CT. Insurance Corporation

IPv6 Strategy and Roadmap

US and European Data Center DMZs

Limited Rights Statement

Applied Methodologies, Inc. (AMI) makes no representations or warranties of any kind, either expressed or implied, with the respect to the contents of this report, including but not limited to typographical errors and technical completeness. Applied Methodologies, Inc. reserves the right to revise this publication and to make changes in the content hereof without obligation of Applied Methodologies, Inc.to notify any person of such revision or changes.

LIMITED RIGHTS NOTICE

This report is submitted with limited rights. This report may be reproduced and used by CT Insurance Corp. with the express limitation that they will not, without written permission of AMI Consulting Solutions, use the contents for purposes except for evaluation of IPv6 Readiness nor disclose the contents outside of the CT Insurance Corp. organization. When permitted; CT Insurance Corp. may disclose this data outside of CT Insurance Corp. provided that the CT Insurance Corp. makes such disclosure subject to the prohibition against further use and disclosure. CT Insurance Corp. organizations are granted unrestricted rights to reproduce and distribute the information contained herein for the sole use of their personnel in the evaluation of future CT Insurance Corp. IPv6 Migration strategies.

This Notice shall be marked on any reproduction of this data, in whole or in part.

(End of notice)

Modification History

AUTHOR	Version	DATE	COMMENTS
	0.001		Document Creation for CT. Ins. Corp.
	1		Edited format – Updated Content
	1.8		Added/updated content – edited format, cleaned up reference flow and sequence, updated toc, initial spell check, work on costing, executive summary translation points, cleaned up additional flow references.
			Need to add remaining costing summary and or put detail costing in separate doc.
	2.5		Quick scan corrected some typos.
	2.5		Corrected typos prior CT. Ins. Corp. Review
	3.0		Added content based on discussion and review with client
Team Acknowledgement			Team Review
Deliver to CT. Ins. Corp. for review			CT. Ins. Corp. IPv6 Strategy Roadmap-DRAFT-v020-Discovery Only.DOCX
			Updates to Costs and Compliance III
			Executive Summary and Recommendations
			Email, all cross references, updates, MS Word grammar
			Peer Review Exec Summary, edit Solution Alternatives
Draft ES, Recommend to Customer			CT. Ins. Corp. IPv6 Strategy Roadmap-DRAFT-v055-Partial Doc.docx
			Review and Edit, Peer review, re-write
			Review and Edit, minor document properties
AMI Team	V1.0		CT. Ins. Corp. IPv6 Strategy and Roadmap_v1.0.DOCX

Table of Contents

Executive Summary

Background

IPv4 address space across the global Internet is reaching a point of near exhaustion due to the accelerating growth in the number of IP enabled devices. In order to address this global exhaustion of usable IPv4 addresses, a new standard known as IPv6 has been developed and is being deployed across the globe which will provide orders of magnitude more usable addresses. The IPv4 address exhaustion process creates an environment where new areas of the globe, which consists of countries and regions that cannot obtain IPv4 spaces, will come online as IPv6 only, cell phones, tablet and public smart devices will too only speak IPv6. These "newcomers", especially regions and countries rising in economic stature provide a new base of potential customers who have their products and services represented via IPv6 enabled web sites.

The global IPv4 condition poses several risks in relation to gaps of internet visibility that CT. Ins. Corp. needs to address over the next 12-18 months. This assessment outlines details of the risks and remediation required. A summary is provided below:

Without support for IPv6 CT. Ins. Corp. risks losing global visibility to its products and services from IPv6 only regions. The initial risk profile is:

- Members, Business to Consumer(B2C) and new customer applications
- Remote access for employees and partners to conduct business
- New or existing Inbound B2B partners
- Internet Email
- Mobility

IPv6 Approach

There are multiple approaches available to deploy IPv6 across an enterprise environment including Dual-Stack, translation and tunneling. Dual-Stack requires the device, system, host, or application to run both IPv4 and IPv6 analogous to being bi lingual. Translation allows for elements of the environment to remain on IPv4 and rely on a bi lingual device to do the translation. The translation or tactical approach will utilize a translation platform that sits at the Internet points of presence, to translate inbound IPv6 traffic to IPv4 for application access, and translation of outbound IPv4 traffic to IPv6 for access to the global Internet. This approach immediately addresses the risks and gaps mentioned above.

- Translation provides a gradual migration to IPv6 by providing seamless Internet experience to greenfield IPv6-only users, accessing IPv4 Internet services.
- Existing content providers and content enablers can provide services transparently to IPv6 Internet users by using translation technology, with little or no change in the existing network infrastructure, thus maintaining IPv4 business continuity.
- Specific protocols such as File Transfer Protocol (FTP) and Session Initiation Protocol (SIP) that embed IP address information within the payload require application-layer gateway (ALG) support for translation.

The strategic approach entails continuing to Dual-Stack the remaining devices at the internet points of presence to enhance IPv6 performance, visibility and scalability resulting in the eventual removal of the translation platform. This approach ensures that IPv6 risks and gaps are never present in the future.

CT. Ins. Corp. IPv6 Readiness

AMIconducted a readiness analysis to determine the impact associated with deploying IPv6 to address the gaps outlined above. The analysis concluded that: ***CT. Ins. Corp. is well positioned to address these gaps from a tactical standpoint with its current IT assets without a large OPEX/CAPEX investment.*** This statement is based on the following key findings:

- Most of the critical assets within the higher risk areas of the network (external and DMZ) currently support IPV6, have software upgrades to provide support, or are currently slated for lifecycle refresh within the next 12-14 months.
- CT. Ins. Corp. maintains a lab environment which already mimics the production environment and will be able to assist with detailed IPv6 planning and preparation.
- The network design facilitates the introduction of functionality such as IPv6 translation to ease the introduction of IPV6 into the CT. Ins. Corp. environment.

Based on CT. Ins. Corp.'s state of readiness, the organization will be able to pursue a phased approach to IPV6 support while allowing for a methodical strategic migration to full compliance over time and as required by the business.

CT. Ins. Corp. IPv6 Roadmap

AMI recommends that CT. Ins. Corp. allocate funding for projects to address the risks outlined earlier and start with a hybrid approach incorporating both dual-stack and translation functionality.

The key recommendations specific to deploying support for IPv6 at CT. Ins. Corp. involve both strategic and tactical capabilities. As described above, the ultimate goal is to deploy dual-stack functionality as the strategic direction and goal. However, based on the timeframes to enable this capability across all areas of the CT. Ins. Corp. external and DMZ environments, the translation capability should be deployed. Translation will provide CT. Ins. Corp. with significant flexibility in supporting IPv6, and can be used to address potential technical gaps that may be discovered during the detailed planning and testing of IPv6 within the environment.

Additional benefits of implementing this approach are:

- Translation platform refreshes IPv4 SLB/DNS/CSS-GSS platforms as a project byproduct
- Provides more immediate visibility for IPv6 only customers to CT. Ins. Corp. globally by addressing risk and gaps outlined earlier
- Low impact, cost and effort to execute, eases CT. Ins. Corp. into IPv6
- Provides time and controlled pace for CT. Ins. Corp. to plan, stage and execute its strategic IPv6 goals

The long term, strategic recommendation is to continue to plan, test and phase in the progression of Dual-Stack protocols either from the edge of the network inward towards the enterprise or from the enterprise out towards the tactical translation point, remove translation and complete the Dual-Stack progression as part of the lifecycle and planned upgrades.

Roadmap Phases

These recommendations are grouped into six phases specifically targeting the US locations reviewed as part of this assessment – Middletown and Windsor. The assessment did glean information regarding CT. Ins. Corp.'s London data center and its recommendations are being provided in a separate section, based on the fact that London can be addressed as a pilot first site, or as a follow on location.

The roadmap provides recommendations for addressing the risks associated with the IPv4 address exhaustion and a structured approach to deploying IPv6. The following table shows the activities by phases, that need to be completed by CT. Ins. Corp.. Although some of these activities can happen in parallel, many of the preliminary planning, design and testing activities will need to occur prior to any configuration or deployment.

- Phase 1 – Detailed Design and Planning - IPv6 Solutions
- Phase 2 – Test and Verify IPv6 Solutions
- Phase 3 – Deploy the Dual-Stack Target Architecture in the External Zone
- Phase 4 – Enable IPv6 Services (DNS/LDAP/NTP/FTP/NMS/VPN) Capability
- Phase 5 - Deploy IPv6 Translation Capability in the External Zone
- Phase 6 – Enable Dual-Stack on Remaining Secure DMZ Infrastructure/Servers

Rough order of magnitude time and cost estimates were assigned to the steps in each phase and task dependencies were identified to project the preliminary roadmap timetable. Additional work will be required to refine these estimates. The following figures represent phases and estimated scheduling to complete in parallel to the lifecycle activities already slated for the coming months.

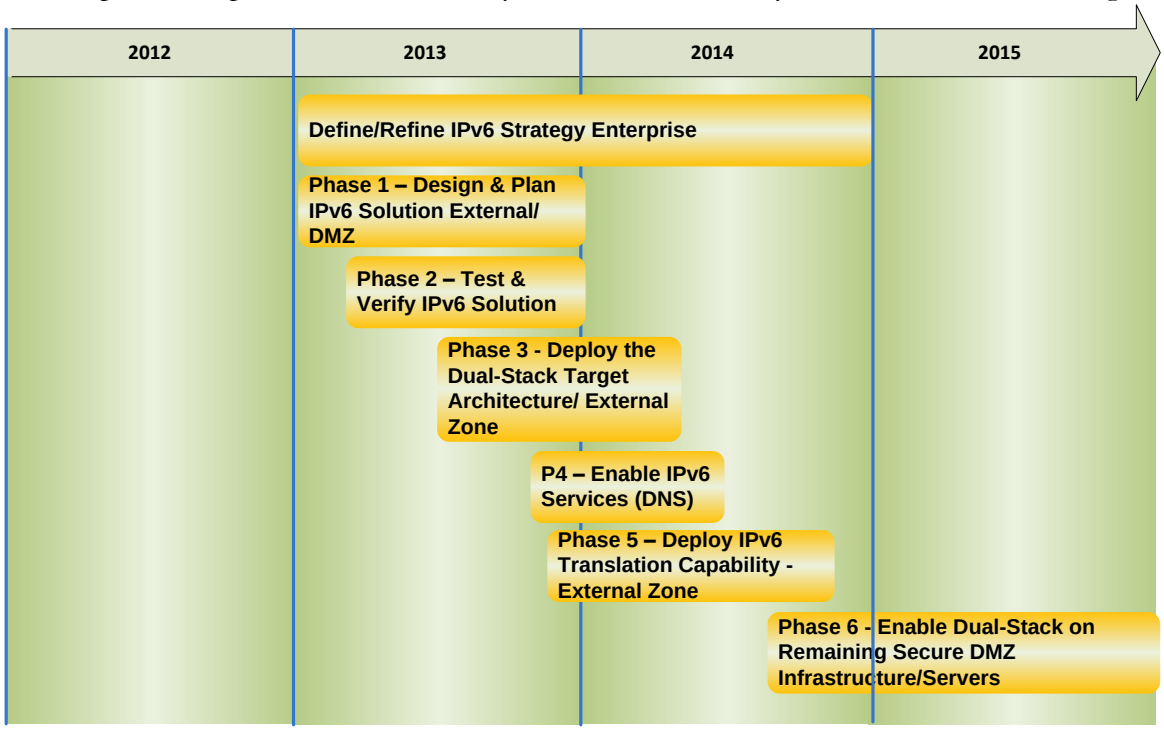

Figure 1 – Ct. Ins Corp. IPv6 Remediation Timeline

IPv6 Architecture

	IPv6 Transiton Technology	Cost	Risk	Complexity	Strategy
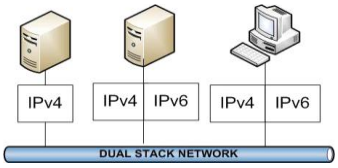	***Dual-Stack*** **End to End IPv6 Communication**	**High**	**Low**	**High**	√ **Strategic Recommended Approach**
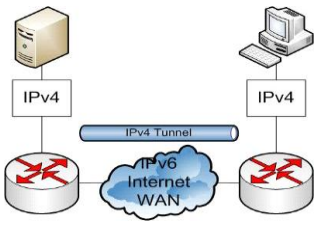	***Tunneling*** **IPv4 over IPv6 Internet** **IPv6 over IPv4 WAN**	**Medium**	**Medium**	**High**	**Transition Enabler**
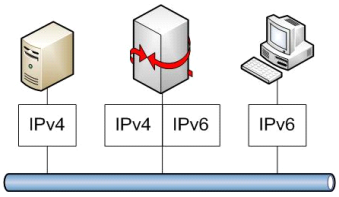	***Translation*** **IPv4 to IPv6** **IPv6 to IPv4**	**Low**	**Medium**	**Medium**	√ **Tactical Recommended Approach**

