White Paper

Cisco Carrier-Grade IPv6 (CGv6) Solution

Delivering on the future of the Internet

April 2012

Contents

Executive Summary

The next version of the Internet Protocol, IP Version 6 or IPv6, was invented in the early 1990s to accomodate the rapidly expanding applications for Internet Protocol (IP) technologies and growth of the Global Internet at large. The design of IPv6 included a number of improvements and optimizations over its predecessor, IPv4, including stateless auto-configuration which eliminates the need for stateful DHCP servers, extended routing headers, and flow labeling. By far, however, the most significant improvement is the increase in the number of bits used for addressing, from 32 bits in IPv4 to 128 bits 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.

On Feb 3, 2011, the Internet Assigned Number Authority (IANA) declared that there were no more public IPv4 addresses left to distribute. The greater Internet community of service providers, vendors, and customers must address the IPv4 run-out problem now.

IPv6 is the only long-term solution the industry has available to continue growth in the manner that the world has come to expect. 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 deployed in networks for many years. Given that the clock is ticking and that large portions of the Internet are yet to be IPv6 ready, a significant challenge confronts the service provider community.

This white paper outlines the Cisco Carrier-Grade IPv6 Solutions or CGv6, for enabling IPv6 support within Service Provider networks. The paper describes a set of approaches, 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)

