Migrating Industrial IPv4 Network to IPv6

Xuechen Yang
Software Architect
Cisco Systems, Inc.

Mitesh Dalal
Product Manager
Cisco Systems, Inc.

Presented at the ODVA
2014 Industry Conference & 16th Annual Meeting
March 11-13, 2014
Phoenix, Arizona, USA

Abstract:

IPv6 has been marketed as a standard enhancement to make the Internet technologies more scalable and capable of addressing more and more devices. Although true, IPv6 offers more advantages than just scalability, including security, improved multicasting, mobility and address abstraction. All of which have an impact on industrial networks.
General adoption of IPv6 has been slow across the board, but it does have growing momentum. Plant floors, being in the adoption phase of standard Ethernet networking, usually exclude IPv6. Migration concepts for industrial protocols and environments are rare.
This paper will review some of the advantages of IPv6 for Industrial applications and some migration concepts from other IT environments on how ODVA members and their customers can adopt the new technology.

Keywords:

IPv6, IPv4, Dual-Stack, NAT64, NAT64 Gateway, DNS64, ALG

Definition of terms:

ALG	Application Level Gateway is an "add-on" component to standard Firewall/NAT function on a router/gateway. It inspects and filters network traffic at application layer by looking into payload data of the TCP/IP packet. In the case where NAT is needed, ALG may need to translate the IP address (and TCP/UDP port number) embedded in the payload. Protocols that need an ALG include FTP, SIP, RTSP, PPTP, and EtherNet/IP... etc.
A Record	DNS resource record that returns a 32-bit IPv4 address, most commonly used to map hostnames to an IP address of the host, but also used for DNSBLs, storing subnet masks in RFC 1101, etc.
AAAA Record	DNS resource record that returns a 128-bit IPv6 address, most commonly used to map hostnames to an IP address of the host.
DNS64	A mechanism that translates between AAAA record and A record depending on client's request. The mapping between AAAA and A records can be either generated dynamically or pre-configured.

DNS64 Server A DNS server that supports DNS64 by, for example, returning the AAAA record of an IPv4 server to an asking IPv6 host, or returning the A record of an IPv6 server to an IPv4 endpoint.

NAT64 A mechanism that allows IPv6 and IPv4 hosts to communicate with each other through network address translation. It translates the IP addresses in the IP header, IP checksum, and TCP/UDP checksum (for TCP/UDP packets) for packets going each direction.

NAT64 Gateway A layer 3 gateway or router that is capable of perform NAT64, i.e. translating IPv4 packet to IPv6 packet and vice versa. Such translation is transparent to the IPv4 and IPv6 endpoints.

While most of the enterprise connected network is moving towards IPv6, we've started seeing the needs of migrating Industrial Control Network as well. For example, more devices are being connected to the Industrial network, which require unique IP addresses; and wireless networks are being designed using IPv6 from ground-up. For example, 6TiSCH, or "IPv6 over the TSCH (Time-slotted Channel Hopping) mode of IEEE 802.15.4e", enables a large number of resource-constrained nodes to form a wireless mesh network and talk IPv6. Below is a quick summary of IPv6 benefits:

- Bigger address space
- Efficient routing and packet processing
- Directed data flows
- Simplified network configuration
- Support for new services
- Built-in security

Three-Staged Approach

To minimize the financial impact and disruption on existing operation, we need a phased approach to migrate the Industrial Control Network to IPv6. Figure 1 shows the three stages:

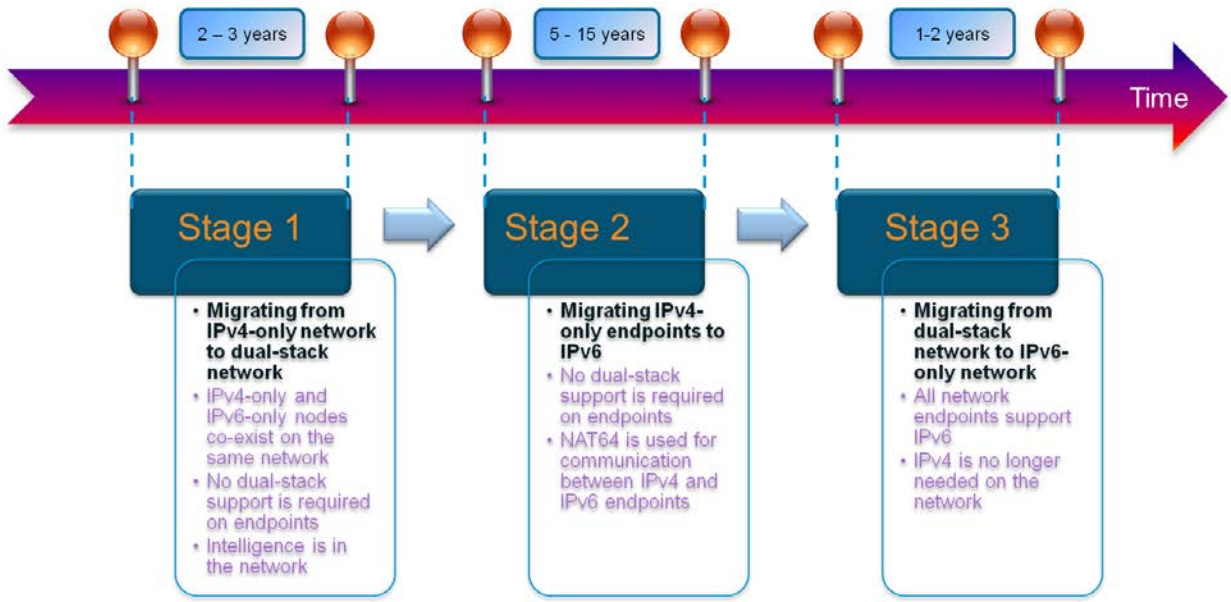

Figure 1: Three-Staged Migration Approach

This approach is driven by several motivational factors:

- Minimize network topology change
- Simplify upgrading process
- Endpoint upgrade is independent from network change

- Protect investment on existing machine endpoints for longer period

Stage 1 – Migrating IPv4-Only Network to Dual-Stack Network

The objectives of stage 1 are:

1. Build intelligent network that supports both IPv4 and IPv6 endpoints
2. Preserve existing network topology and protect existing network investment
3. Separate network migration completely from IPv6 upgrade on endpoints

Instead of transforming the entire Enterprise and Control Network to IPv6 to dual-stack in one single motion, it's better to divide the process into phases. For example, a customer can upgrade the Enterprise IT network and DMZ to dual-stack first, then the Supervisory network, and migrate the Control System network to dual-stack as the final step. The remote facility or field network can be migrated after the main site is upgraded. However there're wireless field network may need to be migrated to IPv6 first for several reasons:

- Massive number of wireless devices need unique IP addresses
- IPv6 is need to support the advanced routing needed for low-powered wireless network
- Manufacturers are already on-board building IPv6-capable wireless devices

Driven by both top-down and bottom-up forces, the actual migration to dual-stack network in stage 1 will look like what is depicted in Figure 2.

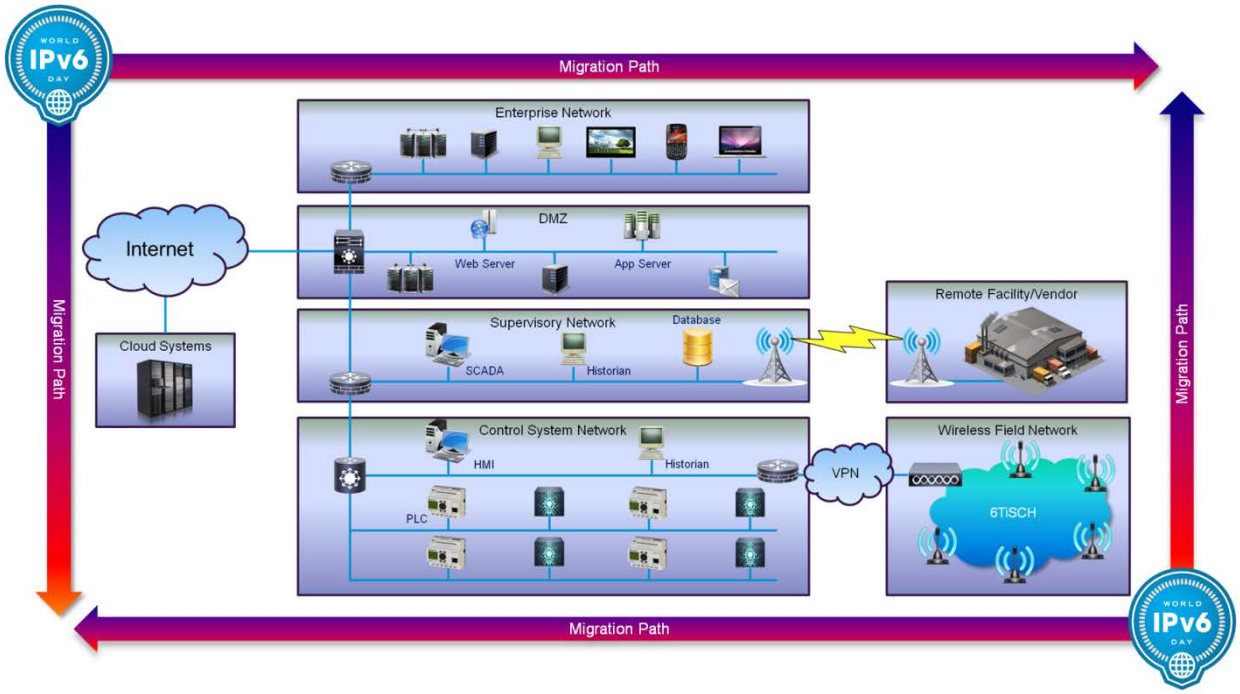

Figure 2: Dual-Stack Network Migration Path

Setting up dual-stack network primarily involves the provisioning and configuration of NAT64 gateways and DNS64 servers, including:

