Managing IPv6 Throughout the Application Delivery Network

Managing IPv6 Throughout the Application Delivery Network

White Paper

Introduction

Organizations of all types are feeling increasing pressure to transition from the well-known and universal Internet Protocol version 4 (IPv4) standard to the newer IPv6 standard, while still supporting both network topologies. There are many reasons for this, not the least of which are the continually shrinking numbers of available IPv4 addresses and the exploding number of devices that require access to Internet applications and services.

Although the IPv6 standard includes many important new features beyond scaling past the limited IPv4 address space, such as increased security and reliability, the world still runs largely on IPv4. As new network technologies continue to drive users and services toward what will eventually be an all-IPv6 network, enterprise IT needs to be ready to adapt, manage, and support a dual-network architecture for the duration of the transition.

The Transition Challenge

Most organizations will not be able to simply flip a switch to make all their applications and equipment IPv6 instead of IPv4. To successfully transition to IPv6, organizations must be able to manage and design applications, network infrastructure, and security systems that simultaneously support both IPv4 and IPv6. But the transition doesn't stop at the network level: services that use the network must also be addressed. Critical network and application services, such as network firewalls, user access management tools, and advanced application delivery tools, must be factored into any IPv6 migration plan.

For most organizations, IT's functional needs will be the primary driver of network migrations to IPv6 as new technologies come into the data center and public IPv4 addresses become scarce. This type of transition will leave some locations and services on IPv4 while other parts of the organization transition to IPv6, affecting not only the core infrastructure, but the users and services that rely on these networks. Isolated IPv4 networks will still need to be able to seamlessly interoperate with the rest of the organization's users and systems on the IPv6 networks, and vice versa.

Making migration plans even more complex, the new technologies and integration challenges associated with a transition can create security challenges and more risk in the data center. Because of this increased security risk, firewalls and other network and application security tools must also be able to simultaneously support IPv4 and IPv6 traffic; otherwise, enterprises could open themselves up to new threats solely due to interoperability issues at the network level.

Read the case study, F5 Supports IPv6 on Core Network, to learn how F5 IT manages the transition to IPv6 for mobile users.

To properly handle the burden of introducing and supporting IPv6, organizations need a smart migration plan and tools to help provide an orderly transition between the two standards. These tools should give the organization the freedom to test, move, and migrate its existing infrastructure at a controlled, secure, and manageable pace. F5 BIG-IP products provide seamless support for both IPv4 and IPv6 networks, allowing organizations to transparently manage application delivery, availability, performance, and security between both network topologies at one central location—all without the need to deploy point products through the infrastructure.

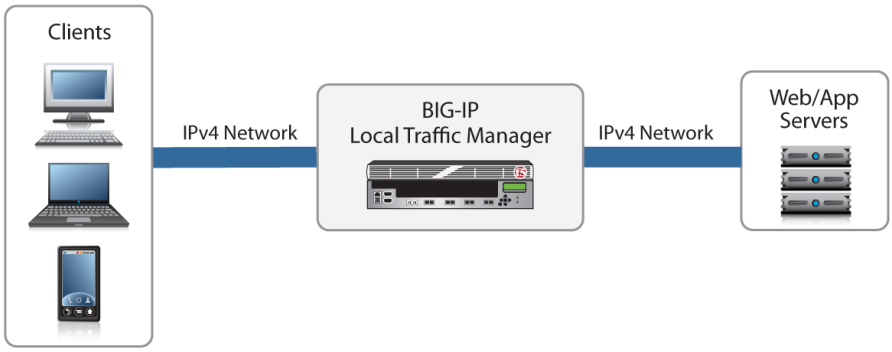

Figure 1: IPv4 clients connect to a standard IPv4 network.

IPv6 Migration Strategies

While an organization can natively introduce IPv6 into the infrastructure at any point, it's much more common to stagger an IPv6 implementation by dealing with one side of the network first. The two main scenarios for a smooth, controlled transition are for an organization to either migrate clients to IPv6 while keeping the servers on IPv4, or it can migrate servers to IPv6 while leaving the clients in an IPv4 environment. Both scenarios, however, have implications for users and applications alike.

