

WHITE PAPER

The Business Case for Delivering IPv6 Service Now

Sponsored by: Cisco Systems

Nav Chander

May 2012

EXECUTIVE SUMMARY

This IDC white paper examines the economics of an Internet Protocol version 4 (IPv4) transition to an Internet Protocol version 6 (IPv6) network for wireline operators providing fixed broadband services. Operators are faced with important decisions about which transition technologies to use, when to implement them, and where to implement them in the network. This paper reviews the costs associated with maintaining and growing current service offerings while dealing with the depletion of public IPv4 address space. This analysis is based upon an interview with a major North American tier 1 fixed wireline operator and includes capex cost considerations in the access, aggregation, and core networks and a discussion of important opex cost considerations.

The key findings include:

- ☒ Fixed wireline operators can expect up to 69% in savings over five years by using IPv6 for new subscribers instead of relying on private IPv4 space and NAT444 exclusively.
- ☒ NAT444 is not a long-term solution to IPv4 address exhaustion. NAT444 should be used only if absolutely necessary due to the long-term cost prospects of maintaining a NAT444 network.
- ☒ By 2016, 95% of all content on the Internet is projected to be reachable by IPv6 subscribers. Building a path to IPv6 via transition technologies such as 6rd will reduce an operator's reliance on NAT444.
- ☒ Stateless transition technologies (e.g., 6rd) provide a cost-effective solution since state awareness and tracking are not required. A centralized implementation in the aggregation/core provides for fewer touch points and further optimizes cost.
- ☒ Opex costs are most directly driven by the requirement for NAT444 to track and maintain session state. Transition technology mix will dictate an operator's opex costs on a case-by-case basis.

IN THIS WHITE PAPER

This IDC white paper examines the costs associated with implementing NAT444 alone for IPv4 address exhaustion versus deploying NAT444 for IPv4 address exhaustion plus 6rd to build a path to IPv6 in a fixed wireline operator's network. Future white papers will cover next steps in the transition to an all-IPv6 network with NAT64 to support legacy IPv4 applications.

SITUATION OVERVIEW


This section outlines the important market trends and issues related to IPv6 adoption and IPv4 address exhaustion that operators have to consider.

Today's broadband, enterprise, and mobile packet networks primarily use IPv4 addressing to provide a unique numerical identifier for each connected endpoint. This address allows the endpoints to originate traffic and receive traffic from the IP network they are connected to. The IPv4 address space, which has a finite limit of approximately 4 billion addresses, has reached its capacity, with IPv4 address exhaustion occurring in February 2011.

Fixed wireline operators providing Internet access are among the first to have to address this depletion as the growth rate of new subscribers and devices is outpacing the available public IPv4 address reserves. Figure 1 depicts a typical scenario in a fixed wireline operator's network, which consists of the home end devices, the CPE, the access network, the aggregation, the core, and, finally, the content.

FIGURE 1

Fixed Wireline Operator Network



Source: Cisco, 2012

All of these pieces of the end-to-end picture have a different timeline on their path to IPv6 enablement. The operator typically has responsibility for the CPE to the core. Most operators have already enabled the aggregation and core for IPv6 transit, with the CPE and access network being the most difficult and costly to upgrade. Operators in most cases have shifted focus to making sure all new equipment is IPv6 capable while the existing install base is left intact.

Eventually, the majority of the Internet and its applications will be on IPv6; however, as operators transition their networks to support it, they will need to maintain connectivity to the IPv4 Internet as well. Because IPv4 and IPv6 are not interoperable, transitional technologies that can address the needs of both protocols are crucial to ensuring the continued reliability of existing data networks, as well as the viability of expanding those networks to account for the massive levels of growth in data access by consumers.

Fixed wireline operators must make important decisions about their network infrastructure to maintain their current service offerings and continue growth while ensuring a profit curve. The long-term objective for most operators is to completely transition from IPv4 to IPv6 to alleviate the IPv4 address exhaustion situation and reduce opex by having to maintain only one protocol stack. Implementing a flash cut to a pure IPv6 environment, though, is not realistic because it requires all network equipment to be IPv6 capable as well as the IP-based content and all applications to be IPv6 compliant from day 1. This is why a well-planned transition strategy is necessary.