- Replace regular layer-3 routers and gateways with **NAT64-capable devices** via HW and/or SW upgrade
- Replace IPv4 DNS servers with **DNS64 servers**
- Install DHCPv6 servers to **serve stateful DHCP requests**.
- Use a **central NMS** (Network Management System) to manage all NAT64 Gateways and DNS64 Servers and ensure **consistent configuration across all systems**

Below is the list of recommendations and guidelines for provisioning and configuring the NAT64 gateways:

- All layer-3 routers and gateways should support NAT64.

- A separate IPv4 address pool should be configured for each NAT64 gateway to facilitate translation.
- To allow IPv4 clients access IPv6 servers, the same static address mappings must be created on NAT64 gateways for IPv6 servers.
- Routing must be configured properly for IPv4 and IPv6 networks to ensure correct path for translated traffic.

Below is the list of recommendations and guidelines for provisioning and configuring the DNS64 servers:

- The DNS64 server must support dual-stack and serve DNS requests from both IPv4 and IPv6 endpoints.
- Every A record should have a corresponding AAAA record with translated address, and vice versa (this mapping may be generated dynamically).
- Configuration on DNS64 servers must be consistent with NAT64 gateways.

This migration strategy assumes the network infrastructure and endpoints use DNS for device naming and address mapping. If DNS is not used or supported, in which case IP addresses are manually configured or distributed via some other protocol, this approach will continue to work as long as the IP address mappings are configured properly on the endpoints. The details of such scenarios are out of scope, and will not be covered by this paper.

Stage 2 – Migrating IPv4-Only Endpoints to IPv6

The objectives of stage 2 are:

1. Support each endpoint to upgrade to IPv6 independently
2. Allow different software and hardware products to be upgraded independently

Because the network supports dual-stack and NAT64 between IPv4 and IPv6 endpoints is transparent to endpoints, IPv4 endpoints can be upgraded to IPv6 at different times and independent from each other. Stage 2 migration actions include:

- Upgrade servers, employee desktops, laptops, and important IT assets to IPv6
- Upgrade HMI, Historian, and other assets on the Supervisory network to IPv6
- Upgrade PLCs, Drives, and other I/O devices to IPv6

During stage 2 migrations, the communication between two IPv4 endpoints and two IPv6 endpoints will remain the same regardless the endpoints on the same local network or not. However the communication between two endpoints running different IP stacks needs to be carefully examined. There're three typical communication scenarios as shown in Figure 3.1, 3.2, and 3.3:

1. IPv6-Only Host accesses IPv4 Server
2. IPv4-Only Host accesses IPv6 Server
3. Remote IPv4-Only Host accesses local IPv6 server via VPN