IPV6 Readiness Cost Summary

As described above, CT. Ins. Corp. has done an outstanding job of maintaining current technology in its Internet DMZ network segments; and a number of refresh plans have already been budgeted which will address many of the aging devices within the higher risk areas of the network. The table below provides a summary the rough order of magnitude (ROM) costs for CT. Ins. Corp. to apply the recommended approach of IPv6 deployment across the target areas of the network reviewed during this engagement. The costs are broken down as follows:

Table 1 - High Level Pricing

Phases 1 and 2 Includes Loc1/Loc2/Europe			
Phases 3 - 5 Implementation costs for Data Centers			
London			
Phase 1	Research costs Estimated 4 months	$ -	Covers tactical and Strategic research activities Strategic research may require additional time
Phase 2	Lab Testing, Verification and Staging Estimated- 3 months Includes building initial sandbox and formal lab	$ -	Covers tactical testing activities - time required for application behavior analysis Uniqueness for DC London staging and testing if possible Time can be used for future phases if ahead of schedule
Phase 3	Deploy Dual Stack in External Zone Estimated - 1 month Includes component upgrades/carrier work	$ -	Per DC
Phase 4	Enable IPv6 Services(DNS etc.) Estimated - 1 month	$ -	Per DC
Phase 5	Translation Zone 1 Deployment Estimated - 1 month	$ -	Per DC
Phase 5	Translation Zone 2 Deployment Estimated - 1 month	$ -	Per DC
Europe	Translation Deployment Estimated - 2 month	$ -	Includes Planning and Lab testing Covers Managed provider uniqueness issues
Training	Allocated under Aetna Budget	$ -	Allocated over 2 years for IPv6 May require additional for F5 if cannot fit within budget
New equipment Production			
BIG IP F5 LTM 6900 Translation appliance	X2 price	$ -	Estimated average of $k F5-BIG-LTM-6900-8G-R BIG-IP SWITCH: LOCAL TRAFFIC MANAGER 6900 8GB ROHS
BIG IP F5 GTM	X2 price	$ -	Estimated average of $k F5-BIG-GTM-1600-4G-R BIG-IP SWITCH: GLOBAL TRAFFIC MANAGER 1600 4GB ROHS $ Estimate on platform and modules could be higher or lower after detailed planning, platform selection and vendor pricing discounts are reviewed
New Equipment Lab			
BIG IP F5 LTM Translation appliance	X2 price for Lab	$ -	Estimated average of $k based on 1600 LTM platform without GTM module

Summary - Same if either Translation Zone 1 OR 2 used		Summary - Two Translations Zones deployed		Notes:
Training	$ -	Training	$ -	
Research and testing (Inc. Lab F5s)	$ -	Research and testing (Inc. Lab F5s)	$ -	
LOC1 complete deployment	$ -	LOC 1 complete deployment	$ -	additionl Phase 5 labor and LTM 6900 pair
LOC 2 complete deployment	$ -	LOC 2 complete deployment	$ -	additionl Phase 5 labor and LTM 6900 pair
Europe complete deployment	$ -	Eurpope complete deployment	$ -	
F5s or QA	$ -	F5s or QA	$ -	
Total without training it was already budgeted	$ -	Total without training it was already budgeted	$ -	

Conclusion

The inevitable transition to IPv6 is not a trivial matter, particularly for a large organization like CT. Ins. Corp., with a global data communications footprint. ***In summary, CT. Ins. Corp. is in a very good position specific to IPv6 readiness based on the relatively low net new costs required to enable IPv6 support utilizing Dual-Stack and translation capabilities.*** This is primarily ***due to the existing CT. Ins. Corp. refresh/lifecycle process*** which will addresses many of the key infrastructure elements in the external and DMZ segments of the network. Many critical existing assets can be refreshed to support IPv6 and introducing the tactical translation devices/capability will allow CT. Ins. Corp. to enable IPv6 support during the completion of the near term refresh plans. Once this capability is enabled, CT. Ins. Corp. will have significantly reduced overall business risk associated with an inability to support IPv6 endpoints across the globe. Additional details and recommendations are outlined in the expanded IPV6 assessment deliverable also compiled as part of this engagement.

Introduction

As stated in the Executive Summary, the need to deploy IPv6 has been driven by the exhaustion of the IPv4 address space across the global Internet. It is expected that access to CT. Ins. Corp. public applications, over the Internet, will be the first users who will have IPv6 only capabilities. This will require CT. Ins. Corp. to have mechanisms in place to allow these users access to their Internet facing servers. Internet facing servers and their applications will be a high priority in the deployment of IPv6.

This document focuses on CT. Ins. Corp.'s strategy for deploying IPv6 across the external segments of the network most likely to be initially impacted by the global IPv6 deployment across the Internet. The document includes specific information that was discovered and analyzed to provide inputs to future budgeting, architectural direction, timelines, and priorities required to implement IPv6.

Each section contains current state findings, analysis of various risks associated with IPv6, approaches, and remediation activities and costs required to integrate IPv6. The scope of these recommendations focuses on the elements of the CT. Ins. Corp. external Internet/DMZ segments in the Middletown, Windsor, and London locations.

Included in the document are inventories, analysis, and recommendations in the following areas:

1. **Current State Discovery** –This section provides a clear comprehension of CT. Ins. Corp. network components and approaches used today including inventory, sites, routing and IP addressing. Use cases are created to determine the level of risk and impacts to the CT. Ins. Corp. Network
2. **IPv6 Asset Readiness and Impact Report** – This section evaluates IPv6 compliance of the current equipment and provides specific IPv6 costs in two approaches. The first approach assumes an immediate, full IPv6 migration. The second describes a phased approach that incorporates CT. Ins. Corp.'s current equipment refresh cycles. This approach allows non-compliant equipment to be refreshed in a typical upgrade path.
3. **IPv6 Architecture and Recommendations** – This section reviews the recommend IPv6 architecture including a summary of technical details, a revisit of the case studies risks and remediation provided from the recommendations presented.
4. **Migration Strategy and Timelines** – This section covers the phases and summary steps within each phase to deploy the recommend IPv6 architecture and roadmap covered in item three above.
5. **IPv6 Planning Resources** - This section provides information to aid CT. Ins. Corp. engineers in the planning, testing and of the recommendations and migration options covered in items three and four. Addressing schema options and steps required to building a formal testing and verification lab are presented. Training and educational information is also included in this section.

IPv4 Exhaustion

Without the impending IPv4 exhaustion, most organizations would not even consider transitioning to IPv6. However, this exhaustion will produce "IPv6 only" users that will not have a way to access traditional IPv4 only applications. This will present risks to organizations that rely on the Internet to do business. Every organization will have different risks associated with IPv6, depending on how they utilize the Internet for business, but suffice to say that eventually all organizations will have some level of risk.

A Regional Internet Registry (RIR) is a not-for-profit organization that oversees Internet Protocol (IP) address space (IPv4 and IPv6) and the Autonomous System (AS) numbers within a specific geographical region. There are five regional RIRs across the globe: ARIN, RIPE, APNIC, LACNIC, and AFRINIC. Together, they are known as the Number Resource Organization (NRO).

The five Regional Internet Registries (RIRs) and the regions they serve are as follows:

1. APNIC – Asia, Australia, New Zealand, and neighboring countries
2. RIPENCC – Europe, the Middle East, and Central Asia
3. ARIN – United States, Canada, several parts of the Caribbean region, and Antarctica
4. LACNIC – Latin America and parts of the Caribbean region
5. AFRINIC – Africa

Table describes IPv4 Global Address Exhaust Projections; predicted exhaustion timelines for the five RIRs responsible for allocating IPv4 addresses are as follows:

Table 2 - IPv4 Global Address Exhaust Projections

www.potaroo.net/tools/ipv4/index.html

	27-Jun	15-Oct	
Updated 10/20	% /8's Left	% /8's Left	Exhaust
APNIC	0.9298	0.9077	**15-Apr-11**
RIPE	1.8217	0.9819	**14-Sept-12**
ARIN	3.5248	3.1564	**27-Aug-13**
LACNIC	3.4243	3.1732	15-June-15
AFRINIC	4.1924	4.0468	22-Sept-19

It is important to note that the RIR's allocate addresses to ISP's, as well as Enterprise customers within the regions that they operate. This means that once the RIR's exhaust their IPv4 address pools, there will still be a period in which the ISP's will have IPv4 addresses. This will result in a staggered exhaustion of IPv4 addresses across the globe. However, as can be seen in the exhaust projections above, the exhaust phases will occur first in Asia Pacific, followed by Europe, then the Americas in the 2012-2013 timeframes. This puts a major emphasis on addressing inbound and outbound connectivity to the Internet moving into 2013.

In contrast to the depletion activity, usage of IPv6 has been increasing in the US.

North America is driving more IPv6 traffic than any other region of the world, according to Akamai. Here are the peak traffic volumes reported by Akamai for each region:

Region Peak IPv6 Traffic Volume Date:

1. North America 92,891 hits/sec 9/11/2012
2. Europe 48,488 hits/sec 9/11/2012
3. Asia 14,540 hits/sec 7/8/2012
4. South America 549 hits/sec 8/24/2012
5. Africa 152 hits/sec 8/23/2012

Risks and Impacts Associated with IPv6

Risk Profiles

The objective of this section is to identify where high, medium, and low priority risks exist for IPv6 in CT. Ins. Corp.'s External/DMZ network. CT. Ins. Corp. has a multitude of services delivered over the Internet that will be impacted by the IPv4 address exhaust. Once identified, the information will be used to create a roadmap to manage and minimize the impact that IPv6 only users have to access resources when doing business with CT. Ins. Corp..

Table describes how an organization can classify IPv6 areas of risk related to the exhaust of the IPv4 address pool as follows:

Table 3 - IPv6 Risk Profile and Classification

Risk Profile	Classification	Description
High	External Segments	All systems, hosts, and network devices used to provide access to internal applications and services. This is usually the Data Centers (DC) and any sites with Internet points of presence.
Medium	Outbound Internet	All systems, hosts, and network devices used to provide outbound access to the global Internet. Internal users who access outbound Internet services are on LAN networks in campus locations or remote site LAN's/WLANs connected over a WAN.
Low	Private Networks	The networks and Infrastructure with no direct Internet services, (do not receive or transmit any Internet traffic)

AMI and CT. Ins. Corp. have identified a series of IPv4 use cases to identify which portions of the IT environment are at risk. Priorities are then assigned on those risks for use in remediation timelines. Firstly, Internet based services are identified as high-risk within in the CT. Ins. Corp. environment for the following reasons:

1. Multiple external user communities utilize connectivity via the Internet for business purposes;
    a. Plan Members, New customers, general information,
    b. Benefits services, Personal Health records, International
    c. B2B vendors
    d. Employee web based services
    e. Employee and remote site connectivity

2. CT. Ins. Corp. heavily utilizes the Internet for external connectivity today and will **to continue the use of the Internet** as an enabling technology.

