Cisco Carrier-Grade IPv6 (CGv6) Solution

Delivering on the future of the Internet

October, 2009

Contents

Executive Summary

The next version of the Internet Protocol, IP Version 6 or IPv6, was invented in the early 1990s. It offered a number of improvements and optimizations over its predecessor, IPv4, including auto-configuration, header simplification, extended routing headers, and a flow labeling. The most significant improvement was the increase in number of bits used for addressing, from 32 in IPv4 to 128 in IPv6. The designers of IPv6 correctly envisioned that some day in the future, the number of addressable IP end-points would exceed the number of addresses supported by IPv4.

That day is coming soon. Current projections, based on IPv4 consumption rates, point to June 2011 as the date when addresses will cease to be available. The greater Internet community of service providers, vendors, and customers must address the IPv4 run-out problem now.

The obvious solution is IPv6. To get from IPv4 to IPv6, a number of different transition methods have been proposed over the years. Some of these methods have been implemented in products and some have been deployed in networks. Given that time is running out and that large portions of the Internet are not IPv6 ready (and will never be), a significant challenge confronts the service provider community.

This white paper outlines the Cisco® framework, called Carrier-Grade IPv6 Solution, or CGv6 for Service Providers. The paper describes a set of tiers, employing known, well-understood, and standardized technologies implemented in various products that will enable them to address the IPv4 run-out problem and commence an orderly, incremental, and safe transition to IPv6. A practical approach to preserve, prepare, and prosper during this transition using various solutions is discussed in detail.

Why Transition to IPv6?

The Internet is a system of globally interconnected networks, powered by the technology of Internet Protocol or IP. The growth of this system has been fueled by a standardization process that allows compatibility between all connected devices. A key function of this standardization is a mechanism to identify and locate each device, known as addressing. IP addresses are unique identifiers assigned to each device. If devices are to be able to communicate with one another, these addresses need to follow the same standard. The current version of the standardized technology is IP version 4 (IPv4). Here we examine a fundamental constraint limiting the further use of this technology on the Internet.

Figure 1. Map of the Internet 2009 (Telegeography, Cisco)

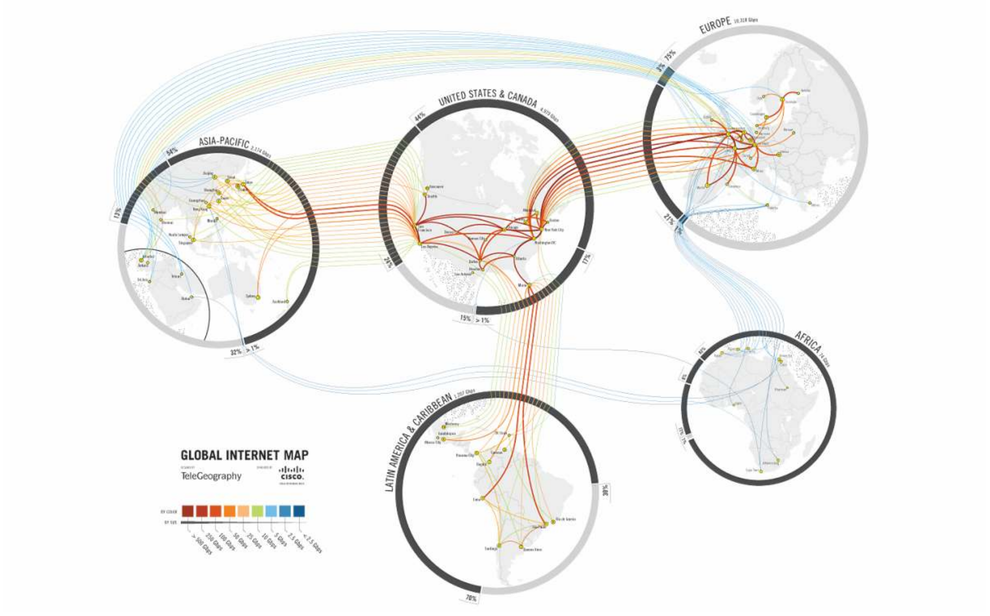

A Fundamental Constraint

The current IP address space based on IPv4 is unable to satisfy the potential huge increase in the number of users, the geographical needs of Internet expansion, or the requirements of emerging applications such as Internet-enabled wireless devices, home and industrial appliances, Internet-connected transportation, integrated telephony services, sensor networks such as RFID, smart grids, cloud computing, and gaming.