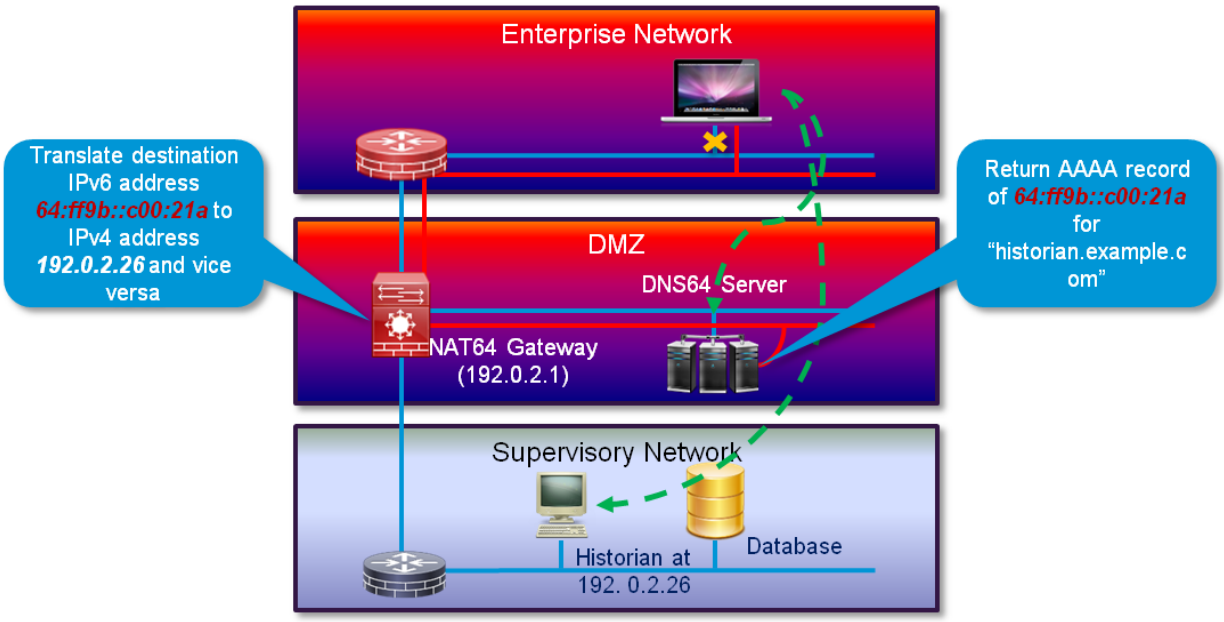

Figure 3.1: IPv6-only Host access IPv4 Historian

Scenario 1 – IPv6-only Host accesses IPv4 Historian (Stateful NAT64 Translation)
As shown above, the IPv4 Historian has the IP address of 192.0.2.26. The DNS64 server maintains the A record for this IP address and maps it to the AAAA record of ***64.ff9b:c00:21a***. Below is the breakdown of the communication:

- Acquire Destination IP
  - IPv6 host queries "historian.example.com"
  - DNS server finds the A record of ***192.0.2.26***.
  - DNS64 server translate A record to AAAA record and returns ***64:ff9b::c00:21a***.
- Contact Destination
  - Host sends first packet to ***64:ff9b::c00:21a***.
  - Packet is routed to the default IPv6 gateway, which is the NAT64 gateway.
  - NAT64 gateway re-encapsulates payload in IPv4 packet with the destination IP address of ***192.0.2.26*** and the source IP address of its own (***192.0.2.1***).
  - NAT64 gateway sends the IPv4 packet to historian.
- Handle Return Traffic
  - Historian accepts request and sends back IPv4 response to NAT64 gateway.
  - NAT64 gateway re-encapsulates payload in IPv6 packet with the destination IP address of the host and source IP address of ***64:ff9b::c00:21a***.

