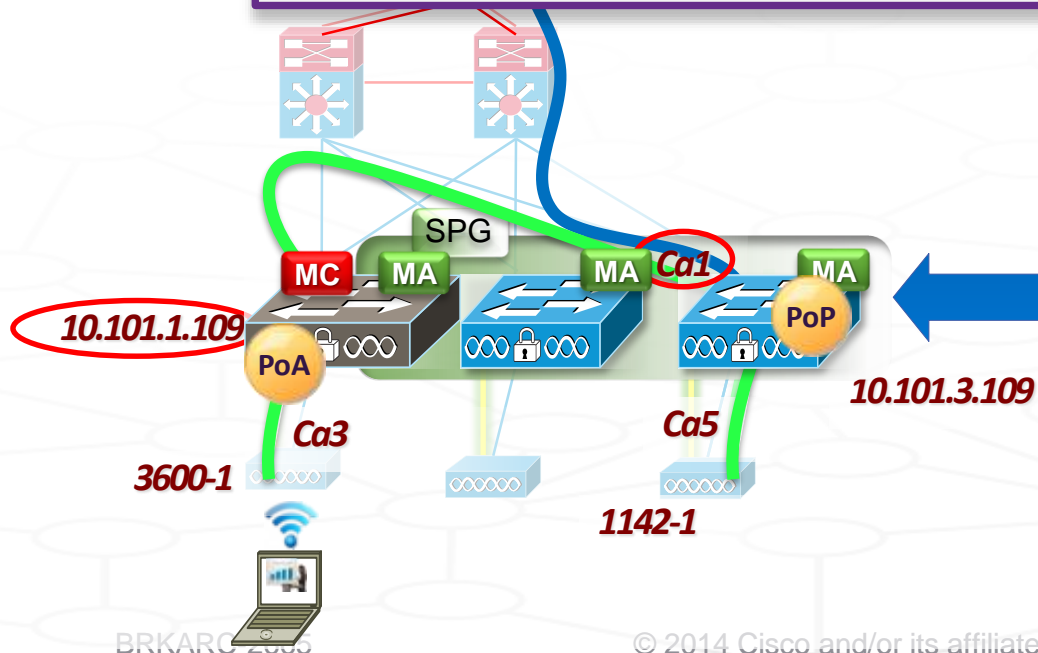


MacAddress	IpAddress	Lease (sec)	Type	VLAN	Interface
00:1E:65:b7:7d:1a	10.101.203.1	10720	dhcp-snooping	2003	Capwap1

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 –

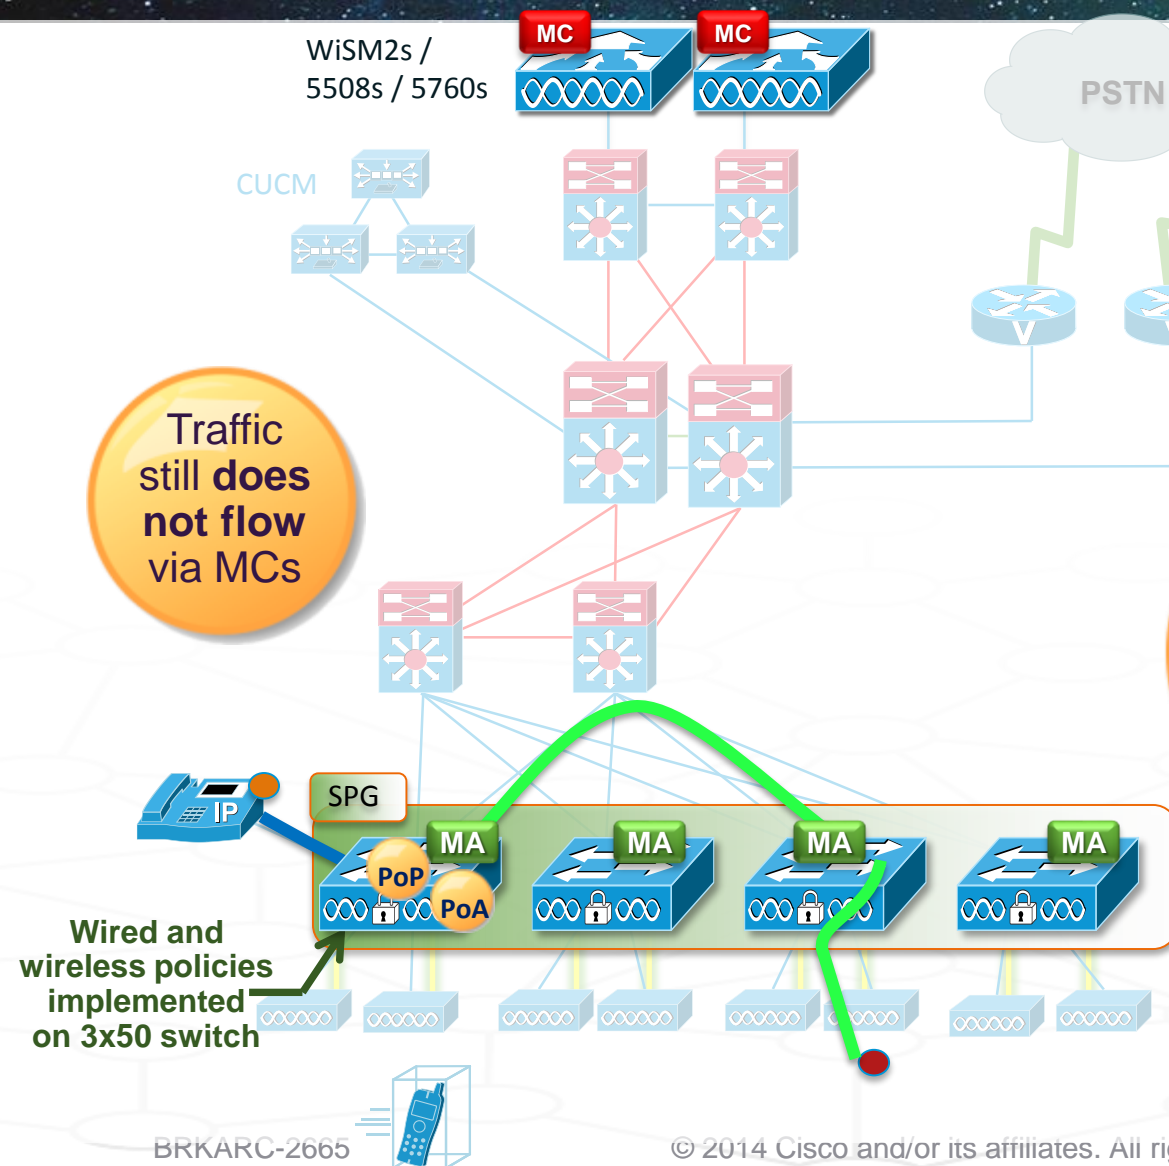
This looks exactly the same as a Layer 3 inter-controller roam in CUWN ...
because it is exactly the same process –
Just distributed, rather than centralized ...

Converged Access –

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



Converged Access – Traffic Flow – with Intra-SPG Roam



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 SPGs)




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



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


50

Converged Access –

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

As Noted –

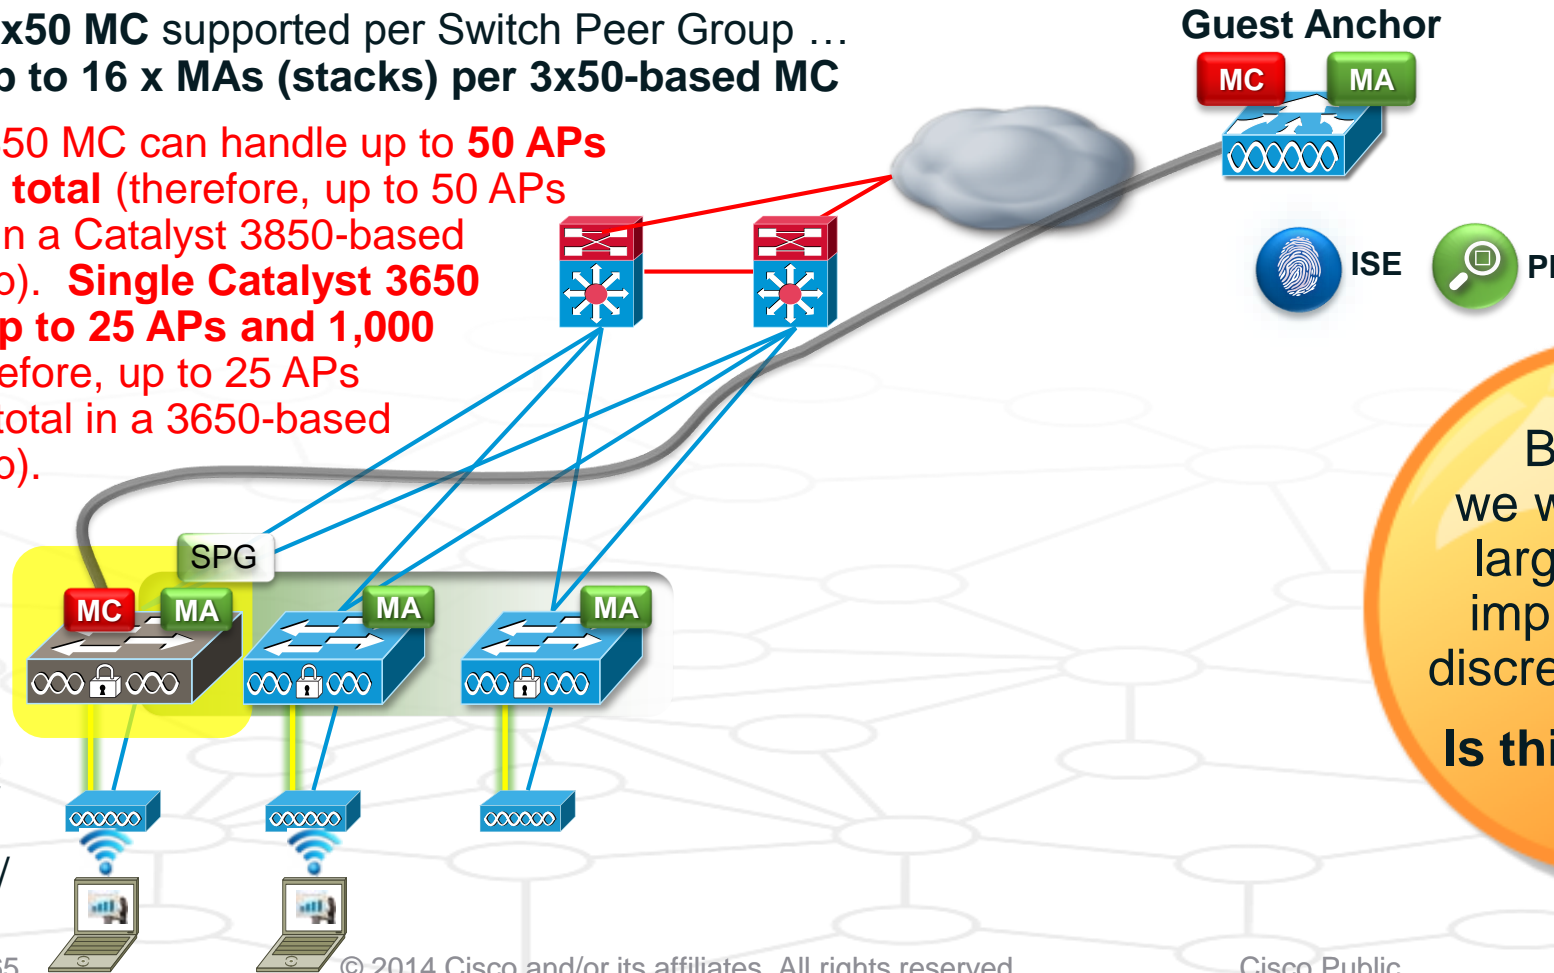
- When a user roams in a L2 environment, an optional setting allows for both the user's PoA and PoP to move.
- The benefits that accrue to a PoP move for an L2 user roam are **reduced end-to-end latency** for the user (less traffic hops), as well as a **reduction of state held within the network** (as 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 3x50-based MCs – Functionality

As we saw previously, we can also optionally use a Catalyst 3x50 switch as an MC + co-located MA for a Switch Peer Group ... let's explore this in more detail –

- Single Catalyst 3x50 MC supported per Switch Peer Group ...
- which can have up to 16 x MAs (stacks) per 3x50-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). **Single Catalyst 3650 MC can handle up to 25 APs and 1,000 clients total** (therefore, up to 25 APs and 1,000 clients total in a 3650-based Switch Peer Group).
- MC handles inter-SPG roaming, RRM, Guest Access, etc.
- More scalable MC capability can be provided by 5760 / 5508 / WiSM2



But what if we want to scale larger, **without** implementing a discrete controller?
Is this possible?

Converged Access – Catalyst 3x50-based MCs – Scaling

Switch Peer Group / Mobility Group Scaling with Catalyst 3x50 –


- Up to 8 x Catalyst 3x50 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 3850s
- Up to 200 APs total and 8,000 clients supported (maximum) across a Mobility Group made up solely of Catalyst 3650s

Guest Anchor

Guest tunnelling is per MC –
to Guest Anchor controller

Licensing
is per MC
not pooled
across MCs

RRM, etc. is
coordinated
across the
MCs in the
same Mobility
Group



Converged Access – Catalyst 3x50-based MCs – Roaming

There are multiple roaming scenarios with Catalyst 3x50-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 3x50 switches operating as MCs
 - Catalyst 3x50-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 – Roaming

Roaming, within a Stack (3850 Switches as MCs)

- Initially, all clients in this example are on their initial, local PoP
- 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

Converged Access – Catalyst 3850-based MCs – Roaming

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

- Now, the client roams to an AP serviced by another switch
- 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)

Converged Access – Catalyst 3850-based MCs – Roaming

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

- Now, let's examine a more complex roam where the user moves to a separate SPG, onto the PoP that is serving as MC for that SPG
- In this example – the user roams to a separate SPG, onto the PoP that is serving as MC for that SPG
- The user 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)

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

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

- Now, let's 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


Converged Access – Catalyst 3x50-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 3x50 switches as MAs, this likely provides the most scalable design option for a larger network build.
- **However, if using 3x50 switches as MCs for smaller builds – and with the scaling limits detailed previously in mind – we can determine where to best use this**
- **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 3x50 switches.
- **Cons –**
 - **OpEx complexity** – due to some additional complexity that comes into roaming situations when using multiple 3x50 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.

Conclusion –

In smaller designs (such as branches), the use of Catalyst 3x50 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.



**Roaming
details
provided on
Reference
slides**

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 3x50-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.

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- ▶ **High Availability**

- Quality of Service

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –

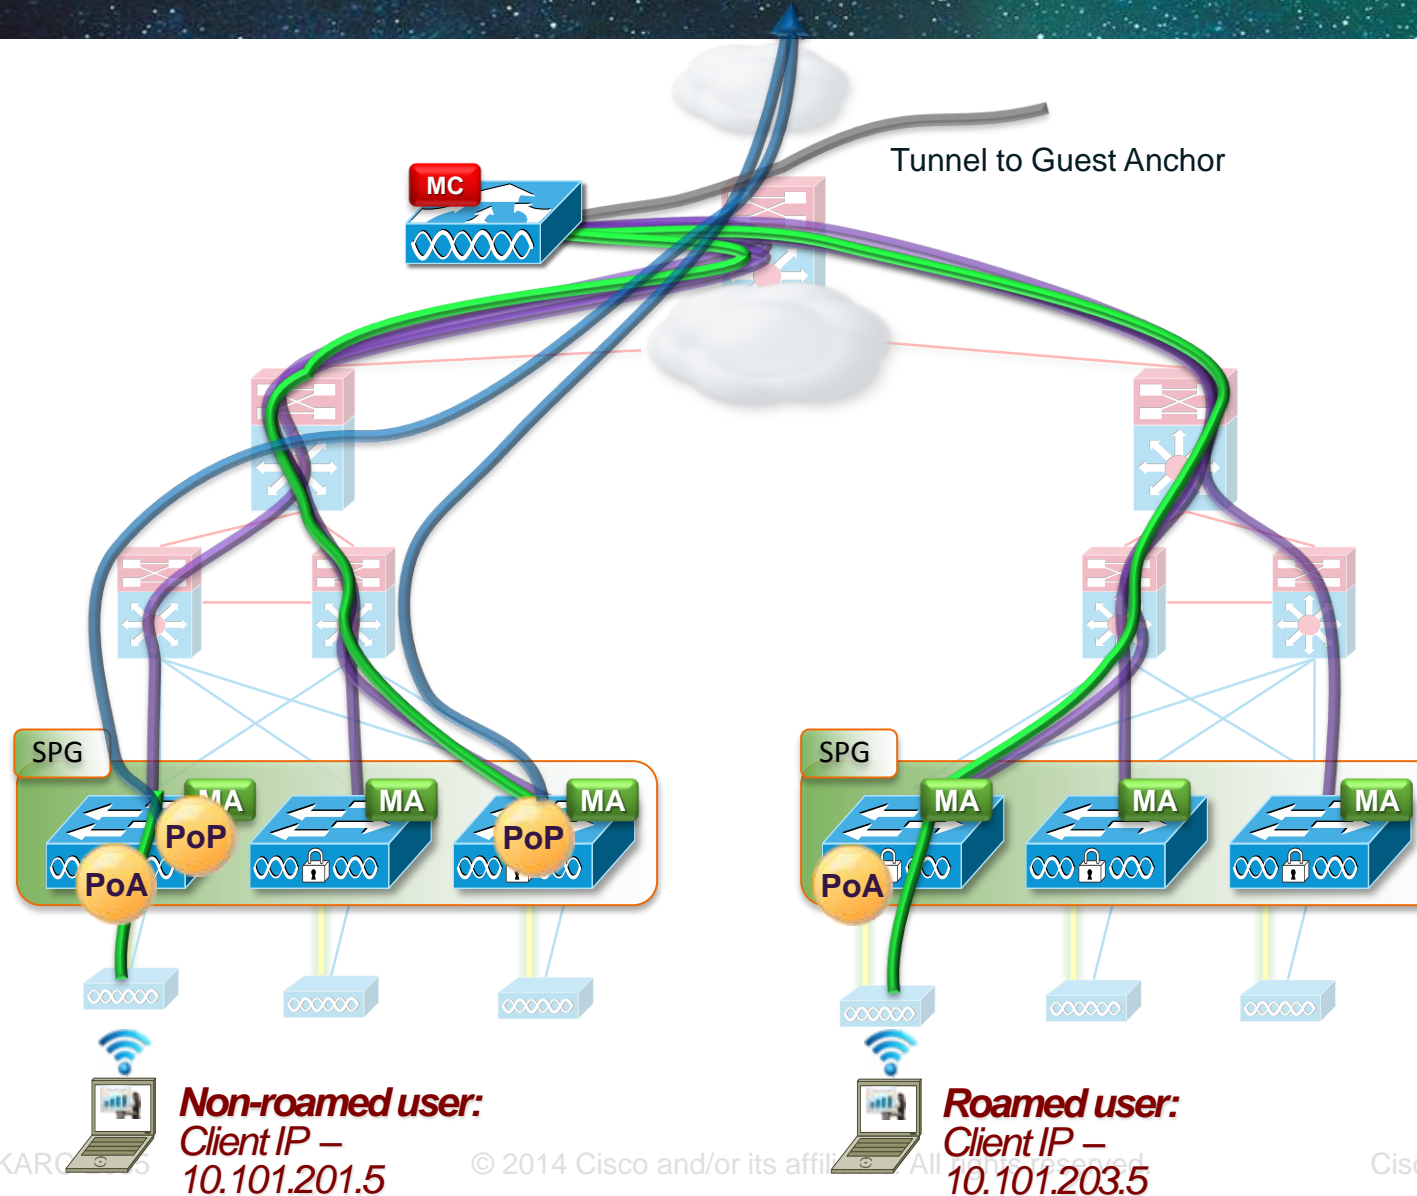
- IP Addressing

- Design Options

- Deployment Examples

Summary

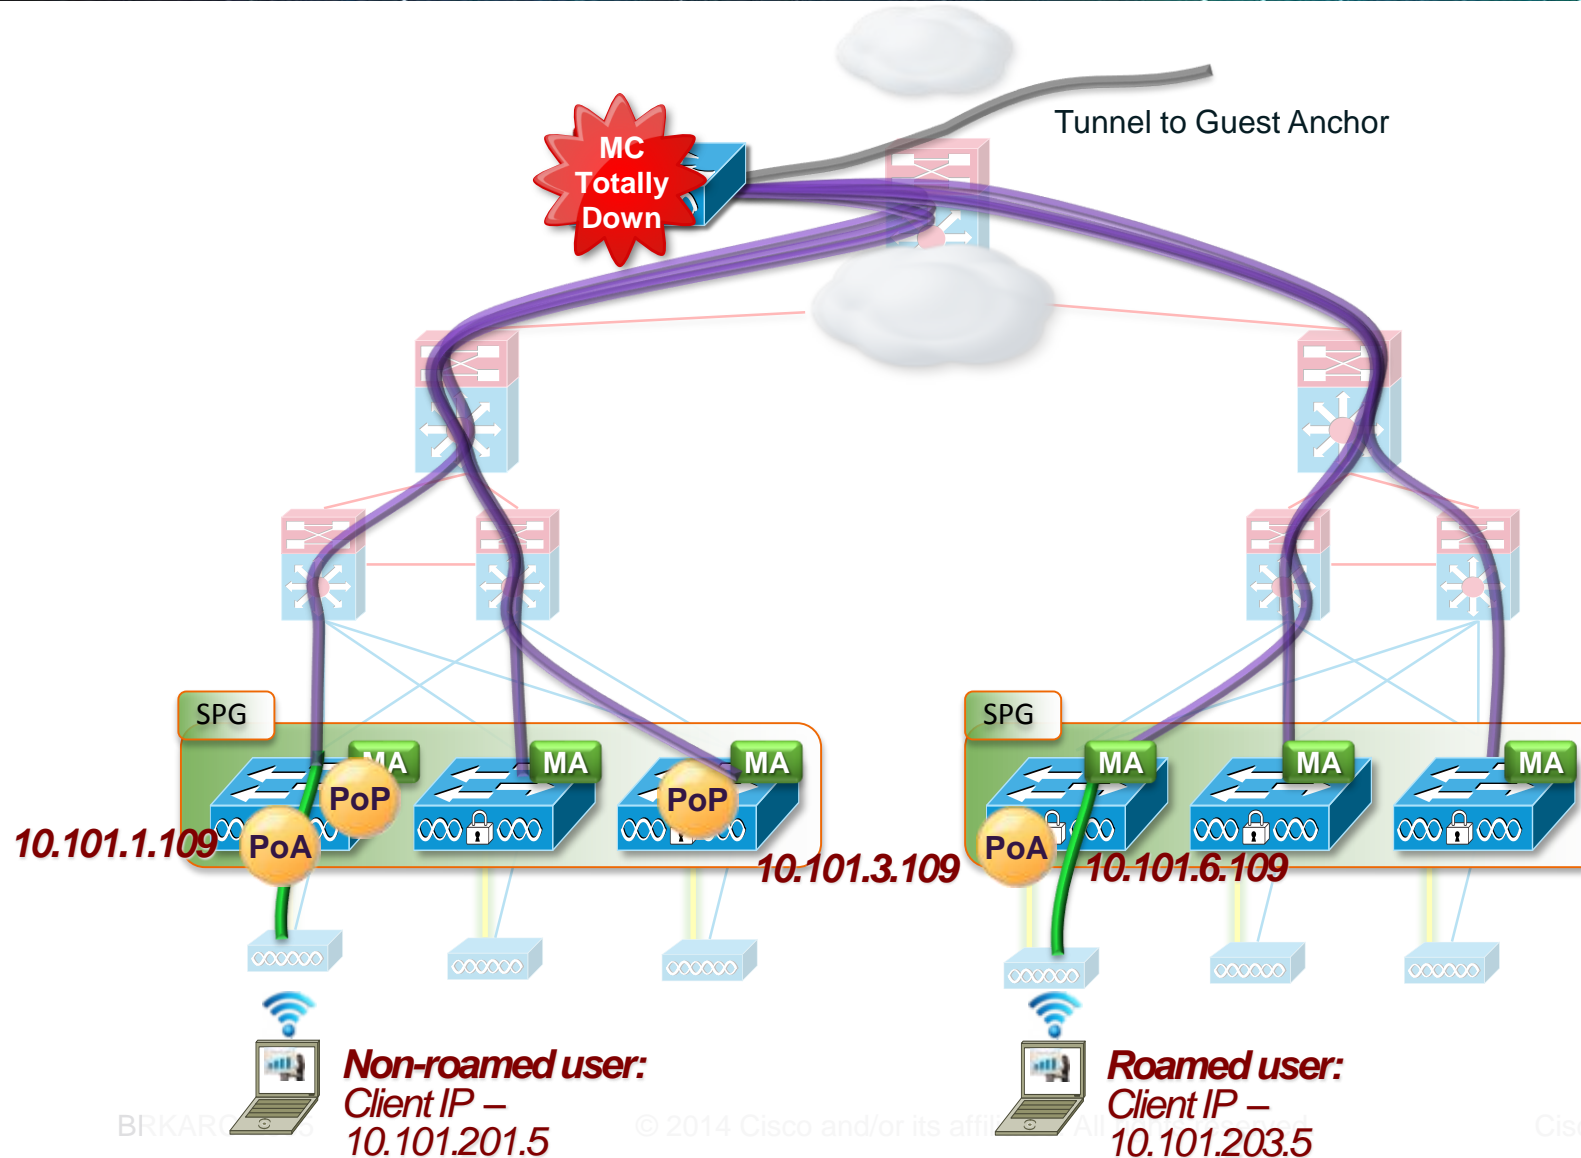
High Availability – State Held within the Network – for Local and Roamed Users



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

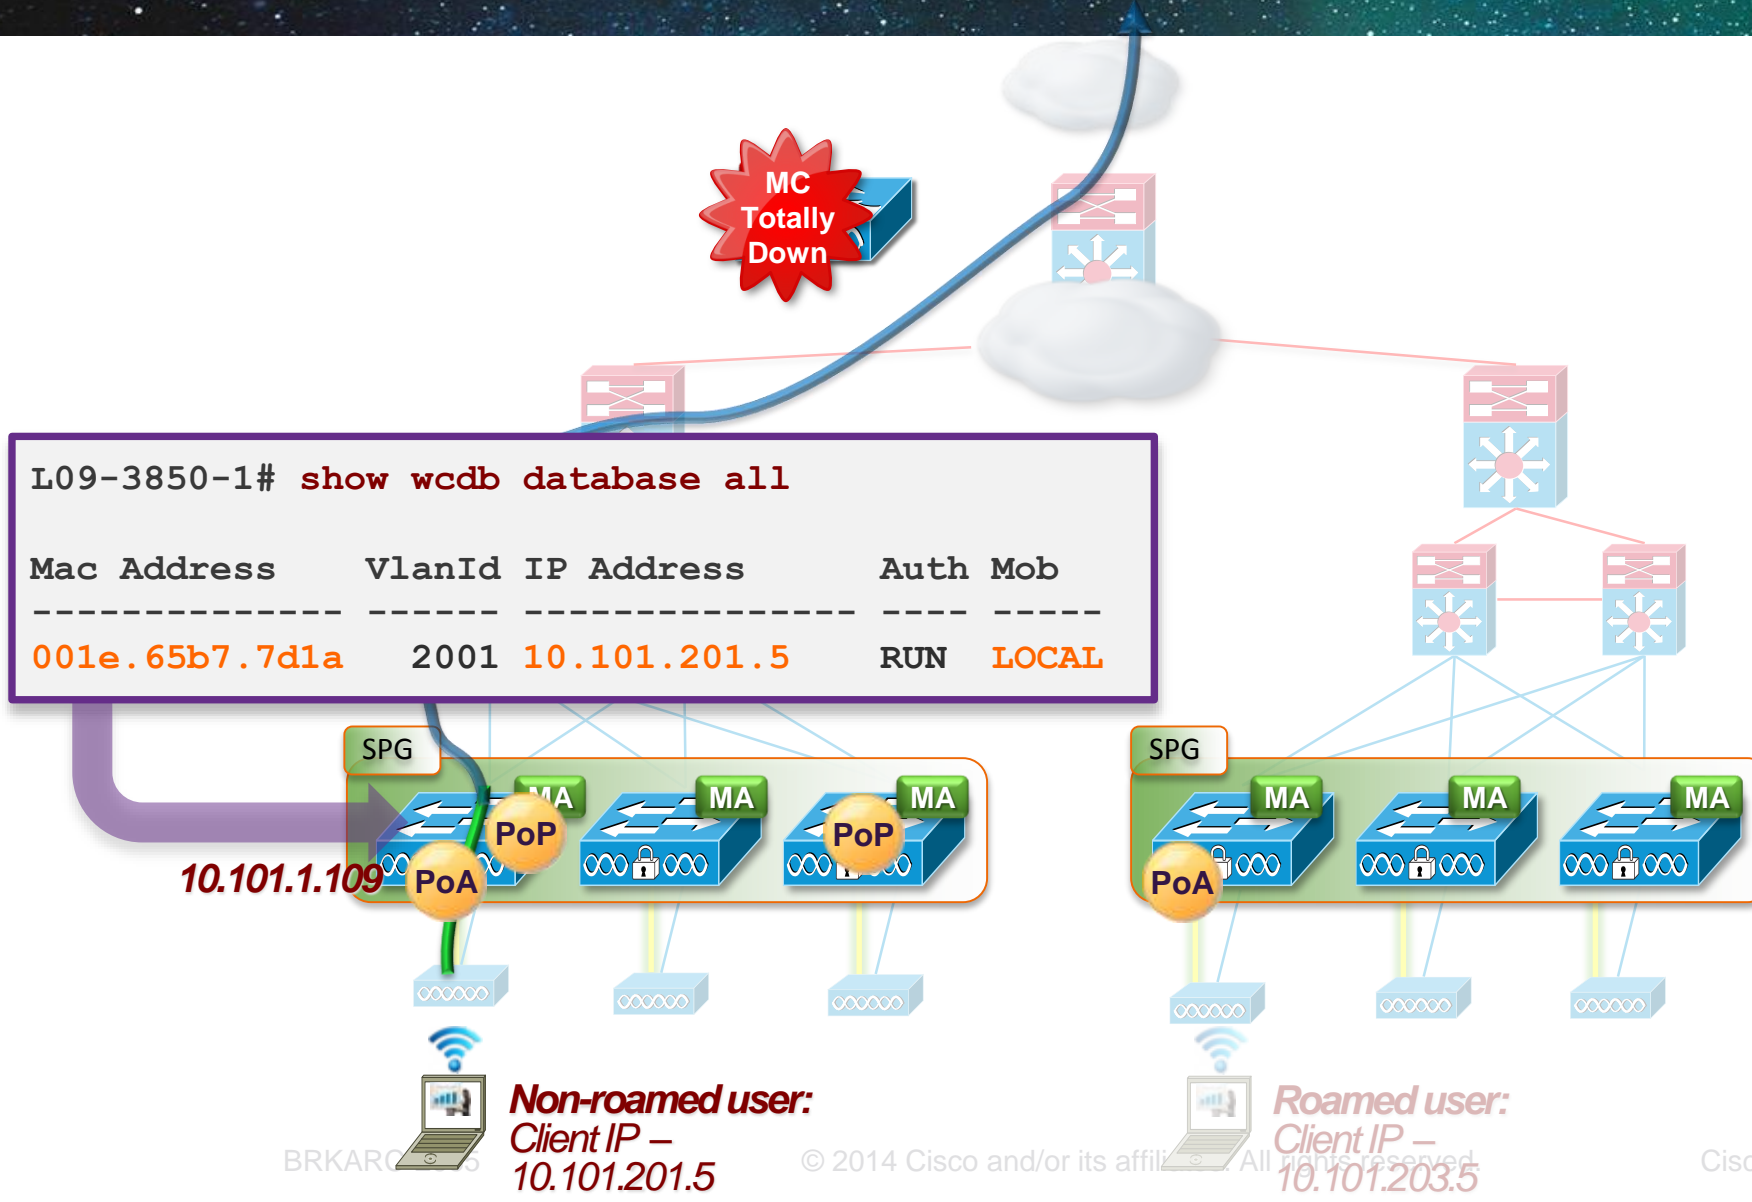


Roamed and Local users, High Availability Considerations –

- 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 ...

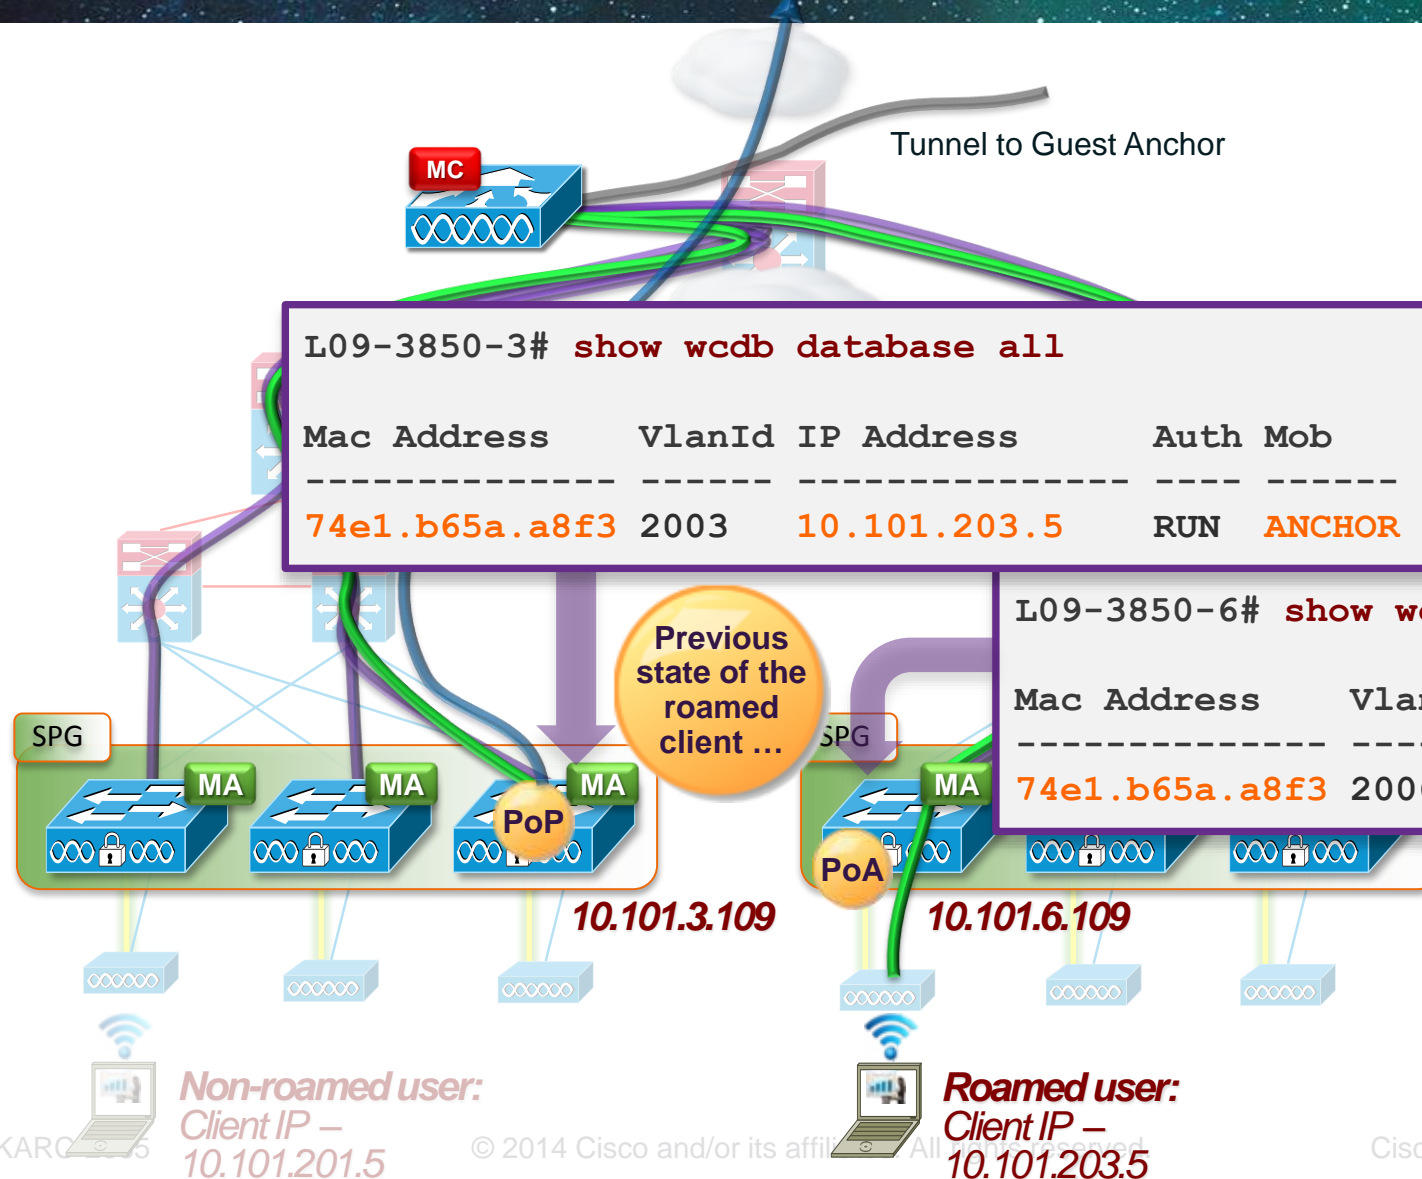
High Availability – MC Failure – Effect on Local (Non-Roamed) Clients