In addition, specific business drivers within CT. Ins. Corp. will increase the risks related to migration to IPv6, which include:

- Internet customers will begin to speak IPv6 only and require CT. Ins. Corp. to support IPv6 to Internet facing applications (e.g. CT. Ins. Corp..com)
- Global and domestic future customers, business partners, service providers etc., will require enterprises to communicate via IPv6
- Patient Protection and Affordable Care Act (HCR) change demand and federal mandates for all agencies to adopt and run IPv6 by 2014, which may require partners to follow suite
- Domestic employees and members need the ability to stay flexible in their choices for Internet Service Providers (ISP)s will require the ability connect via IPv6
- Domestic and Global utilization of the Internet as a transport technology for remote site/employee connectivity will require IPv6
- Mobile applications on smart phones and tablets that speak only IPv6 from their carrier and carrier does not translate

In time, other risks associated with internal user connectivity to the Internet will become evident. These risks are related to the lack of IPv6 related content or the inability to connect to IPv6 content on the Internet. These are deemed medium priority risks.

Finally, the areas of the network that do not require communication either inbound or outbound to the global Internet have virtually zero risk related to the global IPv4 address exhaust.

These risks will have a significant impact on CT. Ins. Corp.'s ability to deliver IT services and to grow service capabilities across a variety of geographies and service delivery needs. AMI has provided mitigating solutions in this assessment to make the IPv6 risk impact acceptable to CT. Ins. Corp..

Current State Discovery

Methodology

This IPv6 Strategy and Roadmap begins with a basic "State of CT. Ins. Corp.'s Union" in relation to the network architecture that comprises inbound internet access and public facing services that CT. Ins. Corp. currently provides via IPv4. This entails their External and DMZ sections only. This strategy and roadmap does not cover CT. Ins. Corp.'s internal network architecture. The current state provides information on the existing IPv4 architecture and the topology used to provide IPv4 connectivity for applications and services in the CT. Ins. Corp. IT environment. This information helps to provide overall context and identify general IPv4 functional requirements. The information gathered during the discovery phase was used during the analysis phase to determine IPv6 readiness.

The current state discovery information is categorized into following areas:

A. **External/DMZ Network Architecture** – describes the currently deployed network that was discovered during our onsite engagement (September/October 2012). This includes how the two data centers are connected and their relation to the various DMZs.
B. **Network Management Systems** - As CT. Ins. Corp. begins to deploy IPv6, network management of these devices will be required. This section briefly describes what is required to manage new IPv6 devices in the CT. Ins. Corp. network.
C. **DMZ Security Infrastructure** – Security related devices identified during the discovery
D. **Servers Discovery** – provides inventory of hardware and operating systems of all servers with a specific priority to Internet facing servers
E. **Applications/External Facing** – describes the public facing applications and services CT. Ins. Corp. provides to the public, members, and employees via the Internet
F. **IPv4 Addressing Plan** – describes current IPv4 addressing used by CT. Ins. Corp. today as it pertains to the External/DMZ sections.
G. **Carriers and Third Party Vendors** – Provides a comparison of the IPv6 capabilities available to CT. Ins. Corp. from its internet providers.
H. **IPv4 Use Cases –** Discusses some inbound Use Cases – User and application Internet inbound access (Users attempting to access CT. Ins. Corp. resources)

External/DMZ Network Architecture

As a component of the IPv6 assessment, AMI reviewed the existing deployed architecture with a focus on the externally facing segments of the environment. This refers to the external Internet connections themselves, as well as the DMZ segments that provide direct termination or services to users/devices on the Internet. Below is a summary of the environment that was discovered and used during the review.

Referring to the figure 1. CT. Ins. Corp. relies on two major points of connections (POC) to the internet. The POCs reside in each of their data centers for redundancy.

CT. Ins. Corp.'s internet facing network architecture comprises of several sections or "DMZ zones". These zones are separated by firewalls. The zone's architecture consists of the following attributes:

- Common Hardware
- Common Network Designs
- Common Software Revisions
- Redundant & Modular Network Designs

As per the Figure 1 below there is an Unsecured DMZ (external) which for this document is simply called the External Zone. This zone compromises of Cisco Catalyst 6506 switches with line cards for the appliances and firewalls (Cisco ASA) to connect into. VLANs and L3 separation are employed. There are redundant 6506 switches connected by a L2 port channel. Each 6500 has an SPA and SIP modules for and OC48 link to their inter-provider. Automatic Path Selection is utilized in this configuration.

There is only ONE OC48 connection per DC POC. The backup 6500 has an OC-48 connection that is in standby mode managed by APS which is still part of the same OC-48 uplink to the ISP, not a separate connection. Failover is achieved by a combination of OC-48's APS and HSRP, DBAR links and default routes.

There is also a pair of gigabit Ethernet links between the data centers. The links denoted as DBAR on the diagram are active traffic paths between the data centers. These links are not really considered for Disaster Backup and Recovery use. If one ISP failed, all traffic would flow through the other data center. The links are used for normal daily traffic. The links are L3 Ethernet ports and EIGRP establishes neighbor relationship through them.

EBGP peering is established from each DC POC's OC48 link to the carrier. A full routing table is provided. There is no IBGP peering between the DC using the DBAR links as discussed earlier.

EIGRP is CT. Ins. Corp.'s core IGP and is present in the External and Secure DMZ zones.

Some of the features employed on the External Zone Cisco Catalyst 6500

APS for OC-48	HSRP
BGP	EIGRP
Port Channel	Sup Redundancy SSO/NSF
SPAN	Standard Access-Lists
Extended Access-Lists	Control Plan Policing
Router Admin	HTTP server
TACACS+	Netflow version 5
NTP client	SNMP

From the External Zone most of the traffic flows into what CT. Ins. Corp. calls the Secure DMZ. This DMZ zone is separated from the External zone by a pair of redundant HA pair Cisco ASAs. The Secure DMZ zone comprises similar hardware and software as in the External zone, redundantly with the addition of Content Switch Module for load balancing. L2/L3 separation is also employed with L2 and L3 6506 based switches. Some of the features employed on the Secure DMZ Zone Cisco Catalyst 6500

CSM Load Balancer Module	VLANs
HSRP	EIGRP
Port Channel	Sup Redundancy SSO/NSF
SPAN	Standard Access-Lists
Extended Access-Lists	Router Admin
HTTP server	TACACS+
NTP client	SNMP

A full detailed list of HW/SW that comprises the zones is listed in section ***IPv6 Asset Readiness and Costs.***

The same platform and HW/SW components are replicated in both DCs.

Within the Secure DMZ zone are other logical zones Such as the Server Load Balancing Zone or SLB zone. This zone is where a majority of the web, application and appliance servers reside.

From the Secure DMZ and SLB zones traffic destined to the internal DC resources flows through a MacAfee firewall into the internal side of CT. Ins. Corp.'s enterprise network called the Surf zone. The MacAfee firewalls are the final demarcation point from the External and Secure DMZ zones into the Enterprise.

Other DMZ zones such as the PerfQA, VPN, Cloud, PCI and the Internet B2B zones comprised of Cisco ASA and MacAfee firewalls separating the zones either from the External zone to its own set of hardware in the DC(outside of the Secure DMZ zone equipment) or connect from the External Zone through the FWs into the Surf Zone.

It is these DMZ zones, demarcation points and the equipment within them where the initial IPv6 deployment planning and technical consideration will be heavily concentrated.

Figure 1 provides a summary view of CT. Ins. Corp.'s DC POC and External Zone orientation.

Figure 2 provides a summary of the various DMZ zones as discussed earlier.