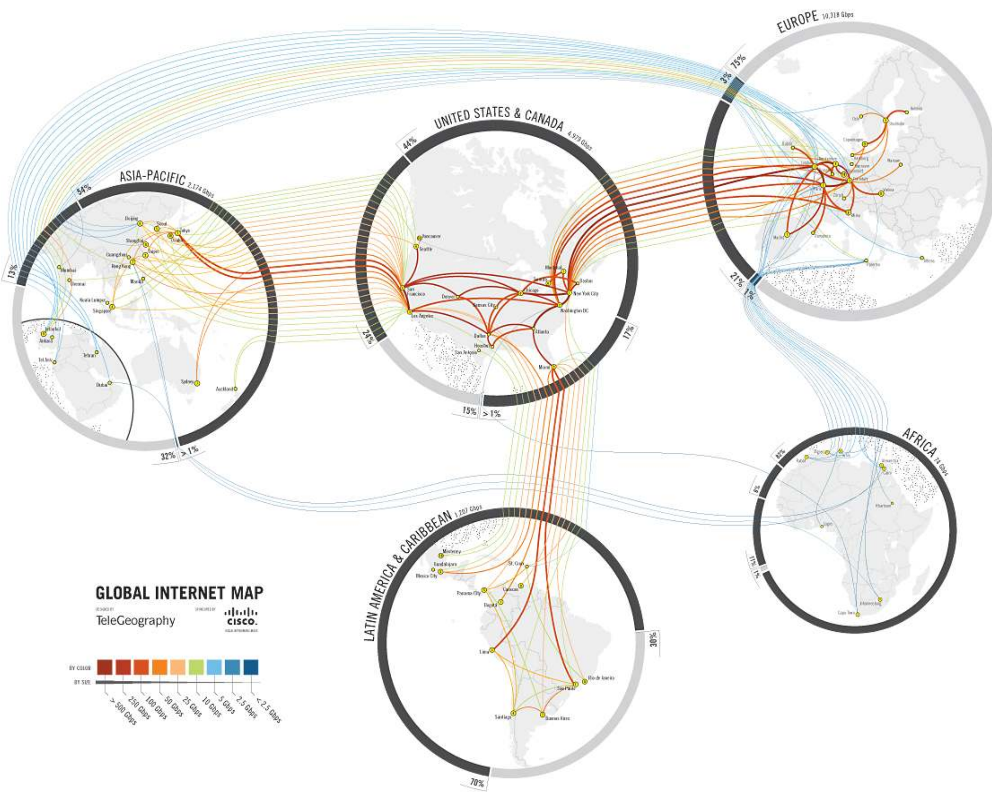

IPv4: A Fundamental Constraint

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

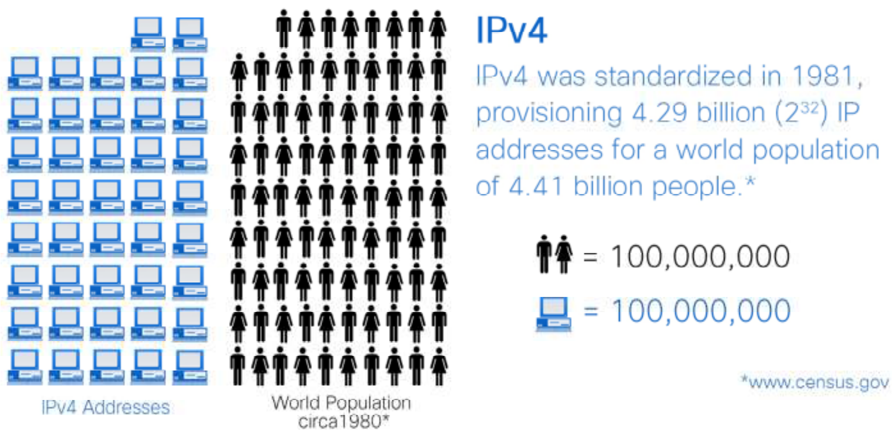

Figure 3. IPv6 Address Space

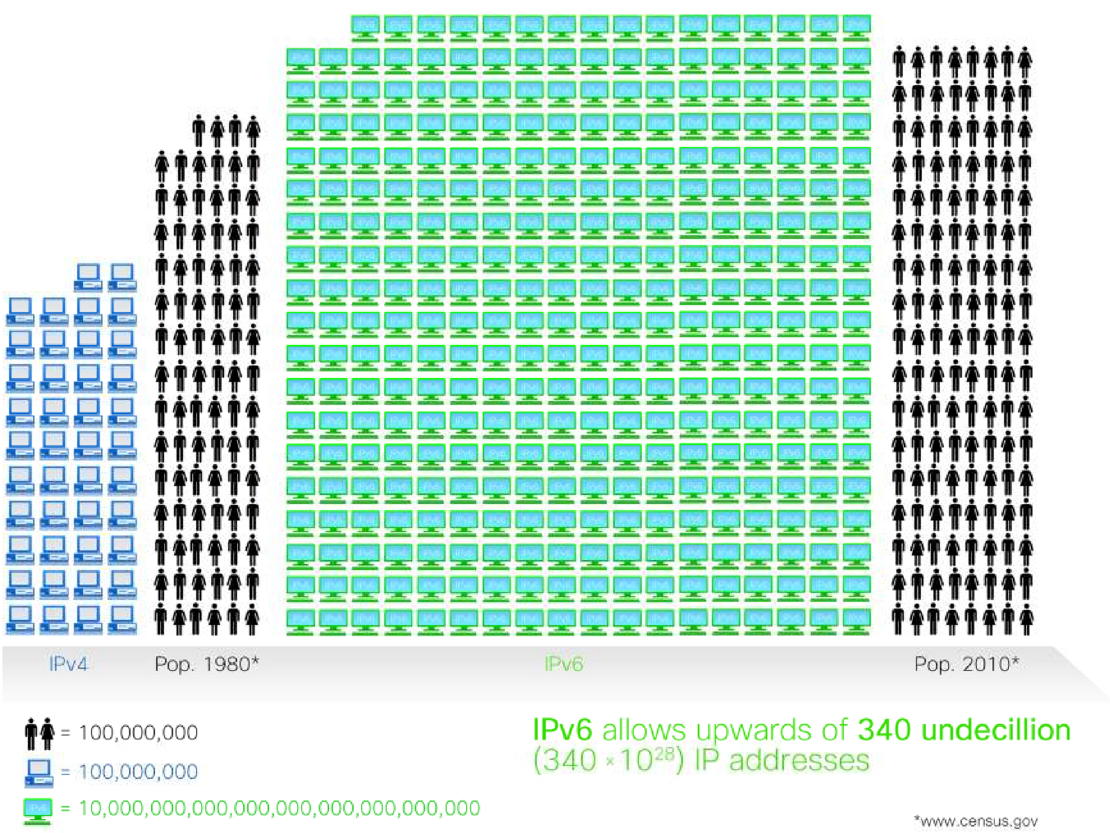

IP version 6 (IPv6) quadruples the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides more than enough globally unique IP addresses for every networked device on the planet. The use of globally unique IPv6 addresses simplifies the mechanisms used for reachability and end-to-end 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. The disjointed approach towards upgrading to IPv6 coupled with the former mentality of 'we will handle the issue when addresses finally do run out' has resulted in a present day Internet that is largely based on IPv4. 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 version, say IPv4, will not be able to directly communicate with an IPv6 network. The design of the IPv6 protocol stack, a design that was standardized in the mid-1990s, specifies a new packet format amongst other additional improvements that make native backwards compatability not possible. To resolve the IPv4 – IPv6 compatability issue, networks would need to run a pair of parallel infrastructures housing both protocols. The operational burden imposed by this scenario is one of the factors that slowed the adoption of IPv6. The wide deployment of IPv4 devices demands a bridging mechanism to evolve and support interoperability between IPv4 and IPv6 devices.

Little Economic Incentive

Prior to the exhaustion of IPv4 addresses, IP addresses were not considered a monetized resource. Customers, including service providers, were allocated addresses by Regional Internet Registries (RIRs) essentially for free provided they had the proper justification. IPv4 (and IPv6) address allocations follow this distribution model, which is predicated on the equitable and open nature of the Internet. With the exhaustion of IANA's pool of free IPv4 addresses (February 3, 2011) a market has been opened for a once free resource. Enterprises and Service Provider's alike are now in this secondary market for IPv4 Address space. While many have speculated on the value of an IP address, Microsoft set the price in March 2011, when it bought 666,624 IPv4 addresses from Nortel's liquidation sale for 7.5 million dollars; roughly $11.25 USD per IPv4 address.

Little Perceived Urgency

The Internet continues to grow at a rapid pace, with the amount of internet enabled devices expected to double by the year 2015. The increasing demand for internet access and the depletion of free IPv4 addresses has spurred mechanisms to reuse IPv4 addresses have grown. A popular mechanism, Network Address Translation (NAT), has gained popularity in home networks, requiring only a single public IPv4 address to be allocated to a home with multiple private IP 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

IPv4 and IPv6 are not interoperable and require the support of both protocol stacks for coexistence. The internet is built on a legacy of diverse networks and devices requiring unique approaches to transitioning to IPv6. Some devices only support the IPv4 protocol stack and others have only partial support for the IPv6 protocol stack. Most of the popular services available on the Internet are enabled on IPv4 making it more difficult for users to shift to IPv6 because continuing to use only IPv4 in the immediate timeframe appears operationally simpler.

A New Urgency

The free pool of IPv4 addresses became officially exhausted on February 3rd, 2011, when the IANA (Internet Assigned Numbers Authority) distributed the last large block of unallocated public IPv4 address space to the Regional Internet Registries (RIRs). With the imminent depletion of all IPv4 addresses, service providers are are facing an ever more urgent need to enable IPv6 within their network infrastructures.

