Creating the Network at Cisco Live San Diego 2012: Technical Case Study

Last updated: October 2012

The Cisco Live conference network is one of the most critical elements of the conference. Built on a 40GE backbone with a wide range of borderless network services, the Cisco Live San Diego network required hundreds of access switches (wired connections), access points (wireless connections) and provided network services like Medianet, security, firewall, load balancing, IPv6, and network monitoring to meet the constant needs of nearly 20,000 conference attendees.

Overall Design and Architecture

The conference network was built by a small team of specialists, including 20 engineers. Several key design considerations were defined:

- Simplicity
- Reliability
- Robustness
- Flexibility
- Demonstration

Simplicity

The entire network was staged in San Jose, California, roughly one month before the event. Given the time constraints, the network had to be simple at multiple levels. The number of hardware architectures was purposely limited, as well as the number of network Operating Systems (OSs) and versions of those networks OSs. The team used a three tier architecture, with a core level, a distribution level and an access level. The core level was composed of two Cisco Catalyst 6500 switches with Supervisor Engines 2T configured in Virtual Switching System (VSS) mode, the distribution level was composed of several Cisco Catalyst 6500 switches either standalone with redundant supervisors or in VSS mode for the critical sections. The distribution level provided connectivity to four different locations: the Hilton® hotel, the Marriott® hotel, the San Diego Conference Center (SDCC), and the World of Solutions, an area within the SDCC where more than 200 hundred partners were showcasing their solutions. The access layer was composed of Cisco Catalyst 3560E and 3550CG switches, depending on the density of the ports required for the different conference rooms. See Figure 1 below.

Using VSS allowed the team to simplify first hop routing by removing the need for first hop redundancy protocols like HSRP or VRRP, while providing full redundancy. It also removed the need for multiple IP addresses required by HSRP or VRRP, and reduced the amount of IGP interfaces used at the control plane level.

Figure 1. Cisco Live San Diego 2012 Network

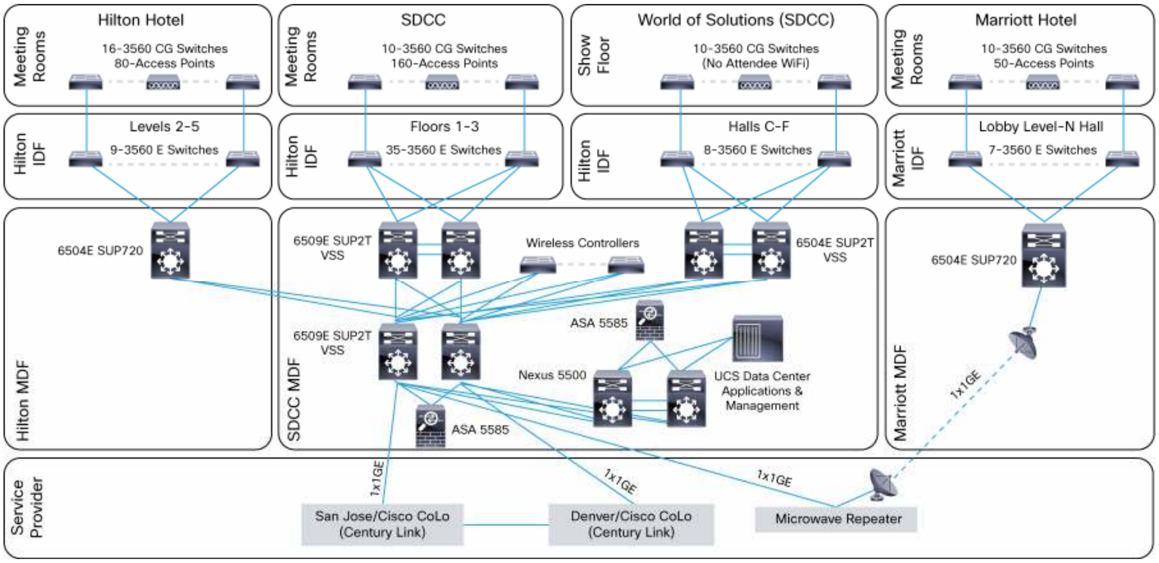

Reliability

All devices were equipped with redundant supervisors or control plane processors, and redundant power supplies. Chassis teaming technologies were used, vPC for the Nexus 5500 and VSS for the Catalyst 6500, to allow the team to redundantly attach access devices without relying on spanning tree, by using an ether channel with member ports on each distribution switch forming the VSS or vPC. These technologies not only reduce CPU resource cycles, but also decrease the risk of outage due to misconfiguration, while providing us more bandwidth in aggregating redundant links in a single port channel.

Robustness

To provide resiliency to the overall network, all devices were configured to use TACACS to authenticate and authorize users using a profile defined in an active directory server. Unauthorized access was logged in the Cisco Prime LAN Management Server (LMS).

A template was designed for each port type depending on its role (phone, access point or media), and Cisco LMS was used to provision each device accordingly.

Demonstration

The Cisco Live event is always a good opportunity to showcase new technologies. This year was no exception. The team introduced a 40 Gigabit Ethernet backbone based on the new Catalyst 6500 WS-X6904-40G-2T linecard. The team also used the 8.4 release of the ASA to provide IPv6 connectivity using NAT64 (see section 10).

Internet Access and Colocations

Having Internet access located within the service provider premises of CenturyLink allowed us to provision connectivity without relying on any preexisting infrastructure. It also enabled us to stage the network within the Cisco labs independently. To improve network reliability, we setup two colocation sites in different geographic regions. These colocation sites were also set up to provide basic services, voice termination and management.

In the colocation center, we deployed a Cisco ASR 1000 router, two Catalyst 6500s in VSS mode, a UCS server and a Cisco Unified Border Element (CUBE).

The ASR 1000 routers were peering using BGP for both IPv4 and IPv6, with routers within the CenturyLink Autonomous System (AS). A redundant connection was established between the colocation sites to ensure maximum uptime. From the San Diego Convention Center, the Catalyst 6500 was peering with both ASR 1000 routers within the same AS. The ASR 1000s were injecting 2 default routes to the show: 0.0.0.0/1 and 128.0.0.0/1, and each of these prefixes was carrying a different weight in order to provide efficient load-balancing between the two links. The team was using a Layer 2 link between the colocation sites and the San Diego Convention Center to provide for redundancy, and also to be able to use both firewalls in active-active mode, thereby avoiding asymmetrical routing. Contrary to IPv4, IPv6 was primarily routed through the Sunnyvale colocation site, with the Denver colocation site as a standby. A backup MPLS connection was also provided in case of a failure of both Layer 2 links. This Layer 3 MPLS connection was idle during the conference. See Figure 2.