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

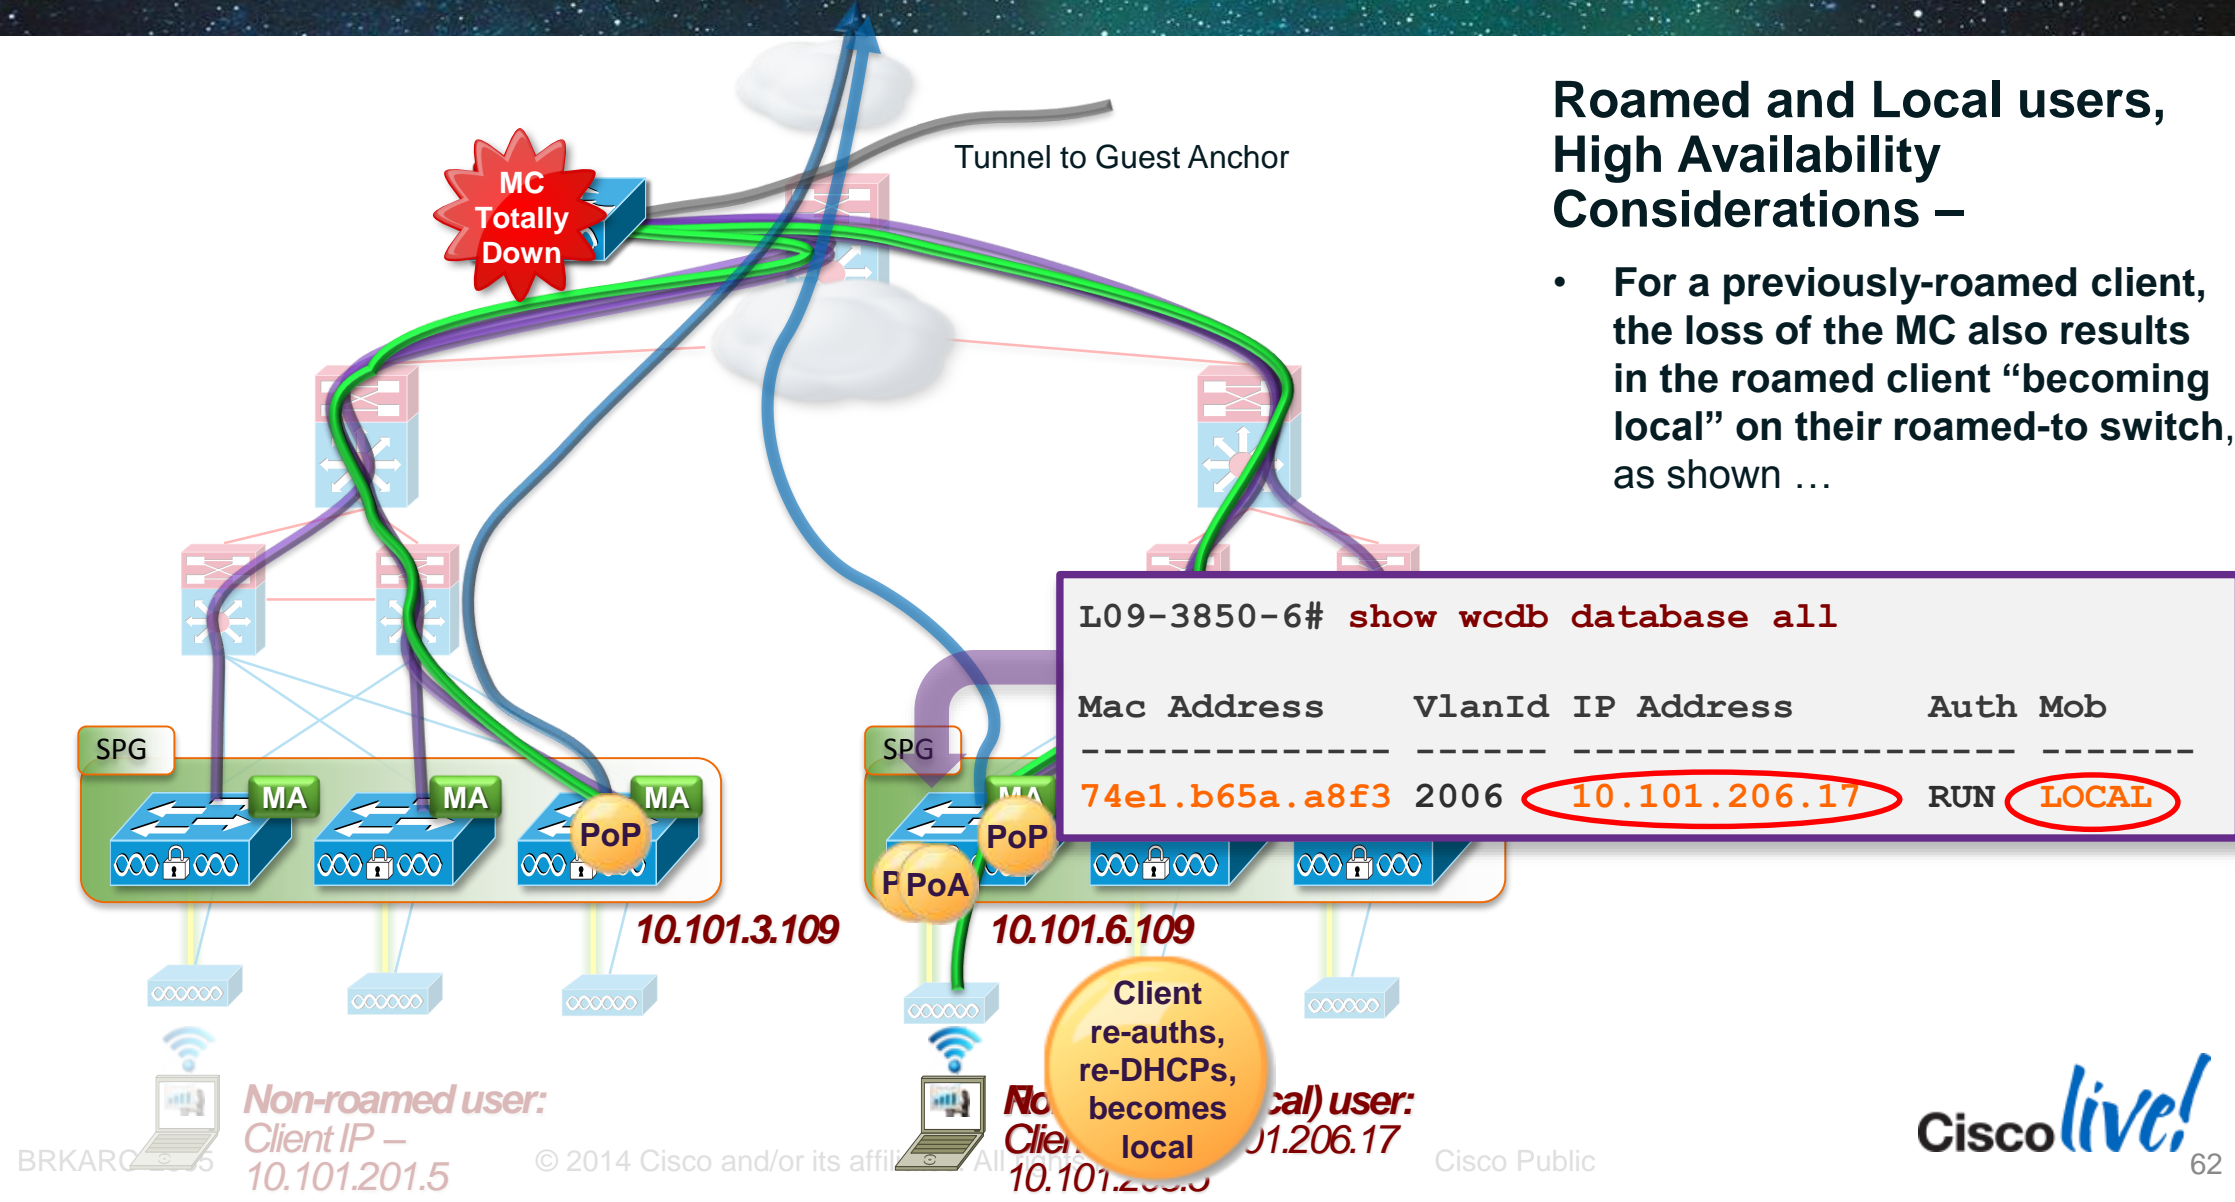
High Availability – MC Failure – Effect on Previously-Roamed Clients



Roamed and Local users, High Availability Considerations –

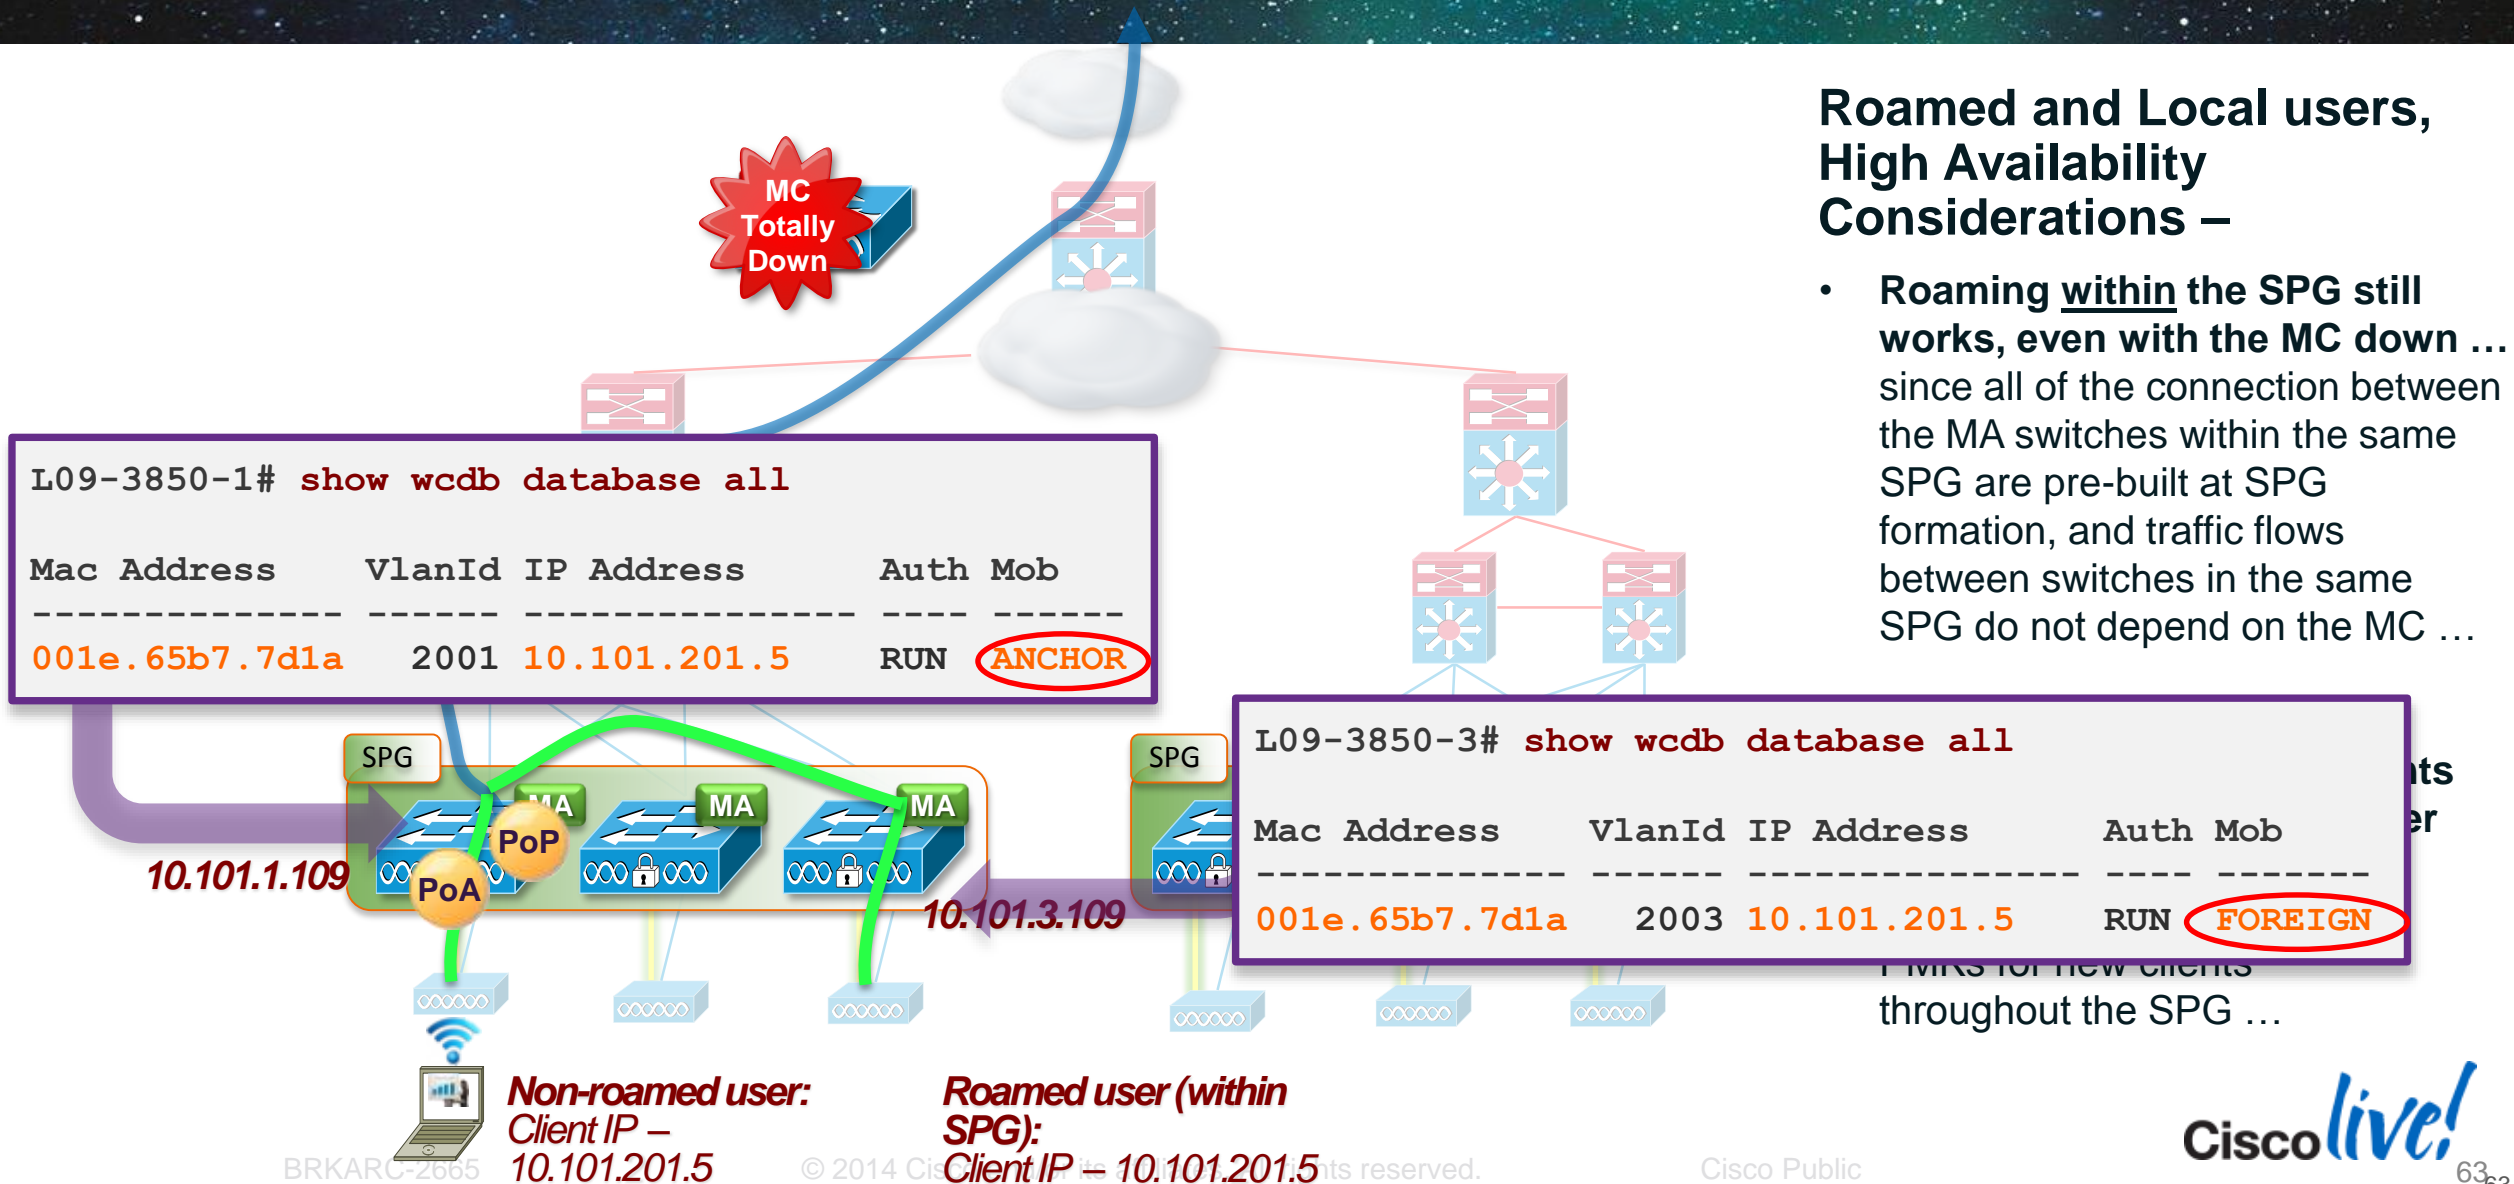
- Here is a client who has roamed from the 10.101.3.109 switch, to the 10.101.6.109 switch, as shown ...

High Availability – MC Failure – Effect on Previously-Roamed Clients



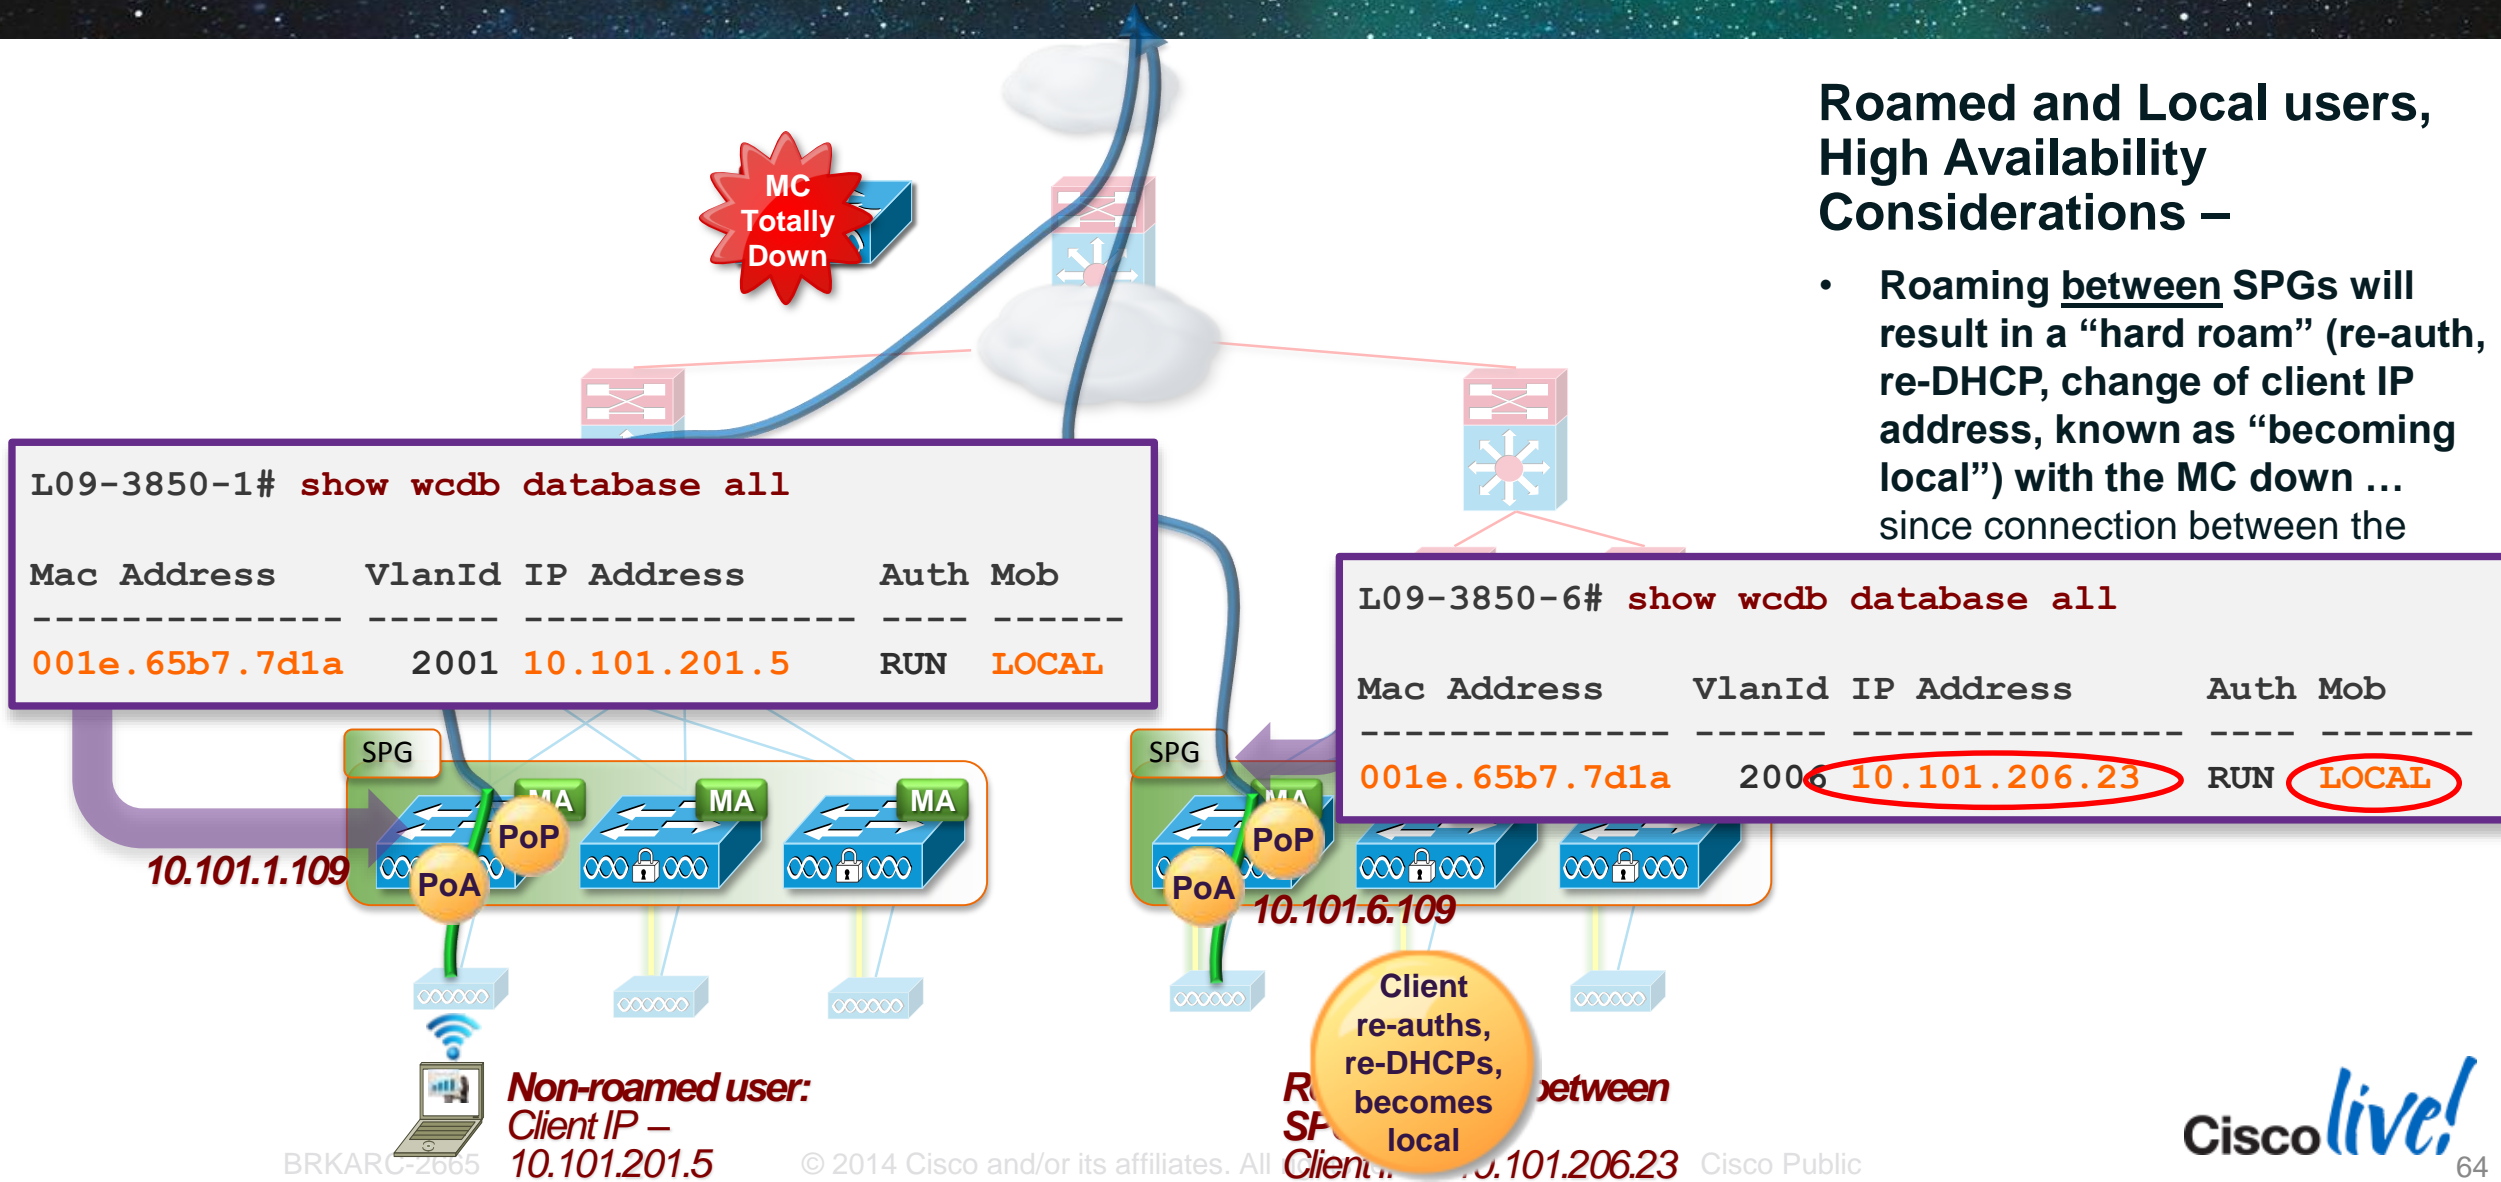
High Availability –

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



High Availability –


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



Roamed and Local users, High Availability Considerations –

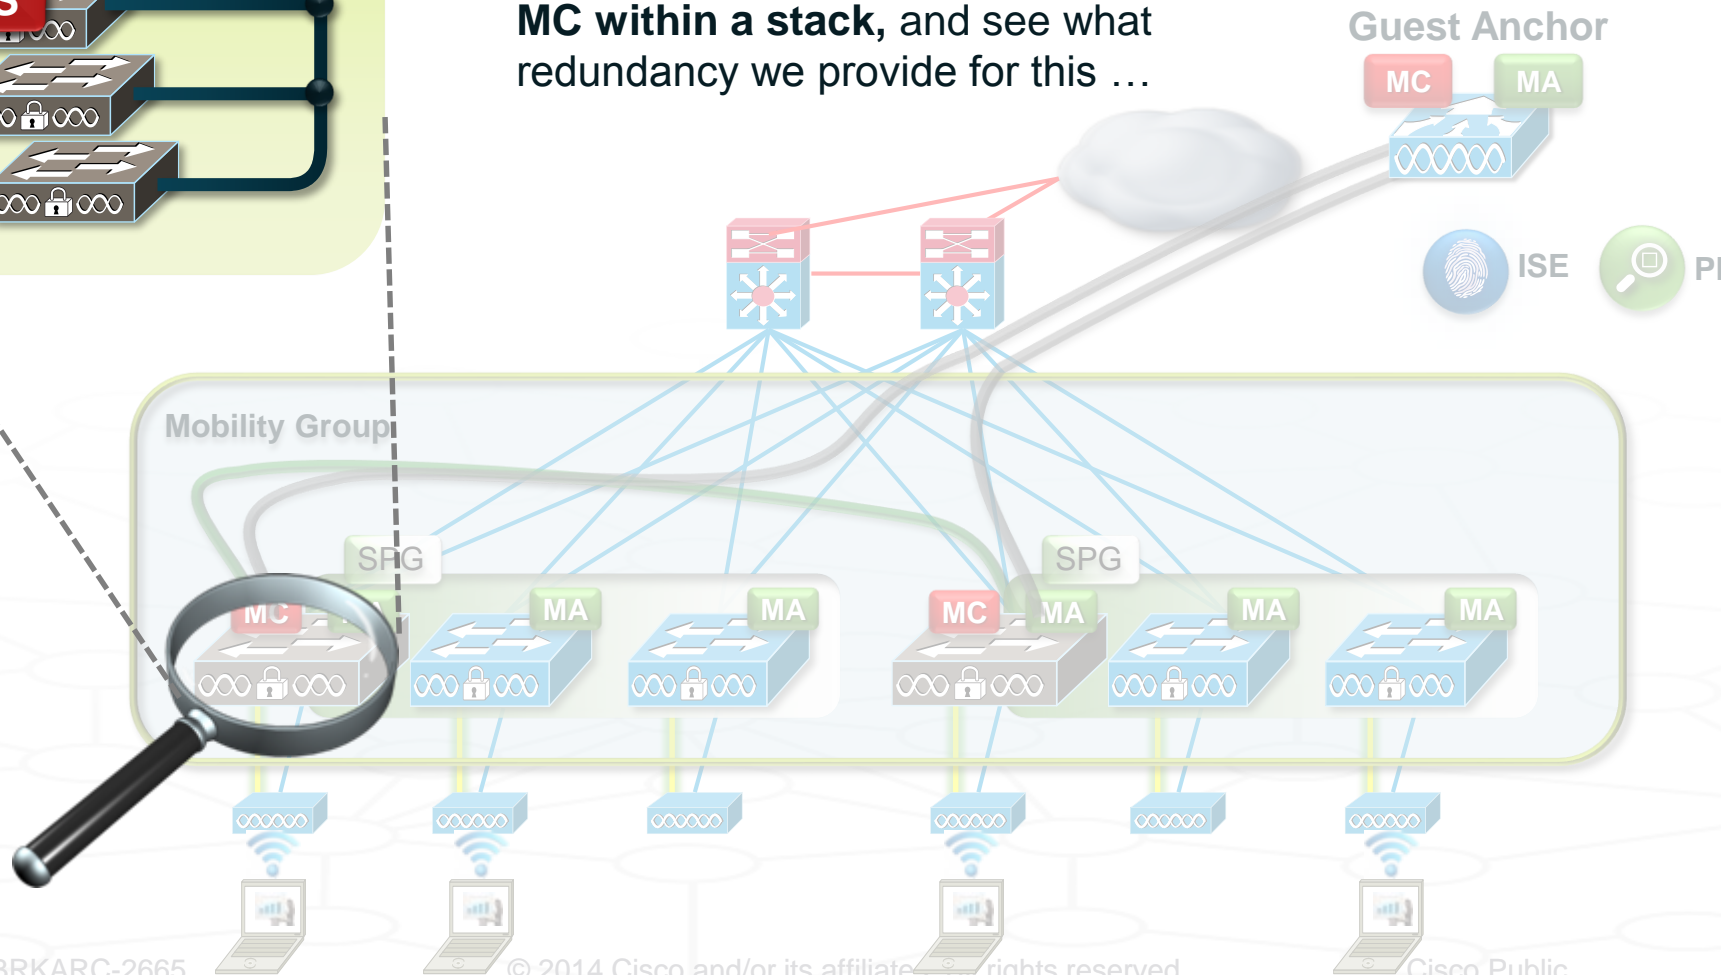
- Roaming between SPGs will result in a “hard roam” (re-auth, re-DHCP, change of client IP address, known as “becoming local”) with the MC down ... since connection between the

High Availability – Catalyst 3x50-based MCs – Fault Tolerance in Stack

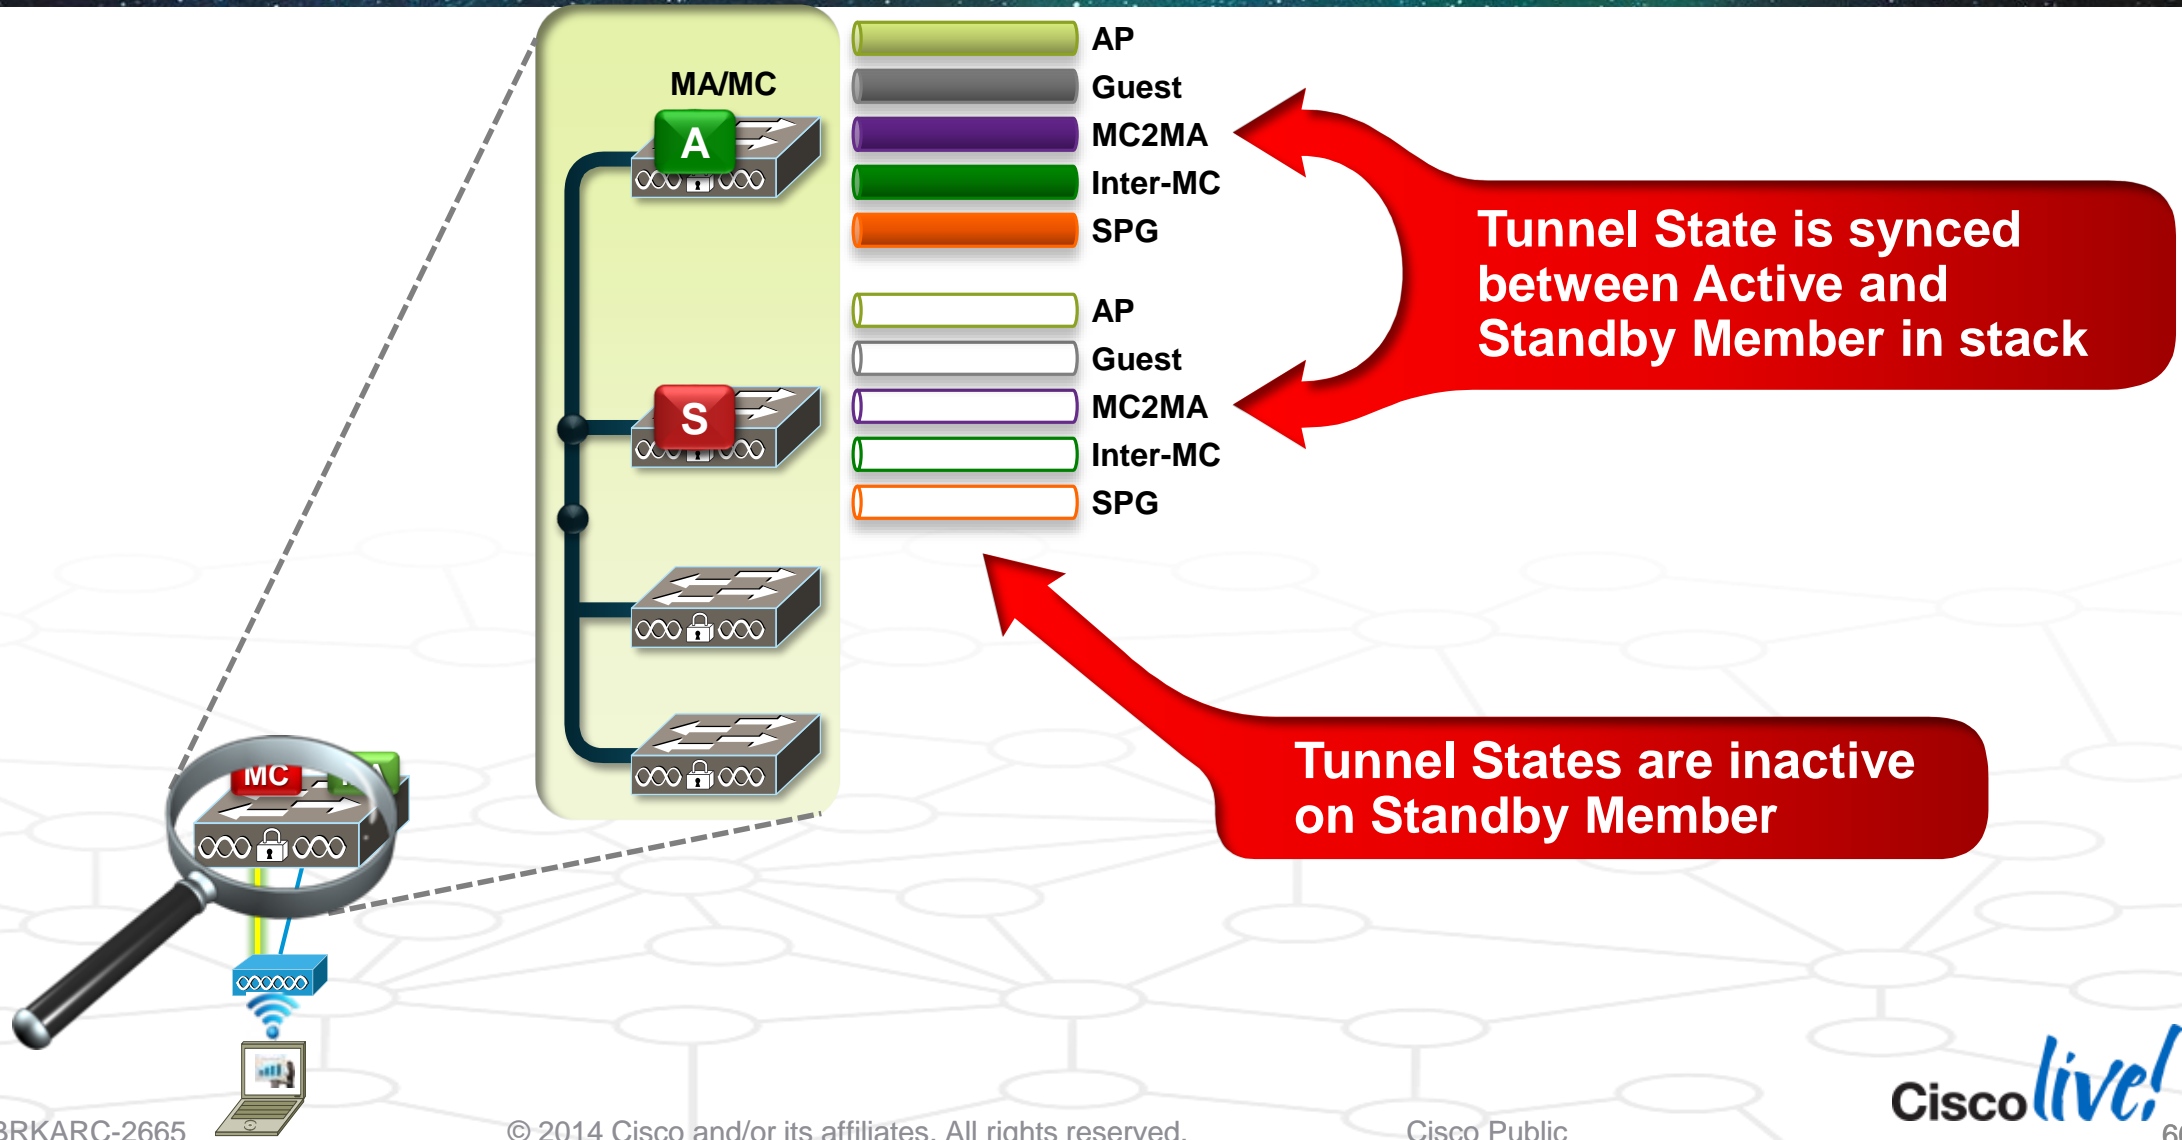


Examining state within the stack (for MC) –

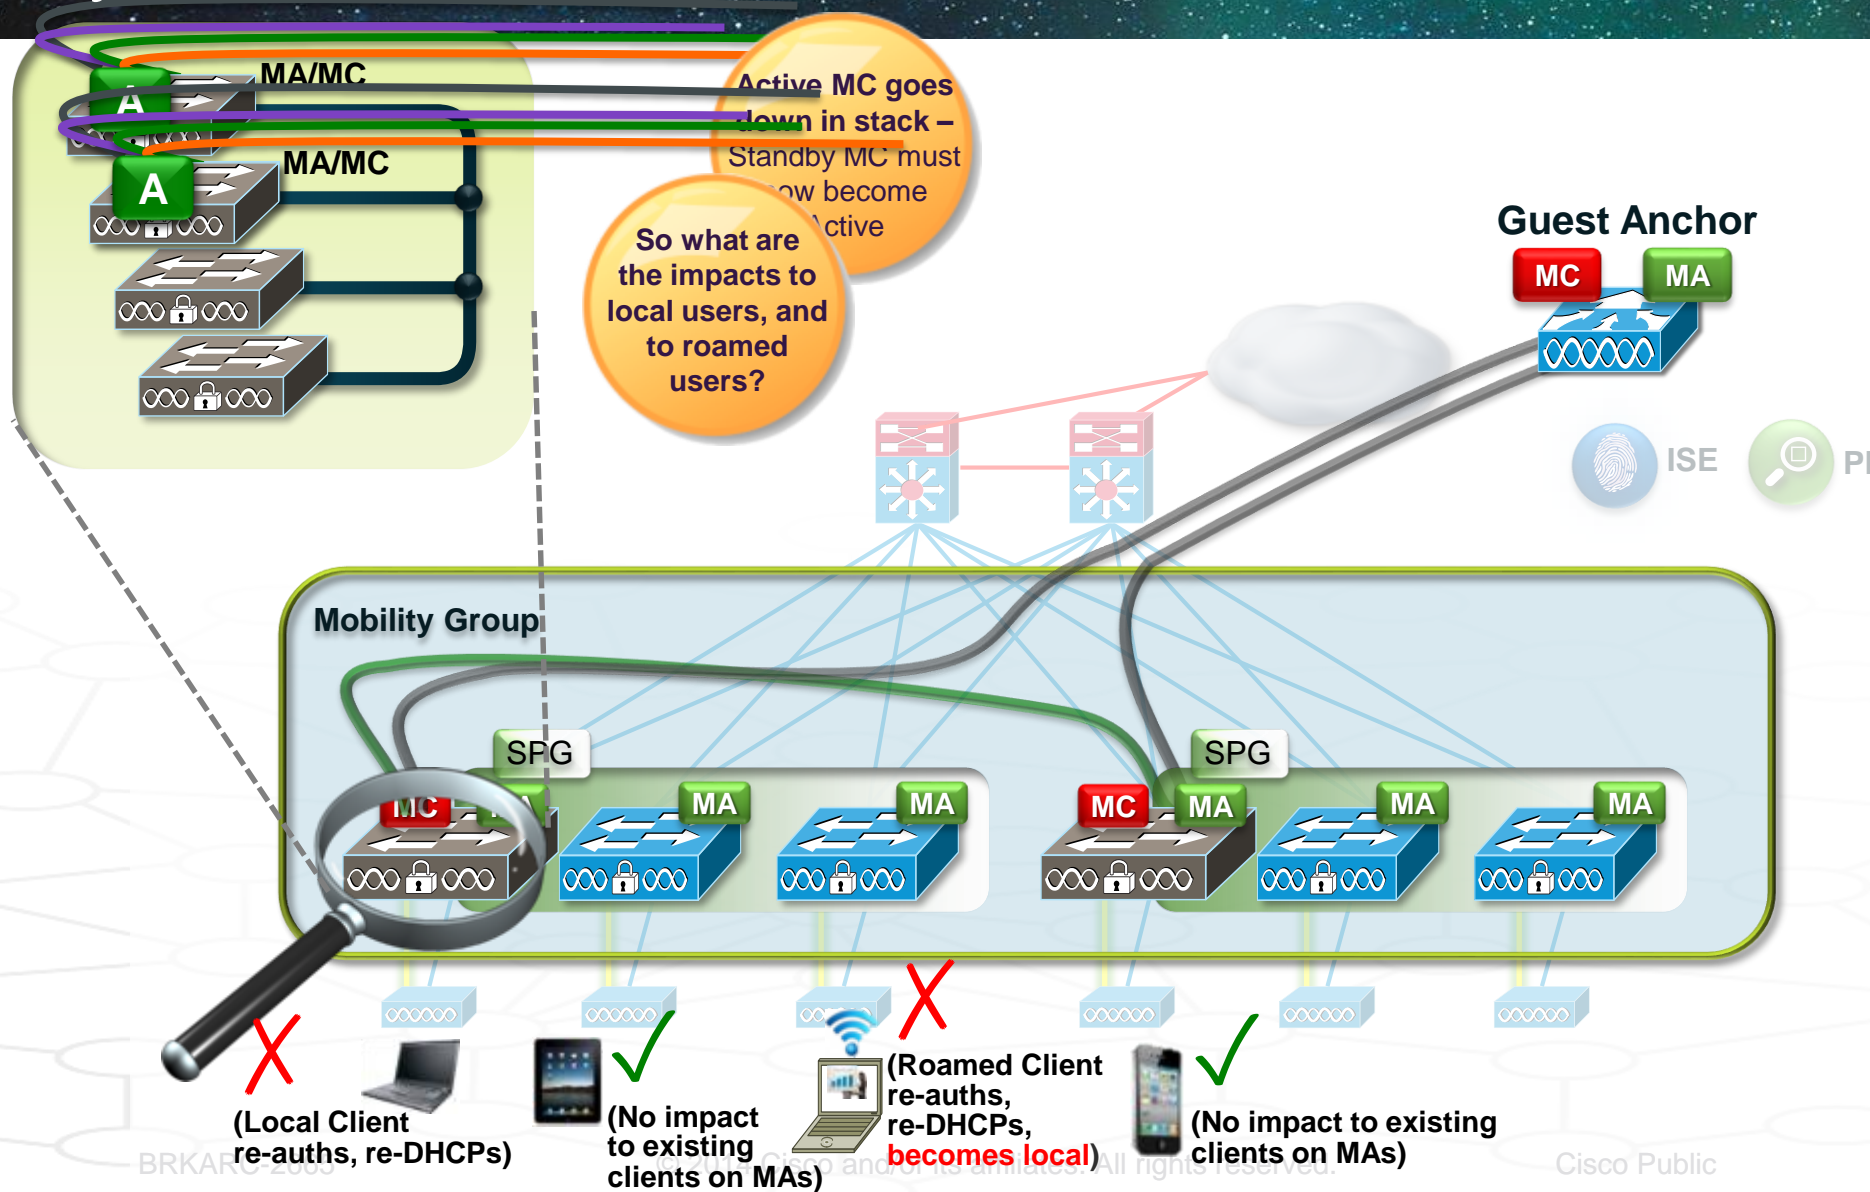
- Let's examine the state maintained by the MC within a stack, and see what redundancy we provide for this ...