Figure 4. Free IPv4 Address Blocks

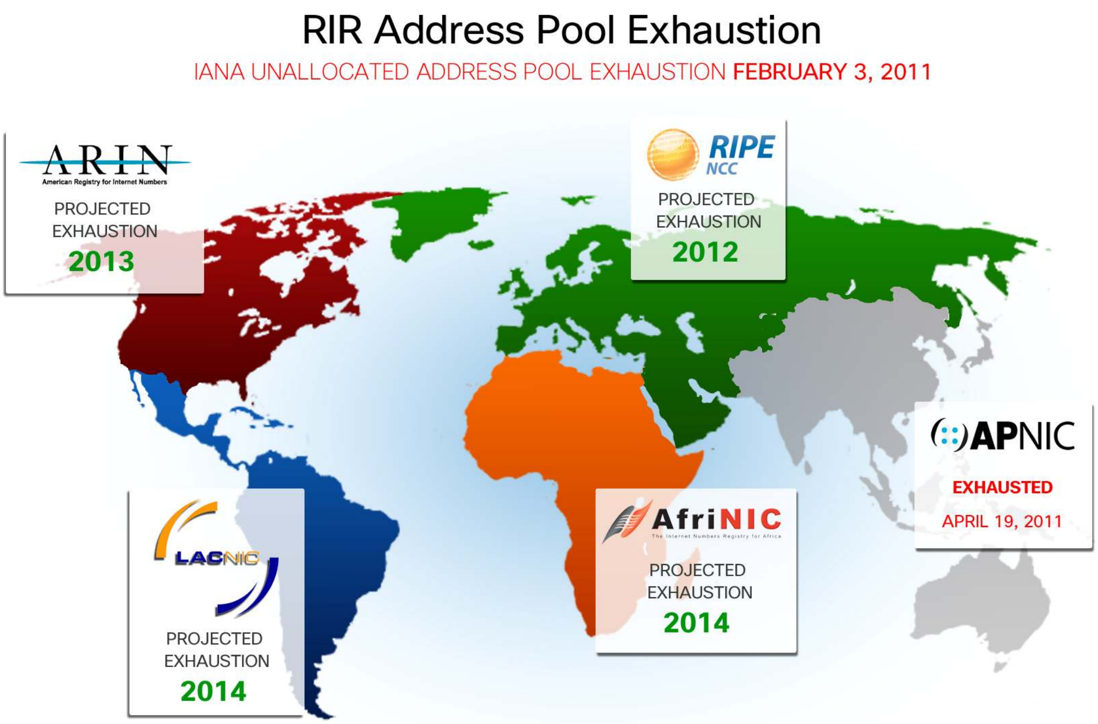

Source: Geoff Huston, http://www.potaroo.net/tools/ipv4/index.html

While the RIR's are systematically running out of IPv4 addresses, the number of IP devices is exponentially increasing, expected to more than double by the year 2015 (see figure 5). The number of IP enabled devices already outnumbers the entire IPv4 address pool and 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.

Figure 5. Number of Internet Enabled Devices

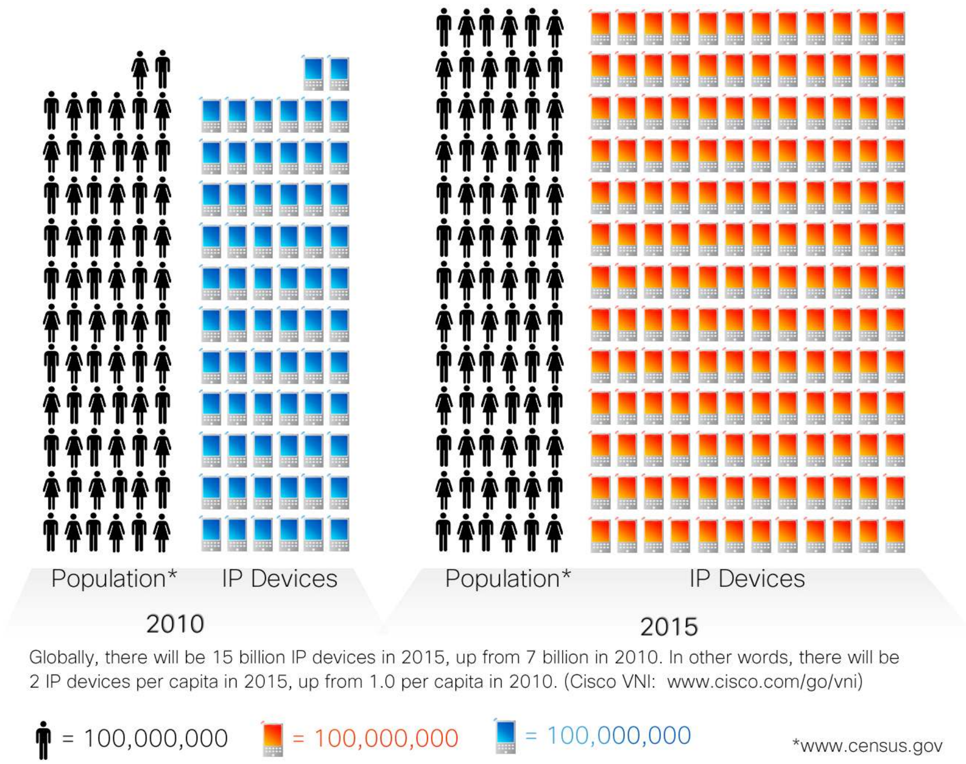

Source: Cisco VNI / www.census.gov

Journey to IPv6: Cisco Leads the Way

As a pioneer in IPv6 technology, Cisco has been a driving force in developing IPv6 standards, working with standards bodies like the Internet Engineering Task Force (IETF), to develop 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 forecasted 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 developing IPv6. While the foundation of IPv6 was set in the 1990s, provider networks only began adopting the technology in the early 2000s. The initial providers that began implementing IPv6 primarily relied on dual-stack or all IPv6 technologies, creating costly parallel infrastructures that operated independently of their existing IPv4 infratructures. For the time being, IPv6 was not considered an urgent matter and the high cost of implementation remained a formidable deterrent. To date, popular end-user applications continue to use IPv4 as providers have extended the life of their IPv4 address pools by requiring end customers to undertake address translation on their home equipment. By employing the use of Network Address Translation technology i.e. NAT44, providers have bought themselves a short term extension to IP address exhaustion. As the exhaustion of IPv4 sinks into reality, the need for a transitional solutions approach to enable IPv6 becomes apparent. For maximum effectiveness, investment in any approach must yield measureable benefits along the way and prepare the network operator to prosper during the transition to IPv6.

Cisco Carrier-Grade IPv6 Solution (CGv6)

Cisco has pioneered the framework for a Carrier-Grade IPv6 Solution (CGv6) with the practical approaches of preserve, prepare, and prosper so as to methodically help operators deploy IPv6 services. CGv6 enables operators to offer IPv6 services while considering one or more of the below approaches in any sequence:

1. Preserve – Leverage the existing IPv4 infrastructure, or
2. Prepare – Use dual-stack infrastructure, and
3. Prosper – Realize the benefits of IPv6 every step of the way towards an all IPv6 infrastructure.

CGv6 allows operators to adopt specific approach(es) based on their own challenges and timelines, since every operator may not need each approach. The challenges may be different depending on the operator and where their network is along the IPv6 deployment continuum. While some operators may be willing to offer IPv6 services while **Preserving** their IPv4 infrastructures, other providers may just more assertively **Prepare** with dual-stack technologies. One thing is common, however, any of these providers will **Prosper** with a well-calculated transition to IPv6 using Cisco CGv6 solutions.