Figure 2. Internet Connectivity Architecture

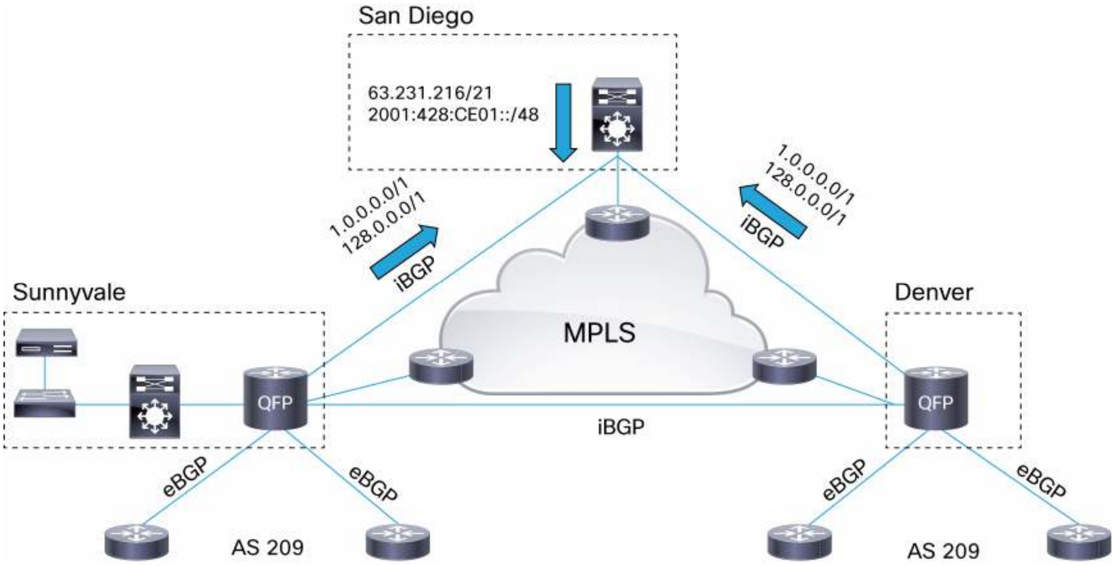

Secure Origin BGP

During Cisco Live San Diego, both ASR 1000 routers were configured with Secure Origin BGP (soBGP). Under soBGP, digital certificates are used to authorize and authenticate packets (see Figure 3). It also proposes a new mechanism to relay information about the security of the routing system outside of the routing system itself.

soBGP verifies the following:

- An Autonomous System (AS) is authorized to advertise a specific prefix.
- A valid path exists to a specific prefix advertised by an AS.
- The peer advertising a route is authorized by the originator, or owner, of the destination, to advertise a path to this destination.
- The path advertised by a peer AS falls within the policies the local network administrators have implemented.

During the conference the team identified 4558 hijacked prefixes and removed them from the local routing table.

Below is an excerpt of the soBGP configuration:

```
bgp rpki server tcp 217.193.137.117 port 30000 refresh 60
bgp rpki server tcp 2001:918:FFF9:0:250:56FF:FE15:159 port 8282 refresh 60
bgp rpki server tcp 2001:918:FFF9:0:250:56FF:FE15:159 port 30000 refresh 60
bgp rpki server tcp 217.193.137.117 port 8282 refresh 600
```

For more information on soBGP, please refer to: Securing BGP Through Secure Origin BGP

Figure 3. soBGP Architecture

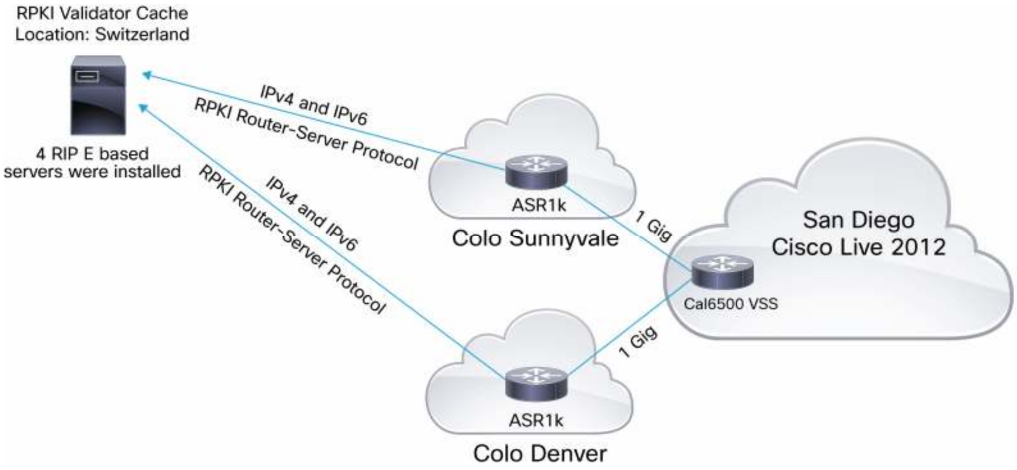

Routing and Switching

The network provided both IPv4 and IPv6 connectivity for both wired and wireless users. And CenturyLink assigned a public address range for both IPv4 and IPv6. That address range was divided into multiple subnets, which were used by internal services that required external access. The rest was used as a pool for Network Address Translations (NAT). Deployed at the convention center network edge, the Cisco ASA Firewall was ideally positioned to provide Carrier Grade Network Address Translation (CGN) services to all users. CGN was administered using a pool of 1022 addresses.

All wireless connections were aggregated on a single subnet and first hop security was set in place on the VLAN Interface that supported the wireless network (see Figure 4). Aggregating all wireless users on a single subnet allowed for easy seamless roaming between APs scattered throughout the convention center.

Layer 2 connections were terminated on each of the distribution switches. OSPFv3 was used for internal routing with access lists to restrict accessibility to specific ranges.

Figure 4. Wireless VLAN Configuration

```
interface Vlan2011
 description Wireless_Users
 ip address 10.211.0.1 255.255.0.0
 ip helper-address 10.63.231.72
 ip helper-address 10.63.231.73
 ip ospf authentication-key <removed>
 ip ospf 100 area 0
 ipv6 address FE80::1 link-local
 ipv6 address 2001:428:CE01:2011:10:211:0:1/64
 ipv6 enable
 ipv6 mtu 1280
 ipv6 nd autoconfig prefix
 ipv6 nd autoconfig default-route
 ipv6 nd cache expire 14400
 ipv6 nd prefix default 1800 1400
 ipv6 nd other-config-flag
 ipv6 nd router-preference High
 ipv6 nd ra interval 150
 ipv6 dhcp relay destination 2001:428:CE01:730:10:63:231:72
 ipv6 ospf 100 area 0
 ipv6 ospf authentication ipsec spi 3735928559 sha1 <removed>
```

Switching Backbone

Cisco Catalyst 6500 switches in the core and distribution layers powered the network at the San Diego Convention Center (see Figure 1). The backbone of the network was running over a 40 Gigabit Ethernet connection in a fully redundant VSS deployment. With dual-stack capability, the network also showcased the feature richness of the Supervisor 2T with Flexible NetFlow, Medianet performance monitor and NAM-3 integration.