IPv6 Adoption Trends

Clear trends show how the growth of IPv6 adoption impacts core, aggregation networks, access networks, subscriber devices (tablets/PCs, mobile devices), content providers, and enterprises. Consider the following:

- ☒ The number of Autonomous Systems advertising global IPv6 reachability in the Global BGP Internet routing table more than doubled from October 2010 to April 2012.¹
- ☒ IPv6 traffic is expected to grow exponentially as the number of IPv6-capable devices such as smartphones, tablets, and routers increases to over 6 billion by 2015 (according to IDC forecasts).
- ☒ The global increase in IPv6 available content is expected to represent 95% of all content by 2016.
- ☒ By 2016, 39% of all mobile devices will be IPv6 capable.²
- ☒ Verizon Wireless 4G LTE has IPv6 built in and all LTE devices must be IPv6 compatible.³
- ☒ IPv6 content providers, such as Google, YouTube, CNN, Microsoft, Yahoo!, and Facebook, are now supporting IPv6 content, which is accessible from Windows 7 and Mac OS 10.7.
- ☒ Many new applications and devices, such as Internet-enabled wireless devices, home and industrial M2M appliances, Internet-connected transportation, integrated telephony services, sensor networks such as RFID, smart grids, cloud computing, and gaming, will be designed for and enabled by IPv6 networks.
- ☒ Applications requiring simple end-to-end reachability by the network, including new services and associated revenue opportunities for service providers, should flourish as IPv6 is deployed and could likewise be inhibited by the introduction of NAT444.

¹ Source: http://v6asns.ripe.net/v/6?s=_ALL;s=_RIR_APNIC;s=_RIR_AfriNIC;s=_RIR_ARIN;s=_RIR_LACNIC;s=_RIR_RIPE_NCC

² Source: *Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011–2016*

³ Source: <http://policyblog.verizon.com/BlogPost/780/VerizonandIPv6.aspx>

IPv6 Transition Considerations and Alternatives

Several service provider segments are facing the IPv4 address depletion issue and the prospect of implementing IPv6 in their network. Each segment has different characteristics that will dictate the approach that can be taken to accomplish this transition. Some of the key factors that contribute to the approach taken are IPv4 address exhaustion status, availability of IPv6 content and applications, and end-device IPv6 readiness. Table 1 provides the status of these factors for each of the possible segments: mobile, fixed enterprise, and fixed wireline broadband. In this paper, we focus on these factors as they relate to the fixed wireline broadband segment.

TABLE 1

Key Considerations for Service Providers by Segment

Key Factor	Mobile	Fixed Enterprise	Fixed Wireline Broadband
IPv4 address exhaustion?	Now — Devices are already deployed with IPv6 addresses	Varies — NAT already used in peering points	Now — Combination of NAT and IPv6-enabled CPE
IPv6 content and application availability?	Rapid growth	Slower growth; custom enterprise applications and development requirement	Rapid growth
What is the device/CPE refresh frequency?	Short refresh cycle	Longer refresh cycle	Longer refresh cycle

Source: IDC, 2012

As the supply of IPv4 addresses is rapidly being depleted, network service providers are employing unique solutions to optimize current networks to preserve the remaining unallocated IPv4 addresses and implement a transition to IPv6.

Today, operators can consider multiple approaches to deal with IPv4 address exhaustion, including IPv4 Carrier-Grade NAT (NAT444), IPv6 Rapid Deployment (6rd), Native Dual Stack, Dual Stack Lite (ds-lite), and NAT64. Of these approaches, IDC compared two of the most widely deployed technologies in the world as a starting point to the transition for fixed wireline broadband networks: NAT444 and 6rd. While the costs associated with 6rd with respect to implementing IPv6 within the access network are different, many of the costs in terms of IPv6 service itself should be the same. However, native IPv6 is not specifically addressed in the paper. As additional operators are interviewed, future white papers will cover the next steps in their transition journey to a native IPv6 network.


FIXED BROADBAND SERVICES: BUSINESS CASE COMPARING CGNAT AND IPV6 6RD

What are wireline service providers implementing today? IDC interviewed a major North American tier 1 fixed wireline service provider to validate the business model for an IPv4 to IPv6 transition strategy using a combination of 6rd and NAT444 and compared it with a NAT444 standalone approach. The 6rd approach is preferred by some fixed wireline operators primarily because it allows a smaller capex investment and is based on incremental growth of IPv6.