High Availability – Catalyst 3x50-based MCs – Tunnel SSO



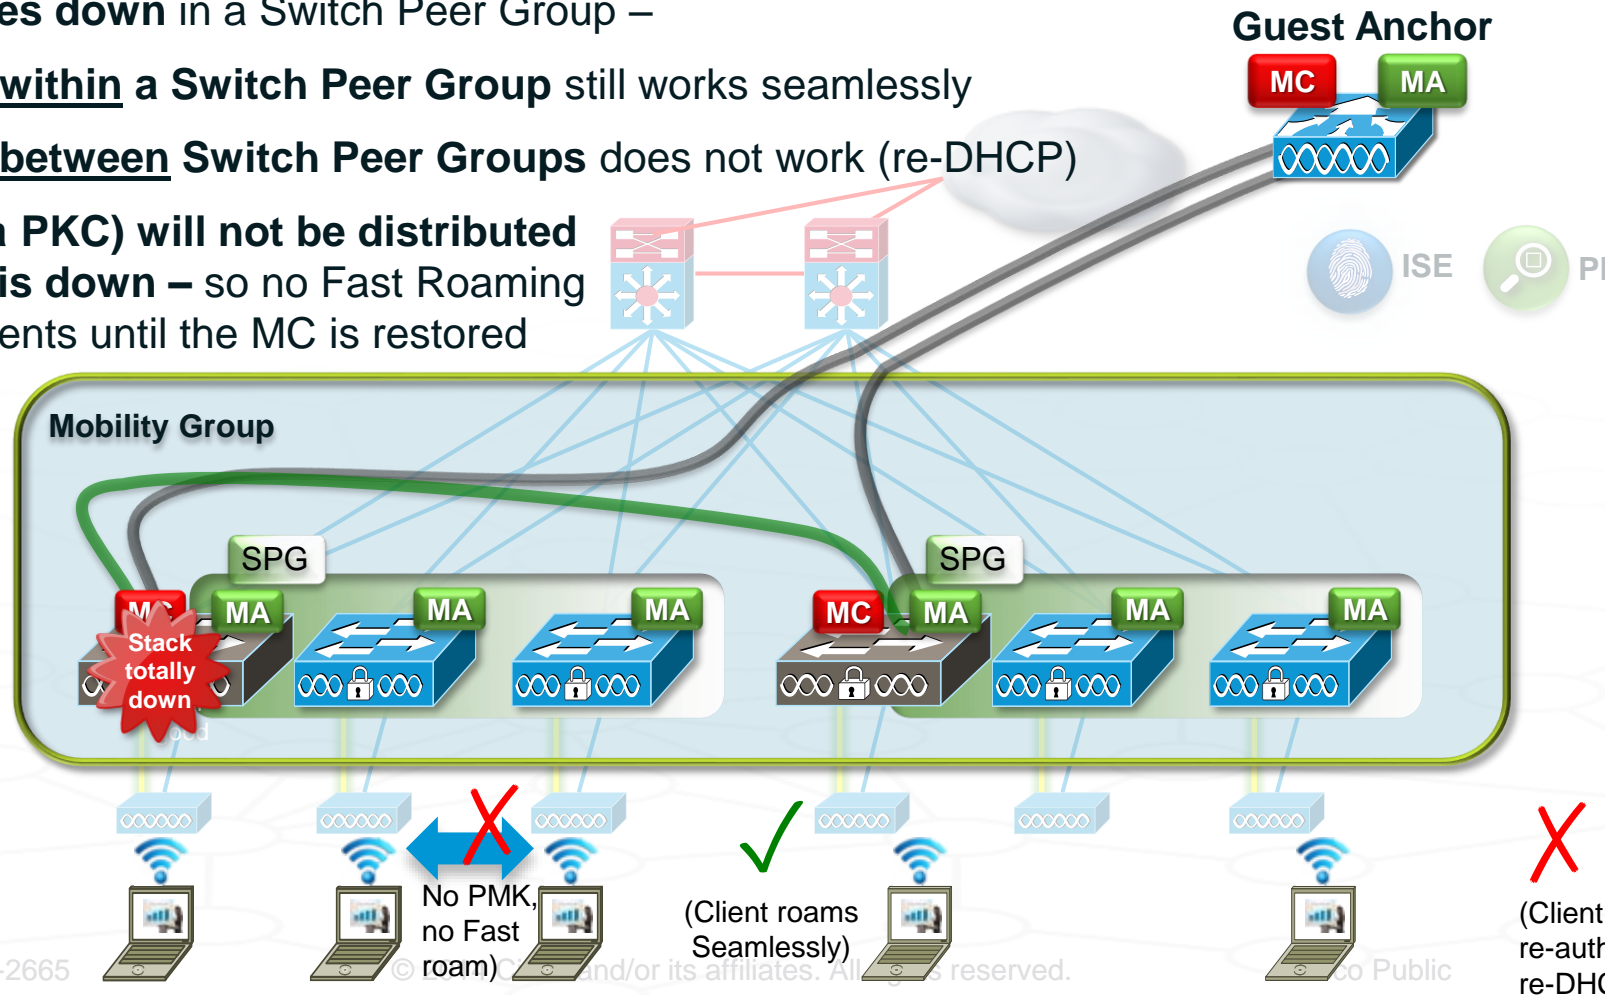
High Availability – Catalyst 3x50-based MCs – Fault Tolerance in Stack



High Availability – Catalyst 3x50-based MCs – Fault Tolerance across Stacks

Switch Peer Group Fault Tolerance with Catalyst 3x50 –

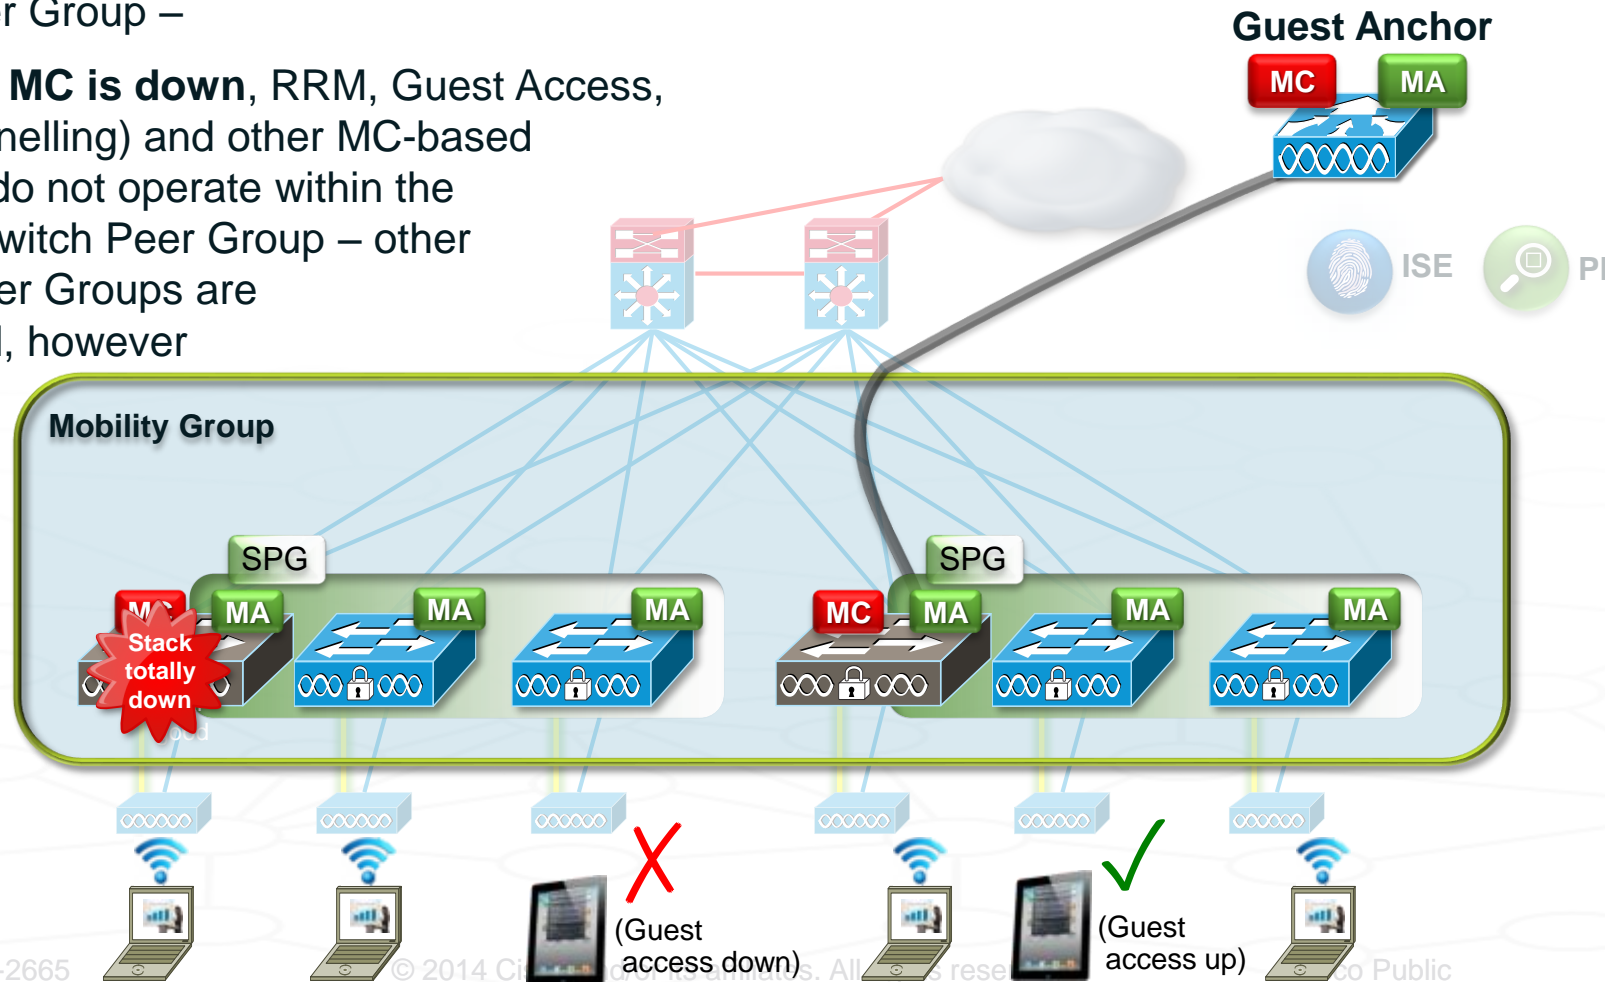
- If an Catalyst 3x50-based stack, operating as an MC, completely goes down in a Switch Peer Group –
 - Roaming within a Switch Peer Group still works seamlessly
 - Roaming between Switch Peer Groups does not work (re-DHCP)
 - PMKs (via PKC) will not be distributed if the MC is down – so no Fast Roaming for new clients until the MC is restored



High Availability – Catalyst 3x50-based MCs – Fault Tolerance across Stacks

Switch Peer Group Fault Tolerance with Catalyst 3x50 –

- If an Catalyst 3x50-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



High Availability – WLC 5760

Two 5760 units can be stacked for 1:1 redundancy, using stack cables

One 5760 elected as Active and the other becomes Hot-Standby

Bulk and Incremental Configuration are synchronised

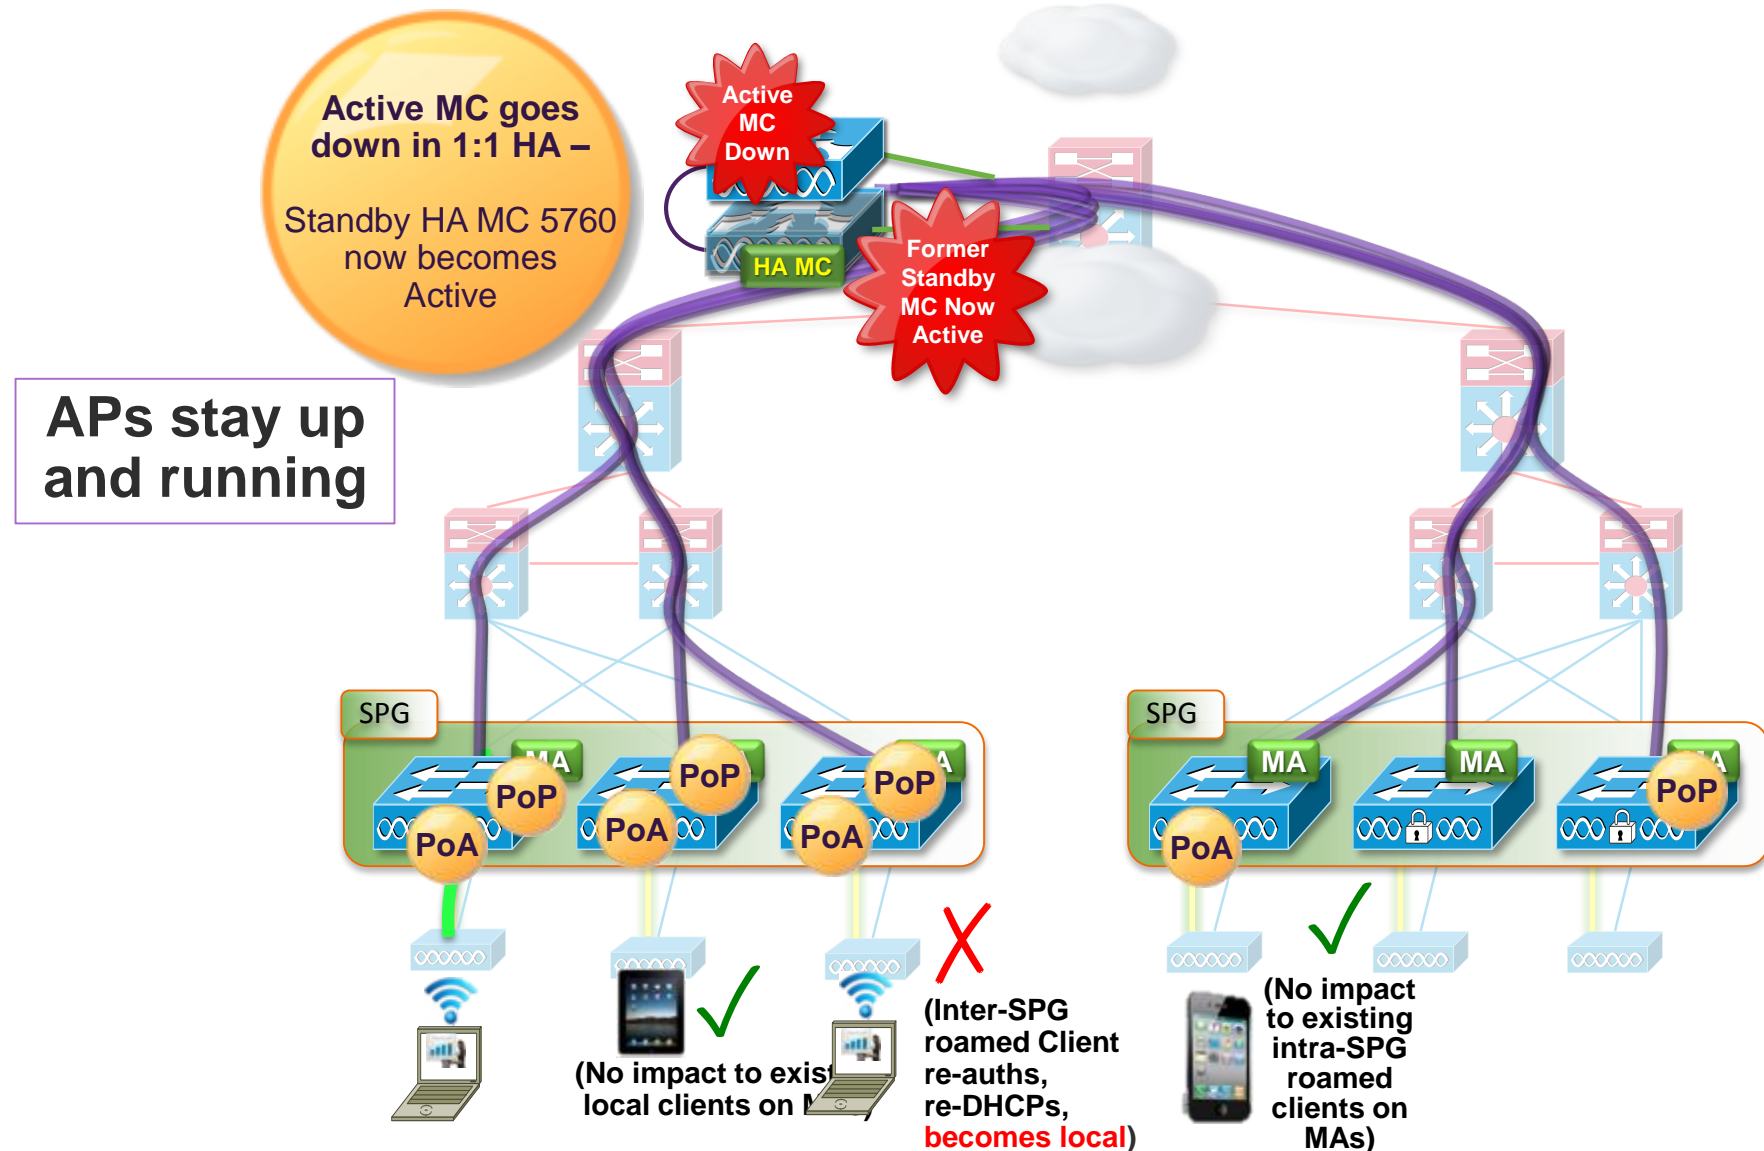
Redundancy is supported both at Port level and System level

AP CAPWAP information is sync'd: APs will not disconnect and continue to be associated to the controller

Significantly reduces network downtime



High Availability – WLC 5760-based MCs – Impact on Clients



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

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- **Quality of Service**

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –

- IP Addressing

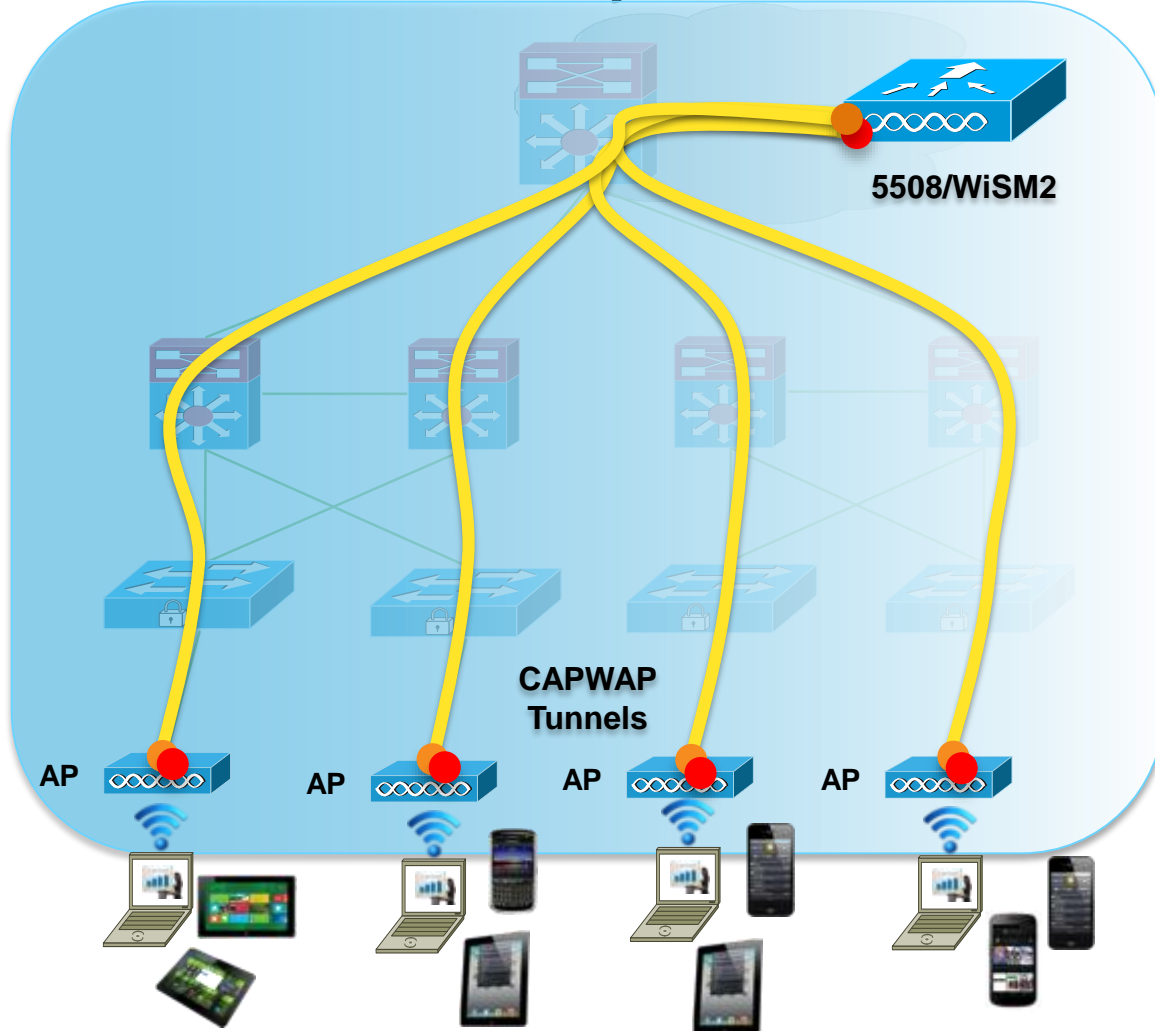
- Design Options

- Deployment Examples

Summary

CUWN Architecture – Overview – Challenges of QoS

Current Mobility Architecture



Challenges –

- Overlay model** with multiple points of policy application*
- Limited **visibility** into applications at the edge
- Lack of **granular classification** at the edge
- Software based **QoS**

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

Wireless QoS Today – How It's Enabled

The screenshot shows the Cisco WLC configuration page for 'Corporate WLAN'. The 'QoS' tab is selected, and the 'Quality of Service (QoS)' dropdown menu is open, showing four options: Platinum (voice), Gold (video), Silver (best effort), and Bronze (background). A red circle highlights the dropdown menu, and a red arrow points from a text box to it.

Under the WLAN one of four QoS profiles can be assigned. By default each profile has a default .1p assigned, but it can be modified using the Wired QoS Protocol options.

*NOTE: Assignment of QoS profile to WLAN

liates. All rights reserved.

Cisco Public

How Do We Enable QoS Today?

Wired – mls-based CLI Exposes Hardware

```
C3750-X(config)# mls qos
C3750-X(config)# interface GigabitEthernet 1/0/1
C3750-X(config-if)# mls qos trust dscp
```

```
C3750-X(config)# mls qos queue-set output 1 buffers 15 30 35 20
C3750-X(config)# mls qos queue-set output 1 threshold 1 100 100 100 100
C3750-X(config)# mls qos queue-set output 1 threshold 2 80 90 100 400
C3750-X(config)# mls qos queue-set output 1 threshold 3 100 100 100 400
C3750-X(config)# mls qos queue-set output 1 threshold 4 60 100 100 400
C3750-X(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46
```

```
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 2 24
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
C3750-X(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 0
C3750-X(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8
C3750-X(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
```

```
C3750-X(config)# interface range GigabitEthernet1/0/1-48
C3750-X(config-if-range)# queue-set 1
C3750-X(config-if-range)# srr-queue bandwidth
C3750-X(config-if-range)# priority-queue
```

NOTE: Only class based policing and marking are available today – last box with mls cli - Cat 3750

QoS – What's New with Converged Access

Wired (Cat 3850)

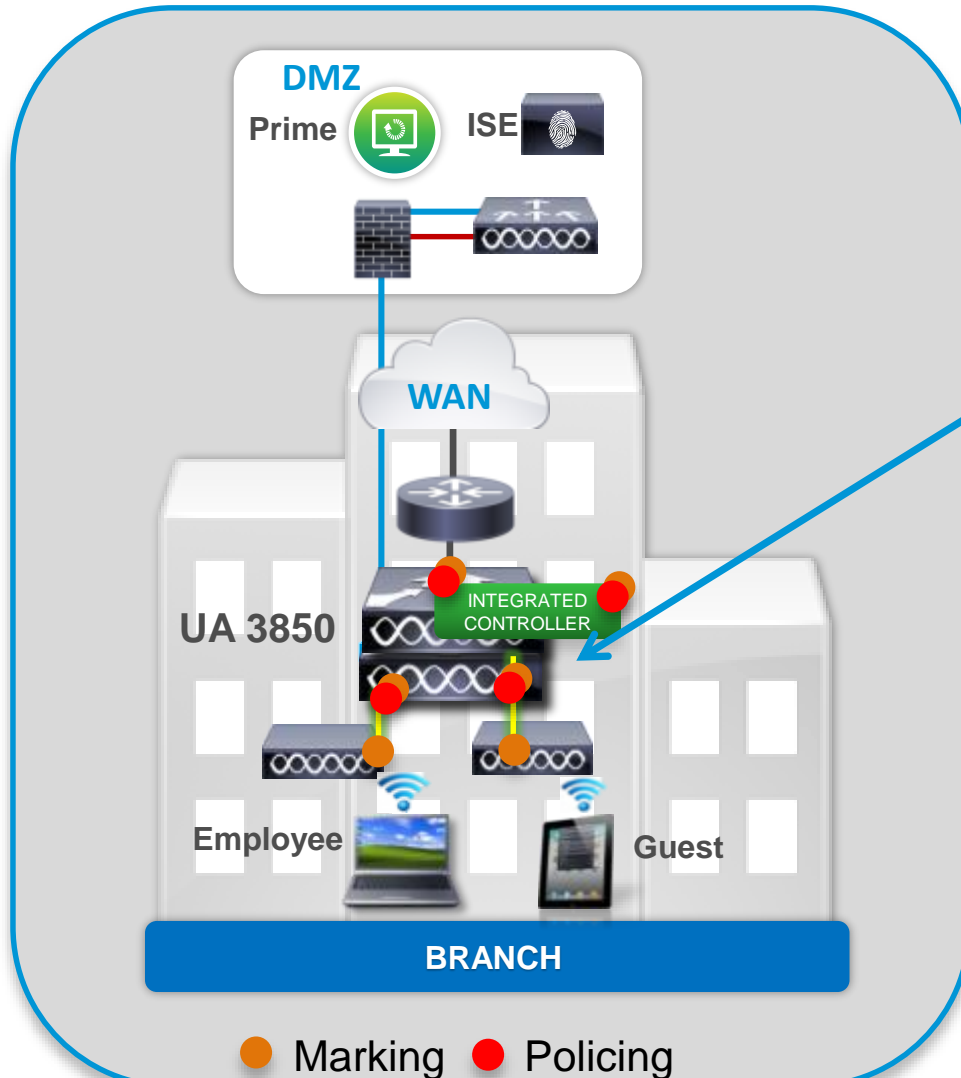
- Modular QoS based CLI (MQC)
 - Alignment with 4500E series (Sup6, Sup7)
 - Class-based Queueing, Policing, Shaping, Marking
- More Queues
 - Up to 2P6Q3T queueing capabilities
 - Standard 3750 provides 1P3Q3T
 - Not limited to 2 queue-sets
 - Flexible MQC Provisioning abstracts queueing hardware

Wireless (Cat 3x50 & 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 3x50 & CT 5760)

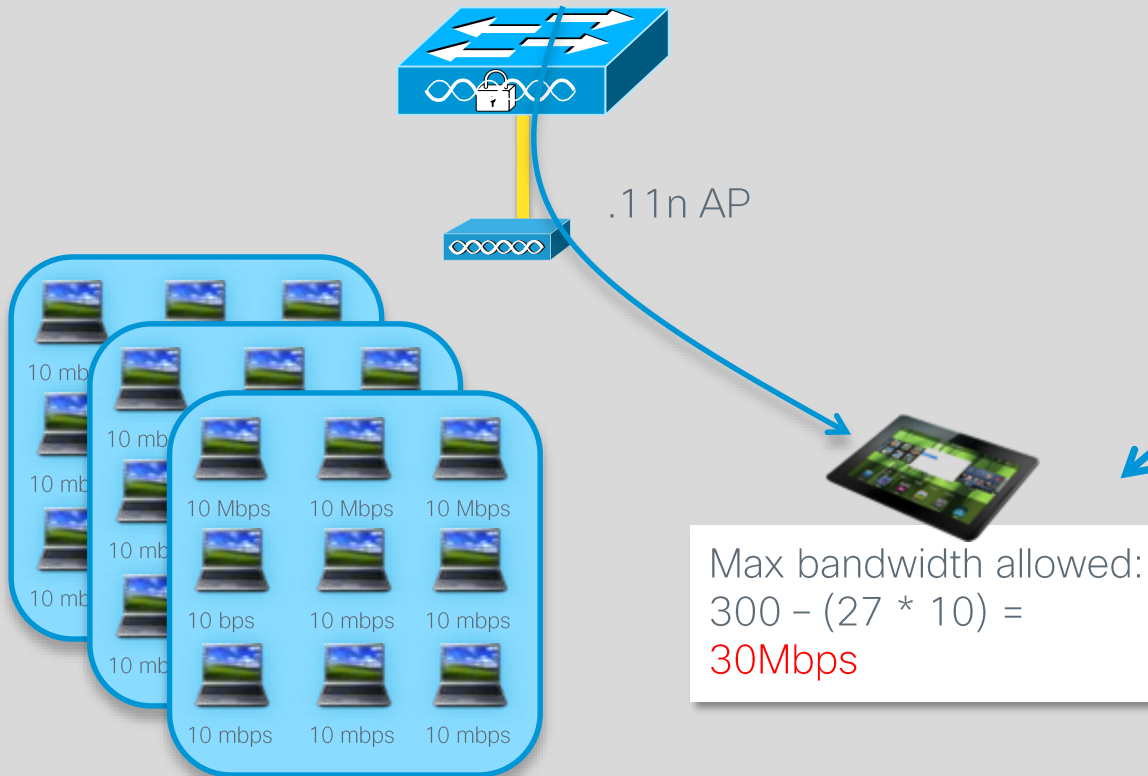
- **Granular QoS control at the wireless edge**
 - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- **Enhanced Bandwidth Management**
 - Approximate Fair Drop (AFD) Bandwidth Management ensures 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 currently supported on the CT 5760

QoS – What's New with Converged Access

With the CT 5760 or CAT 3850 / 3650

Usage based fair allocation **without configuration**



Wireless (Cat 3x50 & 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 currently supported on the CT 5760*

QoS – What's New with Converged Access

With the 3850 / 3650

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 3x50 & CT 5760)

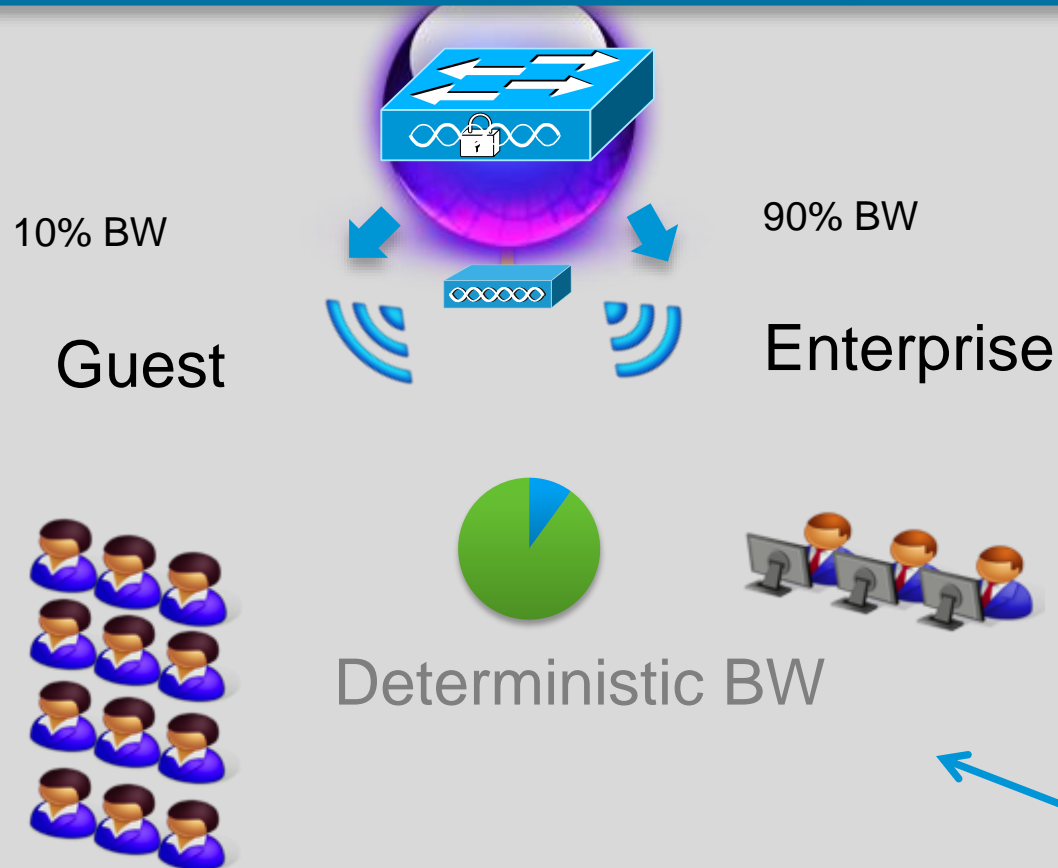
- Granular QoS control at the wireless edge
 - Tunnel termination allows customers to provide QoS treatment per SSIDs, per-Clients and common treatment of wired and wireless traffic throughout the network
- Enhanced Bandwidth Management
 - Approximate Fair Drop (AFD) Bandwidth Management ensures fairness 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 currently supported on the CT 5760

QoS – What's New with Converged Access

With the CT 5760 or CAT 3850 / 3650

Deterministic bandwidth is allocated per SSID



Wireless (Cat 3x50 & 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**

QoS – What's New with Converged Access –

MQC Provides a Unified Provisioning Language

```
C3750-X(config)# mls qos
C3750-X(config)# interface GigabitEthernet 1/0/1
C3750-X(config-if)# mls qos trust dscp
```

```
C3750-X(config)# mls qos queue-set output 1 buffers 15 30 35 20
C3750-X(config)# mls qos queue-set output 1 threshold 1 100 100 100 100
C3750-X(config)# mls qos queue-set output 1 threshold 2 80 90 100 400
C3750-X(config)# mls qos queue-set output 1 threshold 3 100 100 100 400
C3750-X(config)# mls qos queue-set output 1 threshold 4 60 100 100 400
C3750-X(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46
```

```
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 2 24
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
C3750-X(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 0
C3750-X(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8
C3750-X(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
```

```
C3750-X(config)# interface range GigabitEthernet1/0/1-48
C3750-X(config-if-range)# queue-set 1
C3750-X(config-if-range)# srr-queue bandwidth share 1 30 35 5
C3750-X(config-if-range)# priority-queue out
```

policy-map 3850-QoS

```
class PRIORITY-QUEUE
priority level 1
police rate percent 20
```

```
class CONTROL-MGMT-QUEUE
bandwidth remaining percent 30
queue-limit dscp cs2 percent 80
queue-limit dscp cs3 percent 90
queue-limit dscp cs6 percent 100
```

```
class TRANSACTIONAL-DATA-QUEUE
bandwidth remaining percent 5
queue-limit dscp af23 percent 80
queue-limit dscp af22 percent 90
queue-limit dscp af21 percent 100
```

```
class BULKDATA-QUEUE
bandwidth remaining percent 35
queue-limit dscp af13 cs1 percent 80
queue-limit dscp af12 percent 90
queue-limit dscp af11 percent 100
```

*NOTE: Only class based policing and marking are available today – last box with mls cli - Cat 3750

Wireless Queuing and Approximate Fair Drop (AFD)

Into a wired port



Out of a wireless port



Application Visibility – Without Control – Catalyst 3x50

```
flow record fr-avc
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 match flow direction
 match application name
 match wireless ssid
 collect counter bytes long
 collect counter packets long
 collect wireless ap mac address
 collect wireless client mac address
 end
```

```
flow monitor fm-avc
 record fr-avc
 cache timeout inactive 200
 end
```

```
wlan <>
 ip flow mon fm-avc input
 ip flow mon fm-avc output
 end
```