The Network Core

The core network consisted of two Catalyst 6500 switches in VSS mode with Supervisor Engine 2Ts, running the full suite of Layer 3 features with BGP peering sessions set up to the colocation sites. These provided external connectivity to the show. The XL version of both supervisor and line cards were used (VS-S2T-10G-XL, WS-X6904-40G-2TXL, WS-X6908-10G-2TXL and WS-X6848-TX-2TXL) to accommodate larger routing tables. Within the internal network, OSPF was used to route between the distribution blocks. The core switches included 10G/40G modules for connectivity and NAM-3 for network analysis. The core switch had redundant connections to the Cisco ASAs for security, and a VSS – vPC EtherChannel to the Nexus 5000s that aggregated the servers.

Distribution Switches

The distribution layer consisted of the following switches:

- A Catalyst 6500 VSS pair with Supervisor Engine 2Ts powering the San Diego Convention Center
- A Catalyst 6500 VSS pair with Supervisor Engine 2Ts powering the World of Solutions Expo hall (Figure 1)
- Two Catalyst 6500 Switches with redundant Supervisor Engine720s, powering Hilton and Marriott hotels respectively

The distribution block also connected to ten Cisco 5508 Wireless LAN Controllers, which provided CAPWAP termination for all wireless users in the network.

Switching Backbone Feature Highlights

40 Gigabit Ethernet at Cisco Live

As mentioned previously, the network core featured the recently available Catalyst 6500 40GE module, with 4 x 40GE links between the core and distribution VSS switches to form a 160GE backbone.

The EtherChannels configured on the core switch included 10G redundant connections to the Cisco ASAs (Po70 and Po80), 40G connections to the distribution (Po30 and Po40) and a 10G connection on the VSS—vPC link (Po50) to the Nexus 5000s (see Figure 5).

Figure 5. Core Supervisor 2T Channel Summary

```
CORE-6509E#show etherchannel summary

Number of channel-groups in use: 10
Number of aggregators:           10

Group  Port-channel  Protocol    Ports
------+-------------+-----------+---------------------------------------------
10     Po10(RU)         -        Te1/5/4(P)   Te1/5/5(P)
20     Po20(RU)         -        Te2/5/4(P)   Te2/5/5(P)
30     Po30(RU)         -        Fo1/1/1(P)   Fo1/1/2(P)   Fo2/1/1(P)
Fo2/1/2(P)
40     Po40(RU)         -        Te1/2/3(P)   Te1/2/4(P)   Te2/2/3(P)
Te2/2/4(P)
50     Po50(SU)        LACP      Te1/2/5(P)   Te1/2/6(P)   Te2/2/5(P)
Te2/2/6(P)
70     Po70(SU)        LACP      Te1/2/1(P)   Te2/2/1(P)
80     Po80(SU)        LACP      Te1/2/2(P)   Te2/2/2(P)
85     Po85(SU)        LACP      Gi1/8/3(P)   Gi2/8/3(P)
90     Po90(SU)        LACP      Te1/1/13(P)  Te2/1/13(P)
```

On the distribution switch, the WS-X6904-40G-2T line card was set up to work in split-brain mode, with the 40GE uplinks to the core and the 10GE links to the access on the same module (see Figure 6). Ports 1 and 2 (on the left) of the module belong to the port-group 1 and remained in the default 40G mode. Ports 3 and 4 (on the right) of the module were set into 10G mode, using the configuration below for both switches of the VSS. They were then equipped with FourX adapters to form 4 x 10GE ports, and used with regular 10G SFP+ optics to support the access switch connections.

Figure 6. WS-X6904-40-2T 10 Gigabit Ethernet Port Operational Mode

```
hw-module switch 1 slot 1 operation-mode port-group 2 TenGigabitEthernet
hw-module switch 2 slot 1 operation-mode port-group 2 TenGigabitEthernet
```

More information regarding the 40G mixed modes can be found in Cisco Catalyst 6900 Series 40 Gigabit Ethernet Interface Module for Cisco Catalyst 6500 Series Switches Data Sheet.

VSS and vPC Interoperability

The Catalyst 6500 core was set up in VSS mode and connected to a pair of Nexus 5548 switches using a VSS–vPC connection. Best practices for interoperability were adopted according to Cisco Catalyst 6500 VSS and Cisco Nexus 7000 vPC Interoperability and Best Practices.

- The Multichassis EtherChannel (MEC) links on the VSS and the vPC member links were bundled as recommended, using dual homing and a single port-channel ID for all the four member ports.
- The EtherChannels were configured using LACP, with mode active on both ends.
- Dual active detection was enabled on the VSS using fast-hello keepalives (This was configured on all VSS systems in the network). The vPC setup included the vPC keepalive link to detect a dual-active scenario.

Dual Stack Capability

The network was enabled for end-to-end native IPv4 and IPv6 connectivity. Features were configured for IPv6 on the switching backbone as well, including BGP, OSPF, Flexible NetFlow and Medianet. More details are in the IPv6 section of this paper.

PIM SSM

The Catalyst 6500 Switches were enabled with PIM Source Specific Multicast (SSM) to support the Cisco TV deployment at the SDCC. Each display was equipped with Cisco Digital Media Players (DMP). The SSM feature is an extension of IP Multicast, where datagram traffic is forwarded to receivers from only those multicast sources to which the receivers have explicitly joined, significantly reducing the amount of multicast traffic on the network. With PIM SSM, there is no shared tree and an RP is no longer needed, which makes multicast deployments more scalable. PIM SSM requires the client to explicitly define the source in its requests, and that all components support IGMPv3. Unfortunately, at the time of the show, the DMPs were not supporting IGMPv3 and the team used the Catalyst 6500 to translate the (*,G) requests sent by the DMPs to (S,G) forwarded to the video server (see Figure 7).

Figure 7. PIM SSM Configuration

```
interface Vlan352
 description Hilton-DMP-2
 ip address 10.35.2.1 255.255.255.0
 ip pim sparse-mode
 ip igmp version 3
<SNIP>
ip igmp ssm-map enable
no ip igmp ssm-map query dns
ip igmp ssm-map static Cisco_TV 10.35.1.11
ip pim ssm range SSM
<SNIP>
ip access-list standard Cisco_TV
!
ip access-list standard SSM
 remark Cisco_TV
 permit 239.192.1.8 0.0.0.7
 permit 239.100.1.0 0.0.0.255
```

```
ip access-list standard test
```

Live and on-demand video was displayed throughout the SDCC using 60 DMPs connected into the distribution switches. The Cisco TV traffic accounted for an amazing 1 tera byte of multicast traffic in the network, over the duration of the conference.