Carrier-Grade NAT444 (NAT444), sometimes referred to as large-scale NAT (LSN), is an approach to IPv4 network design in which end sites, in particular residential networks, are configured with private network addresses that are translated to public IPv4 addresses by routers embedded in the network operator's network, permitting the sharing of small pools of public addresses among many end sites. This shifts the NAT function and configuration from the customer premise to the Internet service provider network, as shown in Figure 2.

FIGURE 2

Network with Carrier-Grade NAT444




Source: Cisco, 2012

IPv6 Rapid Deployment (6rd) [RFC5969] is a technology used by service providers to give IPv6 access to their customers, even if the access network is not IPv6 capable. A 6rd approach, as shown in Figure 3, can be securely and rapidly deployed without requiring the service provider to upgrade the access network infrastructure. 6rd is not a mechanism to address IPv4 address exhaustion; rather, it is a mechanism to tunnel native IPv6 traffic over an IPv4 infrastructure. For situations where the client, application, or content is not IPv6 enabled, the requirement for NAT444 persists.

FIGURE 3

Wireline Network with IPv6 6rd



Source: Cisco, 2012

Key Business Cost Components for IPv6 Transition

Several important capex and opex business cost considerations exist for the two alternatives. The anticipated increase in opex costs involves the additional state per session, which is part of the NAT444 alternative, but not for the native IPv6 path, and impacts testing and support costs, as shown in Table 2 and Table 3.









TABLE 2

Key Cost Components Covered in the Analysis

Cost	Description	Type	Alternative #1 (NAT444 Only)	Alternative #2 (NAT444 + 6RD)
Core/aggregation/access	Cost of placing NAT444 or 6rd functionality in the network.	Capex	↑	↓
IPv6 6rd-enabled CPE	Cost of placing IPv6 6rd-enabled CPE in the home is cost neutral for new subscribers.	Capex	NA	↔
State per session	As devices are added, the traffic on the network increases, requiring more sessions and state driving additional NAT444 hardware.	Capex	↑	↓

Source: IDC, 2012

TABLE 3**Key Cost Components Not Covered in the Analysis That Should Be Considered**

Cost	Description	Type	Alternative #1 (NAT444 Only)	Alternative #2 (NAT444 + 6RD)
State per session	As devices are added, the traffic on the network increases, requiring more sessions and state driving additional NAT444 hardware.	Opex		
Security	Testing, training, implementation, and troubleshooting.	Opex		
CALEA/lawful intercept	When NAT444 is used, the private IP address must be correlated with the public IPv4 address plus port. This must also be correlated with RADIUS accounting, which can be very laborious.	Opex		
Applications	Testing, performance issues, increased NOC level 1 and 3 support calls, subscriber churn, operator-enforced sessions per application, and the cost of not deploying net-new IPv6-only applications.	Opex		

Source: IDC, 2012

Business Case: Residential Fixed Broadband Service

IDC examined a fixed residential broadband service business scenario by comparing the capex cost of deploying both technologies together — NAT444 to preserve IPv4 addresses and 6rd to build a path to IPv6 Internet — with the cost of deploying only NAT444.

Alternative #1: NAT444 Only

The NAT444-only option is one in which all new users are assigned private IPv4 addresses. There are two key metrics when scoping the NAT444 implementation: sessions per device and throughput per device:

- ☒ **Sessions per device.** The key benefit of using NAT444 is the ability to translate multiple private IPv4 addresses into one public IPv4 address; this is done by dividing the port space for the public IPv4 address. The number of ports that will be used in turn for sessions can be provisioned or dynamically allocated. The number of sessions consumed drives the number of NAT444 engines that are required.
- ☒ **Throughput per device.** The other key metric used when sizing a NAT444 implementation is how much throughput per device is required. It's possible that a NAT444 engine will be able to satisfy the sessions-per-device requirement for a set number of devices, but the throughput cannot be met, thus generating the need for more NAT444 engines and driving more capex. It is imperative to consider both key metrics.

