APPLICATION DELIVERY

Ensuring a Smooth Transition to Internet Protocol Version 6 (IPv6)

As IPv4 addresses dwindle, companies face the reality of a dual-protocol world

BROCADE

The transition from the IPv4 protocol to the IPv6 protocol represents a significant challenge to service providers, hosting services, e-commerce companies, and government agencies. They must find ways to accommodate end users and Customer-Premises Equipment (CPE) devices that may be using either protocol.

This white paper examines four transition technologies along with their advantages and disadvantages. Brocade is offering its customers a transition path to the new wave of Internet technologies while continuing to utilize customers' existing infrastructure.

INTRODUCTION

When the Internet Assigned Numbers Authority (IANA) standardized IPv4 in 1981, no one could have foreseen that its seemingly plentiful pool of 4 billion addresses would become depleted. But according to Internet World Stats, Internet usage grew by 444.8 percent between 2000 and 2010.[1] With the increasing deployment of Internet-enabled mobile devices, smart-grid devices, and cloud-based applications, the spike in usage is only going to increase in this decade.

Indeed, on February 3, 2011, the IANA allocated its last block of IPv4 addresses to five Regional Internet Registries (RIRs) around the world. These registries assign IP addresses to Internet Service Providers (ISPs), which in turn issue IP addresses for home and office machines, smart phones, and other Internet-enabled devices.

Taking the place of the old protocol, IPv6 is designed not only to solve IPv4's address scaling challenge but also to rectify other shortcomings. For instance, it provides a variety of mechanisms to increase both routing efficiency and management simplicity (see box, "What's New in IPv6").

Because of the preponderance of new capabilities, however, the designers of IPv6 were not able to make it backward-compatible with IPv4. This means devices speaking different versions of Internet Protocol can no longer communicate with each other natively, and applications that rely on such communication will fail.

[1] http://www.internetworldstats.com/stats.htm

FINDING THE RIGHT TRANSITION TECHNIQUE

The industry faces a challenging transition while it moves carefully from its current IPv4-capable routers, switches, servers, and applications to IPv6-ready devices. Service providers—whether they are providing content, hosting services, or Internet access—cannot add or accommodate new customers unless their content is equally accessible to both IPv4 and IPv6 users.

Similarly, e-commerce sites need to accommodate customers without knowing which protocol those customers' client devices use. At the same time, multiple federal governments around the world have enacted regulations forcing their agencies to adopt IPv6. As a result, service providers, hosting services, and other content providers need to investigate transition approaches during the evolution period. Today there are several technologies that create a bridge between IPv4 and IPv6; they use such techniques as translation, coexistence, tunneling, overlay, and more. The primary technologies include the following:

- NAT444
- NAT64
- Dual-Stack Lite
- 6rd

Each technology has its advantages and disadvantages, which we'll consider.

Translation: NAT444

In a traditional address-translation mechanism using Network Address Translation (NAT), an IP address from a private address pool is translated to a globally unique, publicly reachable IP address. Optionally, the mechanism translates source port information so that many private IP addresses can share a limited number of global IP addresses.

NAT444 is essentially an extension to the traditional NAT mechanism. It involves two layers of address and port translation. The first takes place at the Customer-Premises Equipment (CPE) and the second at the ISP, which uses a capability known as Large-Scale NAT (LSN). The term NAT444 signifies translation from one IPv4 block to a second IPv4 block, followed by a third IPv4 block. The first IPv4 block is a private address at the CPE. The second one is another private IP address block between the CPE and the ISP, and the third is a globally reachable public address block (see Figure 1).

Figure 1.
How NAT444 works.