Figure 8. DMP Multicast Traffic

```
Vlan351 is up, line protocol is up
<SNIP>
L3 in Switched: ucast: 16420316 pkt, 3340662560 bytes - mcast:777825225 pkt,
1059397956450 bytes
```

Flexible NetFlow

The Supervisor Engine 2T was equipped with Flexible NetFlow configurations to collect NetFlow data over both IPv4 and IPv6. The NetFlow information was used by a variety of tools including the NAM-3, Prime Assurance and by the Lancope booth. This feature can be used to analyze both data plane as well as control plane traffic through the switches. As an example, a Flexible NetFlow monitor applied on the control plane interface was used to identify flows accounting for high CPU utilization on the switch. These “top talkers” could be detected either through the CLI on the switch or through the NAM3 interface, and then rate limited where necessary using Control Plane Policing (CoPP).

Performance Monitor

Performance Monitor is a NetFlow-based feature that measures user traffic performance, generates alerts based on thresholds, and reports through multiple management interfaces. See Figure 9 for the configuration. The team used the feature to monitor audio and video traffic, generating syslog events when the packet loss for these classes of traffic reached a certain threshold. The syslog events were forwarded to our management station (Cisco Prime LMS).

Figure 9. Performance Monitor Configuration

```
flow record type performance-monitor RTP
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 match transport rtp ssrc
 collect routing forwarding-status
 collect ipv4 dscp
 collect ipv4 ttl
 collect transport packets expected counter
 collect transport packets lost counter
 collect transport packets lost rate
 collect transport event packet-loss counter
 collect transport rtp jitter mean
 collect transport rtp jitter minimum
 collect transport rtp jitter maximum
 collect interface input
 collect interface output
 collect counter bytes
 collect counter packets
 collect counter bytes rate
 collect counter packets dropped
 collect timestamp interval
 collect application media bytes counter
 collect application media bytes rate
 collect application media packets counter
 collect application media packets rate
 collect application media event
 collect monitor event
<SNIP>
policy-map type performance-monitor voice-video
  class DSCP_EF
   flow monitor RTP
   monitor parameters
    flows 16
   react 2 transport-packets-lost-rate
    threshold value gt 2.00
    alarm severity critical
    action syslog
  class DSCP_AF41
   flow monitor RTP
   monitor parameters
    flows 16
   react 2 transport-packets-lost-rate
    threshold value gt 2.00
    alarm severity critical
    action syslog
```

MediaTrace

The Supervisor Engine 2T supports MediaTrace, a powerful tool that can follow a particular flow's path and gather various layers of information. MediaTrace uses the IP header of the flow to be traced, and it provides a much better path congruency than a traditional traceroute. The MediaTrace will also not only discover the routers (as with traceroute), but also switches that are only doing Layer 2 forwarding. For more information on Performance Monitor and MediaTrace please refer to Cisco IOS Performance Monitor and Mediatrace QuickStart Guide.

EEM

The Supervisor Engine 2T supports EEM version 3.0. Equipped with an EEM script, which was also tied to a Twitter feed (see Figure 10), the core and main distribution switches were automated to collect network statistics on a periodic basis and post them in real time during the show to a Twitter account (http://twitter.com/CiscoLive2012). The statistics collected included the number of routes, the number of MAC addresses learned by the distribution switch, the number of IPv6 neighbors, etc.

Figure 10. Example of Tweet from the SDCC Distribution Switch

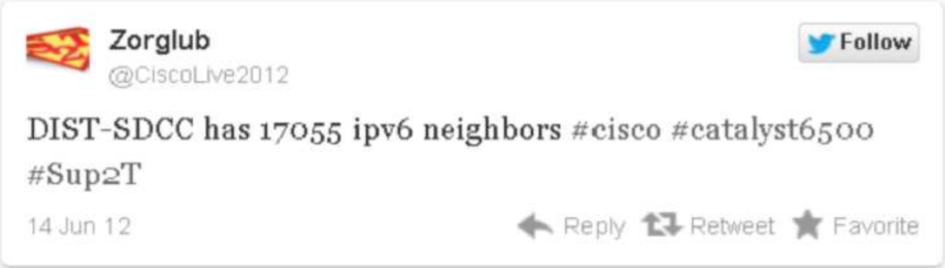

Zorglub
@CiscoLive2012

DIST-SDCC has 17055 ipv6 neighbors #cisco #catalyst6500 #Sup2T

14 Jun 12

The Twitter's API used during the show is available here: https://supportforums.cisco.com/docs/DOC-19363.

EEM was also used to automate the description of equipment connected on the access switches. Using a simple EEM applet (see Figure 11), the team changed the description of the interface to the device name and port ID of the remote end. The applet itself relies on Cisco Discovery Protocol to identify the remote device. This helped the team track cabling changes and reconfigurations during the event.

Figure 11. Interface Description EEM Applet

```
event manager session cli username nmsuser
event manager applet update-port-description
event neighbor-discovery interface regexp GigabitEthernet.* cdp add
action 1.0 cli command "enable"
action 2.0 cli command "config t"
action 3.0 cli command "interface $_nd_local_intf_name"
action 4.0 cli command "description $_nd_cdp_entry_name:$_nd_port_id"
```

Example of the script above in action:

```
SDCC_IDF_1.19#show interface GigabitEthernet0/14
GigabitEthernet0/14 is up, line protocol is up (connected)
<SNIP>
  Description: SDCC_DMS_12.show.ciscolive.com:GigabitEthernet0/10
```

EnergyWise

EnergyWise is a technology that helps manage, monitor and optimize the power consumption of network connected devices (see Figure 12 for configuration example). An EnergyWise domain was configured and all network equipment reported their power consumption to the JouleX management station. We reported a total power consumption of 1272.87 kWh for the duration of the show.

Figure 12. EnergyWise Configuration

```
energywise domain SDCC security shared-secret <removed> protocol udp port 43440
ip 10.34.100.1
energywise role switch
energywise management security shared-secret 0 <removed>
energywise allow query save
```

```
   energywise endpoint security none
```

IPv6

The network provided connectivity for both IPv4 and IPv6. IPv6 connectivity was provided to conference attendees using Stateless Address Autoconfiguration (SLAAC). DHCP was used to provide only the DNS addresses. Some adjusting of the Neighbor Discovery (ND) protocol was required to avoid some issues with mobile devices roaming aggressively throughout the facility. The MTU was set at 1280 to avoid potential issues with sites that tunnel IPv6 and do not have working path MTU discovery. Although MTU issues are rare and typically isolated, they are difficult to troubleshoot, especially in the very limited lifetime networks like the network at Cisco Live. This MTU setting was specifically tailored to the Cisco Live environment. In the typical IPv6 environment, the MTU would be left as it is by default, and the focus would be on troubleshooting the root cause of the PMTUD problems as they arise–because the network is more stationary–and it is always better to troubleshoot the root cause.