An investment in any approach yields measurable benefits at the time of implementation and prepares the provider to prosper until IPv4 is eventually turned off.

Preserve

In this approach, the objective is to scale IP addressing and enable IPv6 offerings, but continue to use IPv4 network infrastructure. This approach is the logical path for those operators whose network infrastructures are not ready, willing, or able to upgrade aggressively to IPv6 in its entirety.

The approach is a low-risk path allowing the provider more time to plan and audit for IPv6 infrastructure. Solutions in this approach include technical solutions such as 6rd (IPv6 Rapid Deployment), 6PE (IPv6 Provider Edge), SLB (Server Load Balancing), LISP business solutions such as IPv4 Address Trading, and operational solutions such as IPv4 Renumbering. 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.

In the realm of technical solutions, 6rd and LISP techniques have been available for tunneling IPv6 over IPv4 network, whereas 6PE has been available for tunneling IPv6 over IPv4 Multiprotocol Label Switching [MPLS] network. IPv6 Rapid Deployment (6rd) standardized by Cisco engineers in RFC 5969 is currently responsible for the largest residential deployment of IPv6 in the world. LISP, also invented by Cisco engineers, is well known for the fastest IPv6 enablement in many networks including the one used by Facebook. SLB technique is available for allowing IPv4 servers usage while serving IPv6 users.

Note: All of these solutions enable IPv6 offerings. An investment to Preserve IPv4 infrastructure by use of just NAT44 translation i.e. Large Scale NAT (LSN /CGN44) in the core of the network does not provide a clear path to IPv6 and should only be implemented if absolutely neccessary. LSN/CGN44 may be the only short-term option for some providers as they run out of IPv4 addresses. *See Appendix A for more information on NAT44/4.

Prepare

In this approach, the objective is to upgrade the infrastructure to support IPv6 while coexisting with IPv4. This approach is the logical path for those operators whose infrastructures can be upgraded to IPv6 and those who want to deploy IPv6 in (existing) IPv4 networks.

Dual-stack running on hosts, switches, and routers is one viable prepare solution, paving the way for native IPv4 and IPv6 packet forwarding. This approach also includes the usage of Multiprotocol Label Switching [MPLS], as/where applicable for optimality (e.g. traffic engineering, fast rerouting [FRR]).

Interestingly, we are quickly getting to the paradigm, where legacy IPv4 packets are carried over an IPv6 network, either using tunneling techniques (e.g. Softwires Mesh, Dual-Stack Lite and MAP-E) or translation techniques (e.g. MAP-T) that are being defined/standardized by Cisco (in collaboration with others) in the IETF Softwires Working Group. Excitingly, Cisco has been leading the IETF standardization of next-generation of 'Stateless 46 solutions' (e.g. MAP-T, MAP-E) that enable sharing of an IPv4 address among multiple subscriber CPEs in a stateless manner, paving the way for operators to solve the IPv4 address exhaustion issue, minimize the complexity in their network and to accelerate the deployment of IPv6.

This approach in its entirety *is not* a prerequisite for operators to move to an all IPv6 network, but rather, to provide IPv4 operators different options to **Prepare** a migration path and deployment plan allowing for entry into IPv6. Ultimately, when all networks, devices and applications are all IPv6 capable the provider can fully prosper with the next-generation Internet.

Prosper

To Prosper is to realize the benefits of Cisco's CGv6 solutions as operators migrate to an all IPv6 infrastructure. It captures the various degrees of investment return as operators implement IPv6 transition approaches. With incremental investment in CAPEX and OPEX, Cisco's CGv6 solutions allow operators to incrementally grow and scale their IPv6 network at a rate inline with their subscriber demands.

In fact, this approach favors IPv6 to be used for delivering one or more servies and applications including the Internet service, as well as existing (voice over IP, email, video) and emerging applications (smart grid, mobility, sensor networks) that will employ IPv6. However, it is quite possible that these applications will not all appear and be ready at the same time; thus may justify the need for the IPv4/IPv6 coexistence and/or technoloigies that accommodate the residual IPv4 endpoints over an IPv6 infrastructure. One example is Stateful and stateless IPv4/IPv6 translator (aka Address Family Translation [AFT]), defined by Cisco engineers (in collaboration with others) in the IETF BEHAVE working group, enables IPv6 end-points to communicate with IPv4 end-points. Other examples are DS-Lite, MAP-E, MAP-T etc.

With a seemingly inexhaustible supply of public IPv6 addresses, every interface connecting every device on the Internet will be able to readily and easily obtain a globally routable IP address. The ambitious goal of achieving a prosperous all IPv6 network is to achieve true global reachability for any connected Internet device, opening up new untapped revenue-yielding service opportunities.

The Journey Begins

With Cisco's CGv6 approach of preserve, prepare, and prosper, service providers can begin their journey toward an all IPv6 infrastructure in a smooth and safe manner. Implementing solutions that preserve 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 approaches of solutions simultaneously or not at all. 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 and challenge, 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. This 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 approaches based on their business and technical priorities. CGv6 supports:

- Address Family Translation (AFT)
- IPv6 Rapid Deployment (6rd)
- Dual Stack Lite (DS-Lite)
- `MAP-E (Mapping of Address and Port - Encapsulation Mode)`
- `MAP-T (Mapping of Address and Port - Translation Mode)`
- All IPv6

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 6. Address Family Translation, AFT (NAT64 shown)