Figure 2. IPv4 Address Space

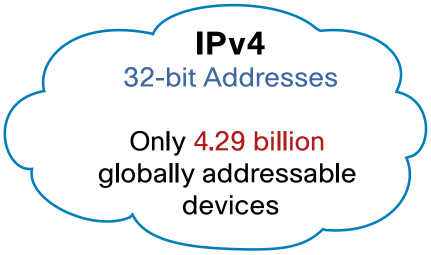

Figure 3. IPv6 Address Space

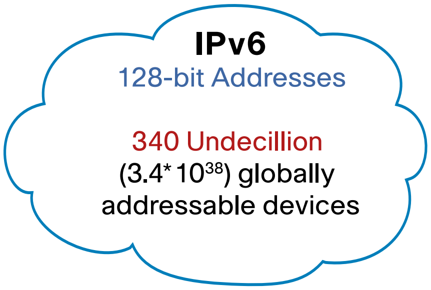

IP version 6 (IPv6) quadruples the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides more than enough globally unique IP addresses for every networked device on the planet. The use of globally unique IPv6 addresses simplifies the mechanisms used for reachability and end-to-end security for network devices, functionality that is crucial to the applications and services that are driving the demand for the addresses.

An Obstructed Journey

To get from IPv4 to IPv6, a number of different transition methods have been proposed over the years. Some of these methods have been implemented in products, some have been deployed in networks, and some do not work the way they were intended. However, today, the Internet is still largely based on IPv4, with few IPv6 implementations. Here we examine the reasons for this state of affairs.

No IPv4 and IPv6 Interoperability

IPv6 and IPv4 are completely separate protocols, with IPv6 not being backwards compatible with IPv4. Thus networks running one of the versions will not communicate with the other directly. This state of affairs is the result of the design of the IPv6 protocol stack, which was standardized sometime in the mid-1990s. As a consequence, networks would need to run parallel infrastructure housing both protocols. The operational burden imposed by this scenario is one of the factors slowing the adoption of IPv6. As devices need to support the same standard to communicate, it is expected that bridging mechanisms will evolve to interoperate between IPv4 and IPv6.

Little Economic Incentive

IP addresses are not a monetized resource. Customers, including service providers, are allocated addresses by Regional Internet Registries (RIRs) essentially for free. However, justification is required for the requested allocation. Both IPv4 and IPv6 address allocations follow this model, which is predicated on the equitable and open nature of the Internet. As the free pool of IPv4 continues to be available, although in a rapid decline, there is no economic incentive for customers to demand and use IPv6.

Little Perceived Urgency

The Internet continues to grow at a rapid pace. With the availability of free IPv4 addresses declining, mechanisms to reuse them have grown. Network Address Translation (NAT) gained popularity in home networks, requiring only a single public IPv4 address for all devices. Moreover, the Internet is not controlled by any single authority; at most, best-practice guidelines are provided for its use. These factors combined to slow down the adoption of IPv6 and the impetus to deal with the known certainty of IPv4 address exhaustion.

Incompatible Infrastructures

Because IPv4 and IPv6 are not interoperable, devices need to support both protocol stacks for coexistence. However, the Internet is built on a legacy of diverse networks and devices. Some of them support the IPv4 protocol stack only. Yet others have only partial support for the IPv6 protocol stack. Most of the popular services on Internet are enabled on IPv4 today. This diverse nature of protocol support makes it harder for users to shift to IPv6, when the alternative of using only IPv4 in the immediate time frame appears operationally simple.

A New Urgency

The day when the free pool of IPv4 addresses is exhausted will soon arrive. Current projections, based on IPv4 consumption rates, point to June 2011 as the date when addresses will cease to be available. To be more specific, that is the projected date when IANA (Internet Assigned Numbers Authority) will distribute the last large block of unallocated public IPv4 address space to the Regional Internet Registries (RIRs). Service providers are being formally notified that IPv4 address space will no longer be available from the registries by mid-2010.