The same type of adjustment was performed on the Router Advertisement (RA) timers and prefix lifetimes–this setting was a necessity due to the way some of the mobile clients were interacting with the code on the wireless controller. A noticeable portion of the wireless devices when using SLAAC were generating temporary addresses. Essentially this means every time a wireless device joined the network, it received a new address. This can potentially overflow the binding table in the wireless controller, which is limited to eight addresses per host. Therefore, relatively short lifetimes were a necessity–alongside a very aggressive RA timer.

If the stateful DHCPv6 for address assignment was used, it would have been possible to have stable addresses for the endpoints and also much less aggressive timers but at present, a significant portion of the devices still do not support stateful DHCPv6 – and for Cisco Live San Diego, the team chose to go with SLAAC. This may change for the subsequent Cisco Live events, as after reviewing the data, the team concluded that the performance overhead of extra ND traffic outweighed the benefits of IPv6 connectivity.

For more information on the Cisco Live IPv6 deployment please visit the Cisco Borderless Networks Blog: IPv6 at Cisco Live! San Diego.

Security

Security was implemented using two Cisco ASA 5585-X-S60s connected in active/standby (A/S) redundancy (see Figure 14). Each firewall was populated with assecurity services processor, performing Intrusion Prevention Services (IPS).

The firewalls were protecting two key services. First, the convention center network was protected from the Internet and second, the FlexPod services were protected from entities within the convention center network.

The firewalls operated in multiple context modes, which enabled the allocation of virtual firewall contexts for each major logical security function; a security context was created for Internet traffic and one for FlexPod traffic.

The advantage of this was efficient use of the installed hardware and simpler security policies for each area. The FlexPod firewall policy only had to deal with FlexPod specific policies. There was no need to mix Internet policies with FlexPod policies. Likewise, Internet policies did not deal with FlexPod policies, only Internet policies.

The Internet firewalls were also providing network address translation services for traffic going in and out of the convention center network. This was necessary since it was anticipated that up to 30,000 wireless devices would be present on the network, but only 2000 public IP address were allocated. We provided Port Address Translation

(PAT) pools so a single external address within the address pools would be overloaded by a factor of 50:1. The team also utilized a carrier grade feature that allowed multiple connections from single hosts to be mapped to the same external PAT address (see Figure 13).

Figure 13. CGN Configuration on the Cisco ASA 5585

```
object network Wireless_1_RANGE
 range xx.231.216.20 xx.231.218.110
object network Wireless_1
 subnet 10.211.0.0 255.255.0.0
 nat (inside,outside)
  dynamic pat-pool Wireless_1_RANGE flat include-reserve round-robin dns
```

The firewalls were also deployed in transparent mode, which allowed the team to deploy the firewalls as invisible bumps on the wire security policy devices. The advantage of transparent mode was that the team could offload all network functions, such as dynamic routing and multicast routing, to the surrounding Catalyst 6500 switches.

Figure 14. Internet Connection and Firewall Architecture

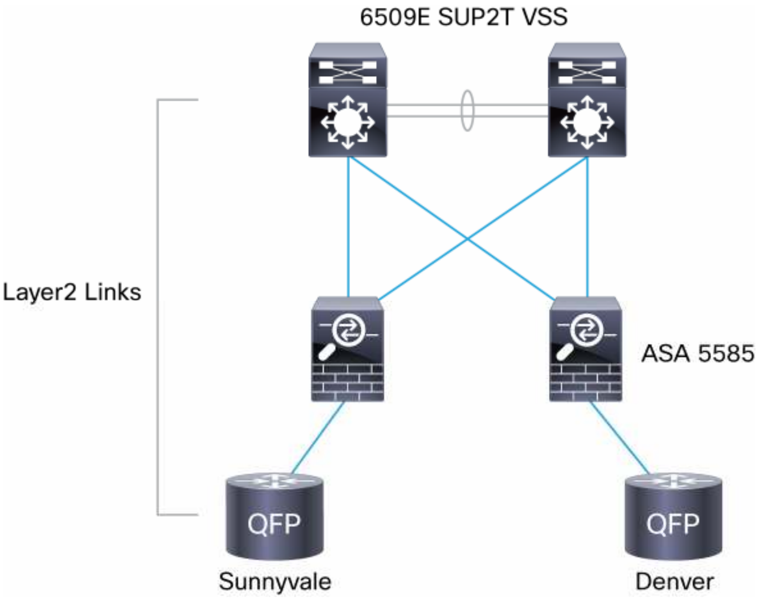

Wireless

The conference was supported by ten Cisco 5508 wireless controllers and over three hundred Cisco 3600 Series Access Points (APs) using Cisco Clean Air technology. The APs provided connectivity on 2.4 Ghz and 5 Ghz. All wireless users were aggregated on a single /16 subnet, to allow for roaming between the different locations. The monitoring of the wireless network was done using Cisco Prime Network Control System (NCS). See Tables 1 and 2 and Figures 15-17 for wireless traffic statistics on the Cisco Live San Diego network. The majority of the hosts (99%) were dual stack capable. The low number of IPv6 only clients is due to the fact that the IPv6 only SSID was provided to limited conference rooms.

Table 1. Wireless Client Summary by IP Address Type

IP Address Type	Average Number of Sessions	Maximum Number of Clients	Average Number of Clients
IPv4	27165	11257	4467
Dual-Stack	26919	9300	3757
Not Detected	1054	2424	779
IPv6	918	2074	698

Figure 15. Average Number of Wireless Clients by Type

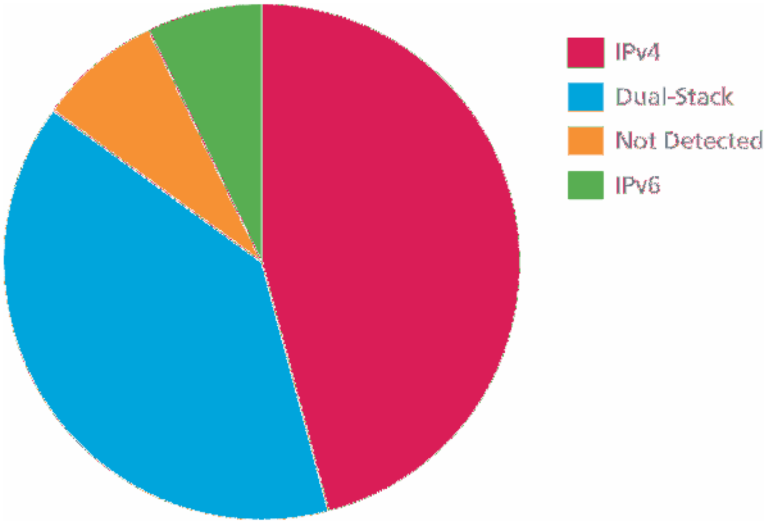

The total amount of traffic measured on the wireless controller per manufacturer (based on the OUI) is shown in Figure 16.