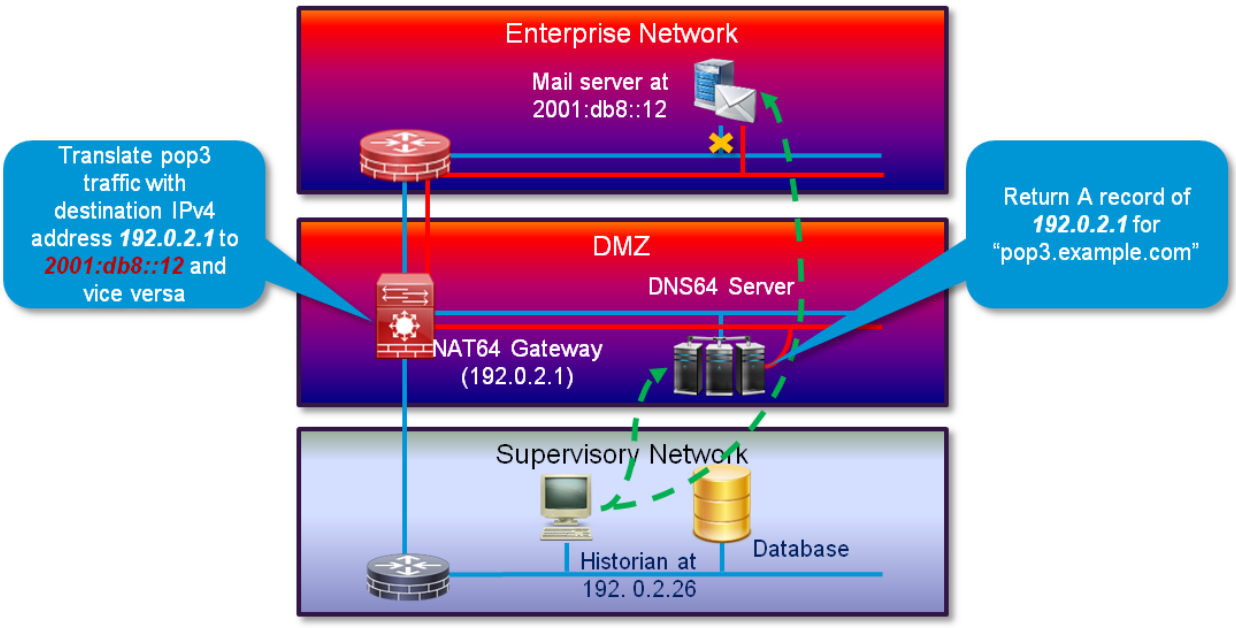

Figure 3.2: IPv4-only Host access IPv6 Server

Scenario 2 – IPv4-only Host accesses IPv6 Server (Stateless NAT64 Translation)
As shown above, the IPv6 Mail Server has the IP address of ***2001:db8:12***. The NAT64 gateway creates and maintains a NAT64 address mapping for the Mail Server (***192.0.2.1→2001:db8:12***), and the DNS64 server creates the corresponding A record of "pop3.example.com" for the mapped IPv4 address. Below is the breakdown of the communication:

- Configure static address mapping on NAT64 gateway and DNS64 server
  - Map ***192.0.2.1***/110 to ***2001:db8:12***/110
  - Create A record of ***192.0.2.1*** for "pop3.example.com"
- Acquire Destination IP
  - IPv4 host queries "pop3.example.com"
  - DNS server returns the A record of ***192.0.2.1***.
- Contact Destination
  - Host sends first POP3 TCP packet to ***192.0.2.1***.
  - Packet is routed to the default IPv4 gateway, which is the NAT64 gateway.
  - NAT64 gateway re-encapsulates payload in IPv6 packet with the destination IP address of ***2001:db8:12*** and the source IP address of ***64.ff9b::c00:21a***.
  - NAT64 gateway sends the IPv6 packet to the mail server.
- Handle Return Traffic
  - Mail server accepts request and sends back IPv6 response to NAT64 gateway.
  - NAT64 gateway re-encapsulates payload in IPv4 packet with the destination IP address of the host (***192.0.2.26***) and source IP address of ***192.0.2.1***.

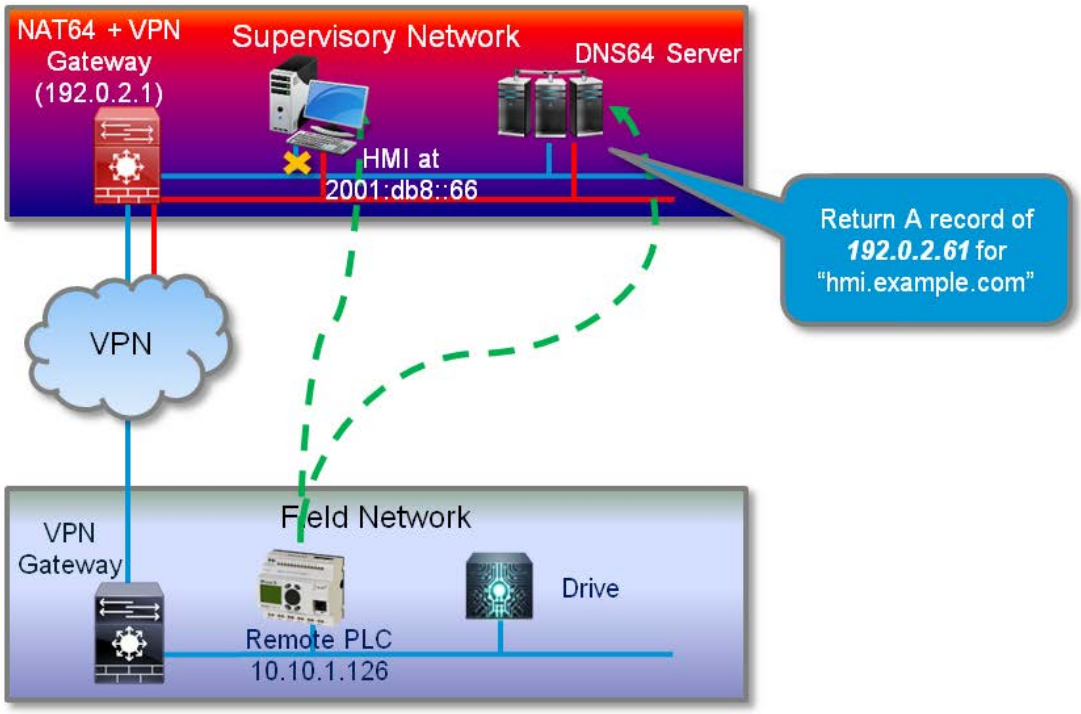

Figure 3.3: Remote IPv4-only Host access IPv6 Server via VPN

Scenario 3 – Remote IPv4-only Host accesses IPv6 Server via VPN
This scenario is very similar to the second scenario where the IPv4-only host is on the same physical network as the IPv6 server. The difference is that in this case the traffic needs to be encrypted and tunneled through IPsec tunnel. Below is the breakdown of the communication:

- Remote VPN gateway establishes IPsec VPN tunnel with the local VPN gateway on the Supervisory Network, which also happens to be the NAT64 gateway.
- Configure static address mapping on NAT64 gateway and DNS64 server
  - Map ***192.0.2.61*** to ***2001:db8:66***
  - Create A record of ***192.0.2.61*** for "hmi.example.com"
- Acquire Destination IP
  - Remote IPv4 host queries "hmi.example.com". Request is sent over VPN to the DNS64 server on Supervisor Network.
  - DNS server returns the A record of ***192.0.2.61***.
- Contact Destination
  - Remote host sends TCP packet to ***192.0.2.61***.
  - Packet is tunneled to the VPN + NAT64 gateway.
  - NAT64 gateway decrypts packet and re-encapsulates payload in IPv6 packet with the destination IP address of ***2001:db8:66*** and the source IP address of ***64.ff9b::c00:21a***.
  - NAT64 gateway sends the IPv6 packet to the IPv6 HMI.
- Handle Return Traffic
  - HMI accepts request and sends back IPv6 response to NAT64 gateway.
  - NAT64 gateway re-encapsulates payload in IPv4 packet with the destination IP address of the host (***10.10.1.126***) and source IP address of ***192.0.2.61***. Packet is encrypted and sent to the remote VPN gateway.

The Industrial Control Network has come a long way to what it is today, and is deeply rooted in IPv4 and related networking technologies. To ensure successful migration to IPv6, we need to understand these dependencies and challenges:

- Dependencies on the IPv4 Infrastructure

    - Using IPv4 address as device and service identifiers
    - Always assume a four-byte IP address
    - Rely on the broadcast and multicast functions of IPv4
- IPv4 Address Embedded in Control Protocols
    - For example, the ListIdentity response message in EtherNet/IP protocol contains the IP address of the responding device
    - Ring protocols even though layer2 (e.g. DLR) may embed IP address as part of the payload.
    - The EtherNet/IP TCP/IP Interface Object and other configuration objects may contain IP address definitions
- IPv4-based Management Tools and Utilities
    - Existing network and automation management tools present IPv4-based management interfaces

To address these issues, it will take collective effort from the customers, networking manufacturers, machine builders, and software vendors. Below is the list of guidelines for stage 2 planning and deployment:

- Industrial protocols may need to be updated to work with IPv6.
- Management tools must be upgraded to support IPv6.
- Control systems that are engaged in the **same control loop** (e.g. using the same protocols) should be **upgraded together** to avoid any issues with NAT64.
- I/O devices on the **same Ring topology** must be **upgraded together**.

Stage 3 – Migrate Dual-Stack Network to IPv6-Only Network

The objectives of stage 3 are:

- Support smooth transition to full IPv6-only network
- Allow different network segments to be migrated independently

Stage 3 begins when all the endpoints on the network supports IPv6. At which point, network operators can simplify the network management and configuration by removing or disabling NAT64 related functionality and infrastructure. To ensure a smooth transition, a phased approach is also recommended:

- Selectively **disable NAT64 functionality** on NAT64 gateways and DNS64 servers and **test drive IPv6-only network.**
- **Create small IPv6 pockets** by replacing NAT64 gateways with regular IPv6 gateways. **Merge small IPv6 pockets into bigger IPv6-only subnets**.
- Remove all IPv4 and NAT64 infrastructure assets (e.g. gateway, DNS server, DHCP server … etc.).

Common NAT64 Issues and Solutions

NAT64 is a mature technique for IPv6 transitioning. Through the course of deploying NAT64 in different network scenario, we've come to understand some common issues related to NAT64 and the corresponding solutions:

- Smaller IP Path MTU (**M**aximum **T**ransmission **U**nit) on IPv6 Network
- Optional UDP Checksum on IPv4 Network
- IPv4 and IPv6 Fragments
- Unnecessary Translation for Dual-Stack Endpoints
- Different Endpoints on the Same Layer-2 Network
- Special Protocols
- Broken IPSec VPN

Smaller IP Path MTU on IPv6 Network

Since both IPv4 and IPv6 network are running on the same physical Ethernet network, the Ethernet MTU value for each endpoint remains the same (e.g. 1500 with no LLC, SNAP, or VLAN). But due to the larger IP address and different header format, IPv6 header is bigger than IPv4, which leads to a smaller IP Path MTU (i.e. maximum size of IP datagram excluding IP header) on the IPv6 network. Mismatched MTU on IPv4 and IPv6 network may cause additional fragmentation. To solve this problem, below are the recommendations on how to solve this problem:

- The IP Path MTU on IPv6 network is smaller due to larger IPv6 header.