Figure 3 - CT. Ins. Corp. Internet Access Overview

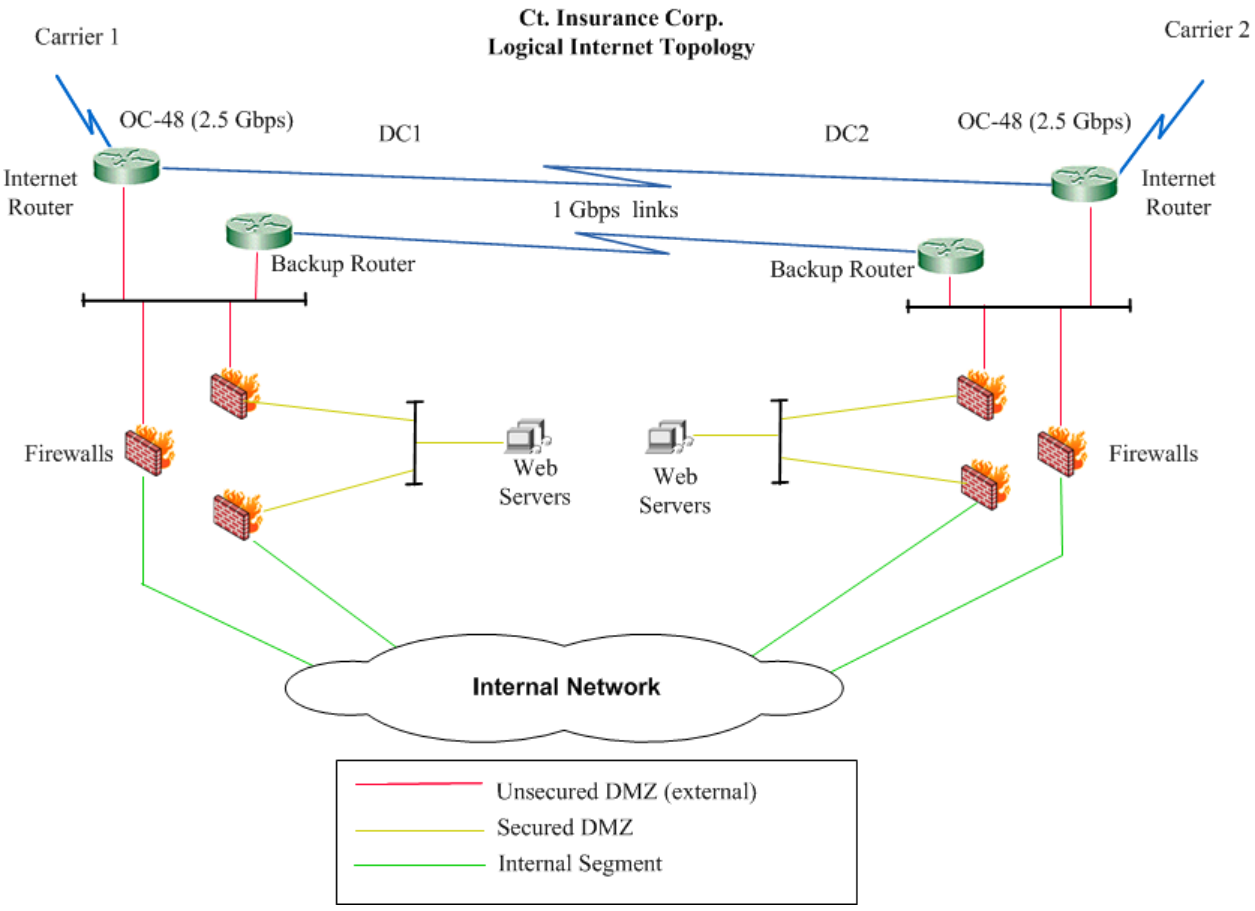

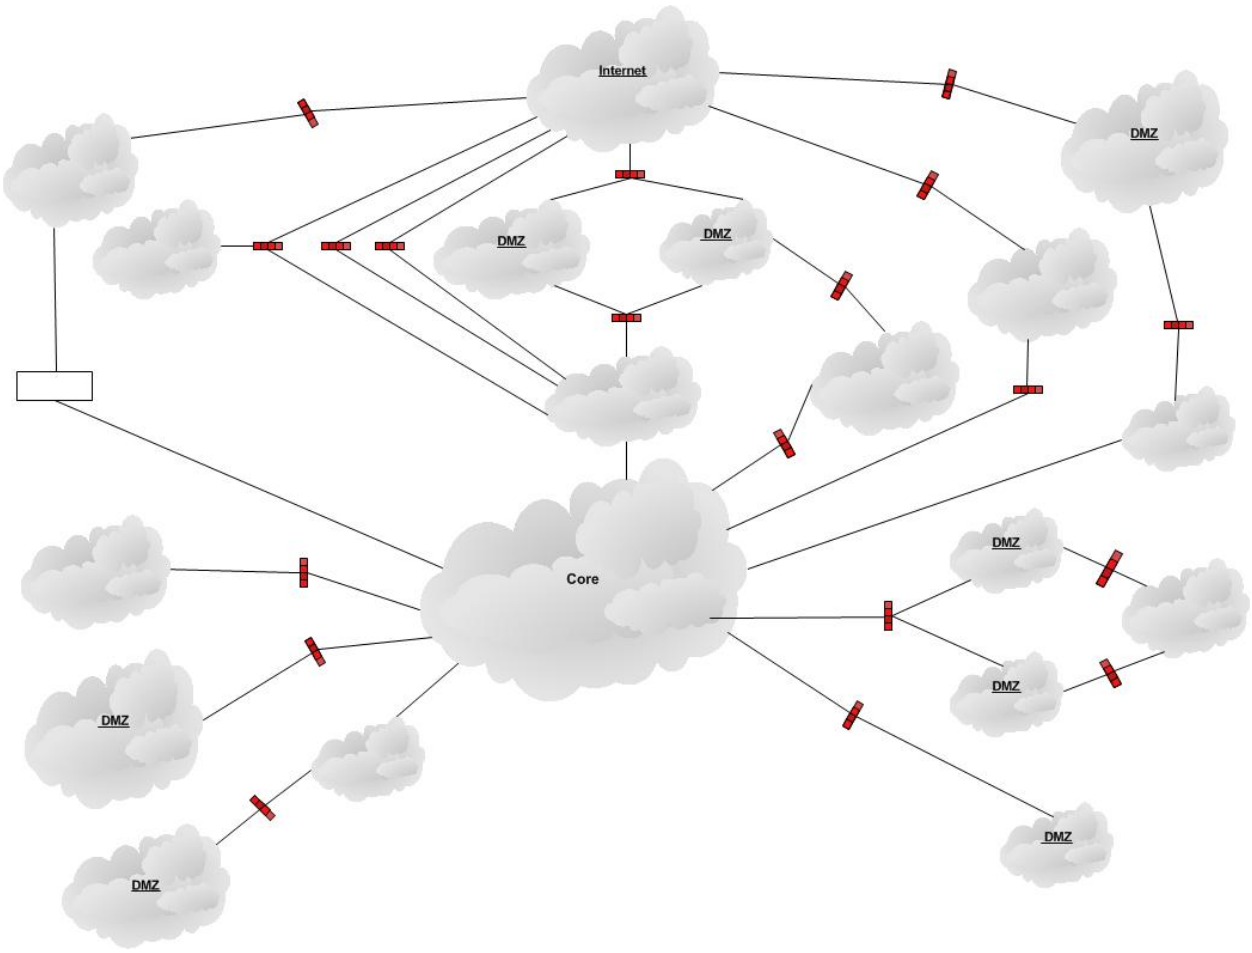

Figure 4 – CT. Ins. Corp. DMZ Map

DNS

CT. Ins. Corp. has roughly 1500 domains names registered and it host 70 locally. The use of NSI Advance DNS services is employed for hosting and consolidation of domestic sties and resolution of common URL misspellings. DNS is supported in the Secure DMZ and External zones with the use of BIND based DNS on the Secure DMZ zone's McAfee. Cisco GSS appliances are also employed in the External zone for resolution load balancing. NSI sends further queries through the External zone towards the McAfee firewall running BIND in the Secure DMZ zone. The McAfee firewall that is the demarcation point for the Secure DMZ and Surf/Internal zone is the SOA, there are 198 zones many primary and are published every 15 minutes.

Figure 3. illustrates CT. Ins. Corp.'s logical view of its DNS architecture as it pertains to the Internet. The shaded area is where the initial focus will be for IPv6.

Figure 5 – CT. Ins. Corp. DNS

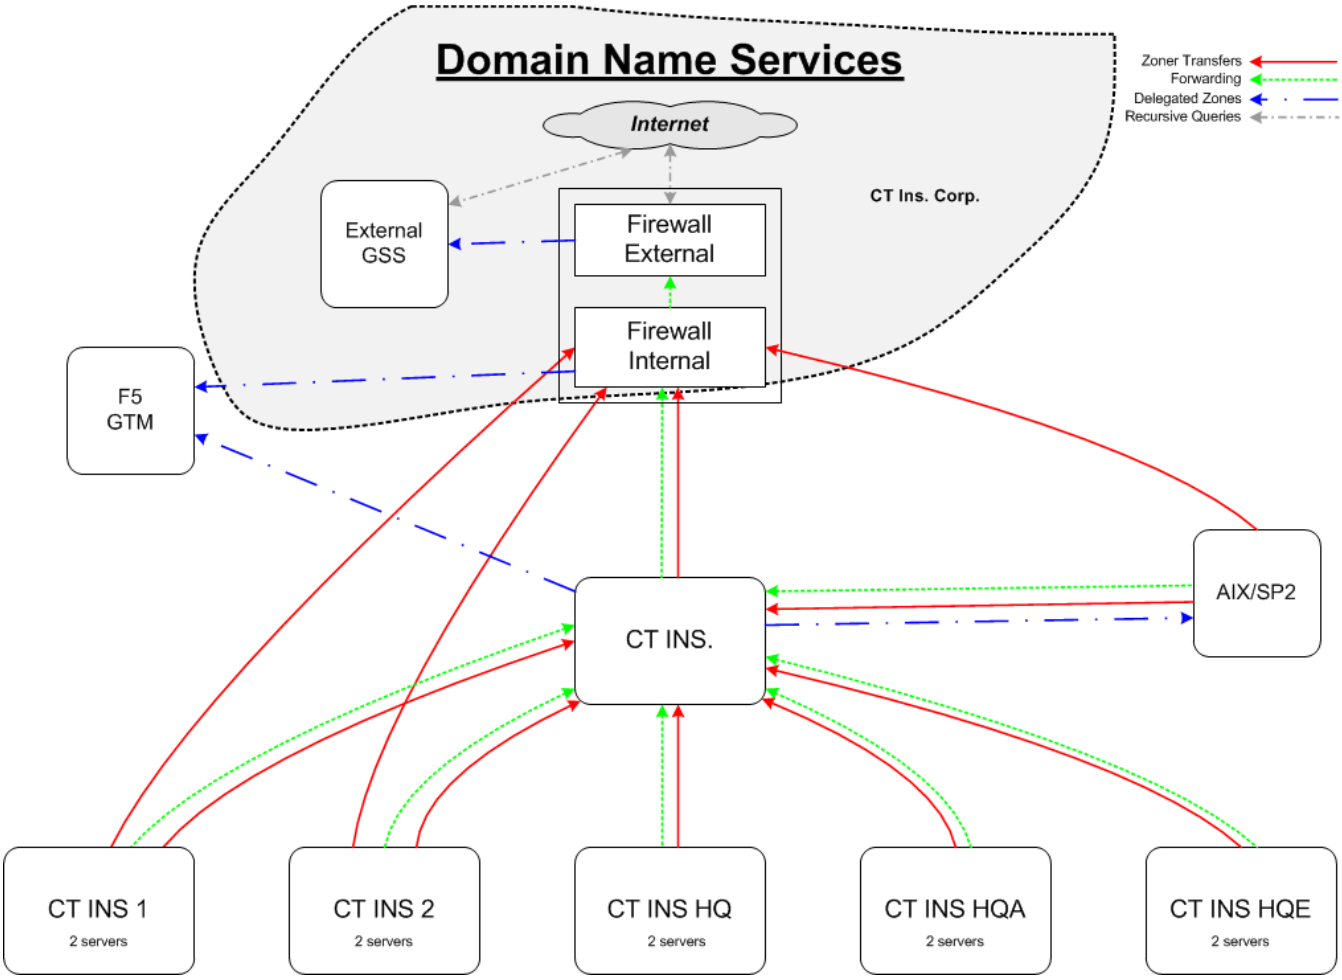

Firewalls

CT. Ins. Corp. utilizes two main firewall products Cisco ASA and McAfee. The details of their HW/SW for IPv6 compliance is covered in section ***IPv6 Asset Readiness and Costs.*** The purpose of the Cisco Firewalls are for DMZ zone separation, VPN and B2B IPsec tunnels. The McAfee products are used for DMZ zone separation, DNS, outbound proxy related functions.

External/DMZ IP Version 4 Address Plan

IPv4 Addressing

CT. Ins. Corp.'s current summarized IPv4 allocation outline is as follows:

There are 2 Class Cs used for external connectivity.

1 Class C was originally allocated by CT. Ins. Corp.
1 Class C was acquired via a merger
1 Class C is split from / into two / for DC 1 and DC 2 use
Only 9 or 10 of the class Cs are actually routable. Rest are used for VIPs in DMZ

There are 2 Class B allocated for the Internal network.

1 Class B is heavily used
1 Class B is lightly to hardly used and was acquired via a merger
The remainder of the addressing and Internal network utilizes private IP addresses.
1 heavily used /16
1 not used often /16 from a merger and provides plenty of room for v4 growth
Remainder of Internal utilizes RFC 1918 ranges.

There is currently no IPAM system deployed or used for the External/DMZ components. A master spreadsheet is the database for IPv4 moves, adds, and changes. There is currently plenty of room in terms of IPv4 routable addresses with no growth issues to be present in the future.

Table 14 below outlines the address ranges relative to the External and DMZ sections.