Figure 4. Free IPv4 Address Blocks

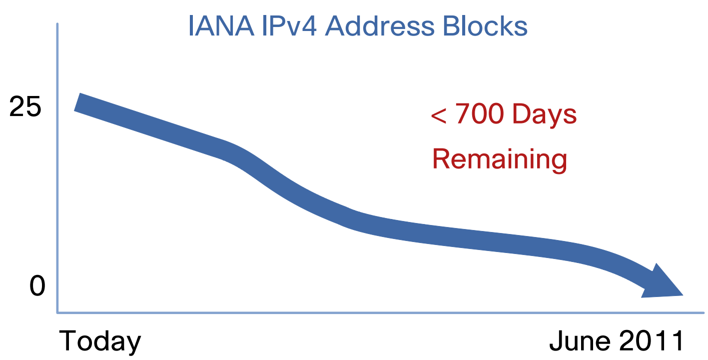

Source: Geoff Huston, http://www.potaroo.net/tools/ipv4/index.html

At the same time, the number of devices being connected to the Internet is rising exponentially, driven by the proliferation of video, mobile, and cloud computing applications. This number will soon exceed the address capacity of IPv4, even after factoring in the use of NAT within homes. At this point, growth of businesses dependent on the Internet is likely to be constrained if no action is taken. Thus, the greater Internet community of service providers, vendors, and customers must address the IPv4 run-out problem now.

Figure 5. Number of Internet Enabled Devices

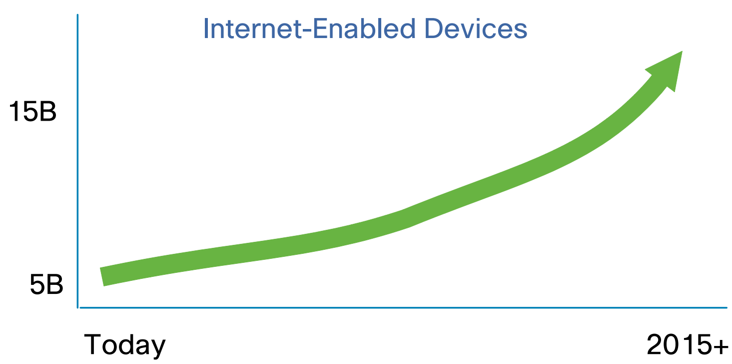

Source: Cisco VNI / Intel Embedded Internet Projections

Journey to IPv6: Cisco Leads the Way

As an early pioneer in IPv6 technology since its inception, Cisco has been a driving force in developing IPv6 standards through various standards bodies, including the Internet Engineering Task Force (IETF), and has been shipping a wide variety of end-to-end IPv6 products and solutions. Selection of a deployment strategy depends on your current network environment, and on factors such as the forecast amount of IPv6 traffic and the availability of IPv6 applications on end systems and appliances, at a given stage in the deployment.

Before Cisco Carrier-Grade IPv6 Solution (CGv6)

Most of the IPv6 protocol was standardized in the mid-1990s. Various education and research networks around the world, such as 6bone, Moonv6, and JGNv6, took the lead on IPv6. However, implementing it in provider networks only began in the early 2000s. These implementations were mainly using dual-stack or all IPv6 technologies, creating parallel infrastructures working independently of IPv4. Government mandates to move to IPv6 remained sporadic, and its adoption varied widely by geography. To date, popular end-user applications largely use IPv4. And providers extended the life of their IPv4 address pools by requiring end customers to undertake address translation on their home equipment. As the exhaustion of IPv4 nears, there is now a need for a multitiered solutions approach to help providers transition to IPv6.

Cisco is pioneering the framework for a Carrier-Grade IPv6 Solution (CGv6) with the practical approaches of preserve, prepare, and prosper.

Preserve

In this approach, the objective is to continue to use IPv4 up to and even after run-out. This approach is the logical path for a provider not ready, willing, or able to upgrade to IPv6. Solutions in this tier include **Large Scale NAT** (LSN – NAT44 or Carrier-Grade NAT), IPv4 **Address Trading**, and **IPv4 Renumbering**. This approach is a safe and low-risk path to take to address run-out in the short to medium term, while providing the operator more time to plan for IPv6. Address trading and renumbering techniques need co-operation within the industry for an equitable solution, because they affect the infrastructure's connection with multiple providers simultaneously. Best-practice guidelines for renumbering or markets for trading have yet to evolve to a mature level for deployment. Meanwhile, the LSN technique affects only the provider using it, and the technology for it (i.e., NAT) is proven.