Key Assumptions

- ☒ IPv4 public addresses are exhausted
- ☒ All traffic is translated through NAT44 at the core node
- ☒ Four devices per subscriber
- ☒ Annual subscriber growth rate of 10%
- ☒ Number of NAT44 translations required = Subscribers x devices per subscriber x sessions per device

Alternative #2: NAT444 + 6rd

The NAT444 + 6rd option means that if the device supports IPv6 and if the requested content is available over IPv6, the traffic is routed via IPv6 and then tunneled through the IPv4-only access network via 6rd at the CPE. If the device does not support IPv6 or the content is not available over IPv6, the traffic is routed through the private IPv4 network. The CGNAT throughput and sessions now only need to account for the traffic that is not IPv6 supported. The business analysis assumes traffic going to IPv6-enabled Web sites is increasing. Over time, if 6rd is deployed, there is less dependence on NAT444 to preserve IPv4 address space, and eventually more and more services can be deployed over the IPv6 network.

Key Assumptions

- ☒ IPv4 public addresses are exhausted
- ☒ IPv4 traffic is translated through NAT44 at the core node
- ☒ Maximize IPv6 traffic to the extent of IPv6 content availability
- ☒ Minimize private IPv4 traffic
- ☒ Four devices per subscriber

- ☒ IPv6 and 6rd-capable CPE for new subscribers
- ☒ Annual subscriber growth rate of 10%
- ☒ NAT44/6rd throughput = Subscribers x devices per subscriber x bandwidth per device

Capex Comparison: 6rd Provides an Easier Transition to Full IPv6

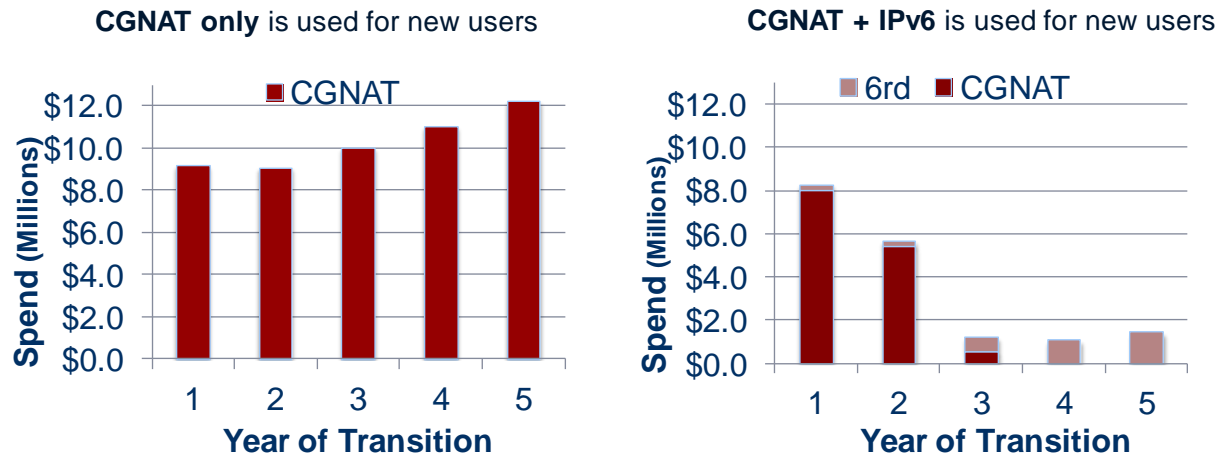
IDC confirms that with 6rd, there is a one-time initial capital investment for aggregation/core router network infrastructure — no session cost, no increase in cost due to churn because of application degradation. As more IPv6 content is available and traffic increases, operational costs using 6rd will stay constant after the initial investment.

Key findings for this five-year capex scenario comparison demonstrate a capex savings of 30–69% depending on the operator's ability to negotiate the cost of 6rd functionality into its IPv6 CPE. Also, the number of sessions per IPv4 will increase over time, adding more hardware cost to the NAT444 solution and in turn more capex.

On the left side of Figure 4, the graph indicates the costs associated with implementing NAT444 over five years. The costs each year are equal to the hardware and software resources needed to support the NAT444 functionality of new customers.

