2

Introduction to IPv6

Contents

Migrating to IPv6

To successfully migrate to IPv6 involves maintaining compatibility with the large installed base of IPv4 hosts and routers for the immediate future. To achieve this purpose, software release K.13.01 supports dual-stack (IPv4/IPv6) operation and connectons to IPv6-aware routers for routing IPv6 traffic between VLANs and across IPv4 networks.

Note

Software release K.13.01 supports traffic connections with IPv6-aware routers, but does not support IPv6 routing operation in the switches covered by this guide.

Beginning with software release K.13.01, the switches covered by this guide support the following IPv6 protocol operations:

- receiving IPv6 traffic addressed to the switch
- transmitting IPv6 traffic originating on the switch
- switching IPv6 traffic between IPv6 devices connected to the switch on the same VLAN
- concurrent (dual-stack) operation with IPv4 traffic and devices on the same VLAN
- using a connection to an external, IPv6-configured router, forward IPv6 traffic intended for devices on other VLANs and for traffic that must traverse an IPv4 network to reach an IPv6 destination

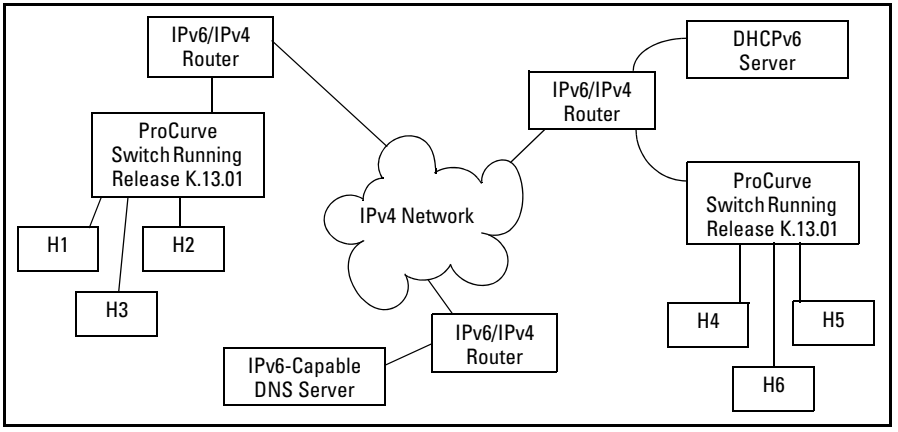

Figure 2-1. Dual-Stack ProCurve Switches Employed in an IPv4/IPv6 Network

IPv6 Propagation

IPv6 is currently in the early stages of deployment worldwide, involving a phased-in migration led by the application of basic IPv6 functionality. In these applications, IPv6 traffic is switched among IPv6-capable devices on a given LAN, and routed between LANs using IPv6-capable routers. Using the IPv6 features in this software release, the switch can operate in an IPv6 network, be managed using an IPv6 management station, and interact with DHCPv6 and IPv6-enabled DNS servers in the same network or accessible through a connection to an IPv6 router.

Dual-Stack Operation

Since most initial IPv6 deployments are in networks having a mixture of IPv6 and IPv4 hosts software release K.13.01 supports dual- stack IPv4/IPv6 operation. This enables the switch to communicate individually with IPv4 and IPv6 devices with their respective protocols. Thus, IPv4 and IPv6 traffic is supported simultaneously on the same VLAN interface. This means that both IPv4 and IPv6 devices can operate at the same time on a given VLAN.

Note

Software release K.13.01 does not include gateways for translation between IPv6 and IPv4 traffic. While IPv4 and IPv6 traffic coexists on the same VLAN, the individual IPv4 and IPv6 devices ignore each other's traffic.

To forward IPv6 traffic from the switch to an IPv6-capable device on a different VLAN, a link to an external IPv6-capable router is needed. Also, IPv6 traffic movement from the switch over IPv4 paths requires routers capable of IPv6 over IPv4 tunneling.