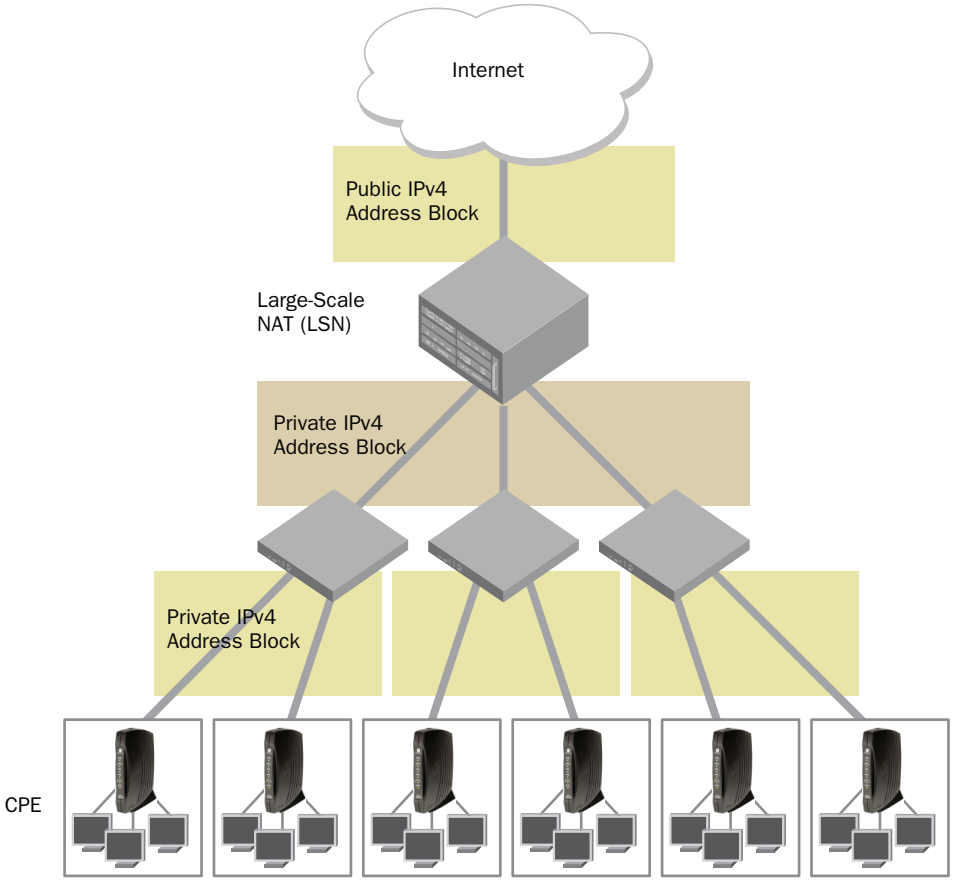

Advantages of NAT444

- It does not require the replacement of existing CPE devices.
- It utilizes proven NAT technology.
- It does not require changing other network elements such as Domain Name System (DNS).

Disadvantages of NAT444

- There could be a potential overlap between the customer's private address block and the private address block used between the CPE and the service provider. This could result in the misrouting of packets.
- The routing of packets between two different customers behind the same LSN is also challenging and may require a change in firewall policies.
- The classic disadvantages of NAT technology, such as the masking of end-user IP addresses and breaking end-to-end transparency, apply to NAT444 as well. In fact, NAT444 complicates communication even more for applications that embed IP addresses inside the payload of the packet, such as media applications, because it involves invoking the NAT mechanism twice for every transmission.
- As an IPv4 address-preservation technique, it does not offer any longer-term IPv6 transition benefits.
- Doing NAT at two locations negatively affects performance.

Coexistence: NAT64

The Network (and Port) Address Translation between IPv6 to IPv4 model—generally known as NAT64—is a mechanism for both the transition to IPv6 and the coexistence of IPv4 with IPv6. It works together with DNS64, essentially a DNS translation service, to enable client-server communication between an IPv6-only client and an IPv4-only server and vice versa. It allows for peer-to-peer communication where communication can originate from an end-node running either of the two protocols. NAT64 utilizes a preassigned IPv6 prefix to algorithmically translate IPv4 addresses of IPv4 servers. Similarly, it translates the IPv6 addresses of IPv6 servers to and from IPv4 addresses by installing mappings (see Figure 2). Overall, the NAT64 model offers a non-intrusive and seamless transition path for organizations looking to explore IPv6.

Advantages of NAT64

- It is completely transparent to end-users because address translation occurs at the service provider network edge and it involves no change in client-end CPE devices. Thus it can be an extremely cost-effective and practical solution.
- It allows for transition to IPv6 while preserving existing IPv4-based infrastructure. It facilitates coexistence of IPv4-only and IPv6-only devices while ensuring seamless communication between the two during the transition period.

Disadvantages of NAT64

- It involves translating addresses between IPv4 and IPv6, resulting in potential loss of originating client IP addresses unless they are captured through some other means.
- Address translation requires processing power and can negatively affect overall system performance if it's not designed correctly.

Figure 2.
How NAT64 works.