FIGURE 4

Capex Comparison of IPv6 Strategies



69%¹ capex savings by turning on 6rd + CGNAT (6rd solution eases CGNAT requirements and paves path to Dual Stack)

1 SP with 5M residential subs and a 10% yoy growth; no additional cost is incurred for turning on 6rd in CPE
2 Each device uses an average of 500 sessions due to high session applications

Source: IDC, 2012

On the right side, the graph indicates the costs associated with the NAT444 + 6rd strategy. The capex analysis assumes that 25% of the content will be available via IPv6 in year 1 and will increase to 95% by year 5, which is consistent with what most tier 1 fixed wireline operators expect. Because of this, the network is able to relieve the private IPv4 network of any traffic that is being originated by an IPv6 device and is pointing to content that is available over IPv6.

For the NAT444-only strategy, the NAT444 requires increasing amounts of hardware and software resources and results in substantial costs, including additional opex costs described in Table 3 and not shown in this analysis. The NAT444 + 6rd strategy, in comparison, leverages the growth in the IPv6 traffic during this transition process to control the CGNAT costs while letting customers use their IPv6 devices.

For the NAT444 + 6rd strategy, once a significant amount of content is available over IPv6, the NAT444 resource requirement decreases significantly. New IPv6 CPE will be equipped with 6rd with no additional spend required.

CHALLENGES/OPPORTUNITIES

IDC's view is that fixed wireline operators face challenges in planning an economical and pragmatic transition of their networks, subscribers, and content to IPv6. This requires careful planning and support from key vendors such as Cisco. Cisco's support for delivering IPv6 service alongside preserving IPv4 service in the face of IPv4 exhaustion mitigates the risk to operators and also provides both short-term and long-term opex and capex savings.

CISCO'S POSITION

Cisco supports IPv6 in the core, aggregation, and subscriber equipment as well as in the network management software it supplies to operators, which ensures an end-to-end, reliable, and fully supportable IPv4 to IPv6 transition. Cisco has a full suite of options for placement of the appropriate transition technology in the network, with the CRS in the core and the ASR 9000 and 1000 series routers in the aggregation. This allows for strategic placement of the Carrier-Grade IPv6 solutions in the network.

The Cisco Carrier-Grade Services Engine (CGSE) is an industry-leading solution for IPv6 support in the core and aggregation network. 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 CRS-1/3. 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.

The ASR 9000 System is a high-end aggregation device that is equipped with the Cisco IOS XR operating system. The ASR 9000 leads the industry, boasting 96Tbps of total capacity alongside Carrier-Grade high availability designed for continuous operations. The ASR 9000 is also outfitted with a multifunctional service module used

for CGNAT functions to provide a full selection of available transition technologies at the edge of the network.

The Cisco ASR 1000 Series is a highly scalable WAN and Internet edge router platform. 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.

The Cisco Linksys Router Family will support IPv6 by default on a range of new Linksys routers and has IPv6 firmware updates available for the following Linksys E-series Wireless Routers:

- Linksys E1200 Wireless-N300 Router
- Linksys E1500 Wireless-N300 Router
- Linksys E1550 Wireless-N300 Router with Speed Boost
- Linksys E2500 Advanced Dual-Band N600 Router
- Linksys E3200 High Performance Dual-Band N600 Router
- Linksys E4200 Maximum Performance Dual-Band N750 Router
- Linksys E4200v2 Maximum Performance Dual-Band N N900 Router

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 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 life cycle, from preparing and planning to operations and optimization.

As new IPv6 applications emerge and existing services such as Netflix and other video-based services migrate to IPv6, it is important and timely for fixed broadband communication providers to become more aggressive with their network transformation and transition to IPv6.

As the World IPv6 Launch approaches, IDC believes that Cisco is well positioned to help operators manage the IPv4 to IPv6 transition costs and to advise operators to plan and implement new innovative service opportunities as they open up with the global reachability of all Internet devices.

CONCLUSION

IDC believes that service providers have to take a longer-term economic view with regard to preserving IPv4 addressing services while introducing new IPv6 addresses. Though the use of a technology such as NAT444 is important in the next 12–24 months, it is equally important to select a vendor that can allow for gradual IPv6 introduction in parallel to maintaining its existing IPv4-based services and minimize capex. Technologies to support this IPv6 introduction include options such as 6rd that can also provide capex advantages with early implementation alongside NAT444 and enable plans for future value-added services based on IPv6.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2012 IDC. Reproduction without written permission is completely forbidden.