Prepare

In this approach, the objective is to upgrade the infrastructure to support IPv6 while coexisting with IPv4. This tier is a mandatory prerequisite to prosper from the IPv6 services and applications tier. **Dual-stack** running on hosts, switches, and routers is one solution. A number of mature and emerging techniques are available for tunneling IPv6 over IPv4 (or Multiprotocol Label Switching [MPLS]), with 6PE being a well-known example. The often overlooked case is the inverse, where legacy IPv4 packets are tunneled across an IPv6 network. **Softwires Mesh** and **Dual-Stack Lite** are two examples being defined in the IETF Softwires Working Group. Also associated with tunnels are the solutions in the operator's network that **terminate** such tunnels and relay their payloads to the next-hop in the network. A stateful and stateless IPv4/IPv6 translator (aka **Address Family Translator** [AFT]), which enables IPv6 end-points to communicate with IPv4 end-points, is yet another important solution contained in this tier. The IETF BEHAVE Working Group has been formed to standardize the functions and behavior for NAT64 and NAT46 AFTs.

Prosper

In this approach, the objective is to get most of the Internet services and applications to run over IPv6. It includes a suite of existing (voice over IP, email, video) and emerging applications (smart grid, mobility, sensor networks) that will employ IPv6. Obviously these applications will not all appear and be ready at the same time; thus the need for the IPv4/IPv6 coexistence infrastructure. With a seemingly inexhaustible supply of public IPv6 addresses, every device needing an address will get one. True global reachability can be achieved for any connected Internet device, opening up new untapped revenue-yielding service opportunities.

The Journey Begins

With Cisco's multitiered approach of preserve, prepare, and prosper, service providers can begin their journey toward an all IPv6 infrastructure in an orderly, incremental, and safe manner. Implementing solutions from the preserve approach like CGN does not mean IPv6 can be avoided altogether. With the exponential growth of Internet devices, operators still have a finite amount of time that they can exploit their IPv4 infrastructure, though the immediate time pressure of IPv4 address run-out is mitigated. Additionally operators can deploy multiple tiers of solutions simultaneously. For example, NAT44 can relieve public IPv4 address consumption for the access portion of the network for residential customers, while dual-stack 6VPE routers are deployed in the backbone to support IPv6 VPN service for enterprise customers. A key question to address is the place in the network to begin this journey. Although this decision is subject to each customer's situation, general guideline can be given on introducing CGv6. This is at places where the maximum coverage can be obtained for IPv4 address run-out and where adjacent network realms abound; both these places converge on the backbone section of operator networks. Finally, the journey itself can be eased with professional services support, which undertakes assessment, planning, and deployment activities.

Cisco Carrier-Grade IPv6 Solution (CGv6)

With its wealth of IP experience, Cisco is leading the shift to IPv6 with CGv6. The multitiered approach to preserve, prepare, and prosper during the transition is based on well-known and standardized technologies. These technologies are delivered using the components of platforms and service offerings for a comprehensive solution that providers can deploy.

CGv6 Technologies

A plethora of technologies are available for service providers to move to IPv6. These technologies use the established techniques of address translation, tunneling, and encapsulation. CGv6 does not mandate a serial adoption of these technologies; customers can choose any of the multitiered approaches based on their business and technical priorities.

Double NAT 444

NAT refers to translation of one IP address into another IP address. NAPT (Network and Address Port Translation) refers to a NAT that translates multiple IP addresses on one side into a single IP address on the other side, where the TCP/UDP port number distinguishes the different packet flows. Large Scale NAT (LSN) is the service provider version of a subscriber NAT device. The latter can comfortably handle the needs of a household or small business; the former is designed to handle millions of translations, and is intended for the backbone of the provider network. LSN is not limited to IPv4 NAT, though; it is also used in the context of translating between IPv4 and IPv6. Double NAT 444 is a scenario when the subscriber uses IPv4 NAT in addition to the service provider using LSN with NAT44 within its network. Carrier Grade NAT (CGN) is a synonym for LSN.