- To avoid fragmenting translated payload on IPv6 network, **the IP Path MTU on IPv4 network should be set at a lower value (e.g. 1460)**.
- If not possible to configure the MTU on IPv4 endpoints, **NAT64 gateways must participate in Path MTU Discovery** and notify IPv4 endpoints of the new MTU value.
- NAT64 gateways **must be able to fragment big IPv4 datagrams** that exceed MTU on IPv6 network.

Optional UDP Checksum on IPv4 Network

UDP checksum is mandatory on IPv6 network, but is optional in IPv4. This means that the NAT64 gateway may receive an IPv4 UDP datagram with UDP checksum set to zero. In this case, the NAT gateway must recalculate the UDP checksum by following the guidelines below:

- NAT64 gateway must **recalculate TCP and UDP checksums**, which is typically done by calculating the difference between the two different pseudo-headers.
- NAT64 gateway must **recalculate UDP checksum using the entire payload data** when it receives an UDP datagram with zero checksum.
- If the zero-checksum UDP datagram is also a fragment, NAT64 gateway must **reassemble the UDP datagram before re-calculating the checksum**. If the fragments arrive out of order, the UDP datagram may end up being dropped.

IPv4 and IPv6 Fragments

This is a more general scenario than the first two. Even if you could re-configure MTU on IPv4 network, NAT64 gateways still need to be able to handle IP fragments received on either network:

- IPv6 fragments are handled end-to-end. IPv6 router shall never fragment an IPv6 datagram.
- If IP fragments arrive in order, NAT64 gateway will **translate fragments as they arrive**. States will be maintained on the gateway in order to translate following fragments.
- If IP fragments arrive out of order, NAT64 gateway **queues fragments until the first fragment arrives**, at which time translation will be done for queued fragments as well.

Unnecessary Translation for Dual-Stack Endpoints

The proposed IPv6 migration strategy does not require endpoints to support dual-stack. The dual-stack network infrastructure tries to hide the involvement of NAT64 and DNS64 from the endpoints. If an endpoint does support dual-stack, it needs to know which stack to use in order to avoid unnecessary translation:

- To talk to another endpoint, the dual-stack host needs to know whether to use IPv4 or IPv6. If wrong stack is chosen, unnecessary translation may occur.
- To solve this problem, **separate DNS servers** (independent from DNS64 server) should be configured for IPv4 and IPv6 network independently. The host should always **use the stack on which a valid DNS record was returned**. The host should send the query on IPv6 network first. If a valid AAAA record is returned on IPv6 network, the host must not send the DNS query on IPv4 network. If for some reason the host receives valid DNS records on both networks (e.g. timeout the first query too quickly), it's up to the host to decide which stack to use for talking to the other endpoint.

Different Endpoints on the Same Layer-2 Network

Industrial Control Network typically employs large complex layer-2 network with mixed linear and ring topologies. Since NAT64 is a layer-3 function, supporting communication between IPv4 and IPv6 endpoints on such layer-2 network requires careful planning:

- Both IPv4 and IPv6 endpoints may be on the same layer-2 network (e.g. same VLAN and connected by same switches).
- NAT64 gateway must be able to **perform translation on any physical or logical interface**, and must handle the scenario where **source and destination endpoints are connected to the same physical or logical port**.
- IPv4 and IPv6 endpoints should be **grouped into separate layer-2 networks** whenever is possible.

Special Protocols

There're certain application protocols that embed IP addresses in the protocol payload. For example, FTP PORT command and PASV response contain IP address and port information. Below is the guideline on how to address this problem:

- Some protocols embed IP addresses in the protocol payload, e.g. FTP, RTSP, PPTP, SIP, EtherNet/IP… etc.
- You shall only install NAT64 gateways that support **ALG** (**Application Level Gateway**) functionality for these protocols.
- If you can't or don't use a NAT64 gateway supporting such features, to avoid service disruption, client and server endpoints using these protocols should be migrated to IPv6 at the same time.

IP Multicast

Certain protocols use IP multicast for communications, e.g. EtherNet/IP. In order to forward such multicast traffic across the NAT64 boundary, you must install NAT64 gateways that are able to translate between IPv4 and IPv6 multicast packets. Below is how it can be achieved:

- NAT64 gateway should maintain the one-to-one mappings between IPv4 and IPv6 multicast addresses. The mapping entries can be manually configured or hard-coded.
- NAT64 gateway should support two translation modes for each mapping:
    - Always-on mode – In this mode, whenever a multicast message is received on either network, the gateway shall translate the packet by replacing the IP header, the multicast destination IP address, and the source IP address, and then forward it to the interfaces on the other network. The gateway must allow the forwarding on each network interface to be turned on and off independently.
    - Learning mode – In this mode, the gateway must function as a MLD (for IPv6 interface) or IGMP (for IPv4) proxy on each interface. For example, when the gateway detects a new host has joined a multicast IGMP group on IPv4 network, it must compose a corresponding MLD message and send it to the IPv6 network. The gateway should maintain a list of member host information for each multicast group, and only translate and forward a multicast message when there's at least one active member on the other side of the network.
- Because the multicast IP address is embedded in the CIP Explicit message setting up the multicast CIP I/O exchange, NAT64 gateway must implement the CIG ALG that is capable of translating IP address contained in such messages.
- When SSM (Source-Specific Multicast) was specified by an endpoint, NAT64 gateway should translate the source IP address of the multicast if it knows the mapping; otherwise, NAT64 gateway will have to drop the source in the translated messages. Note that when the IGMP version on the IPv4 network is not 3, all SSM information from MLD will be lost in translation.
- In a CIP environment, if the installed NAT64 gateway cannot meet the requirements of translating CIP multicast packets, you should disable multicast on EtherNet/IP CIP endpoints.

Broken IPSec VPN

IPSec (IP Security) VPN has different setups, for example, Point-to-Point, Site-to-Site, and Remote Access. It's a great way to protect traffic on Industrial Control Network. Special considerations are needed when IPSec tunnels go through NAT64 gateway:

- **IPSec AH and IPSec ESP in Tunnel Mode break** because they protect the outer IP header.
- **A single IPSec ESP session in Transport Mode generally works** across NAT64 gateway. Multiple sessions between an endpoint and the NAT64 gateway may not work because PAT cannot be performed on IPSec ESP. You should install NAT gateways that support IPSec pass-through if you need to pass IPSec VPN traffic through the gateways.

- **NAT-T negotiation or IPSec-over-UDP** (i.e. NAT-T without negotiation) should be always enabled for all IPSec endpoints.

Conclusion

The key to IPv6 migration in Industrial Control Network is to build an intelligent network that supports communication between different endpoints, including IPv4-only, IPv6-only, and dual-stack hosts. By deploying NAT64 and DNS64 support across the network and enforcing consistent configuration and operation on all network systems, we can hide most of the complexity from the endpoints and ensure a smooth transition to IPv6-based network infrastructure.

References:

[1] C. Partridge, F. Kastenholz, Technical Criteria for Choosing IP The Next Generation (IPng), RFC 1726, Internet Engineering Task Force (IETF), December 1994

[2] E. Nordmark, Stateless IP/ICMP Translation Algorithm (SIIT), RFC 2765, Internet Engineering Task Force (IETF), February 2000

[3] A. Durand, P. Fasano, I. Guardini, D. Lento, IPv6 Tunnel Broker, RFC 3053, Internet Engineering Task Force (IETF), January 2001

[4] C. Bao, C. Huitema, M. Bagnulo, M. Boucadair, X. Li, IPv6 Addressing of IPv4/IPv6 Translators, RFC 6052, Internet Engineering Task Force (IETF), October 2010

[5] F. Baker, X. Li, C. Bao, K. Yin, Framework for IPv4/IPv6 Translation, RFC 6144, Internet Engineering Task Force (IETF), April 2011

[6] X. Li, C. Bao, F. Baker, IP/ICMP Translation Algorithm, RFC 6145, Internet Engineering Task Force (IETF), April 2000

[7] M. Bagnulo, P. Matthews, and I. van Beijnum, Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, RFC 6146, Internet Engineering Task Force (IETF), April 2011

[8] R. Despres, IPv6 Rapid Deployment on IPv4 Infrastructures (6rd), RFC 5569, Internet Engineering Task Force (IETF), Januray 2010

[9] W. Townsley, O. Troan, IPv6 Rapid Deployment on IPv4 Infrastructure (6rd) -- Protocol Specification, RFC 5969, Internet Engineering Task Force (IETF), August 2010

[10] M. Bagnulo, A. Sullivan, Shinkuro, P. Matthews, I. van Beijnum, DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers, RFC 6147, Internet Engineering Task Force (IETF), August 2010

[11] Mark Townsley (September 24, 2012). "Mapping Address + Port". Cisco. Retrieved 2012-09-25.

[12] N. Kushalnagar, G. Montenegro, and C. Schumacher, IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals, RFC 4919, Internet Engineering Task Force RFC 4919, August 2007.

[13] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. P. Vasseur, and R. Alexander, RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks, RFC 6550, Internet Engineering Task Force RFC 6550, March 2012.

[14] T. Narten, E. Nordmark, W. Simpson, and H. Soliman, Neighbor Discovery for IP version 6 (IPv6), RFC 4861, September 2007.

****