Connecting to Devices Supporting IPv6 Over IPv4 Tunneling

The switches covered by this guide can interoperate with IPv6/IPv4 devices capable of tunneling IPv6 traffic across an IPv4 infrastructure. Some examples include:

- traffic between IPv6/IPv4 routers(router/router)
- traffic between an IPv6/IPv4 router and an IPv6/IPv4 host capable of tunneling (router/host)

Note

Tunneling requires an IPv6-capable router. A switch running software release K.13.01 does not route or tunnel IPv6 traffic. To enable IPv6 traffic from the switch to be routed or to be tunneled across an IPv4 network, it is necessary to connect the switch to an appropriate IPv6-capable router. For more information, refer to the documentation provided with the dual- stack (IPv4/IPv6) routers you plan to use for this purpose.

IPv6 tunneling eases IPv6 deployment by maintaining compatibility with the large existing base of IPv4 hosts and routers. Generally, the various IPv6 tunneling methods enable IPv6 hosts and routers to connect with other IPv6 hosts and routers over the existing IPv4 Internet.

Information Sources for Tunneling IPv6 Over IPv4

For more information on IPv6 routing and tunneling, refer to the documentation provided with the IPv6/IPv4 routing and tunneling-capable devices in your network. Some other sources of information are:

- RFC 2893: “Transition Mechanisms for IPv6 Hosts and Routers”
- RFC 2401: “Security Architecture for the Internet Protocol”
- RFC 2473: “Generic Packet Tunneling in IPv6 Specification”
- RFC 2529: “Transmission of IPv6 via IPv4 Domains without Explicit Tunnels”
- RFC 3056: “Connection of IPv6 Domains Over IPv4 Clouds”

Use Model

Adding IPv6 Capability

IPv6 was designed by the Internet Engineering Task Force (IETF) to improve on the scalability, security, ease of configuration, and network management capabilities of IPv4.

IPv6 provides increased flexibility and connectivity for existing networked devices, addresses the limited address availability inherent in IPv4, and the infrastructure for the next wave of Internet devices, such as PDAs, mobile phones and appliances.

Where IPv4 networks exist today, IPv6 will be phased in over a period of years, requiring an interoperability among the devices using the two protocols. Beginning with software release K.13.01, the switches covered by this guide offer IPv4/IPv6 dual stack operation. This allows full ethernet link support for both IPv4 and IPv6 traffic to move on the same interface (VLAN) without modifying current IPv4 network topologies. This enables you to use IPv6 devices on existing VLANs, manage the switch and other devices from IPv6 management stations, and create "islands" of IPv6 devices as needed to accomodate the need for the IPv6 network growth anticipated for the future.

Supported IPv6 Operation in Release K.13.01

Software release K.13.01 provides IPv6 protocol and addressing to support host-mode (endpoint) IPv6 operation, including basic layer-2 functionality. IPv6 routing features are not available in this release. However, using a dual-stack (IPv4/IPv6-capable) router, IPv6 traffic can be routed between VLANs and sent across an IPv4 network to another IPv6 device.

(For general information on sending IPv6 traffic across an IPv4 network, refer to “Connecting to Devices Supporting IPv6 Over IPv4 Tunneling” on page 2-5.)

The IPv6 features available in release K.13.01 belong to these general categories:

- switch configuration and management
- security
- IPv6 multicast traffic
- diagnostic and troubleshooting

The next three sections outline the IPv6 features supported in software release K.13.01.

Configuration and Management

This section outlines the configurable management features supporting IPv6 operation on your ProCurve IPv6-ready switch.

Management Features

Software release K.13.01 provides host-based IPv6 features that enable the switches covered in this guide to be managed from an IPv6 management station and to operate in both IPv6 and IPv4/IPv6 network environments.

Note

Software release K.13.01 does not include IPv6 routing, but interoperates with routers that support IPv6 and IPv4/IPv6 router applications.

IPv6 Addressing

The switch offers these IPv6 address configuration features:

- SLAAC (stateless automatic address configuration)
- DHCPv6 (stateful automatic address configuration)
- static address configuration

SLAAC (Stateless Automatic Address Configuration)

Enabling IPv6 on a VLAN automatically enables configuration of a link-local unicast IPv6 address on the VLAN. (No DHCPv6 server is needed.) This address begins with the hexadecimal prefix **fe80**, which is prepended to the interface identifier part of the address. (The interface identifier is generated from the MAC address of the VLAN itself, using the 64-bit extended unique identifier (EUI) method.) This enables the IPv6 nodes on the VLAN to configure and manage the switch.