Figure 6. Double NAT 444

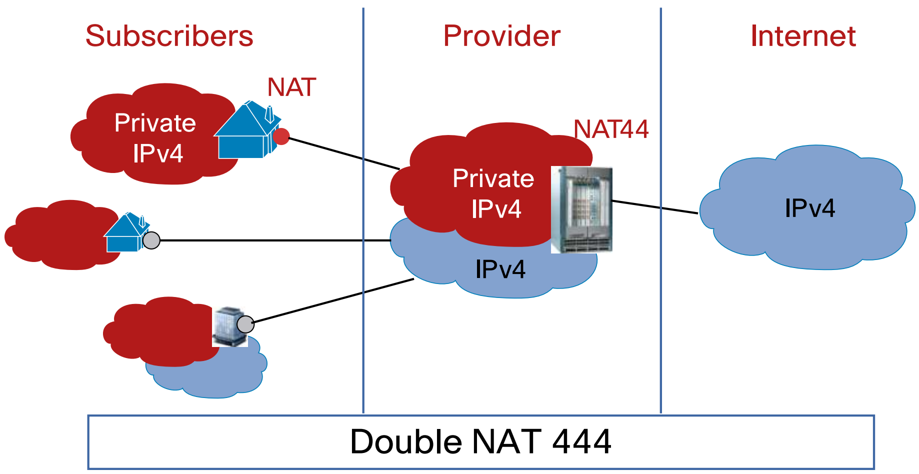

Address Family Translation (AFT)

Address Family Translation (AFT) refers to the translation of one IP address from one address family into another IP address of another address family; for instance, from one IPv4 address into an IPv6 address or vice versa. This translation is sometimes denoted as NAT46 (when the initiator is on the IPv4 side) or NAT64 (when the initiator is on the IPv6 side). AFT can be stateful or stateless. Stateless AFT is also known as IVI (in Roman numerals, IV = 4 and VI = 6); note, IVI can be IPv4 or IPv6 initiated.

Figure 7. Address Family Translation, AFT (NAT64 shown)

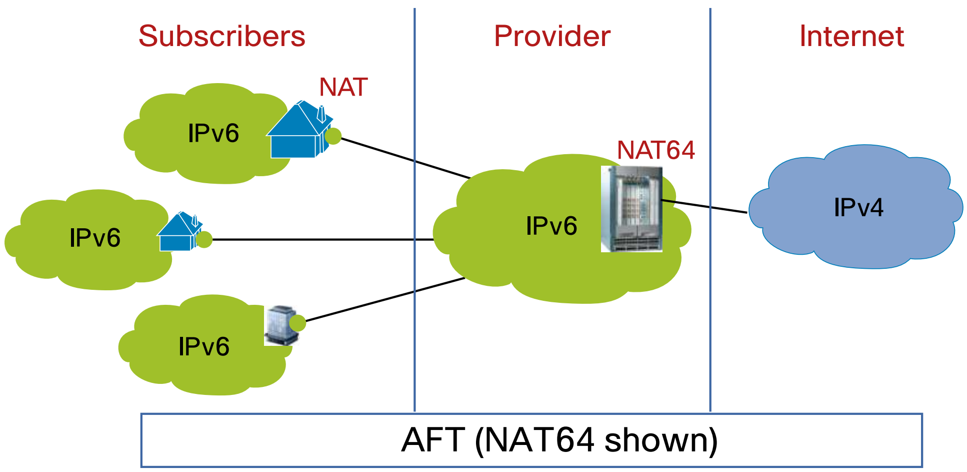

IPv6 Rapid Deployment (6rd)

IPv6 rapid deployment (6rd) enables a service provider to rapidly deploy IPv6 services to existing IPv4 sites to which it provides customer premise equipment (CPE). This approach utilizes stateless IPv6 in IPv4 encapsulation in order to transit IPv4-only network infrastructure. The encapsulation (aka softwires) must be supported by the CPE, while the CGv6 solution (6rd Border Relay) must support tunnel termination to route packets to Internet hosts on IPv6. The provider access network continues to be on IPv4, while customers see IPv6 and IPv4 service simultaneously. One of the leading deployments of this technology has been Free (Iliad of France).