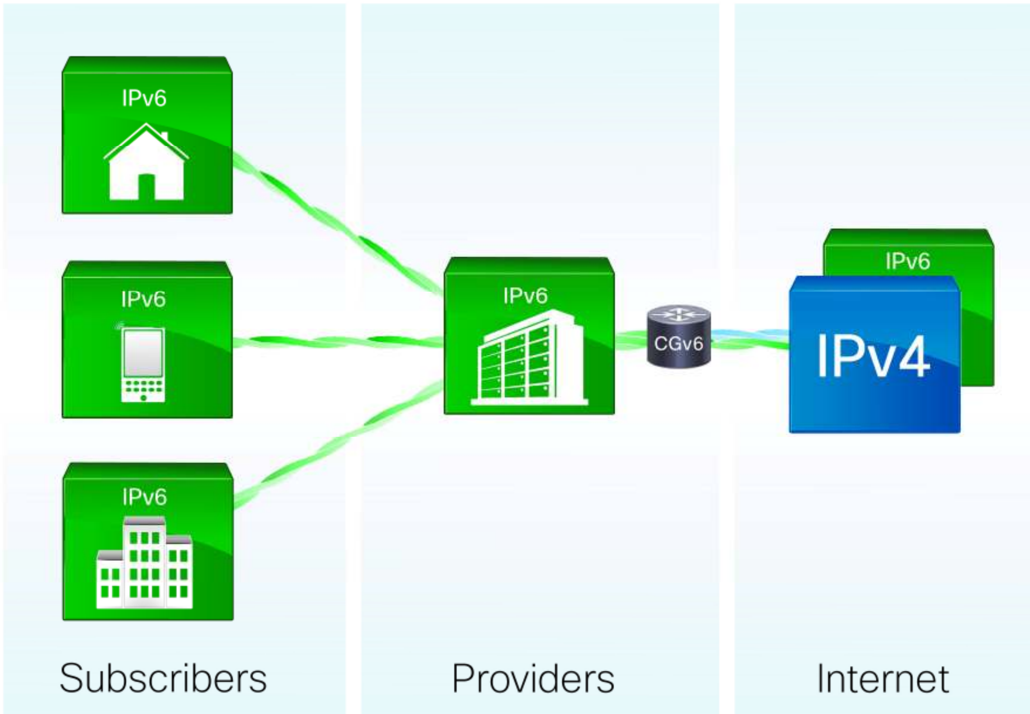

IPv6 Rapid Deployment (6rd)

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 7. IPv6 Rapid Deployment (6rd)

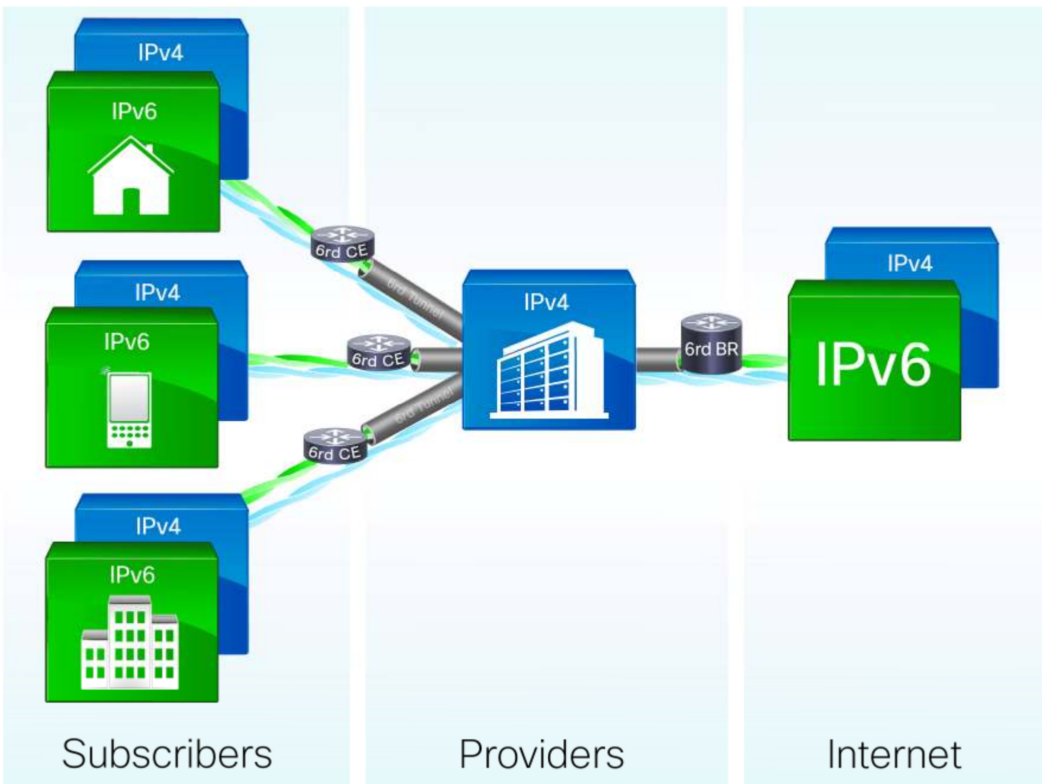

Dual-Stack Lite (DS-Lite)

DS-Lite is a combination of tunnel and translation (CGN) technologies. With DS-Lite, 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 8. Dual Stack Lite (DS-Lite)

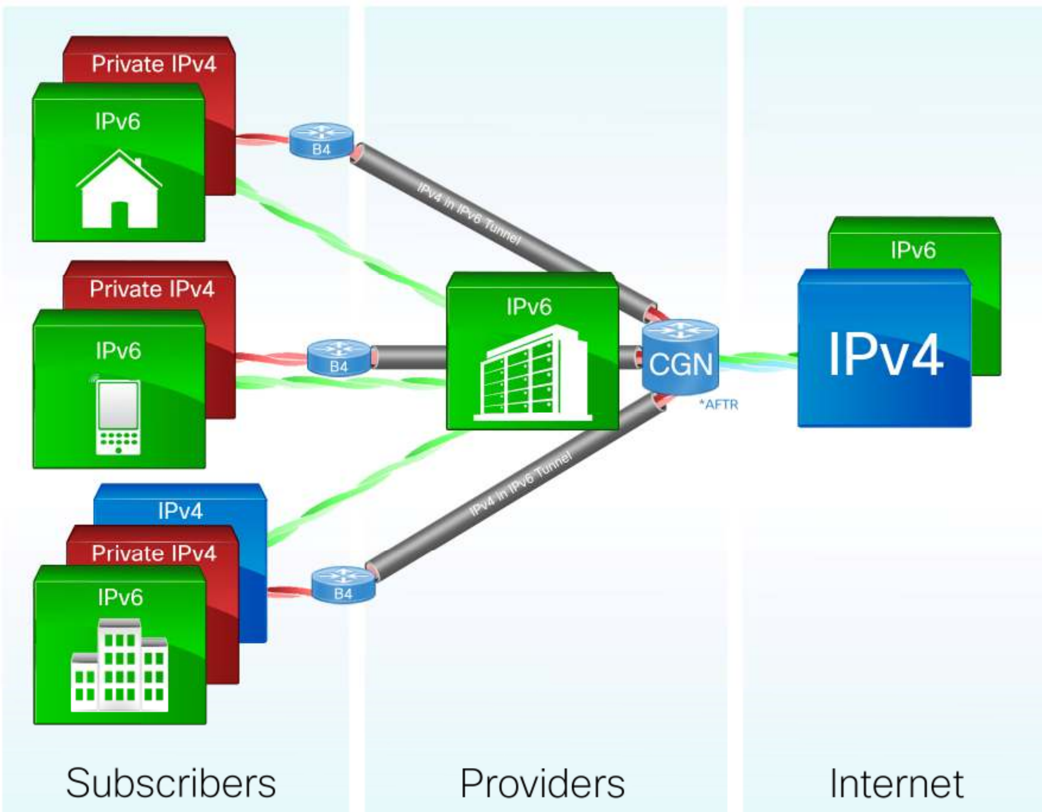

MAP-E (Mapping of Address and Port – Encapsulation Mode)