Enabling IPv6 address autoconfiguration on a VLAN automatically enables automatic configuration of global unicast addresses on the VLAN. After enabling autoconfiguration, a router advertisement (RA) containing an assigned global address prefix must be received on the VLAN from an IPv6 router on the same VLAN. The resulting address is a combination of the prefix

and the interface identifier currently in use in the link-local address. Having a global unicast address and a connection to an IPv6- aware router enables IPv6 traffic on a VLAN to be routed to other VLANs supporting IPv6-aware devices. (Using software release K.13.01, an external, IPv6- aware router is required to forward traffic between VLANs.)

Multiple, global unicast addresses can be configured on a VLAN that receives RAs specifying different prefixes.

DHCPv6 (Stateful) Address Configuration

The IPv6 counterpart to DHCP client for IPv4 operation is DHCPv6. Global unicast addresses of any scope can be assigned, along with NTP (timep) server addressing when DHCPv6 server support is available through either of the following modes:

- accessible on a VLAN configured on the switch
- accessible through a connection to a router configured with DHCP relay

IPv6 also allows the option of using stateless autoconfiguration or static configuration to assign unicast addresses to a VLAN, while using a DHCPv6 server for time server addressing.

Static Address Configuration

Statically configuring IPv6 addresses provides flexibility and control over the actual address values used on an interface. Also, if a statically configured link-local address is configured on a static VLAN, the global addresses configured on the VLAN as the result of router advertisements uses the device identifier included in the link-local address. Statically configuring an IPv6 address on a VLAN enables IPv6 on the VLAN if it has not already been enabled.

Default IPv6 Gateway

Instead of using static or DHCPv6 configuration, a default IPv6 gateway for an interface (VLAN) is determined from the default router list of reachable or probably reachable routers the switch detects from periodic multicast router advertisements (RAs) received on the interface. For a given interface, there can be multiple default gateways, with different nodes on the link using different gateways. If the switch does not detect any IPv6 routers that are reachable from a given interface, it assumes (for that interface) that it can reach only the other devices connected to the interface.

Note

In IPv6 for the switches covered in this guide, the default route cannot be statically configured. Also, DHCPv6 does not include default route configuration.)

Refer to “Default IPv6 Router” on page 4-28 and “View IPv6 Gateway, Route, and Router Neighbors ” on page 4-29.

Neighbor Discovery (ND) in IPv6

The IPv6 Neighbor Discovery protocol operates in a manner similar to the IPv4 ARP protocol to provide for discovery of IPv6 devices such as other switches, routers, management stations, and servers on the same interface. Neighbor Discovery runs automatically in the default configuration and provides services in addition to those provided in IPv4 by ARP. For example:

- Run Duplicate Address Detection (DAD) to detect duplicate unicast address assignments on an interface. An address found to be a duplicate is not used, and the **show ipv6** command displays the address as a **duplicate**.
- Quickly identify routers on an interface by sending router solicitations requesting an immediate router advertisement (RA) from reachable routers.
- If a default router becomes unreachable, locate an alternate (if available on the interface).
- Learn from reachable routers on the interface whether to use DHCPv6 or stateless address autoconfiguration. In the latter case, this also includes the address prefixes to use with stateless address autoconfiguration for routed destinations. (A DHCPv6 server can also be used for "stateless" service; that is, for configuring the interface for access to other network services, but not configuring a global IPv6 unicast address on the interface. Refer to “Neighbor Discovery (ND)” on page 4-17.)
- Use multicast neighbor solicitations to learn the link-layer addresses of destinations on the same interface and to verify that neighbors to which traffic is being sent are still reachable.
- Send a multicast neighbor advertisement in response to a solicitation from another device on the same interface or to notify neighbors of a change in the link- layer address.
- Advertise anycast addresses that may be configured on the device.
- Determine the MTU (Maximum Transmission Unit) for the interface from router advertisements.

For more on IPv6 neighbor discovery applications, refer to “Neighbor Discovery (ND)” on page 4-17.

IPv6 Management Features

The switch's IPv6 management features support operation in an environment employing IPv6 servers and management stations.With a link to a properly configured IPv6 router, switch management extends to routed traffic solutions. (Refer to the documentation provided for the IPv6 router.) Otherwise, IPv6 management for the switches covered by this guide are dependent on switched management traffic solutions.