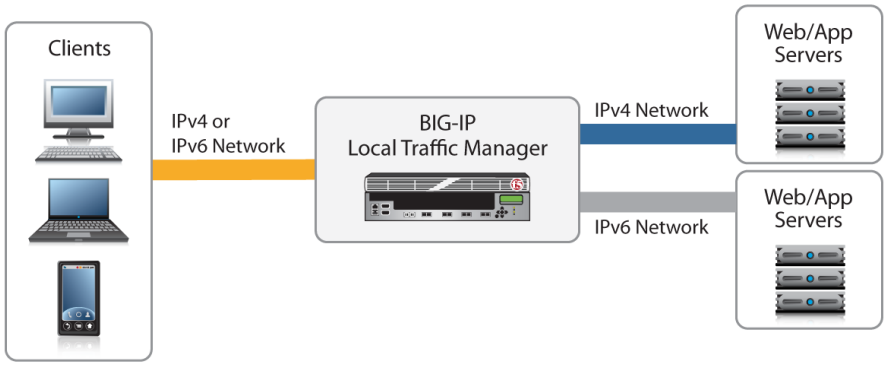

Figure 2: When both IPv4 and IPv6 networks exist, both client types can connect to any server.

The first scenario, migrating clients to IPv6 or supporting native IPv6 clients such as smartphones and other new devices, requires that all the clients be able to directly attach to the network and access services via IPv6-enabled pathways. Native IPv6 support for these new devices on the client side involves touching potentially every client device, implementing new access policies across the network, and incorporating new infrastructure services such as Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS). User access and management tools also need to support an IPv6 environment, or users may lose access to application services that are strictly dependent on an IPv4 network. While the goal may be to support only IPv6-enabled clients, the end result often has implications across the front- and back-end networks, as well as every part of the infrastructure.

In the second scenario of migrating application services to IPv6, it is likely that moving the servers to IPv6 will reveal some application dependencies that potentially affect all users simultaneously if something goes wrong. Even so, most organizations will find it much easier to begin migrating their applications before their clients, simply because the servers are completely under their control whereas client devices often are not. In addition, clients will continue for some time to exhibit a need to use IPv4 communication for public resources, as will many of the applications that are delivered from those IPv6-enabled servers.

In reality, migrations are very seldom as clear-cut as either of the scenarios above. More often a combination of IPv6 requirements drives the need to migrate: supporting new IPv6-enabled devices while also upgrading application services and the core infrastructure to support these new devices and new services enabled by IPv6. In meeting all of these requirements, maintaining seamless support for all clients across all services is paramount. As organizations across the globe struggle to transition to IPv6, the critical nature of access and connectivity between devices and services across all networks will increase the complexity of application services.

The BIG-IP System: A Gateway for Transition

At the core of the Application Delivery Network (ADN) is the F5 BIG-IP platform, which provides application delivery services ranging from high availability, SSL processing, and caching and compression, to more advanced services such as application security, user access management, and WAN optimization. For all application delivery services, the BIG-IP platform functions as a native IPv4-to-IPv6 gateway by transparently managing application delivery in both networking topologies, which enables it to continue supporting advanced services as applications traverse both networks.

IPv6 Migration at the Strategic Point of Control

In a typical BIG-IP deployment, the BIG-IP device is situated between the clients and the servers to provide the applications the clients use. In this position—the strategic point of control—the BIG-IP device provides virtualization and high availability for all application services, making several physical servers look like a single entity behind the BIG-IP device. This virtualization capability provides an opportunity to start migrating either clients or servers—or both simultaneously—to IPv6 networks without having to change clients, application services, and both sides of the network all at once.