Table 4 - CT. Ins. Corp. External and DMZ Addressing Summary

DC Location	DMZ section	Address Range	Notes:
DC 1/DC 2	External	XXXX.XXXX.XXXX.XXX	OC3 links / Loopbacks
DC 1	External	XXXX.XXXX.XXXX.XXX	
DC 1	Secure public	XXXX.XXXX.XXXX.XXX	
DC 1	Secure private	XXXX.XXXX.XXXX.XXX	SLB
DC 1	Secure private	XXXX.XXXX.XXXX.XXX	SLB
DC 2	External	XXXX.XXXX.XXXX.XXX	
DC 2	Secure public	XXXX.XXXX.XXXX.XXX	
DC 2	Secure private	XXXX.XXXX.XXXX.XXX	SLB
DC 2	Secure private	XXXX.XXXX.XXXX.XXX	SLB
DC 1	Surf	XXXX.XXXX.XXXX.XXX	Internal
DC 2	Surf	XXXX.XXXX.XXXX.XXX	Internal
DC 2	Surf	XXXX.XXXX.XXXX.XXX	AVE
DC 1	Surf	XXXX.XXXX.XXXX.XXX	AVE
DC 1	VPN DMZ	XXXX.XXXX.XXXX.XXX	VPN
DC 1	VPN DMZ	XXXX.XXXX.XXXX.XXX	VPN
DC 2	VPN DMZ	XXXX.XXXX.XXXX.XXX	VPN
DC 2	VPN DMZ	XXXX.XXXX.XXXX.XXX	VPN
DC 1	DMZ	XXXX.XXXX.XXXX.XXX	VIP
DC 2	DMZ	XXXX.XXXX.XXXX.XXX	VIP
DC 1	PerfQA-DMZ	XXXX.XXXX.XXXX.XXX	QA Real
DC 1	PerfQA-DMZ	XXXX.XXXX.XXXX.XXX	QA NAT
DC 1	PerfQA-DMZ	XXXX.XXXX.XXXX.XXX	QA VIP
DC 1	PerfQA-DMZ	XXXX.XXXX.XXXX.XXX	QA Secure
DC 1	PerfQA-DMZ	XXXX.XXXX.XXXX.XXX	QA SLB 1 or old 3
DC 1	PerfQA-DMZ	XXXX.XXXX.XXXX.XXX	QA SLB 2 or old 4
DC 1/DC 2	Internal Enterprise	XXXX.XXXX.XXXX.XXX	Internal
DC 1/DC 2	Internal Enterprise	XXXX.XXXX.XXXX.XXX	Internal
DC 1/DC 2	Internal Enterprise	XXXX.XXXX.XXXX.XXX	Internal
DC 1/DC 2	Internal Enterprise	XXXX.XXXX.XXXX.XXX	Internal
DC 1/DC 2	Internal Enterprise	XXXX.XXXX.XXXX.XXX	Internal
DC 1/DC 2	Internal Enterprise	XXXX.XXXX.XXXX.XXX	Internal
Europe	DMZ	XXXX.XXXX.XXXX.XXX	Managed Svs. provided
Europe	DMZ	XXXX.XXXX.XXXX.XXX	SLB
Europe	DMZ	XXXX.XXXX.XXXX.XXX	VIP
Europe	DMZ	XXXX.XXXX.XXXX.XXX	Legacy
		XXXX.XXXX.XXXX.XXX	

Network Management Systems (NMS)

As CT. INS. CORP. begins to deploy IPv6, network management of these devices will be required. Out-of-the-box applications will be capable of capturing and reporting of IPv6 devices and capturing IPv6 MIBs. These MIBs could be IPv6 only, or protocol-version independent (PVI). (SNMPv3) protocol support will be required. It is unclear at this time which tools will allow the ability to capture and report mixed IPv4 and IPv6 environments. There are also in-house or custom management applications that may or may not be able to glean IPv6 related information.

There are currently some gaps in the industry where NMS is concerned resulting in some products may not have full support of IPv6 from a protocol lingual level or device querying level. The progression of NMS in the enterprise is outlined below.

NMS platform speaks IPv4 only at L3 and can glean IPv4 but cannot glean IPv6 statistics from devices.
NMS platform speaks IPv4 only at L3 and can glean IPv4 and IPv6 statistics from devices.
NMS Platform speaks Dual-Stack IPv4/6 natively at L3 and can glean IPv4/IPv6 statistics from devices.
NMS platform just speaks IPv6 only at L3 and can glean IPv6 and Legacy IPv4 statistics from devices.

A general review of the network management environment was conducted just to determine the products used in the DMZ or from the internal side and their reach into the DMZ and if they support IPv6 from the perspective of gleaning IPv6 statistics. CT. Ins. Corp. relies on Tivoli and HP Opsware to monitor and change devices in the DMZ. These platforms reside in the Internal section of the network and are thus out of scope for IPv6 integration. However, they are important for they reach into the DMZ and External zones to glean operational data. The Tivoli system has a rich history of supporting IPv6.

A full list of supported modules can be found here:

http://www-01.ibm.com/software/info/ipv6/compliance.jsp

It is recommended that CT. Ins. Corp. perform a more detailed review of the NMS capabilities once the IPv6 deployment strategy is finalized to ensure they have the latest supported modules.

Opsware is comprehensive configuration management solution supporting both network and system. Opsware is a configuration management solution from HP. The key components of Opsware include NAS for network and SAS for server. Opsware can manage and maintain network configuration file history and has change management and scheduling capability at the same time.

For the initial deployment it is recommended that CT. Ins. Corp. continue to monitor and manage IPv4 only and IPv4/v6 devices via IPv4. As IPv6 products mature for NMS CT. Ins. Corp. can slowly deploy a v6 based NMS overlay over its v6 assets when the products are ready.

Security Services in the DMZ

CT. Ins. Corp. utilizes the following scanning and IDS tools in the DMZ:

The Ixia Anue Net Tool OptimizerTM 5236 network monitoring switch is ideally suited to provide high-performance 10G visibility for network monitoring tools in all parts of a fiber network. This includes the access, distribution and backbone layers as well as the server farm and data center environments. The 5236 aggregates, replicates and filters network traffic using the industry's most advanced, easy to use, drag-and-drop Control Panel. The 5236 supports IPv6 and since it's role is to aggregate and replicate traffic at L2 and above it will recognize IPv6 traffic. However, any platforms where the data is sent for further review must be able to understand IPv6 Ethertype and higher layer headers.

QualysGuard is an on demand vulnerability management, policy compliance and asset management solution that enables organizations to assess and manage business risk. QualysGuard automates the network security auditing process across the enterprise both inside and outside the firewall, and across distributed networking environments. QualysGuard provides network discovery and analysis, asset prioritization, centralized reporting, and remediation workflow and verification. Executive-level reports allow security professionals to demonstrate effective security practices and verify compliance with data protection laws and regulations. QualysGuard's SaaS technology is far more accurate, cost effective, and easier to deploy than software-based alternatives. QualysGuard is IPv6 compliant and ready for CT. Ins. Corp.'s deployment use.

CT. Ins. Corp. also uses Corero Network Security (Top Layer) E, EC and ES IPS and DDoS appliances reside in the enterprise but reach into the DMZ for monitoring. Since they reside in the internal section of the network they are out of scope for initial IPv6 integration. However, there upgrade information has been captured the main issue is how well the appliances detect IPv6 related issues.

Server Discovery

External Facing Servers

This section defines the hardware types and operating systems of the servers currently deployed within CT. Ins. Corp.'s Secure DMZ/SLB DMZ zones. Only the external facing servers which are those accessed by external Internet users were discovered during this engagement. This distinction is important because Internet facing servers will have a higher priority for IPv6 implementation to remediate IPv6 access risk as opposed to internal servers as the Dual-Stack environment progresses through the Secure DMZ and SLB DMZ to the internal sections of CT. Ins. Corp.'s network.

The following tables outlines the server's currently deployed based on queries sent to respected points of contact who either support "own" the device, or who have a related interest and provided the data. In addition, a report that gleaned what was actually on the "wire" in the DMZ was used and the servers were reconciled and checked for IPv6 compliance.

Note: this is not an exact but close accounting of the servers due to time constraints and support resource changes within CT. Ins. Corp.. There are some servers/appliances scheduled to be deprecated/removed that were noted on the wire. There are 112 servers.

Table 6 – IPv6 Compliant Servers

Hostname	Associated HW Model	Device Group	Apps suported	Location at site External/DMZ/Surf/Internal	DC Location	Operating System and Revison	IPv6 Compliant
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2008 R2	Y
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2008 R2	Y
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2008 R2	Y
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2008 R2	Y
Serv x	HP DL385G2	Backup	IBM / Tivoli Storage Mana	DMZ	Loc 1	RHEL 5.6	Y
Serv x	F5 BIG IP 1600	Load Balancer	BIG IP	DMZ	Loc 1	BSD Unix 10.2	Y
Serv x		Office Comm		DMZ	Loc 1	Server 2008 R2	Y
Serv x		Office Comm		DMZ	Loc 1	Server 2008 R2	Y
Serv x		Patches	MS WSUS 3.2	DMZ	Loc 1	2008 SP1	Y
Serv x	HP DL380G6	PCI	PCI	DMZ (PCI)	Loc 1	RHEL5.8	Y
Serv x	HP DL380G6	PCI	PCI	DMZ (PCI)	Loc 1	RHEL5.8	Y
Serv x		Scanner	Qualysguard 2.6	DMZ	Loc 1		Y
Serv x		Webtrend	Webtrend	DMZ	Loc 1	Server 2008 R2	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 1	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 1	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 1	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 1	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 1	RHEL5.8	Y
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2008 R2	Y
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2008 R2	Y
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2008 R2	Y
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2008 R2	Y
Serv x	HP DL385G2	Backup	IBM / Tivoli Storage Mana	DMZ	Loc 2	RHEL 5.6	Y
Serv x	F5 BIG IP 1600	Load Balancer	BIG IP	DMZ	Loc 2	BSD Unix 10.2	Y
Serv x		Office Comm		DMZ	Loc 2	Server 2008 R2	Y
Serv x		Office Comm		DMZ	Loc 2	Server 2008 R2	Y
Serv x		Patches	MS WSUS 3.2	DMZ	Loc 2	2008 SP1	Y
Serv x	HP DL380G6	PCI	PCI	DMZ (PCI)	Loc 2	RHEL5.8	Y
Serv x	HP DL380G6	PCI	PCI	DMZ (PCI)	Loc 2	RHEL5.8	Y
Serv x		Scanning	Qualysguard 2.6	DMZ	Loc 2		Y
Serv x		Webtrend	Webtrend	DMZ	Loc 2	Server 2008 R2	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 2	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 2	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 2	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 2	RHEL5.8	Y
Serv x	HP 380 G6	Webeng		DMZ	Loc 2	RHEL5.8	Y
Serv x		App	VONTU	DMZ	QA	Server 2008 R2	Y
Serv x		App	Password Auto Repository	DMZ	QA	Server 2008 R2	Y
Serv x		Office Comm		DMZ	QA	Server 2008 R2	Y
Serv x		Office Comm		DMZ	QA	Server 2008 R2	Y
Serv x		Office Comm		DMZ	QA	Server 2008 R2	Y
Serv x		Office Comm		DMZ	QA	Server 2008 R2	Y
Serv x		Office Comm		DMZ	QA	Server 2008 R2	Y
Serv x		Office Comm		DMZ	QA	Server 2008 R2	Y
Serv x	DL385G1	PCI		DMZ	QA	RHEL 5.8	Y
Serv x	DL385G5	PCI		DMZ	QA	RHEL 5.8	Y
Serv x		WebEng		DMZ	QA	RHEL 5.7	Y
Serv x	DL380 G6	WebEng	Shared Web Environment	DMZ	QA	RHEL 5.8	Y
Serv x	DL380 G6	WebEng	Shared Web Environment	DMZ	QA	RHEL 5.8	Y
Serv x	DL380 G6	WebEng	Shared Web Environment	DMZ	QA	RHEL 5.8	Y
Serv x	DL380 G6	WebEng	Shared Web Environment	DMZ	QA	RHEL 5.8	Y
Serv x	DL380G7	WebEng	WebSphere	DMZ	QA	RHEL 5.8	Y
Serv x	DL380G7	WebEng	WebSphere	DMZ	QA	RHEL 5.8	Y
Serv x	DL380G7	WebEng	WebEng - I H S	China	Europe	RHEL 5.8	Y
Serv x	DL380G7	WebEng	WebEng - I H S	China	Europe	RHEL 5.8	Y