TFTPv6 Transfers

The switch supports these downloads from an IPv6 TFTP server:

- automatic OS download
- manual OS download
- command script download and execution
- configuration file downloads
- public key file downloads
- startup configuration file downloads

The switch supports these uploads to an IPv6 TFTP server

- startup or running configuration upload
- OS upload from flash in current use (primary or secondary)
- event log content upload
- crash log content upload
- output of a specified command

Refer to “TFTP File Transfers Over IPv6” on page 5-15.

IPv6 Time Configuration

The switch supports both Timepv6 and SNTPv6 time services. Refer to “SNTP and Timep” on page 5-9.

Telnet6

The switch supports both of the following Telnet6 operations:

- Enable (the default setting) or disable Telnet6 access to the switch from remote IPv6 nodes.
- Initiate an outbound telnet session to another IPv6 networked device.

Refer to “Telnet6 Operation” on page 5-6

IP Preserve

IP Preserve operation preserves both the IPv4 and IPv6 addresses configured on VLAN 1 (the default VLAN) when a configuration file is downloaded to the switch using TFTP. Refer to “IP Preserve for IPv6” on page 5-23.

Multicast Listener Discovery (MLD)

MLD operates in a manner similar to IGMP in IPv4 networks. In the factory default state (MLD disabled), the switch floods all IPv6 multicast traffic it receives on a given VLAN through all ports on that VLAN except the port receiving the inbound multicast traffic. Enabling MLD imposes management controls on IPv6 multicast traffic to reduce unnecessary bandwidth usage. MLD is configured per- VLAN. For information on MLD, refer to the chapter titled “Multicast Listener Discovery (MLD) Snooping”.

Web Browser Interface

For the web browser interface, software release K.13.01 adds the following IPv6 functionality:

- configure and display IPv6 addressing
- ping6 diagnostic operation

Configurable IPv6 Security

This section outlines the configurable IPv6 security features supported in software release K.13.01. For further information on these features, refer to the indicated pages.

SSHv2 on IPv6

SSHv2 provides for the authentication between clients and servers, and protection of data integrity, and privacy. It is used most often to provide a secure alternative to Telnet and is also used for secure file transfers (SFTP and SCP). Software release K.13.01 with SSHv2 on IPv6 extends to IPv6 devices the SSH functionality that has been previously available on ProCurve switches running IPv4. This means that SSH version 2 connections are

supported between the switch and IPv6 management stations when SSH on the switch is also configured for IPv6 operation. The switch now offers these SSHv2 connection types:

- IPv6 only
- IPv4 only
- IPv4 or IPv6

The switch supports up to six inbound sessions of the following types in any combination at any given time:

- SSHv2
- SSHv2 IPv6
- Telnet-server
- Telnet6-server
- SFTP/SCP
- Console (serial RS-232 connection)

For more information, refer to “Secure Shell for IPv6” on page 6-15.

IP Authorized Managers

The IPv6 Authorized IP Managers feature, like the IPv4 version, uses IP addresses and masks to determine which stations (PCs and workstations) can access the switch through the network, and includes these access methods:

- Telnet, SSH, and other terminal emulation applications
- the switch's web browser interface
- SNMP (with a correct community name)

Also, when configured in the switch, the access control imposed by the Authorized IP Manager feature takes precedence over the other forms of access control configurable on the switch, such as local passwords, RADIUS, and both Port-Based and Client-Based Access Control (802.1X). This means that the IP address of a networked management device must be authorized before the switch will attempt to authenticate the device by invoking any other access security features. Thus, with Authorized IP Managers configured, having the correct passwords or MAC address is not sufficient for accessing the switch through the network unless an IPv6 address configured on the station attempting the access is also included in the switch's Authorized IP Managers configuration. This presents the opportunity to combine the Authorized IP Managers feature with other access control features to enhance the security fabric protecting the switch.

Caution

The Authorized IP Managers feature does not protect against unauthorized station access through a modem or direct connection to the Console (RS-232) port. Also, if an unauthorized station "spoofs" an authorized IP address, then the unauthorized station cannot be blocked by the Authorized IP Managers feature, even if a duplicate IP address condition exists.

To configure authorized IPv6 managers, refer to "Authorized IP Managers for IPv6" on page 6-3.

For related information, refer to:

- RFC 4864, "Local Network Protection for IPv6".

Diagnostic and Troubleshooting

Software release K.13.01 includes the IPv6 diagnostic and troubleshooting features listed in this section.