Cat3850# **sh avc client 8c70.5a20.35b4 top 10 application agg**
Cumulative Stats:

No.	AppName usage%	Packet-Count	Byte-Count	AvgPkt-Size	

	--				
1	http	69451	72146465	1038	67
2	youtube	16284	17117601	1051	15
3	rtmpe	9349	9266013	991	8
4	hulu	8096	7974952	985	7
5	unknown	1686	126067	74	0
6	rtmp	1593	1723269	1081	2
7	netflix	1305	1371679	1051	1
8	ssl	937	530577	566	0
9	dns	748	70418	94	0
10	facebook	512	372629	727	0


Last Interval (90 seconds) Stats:

No.	AppName usage%	Packet-Count	Byte-Count	AvgPkt-Size	

1	http	65410	68322192	1044	78
2	rtmpe	8812	9242082	1048	11
3	youtube	5752	6262985	1088	7
4	rtmp	1593	1723269	1081	2
5	netflix	1305	1371679	1051	2
6	unknown	797	76004	95	0
7	dns	265	29420	111	0
8	flash-video	206	196639	954	0
9	ssl	148	62384	421	0
10	hulu	82	25238	307	0

Quality of Service – Recommended Reading

- **Comprehensive QoS design guidance for PINs and platforms –**
 - Campus **Catalyst 3750 / 4500 / 6500**
 - WLAN **WLC 5508 / Catalyst 3850/3650**
 - Data Centre **Nexus 1000V / 2000 / 5500 / 7000**
 - WAN & Branch **Cisco ASR 1000 / ISR G2**
 - MPLS VPN **Cisco ASR 9000 / CRS-3**
 - IPSec VPNs **Cisco ISR G2**
- **ISBN: 1-58714-369-0**



Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service



- Security**

- Multicast

- NetFlow

Converged Access Design and Deployment –

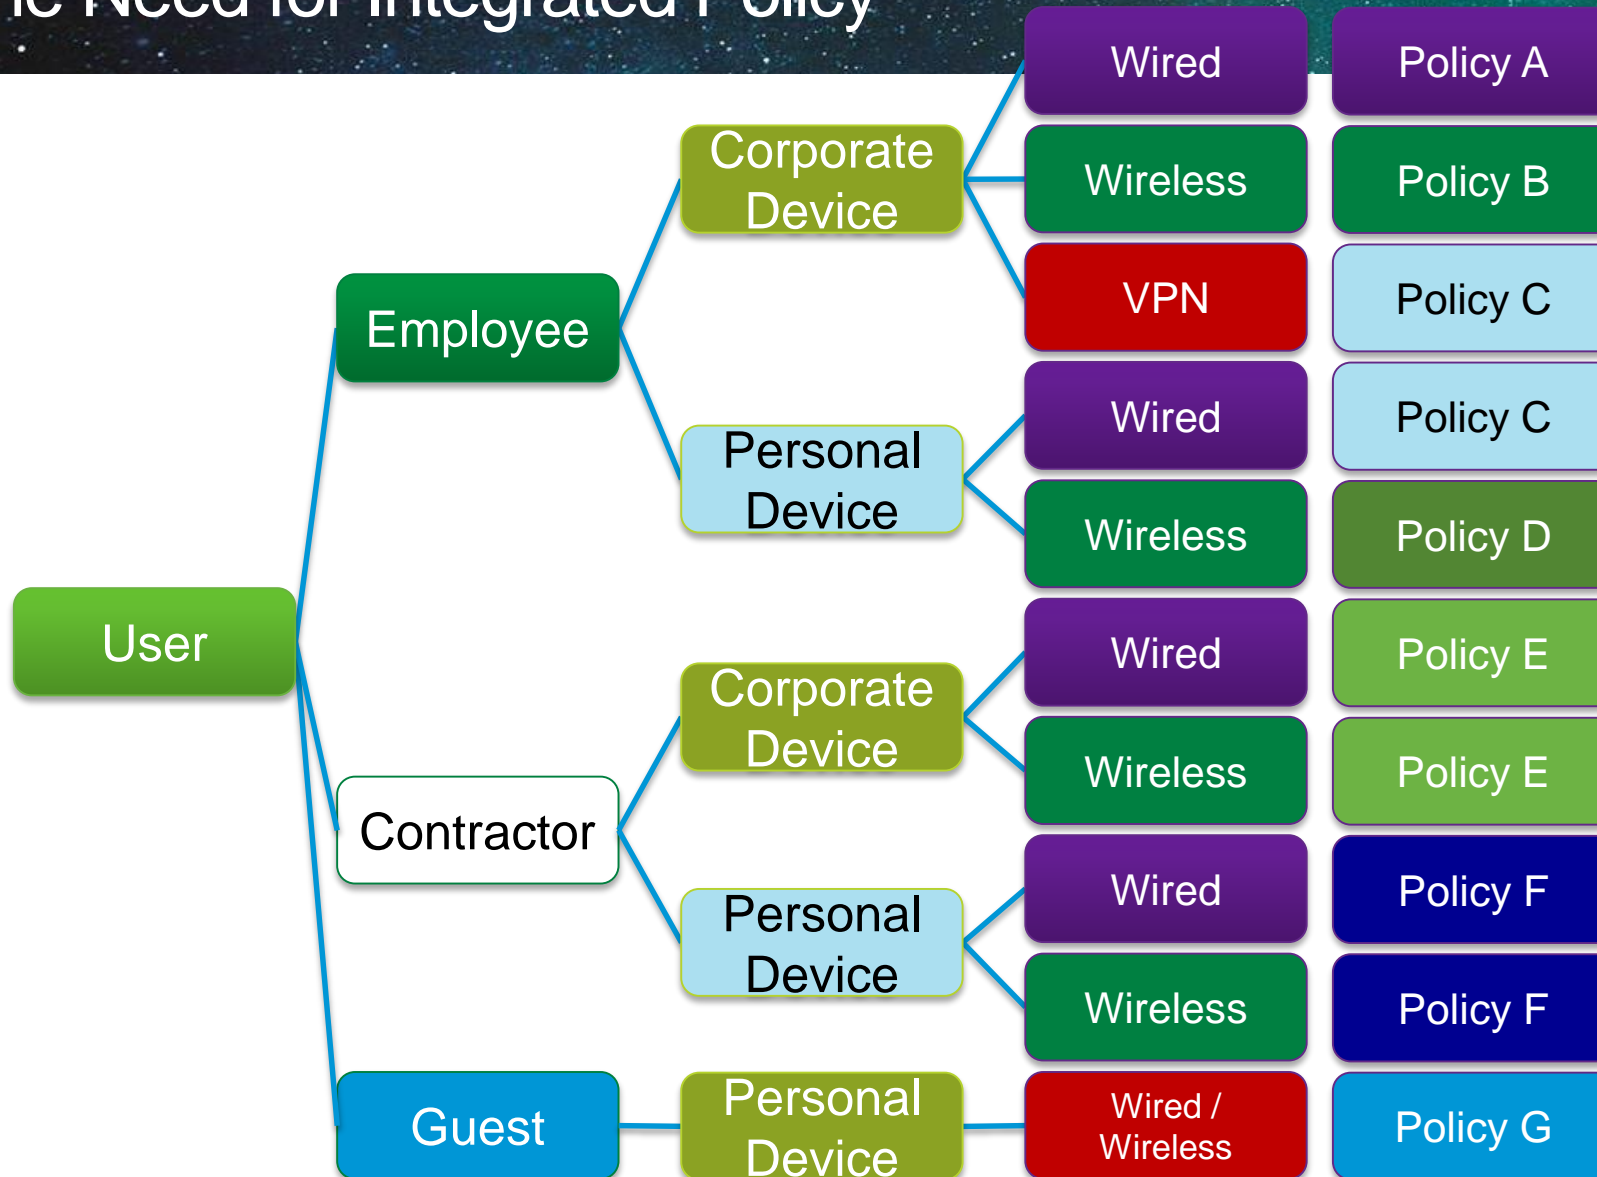
- IP Addressing

- Design Options

- Deployment Examples

Summary

Converged Access – The Need for Integrated Policy



How to **define and apply** security policy **consistently** across every device on the network?

Policy Definition – Where?


Distributed and/or Centralised

- On-Device Policy
 - AAA services (mandatory)
 - Local Policy Objects
 - Local Policy
 - Users
- Central Policy
 - Users / External Databases
 - Central Policy Objects
 - Central Policy and Control
 - Profiling
- Typically a Combination of both




Today – Inconsistent Central Policy Definition

One Policy to the Rescue!

Feature		
		
ACL Application	dACL, Filter-ID, per-User ACL	Airspace-ACL-Name
VLAN Assignment	Tunnel-Type, Tunnel-Medium-Type, Tunnel-Private-Group-ID	As with wired but PLUS Airspace-Interface-name
QoS	Platform dependent ☹️ (C3PL, MQC, ...)	Airspace-QoS-Level, Airspace-DSCP

C3PL: Cisco Classification Configuration Policy Language
MQC: Modular QoS CLI

One Policy – Wired and Wireless



- Employee using the same SSID, can be associated to different VLAN interfaces and policy after EAP authentication
- Employee using corporate wired and wireless device with their AD user id can be assigned to same VLAN 30 to have full access to the network
- Employee using personal iDevice with their AD user id can be assigned to VLAN 40 and policy to access internet only

ISE Policy Definition Example –

Same Authorisation Policy for Wired AND Wireless

Employee-Personal-Device if **RegisteredDevices** AND (Radius-Service-Type-Frame AND **Wired-OR-Wireless-802.1x** AND Radius:Called-Station-ID EQUALS CERTIFICATE:Subject Alternative Name AND Network Access:EapAuthentication EQUALS EAP-TLS AND AD1:ExternalGroups EQUALS WS2008er2.corp1.rf-demo.com/Users/byod_user) then Restricted-Access-Employee

Contractor-Personal-Device if **RegisteredDevices** AND (Radius-Service-Type-Frame AND **Wired-OR-Wireless-802.1x** AND Radius:Called-Station-ID EQUALS CERTIFICATE:Subject Alternative Name AND Network Access:EapAuthentication EQUALS EAP-TLS AND AD1:ExternalGroups EQUALS WS2008er2.corp1.rf-demo.com/Users/Domain Users) then Restricted-Access-Contractor

Guest-Personal-Device if **RegisteredDevices** AND (Radius-Service-Type-Frame AND **Wired-OR-Wireless-802.1x** AND Radius:Called-Station-ID EQUALS CERTIFICATE:Subject Alternative Name AND Network Access:EapAuthentication EQUALS EAP-TLS AND AD1:ExternalGroups EQUALS WS2008er2.corp1.rf-demo.com/Users/Guest) then Internet-Access-Policy

Compound Condition

* Name: **Wired-OR-Wireless-802.1x**

Description: A Condition To Match An 802.1X Based Authentication Request From Cisco Converged Access Platform.

*Condition Expression

Condition Name	Expression	Operator	Value
	Radius:NAS-Port-Type	Equals	Ethernet
	Radius:NAS-Port-Type	Equals	IEEE 802.11

Attributes Details

- Access Type = ACCESS_ACCEPT
- Tunnel-Private-Group-ID = 1:101
- Tunnel-Type=1:13
- Tunnel-Medium-Type=1:6
- DAACL = corp-policy-1
- cisco-av-pair = ip:sub-qos-policy-in=Standard-Employee
- cisco-av-pair = ip:sub-qos-policy-out=Standard-Employee

Converged Access – Security Features



	Catalyst 3650 & 3850	CT5760	CT5508
BYOD Functionality	YES	YES	YES
Rogue detect / classify / contain, RDLP	YES	YES	YES
Port Security	YES	YES	NO
IP Source Guard	YES	YES	NO
Dynamic ARP Insp.	YES	YES	NO
LDAP, TACACS+, RADIUS	YES	YES	YES
LSC and MIC	YES	YES	YES
AP dot1x EAP-FAST	YES	YES	YES
Secure Fast Roaming	YES	YES	YES
802.1X-rev-2010 (MACsec / MKA)	H/W Ready	H/W Ready	NO

Converged Access – Security Features, continued



	Catalyst 3650 & 3850	CT5760	CT5508
IP Theft, DHCP Snooping, Data Gleaning	YES	YES	YES
IOS ACL	YES	YES	YES
Adaptive wIPS, WPS	YES	YES	YES
CIDS	YES	YES	YES
TrustSec SGT / SGACL	YES	YES	SXP
Guest Access	YES	YES	YES
IPv6 RA Guard	YES	YES	NO
MFP	YES	YES	YES
IP Device Tracking	YES	YES	NO
CoPP	Static	Static	NO

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security



- Multicast**

- NetFlow

Converged Access Design and Deployment –

- IP Addressing

- Design Options

- Deployment Examples

Summary

Multicast on Converged Access WLC – Default State

- Multicast forwarding is disabled on CA Controller by default –
Multicast frames received at the controller level are not forwarded to / from APs
- IGMP Snooping is disabled –
The Controller does not know which AP needs multicast forwarding
- Both Multicast forwarding and IGMP snooping are needed, use the following commands –

```
(config)# wireless multicast
```

 enables IP mcast

```
(config)# ip igmp snooping
```

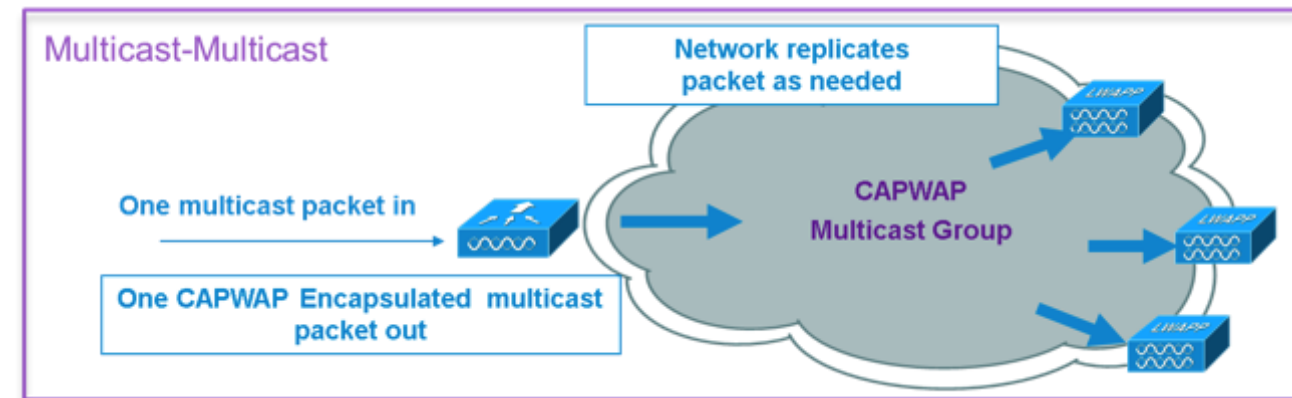
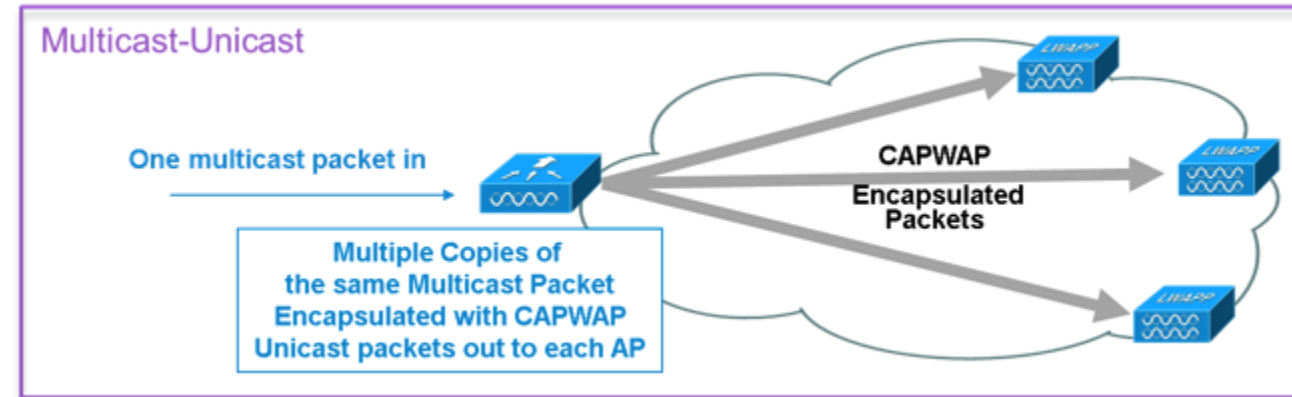
 enables snooping for IPv4

```
(config)# ipv6 mld snooping
```

 enables snooping for IPv6
- As in CUWN, CAPWAP multicast-multicast mode can be configured on the Controller –


```
(config)# ap capwap multicast <ip mcast >
```


 By default multicast-unicast is used



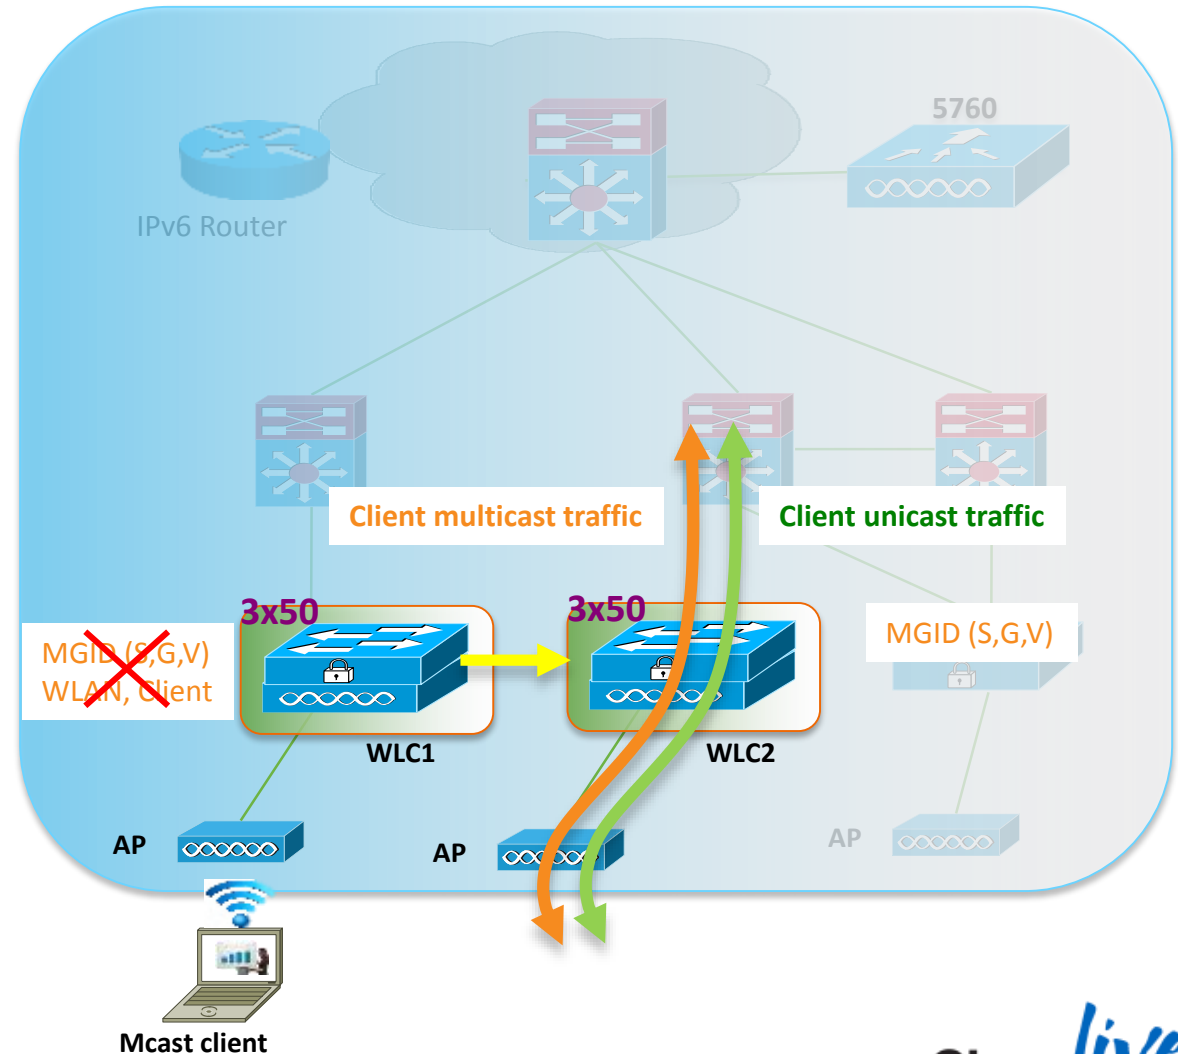
Multicast on Converged Access WLC – Roaming – Intra-WLC

- Intra-WLC roaming –
 1. WLC notifies IOS with client move notification
 2. IOS updates the capwap ports for MGIDs (groups) to which the client was subscribed
 3. WLC checks with the Medianet configuration and AP capability to see if the group should be allowed in the new AP.
 4. If allowed, WLC adds all the new MGID to the new AP and deletes the client reference from the MGID from the old AP.



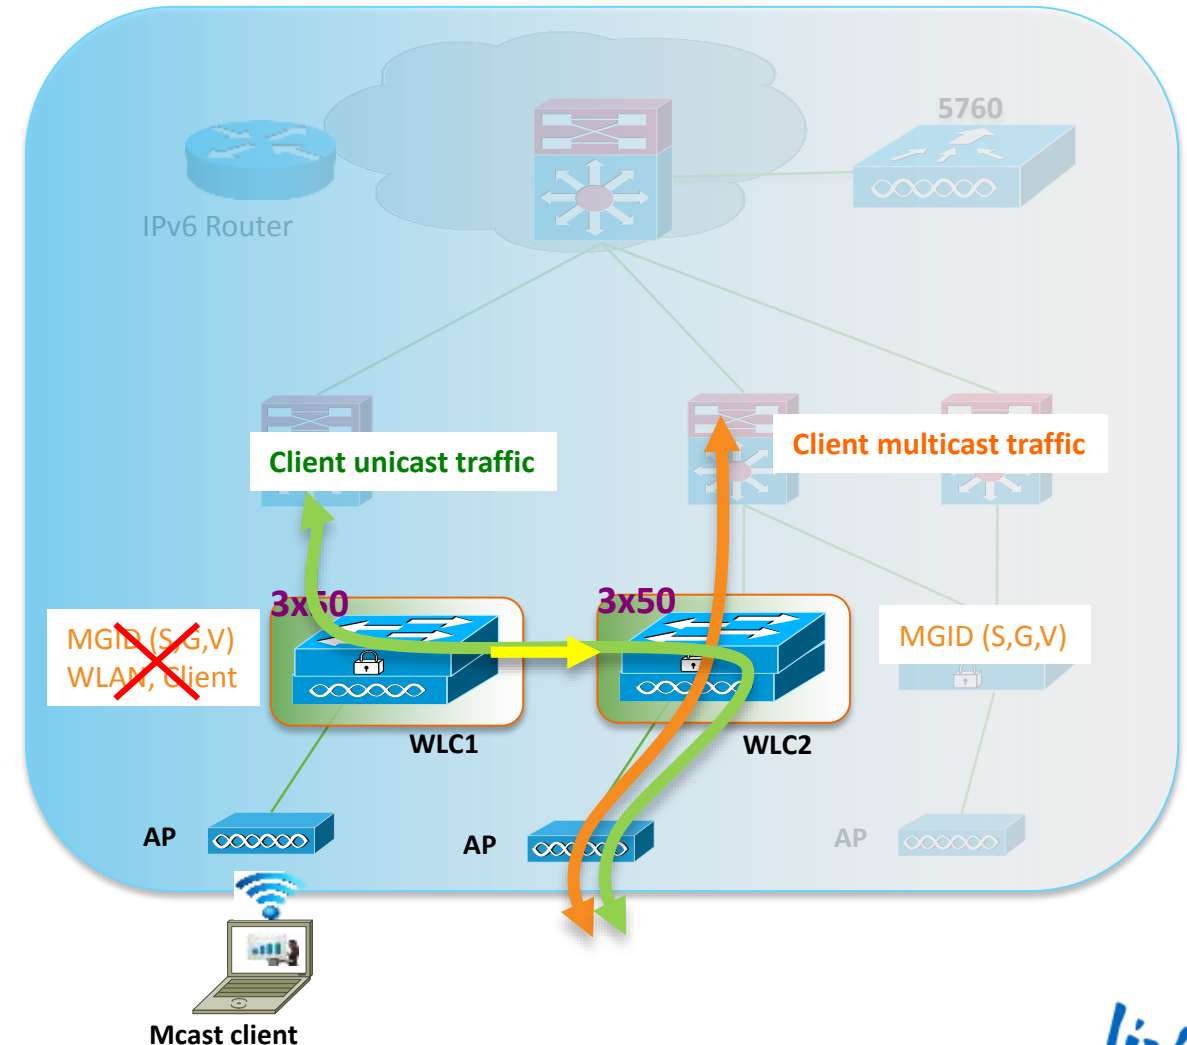
Multicast on Converged Access WLC – Roaming – Inter-WLC, Layer 2

- Inter-WLC Layer 2 roaming, **Sticky Anchoring disabled** –
- WLC in switch1/WLC1 **moves** all the group info in the mobility handoff payload to the switch2 /WLC2
 - WLC in switch2/WLC2 creates new MGID as if igmp report packets from the client has arrived
 - Old switch1/WLC1 removes all the client references as if a leave message was received
 - Multicast traffic flows through the new controller



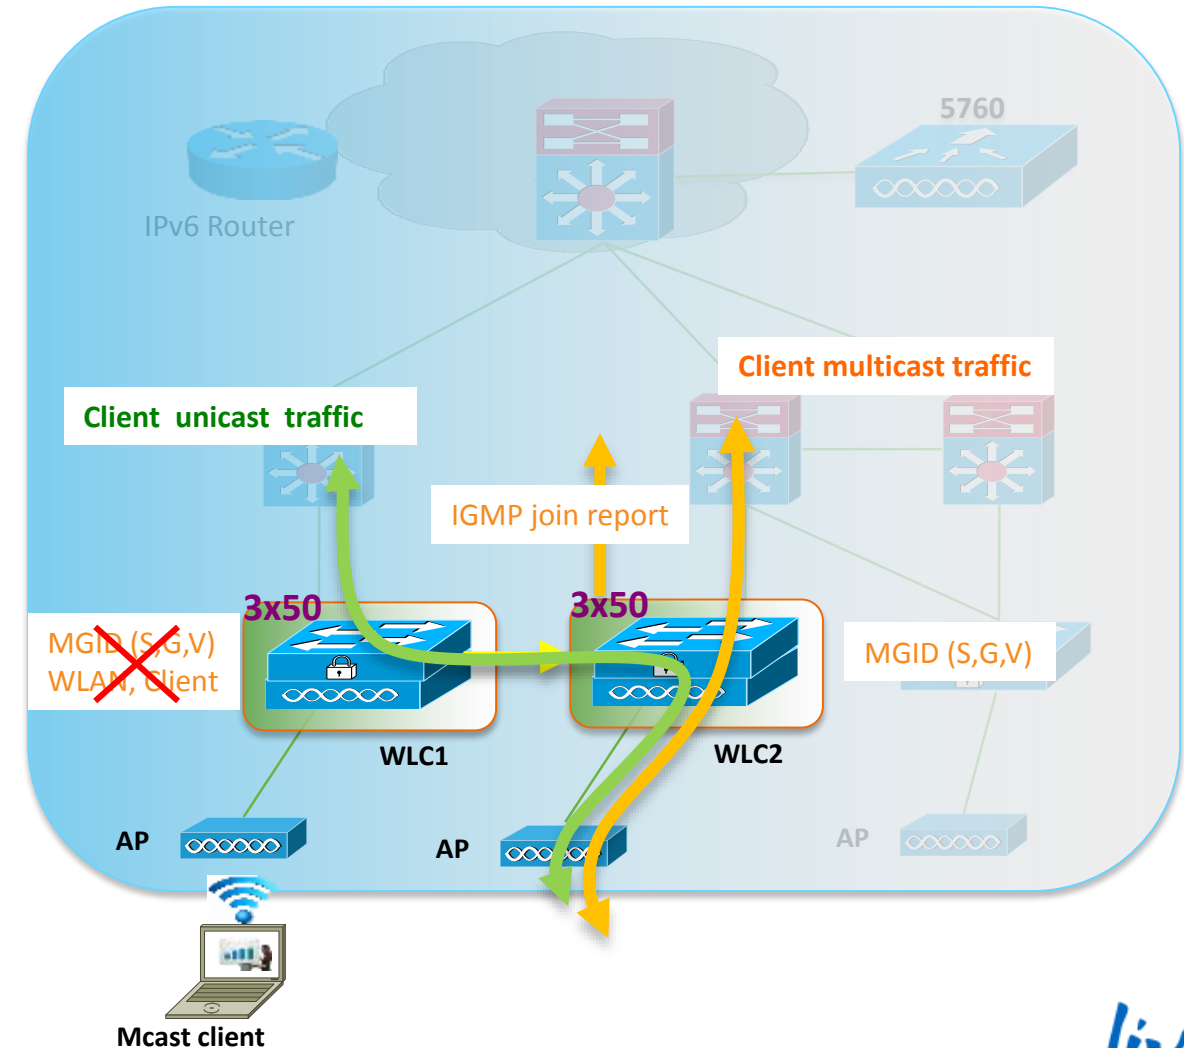
Multicast on Converged Access WLC – Roaming – Inter-WLC, Layer 2

- Inter-WLC Layer 2 roaming,
Sticky Anchoring enabled (default) –
- 1. WLC in switch1 / WLC1 **moves** all the multicast group info in the mobility handoff payload to the switch2 / WLC2. **Unicast info is copied.**
- 2. WLC in switch2 / WLC2 creates new MGID as if igmp report packets from the client has arrived
- 3. Old switch1 / WLC1 keeps client unicast references, but drops multicast references as if a leave message was received
- 4. Multicast traffic flows through the new controller ... but unicast traffic still goes through switch1 / WLC1 to preserve client IP



Multicast on Converged Access WLC – Roaming – Inter-WLC, Layer 3

- Inter-WLC Layer 3 roaming –
 1. WLC in switch1 / WLC1 **moves** all the multicast group info in the mobility handoff payload to the switch2 / WLC2. **Unicast info is copied.**
 2. WLC in switch2 / WLC2 creates new MGID
 3. Old switch1 / WLC1 removes all the client references
 4. Switch2 / WLC2 sends IGMP reports based on its WLAN/VLAN mapping
 5. Multicast traffic flows through the new controller ... but unicast traffic still goes through switch1 / WLC1 to preserve client IP



Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast



- NetFlow**

Converged Access Design and Deployment –

- IP Addressing

- Design Options

- Deployment Examples

Summary

Troubleshooting – From Simple to Complex

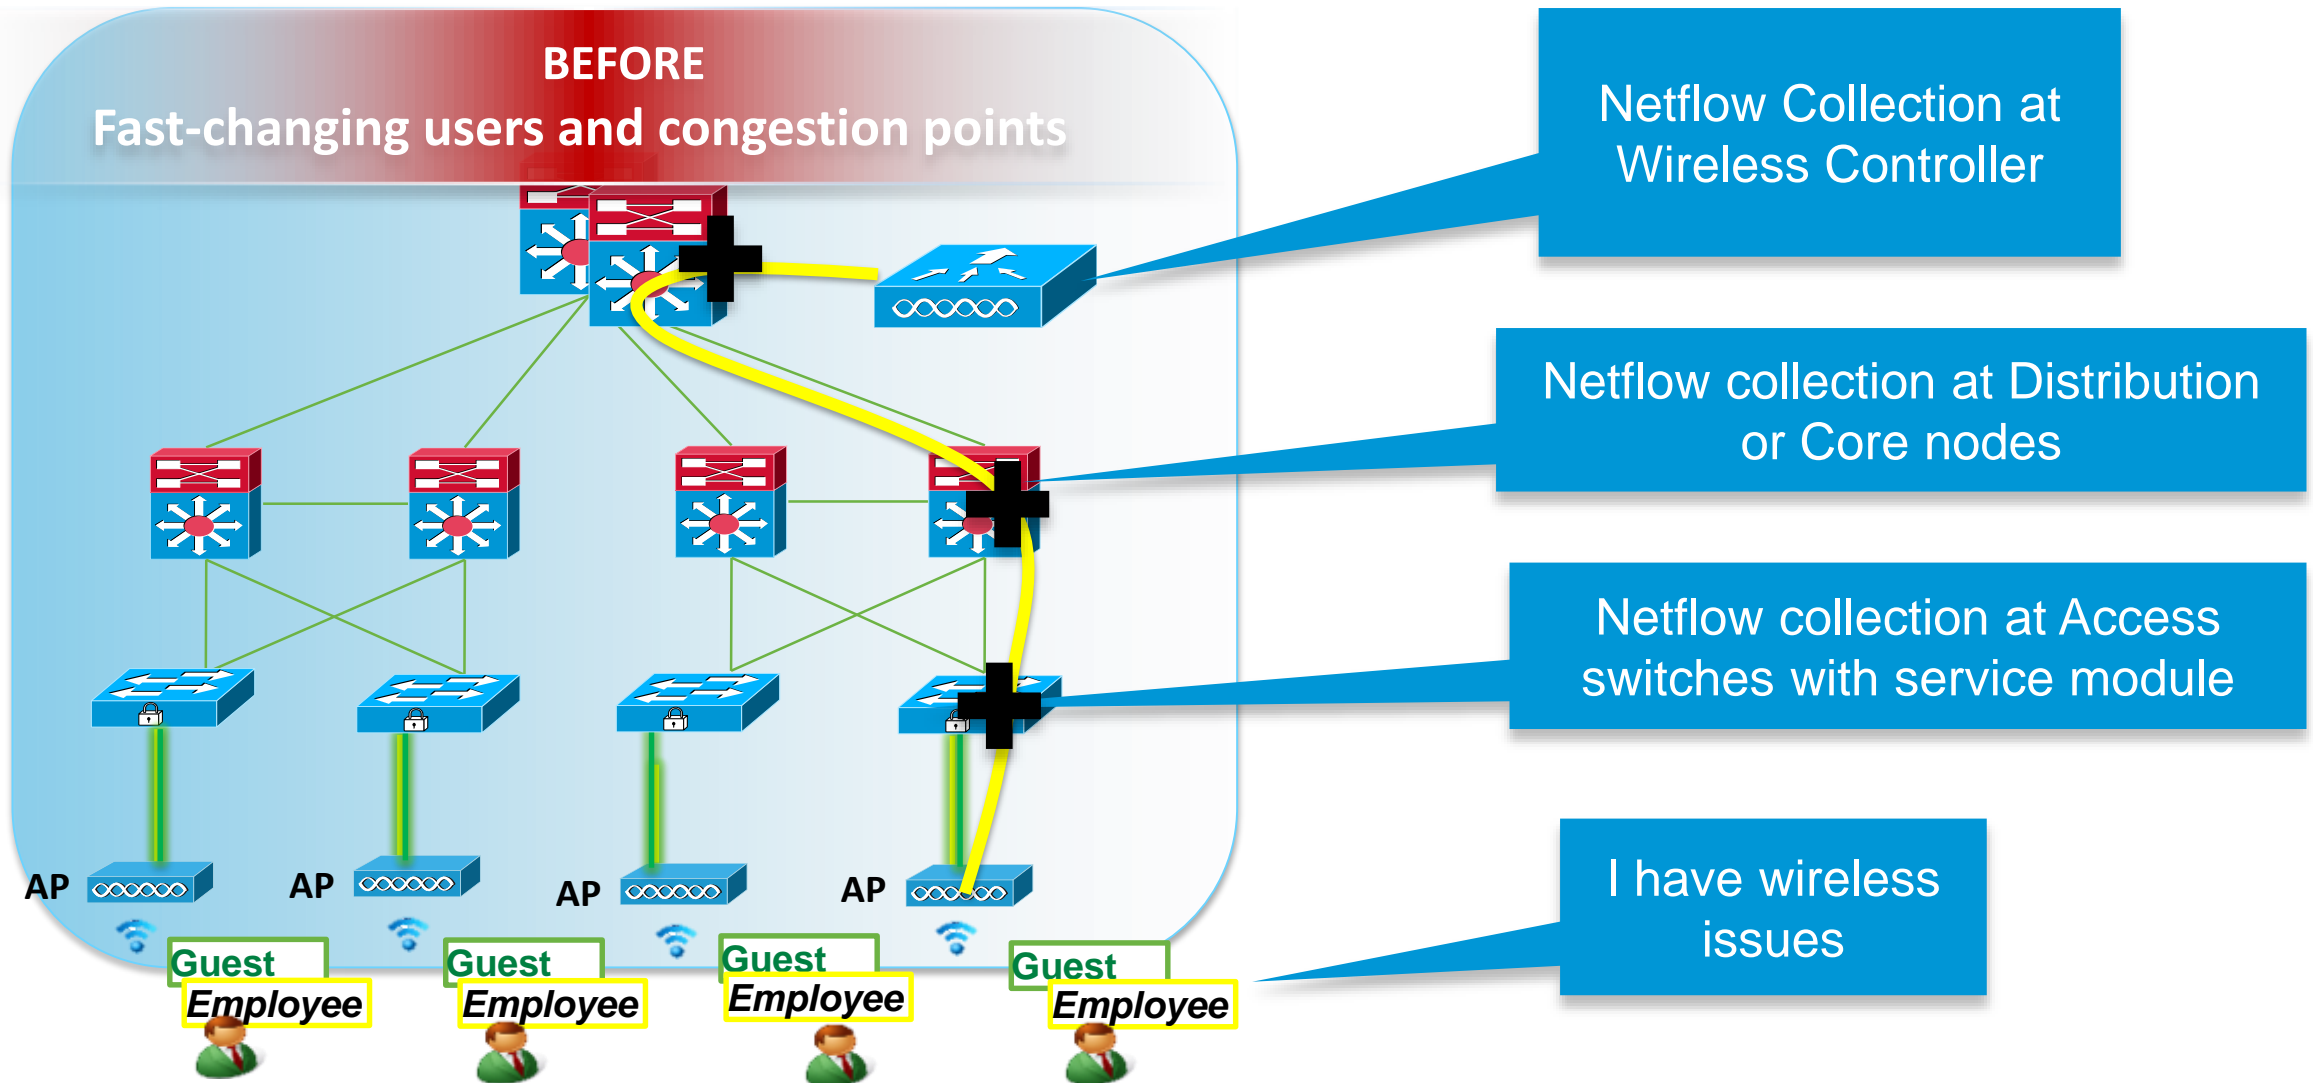


The wireless network is slow

My application is taking too long
to load

I have wireless issues

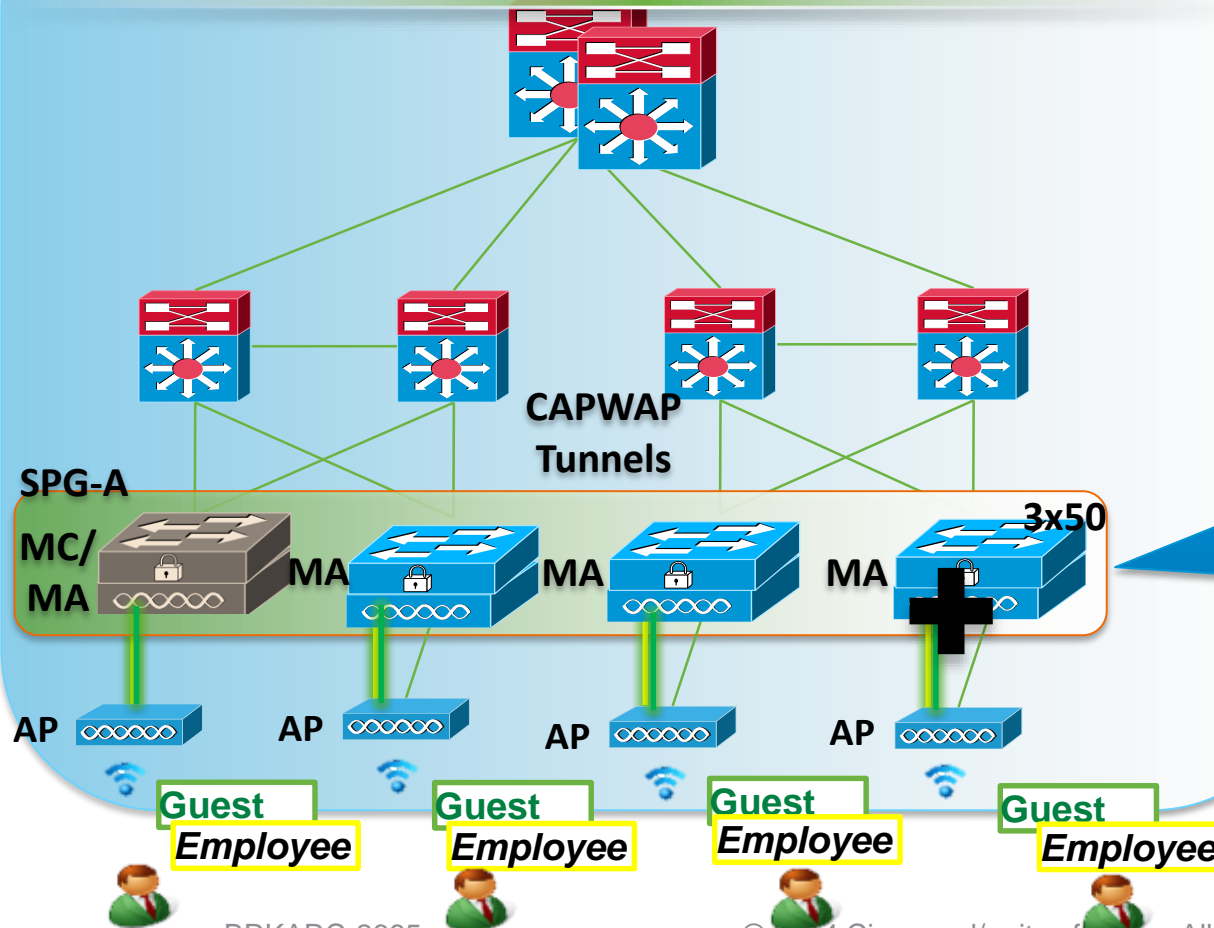
Troubleshooting with NetFlow – Before Converged Access