Mapping of Address and Port – Encapsulation Mode (MAP-E) is a combination of tunnel and translation technologies. Prior to being standardized, MAP-E was commonly referred to as IPv4 Residual Deployment (4rd). MAP-E enables a service provider to rapidly allow IPv4 services to IPv6 (customer) sites to which it provides customer premise equipment (CPE). This approach utilizes stateless IPv4-in-IPv6 encapsulation (i.e. tunneling) to transit IPv6-enabled network infrastructure. The encapsulation must be supported by the CPE and MAP-E Gateway/Border Relay, which removes the IPv6 encapsulation from IPv4 packets while forwarding them to the Internet. The provider access network can now be on IPv6, while customers see IPv6 and IPv4 service simultaneously. MAP-E also helps manage IPv4 address exhaustion by keeping the stateful NAT44 on CPE, as usual.

MAP-E is attractive to those SPs who have deployed or are planning to deploy IPv6 end-to-end. It is a superior alternative to DS-Lite, given the stateless tunneling and no CGN/44 in the service provider network.

Figure 9. MAP-E (Mapping of Address and Port – Encapsulation Mode)

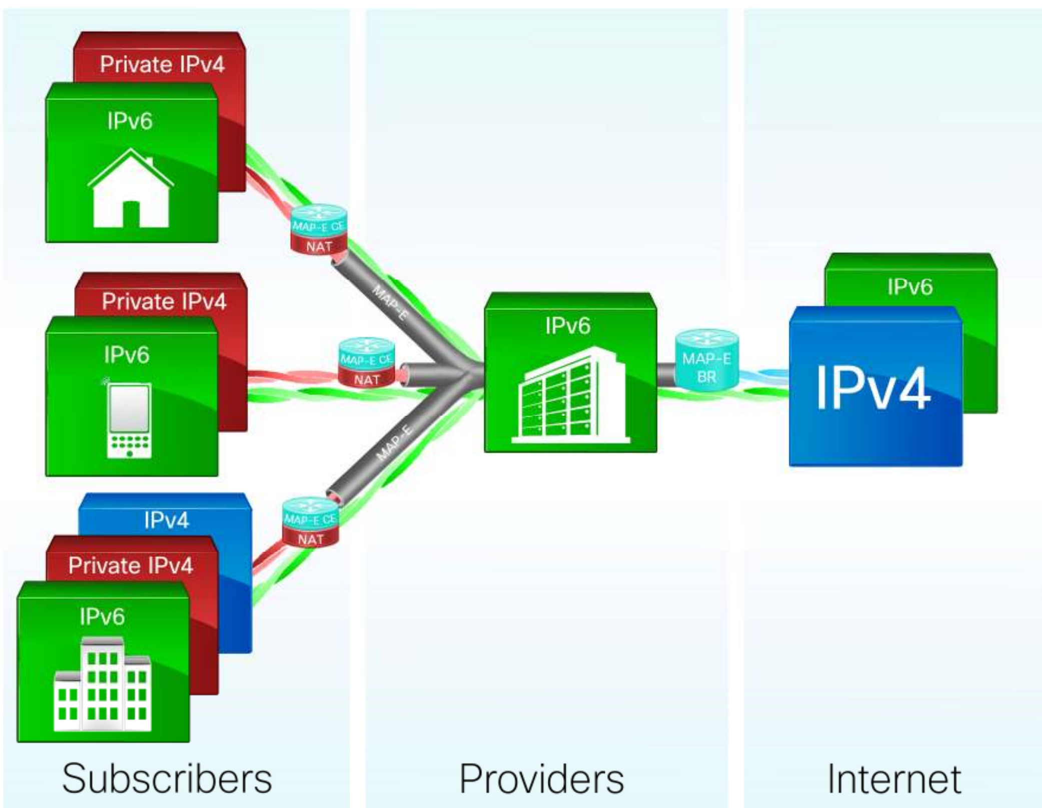

MAP-T (Mapping of Address and Port – Translation Mode)

Mapping of Address and Port – Translation Mode employs a combination of translation technologies. Prior to being standardized, MAP-T was often referred to as Dual IVI (dIVI) or 4464. enables a service provider to offer IPv4 services to IPv6 enabled (customer) sites to which it provides customer premise equipment (CPE). This approach utilizes stateless IPv4 to IPv6 translation (i.e. NAT64) to transit IPv6-enabled network infrastructure. The provider access network can now be on IPv6, while customers see IPv6 and IPv4 services simultaneously. MAP-T keeps the stateful NAT44 on CPE, as usual, to handle IPv4 address exhaustion, in addition to stateless NAT64 on CPE and Border Router.

MAP-T is attractive to those SPs who have deployed or are planning to deploy IPv6 end-to-end and want to manage IPv4 address exhaustion with utmost predictability. MAP-T is a superior alternative to DS-Lite, given no tunneling and no CGN/44 in the service provider network.