ICMP Rate-Limiting

Controlling the frequency of ICMPv6 error messages can help to prevent DoS (Denial- of- Service) attacks. With IPv6 enabled on the switch, you can control the allowable frequency of these messages with ICMPv6 rate-limiting. Refer to "ICMP Rate-Limiting" on page 8-2.

Ping6

Implements the Ping protocol for IPv6 destinations, and includes the same options as are available for IPv4 Ping, including DNS hostnames. Refer to "Ping for IPv6 (Ping6)" on page 8-4.

Traceroute6

Implements Traceroute for IPv6 destinations, and includes the same same options as are available for the IPv4 Traceroute, including DNS hostnames. Refer to "Traceroute for IPv6" on page 8-6.

Domain Name System (DNS) Resolution

This feature enables resolving a host name to an IPv6 address and the reverse, and takes on added importance over its IPv4 counterpart due to the extended length of IPv6 addresses. With DNS-compatible commands, CLI command entry becomes easier for reaching a device whose IPv6 address is configured with a host name counterpart on a DNS server.

Software release K.13.01 includes the following DNS-compatible commands:

- **ping6**
- **traceroute6**

The switches covered by this guide now support a prioritized list of up to three DNS server addresses. (Earlier software releases supported only one DNS server address.) Also, the server address list can include both IPv4 and IPv6 DNS server addresses. (An IPv6 DNS server can respond to IPv4 queries, and the reverse.)

Note

If an IPv6 DNS server address is configured on the switch, at least one VLAN on the switch (and in the path to the DNS server) must be configured with an IPv6 address.

For information on configuring DNS resolution on the switch, refer to “DNS Resolver for IPv6” on page 8-9.

IPv6 Neighbor Discovery (ND) Controls

The neighbor discovery feature includes commands for:

- increasing or decreasing the frequency of Duplicate Address Detection searches
- displaying the IPv6 neighbor cache
- clearing dynamic entries from the neighbor cache

Refer to “Neighbor Discovery (ND) in IPv6” on page 2-9.

Event Log

Messages returning IP addresses now include IPv6 addresses where applicable.

SNMP

When IPv6 is enabled on a VLAN interface, you can manage the switch from a network management station configured with an IPv6 address. Refer to "SNMP Management for IPv6" on page 5-20.

Loopback Address

Like the IPv4 loopback address, the IPv6 loopback address (::1) can be used by the switch to send an IPv6 packet to itself. However, the IPv6 loopback address is implicit on a VLAN and cannot be statically configured on any VLAN. Refer to "Loopback Address" on page 3-24.

Debug/Syslog Enhancements

Includes new options for IPv6. Refer to "Debug/Syslog for IPv6" on page 8-12.

IPv6 Scalability

As of software release K.13.01, the switches covered by this guide support the following:

- Dual stack operation (IPv4 and IPv6 addresses on the same VLAN).
- Maximum of 512 VLANs with IPv4 and IPv6 addresses in any combination.
- Up to 2048 VLANs configured on the switch.
- Maximum of 2048 active IPv6 addresses on the switch, in addition to a maximum of 2048 IPv4 addresses. ("Active IPv6 addresses" includes the total of all preferred and non-preferred addresses configured statically, through DHCPv6, and through stateless autoconfiguration. Excluded from "Active IPv6 Addresses" is the link-local address assigned to each VLAN, and "on- link" prefixes received as part of a router advertisement.)
- Maximum of 32 IPv6 addresses on a VLAN.
- Maximum of 10,000 IPv6 routes.

For more information on VLAN and route scalability on the switches covered by this guide, refer to the appendix titled "Scalability: IP Address, VLAN, and Routing Maximum Values" in the *Management and Configuration Guide* for your switch.

Path MTU (PMTU) Discovery

IPv6 PMTU operation is managed automatically by the IPv6 nodes between the source and destination of a transmission. For Ethernet frames, the default MTU is 1500 bytes. If a router on the path cannot forward the default MTU size, it sends an ICMPv6 message (PKT_TOO_BIG) with the recommended MTU to the sender of the frame. If the sender of the frame is an IPv6 node that supports PMTU discovery, it will then use the MTU specified by the router and cache it for future reference.

For related information, refer to:

- RFC 1981: “Path MTU Discovery for IP version 6”