Table 2. Wireless Traffic by Vendor (in MB)

Vendor	
Apple	800.79
Asustek	9.88
Azureware	5.24
Cisco	0.07
Gemtek	6.04
Hewlett-Packard	0
Hon Hai Precision	45.04
Htc	8.25
Intel	485.59
Liteon	19.95
Motorola	9.4
Murata	5.21
Nokia	0.2
Others	24.15
Palm	0.64

Vendor	
Private	0.77
Research in Motion	5.32
Samsung	13.93
Sony	0.38
Unknown	100.28

Figure 16. Wireless Traffic by Vendor (in MB)

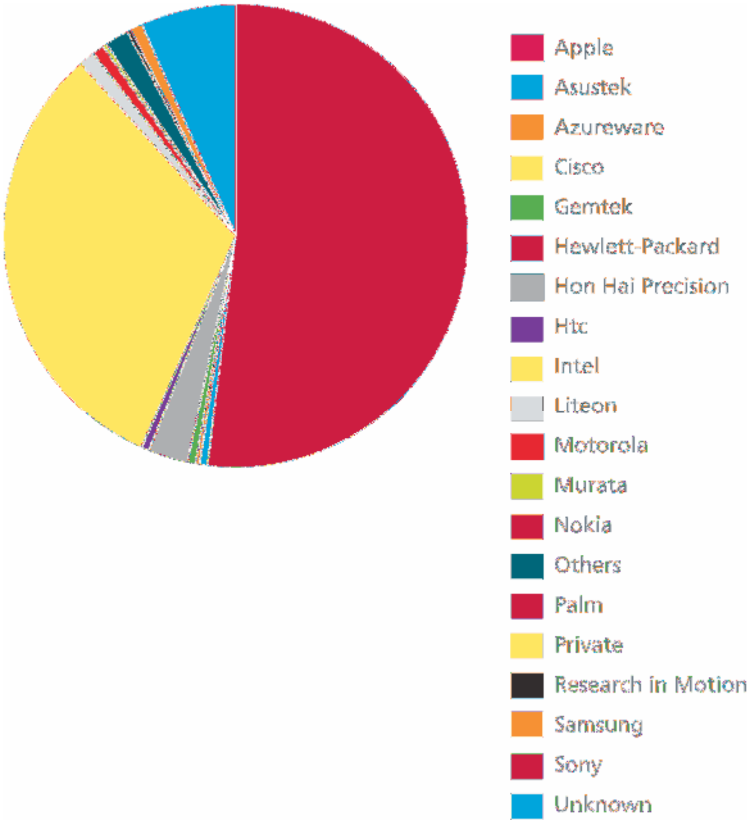

Figure 17. Associated Wireless Client Count in the Network Over the Course of Cisco Live

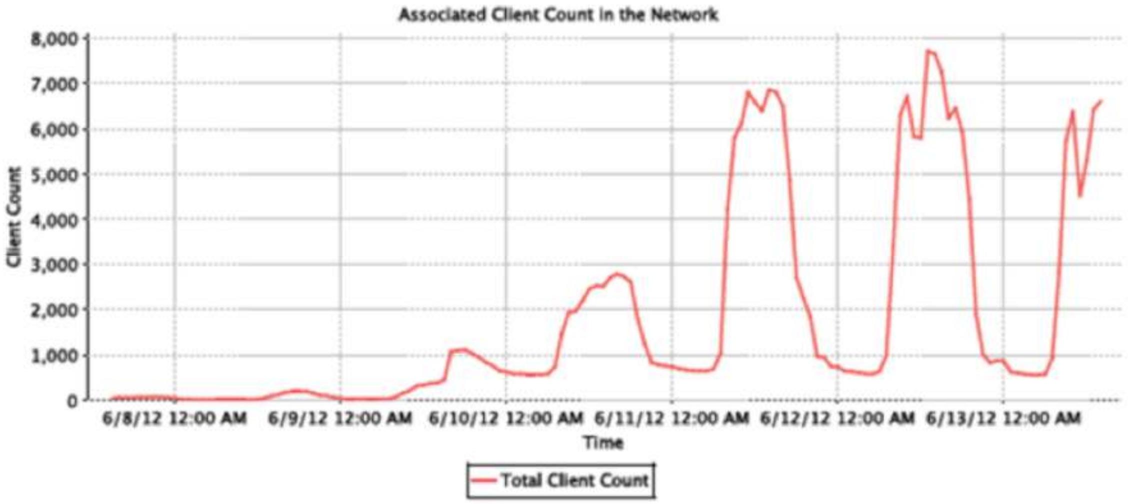

Figure 18. Associated Client Count by IP Address Type During the Course of Cisco Live

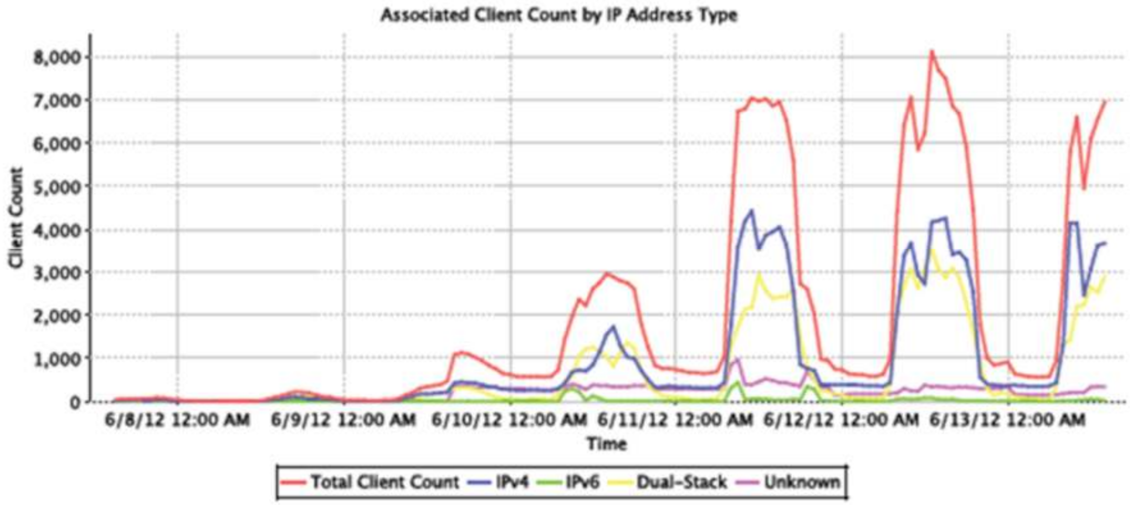

Servers

All services (DHCP, DNS, Network Management, etc.) were hosted within a FlexPod, an architecture designed jointly with NetApp and virtualization software providers that provisions a flexible environment with increased efficiencies and reduced risks (Figure 19).

This particular FlexPod architecture is a predesigned configuration that is built on the Cisco Unified Computing System (UCS), the Cisco Nexus family of data center switches, NetApp FAS storage components, and VMware virtualization software.