Table 7 – Non IPv6 Compliant Servers

Hostname	Associated HW Model	Device Group	Apps suported	Location at site External/DMZ/Surf/Internal	DC Location	Operating System and Revison	IPv6 Compliant
Serv x		App		DMZ	Loc 1	Server 2003	N
Serv x		Webeng	IIS 6	DMZ	Loc 1	Server 2003	N
Serv x		Webeng	EPSM, IIS 6 claims	DMZ	Loc 1	Server 2003	N
Serv x		Webeng	EPSM, IIS 6	DMZ	Loc 1	Server 2003	N
Serv x		Webeng	EPSM, IIS 6	DMZ	Loc 1	Server 2003	N
Serv x		App		DMZ	Loc 1	Server 2003	N
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2003	N
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2003	N
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2003	N
Serv x		Auth	SiteMinder/IIS	DMZ	Loc 1	Server 2003	N
Serv x		Auth	Ping Federate	DMZ	Loc 1	Server 2003	N
Serv x	Hp dl380 g7	Backup	Symantec/Netbackup/6.5	DMZ	Loc 1	Server 2003	N
Serv x		ESM	SAV Management	DMZ	Loc 1	Server 2003	N
Serv x		Ipass		DMZ	Loc 1	Server 2003	N
Serv x		Webeng	PAL (Salseweb), AGB , IIS	DMZ	Loc 1	Server 2003	N
Serv x		Webtrend	Webtrend	DMZ	Loc 1	Server 2003	N
Serv x		App	AHIA	DMZ	Loc 2	Server 2003	N
Serv x		Webeng	IIS 6	DMZ	Loc 2	Server 2003	N
Serv x		Auth		DMZ	Loc 2	?	N
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2003	N
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2003	N
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2003	N
Serv x		Auth	Siteminder/IIS	DMZ	Loc 2	Server 2003	N
Serv x		Auth	Ping Federate	DMZ	Loc 2	Server 2003	N
Serv x	Hp dl380 g7	Backup	Symantec/Netbackup/6.5	DMZ	Loc 2	Server 2003	N
Serv x		ESM	Anti-Virus	DMZ	Loc 2	Server 2003	N
Serv x		Ipass		DMZ	Loc 2	Server 2003	N
Serv x				DMZ	Loc 2		N
Serv x		Webeng	PAL (Salseweb), AGB , IIS	DMZ	Loc 2	Server 2003	N
Serv x		Webtrend	Webtrend	DMZ	Loc 2	Server 2003	N
Serv x		App	AGB/IIS	DMZ	QA	Server 2003	N
Serv x		App	AGB/IIS	DMZ	QA	Server 2003	N
Serv x		App	AGB/IIS	DMZ	QA	Server 2003	N
Serv x		Webeng	ClaimXTen/IIS	DMZ	QA	Server 2003	N
Serv x		Webeng	ClaimXTen/IIS	DMZ	QA	Server 2003	N
Serv x		Webeng	ClaimXTen/IIS	DMZ	QA	Server 2003	N
Serv x		App	EPSM\IIS	DMZ	QA	Server 2003	N
Serv x		App	ClaimXTen/IIS	DMZ	QA	Server 2003	N
Serv x		App	EPSM\IIS	DMZ	QA	Server 2003	N
Serv x		Auth	SiteMinder	DMZ	QA	Server 2003	N
Serv x		Auth	SiteMinder	DMZ	QA	Server 2003	N
Serv x		Auth	SiteMinder	DMZ	QA	Server 2003	N
Serv x		Auth	Ping Federate	DMZ	QA	Server 2003	N
Serv x		Auth		DMZ	QA	Unknown	N
Serv x		Backup		DMZ	QA	Server 2003	N
Serv x		DNS		DMZ	QA		N
Serv x		Email	Exchange	DMZ	QA	Server 2003	N
Serv x		Email	Exchange	DMZ	QA	Unknown	N
Serv x		WebEng	Shared Web Environment	DMZ	QA	Server 2003	N
Serv x		Webtrends		DMZ	QA	Server 2003	N
Serv x		ESM			Europe	Unknown	N
Serv x		WebEng			Europe	Server 2003	N
Serv x		WebEng			Europe	Server 2003	N
Serv x		WebEng			Europe	Server 2003	N
Serv x		WebEng			Europe	Server 2003	N
Serv x		SQL server			Europe	Unknown	N

Applications

External Facing

CT. Ins. Corp. relies on a set critical web based applications that support users which comprise of the general internet public, potential customers, policy members, business partners and healthcare professionals. This section is focused on the customer, partner and employee services applications outside of the employee support applications such as email, and VPN. These applications provide the necessary information for these users to conduct business with CT. Ins. Corp.. AMI has conducted a cursory review of the applications and supporting platforms based on interviews identified a few of the critical applications necessary at a minimum to be present in an IPv6 enabled environment.

- The main web server platforms are Linux based Apache 2.2 and IBM's HTTP Server
- The Windows server based platform utilize IIS 6 and 7
- There are some 1 for 1 server, appliances and shared web resources deployed – refer to section ***Server Discovery External Facing Servers*** for a breakdown.
- User data is not stored on these servers but in the internal enterprise DC
- All transactions flow from the External/DMZ sections into the enterprise DC
- London application transactions flow must be revisited – refer to section ***London***

Summary of applications discussed during the analysis:

Computer Associates SiteMinder for Single Sing On (SSO) is used to provide secure and controlled access to Web applications and portals for employees, customers and business partners efficiently with powerful Web access management supporting access to CT. Ins. Corp. resources. CA SiteMinder Federation servers provide secure SSO into CT. Ins. Corp. for various CT. Ins. Corp. business constituents using SAML. These servers communicate to back-end SiteMinder policy servers where rules reside regarding access. The policy servers use various directories to authenticate and provide authorization to users.

Navigator is one of the larger external facing applications that also connect with Pharmacies such as CVS for additional customer data.

Webtrends provides web, social, and mobile analytics and other software solutions related to marketing intelligence such as Client/Server Side Data, Webtrends Analytics, On Demand & On Premises, Visitor Data Mart, Optimization and search, User behavioral Data Social Measurement, Marketing Channel Data and Webtrends Ads.

Websphere is an application server that can go from single server to a moderately sized configuration or to dynamic web applications requiring web tier clustering and fail over across multiple application server instances. There are plugins to IBM HTTP server utilizing ODR for transactions towards web servers. ODR is a separate appliance server on inside and routs sessions to JVMs on the physical servers. The ODR plugin in the external IBM HTTP servers keep session state with the internal ODR instances. Figure 6 below depicts one architecture.

Figure 6 – IPv6 Webshpere ODR

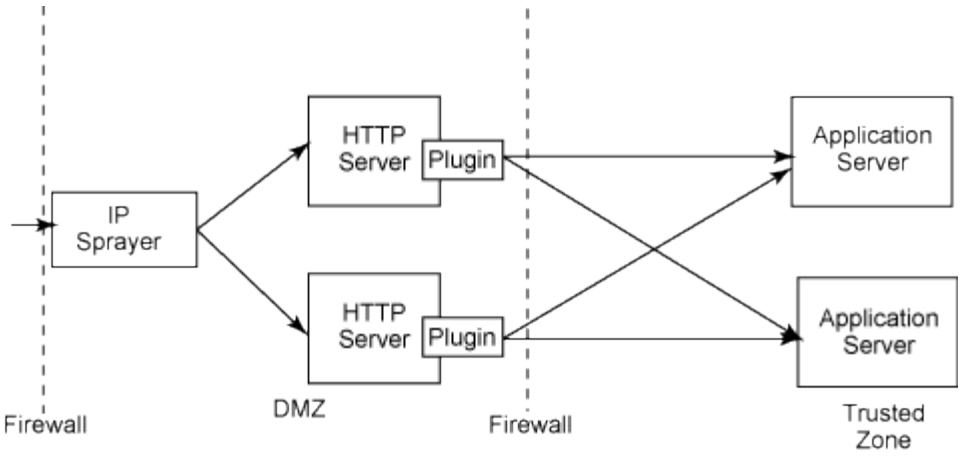

Additional details can be found here for Network and Web engineering to review.
http://www.ibm.com/developerworks/websphere/techjournal/0509_brown/0509_brown.html

Risk – *when the Dual-Stack deployment moves into the DMZ additional research must be conducted with the application development teams in terms of IPv6 API support of the Webshpere and ODR plugin to ensure protocol to application parity.*

PCI is also a critical component but due to the scope of the initial analysis was not covered.

The following is an informal prioritized table of external facing applications, services, instances CT. Ins. Corp. relies on to support its customers, partners and employees via the internet.

Table 8 – Public Facing Applications

Application/Service Name/type	Priority to company High or Normal	External Facing
-IBM HTTP Server Instances		
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y

Application/Service Name/type	Priority to company High or Normal	External Facing
Instances		
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	High	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y
	Normal	Y

Application/Service Name/type	Priority to company High or Normal	External Facing
Apache	High	**Y**
IBM HTTP Server I-H-S	High	**Y**
Microsoft IIS 6 up	High	**Y**
Navigator	High	**Y**
Eservices	High	**Y**
SFTP	High	**Y**
Email SMTP	High	**Y**
Secure Email	High	**Y**
Anyconnect VPN(SSL)	High	**Y**
VPN IPSEC	High	**Y**
PAL (Salseweb)	High	**Y**
LDAP	High	**Y**
Citrix SSL	High	**Y**
WebTrends	High	**Y**
Mircosoft ISA 2003	High	**Y**
Avaya voive portal	High	**Y**
DNS(GSS)	High	**Y**
BIND 9.4.3	High	**Y**
Data Power	High	**Y**
Ipass	High	**Y**
MS LYNC UNC	Normal	**Y**
MS LYNC UNC	Normal	**Y**
Federated Services	Normal	**Y**
Claims Xten	Normal	**Y**
EPSM	Normal	**Y**
Microsoft WSUS	Normal	**N**
PCI	Normal	**Y**
B2B services Internet	Normal	**Y**
Qualys Scanning	Normal	**N**

The key concern for IPv6 in regards to applications and their platforms is what AMI deems as **"agnostic parity"** across applications. The applications above layer 4 of the OSI model should be completely agnostic to the delivery L3 protocol used. However some applications may embed L3 addresses for various higher level application specific functions, or the various APIs used by the underlying web or application server do not entirely interface properly to the L3/4 server OS stack for IPv6. These types of concerns should be considered and tested before or during the lab or pilot testing phase of the IPv6 deployment.

AMI recommends CT. Ins. Corp. to encourage its application, web and network engineering resources as part of theIPv6 project initiate a project phase which entails prioritizing the top critically exposed applications with the greatest impact to the organization and impact and conduct the necessary research and lab testing for IPv6 use and agnostic parity. Case studies from certain vendors may already have customers that are using similar application platforms in an IPv6 environment that CT. Ins. Corp. may not be aware of.

Carriers and Third Party Vendors

Transport/Circuits

For IPv6 contingency planning and discovery purposes the ISP providers were asked about their IPv6 roadmaps using the AMI prepared ISP Questionnaire. This questionnaire focuses on key areas that will affect CT. Ins. Corp.'s ability to transition to an IPv6 infrastructure. The high- level results outlined in each carrier's questionnaire response are provided in tables 9 and 10.