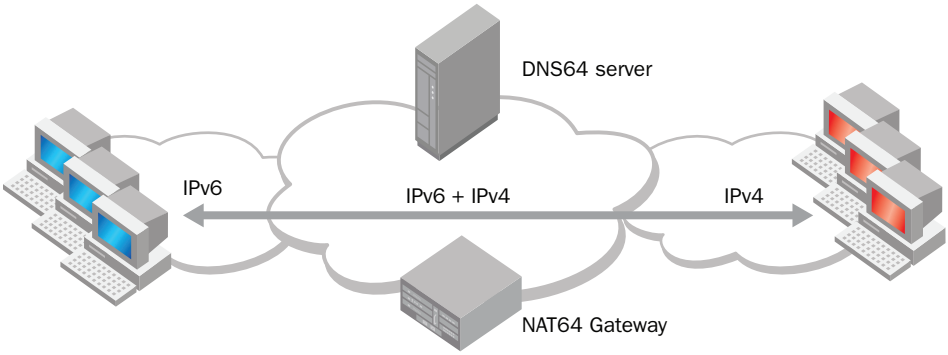

Tunneling: Dual-Stack Lite

Dual-Stack Lite (DS Lite) is a tunneling technique that carries IPv4 traffic from the CPE over an IPv6 link between the CPE and its service provider. When traffic originates from an IPv4 device at a customer location, the DS Lite mechanism encapsulates it inside an IPv6 packet for transport to the provider network. At the provider's end, the mechanism reverses the process on the packet and uses LSN to reach the appropriate IPv4 destination on the Internet (see Figure 3). Because the DS Lite model relies on IPv6 technology in the provider core, it can also route IPv6 packets between IPv6-only end points.

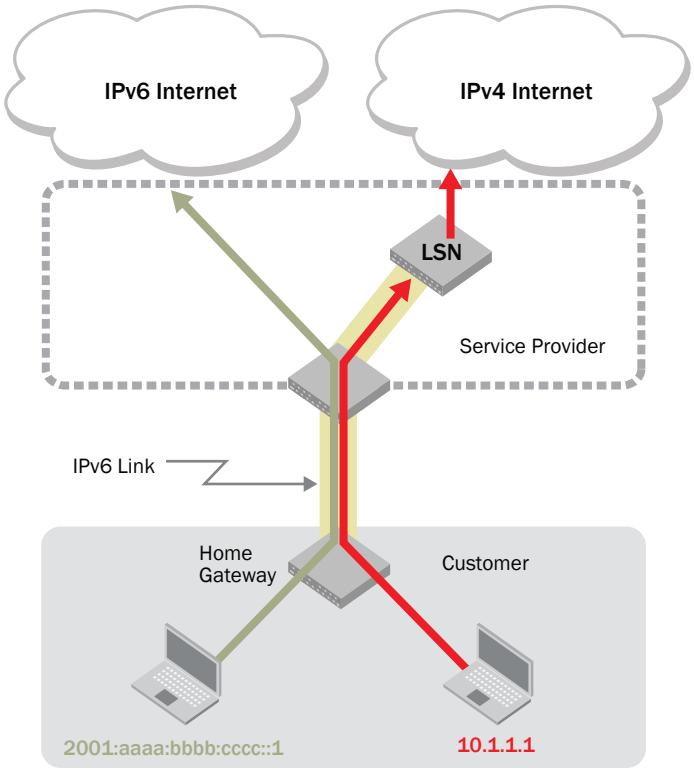

Figure 3.
How Dual-Stack Lite works.

Advantages of DS Lite

- It mitigates the routing challenge between two customers behind the same LSN device.
- It enables simultaneous connectivity for IPv6 clients communicating with IPv6 resources on the Internet and IPv4 clients communicating with IPv4 resources on the Internet.

Disadvantages of DS Lite

- It requires expensive upgrading of CPE devices, a cost service providers may not want to cover.
- It does not provide interconnectivity between IPv4 devices and IPv6 devices, which will eventually be required for IPv4-to-IPv6 transition.

Overlay: 6rd

The basic principle of the 6rd model is to deploy IPv6 over the existing IPv4 infrastructure of service providers (the rd in 6rd stands for "rapid deployment"). The mechanism relies upon algorithmic mapping between IPv6 and IPv4 addresses assigned for use within the service provider network. A 6rd mechanism requires deploying 6rd-aware CPE and one or more 6rd-aware border relay routers. The CPE device encapsulates IPv6 packets, which are then carried over the service provider's IPv4 network to border relay routers. The 6rd border relay router then decapsulates the packet and forwards it natively to an IPv6 network (see Figure 4). The model enables service providers to offer IPv6 services alongside IPv4 services, while making minimal upgrades to their existing IPv4 infrastructure. The model can be decommissioned upon completion of a service provider's IPv4 network migration to a dual-stack model.