Figure 10. MAP-T (Mapping of Address and Port – Translation Mode)

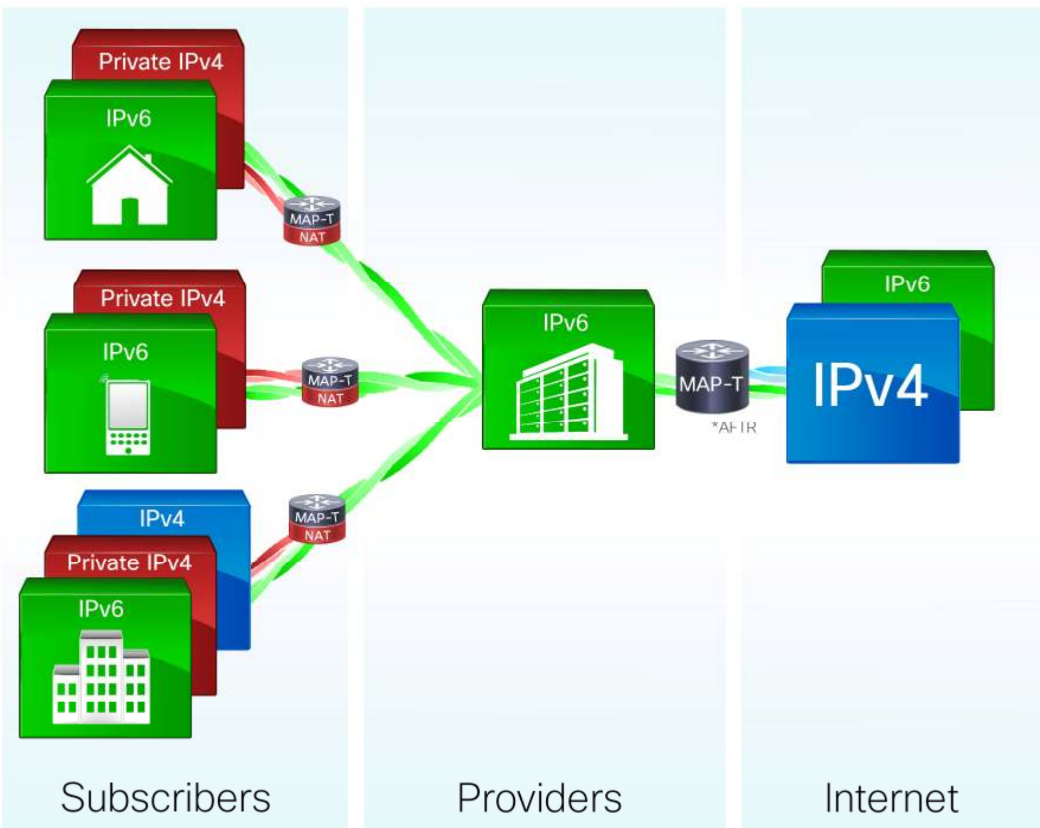

All IPv6

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 11. All IPv6

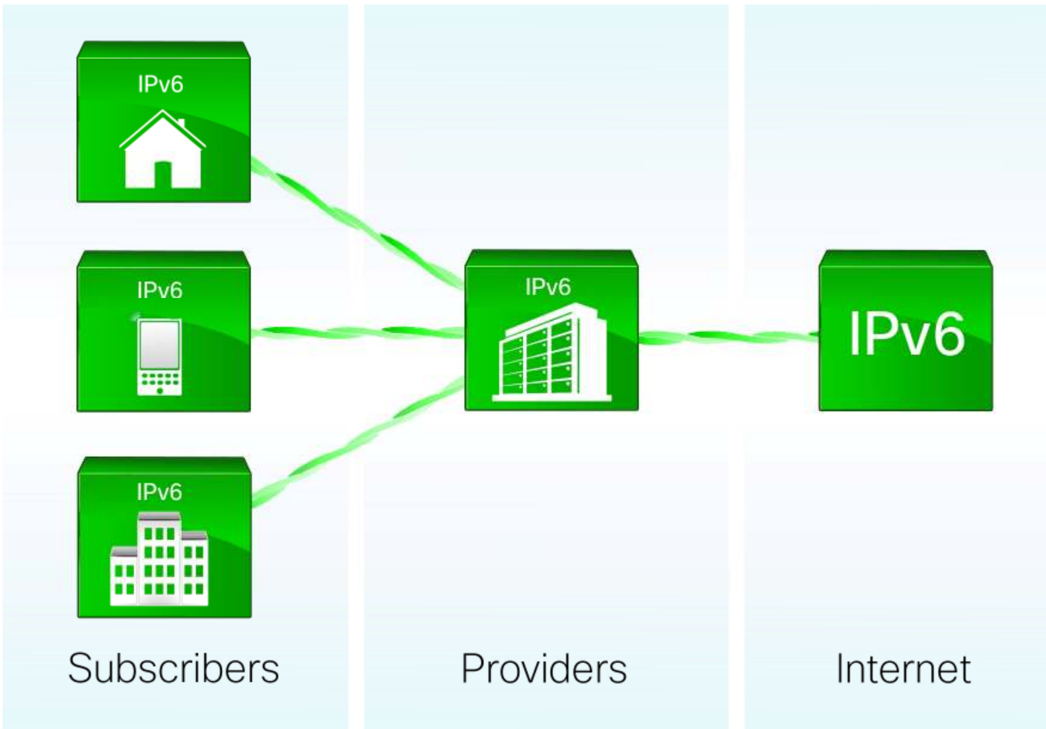

CGv6 Components

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.

Cisco Carrier Routing System

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, the CRS. 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. The CRS also allows for centralized placement of the 6rd border relay function and takes advantage of its stateless nature to remove session limit bottlenecks.

Figure 12. Cisco Carrier Routing System

Carrier-Grade Services Engine (CGSE) for CRS

Figure 13. Cisco Carrier-Grade Services Engine (CGSE) for CRS

Cisco Aggregation Services Routers

The Cisco ASR 9000 and ASR 1000 platforms support CGv6 and provide dual stack 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.

Figure 14. Cisco ASR Family

The Cisco ASR 9000 is ready for the transition to IPv6. The ASR 9000 represents efforts Cisco is making to simplify design, deployment and management of services for global service provider’s IPv6 next-generation networks. The ASR 9000 Series Integrated Service Module (ISM) provides a single touch point for carrier-grade IPv6 deployment across thousands of devices. The ASR 9000 is also outfitted with a multi-functional service module used for carrier grade NAT functions to provide a full selection of available transition technologies at the edge of the network. With release 4.2.1 the ASR 9000 also supports DS-Lite.