Initially, most IPv6 migration plans focused on the back-end network moving to IPv6 addresses for application servers and the corporate network. This involved a lot of migrating the private network and on-site application servers in preparation for future IPv6 requirements in applications and with new devices. But as newer IPv6-enabled devices begin to penetrate the enterprise, the focus is shifting to supporting those devices in a native IPv6 client-side network. The BIG-IP platform provides centralized IPv6 gateway functions across the entire product suite and on both sides of the network, offering independent and flexible migration solutions regardless of where the organization needs to focus immediate IPv6 support.

Migrating servers

With a BIG-IP appliance strategically located between clients and servers, a network administrator can simply add a new "server" network to the BIG-IP device—one that is IPv6-capable. The result is that the network will have IPv4 on the front (client) side of the BIG-IP device and simultaneously support both an IPv4 and IPv6 network behind it.

Most public networks only support IPv4. However some parts of the world, such as parts of Asia, and some mobile networks have already migrated to IPv6 and can natively support both clients and applications services—both sides of the network—in an all-IPv6 environment.

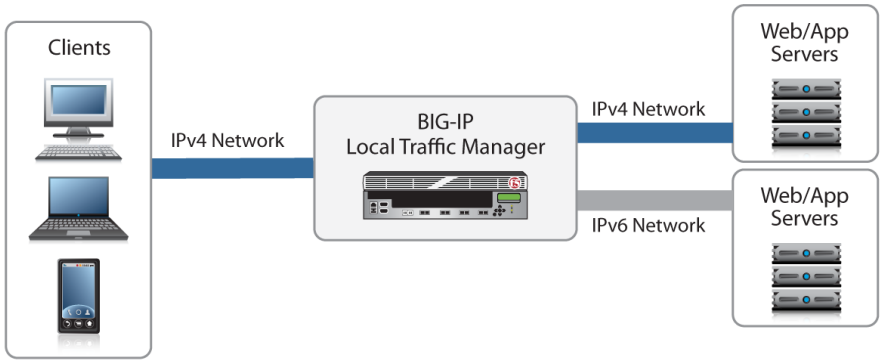

Figure 3: BIG-IP LTM enables organizations to gradually migrate servers to IPv6.

Once the IPv6 network is established, the organization can start migrating its servers from the IPv4 network. For example, if it has several back-end servers providing services for an application, it can simply take one IPv4 server offline and switch it to IPv6. The BIG-IP platform will handle application delivery requests by load balancing among all servers while continuing to provide IPv4 virtual addresses to the clients. The clients themselves won't realize any difference because they are still contacting and using the IPv4 virtual server being serviced by BIG-IP Local Traffic Manager (LTM) to access their applications.

After the offline server is switched to IPv6, administrators can bring it back online and add it back to the original load balancing pool with its IPv6 address instead of the old IPv4 address. BIG-IP will incorporate new servers in the IPv6 environment, simultaneously providing application availability in this mixed environment, but clients will still use the old IPv4 virtual address to connect to those services. Client requests will now be load balanced across all of the IPv4 and IPv6 servers.

As back-end servers are migrated to the new IPv6 environment, additional services, such as network and application firewalls, can also be transparently migrated to the IPv6 network. A device that can only provide security for each network topology independently requires that the organization deploy multiple point solutions throughout the mixed networking environment and keep security zones isolated between IPv4 and IPv6 networks. With the BIG-IP platform, organizations can provide simultaneous security for both IPv4 and IPv6 servers and applications in a mixed environment, where some servers are IPv4-enabled and some are IPv6, regardless of whether those servers and applications require one type of network or the other. With this strategy, the organization can complete the entire server migration with no effect on clients and no downtime.

Throughout the transition to IPv6, IT organizations must also maintain application and network security. As more types of devices and applications are developed, more security threats develop as well. To prevent new security vulnerabilities and provide network continuity, organizations must preserve seamless and secure application access. These increasing threats can also cause more complexity in IPv6 migration. The BIG-IP platform enables administrators to secure access to both IPv4 and IPv6 networks with top-level, high-performance management, including integrated and simultaneous high-speed VPN and SSL connections on both sides of the IPv4 and IPv6 networks. To achieve the same level of security that existed on the legacy IPv4 network, the BIG-IP platform acts as an application firewall for both IPv4 and IPv6 networks by applying granular, application-level policies. It is a flexible and adaptive solution to the increasing security risks at an application level for both IPv4 and IPv6 networks.

