

# IPv6 Network Infrastructure Overview in Cisco Live London 2012

**Last updated: June 2012**

Cisco Live is considered a flagship event that gathers Cisco partners and customers worldwide, and is also a premier showcase for new Cisco technologies. The event runs for five days and offers a combination of education and training on the latest technologies and trends; testing and certification on Cisco products; and the opportunity for attendees to extend their professional networks by meeting industry peers.

**Figure 1.** Cisco Live London 2012 Network Fact Sheet

- >8500 end user devices
  - About 75% of all end user devices enable IPv6
  - 3600 rogue clients and 2280 rogue APs
- Over 26km (90k ft!) of cabling
- Network Infrastructure
  - 2 x Cat 6500
  - 15 x Cat 3750X
  - >200 access switches
  - >130 wireless APs

Cisco Live London 2012 was powered by a network running both IPv4 and IPv6 protocols, supporting over 8500 end-user devices over wired and wireless connections. During any flagship event—such as Cisco Live—it is paramount to provide a scalable, reliable, secure and manageable environment that provides vital connectivity for attendees, exhibitors and Cisco staff delivering techtorials, breakout sessions, customer demonstrations and hands-on labs.

The drivers for IPv6 are well known and given the size and scope of the event, IPv6 enablement was a natural evolution of the design used in previous years. Each year sees an increasing number of devices per user, and as of Cisco Live London 2012, the majority of end user devices are IPv6-enabled, whether the users are aware of it or not. Providing native IPv6 connectivity helped reduce the administration complexity, offered better scaling and inevitably resulted in better service to end-users.


## Design of the Network and Addressing Plan

As the event took place in the Excel Center in London for the second time, the same physical topology was re-used from Cisco Live 2011. But while in 2011 the network was IPv4, for 2012 full IPv6 support was added in preparation for important events such as the World IPv6 launch.

The design emulates a typical hierarchical enterprise campus network with core, distribution and access layers, in addition to a data center module (see Figure 1). Two Catalyst 6500 switches provided the core, while 15 Catalyst 3750X switches implemented the distribution. Core and distribution equipment was housed in secure equipment

rooms, while access switches were deployed as needed around the entire venue to support end-user connectivity. Nearly 200 wired and over 130 wireless access points constituted the access layer.

**Figure 2.** Cisco Live London 2012 Network Infrastructure



## IPv6 Design Considerations

The guiding principle for the IPv6 network design is reflected in the following quote by Veronika Storkova, a Systems Engineer in the UK&I IPv6 Team: "The major inhibitor of IPv6 adoption is the lack of familiarity with the technology. There are many best practices and recommendations published by Cisco which can be used to plan and prepare for a successful IPv6 deployment. The dual-stack network at Cisco Live was built, secured and operated leveraging the same best practices."

To describe the IPv6 addressing plan we first need to point out that, as no native IPv6 WAN connectivity was available by a service provider, a 6-in-4 tunnel was provided by BT and provisioned over the IPv4 WAN Links.

BT provided a /48 prefix and this was assigned in IPv6 according to the same design principles that guided the IPv4 addressing design (see next sections). The fact that Cisco Live provides the network designer with a green field environment luckily eliminates some constraints that may otherwise limit the options for an enterprise network designer.

## IPv4 Addressing

The IPv4 addressing scheme was designed around the requirements for an event of this nature, namely a standard private address block (10.0.0.0/8), and then divided further to reflect the network topology and thereby ease troubleshooting. This resulted in a 10.X.Y.H/24, where X represents the VLAN, Y the distribution switch, and the final H octet was assigned to the host. The one exception to this rule was for the VLANs dedicated to the

---

wireless domain, where a single octet would unnecessarily limit scalability. Thus, 10.X.H.H/16 was dedicated to the wireless domain, giving the five combined wireless VLANs the ability to support –in theory- over 300k end devices.

## IPv6 Addressing

There were two major design decisions that governed the assignment of IPv6 addresses:

- Global transparency by reusing the prefix allocated to the 6-in-4 tunnel by BT.
- The ability to easily derive the IPv6 address from the IPv4 address, and vice-versa. While this requirement may result in sub-optimal use of address space, the inherent ease of use represents a huge operational advantage in an event such as Cisco Live.

The IPv6 address space was structured in the following format (hex): 2A00:2000:4004:XXYY:H/64

- The first three 16-bit pieces of the address (2A00:2000:4004) represent the prefix assigned by BT.
- XXYY stands—just as in the IPv4 addressing space- for the VLAN (XX) and the distribution switch (YY). Note that hex conversion needs to take place for translation between IPv4 and IPv6 addresses, meaning that –for example an IPv4 VLAN-id of 64 would result in an IPv6 VLAN-id of 40.
- The 64-bit host portion of the address was self-assigned via SLAAC (Stateless Address Auto-Configuration), depending on the requirements of the supported device.