Within the FlexPod, the UCS manager performs the role of abstracting hardware resources into service templates that simplifies configuration and facilitates operational management, and the NetApp OnCommand Insight provides a holistic view of the storage infrastructure as a unified set of services.

VMware vSphere, coupled with the NetApp Virtual Storage Console (VSC), serves as the foundation for VMware virtualized infrastructures. The VMware vCenter domain manages and provisions resources for all the ESX hosts in the data center. For more information on the FlexPod architectures, go to Data Center Designs: Cloud Computing, DesignZone for FlexPod.

Figure 19. FlexPod Architecture

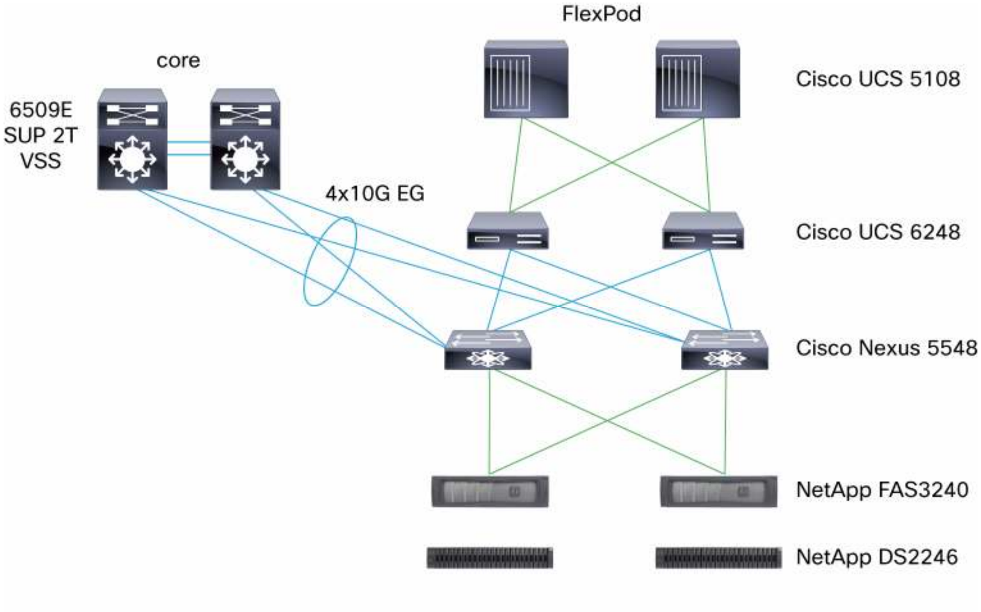

IP Telephony

Two redundant Cisco Unified Called Managers (CUCM) were used, one of which was located within the Centurylink colocation in Sunnyvale and the other was on the Cisco Live show floor. A Cisco Unified Border Element (CUBE) was used to terminate the PSTN services in Denver. Since the phones were residing behind the NAT on private address space, a second CUBE was added and a SIP trunk was established between the two (Figure 20).

Figure 20. Telepresence and IP Telephony Architecture

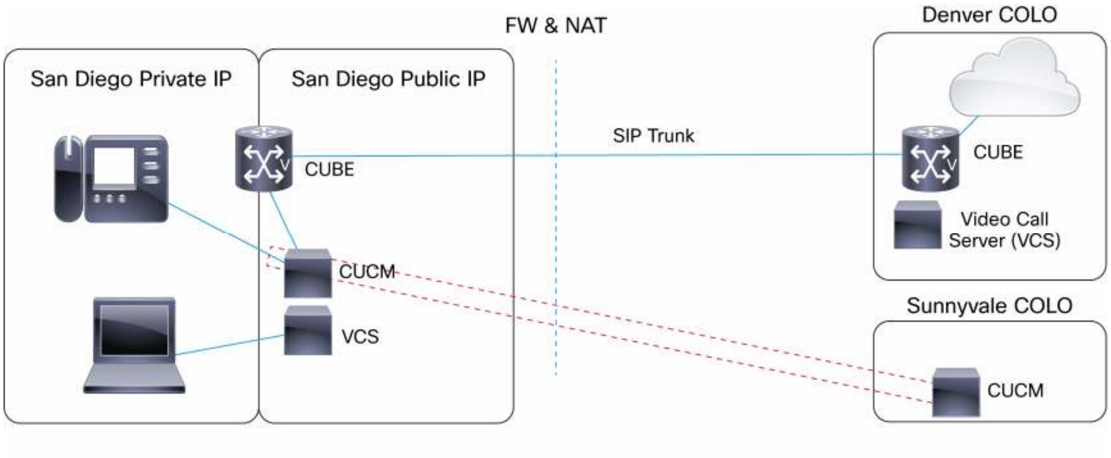

Services

DNS/DHCP

Cisco Prime Network Registrar was running on a virtual machine within the FlexPod. The team had two types of DNS servers: two authoritative servers used for inbound traffic and servicing four domains (noc.ciscolive.com, show.ciscolive.com, denver.ciscolive.com and sun.ciscolive.com) and two caching DNS servers sitting on the network operations center (NOC) internal VLAN for servicing conference attendees (Figure 21).

Figure 21. DHCP and DNS Architecture

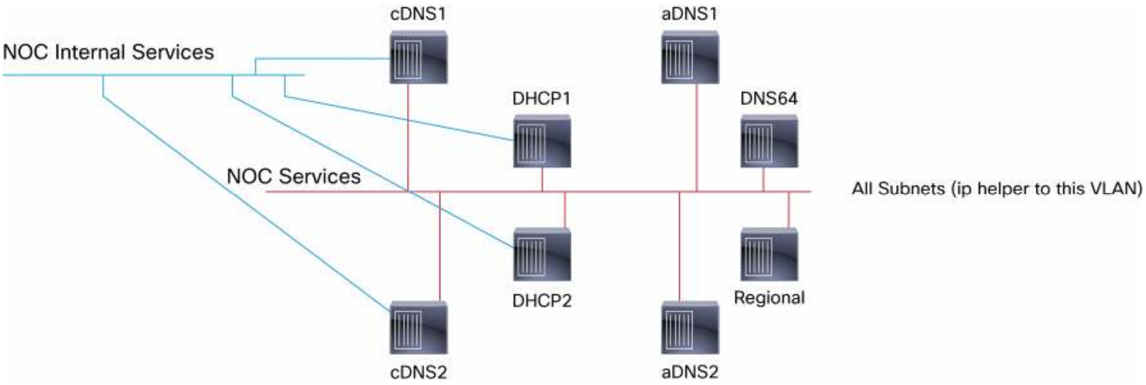

Management

Cisco Prime LMS 4.2 was used during the show to monitor devices and services, and also for provisioning the access switches with template center. Smart Call Home was used to report any issues that occurred on the access switches.