NetFlow Visibility – For Wired and Wireless Traffic

AFTER

Granular visibility and real-time feedback



Prime Assurance Manager



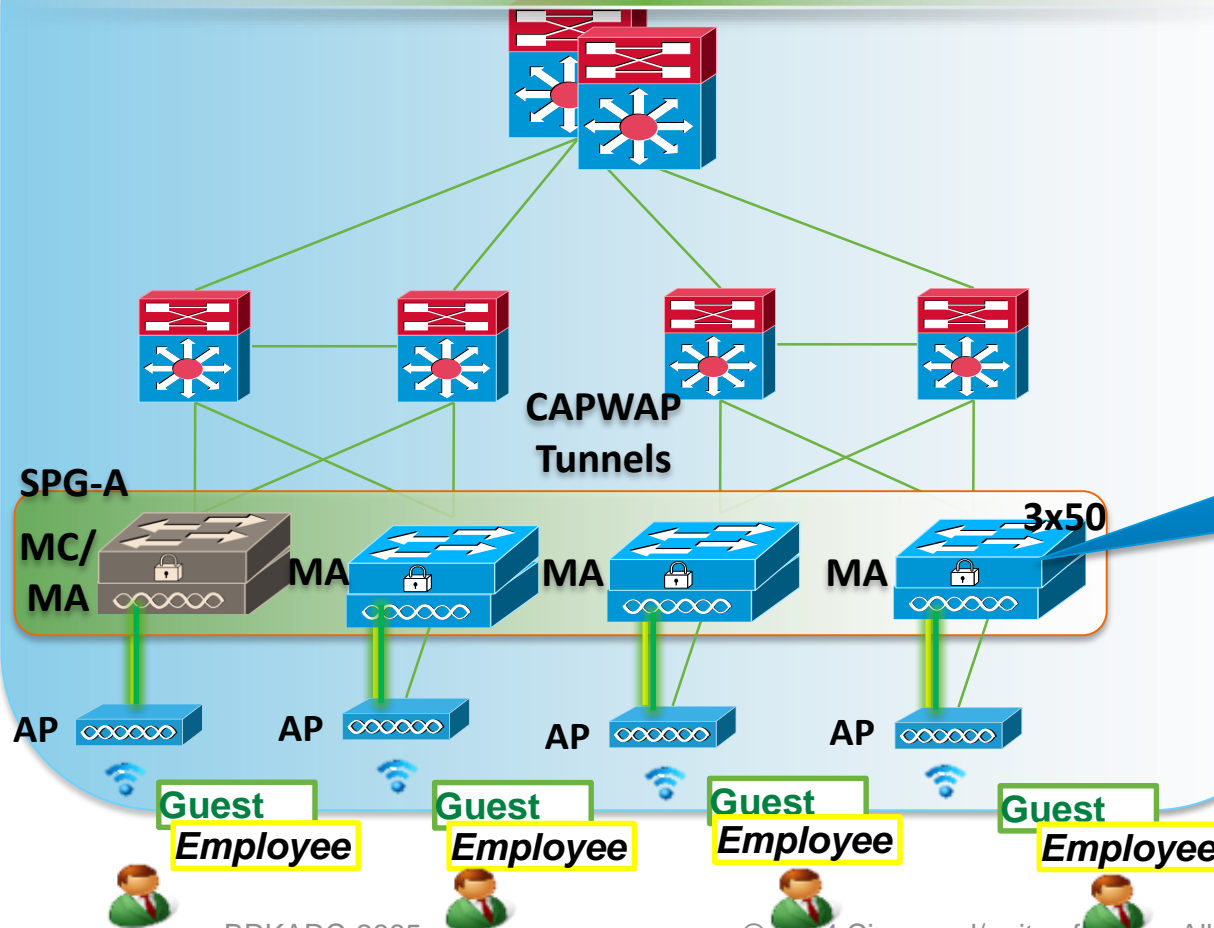
Real-time Monitoring

- Unprecedented visibility to wireless traffic
 - Client
 - AP
 - SSID view
- SSID VLAN Mapping
- Top talkers

I have wireless issues

AFTER

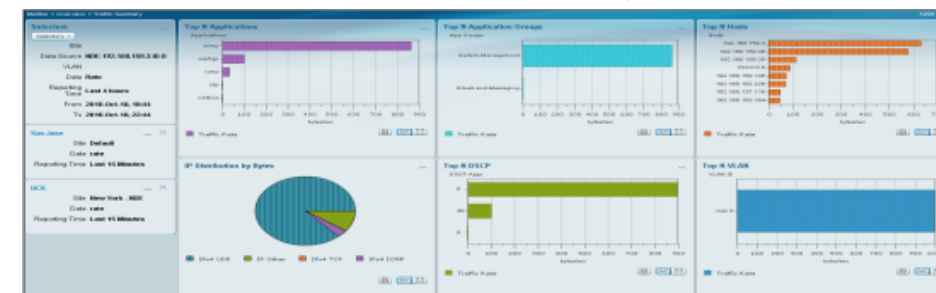
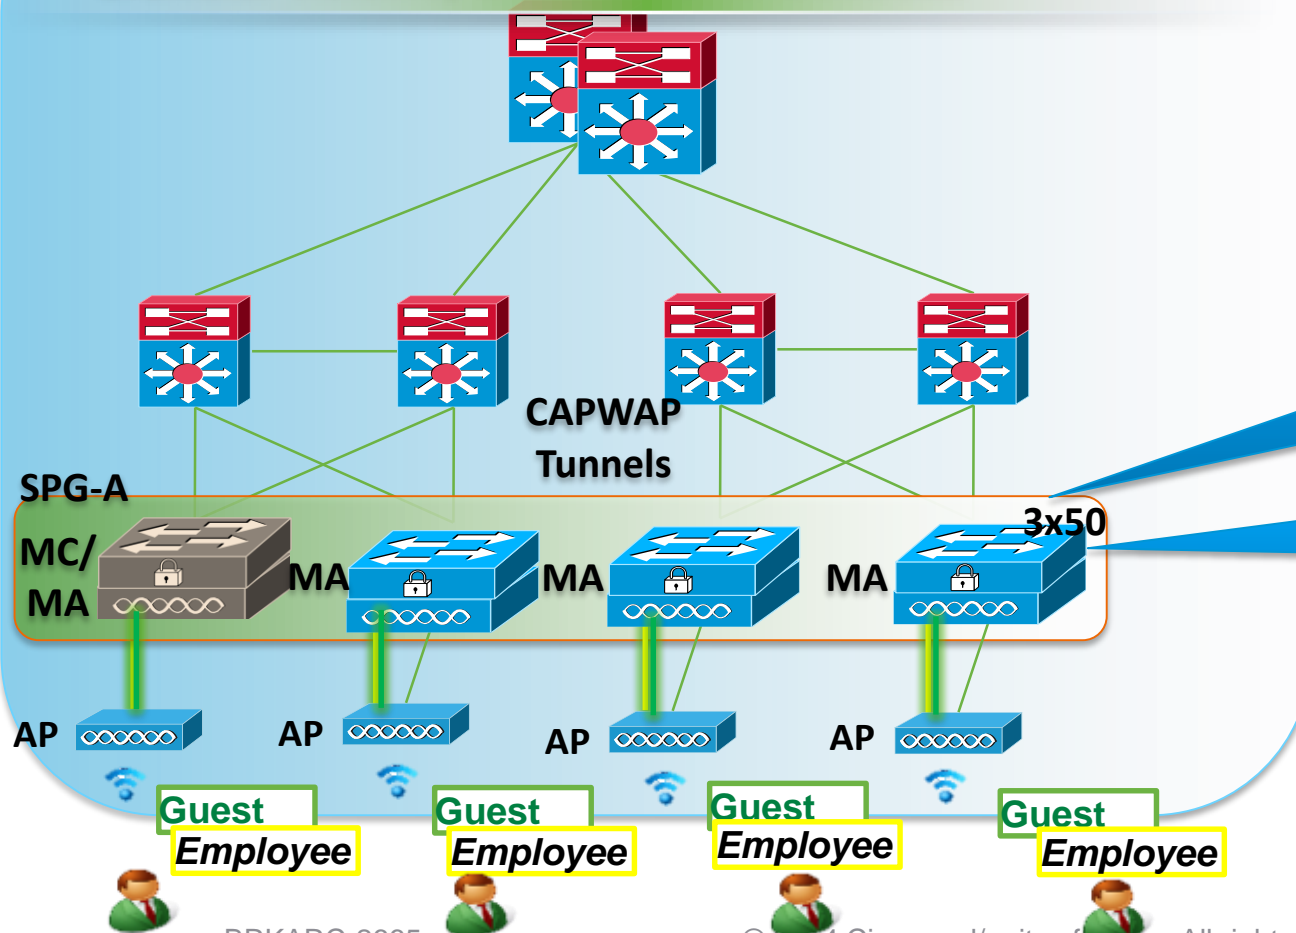
Granular visibility and real-time feedback



```
# flow exporter CISCO-PI- COLLECTOR
# description CISCO PRIME INFRA FNF COLLECTOR
# destination 10.100.1.82 transport udp 2055 source vlan 4093
```

AFTER

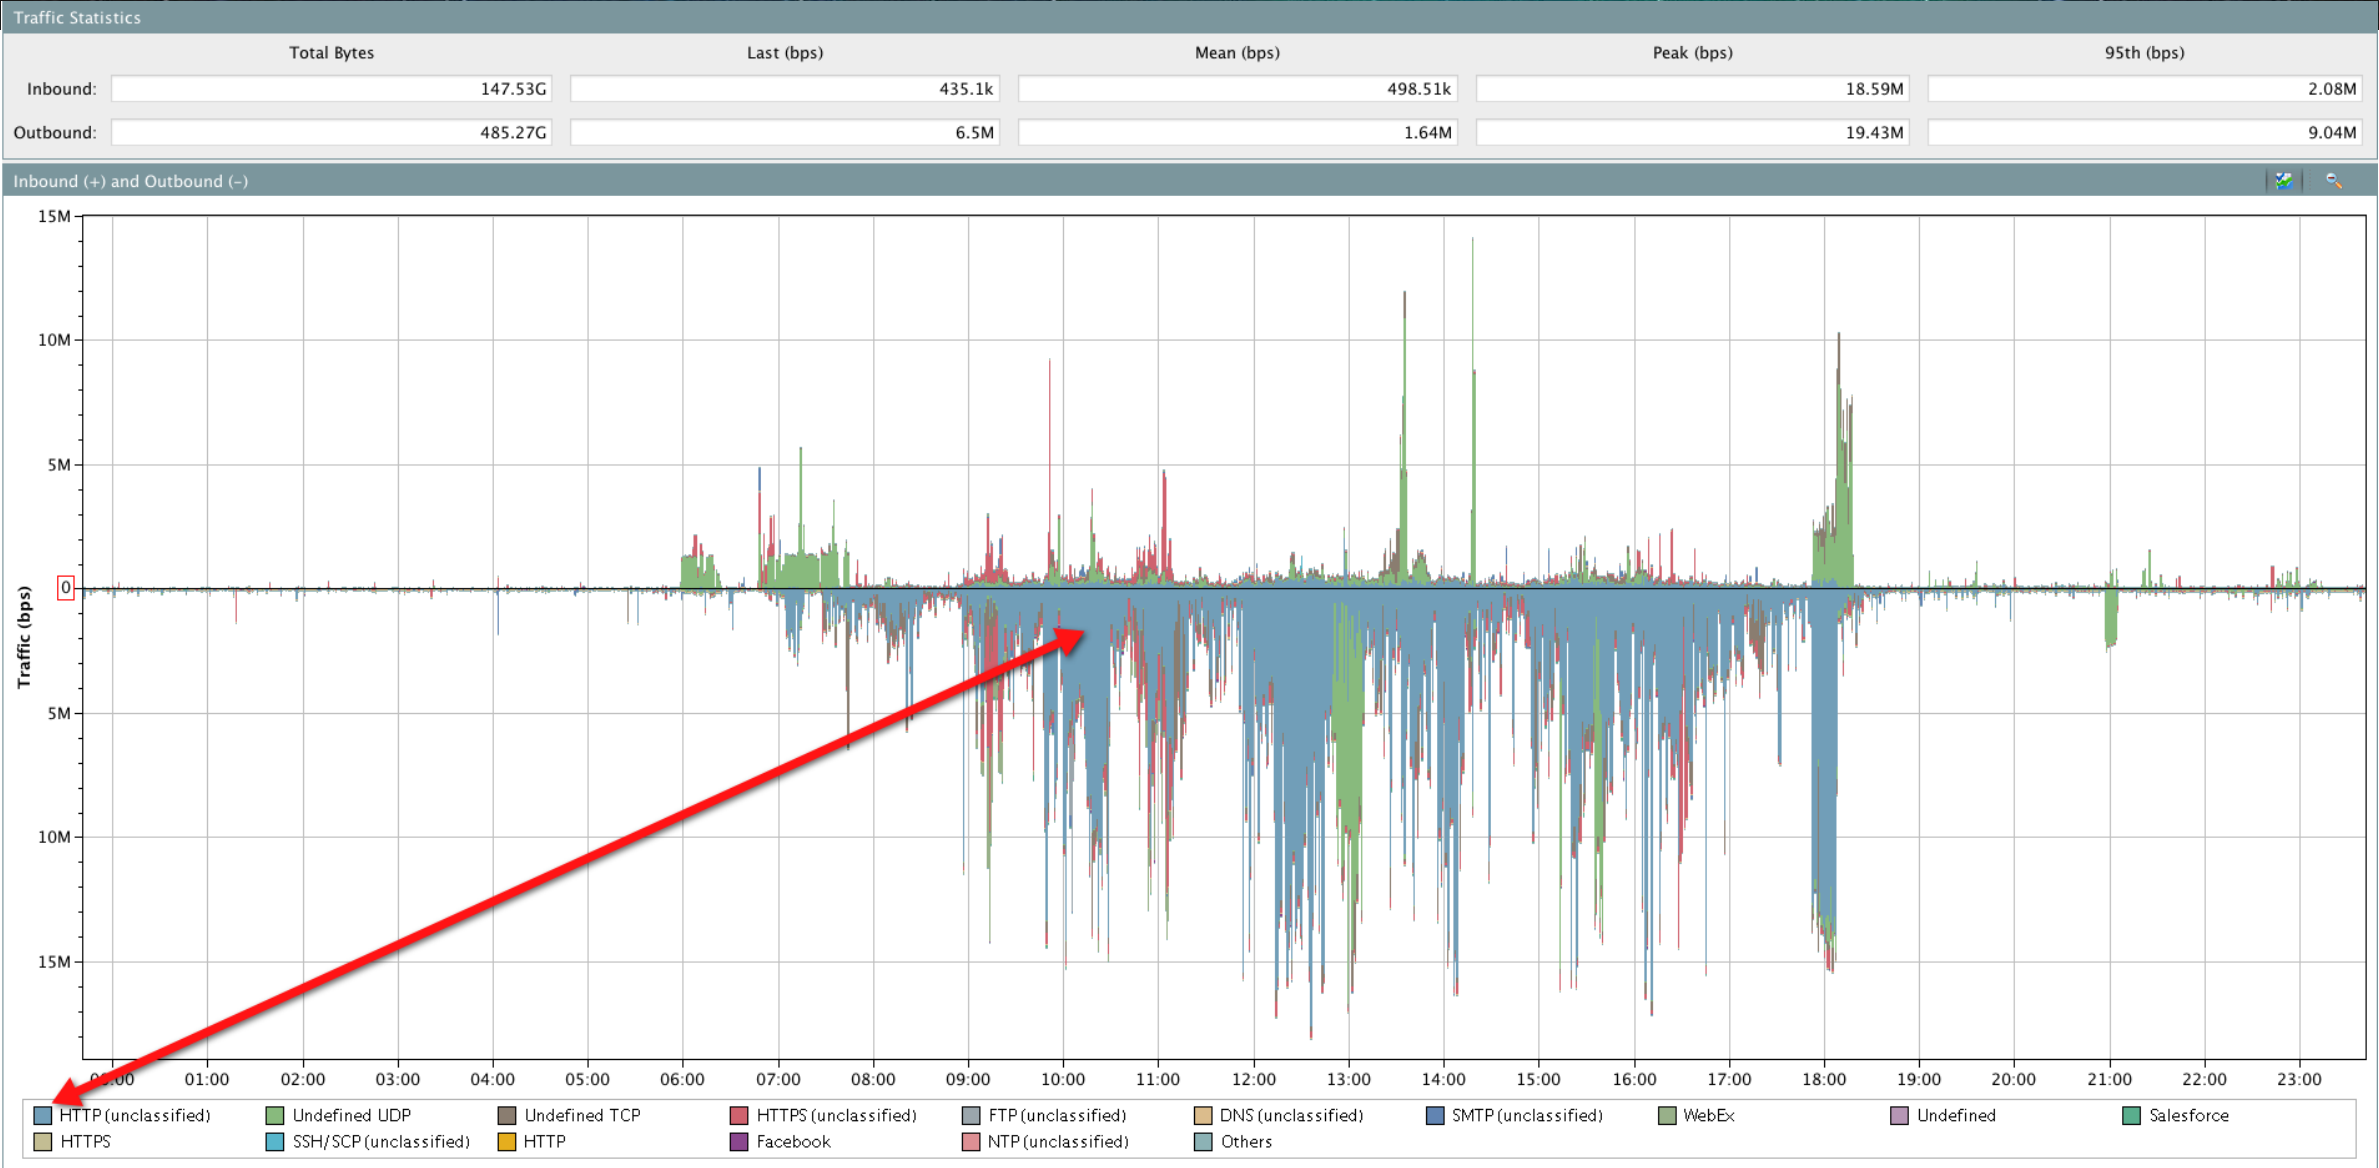
Granular visibility and real-time feedback



```
# flow monitor WIRED-PHONE-FNF- MONITOR
# record GLOBAL-FNF-POLICY
# exporter CISCO-PI-COLLECTOR
```

```
# flow monitor CA-WiFi-L3-SSID- FNF-MONITOR
# record GLOBAL-FNF-POLICY
# exporter CISCO-PI-COLLECTOR
```



Capacity Planning – NetFlow-based



Comprehensive Visibility – NetFlow-based

Host	Host Role	Peer	Port	Bytes
spyglass.lancope.com (209.182.184.2)	Client	vip1.g-anycast1.cachefly.net (205.234.175.175)	80/tcp (http)	76.73M
spyglass.lancope.com (209.182.184.2)	Client	mediaserver-sv5-t1-2.pandora.com (208.85.42.22)	80/tcp (http)	75.73M
spyglass.lancope.com (209.182.184.2)	Client	ragana.canonical.com (91.189.91.13)	80/tcp (http)	73.18M
spyglass.lancope.com (209.182.184.2)	Client	s3-1.amazonaws.com (207.171.163.151)	80/tcp (http)	69.94M
spyglass.lancope.com (209.182.184.2)	Client	mediaserver-dc6-t1-3.pandora.com (208.85.46.23)	80/tcp (http)	69.01M
spyglass.lancope.com (209.182.184.2)	Client	mediaserver-sv5-t1-3.pandora.com (208.85.42.33)	80/tcp (http)	62.83M
spyglass.lancope.com (209.182.184.2)	Client	mediaserver-sjl-t1-2.pandora.com (208.85.41.12)	80/tcp (http)	HTTP (unclas...
spyglass.lancope.com (209.182.184.2)	Client	65.121.209.25	80/tcp (http)	
spyglass.lancope.com (209.182.184.2)	Client	91.197.45.9	80/tcp (http)	
spyglass.lancope.com (209.182.184.2)	Client	mediaserver-sjl-t1-1.pandora.com (208.85.41.11)	80/tcp (http)	
spyglass.lancope.com (209.182.184.2)	Client	mediaserver-sv5-t1-1.pandora.com (208.85.42.21)	80/tcp (http)	
spyglass.lancope.com (209.182.184.2)	Client	cds56.mia9.msecn.net (65.54.93.59)	80/tcp (http)	
spyglass.lancope.com (209.182.184.2)	Client	cds115.mia9.msecn.net (65.54.93.118)	80/tcp (http)	


Lancope
Network Performance + Security Monitoring™



NetFlow –


Evolution for Wired and Wireless

NetFlow Today



- Used in Distribution/WAN
- Use Cases: Capacity Planning, Application Visibility

NetFlow in Wired Access



- End-to-End Visibility: L2, East-West Traffic, App ID, **UserID**, **PoE**, **Device Type***
- Use Case: Collaboration, Security, Capacity Planning

NetFlow in Converged Access



- Converged Access Visibility: **SSID**, **AP Name**, **Client ID**, **Device Type***
- Use Case: BYOD, Mobility, Collaboration, Security, Capacity Planning

* Some enhancements on roadmap
BRKARC-2665

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –



- IP Addressing

- Design Options

- Deployment Examples

Summary

Converged Access – IP Addressing – Options

Multiple options exist for how to assign user subnets in Converged Access

Several possible IP addressing deployment models exist for wired / wireless use ...

- Option 1** – Separate wired and wireless VLANs, per wiring closet
- Option 2** – Merged wired and wireless VLANs, per wiring closet
- Option 3** – Separate wired VLANs per wiring closet, spanned wireless VLAN across multiple wiring closets (below a single distribution)

There are trade-offs between each of these IP addressing design models

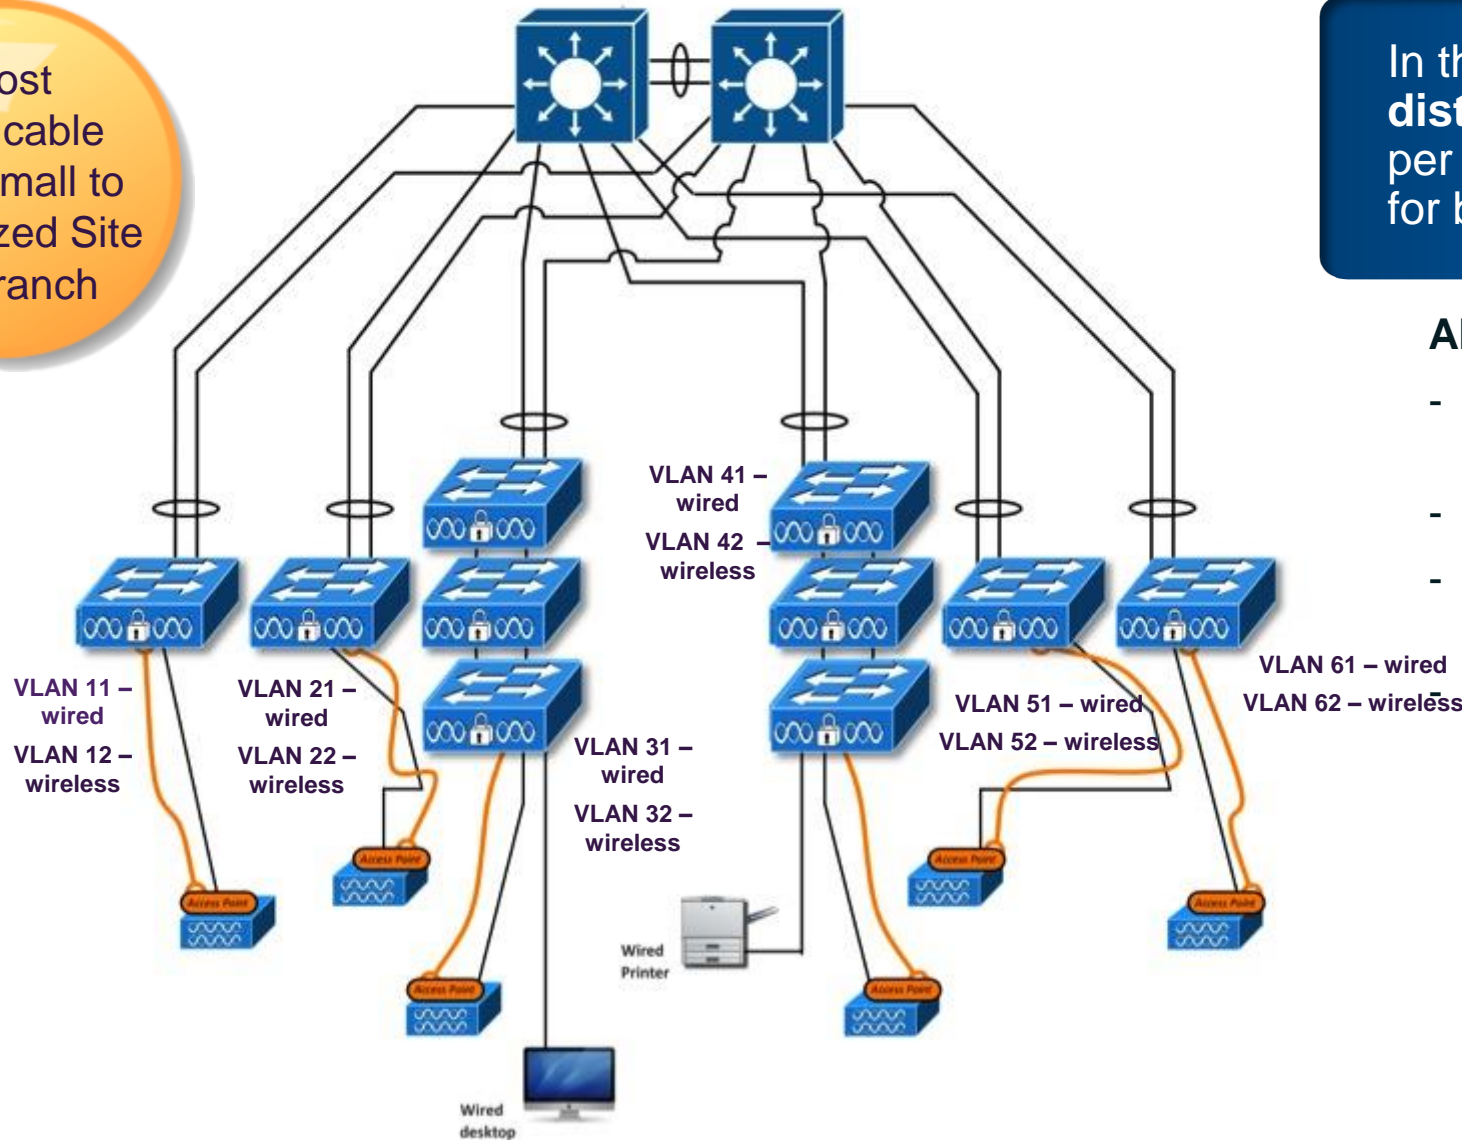
All of the models presented here have been tested and validated.

This is a customer / partner design choice based on current deployment and requirements.

Converged Access – IP Addressing – Option 1

OPTION 1 – Separate VLANs / subnets per wiring closet, for wired and wireless

Most
Applicable
to a Small to
Mid-Sized Site
or Branch



In this design option, **separate and distinct subnets** are configured per Converged Access wiring closet, for both wired and wireless users

ADVANTAGES –

- Easy to understand – maps well to user expectations for wired design
- Can match any wired deployment (L2 / L3)
- Can create separate wired and wireless policies based on VLAN

Eliminates DHCP contention wired/wireless


CONSIDERATIONS –

- May lead to more subnets required
- May be hard to size wireless subnets for number of anticipated wireless clients, per wiring closet (may lead to wasted IP address space for wireless use, potentially)

Converged Access – IP Addressing – Option 2

OPTION 2 – Merged VLANs / subnets per wiring closet, for wired and wireless

Most
Applicable
to a Small to
Mid-Sized Site
or Branch



In this design option, wired and wireless users and devices **share common subnets** per Converged Access wiring closet (i.e. one or more wired / wireless VLANs per wiring closet)

ADVANTAGES –

- Leads to fewer subnets req'd vs. Option 1


CONSIDERATIONS –

- Potential dual-attached device issues (possible client-side bridging issues)
- No longer possible to apply separate per-VLAN policies for wired / wireless
- May be hard to size combined subnets appropriately for number of wired / wireless clients, per wiring closet (may be slightly more efficient vs. Option 1)
- Possible DHCP contention, wired / wireless

Converged Access – IP Addressing – Option 3

OPTION 3 – Separate wired VLANs / subnets per wiring closet, with wireless VLAN spanned

Most Applicable to a Larger Site or Campus



In this design option, **separate and distinct subnets** per Converged Access wiring closet for both wired and wireless users, with wireless spanned below the distribution layer

ADVANTAGES –

- Can create separate wired and wireless policies based on VLAN
- Leads to fewer subnets req'd vs. Option 1 (only one wireless subnet below dist.)
- Easier to size wireless subnet(s) below distribution layer (closer correspondence to IP addressing in the CUWN model)

CONSIDERATIONS –

- Optimised with VSS, or other similar single-switch-equivalent model, at distribution (to avoid L2 loops)
- Topology differs, wired vs. wireless

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –

- IP Addressing

- ▶ Design Options

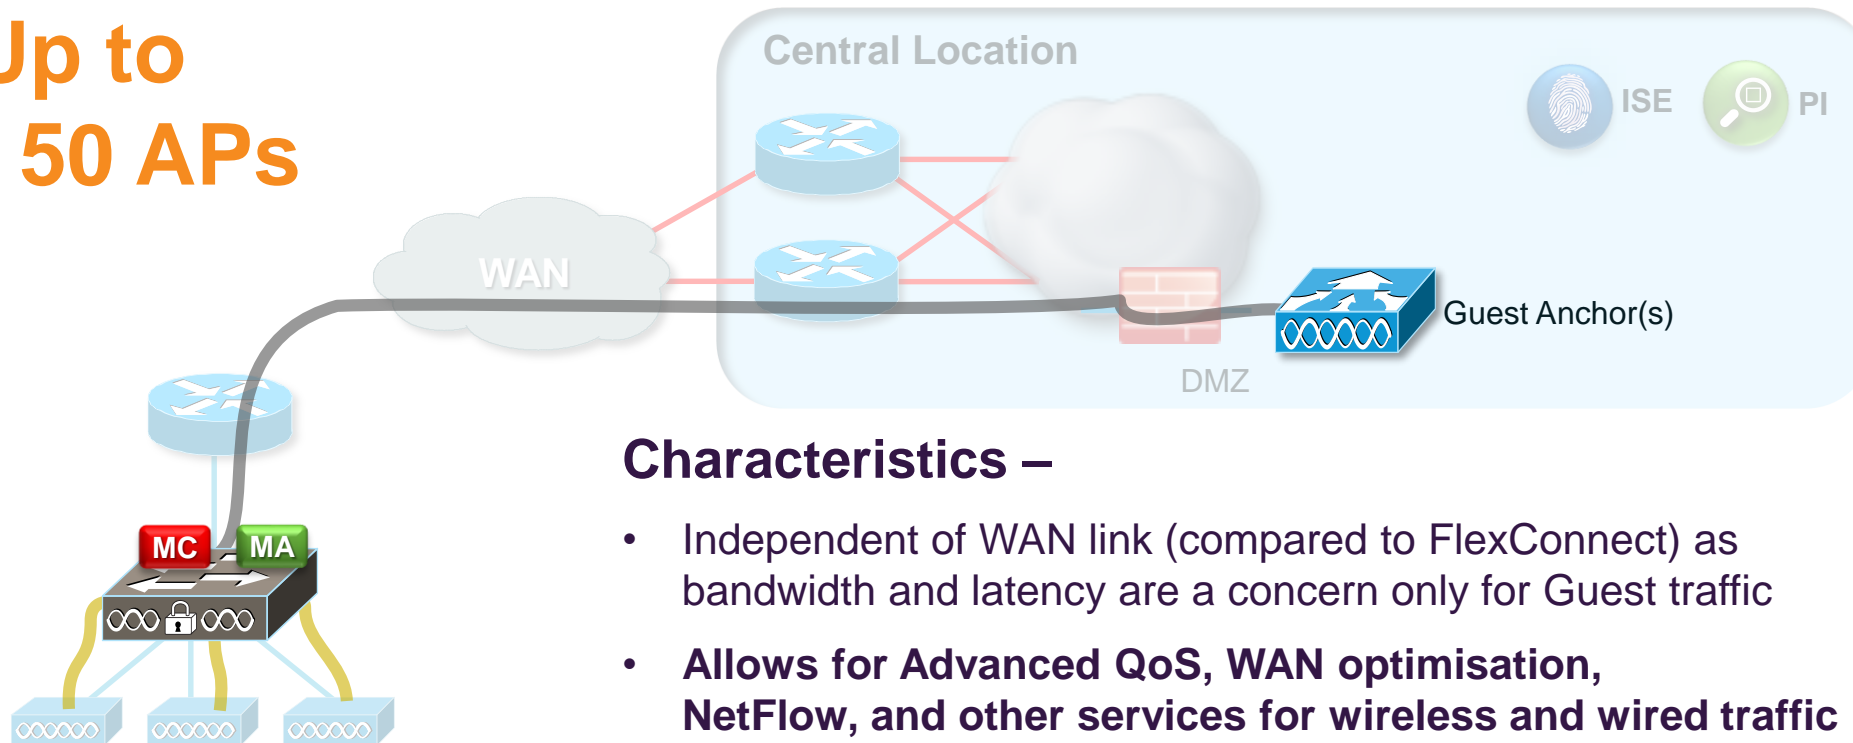
- Deployment Examples

Summary

Converged Access – Small Branch

No Discrete Controllers, Catalyst 3x50s as MCs / MAs

Up to
25 / 50 APs



Applicable
to a Small
Branch
Deployment

Characteristics –

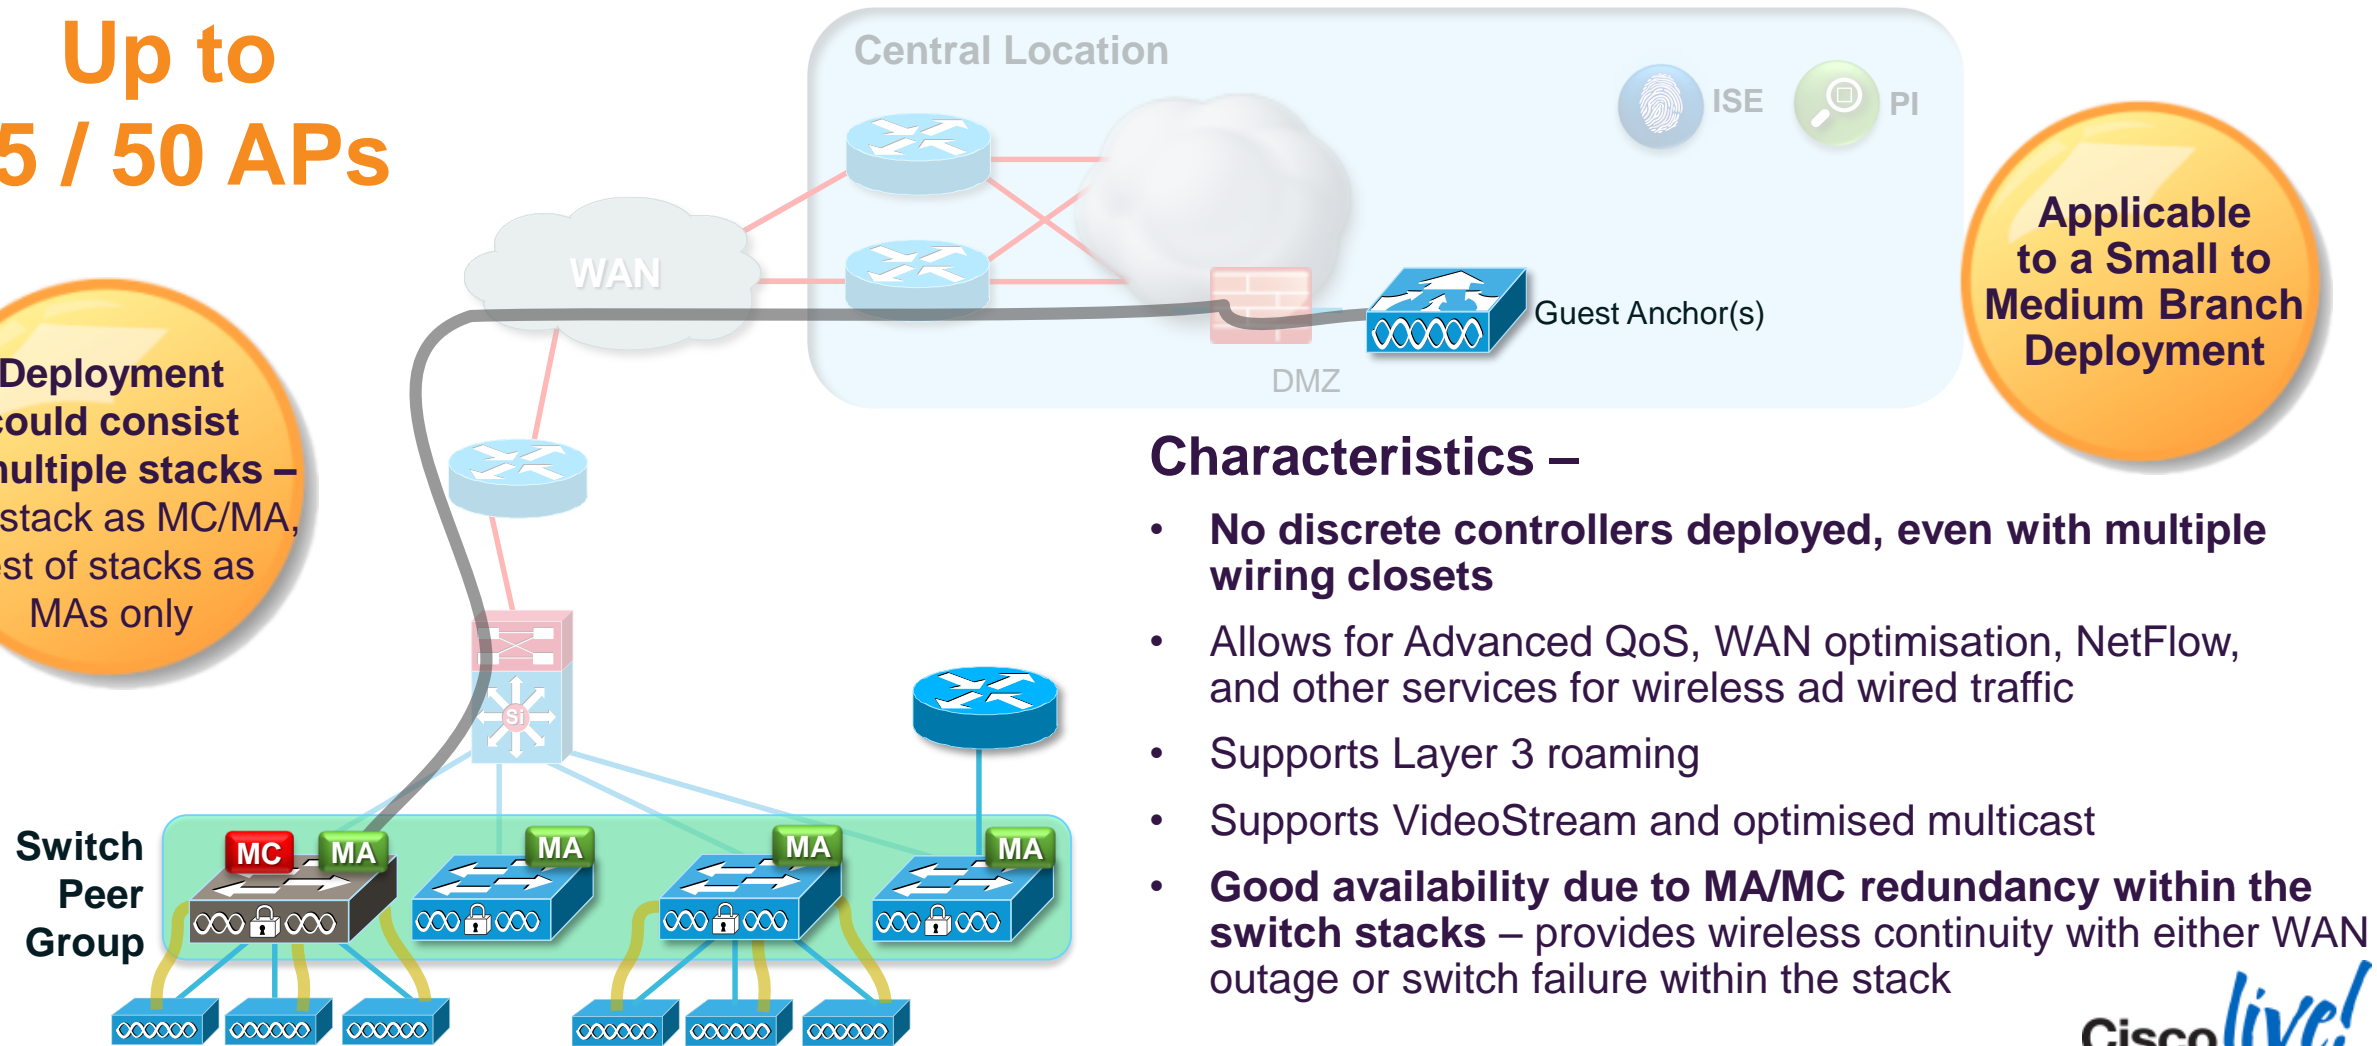
- Independent of WAN link (compared to FlexConnect) as bandwidth and latency are a concern only for Guest traffic
- **Allows for Advanced QoS, WAN optimisation, NetFlow, and other services for wireless and wired traffic**
- **Supports Layer 3 roaming**
- **Supports VideoStream and optimised multicast**
- **Good availability due to MA/MC redundancy within the 3x50 stack** – provides wireless continuity with either WAN outage or switch failure within the stack