## Visual Mapping of IPv4 to IPv6 Addresses

Despite the straight forward nature of translation between the IPv4 and IPv6 address spaces, a very helpful tool proved to be a spreadsheet that provided immediate translation between IPv4 and IPv6 addresses and vice-versa.

### *Other considerations*

#### Security

Just because Cisco Live takes place outside of the primary enterprise parameter does not in any way mean that security can be neglected as a key design consideration. Therefore, best design practice security measures were enabled throughout the network. The primary line of defense is the access layer, of course, with the Layer 3 distribution layer also being thoroughly secured.

When it comes to rogue devices, it does not really matter if such behavior comes about through inadvertent misconfiguration or malicious intent. In either case, the online experience of event attendees can be severely impacted, and the entire event disrupted, if the network is brought down by security breaches.

The number of rogue clients and endpoints –over 5000 combined!- experienced at Cisco Live London 2012 made first hop security of paramount importance. Dynamic Host Configuration Protocol (DHCP) attacks and false route advertisements were successfully repelled, and the integrity of the wireless network maintained throughout the event.

This was accomplished through various IPv6 security best practices, including:

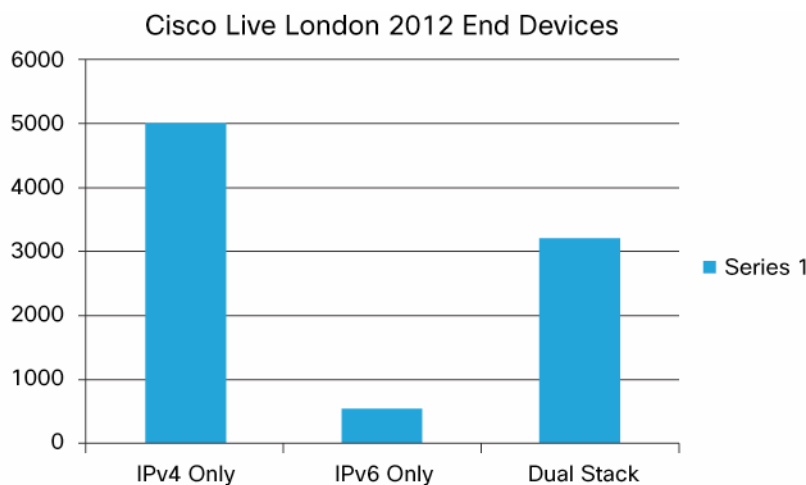
- Disabling Route Advertisement (RA) announcements on Ethernet infrastructure links.
- Disabling IPv6 redirects and IPv6 unreachables on infrastructure links.
- IPv6 Unicast RPF (uRPF) checking.

- Restricting IPv6 VTY access.
- Deploying First Hop Security (FHS) on the access switches by filtering out rogue DHCPv6 or RA packets.

## Operating the Network

As mentioned in Figure 1, over 8500 clients were connected to the network. A very interesting fact was the finding that, to most users, the operation of IPv6 was totally transparent. Users' devices autonomously tend to configure dual stack capabilities. In fact, about 75% of users' devices were IPv6 enabled, while only about 20% of users reported awareness of the presence of an IPv6 stack on their devices.

**Figure 3.** Number of Connected Devices



One of the reasons why the experience may have been so transparent to IPv6 users may have been the performance. RTT differences between IPv4 and IPv6 websites were negligible; despite the often repeated preconception that IPv6 traffic incurs higher delays.

Furthermore, and as a positive side effect of IPv6, when a shortage of addresses because of initial misconfiguration of NAT impacted IPv4 address space on the first day of the event, IPv6 users remained unaffected and were able to reach the Internet.

On the other hand, this is not to say the IPv6 environment was totally free of issues. The combination of a number of new technologies (Bring Your Own Devices (BYOD), new features on wireless controllers etc.) resulted in non-trivial trouble-shooting. Thus, the best way to prepare for IPv6 is to start deploying it now to gain valuable operational experience.

## Conclusion

The large and highly successful deployment of IPv6 during Cisco Live London 2012 served to emphasize that IPv6 has become a mature technology to be deployed in enterprise environments. At this stage, enterprises should start implementing their IPv6 strategy. Good planning, knowledge of the technology and testing are all vital for a successful deployment.

To sum it up, a quote by Ian Foddering, the CTO of Cisco UK&I, puts it succinctly: "IPv6 is an evolutionary step to support the next wave of applications and Internet connected devices. The Cisco Live dual-stacked network in

---

Cisco Live London 2012 demonstrated that IPv6 is ready for deployment in production environments, both within service provider and enterprise infrastructures.”

For additional information on IPv6, please visit: <http://www.cisco.com/web/solutions/trends/ipv6/index.html>



---

**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](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA

C36-710810-01 07/12