Session Recording

For those who couldn’t attend the event in San Diego, the team decided to extend live recordings of breakout sessions. The live content was collected for 270 sessions in 36 separate rooms and is available at: https://www.ciscolive365.com. To provide this service, the team used a Cisco TelePresence Codec C90 interfaced with high definition cameras, microphones and video input from presenter laptops. These high definition streams were sent over the Cisco Live show network to the TelePresence Content Servers (TCS), where everything was recorded. As soon as the sessions were over, the TCS immediately transcoded the session to specified output settings and a custom template. The TCS created a composite MP4 file that is available on the ciscolive365 web site. The Video Communication (VCS) was also used to upload the content to the Content Delivery System (CDN). A total of 600 GB of data was generated by the session recording processes. The Telepresence Management System (TMS) was used to automate the start and stop of the sessions based on an uploaded schedule.

Figure 22. Session Recording Solution

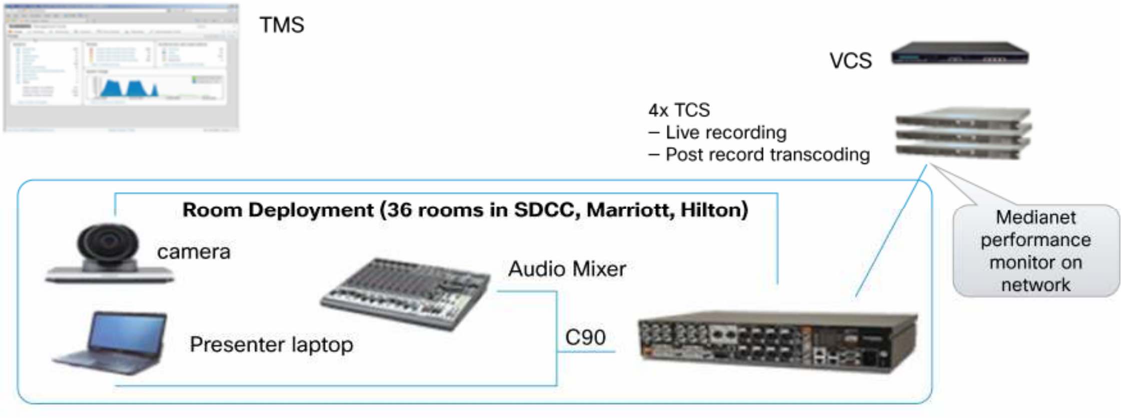
TMS
VCS
4x TCS
– Live recording
– Post record transcoding
Room Deployment (36 rooms in SDCC, Marriott, Hilton)
camera
Audio Mixer
Medianet performance monitor on network
Presenter laptop
C90

Special Services

NAT64/DNS64

The team used NAT64 on the Cisco ASA, with a Cisco ASR 1000 router in warm-standby as a backup. Both of the devices were in a “two-armed” setup—with an IPV4-only arm on one side, and an IPv6-only arm on the other side. The stateful NAT64 had different addresses used for the overloaded pool of IPv4 – this took care of the return routing in case of a failover event.

The failover was designed by means of routing. The next hop on the IPv6 side of the translators had two routes setup: one route for the NAT64 prefix with the subnet mask of /96, pointing towards the ASA, and the other route for the NAT64 prefix with the subnet mask of /95, pointing to the Cisco ASR 1000 router (Figure 23). This meant that the failover could be performed in a very efficient fashion by a single command. To avoid any loops, the Cisco ASR 1000 had a null route for the entire /64 which was dedicated for the NAT64.

Figure 23. NAT64 Static Routes Configuration

```
ipv6 route 2001:428:CE01:CAFE::43CA:6B20/128 2001:428:CE01:3031::4
ipv6 route 2001:428:CE01:CAFE::/96 2001:428:CE01:3031::2
ipv6 route 2001:428:CE01:CAFE::/95 2001:428:CE01:3031::4
ipv6 route 2001:428:CE01::/48 Null0
```

Every client in the network could use the NAT64, provided they had the corresponding function of DNS64 to create the synthetic AAAA records in the replies. This functionality was provided by the CNR, with the backup of BIND9 – so as soon as the clients were authenticated, they would automatically start using NAT64.

Because NAT64 is the transition mechanism for the IPv6-only networks, this functionality was primarily used by the experimental IPv6-only SSID. IPv6-only wireless networking on many portable devices is still a work in progress, so the team kept this SSID hidden by default, to minimize potential user confusion and advertised the service only to attendee of IPv6 sessions (Figure 24 and 25).

During the preparation and testing, the team discovered that certain IPv4-only applications trigger a burst of ARP messages from the devices on the IPv6-only network – the team decided to go with the “synthetic” 100.64.0.0/16 network, which was blocked by an access list at the first hop. The RFC6555 implementations in the end hosts would take care of not using IPv4, by keeping the experience mostly transparent to the end user.

Figure 24. IPv6 Only Wireless Vlan Configuration

```
interface Vlan2014
 description Wireless_4_IPv6_Only
 ip address 100.64.0.1 255.255.0.0
 ip access-group v6-only-block in
 no ip proxy-arp
 ipv6 address FE80::6 link-local
 ipv6 address 2001:428:CE01:2014:10:214:0:1/64
 ipv6 enable
 ipv6 mtu 1280
 ipv6 nd autoconfig prefix
 ipv6 nd autoconfig default-route
 ipv6 nd cache expire 14400
 ipv6 nd prefix default 1800 1400
 ipv6 nd other-config-flag
 ipv6 nd router-preference High
 ipv6 nd ra interval 150
 ipv6 dhcp server v6-only
 ipv6 ospf 100 area 0
```

Figure 25. IPv6 Only Access-list

```
ip access-list extended v6-only-block
 permit udp any any eq bootps
 permit ip host 10.214.0.18 host 10.214.0.1
 permit ip host 100.64.0.16 host 100.64.0.1
 permit udp any host 63.231.220.84 eq domain
 permit tcp any host 63.231.220.84 eq domain
 deny   ip any any
```

IPv6-only wireless networks are still in their infancy, and the experiments like this help engineers advance the state of the art in this area – as a consequence of the Cisco Live 2012 network, there were very successful follow-ups with the other vendors, resulting in improvements to the behavior of their IPv6 stacks.

Conclusion

With more than 20,000 registered attendees, the show broke some records. A total of 51.7 terabits of traffic crossed the internet connection. This was an increase of 291.6 % from Cisco Live 2011.

The break up was the following: 50.5 terabits of IPv4 and 1.2 terabits of IPv6, so even if the number of dual-stack hosts was high (see section 6.1.1) the amount of IPv6 traffic was only 2.37 % of the total which is still an increase of 400 % from Cisco Live 2011.

Traffic peaked on Tuesday 6/12 when the rate reached a maximum of 410.9 Megabits/s over the Internet connection. That day also saw the highest number of DHCP leases on the wireless with a peak at 10,011.

In terms of applications, http was the on top followed by https, imaps, BitTorrent and dns.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-721661-00 12/12