Supporting IPv6-capable clients

Organizations can use a similar strategy to move their clients to IPv6 without migrating back-end servers. In this case, along with its IPv4 virtual address that points to its IPv4 servers, the organization creates a client-facing IPv6 interface with a new IPv6 virtual address that points to those same IPv4 servers. Then, as clients are transitioned to the new IPv6 client network, BIG-IP appliances, which provide IPv6 DNS resolution services, will "hand out" the new IPv6 address of the virtual server, using the same DNS name that previously pointed to the IPv4 address.

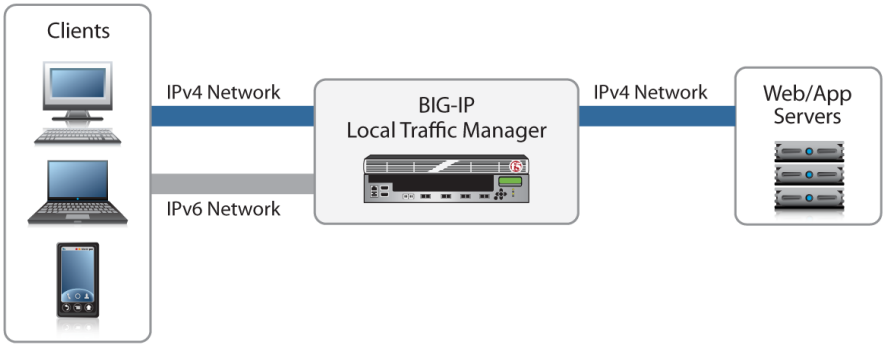

Figure 4: BIG-IP LTM enables IPv6 clients to connect to the IPv4 network.

By using the same host name and back-end servers on both the IPv4 and IPv6 networks, in most cases the clients will be able to start using their old applications as if nothing has changed.

IPv6 application awareness

One challenge that organizations often don't deal with until after an application is moved to an IPv6 network is how applications communicate with each other across the network. Many IP-based applications are hard-wired to use IPv4 addresses and can't work in a mixed IP environment; web-based widgets provide an excellent example of application components that are often explicitly bound to one network type or the other. It's not uncommon for a web application to make external calls to other application services for additional functionality, such as "liking" or sharing content, pulling dynamic content from external sources, or calling a specific web-based function such as a java-based chat box. Once the base web application server is transitioned to an IPv6 network, it begins making those external application calls across the same IPv6 network, but the destination applications may not be IPv6aware or able to function at all on an IPv6 network.

By using the BIG-IP platform, inter-application traffic can be fully proxied between the IPv6 and IPv4 networks, allowing the web application server to make outbound calls over an IPv6 network that are translated to the IPv4 network where the external application data resides. Using a BIG-IP appliance as a full IPv6 proxy also allows IT to continue supporting legacy IPv4 applications that would otherwise never work on an IPv6 network; in this way, organizations can perform a graceful migration from legacy applications rather than forcing a last-minute crisis as end-users lose access to those applications.

Conclusion

The transition to IPv6 is something that every enterprise IT department will have to deal with. This is due to IPv4 depletion, but also to new, IPv6-native application services and devices becoming available to end-users. The BIG-IP system provides the flexibility for an organization to securely migrate IPv4 network services and clients at its own pace, while maintaining control of the application and network. If some applications can't be moved or don't support IPv6, they can be left on IPv4 until they are replaced (assuming the organization will be required to use all IPv6) or retired. In the same manner, clients that still need to maintain their IPv4 identity can either utilize both networks (as applicable) or can simply continue to use the IPv4 network and access the organization's IPv6 application services via the BIG-IP platform.

Regardless of when or why IPv6 moves into the enterprise, F5 provides a seamless, integrated solution for managing a controlled IPv4-to-IPv6 rollout across the entire Application Delivery Network.

F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com