Figure 8. IPv6 Rapid Deployment (6rd)

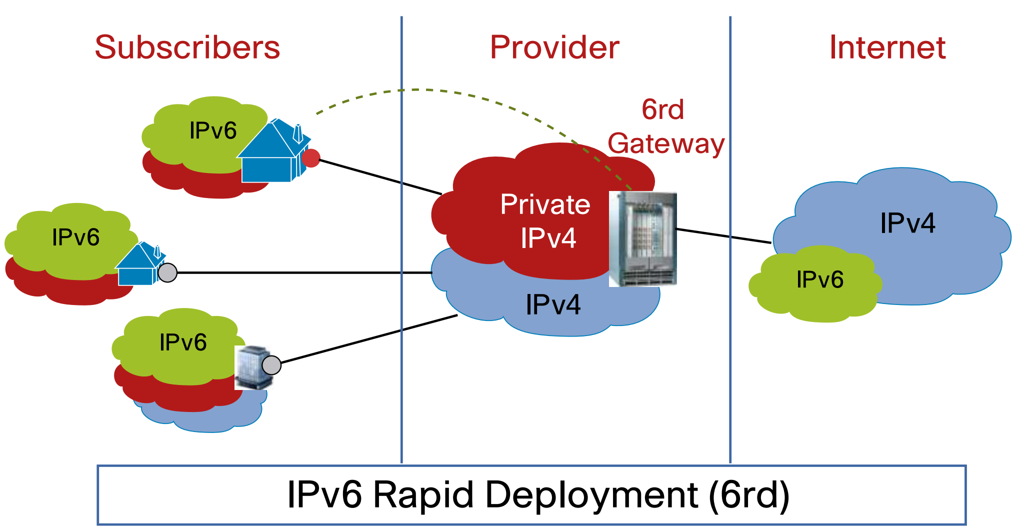

Dual-Stack Lite (DS-Lite)

In this scenario, at least part of the service provider network (e.g., access, aggregation network) only supports IPv6 routing. The CPE is provisioned only and natively with IPv6. Any IPv4 traffic on its local LAN is tunneled by the CPE over the IPv6 infrastructure to the CGv6 Gateway. The encapsulation (aka softwires) must be supported by the CPE. The IPv4 address space of the subscriber is a private one. The CGv6 Gateway terminates the tunnel and translates the IPv4 local addressing into globally routable IPv4 (NAT44). If the subscriber network has the capability of using IPv6, the IPv6 traffic is routed natively through the SP infrastructure. There is a single IPv4 NAT operation applied in the service provider network to the subscriber traffic.

Figure 9. Dual Stack Lite (DS-Lite)

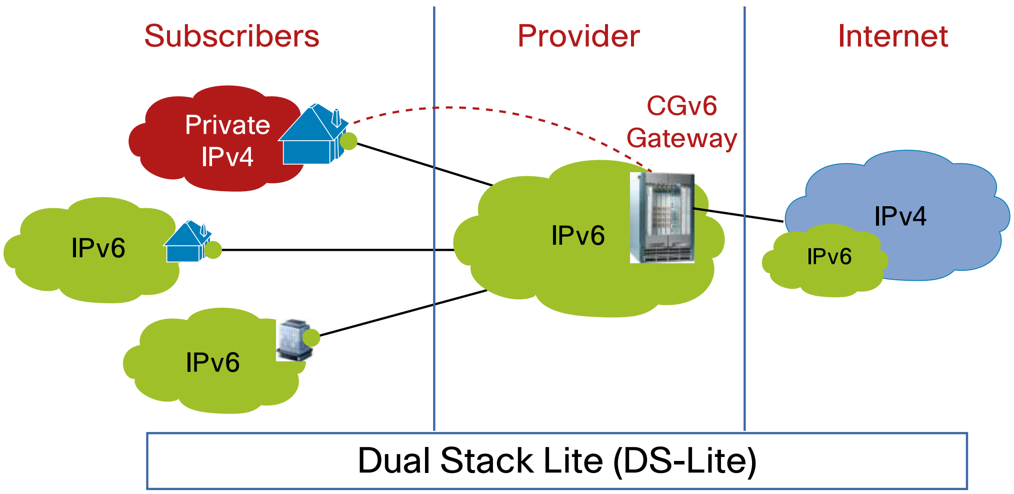

All IPv6

The end goal of an all IPv6 Internet infrastructure is probably some time away. In this scenario, IPv4 is sparingly used, if at all. All sections of the network support IPv6 natively. With this technology, all devices are globally reachable because they have unique addresses.This arrangement enables new revenue-yielding service opportunities.