Figure 4.
How 6rd works.

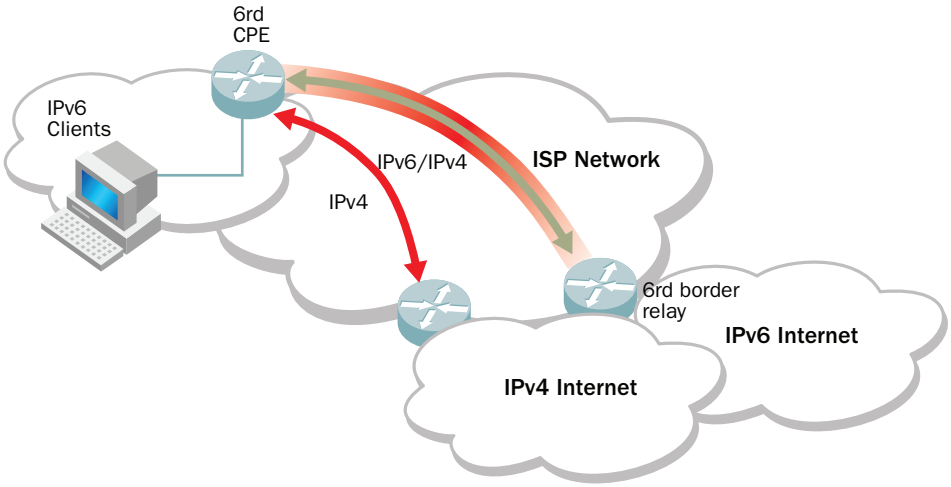

Advantage of 6rd

- It enables service providers to quickly offer IPv6 services without making forklift upgrades to existing IPv4 infrastructures.

Disadvantages of 6rd

- It requires upgrading CPE devices. Service providers can easily accommodate new customers with new equipment, but it may not be economical to upgrade existing customers.
- It does not offer interconnectivity between IPv4 and IPv6 protocols.

WHICH PATH TO FOLLOW?

One question continues to dawn upon service providers and enterprises as they face the transition from IPv4 to IPv6: How soon does the transition decision need to be made?

IPv4 pundits argue that there are still millions of available IPv4 addresses, and technologies such as NAT444 and LSN will help service providers accommodate new users for several years. Alternatively, IPv6 loyalists—citing the increasing deluge of new mobile devices, IPv6-only smart-grid capabilities, and cloud applications—recommend a more urgent move to IPv6.

Regardless of these two views, it's clear that a two-protocol world of IPv4-only and IPv6-only users is imminent. No matter which path organizations choose, they need to start planning a pragmatic transition to IPv6 using models and methods that keep their costs down and upgrade options open.

HOW BROCADE CAN HELP

For organizations to make the transition smoothly and without interruption, Brocade recommends that organizations impacted by the move to IPv6 should have their vendors address these questions:

- Which features are available with IPv6?
- Does the equipment meet performance requirements for IPv6?
- What is the impact of adding IPv6 to existing network devices?
- Can the device provide the necessary intelligence when an address is translated from one format to another?

Traditional Layer 3 routers can provide simple address translation between IPv4 and IPv6 addresses; however, they do not offer any mechanism to maintain address intelligence between two end points. For example, some applications rely heavily on source IP address of the originating client for auditing and tracking transactions, and for maintaining integrity of transactions.

An application delivery switch not only provides address translation, but also offers a mechanism to retain address information between two end points. For example, an application delivery switch can leverage HTTP capabilities to insert the source-IP address of an originator into a custom header so that back-end application servers can utilize it for tracking purposes.

To facilitate seamless communication with the new breed of IPv6-only customers in addition to IPv4 customers, the Brocade ServerIron ADX Series offers a simple and cost-effective transition path to IPv6 using a standards-based NAT64 gateway.

Topology A: The NAT64 gateway capabilities of Brocade ServerIron ADX enables organizations to bring new IPv6 customers onboard while utilizing their existing IPv4-based infrastructures.

Topology B: Similarly, Brocade ServerIron ADX enables organizations to offer new IPv6-based services to their existing IPv4 clients.

Topology C: Organizations can leverage the IPv6 application-delivery capabilities of Brocade ServerIron ADX to ensure equal service access by IPv4 and IPv6 users while distributing traffic among existing IPv4-only application servers. The IT team at Brocade is currently using this technology to deliver three critical Internet services—HTTP, DNS, and e-mail—in native IPv6 format.