Converged Access – Small / Medium Branch

No Discrete Controllers, Catalyst 3x50s as MCs / MAs, Single SPG

Up to
25 / 50 APs

Deployment could consist of multiple stacks – one stack as MC/MA, rest of stacks as MAs only



Applicable to a Small to Medium Branch Deployment

Characteristics –

- No discrete controllers deployed, even with multiple wiring closets
- Allows for Advanced QoS, WAN optimisation, NetFlow, and other services for wireless and wired traffic
- Supports Layer 3 roaming
- Supports VideoStream and optimised multicast
- **Good availability due to MA/MC redundancy within the switch stacks** – provides wireless continuity with either WAN outage or switch failure within the stack

Converged Access – Large Branch

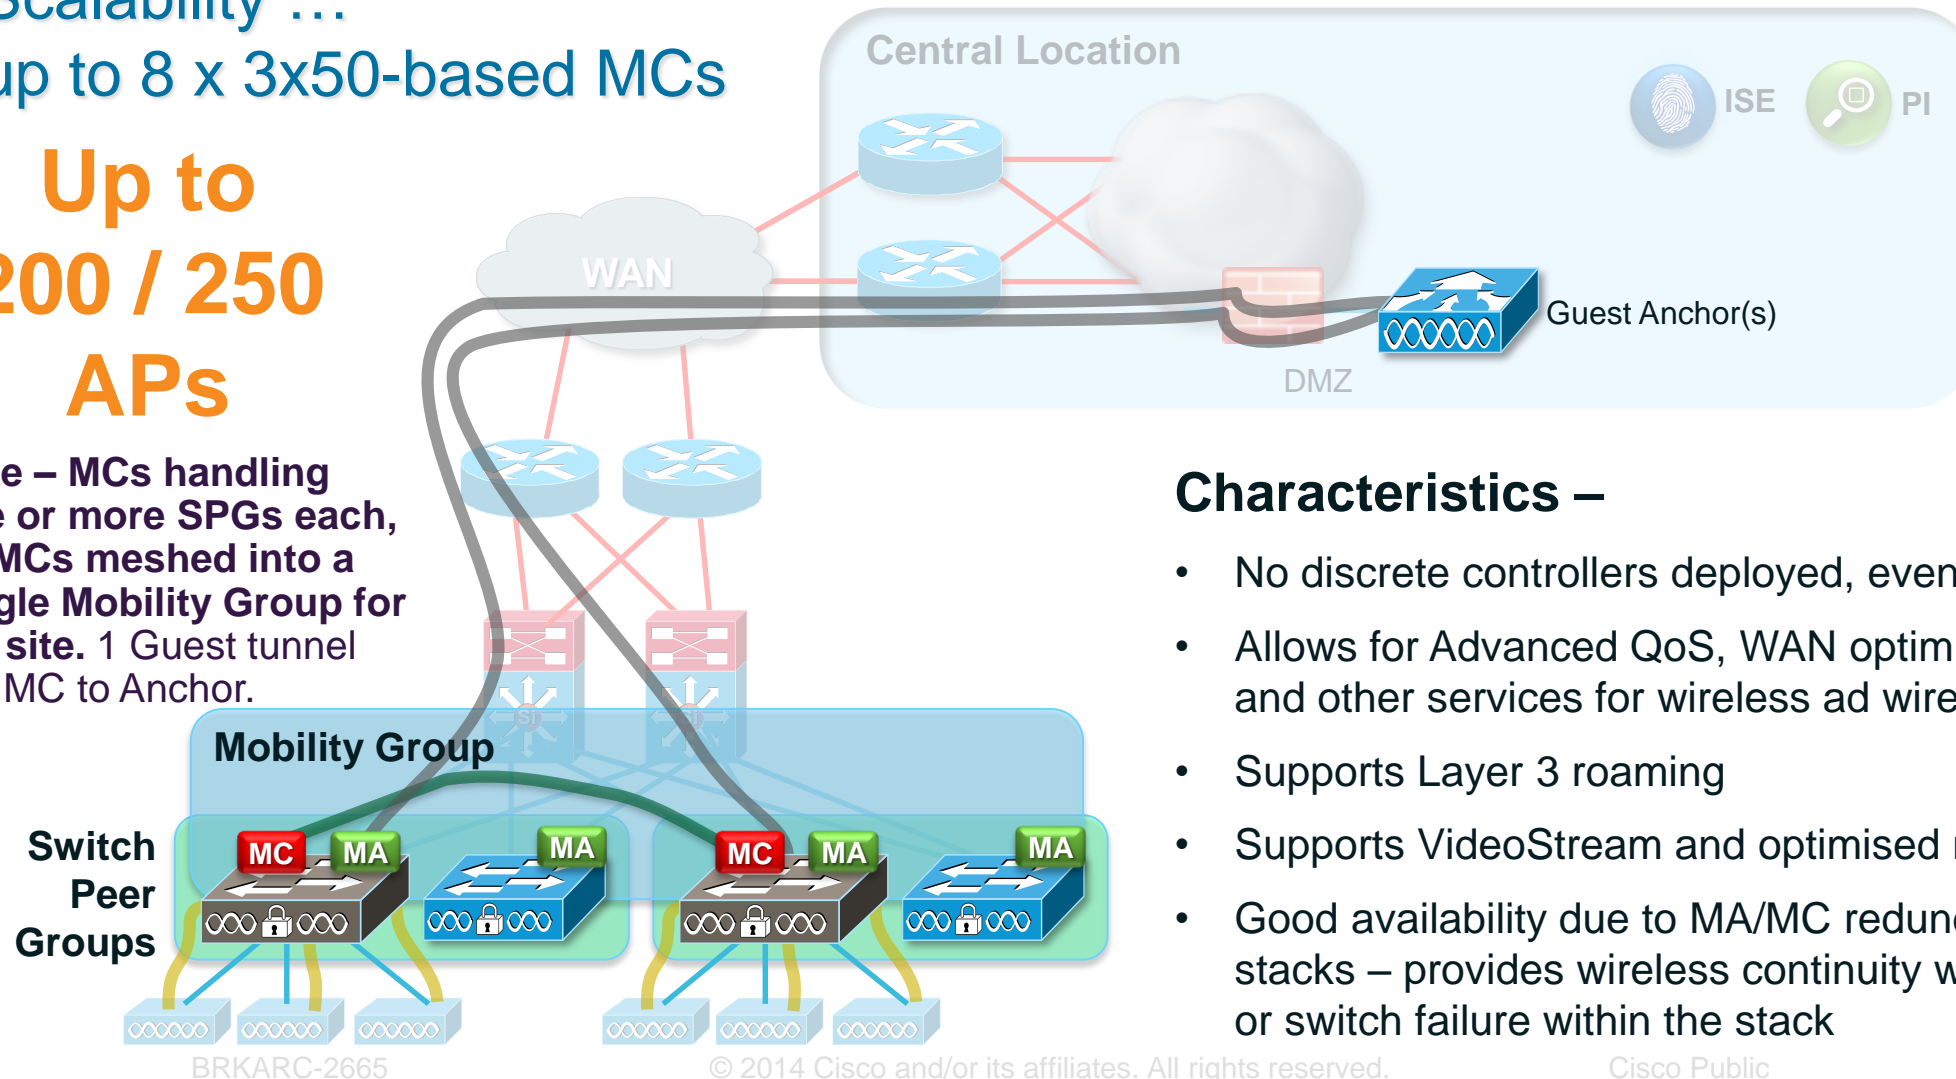
No Discrete Controllers, Catalyst 3x50s as MCs / MAs, Multiple SPGs

Scalability ...

up to 8 x 3x50-based MCs

**Up to
200 / 250
APs**

Note – MCs handling one or more SPGs each, all MCs meshed into a single Mobility Group for the site. 1 Guest tunnel per MC to Anchor.



**Applicable
to a Larger
Branch
Deployment**

Characteristics –

- No discrete controllers deployed, even at a larger branch
- Allows for Advanced QoS, WAN optimisation, NetFlow, and other services for wireless and wired traffic
- Supports Layer 3 roaming
- Supports VideoStream and optimised multicast
- Good availability due to MA/MC redundancy within the switch stacks – provides wireless continuity with either WAN outage or switch failure within the stack

Converged Access – Small Campus

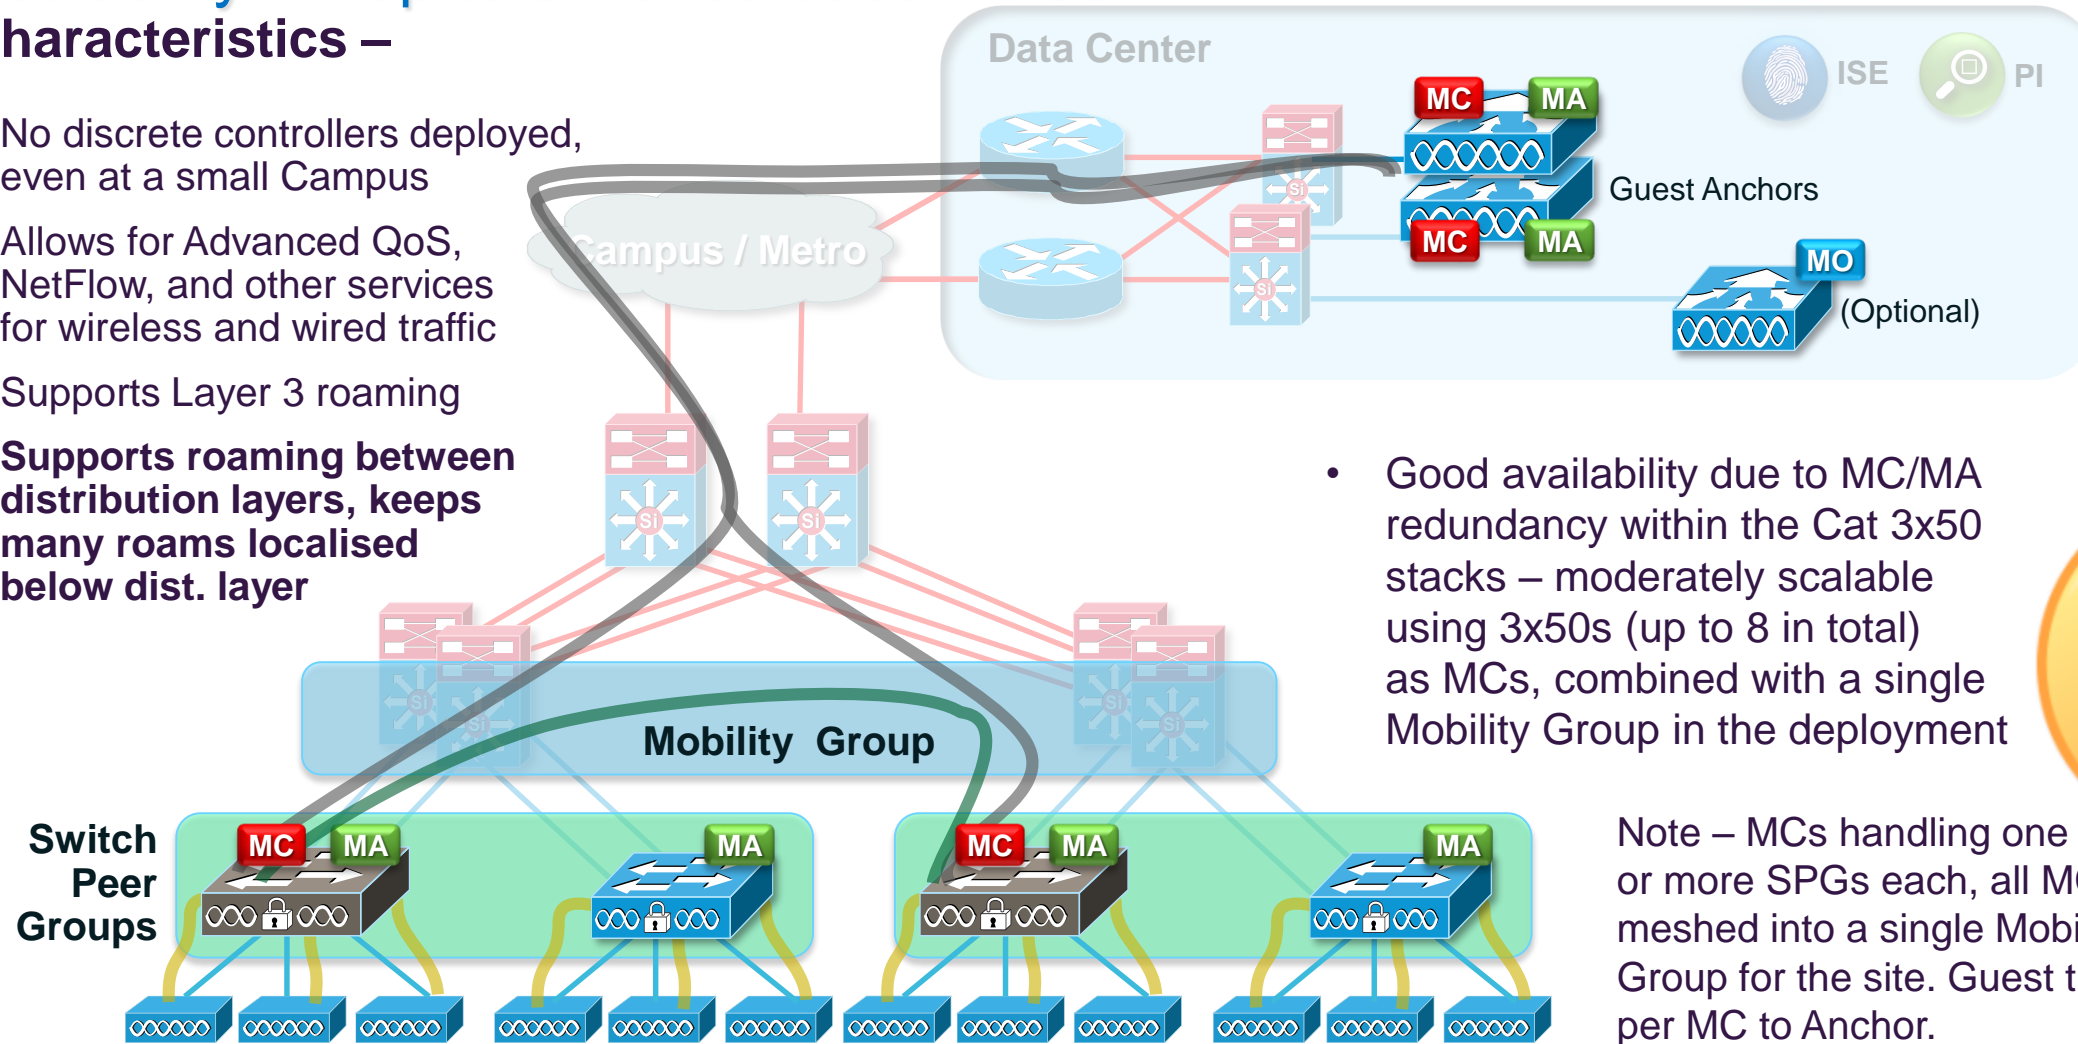
Catalyst 3x50s as MC s / MAs, Multiple SPGs

Up to 200 / 250
APs

Scalability ... up to 8 x 3x50-based MCs

Characteristics –

- No discrete controllers deployed, even at a small Campus
- Allows for Advanced QoS, NetFlow, and other services for wireless and wired traffic
- Supports Layer 3 roaming
- **Supports roaming between distribution layers, keeps many roams localised below dist. layer**



- Good availability due to MC/MA redundancy within the Cat 3x50 stacks – moderately scalable using 3x50s (up to 8 in total) as MCs, combined with a single Mobility Group in the deployment

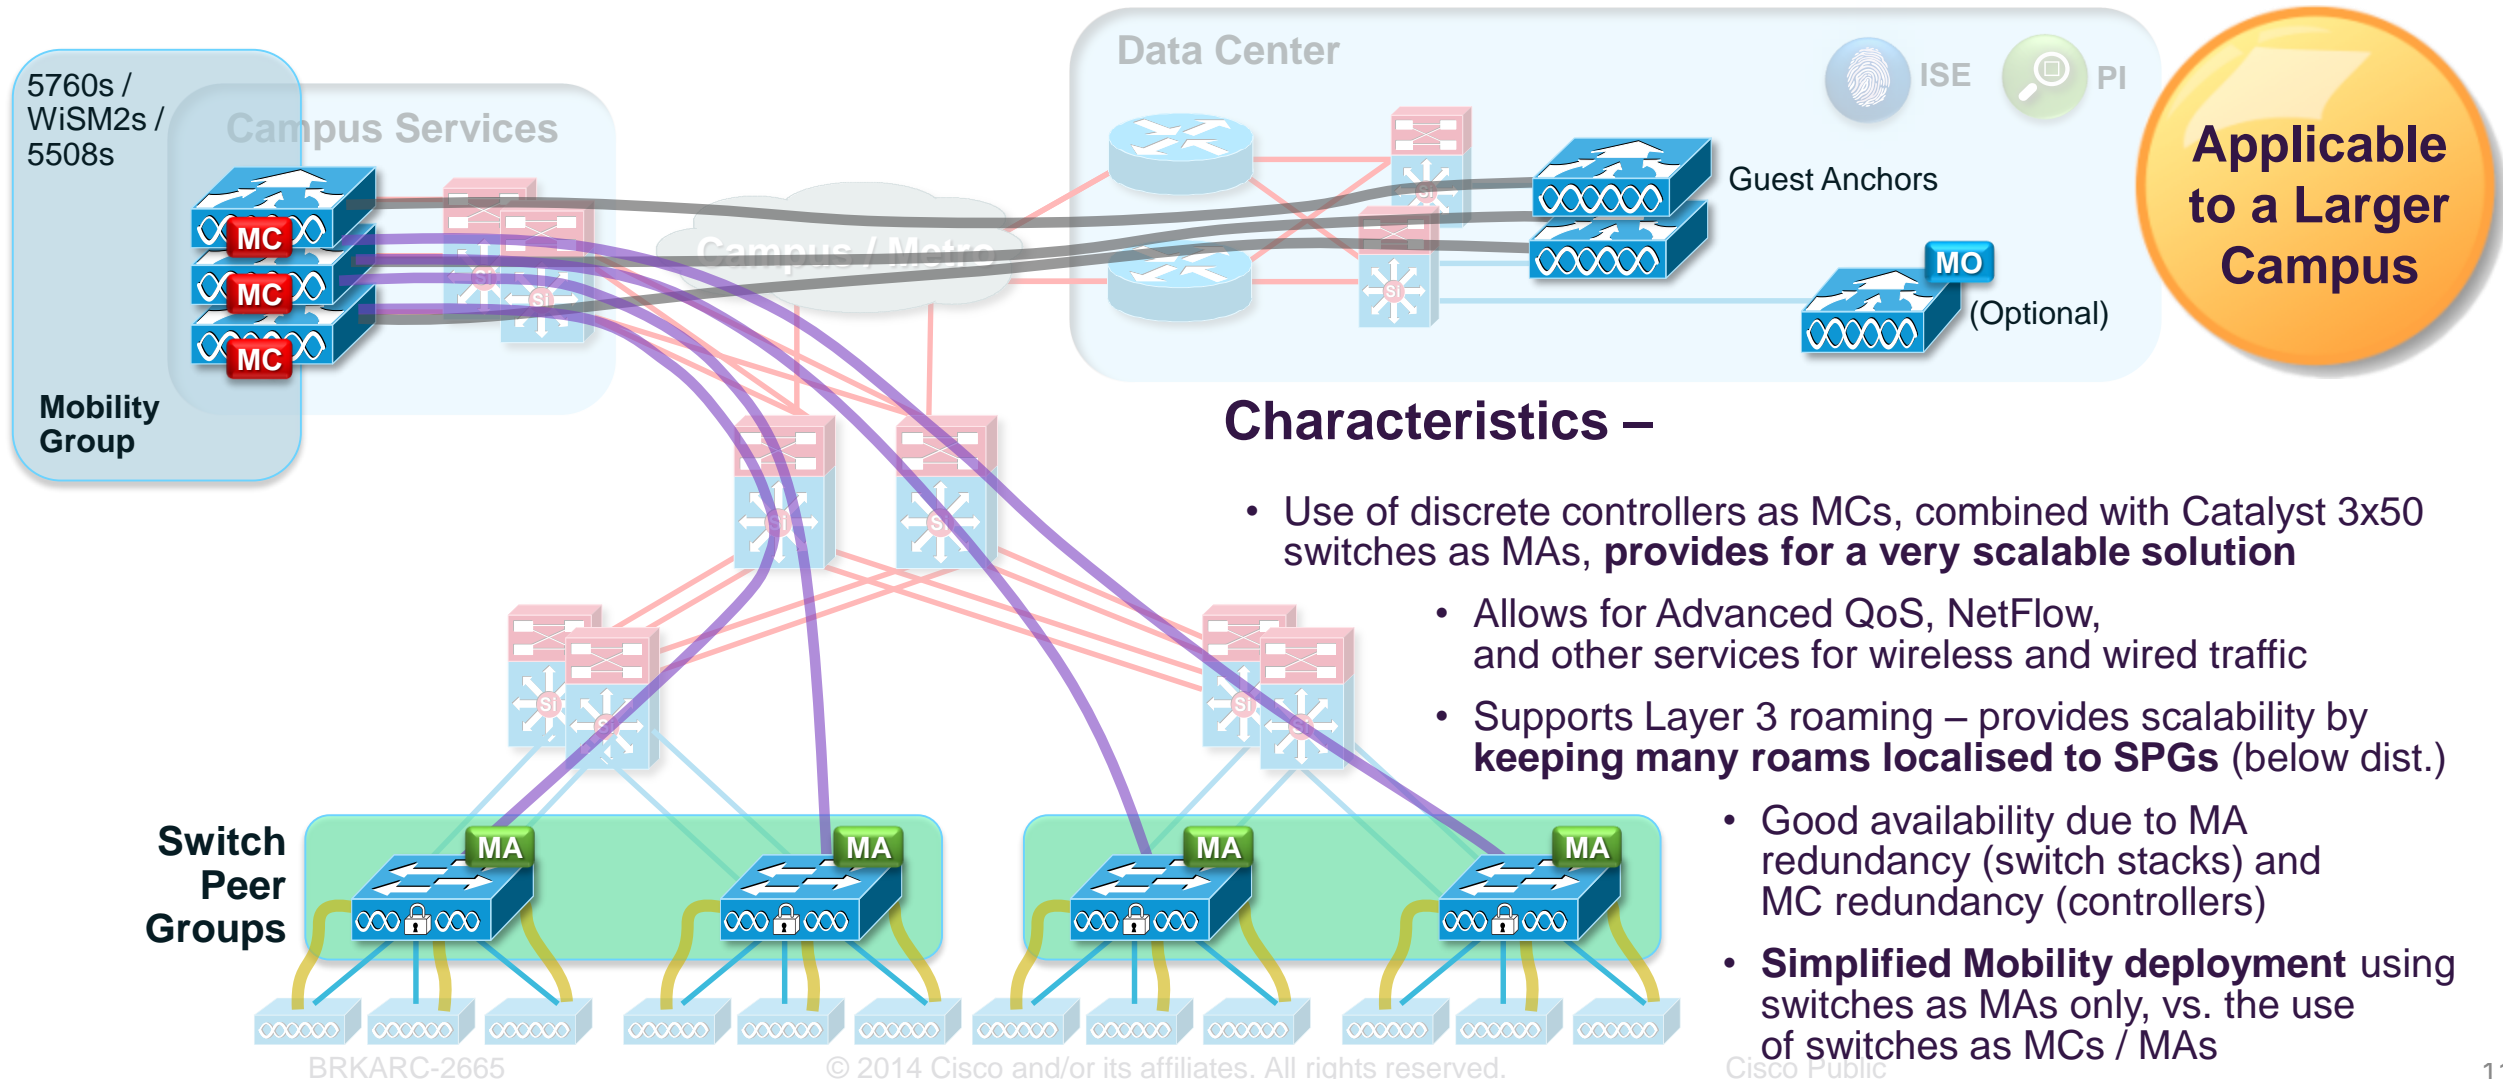
Applicable
to a Small
Campus
Deployment

Note – MCs handling one or more SPGs each, all MCs meshed into a single Mobility Group for the site. Guest tunnel per MC to Anchor.

Converged Access – Large Campus

Centralised Controllers as MCs, 3x50s as MAs Only

>250 APs



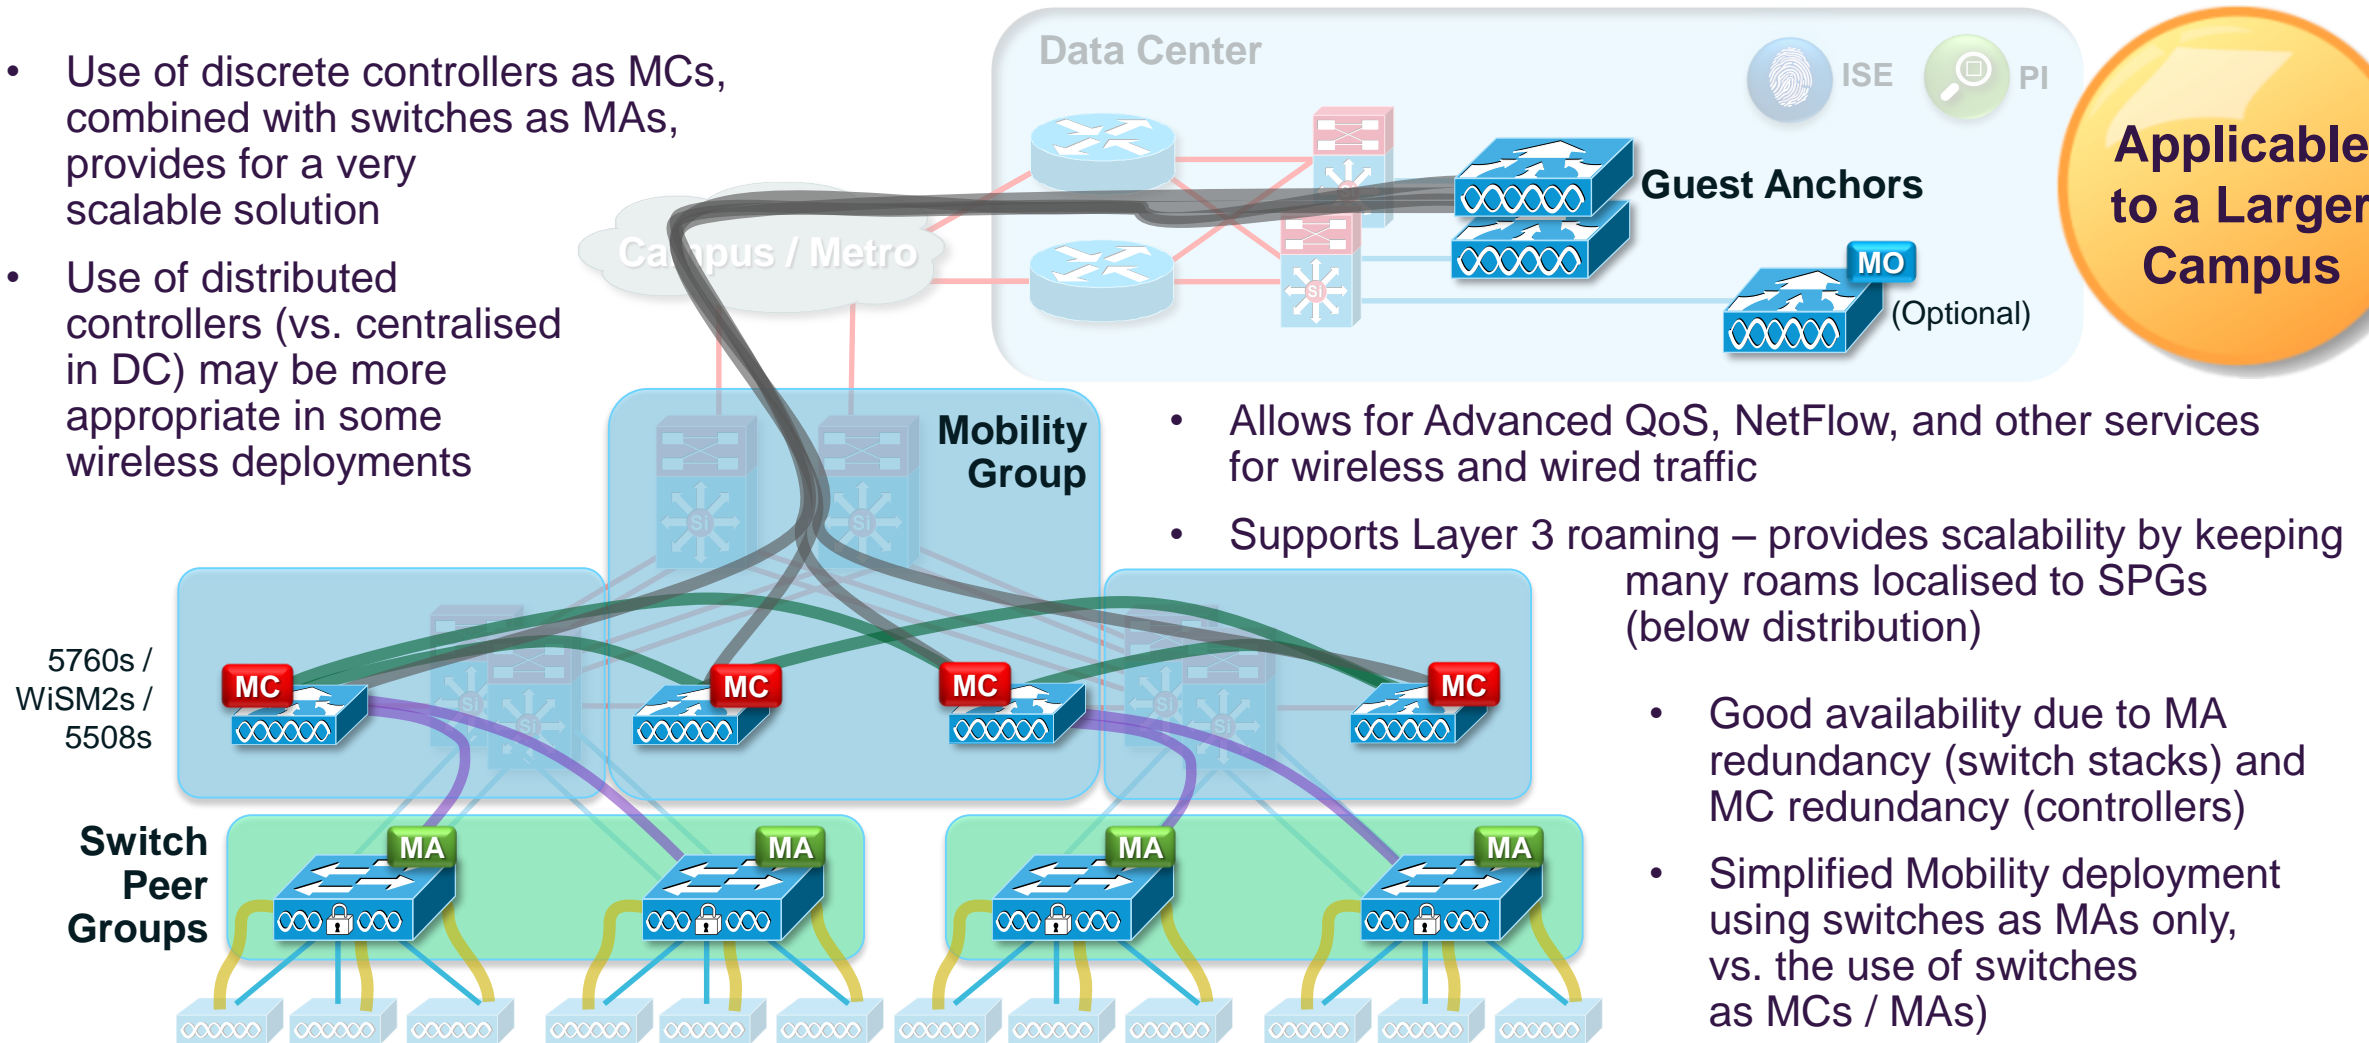
Converged Access – Large Campus

Distributed Controllers as MCs, 3x50s as MAs Only

>250 APs

Characteristics –

- Use of discrete controllers as MCs, combined with switches as MAs, provides for a very scalable solution
- Use of distributed controllers (vs. centralised in DC) may be more appropriate in some wireless deployments



- Allows for Advanced QoS, NetFlow, and other services for wireless and wired traffic
- Supports Layer 3 roaming – provides scalability by keeping many rooms localised to SPGs (below distribution)

- Good availability due to MA redundancy (switch stacks) and MC redundancy (controllers)
- Simplified Mobility deployment using switches as MAs only, vs. the use of switches as MCs / MAs)

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast

- NetFlow

Converged Access Design and Deployment –

- IP Addressing

- Design Options

▶ **Deployment Examples**

Summary

Converged Access Deployment – Before You Begin – AP Licensing

- AP licenses are applied at the MC level only (not at MA)
- As with CUWN, a valid license on the Controller is needed for an AP to register
- If MA and MC functionality are not co-located (ex MA on 3x50 and MC on 5760), the communication between MA and MC needs to be UP for AP to join
- Licenses need to be activated on the MC –
 - Converged Access adopts a honor based license mechanism
 - User needs to accept a End User License Agreement (EULA)
 - Use the following command to activate AP licenses:

```
5760-# license right-to-use activate apcount ?
```

```
<5-1000>      configure the number of adder licenses in multiples of 5  
evaluation    activate evaluation license
```

Converged Access Deployment – Before You Begin – AP Licensing

- Must run **ipservices** or **ipbase** license to activate wireless services on 3650 / 3850 –

```
3850# sh license right-to-use
```

Slot#	License name	Type	Count	Period left
1	ipbase	permanent	N/A	Lifetime
1	apcount	base	0	Lifetime
1	apcount	adder	50	Lifetime

```
License Level on Reboot: ipbase
```

- The 5760 does not have activated license levels, the image is already ipservices
- Licensing on a 3850 / 3650 stack –
 - Each switch member in the stack can be licensed independently, up to 25 APs total on a 3650 stack, and up to 50 APs total on a 3850 stack
 - For best redundancy, it is ideal to enable the proper license count for each stack member, based on the number of APs that will be connected to it. So if you lose a switch, you only lose those licenses

Converged Access Deployment – Before You Begin – How to Connect APs

- The Catalyst 3850 and 3650 support only **directly attached APs**

APs need to be in the same VLAN as the Wireless Management interface:

```
interface GigabitEthernet1/0/1
description to_AP
switchport access vlan 31
switchport mode access
```

```
interface Vlan31
ip address 192.168.31.42 255.255.255.0
!
wireless management interface Vlan31
```

If you do not define a wireless management VLAN on the 3x50, the switch will then be transparent to AP attachment and everything will continue to operate as it does today on a 3750-X.

As soon as you define a «wireless management interface VLAN», the Catalyst 3x50 will intercept all incoming AP CAPWAP requests, and terminate / process them at the local ASIC.