Table 9 – Verizon IPv6 Services

ISP IPv6 Questionnaire			
Question	Verizon OC-x	Verizon Ethernet	Response
Do you currently have an IPv6 roadmap for IPv6 Dual Stack Services that support <CLIENT>?			Our responses will be for Public IP only. We don't really break this out OCx v. Ethernet. My answers are the same for both.
1. What is your roadmap for IPv6 dual stack circuit support for Internet based services?			Verizon (UUNET/MCI) has a lengthy history with IPv6 dating back to 1998. We deployed Public IP IPv6 in AS701 (North America) in 2006 as a GRE implementation, and then in 2007 as dual-stack. In 2011, we extended this to AS702 (Europe) and AS703 (Asia/Pacific).
1. Can you provide IPv6 and IPv4 services on the same physical circuit or is a separate IPv6 connection required?			In most situations we can provide IPv6 and IPv4 on the same physical circuit. This is referred to as "dual-stack". There are a scant few locations where IPv6 is not available. In this situation, a circuit back-haul would be required, or there is a GRE option. The GRE option would ride the IPv4 circuit.
2. Locations (by region) and deployment timelines (How much of your network supports IPv6 now?)			IPv6 is available in nearly all locations throughout the United States, Europe, and Asia-Pacific. There are some locations in Canada, with more coming in early 2013. Latin/South America is scheduled for Q4'2012.
c. Speeds and feeds offered			The entire product set (in terms of port size) is offered with the exception of 100MB via Fast Ethernet port and NxT1 using MLFR. 100MB port terminated via GigE is available. NxT1 using MLPPP is available.
2. What routers are required/supported by your internet based services?			
a. Vendor			There are no specific vendors or models required. Customers can choose their own CPE so long as it meets the IETF standards for IPv6 and has appropriate code. We do not maintain a specific list. Managed router is optional and only available is the customer is running dual-stack. We do not currently manage devices on IPv6-only circuits.
b. Software version			We do not have a CGN option in the network. There are no plans to offer this at the present time. Our recommendation is for customers to go dual-stack.
3. Is a managed router required?			No
4. Do you support some form of NAT solution (CGN/LSN) in the Internet cloud that allows for translation of IPv4 to IPv6 and vice versa?			Not on Verizon network, possible at customer premise or customer managed network
a. If so, when is this available?			NA
b. If not are there plans to provide this type of service?			No plans at this moment, except customer premise or customer managed
5. Do you support any other translation type services?			Not on Verizon network, possible at customer premise or customer managed network
a. Proxy based services			Not on Verizon network, possible at customer premise or customer managed network
b. Load balancer services (such as F5)			Not on Verizon network, at customer premise or customer managed network or Verizon managed as managed load-balancing device
c. NAT appliance on customer premise			Yes customer managed devices
d. Router based NAT			Not on Verizon network, possible at customer premise or customer managed network
6. Do you provide internet based firewall services (in the cloud)?			We have managed FWs, not IPv6 ready though.
a. If so, what is your roadmap for locations (by region) and deployment timelines			IPv6: History at Verizon Verizon Business has many years of IPv6 experience with actual customers and applications: · 1998: Experimental IPv6 Service through vBNS+ network · 2002: MAE Internet Exchange IPv6 peering service · 2004: Global Public IPv6 Internet Service · 2006: Public IPv6 Internet Service on AS701 · 2007: Dual-stack public IPv6 Internet Dedicated Services on AS701 · 2011: Native, Dual-stack, and Tunneled IPv6 Internet Dedicated Services on AS701, AS702 & AS703 IPv6: Access Methods · Native - Same as IPv4 today · Dual Stack - Transportation of both IPv6 and IPv4 over a single access line · Encapsulated within IPv4 Tunnel - Using a Tunnel Mechanism, GRE for Public IP, to transport an IPv6 packet to an IPv6 destination router IPv6: Future Plans · 4Q2012: Native, Dual-Stack, and Tunneled IPv6 Internet Dedicated Services in LATAM · 4Q2012: Dual Stack Deployment for FiOS and Internet DSL · 1H2013: Native, Dual-Stack, and Tunneled IPv6 Internet Dedicated Services in Canada In addition to our network support for IPv6, Verizon also provides IPv6 Transition Professional Services.
7. Do you provide DNS services via the Internet with IPv6 records?			Yes, depends on features.
a. What IPv6 services are offered by region and what are the deployment timelines			IPv6 transport on Verizon networks and customer premise and customer managed network, most of IPv6 vendor available features.
b. What are the IPv6 integration requirements			There are no specific vendors or models required. Customers can choose their own CPE so long as it meets the IETF standards for IPv6 and has appropriate code. We do not maintain a specific list. Managed router is optional and only available is the customer is running dual-stack. We do not currently manage devices on IPv6-only circuits.
8. What is the minimum IPv6 prefix announcement that is allowed via the Internet peering? (/30, 32, /36, /40, /44 or /48)			
a. Do you accept prefix announcements larger than a /48 such as a /49, /50 etc?			Verizon will accept IPv6 announcements all the way to /128. Prefixes longer than /48 will not be advertised to other regions, customers, or Peers.
b. Do you filter out prefix announcements larger than a /48 from other carriers?			Yes. This is the current industry best practice.
9. Will you accept an IPv6 prefix assigned from any RIR (such as ARIN or RIPE) or is it restricted to the RIR within the region where the internet point of presence is located?			Yes. We have no advertising restriction based on which RIR the IP space originated from.
10. Do your downstream peering partners accept provider independent IPv6 announcements from all RIR's?			Yes, to the best of our knowledge.
a. Are there known filters for prefix sizes?			Yes.
b. What is the minimum prefix announcement allowed?			The current industry best practice is /48.
11. Do you have configuration guides that provide best practices for integrating IPv6 services?			We can provide a very basic interface configuration tempate for Cisco and Juniper. Any further integration assistance would require Professional Services or Managed Services depending on the specific requirements.
12. Do your peering partners support IPv6?			Yes (any Peers that support IPv6 are supporting IPv6 Peering).
13. Are there any unique or special requirements for connecting a dual stack router to your Internet points of presence?			Pre-verification of any existing IPv4 circuit is required to make sure that the router and interface are enabled for IPv6.
14. Can you provide you standardized Internet architectures, topology diagram for your Internet connections? URL's to standard policies and maps is acceptable.			We do not publish this on our public website, but we can make this available under NDA.

Table 10 – AT&T IPv6 Services

ISP IPv6 Questionnaire		
Question	AT&T OC48	AT&T TDM up to OC12 / Ethernet up to 10G
Do you currently have an IPv6 roadmap for IPv6 Dual Stack Services that support <CLIENT>?	Yes	Yes
1. What is your roadmap for IPv6 dual stack circuit support for Internet based services?	OC48 available 2014. All information provided for TDM / Ethernet will apply for OC48 when the service becomes available.	Available today.
1. Can you provide IPv6 and IPv4 services on the same physical circuit or is a separate IPv6 connection required?		AT&T can provide IPv6 and IPv4 services on the same physical circuit using a "dual stack" configuration. "Dual stack" refers to a network stack that supports both IPv4 and IPv6. Dual stack nodes have the ability to send and receive both IPv4 and IPv6 packets. They can directly interoperate with IPv4 nodes using IPv4 packets, and also directly interoperate with IPv6 nodes using IPv6 packets.
2. Locations (by region) and deployment timelines (How much of your network supports IPv6 now?)		US Domestic, Most of World; 100%
c. Speeds and feeds offered		US Domestic: Available TDM up to OC12 and Ethernet up to 10G, MOW: Available TDM up to T3/E3 and Ethernet speeds / availability are country dependent
2. What routers are required/supported by your internet based services?		AT&T Managed routers are Cisco devices sized according to circuit. Customer provided routers will have no specific AT&T requirement as they are the customers responsibility.
a. Vendor		
b. Software version		
3. Is a managed router required?		No
4. Do you support some form of NAT solution (CGN/LSN) in the Internet cloud that allows for translation of IPv4 to IPv6 and vice versa?		Not at this time.
a. If so, when is this available?		This function is not on the roadmap.
b. If not are there plans to provide this type of service?		Not at this time.
5. Do you support any other translation type services?		Not at this time.
a. Proxy based services		
b. Load balancer services (such as F5)		
c. NAT appliance on customer premise		
d. Router based NAT		
6. Do you provide internet based firewall services (in the cloud)?		Yes. Network-Based Firewall (NBF) is designed to protect a customer's MPLS WAN while providing a path to the Internet. NBF provides a customer with outbound and inbound Internet access through one of seven AT&T data centers around the globe. It also provides centralized firewall management and consistency of policy enforcement across multiple sites without installing any equipment at any of the customer's sites. The traffic is enforced according to a customer-defined security policy.
a. If so, what is your roadmap for locations (by region) and deployment timelines		Globally available.
7. Do you provide DNS services via the Internet with IPv6 records?		Yes
a. What IPv6 services are offered by region and what are the deployment timelines		Primary and Secondary DNS hosting is globally available for IPv4 and IPv6 address resolution.
b. What are the IPv6 integration requirements		
8. What is the minimum IPv6 prefix announcement that is allowed via the Internet peering? (/30, 32, /36, /40, /44 or /48)		AT&T will advertise subnets to peers only if they are /48 or larger blocks (/48, /40, etc.).
a. Do you accept prefix announcements larger than a /48 such as a /49, /50 etc.?		AT&T will accept customer route announcements for use in our core between customer locations of variable subnet length up to /64, however, AT&T will advertise subnets to peers only if they are /48 or larger blocks (/48, /40, etc.).
b. Do you filter out prefix announcements larger than a /48 from other carriers?		No. We will accept blocks larger than /48. (/40, /36, etc.)
9. Will you accept an IPv6 prefix assigned from any RIR (such as ARIN or RIPE) or is it restricted to the RIR within the region where the internet point of presence is located?		AT&T will accept PI prefixes from any RIR in all of the global regions we operate. However the RIRs are encouraging the use of regional PI prefixes and not across regions. Therefore when this policy becomes more mature, AT&T may change its policy to reflect the new standard.
10. Do your downstream peering partners accept provider independent IPv6 announcements from all RIR's?		Yes, at this time.
a. Are there known filters for prefix sizes?		Minimum of /48
b. What is the minimum prefix announcement allowed?		/48
11. Do you have configuration guides that provide best practices for integrating IPv6 services?		Yes.
12. Do your peering partners support IPv6?		AT&T has IPv6 peering with 21 of its existing 23 US peers using IPv6. This peering takes place across a subset of the existing IPv4 peering links where AT&T has enabled dual-stack connectivity. AT&T has the technical capability to support IPv6 peering at all its peering cities in the United States.
13. Are there any unique or special requirements for connecting a dual stack router to your Internet points of presence?		No.
14. Can you provide you standardized Internet architectures, topology diagram for your Internet connections? URL's to standard policies and maps is acceptable.		The attached guide has all supported architectures in section 3.

IPv4 Use Cases

AMIutilizes Use Cases as its approach to identifying key areas of the network that require IPv6 support. The intent of the use cases is to identify users and traffic flows through the network. This assists in identifying what user communities connect to what internal and external applications over the Internet service points.

The use cases provide a framework for understanding the flow of information and application connectivity into and out of the CT. Ins. Corp. intranet and determining the possible impact of global IPv4 exhaustion on these flows. The use cases also assist in defining where dual-stack, translation, and tunneling techniques should be applied to enable IPv6 communication. The use cases further provide the foundation for categorizing high, medium, and low risks, provide inputs to the architectural strategy, and facilitate the classifications of IPv6 impact in following areas.

IPv4 Internet Access

- **Inbound Internet** - external user communities are connecting to CT. Ins. Corp. internal applications and what IT systems, hosts, and networks are required to deliver these communications. This includes use cases for access to .COM domains, VPN users, remote site IPSEC connections, and partner connectivity.