In addition, Brocade delivers a comprehensive portfolio of products that deliver high-performance, and feature-rich IPv4 and IPv6 switching and routing solutions.

Figure 5. ServerIron ADX NAT64 gateway enables a seamless transition to IPv6.

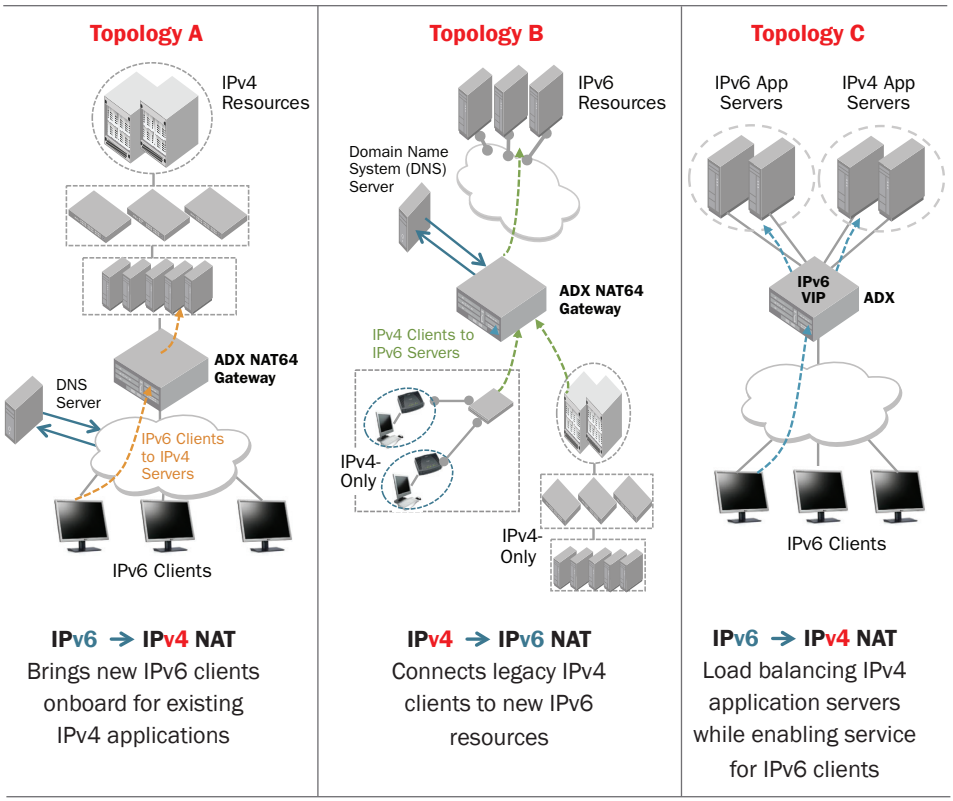

CONCLUSION

The Brocade ServerIron ADX Series helps facilitate seamless migration to strategic Internet technologies such as IPv6, helping organizations achieve continuity and infrastructure simplicity. It is a purpose-built platform that delivers unmatched performance value.

WHAT'S NEW IN IPV6

The IPv6 protocol brings a variety of new mechanisms to improve Internet reliability. These include the following:

- Significantly large address space
- Simplified network management using stateless auto-configuration of nodes
- Routing efficiency due to use of fixed-length headers
- Reduction in network processor overhead due to reduced fragmentation
- Improved security (IPsec is built-in)
- Well-defined flow labels for Quality of Service
- End-to-end address transparency

For more information about Brocade solutions, visit www.brocade.com.

SEE ALSO

http://www.brocade.com/adx

http://www.brocade.com/downloads/documents/at_a_glance/SI_ADX_12.2_IPv6_GA-AG-00.pd

REFERENCES

(sources and more information)

IPv4 Address Depletion
http://www.nro.net/news/ipv4-free-pool-depleted

NAT444
http://tools.ietf.org/html/draft-shirasaki-nat444-03

NAT64
http://datatracker.ietf.org/doc/draft-ietf-behave-v6v4-xlate-stateful/

Dual-Stack Lite
http://datatracker.ietf.org/doc/draft-ietf-softwire-dual-stack-lite/

6rd
http://datatracker.ietf.org/doc/rfc5969/

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

© 2011 Brocade Communications Systems, Inc. All Rights Reserved. 03/11 GA-WP-1574-00

Brocade, the B-wing symbol, BigIron, DCFM, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, TurboIron, and Wingspan are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, Extraordinary Networks, MyBrocade, VCS, and VDX are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

BROCADE