Figure 10. All IPv6

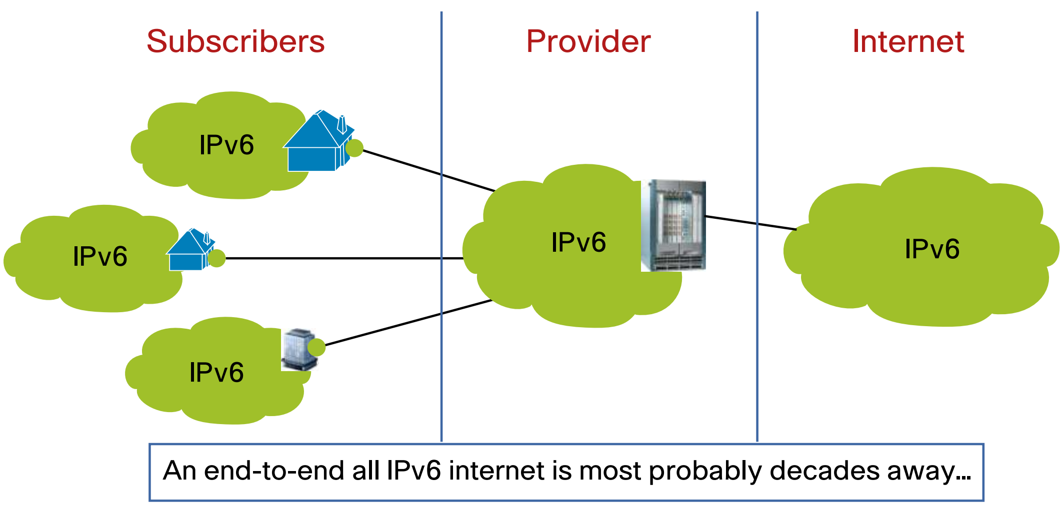

CGv6 Components

The Cisco Carrier-Grade IPv6 Solution (CGv6) comprises various product and service components. These components range from new hardware and new capabilities on existing platforms to existing capabilities on the service provider product portfolio. The exceptional scalability and feature richness of these components help ensure providers of a comprehensive tool set for the transition.

Carrier-Grade Services Engine (CGSE) for CRS-1

Figure 11. Cisco Carrier-Grade Services Engine (CGSE) for CRs-1

The Cisco Carrier-Grade Services Engine (CGSE) is an industry-leading solution for CGv6. Powered by a multicore CPU complex, it scales to tens of millions of address translations and gateway functionality with gigabits of throughput for hundreds of thousands of subscribers. In addition, rapid connection setup time boosts performance significantly. The CGSE is a single-slot module supported on all models of Cisco's proven high-end routing system: CRS-1. Several modules can be populated within a chassis for a high-performance solution that is deployable at places in the network where maximum CGv6 coverage can be obtained. The CGSE supports a highly available architecture, with line rate accounting and statistical logging for superior lawful intercept applications.

ASR Family

Figure 12. Cisco ASR family

The Cisco ASR Family Series comprises ASR9000 and ASR1000 platforms providing intelligent IP/MPLS edge services. The Cisco ASR 1000 Series is a highly scalable WAN and Internet edge router platform that delivers embedded hardware acceleration for various services such as address translation, VPN, firewall, network-based application recognition (NBAR), Flexible Packet Matching (FPM), NetFlow, quality of service (QoS), IP Multicast, access control lists (ACLs), Unicast Reverse Path Forwarding (uRPF), and Policy-Based Routing (PBR), without the need for separate service blades. Powered by the Cisco QuantumFlow Processor, with parallel processing, it scales to millions of address translations, with gigabits of throughput for CGv6 functionality. Being highly programmable, it delivers fast feature velocity in a compact form factor. In-service software upgrades make the platform highly reliable.