- WLC 5760 supports only NON-directly attached APs

Same as it works today in CUWN: AP attached to a local switch (3750-X or alike) finds the centralised controller through DHCP option 43 or other methods and registers

Converged Access Deployment – Branch Use Case – Mobility Configuration

Management VLAN Configuration

```
interface Vlan31
  description MANAGEMENT VLAN
  ip address 192.168.31.42 255.255.255.0
```

SVIs for client VLANs defined locally on the switch

```
interface Vlan32
  description Client VLAN32
  ip address 192.168.32.2 255.255.255.0
```

```
interface Vlan33
  description Client VLAN33
  ip address 192.168.33.2 255.255.255.0
```

Wireless Management Interface Configuration

```
3850(config)# wireless management interface VLAN31
```

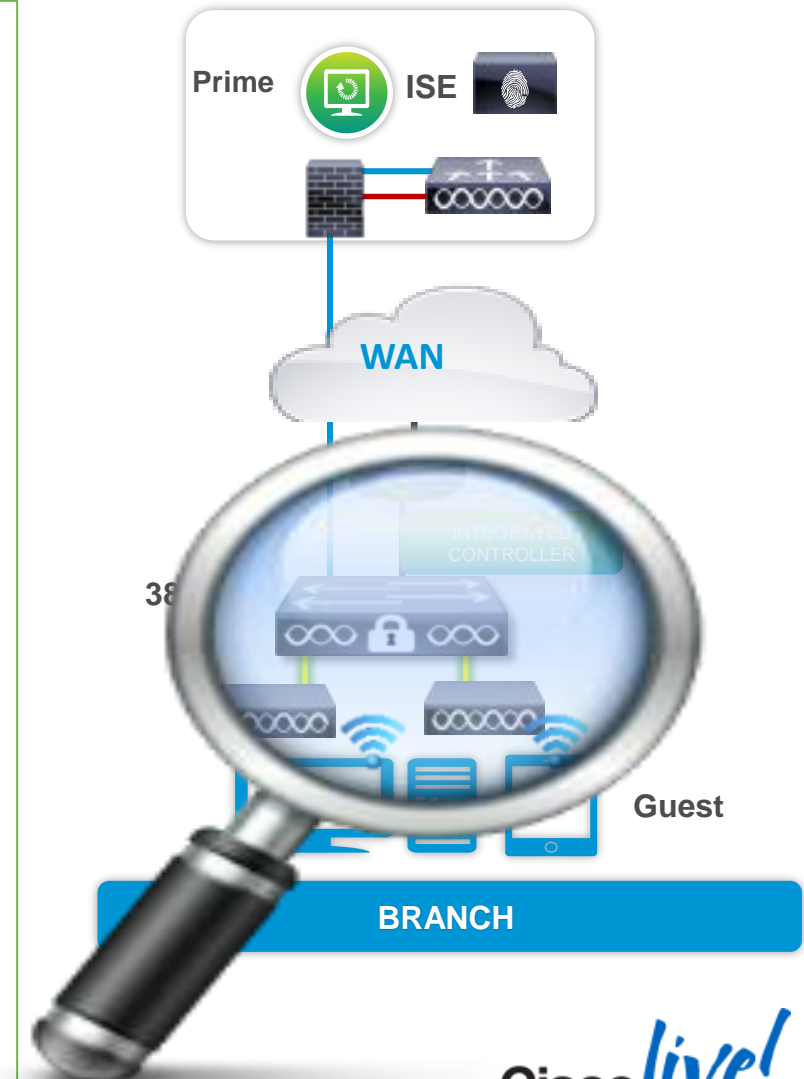
This activates the
MA functionality

```
3850# show wireless Interface summary
```

```
Wireless Interface Summary
```

```
AP Manager on management Interface: Enabled
```

Interface Name	Interface Type	VLAN ID	IP Address	IP Netmask	MAC Address
Vlan31	Management	31	192.168.31.42	255.255.255.0	2037.06ce.0a55



Converged Access Deployment – Branch Use Case – Mobility Configuration, continued

■ Configuring Mobility Controller

```
3850(config)# wireless mobility controller
```

This activates the
MC functionality

```
Mobility role changed to Mobility Controller  
Please save config and reboot the whole stack
```

```
3850# sh wireless mobility summary
```

After reboot


```
Mobility Controller Summary:
```

```

Mobility Role                : Mobility Controller
Mobility Protocol Port       : 16666
Mobility Group Name          : default
Mobility Oracle IP Address   : 0.0.0.0
DTLS Mode                    : Enabled
Mobility Domain ID for 802.11r : 0xac34
Mobility Keepalive Interval  : 10
Mobility Keepalive Count     : 3
Mobility Control Message DSCP Value : 0
Mobility Domain Member Count : 1
Link Status is Control Path Status : Data Path Status
  
```

```
Controllers configured in the Mobility Domain:
```

IP	Public IP	Group Name	Multicast IP	Link Status
192.168.31.42	-	default	0.0.0.0	UP : UP



Converged Access Deployment – Branch Use Case – AP Port and WLAN Configuration

Access Point port configuration

```
interface GigabitEthernet1/0/15
  description - Access port for Access points
  switchport access vlan 31
  switchport mode access
```

Access Points need to be configured on Wireless Management VLAN

```
3850# show ap summary
Number of APs: 1
```

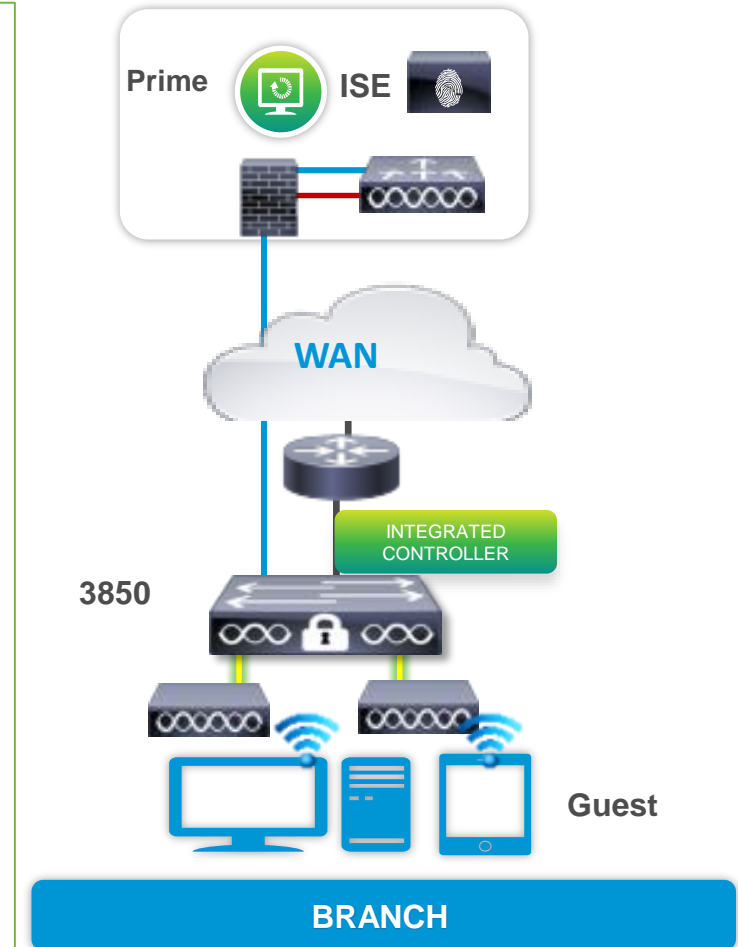
```
Global AP User Name: Not configured
Global AP Dot1x User Name: Not configured
```

AP Name	AP Model	Ethernet MAC	Radio MAC	State
AP3502I	3502I	c47d.4f3a.ed80	04fe.7f49.58c0	Registered

WLAN Configuration

```
3850(config)# wlan WPA-PSK 4 wpa-psk
3850(config-wlan)# client vlan 32
3850(config-wlan)# no security wpa akm dot1x
3850(config-wlan)# security wpa akm psk set-key ascii 0 Cisco1234
3850(config-wlan)# no shut
```

WLAN sample configuration



Converged Access Deployment – Branch Use Case – Client Connectivity

Client Connectivity

```
3850# sh wireless client summary
```

```
Number of Local Clients : 1
```

MAC Address	AP Name	WLAN State	Protocol
f81e.dfe2.e80e	AP3502I	4 UP	11n (5)


```
3850# sh wcd b database all
```

```

Total Number of Wireless Clients = 1
  Clients Waiting to Join         = 0
  Local Clients                   = 1
  Anchor Clients                  = 0
  Foreign Clients                 = 0
  MTE Clients                     = 0

```

Mac Address	VlanId	IP Address	Auth	Mob
f81e.dfe2.e80e	32	192.168.32.57	RUN	LOCAL



Converged Access Deployment – Large Campus Use Case – Mobility Configuration

- Configure 5760 as MC and member of SPG

```
interface Vlan100
  description WIRELESS MANAGEMENT VLAN
  ip address 192.168.100.42 255.255.255.0
```

```
5760(config)# wireless management interface VLAN100
```

```
5760(config)# wireless mobility controller peer-group WestBldg
```

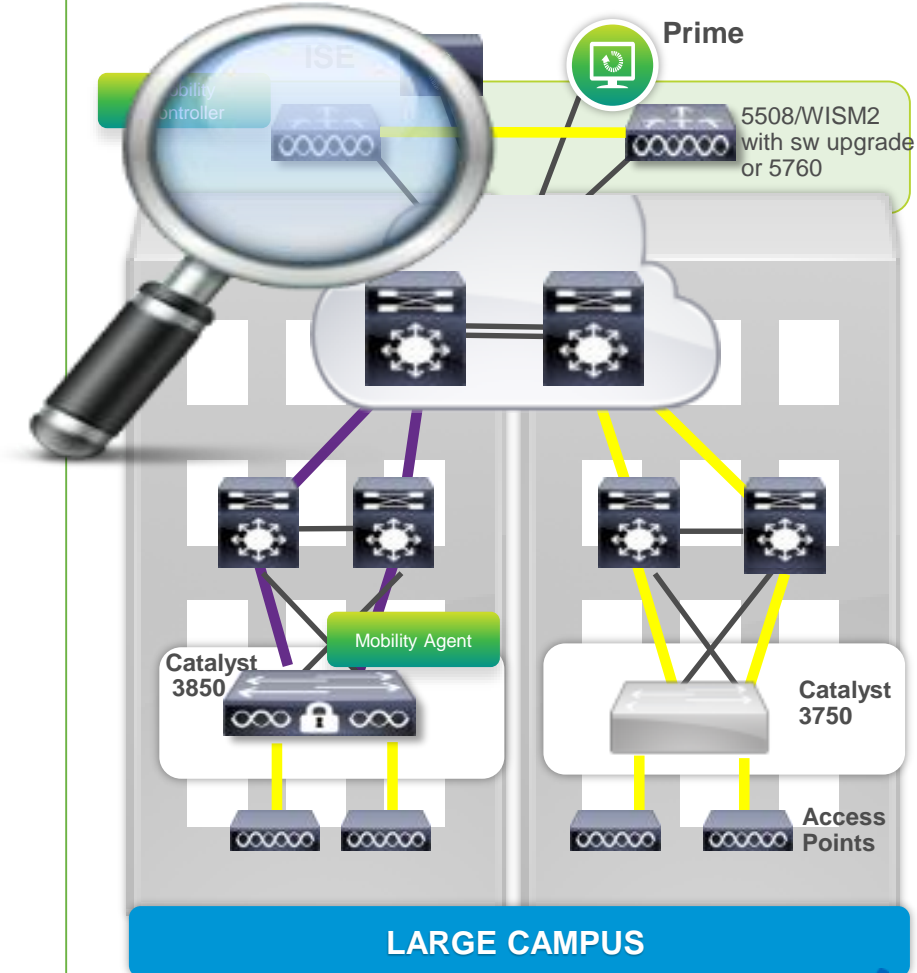
```
5760(config)# wireless mob contr peer-group WestBldg member ip 192.168.41.44
```

- Configure 3850 as MA

```
interface Vlan41
  description MANAGEMENT VLAN
  ip address 192.168.41.44 255.255.255.0
```

```
3850(config)# wireless management interface VLAN10
```

```
3850(config)# wireless mobility controller ip 192.168.100.42
```



Converged Access Deployment – Large Campus Use Case – Mobility Configuration, continued

- Mobility Group configuration

```
5760(config)# wireless mobility group name cisco-live
```

```
5760(config)# wireless mobility group member ip 10.1.1.5
```

- Verify the configuration

```
5760# sh wireless mobility summary
```

Mobility Controller Summary:


Mobility Role : Mobility Controller
 Mobility Protocol Port : 16666
 Mobility Group Name : cisco-live

Controllers configured in the Mobility Domain:

IP Address	Public IP Address	Group Name	Multicast IP	Status
192.168.100.42	-	cisco-live	0.0.0.0	UP
10.1.1.5	10.1.1.5	cisco-live	0.0.0.0	UP

Switches configured in WestBldg switch Peer Group: 1

IP Address	Public IP Address	Status
192.168.41.44	192.168.41.44	UP



LARGE CAMPUS

Cisco live!

Agenda – BRKARC-2665

Evolution ... Towards One Policy, One Management, One Network

Converged Access – Platform Overviews

Existing Wireless Deployments – Architecture Refresher

Converged Access Architecture –

- Terminology and Building Blocks

- Traffic Flows and Roaming

- High Availability

- Quality of Service

- Security

- Multicast

- NetFlow

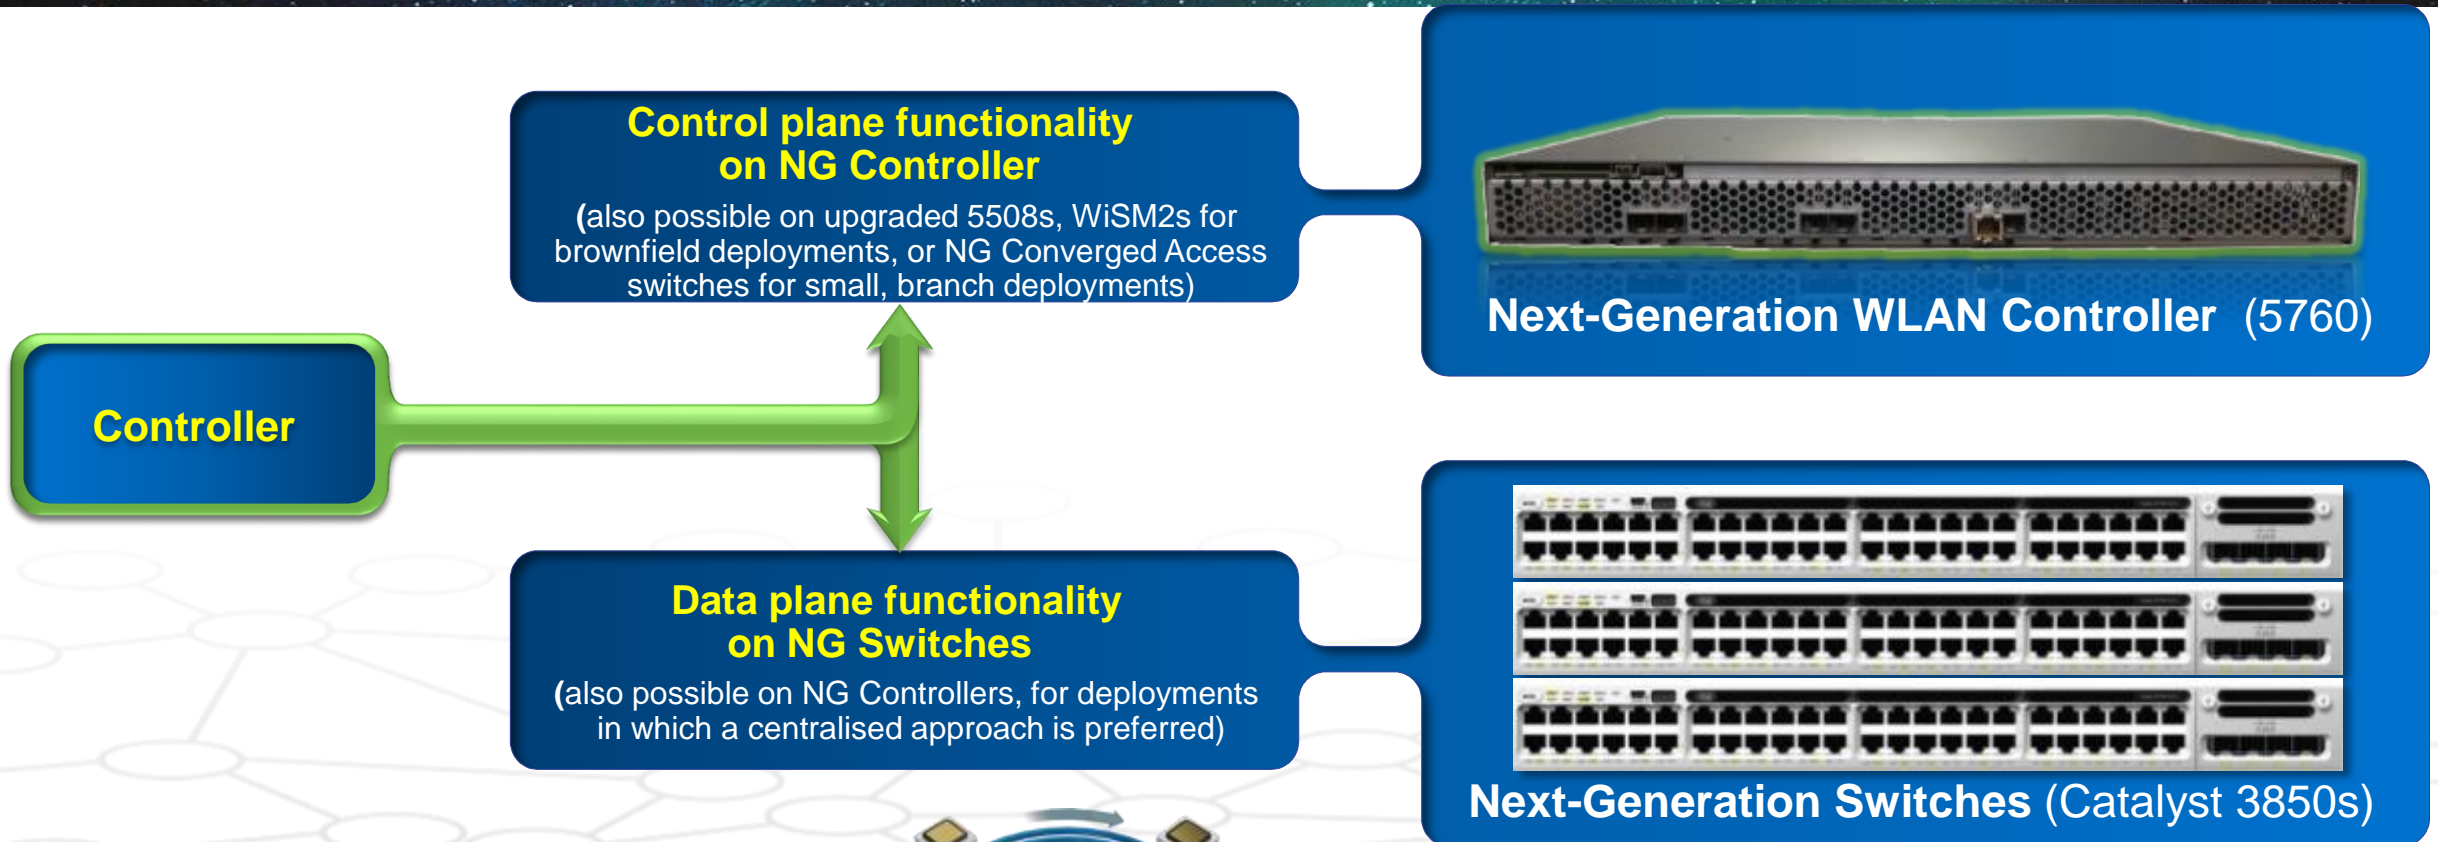
Converged Access Design and Deployment –

- IP Addressing


- Design Options

- Deployment Examples

Bringing Together Wired and Wireless – How Are We Addressing This Shift?



**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



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

Converged Access—

Tell Me How I Did!

Did I Achieve My Objectives?

Do You Have a Better Understanding ...

of what Converged Access is ...

of how Converged Access works ...

and how you would use it
in your network designs?

Don't Forget
to fill out your evaluations!





Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 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 21 March 12:00pm - 2:00pm



Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.

www.CiscoLiveAPAC.com



CISCO TM




REFERENCE MATERIAL

SCALABILITY


CUWN Scalability –

With CT5508 – Mobility Groups and Domains

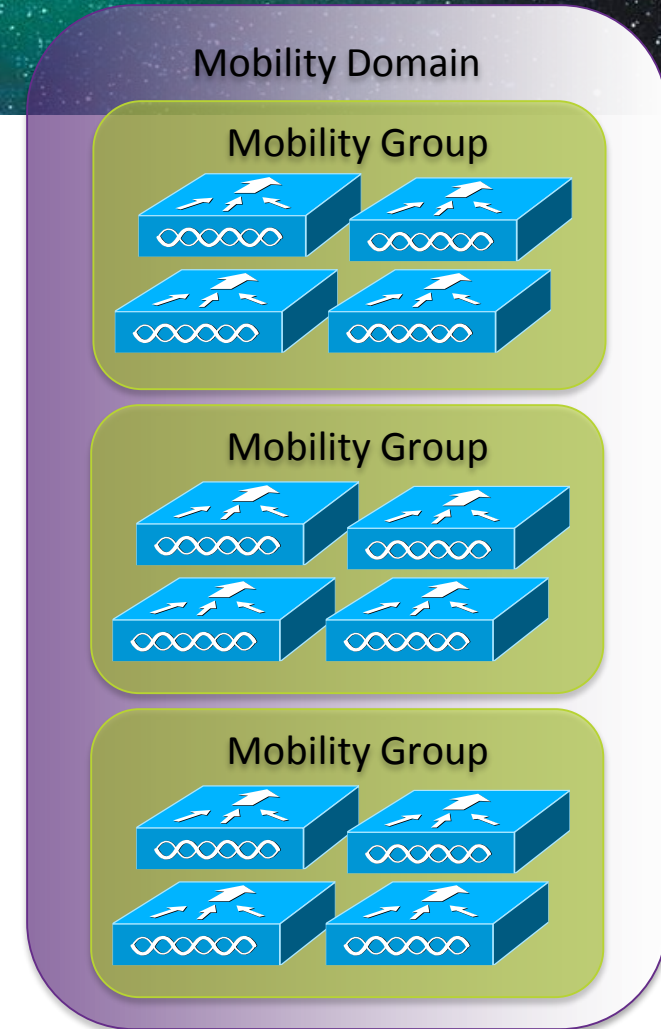


- Up to 500 APs
- Up 7K Clients
- Up to 8 GB I/O for AP Traffic

- CT5508 rel 7.6
- Max theoretical scalability numbers
- Without Considering FlexConnect



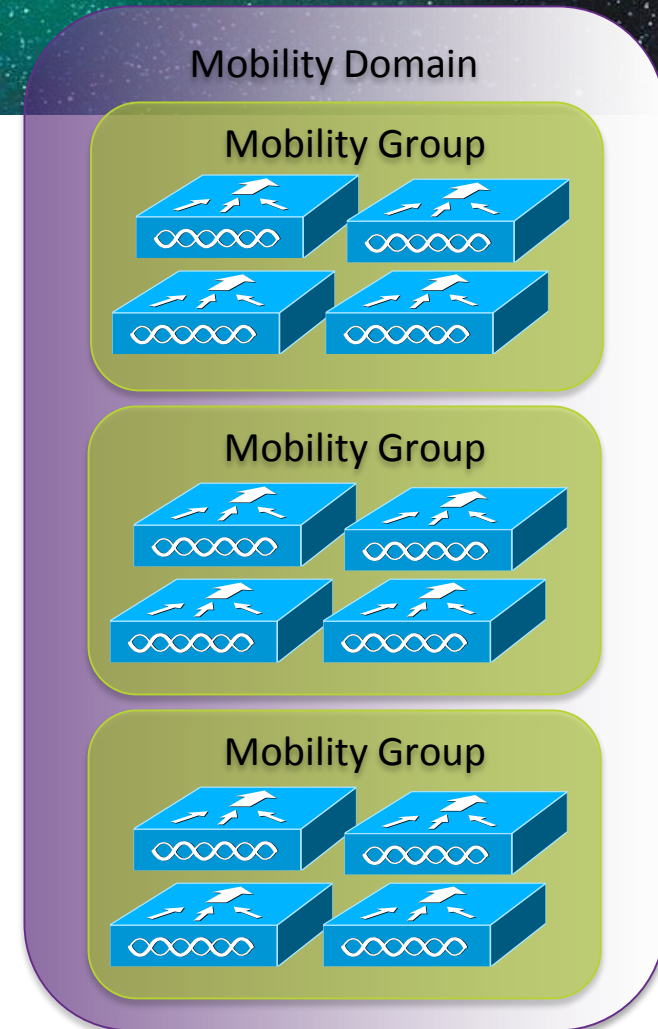


- Up to 12K APs
- Up 168K Clients
- Up to 24 WLCs in a MG
- Up to 192 GB I/O for AP Traffic



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


CUWN Scalability –

With WiSM-2 – Mobility Groups and Domains




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



Converged Access Scalability – With Catalyst 3650

Cat3650
MACat3650
MC

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



- 1 MC = 1 SD
- Up to 25 APs
- Up to 1K Clients
- Up to 20 GB I/O for AP Traffic




25AP
1K Clients
16 MA

...

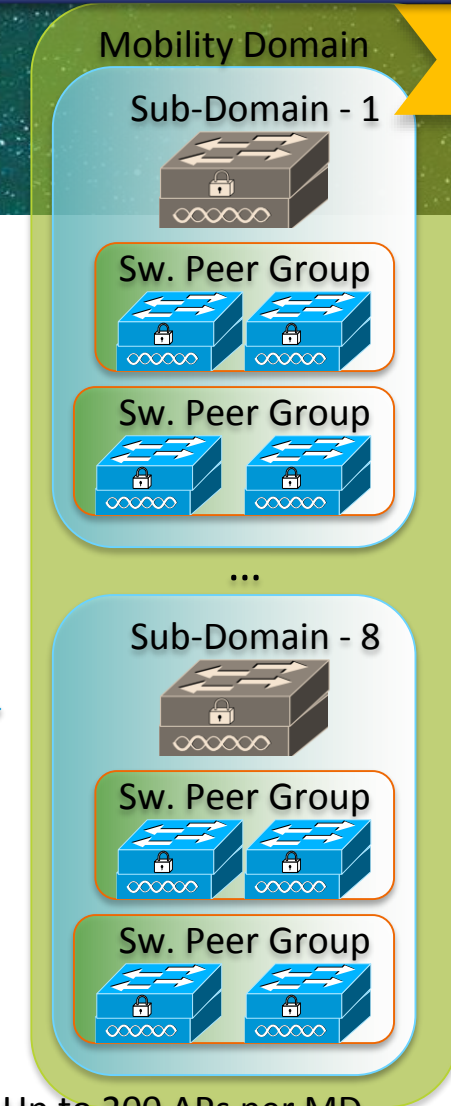
...

...



25 AP
1K Clients
8 SPG


- Up to 25 APs per SPG/MC
- Up to 1K Clients per SPG/MC
- Up to 16 MAs in a SPG/MC
- Up to 8 SPGs in a SD
- Up to 25 GB I/O for AP Traffic




200 AP
8K Clients
8 SD

- Up to 200 APs per MD
- Up to 8 SDs per MD
- Up to 128 MAs per MD
- Up to 8K Clients per MD
- Up to 250 GB I/O for AP Traffic


Converged Access Scalability – With Catalyst 3850

Cat3850
MACat3850
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




50 AP
2K Clients
16 MA

...

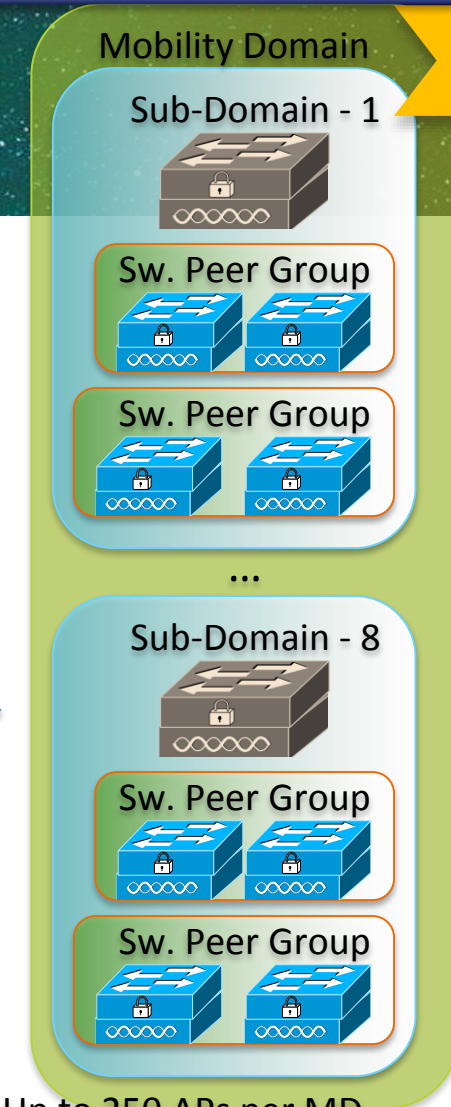
...

...



50 AP
2K Clients
8 SPG

- 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




250 AP
16K Clients
8 SD

- 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

Converged Access Scalability –

With WLC 2504 as MC and 3850s as MAs

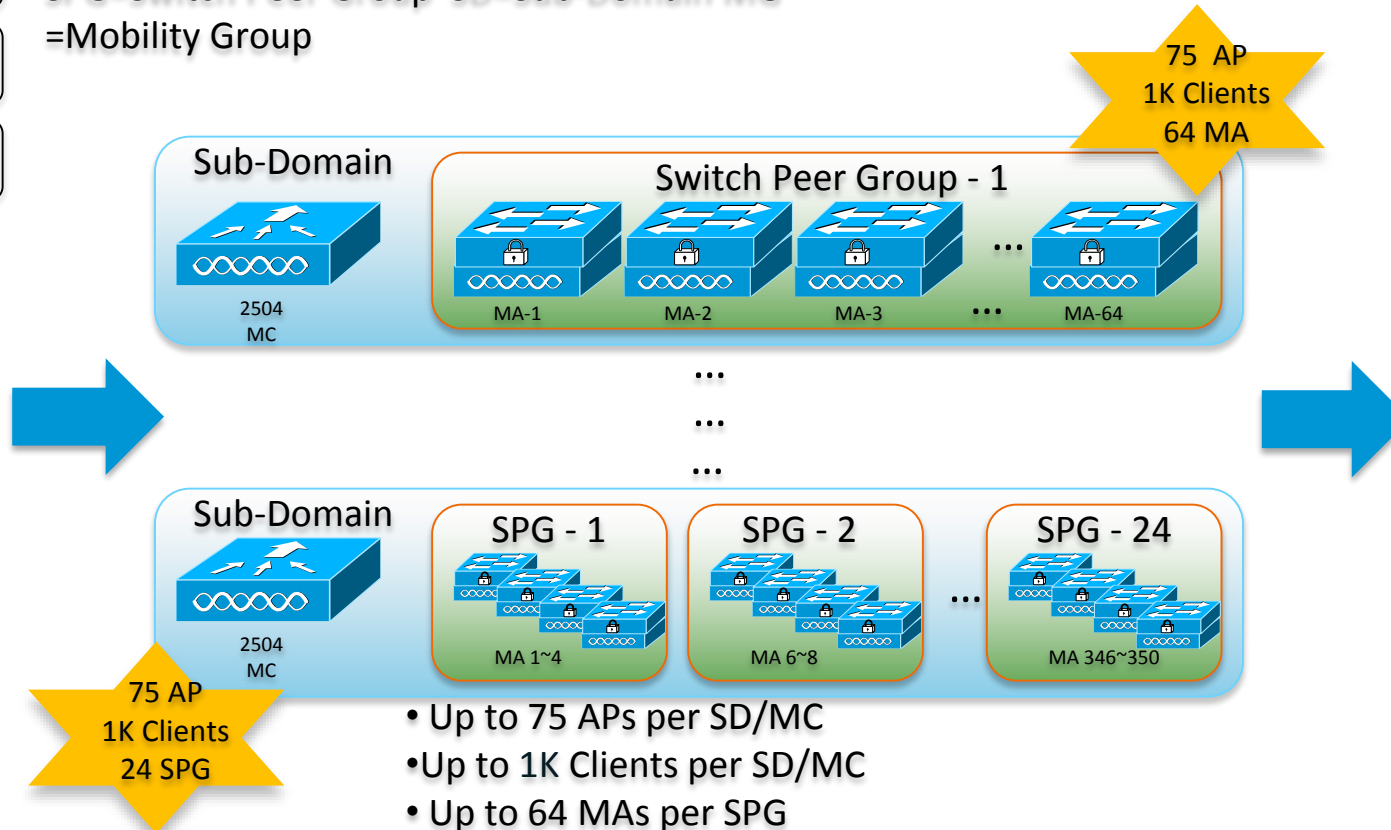


MC/MA
on one Switch

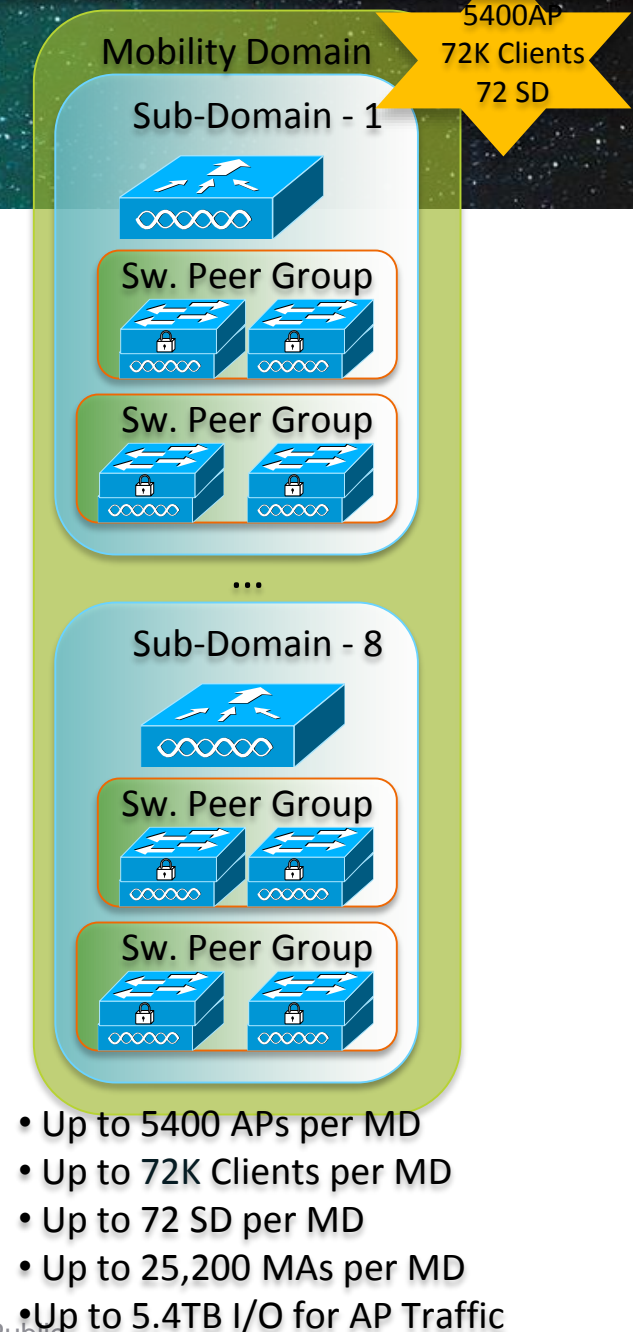
Sub-Domain

- 1 MC = 1 SD
- Up to 75 APs
- Up to 1K Clients
- Up to 1GB I/O for AP Traffic

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




- Up to 75 APs per SD/MC
- Up to 1K 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 75GB I/O for AP Traffic



- Up to 5400 APs per MD
- Up to 72K Clients per MD
- Up to 72 SD per MD
- Up to 25,200 MAs per MD
- Up to 5.4TB I/O for AP Traffic

Converged Access Scalability –

With CT5760 as MC and 3850s as MAs

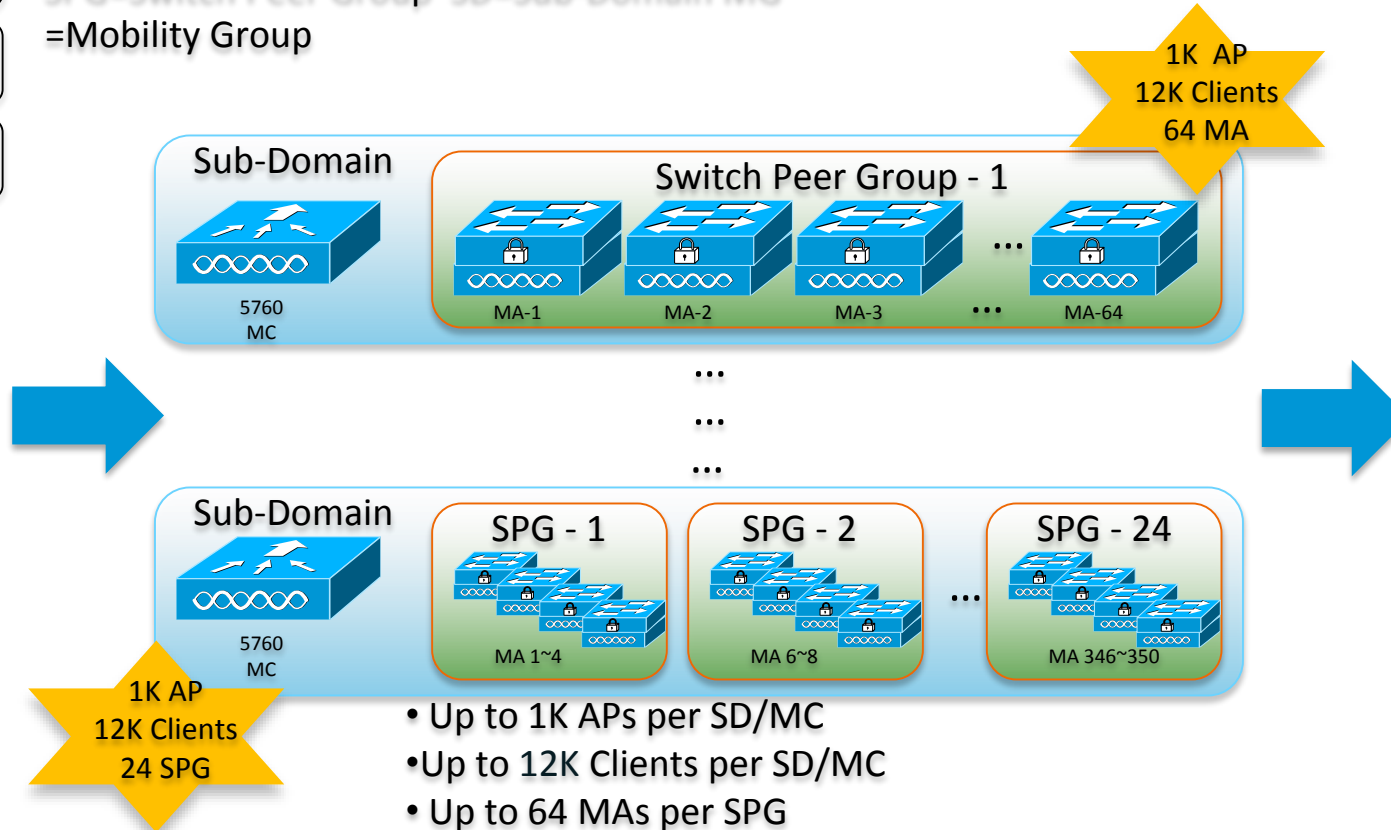


MC/MA
on one Switch

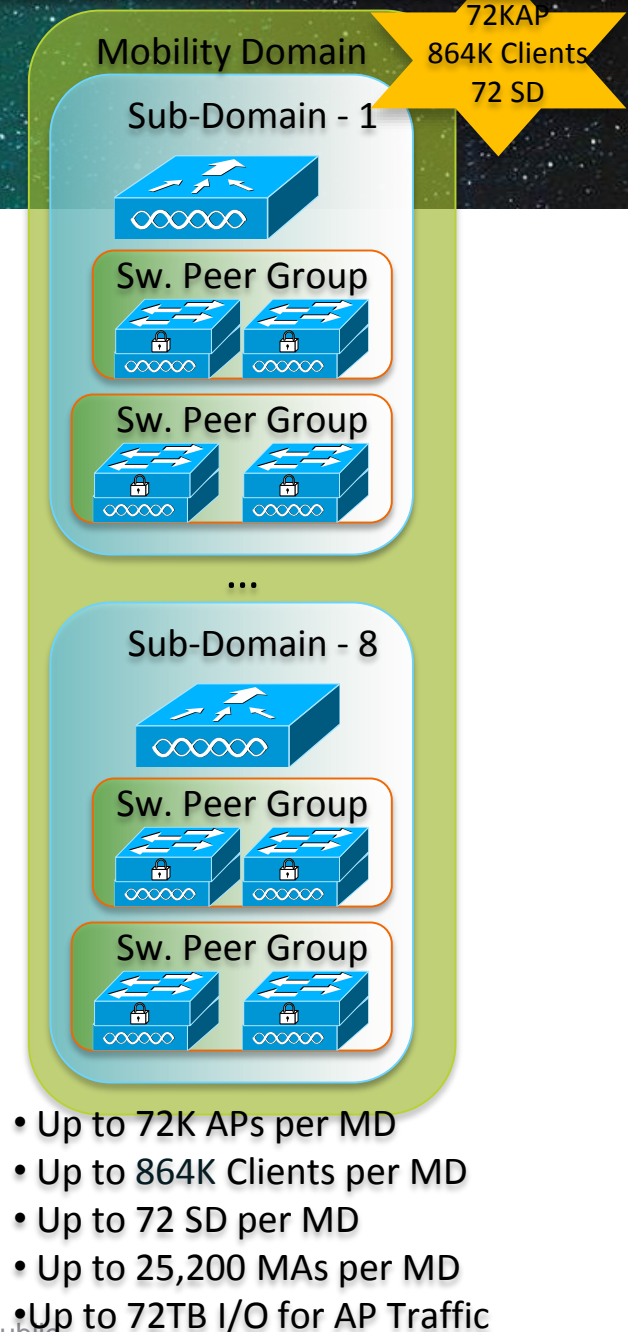
Sub-Domain

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

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




- 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



Converged Access Scalability –

With CT5508 as MC and 3850s as MAs




MC/MA
on one Switch

Sub-Domain

- 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

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



- Up to 500 APs per SD/MC
- Up to 7K Clients per SD/MC
- Up to 64MAs 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

500 AP
7K Clients
64 MA

Mobility Domain

Sub-Domain - 1



Sw. Peer Group



Sw. Peer Group



...