Service Provider Product Portfolio

Figure 15. Service Provider Portfolio with Dual Stack Capabilities

Core Routing

Cisco CRS

Edge Routing

Cisco ASR 9000

Access Routing

Cisco ISR Family

Cisco ASR 1000

Security/Data Center

Cisco ASA/PIX/Nexus/MDS Platforms

Catalyst 6500

7600

Cisco Prime Network Registrar

The potential for 340 Undecillion IP devices connected to the network presents a challenge for network management. Service providers in particular face challenges with IP address management to accelerate provisioning, simplify service activation, transition from IPv4 to IPv6, and reduce OPEX for administration and management. Current operational processes are largely manual and are unable to scale with the increase in demand. Operators require integrated and comprehensive solutions to simplify management of IP addresses and the transition to IPv6.

Cisco has developed Cisco Prime, a network management solution that simplifies the design, provisioning and management of carrier-grade networks. Prime is a comprehensive solution that centralizes and automates service design, fulfillment, assurance, and performance analysis – to help service providers and enterprises lower their costs while meeting high customer expectations. Cisco Prime Network Registrar provides integrated, scalable, reliable Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and IP Address Management (IPAM) services for both IPv4 and IPv6 transition planning. Cisco Prime is USGv6 (US Government IPv6) certified.

Cisco's service provider product portfolio has been supporting IPv6 as a dual-stack technology for several years. Customers can enable IPv6, in addition to running an IPv4, without negatively affecting performance. IPv4/IPv6 coexistence is part of the "preserve and prepare" approach where providers can use their existing equipment to support both protocols. 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

Cisco Advanced Services has developed a service solution to help operators meet the challenges of IPv4 exhaustion and smoothly integrate IPv6 into their environments. Cisco Advanced Services offer a complete portfolio designed to meet operator needs across all phases of the network lifecycle, from preparing and planning to operations and optimization. Service providers worldwide can attest to the expertise that Cisco provides, developing solutions that address the needs of the customer, assured by Cisco's commitment to excellence.

Benefits of Cisco CGv6

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

Service Provider Benefits

The approach of preserve, prepare, and prosper is designed for operators to execute an orderly and gradual transition, instead of implementing hurriedly assembled solutions when IPv4 run-out occurs. CGv6 allows service providers "preserve" their IPv6 network infrastructure by maintaining network services throughout the transition to IPv6. With the "prepare" approach, service providers can begin IPv4/IPv6 coexistence and ramp up IPv6 operations as they enter into the early stages of a full IPv6 network. And finally, the Cisco CGv6 solutions allow operators to "prosper", through a seemless transition from IPv4 into a world of 360 undecillion IP addresses with limitless appliications.

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 is 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 an approach to preserve and prepare, while allowing operators to prosper during this migration. This framework ensures an orderly and efficient, incremental transition without being rigid or linear. It's widely applicable to Mobile, Cable, Wireline and Over-The-Top providers. CGv6 allows service proivders to adopt specific approaches based on their own timelines. With preserve, service providers can prolong the life of their IPv4 network. With prepare, service providers introduce IPv6 in the IPv4 infrastructure for them to coexist while continuing to grow as they enter into the early stages of a full IPv6 network. Every step of the way, CGv6 allows the operator to prosper. New revenue-yielding service opportunities are created while controlling IPv4 to IPv6 transition CAPEX and OPEX costs. Each incremental investment is based on the service providers specific timeline and challenges and posititions the provider to fully prosper when an all IPv6 environment is realized in the future. The technologies of translation, tunneling, and encapsulation underlie CGv6's approach, and these technologies are delivered via new and existing offerings from Cisco: Carrier-Grade Services Engine (CGSE) for CRS, 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.

Appendix A

Double NAT 444 (Carrier Grade NAT)

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 16. Double NAT 444 (Carrier Grade NAT)

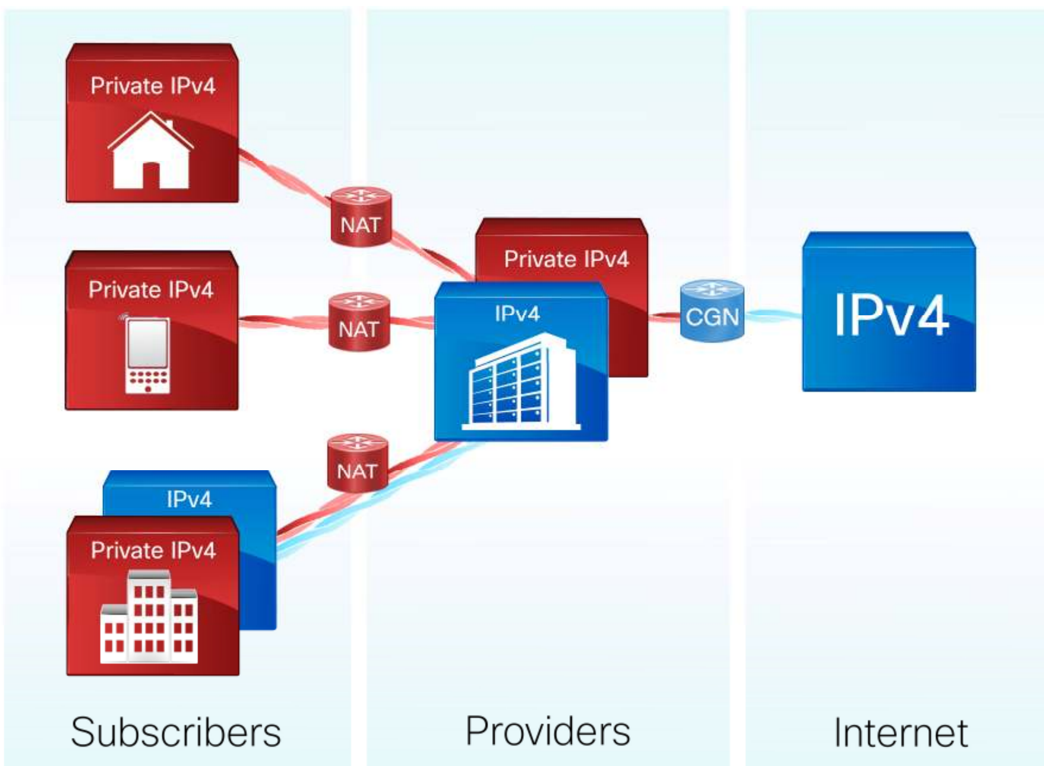

For More Information

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

C11-558744-02 04/12