SP Product Portfolio

Figure 13. SP Portfolio with Dual Stack

Cisco's service provider product portfolio has been supporting IPv6 as a dual-stack technology for some time. Customers can enable IPv6, in addition to running an IPv4 protocol stack, without affecting performance. IPv4/IPv6 coexistence is part of the prepare approach where providers can use the same platform. Support for dual-stack technology includes IPv6 routing (Static, RIPng, IS-IS, OSPFv3, MP-BGP, EIGRP), mobility (Mobile IPv6), QoS (packet classification, queuing, traffic shaping, WRED, ACLs, NBAR), VPN (6PE, 6VPE), multicast (PIM-SM, PIM-SSM, PIM-Bidir), and management (SNMP, Netflow, Ping, Traceroute) among other protocols.

Cisco Services for IPv6

The Cisco Services for IPv6 offering enables service providers to transition to IPv6 in a controlled, safe, and cost-effective manner, thereby reducing the risk to their business. Cisco's proven track record in successful IPv6 implementations across the world led to development of global best practices. The service is modular, with key components being: readiness assessment, business planning assessment, implementation planning, and enablement.

Benefits of Cisco CGv6

Cisco's Carrier-Grade IPv6 Solution (CGv6) is specifically engineered to help service providers on their journey to an IPv6 infrastructure. The benefits of CGv6 extend beyond service providers to end customers as well.

Service Provider Benefits

The multitiered approach of preserve, prepare, and prosper is designed for an orderly and gradual transition, instead of implementing hurriedly assembled solutions when IPv4 run-out occurs. With the "preserve" approach, service providers can slow down the depletion of valuable IPv4 addresses and exploit their IPv4 infrastructure for some more time. Meanwhile subscriber and device growth can continue unconstrained. With the "prepare" approach, service providers can begin IPv4/IPv6 coexistence and ramp up IPv6 operations. And finally, with the "prosper" approach, new innovative service opportunities open up with the global reachability of all Internet devices.

End Customer Benefits

End customers are usually dependent on their service providers for addressing; thus Cisco CGv6 benefits them indirectly. With the extension of IPv4 availability within providers adopting the "preserve" approach, addresses for new services and devices can be easily obtained by subscribers. Inbound access to Connected Home devices is simplified. With the proliferation of IPv6 addressing, global accessibility and mobility can drive innovation of new services. Finally, the growth of networks within emerging markets and mobile vertical is no longer affected by the dwindling IPv4 address space. Other markets and verticals have been traditionally ahead in time for address allocation, but now with IPv6, the expanding growth in emerging and mobile can be equitably addressed.

Summary

The day for the last IPv4 address to be handed out will be upon us very soon. Procrastination is no longer an option. To cope with the run-out, the industry needs to move to IPv6. However, IPv4 and IPv6 protocols are not directly compatible, and hence various techniques are necessary for their coexistence. End customers are largely dependent on service providers for IP addressing, and thus operators need to promptly plan for the eventual exhaustion of addresses.

Cisco is leading the journey to IPv6 with the innovative Carrier-Grade IPv6 Solution (CGv6). CGv6 adopts a multitiered approach to preserve, prepare, and prosper during this migration. With preserve, service providers can prolong the life of their IPv4 addresses. With prepare, service providers introduce IPv6 in the IPv4 infrastructure for them to coexist while continuing to grow. With prosper, service providers eventually move to a mostly IPv6 network, opening up new revenue-yielding service opportunities. The technologies of translation, tunneling, and encapsulation underlie CGv6's multitiered approach, and these technologies are delivered via new and existing offerings from Cisco: Carrier-Grade Services Engine (CGSE) for CRS-1, ASR Series, and the entire service provider product set. This comprehensive solution set can help ensure that customers make an orderly and gradual transition toward IPv6.

The imperative to act is now.

For More Information

For more information about Cisco Carrier-Grade IPv6 Solution (CGv6), visit www.cisco.com/go/cgv6.

For more information about Cisco CRS-1, visit: http://www.cisco.com/go/crs

For more information about Cisco ASR Series, visit: http://www.cisco.com/go/asr

For more information about Cisco SP Solutions, visit: http://www.cisco.com/go/sp

For more information about IPv6, visit: http://www.cisco.com/go/ipv6

For more information about Cisco's IPv6 implementation on a satellite in space, visit: http://www.cisco.com/web/strategy/docs/gov/wood-iac-07-B-2-6-06-paper.pdf

Additionally, you may also contact your local Cisco account representative.

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 C11-558744-00 10/09