Sub-Domain - 8



Sw. Peer Group



Sw. Peer Group




- 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

36K AP
504K Clients
72 SD

Converged Access Scalability – With WiSM-2 as MC and 3850s as MAs

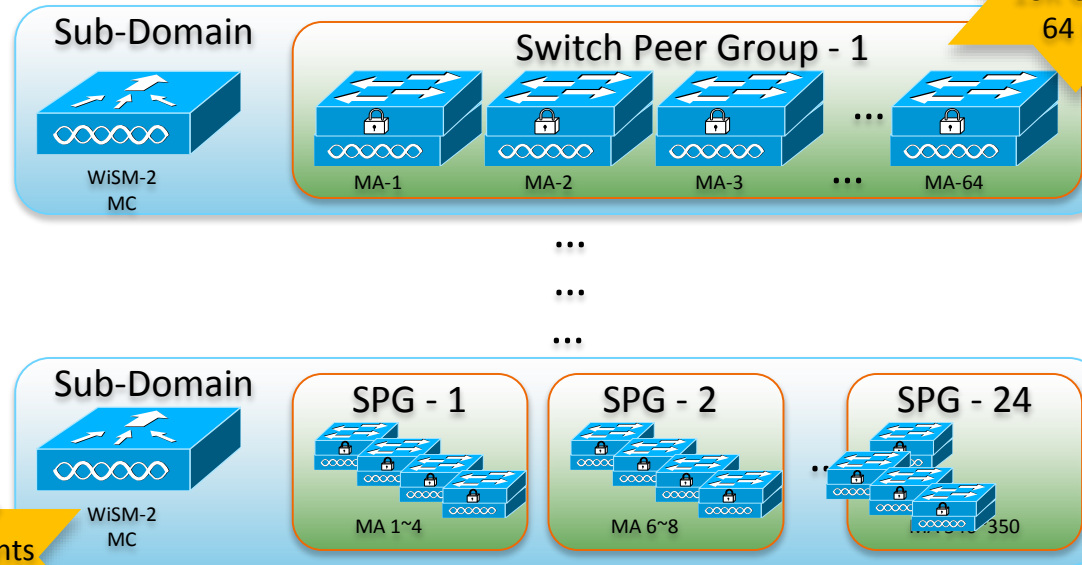

MA=Mobility Agent MC=Mobility Controller
SPG=Switch Peer Group SD=Sub-Domain MG
=Mobility Group



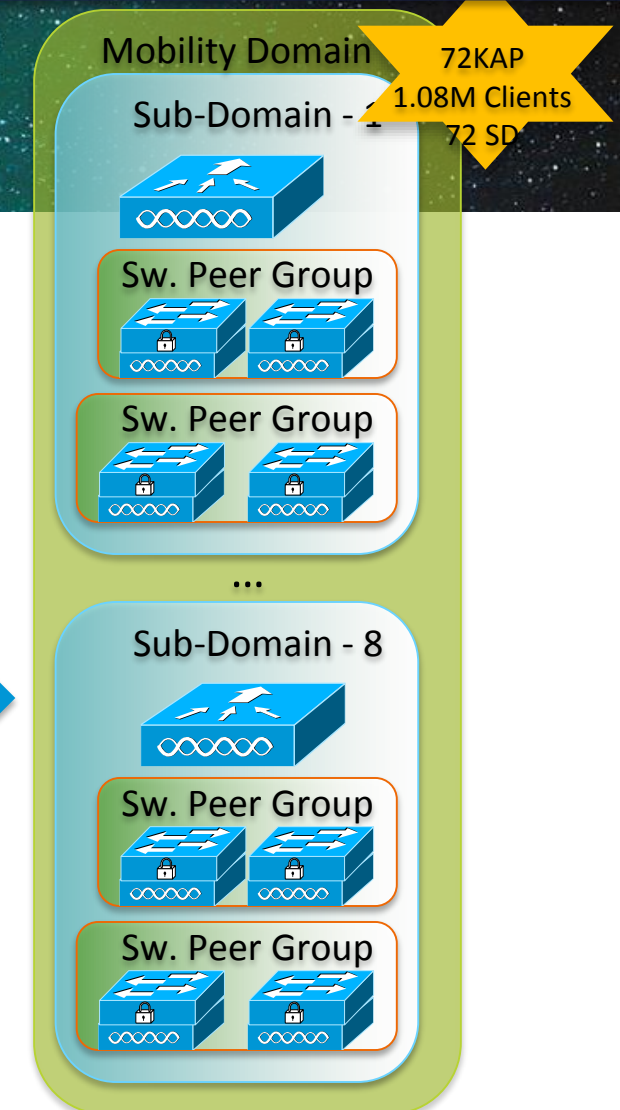
MC/MA
on one Switch

Sub-Domain

- 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

Converged Access Scalability – Summary



For Your
Reference

Scalability	3650 as MC (3.3.1SE)	3850 as MC (3.3.1SE)	WLC2504 (7.6)	WLC5760 (7.6)	WLC5508 (7.6)	WiSM2 (7.6)
Max APs Supported per MC	25	50	75	1000	500	1000
Max APs Supported in overall Mobility Domain	200	250	5400	72000	36000	72000
Max Clients Supported per MC	1000	2000	1000	12000	7000	15000
Max Clients Supported in overall Mobility Domain	8000	16000	72000	864000	504000	1.08M
Max number of MC in Mobility Domain	8	8	72	72	72	72
Max number of MC in Mobility Group	8	8	24	24	24	24
Max number of MAs in Sub-domain (per MC)	16	16	350	350	350	350
Max number of SPGs in Mobility Sub-Domain (per MC)	8	8	24	24	24	24
Max number of MAs in a SPG	16	16	64	64	64	64
Max number of WLANs	64	64	16	512	512	512

REFERENCE MATERIAL

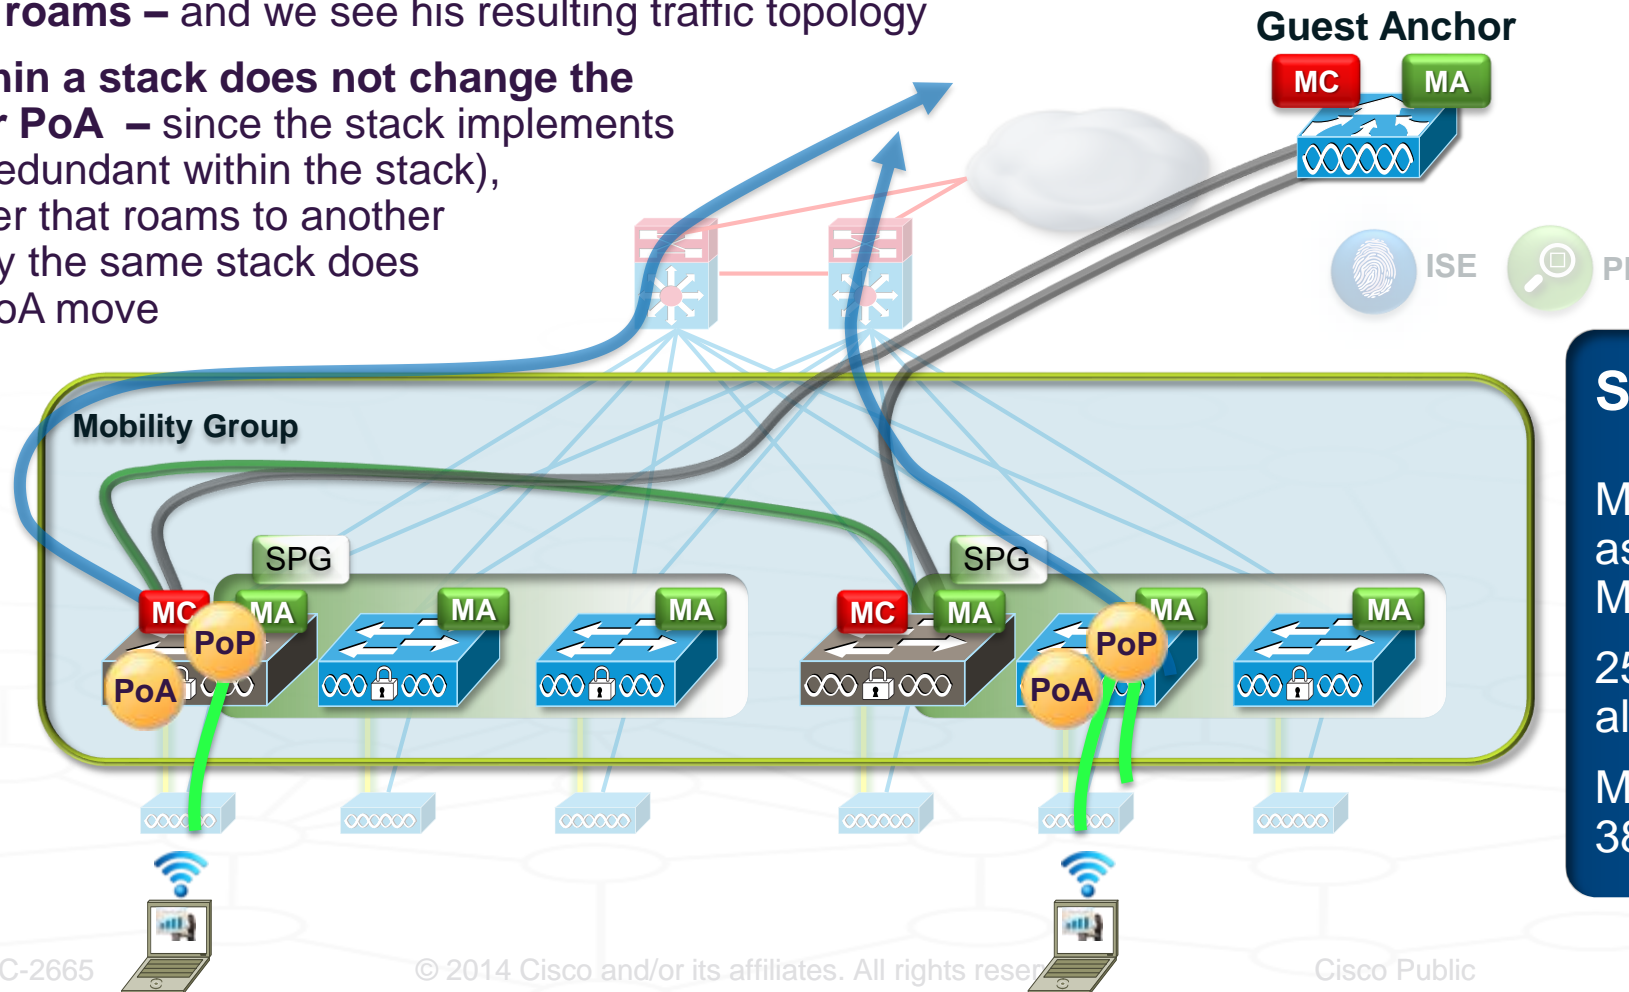
CATALYST 3850-BASED MCs – ROAMING DETAILS

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-based MCs

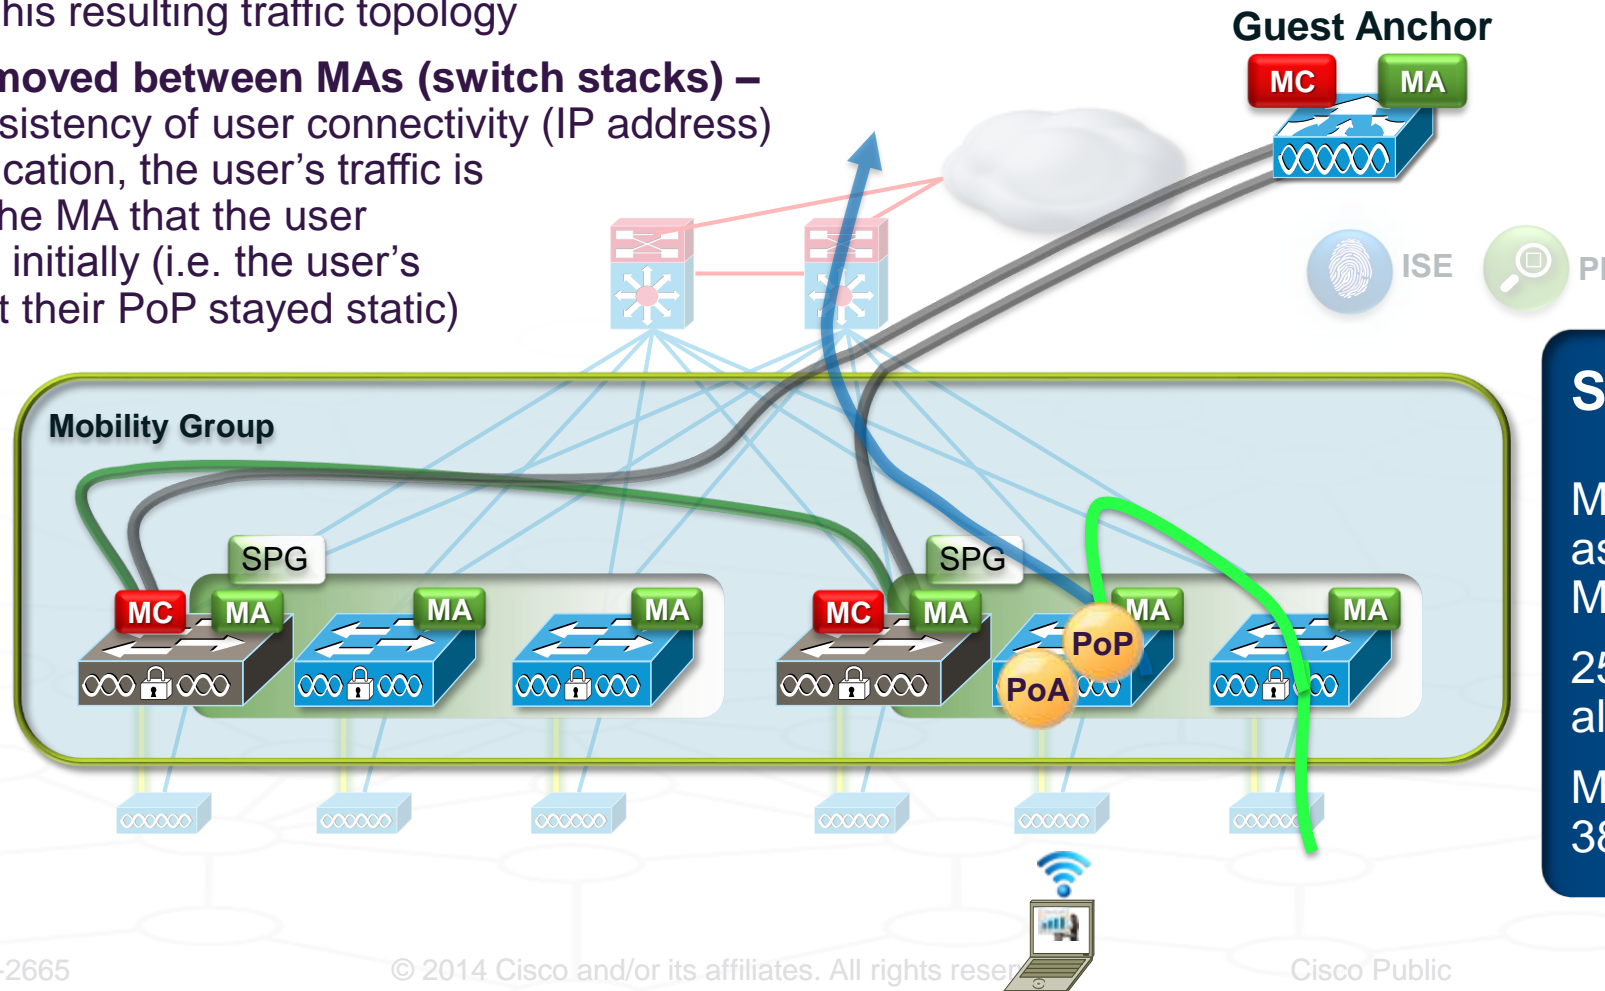
Max. 50 APs per 3850 stack / SPG

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-based MCs

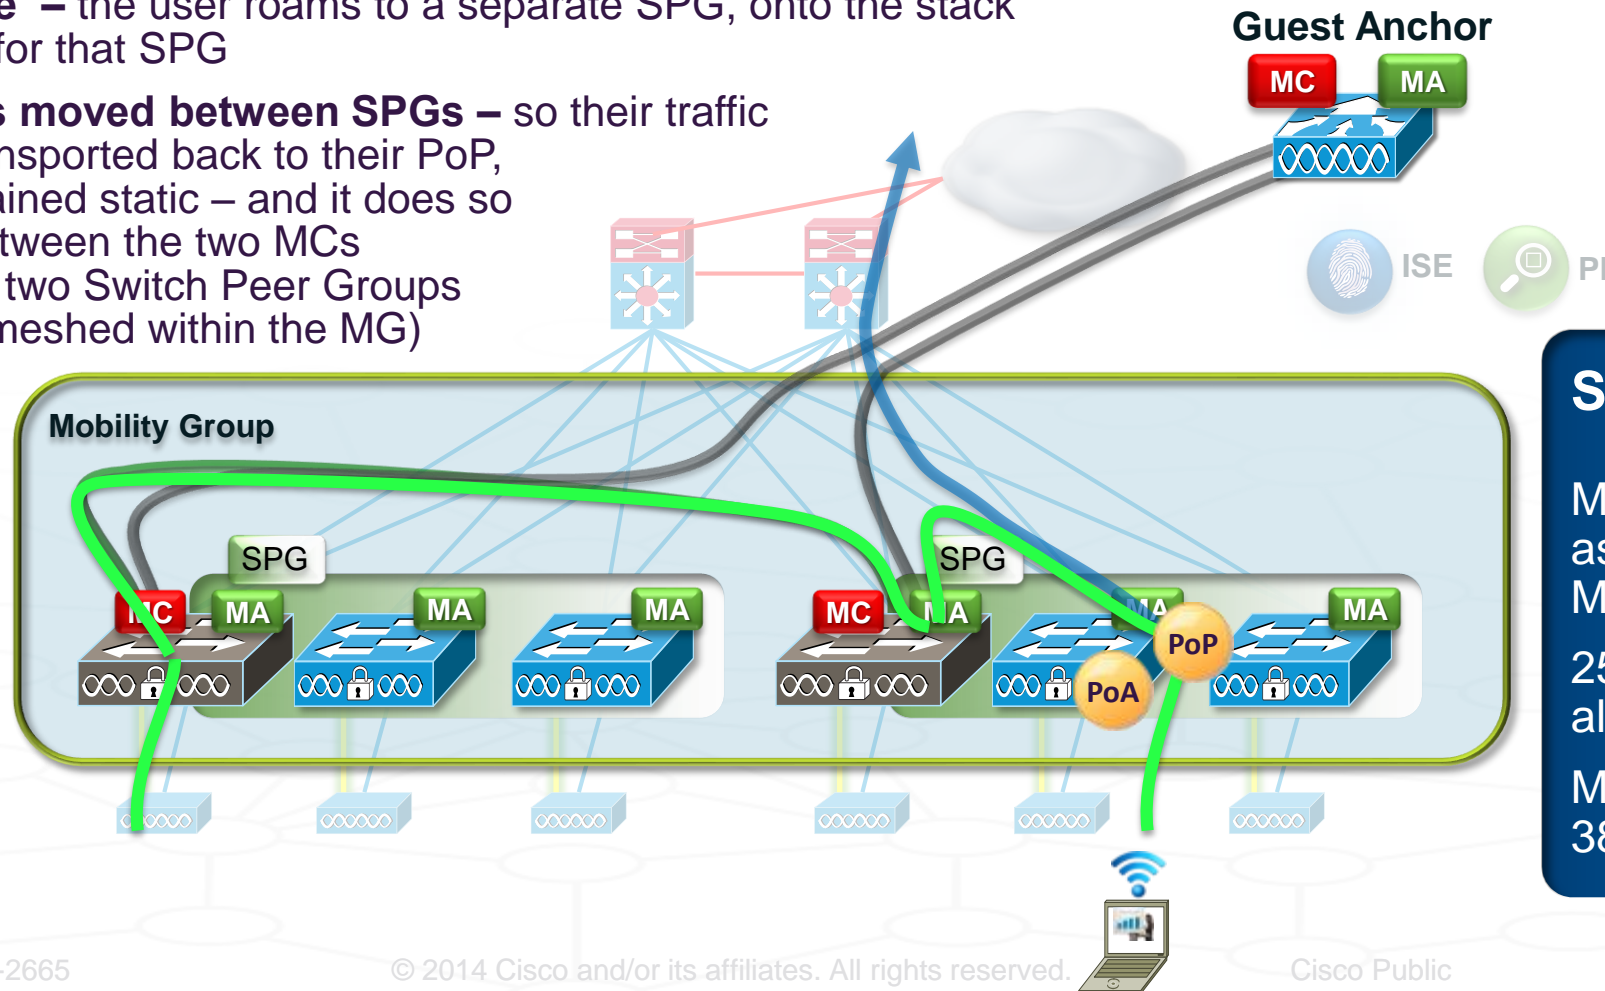
Max. 50 APs per 3850 stack / SPG

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-based MCs

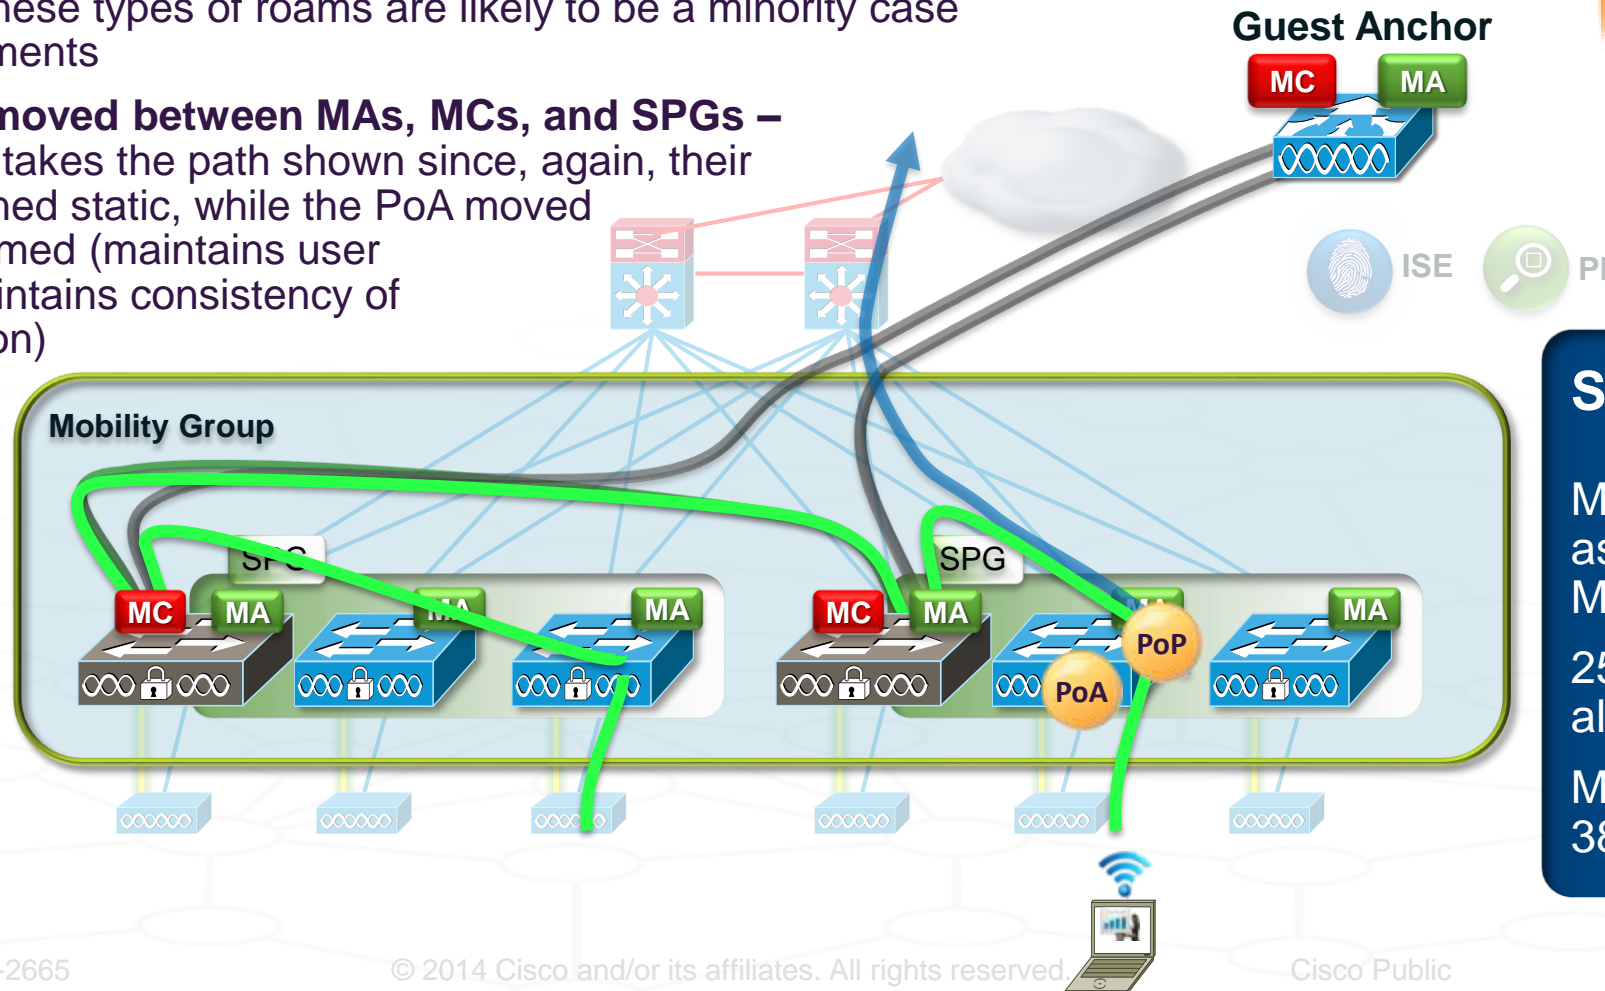
Max. 50 APs per
3850 stack / SPG

Converged Access –

Catalyst 3850-based MCs – Roaming across SPGs and MCs

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

- Now, let's 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-based MCs

Max. 50 APs per
3850 stack / SPG



REFERENCE MATERIAL

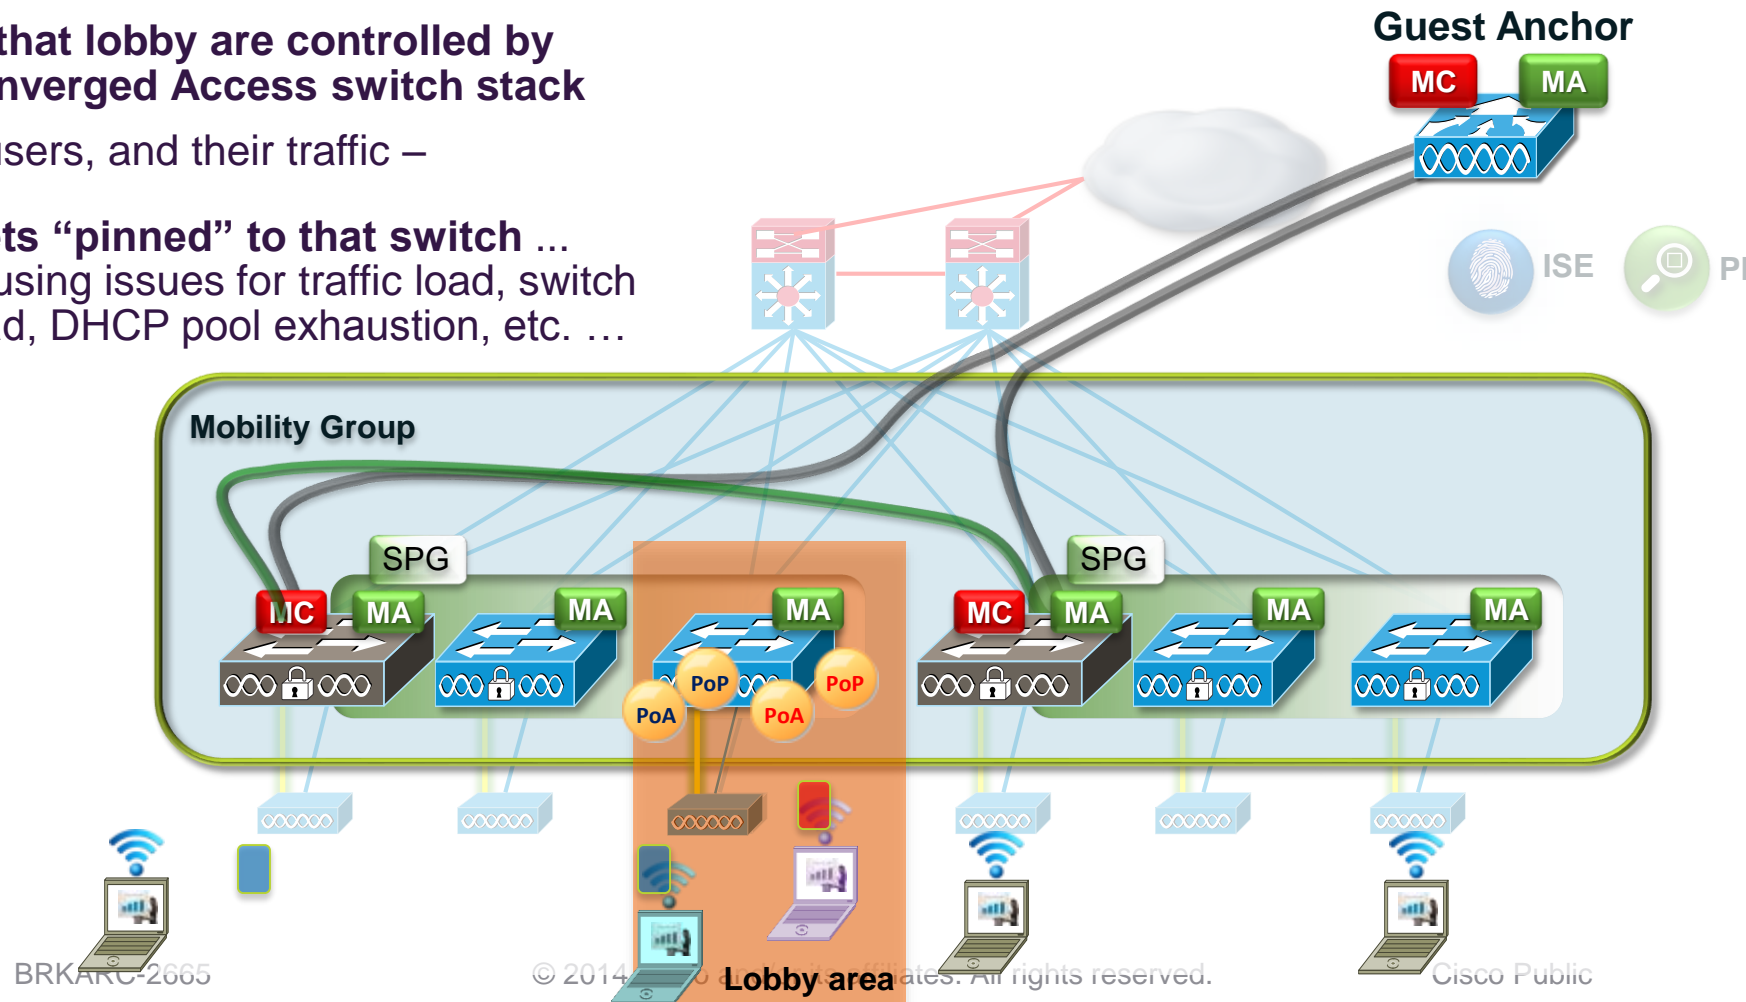
LOBBY ISSUE / SOLUTION

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

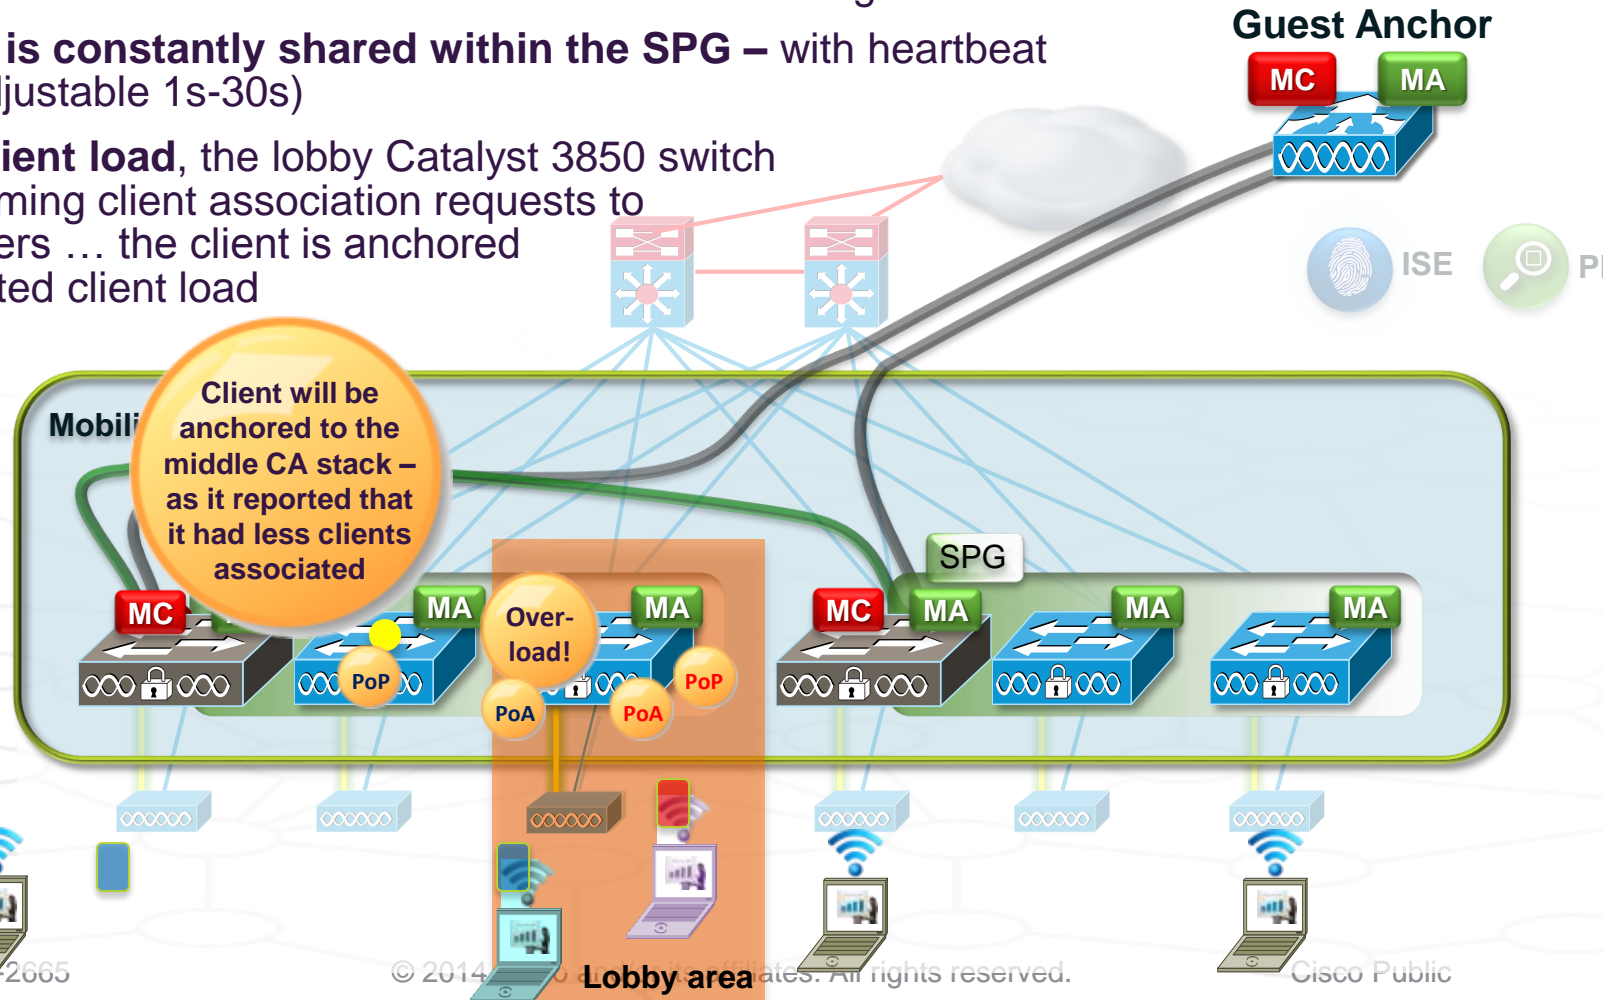


Converged Access – Common Building Access – The “Lobby Solution”

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 its 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
```



REFERENCE MATERIAL

DEPLOYMENT

Converged Access Deployment – IOS-XE 3.2.0 (FCS) vs. AireOS – Feature Comparison

Additional Features included

7.2

AP 3600 support
IPv6/dual stack 'client
Mobility'
ISE 1.1 MnR
OKC/PKC

7.3

AP 2600 support
Right-To-Use Adder
Licenses

7.4

AP 1600 support

IOS XE 3.2.0 features are based on AireOS features 7.0.116.0

Features NOT included

7.2

CleanAir enhancement
Limit # of Clients per radio
and per SSID
Wi-Fi Direct
AP Groups/Profile ph2
SXP

7.2 MR1

802.11r
HTTP sensor

7.3

AP SSO
Bid. Rate Limiting
11n Voice CAC
Video CAC
ISE 1.2: DHCP sensor
Hot spot 2.0
PMIPv6 MAG

7.4

AVC
Bonjour Services Dir.
Neighbour List (11k)
N+1 with HA SKU
Modules on AP3600
802.11w for local mode

Converged Access Deployment – IOS-XE 3.2.X (Maintenance) vs. AireOS – Feature Comparison

Additional Features included

7.2

AP 3600 support
IPv6/dual stack 'client
Mobility'
ISE 1.1 MnR
OKC/PKC

7.3

AP 2600 support
Right-To-Use Adder
Licenses

7.4

AP 1600 support

Maintenance Releases

PI 2.0 support (3.2.3)
Captive Portal Bypassing
(3.2.3)
GUI enhancements (3.2.2)
Fast SSID change (3.2.2)
CoA for BYOD support (3.2.2)

IOS XE 3.2.x features are based on AireOS features 7.0.116.0

Features NOT included

7.2

CleanAir enhancement
Limit # of Clients per radio
and per SSID
Wi-Fi Direct
AP Groups/Profile ph2
SXP

7.2 MR1

802.11r
HTTP sensor

7.3

AP SSO
Bid. Rate Limiting
11n Voice CAC
Video CAC
ISE 1.2: DHCP sensor
Hot spot 2.0
PMIPv6 MAG

7.4

AVC
Bonjour Services Dir.
Neighbour List (11k)
N+1 with HA SKU
Modules on AP3600
802.11w for local mode

Converged Access Deployment – IOS-XE 3.3.0 vs. AireOS – Feature Comparison

Additional
Features included

IOS XE 3.3 features are based on AireOS features 7.4

Features **NOT**
included

7.4

The “C” in AVC
Device Sensor (HTTP and DHCP)
Hot Spot 2.0
IPv6 Source Guard
LSC (Local Signed Certificate)
CMX (see next slide)

7.5

Client SSO
Wireless Policy Classification engine
Bonjour phase 2
Guest Access sleeping clients
OEAP split tunnelling

Converged Access Deployment – IOS-XE 3.3.1 vs. AireOS – Feature Comparison

Additional Features included	IOS XE 3.3 features vs. AireOS features 7.4	
	7.5	7.6
	802.11ac Module for 3600	AP 3700 802.11ac Module for 3600
Features NOT included	3.3.1 is the Recommended release*	
	7.5	
	The “C” in AVC Device Sensor (HTTP and DHCP) Hot Spot 2.0 IPv6 Source Guard LSC (Local Signed Certificate) CMX (see next slide)	Client SSO Wireless Policy Classification engine Bonjour phase 2 Guest Access sleeping clients OEAP split tunneling



CISCO TM