  Case 1: Internet web presence – B2C Business to Consumer
  - Plan Members, New customers, general information
  - Benefits services, Personal Health records, International
  - Employee web based services

  Case2 : Inbound Employee and Contractor Access
  - Travelers
  - Telecommuter (VPN)

  Case 3: Inbound B2B Applications (Business to Business)

  Case 4: Inbound Email

- **Other** – These use cases may not have a direct correlation to inbound Internet connectivity but were important in the overall readiness analysis.

The use cases serve as a foundation for the development of test strategies, because use case remediation solutions will need to be tested and validated to ensure that IPv6 communications are functioning properly.

The remainder of this section provides a summary related to the Inbound and Other use cases. The same technology may be used as the access method across multiple use cases. Therefore, uses cases sometimes are a combination of multiple technologies to comprise one connectivity model. For example, SSL VPN may use the same architecture for a or an employee, but the specific devices that the data traverses may be different.

Case 1: Internet Web Presence

CT. Ins. Corp. allows inbound internet access via its two major points of connections to the internet in each of its data centers outlined earlier in the previous section. The data centers each have a secure DMZ with firewalls that separate other services (VPN, B2B, QA) deemed as their own DMZs. There are customer facing servers present within the DMZ that provide content and services for members, potential customers, employees, business partners, affiliates and the related services for single sign on, directory services, authentication for secure access to content. Members and customers access CT. Ins. Corp. products and services through these "front facing" servers. The emphasis here is "through" for the content and initial portal the customer sees is mostly static on these servers. Customer/Member data is processed and forwarded into the internal network's data center server farms for further processing and response.

Risk

Without support for IPv6 CT. Ins. Corp. is in risk of losing additional global visibility to its products and services from IPv6 only site visitors.

Case 2: Inbound Employee/Contractor Access

CT. Ins. Corp. uses several methods for its users and business partners to access extranet services.

For CT. Ins. Corp. employees there are several options for extranet VPN and Remote Access services.

IPsec VPN connection deployed today using the Cisco AnyConnect software or classic VPN client and authentication through non-clustered individual Cisco ASA firewall appliances in both data centers. Currently about 60% of the user base uses AnyConnect. CT. Ins. Corp. estimates to be off IPsec based services by the end of the year. However, this is just an estimate and any IPv6 planning should still consider IPsec based services.

CT. Ins. Corp. also uses a Web based VPN for public access computers, users who do not require full access to the CT. Ins. Corp. Network, or users who cannot load the Cisco AnyConnect CT. Ins. Corp. software.

The Citrix Netscaler platform is deployed redundantly in both datacenters. The Secure Socket Layer (SSL) based Virtual Private Network (VPN) provides users with secure access to specific enterprise applications, such as e-mail, web applications, and web browsing, without requiring them to have VPN software installed on their end-user devices.

Aruba – CT. Ins. Corp. is currently in a pilot phase of project to extend full connectivity to "micro offices", and small offices with cable based connections as an alternative to MPLS wan based services. This IPsec L2 based tunnel solution targets offices with 1 to roughly 20 users in a home or small office. The Aruba VPN provides network connectivity to the small office and extends the internal routing domain to the small office as if they are connected via a traditional WAN links. Clients can access the network though their laptop or utilize thin client products.

Sonicwall – 8 appliances 4 in each DC in clusters.
Used for CT. Ins. Corp. OWA email clients and Affiliates – low use and is expected to be deprecated by end of year and users moved to Citrix.

Cisco 7200VXR based VPN - 7 units distributed across both DCs and are expected to be deprecated soon since they do not support Any Connect.

Risk

CT. Ins. Corp.'s extranet VPN/Remote access is a critical productivity tool for its employees and partners. Based on the criticality of this service, the inability for CT. Ins. Corp. to support IPv6 connectivity from broadband endpoints (cable/dsl/cellular) domestically or globally; could have significant productivity and potential financial impact.

Case 3: Inbound B2B Application support

There is a separate "DMZ" carved out of VLANs for Cisco ASA firewalls that support various business partner access to CT. Ins. Corp.. These Cisco ASAs provide traditional IPSec tunnels into the Surf/Internal side of CT. Ins. Corp.'s network for direct partner connectivity via the internet. The B2B firewalls usually have 50 plus tunnels but will dwindle down due to growing Citrix use for B2B connectivity.

Risk

The domestic risk depends on CT. Ins. Corp.'s agreements with its business partners. The domestic business partners already on IPv4 have time to adopt IPv6 on their end. CT. Ins. Corp. is not forcing existing B2B vendors to embrace IPv6 immediately but new vendors are queried for their IPv6 capabilities. The risk is for international and new global partners that may only communicate via IPv6. CT. Ins. Corp. must provision to, at minimum, have the capability in place when that first IPv6 only request arrives.

Case 4: Inbound Email

CT. Ins. Corp. provides inbound SMTP based email services for internet users by utilizing a pair of un-clustered MacAfee Ironmail Edge appliances in each DMZ. These appliances have Anti-Virus engines built in. They send traffic to the internal mail servers/appliances and reverse. The email environment is currently under re-evaluation.

Secure encrypted email system based on Cisco IDA appliances are also in use but these are near EOL and are on the internal portion of the network thus not covered in this analysis.

CT. Ins. Corp. bulk mail – newsletters, mailings, marketing, member information is processed by two Cisco appliances, one in each DC, C370 Ironports on the internal portion of the network and are not covered by this analysis.

Risk

The global need to process email from native IPv6 packets is important to CT. Ins. Corp.'s ability to remain flexible in its communications stature. External email system today can support Dual-Stack IPv4/6 however if IPv6 is not deployed into the internal zones of the network to talk to the DC mail servers so the external IPv6 paradigm is broken because the email appliances in the DMZ do not conduct translation services. Additional HW/SW solution may need to in place thus adding to the cost and complexity of supporting IPv6 in the External and DMZ Zones.

Other Applications

Some of the other applications and services were discovered but due to scope are not covered in detail. A partial listing is below. A list of applications discovered is included in section ***Applications - External Facing.*** When planning for IPv6 integration these applications and their use cases should be reviewed to determine impact and risk levels.

SFTP, ESM, IPass, Pal, Lync, PCI, Webtrends, Siteminder Federated Services, Claims Xten, Navigator.

IPv6 Asset Readiness and Costs

Scope and Methodology

The data collected during the Discovery Phase of this engagement was analyzed to determine the extent to which current network hardware and software support IPv6 features. This process involved analyzing vendor documentation and communicating directly with vendors to compare the IPv6 features currently offered, or on the vendors' roadmap, to the features required by the CT. Ins. Corp. Use Cases outlined in this document.

Support of IPv6 is not a simple check box requirement. Experience has shown that some vendors will indicate that their equipment is IPv6 compliant, when in fact it has limited IPv6 functionality. In most cases, full compliance will not be determined until the device has been configured with all of the required features and tested in a lab environment. Outside of testing every device, which is not in scope for this engagement, AMI has had to rely on the vendor claims for IPv6. This puts an additional emphasis on testing in a lab or pilot environment to validate IPv6 and IPv4 feature parity.

Asset Readiness

The following table is a summary of the IPv6 compliance of the assessment's technologies:

Area of Technology	Product	IPv6	Can be	Can be	Risk	What Needs To Happen
			Dual Stack	Translation	Caveats	Dual Stack -GREEN
		Ready	Candidate	Candidate		Translation - Blue
Client VPN	Cisco ASA	Y	Y	Y	Code base is IPv6 compliant but outdated	For Dual Stack
B2B						client should upgrade to latest tested version
Ipsec FW						for IPv6 enhancements and bug remediation
Proxy FW						has occurred in later revisions
DMZ FW						For Translation NO ACTION
Client VPN	Citrix 10500MPX	Y	Y	Y	None	For Dual Stack NO ACTION
					Citrix can act as translation proxy too	For Translaiton NO ACTION
Client VPN	Cisco 7206VXR	N	Y	Y	Code base is outdated	For Dual Stack upgrade to 15.0(1)M3 +
					Anyconnect issue	For Translation NO ACTION
Client VPN	Sonicwall	N	N	Y	Deprecated clients moved to Citrix	Move clients to citrix platform
						For Translation No Action
Client VPN	Aruba	Y	Y	Y	End clients cable op. provides IPv6 only	For Dual Stack No Action
						For Translation No Action
Email	McAfee Ironmail	Y	Y	Y	Internal systems do not speak IPv6	For Dual Stack must extend IPv6 into enterprise
						For Translation No Action
						Verify if Mcafee can translate email 6/4
DNS	McAfee BIND	N	N	Y	Code is outdated does not support IPv6	For Dual Stack must upgrade
						For Translation NO ACTION
						BIND 9.7.3 for 8.3 but 9.8+ meets NIST USGv6
	Cisco GSS	N	N	Y	Code is outdated does not support IPv6	For Dual Stack must upgrade
						For Translation NO ACTION
						Replaced with F5
SFTP	Axway 5620	Y	Y	Y	None	For Dual Stack No Action
						For Translation No Action
Outbound Firewalls	MacAfee 2150e	N	N	Y	Code is outdated does not support IPv6	For Dual Stack must upgrade
Proxies	MacAfee WG-5500					For Translation NO ACTION
Ext Cloud	McAfee S5032				None cloud not used often	
DMZ FW						
External Router	Catalyst 6506	N	N	Y	Code is outdated does not support IPv6	For Dual Stack must upgrade software
					Required for first contact	For Translation must upgrade software
					Will Migrate to ASR platform in future	
					Will SPA be compatible with ASR	Cisco verification
					Port channels don't support IPv6 at L3	
					Port channel L2 load balance verification	Cisco verification
					on IPv6 addresses	
					TCAM and dCEF verify	Cisco verification
					Cisco conducting bug scrub	Cisco verification
					Netflow	Upgrade to Version 9 for IPv6
Secure DMZ Router	Catalyst 6506	N	N	Y	Code is outdated does not support IPv6	For Dual Stack must upgrade software
					Will Migrate to ASR platform in future	For Translation must upgrade software
Servers in DMZ					If used NIC teaming/Bonding TCP offloading	For Dual Stack must test
					May not work with Dual stack	For Translation NO ACTION
Content Switching LB	Cisco CSS11503	N	N	Y	Code is outdated does not support IPv6	For Dual Stack must upgrade
						For Translation NO ACTION
						To be replaced with F5
Applications	IBM HTTP Server	N	Y	Y	ODR plugins must be tested	For Dual Stack must test
	ODR plugin					For Translation NO ACTION
	Apache	Y	Y	Y	None	For Dual Stack No Action
						For Translation No Action
	Microsoft IIS 6	Y	Y	Y	Only with Windows Server 2003	For Dual Stack must upgrade 2003 servers
						For Translation No Action
Network Management	Tivoli	Y	Y		May not glean IPv6 statistics	Check with NMS vendor on MIBS supported
	HP OPSware	Y	Y		May not glean IPv6 statistics	Check with NMS vendor on MIBS supported
	Satellite	N	N		Current Red hat versions deployed do not	For Dual Stack must upgrade
					support IPv6	For Translation NO ACTION
Security Services	IXIA ANUE	Y	Y	Y	Outside services that parse data from ANUE	For Dual Stack NO ACTION
					does not understand IPv6	For Translation NO ACTION
	QualysGuard	Y	Y	Y	None	For Dual Stack NO ACTION
						For Translation NO ACTION
	Corero Network Secur	N	N	Y	Outside services that parse data from Corero	For Dual Stack must upgrade
					does not understand IPv6	For Translation NO ACTION

Table 11 – Summary IPv6 compliance assessment

In the following sections, AMI used the data collected during Discovery to ascertain the IPv6 readiness compliance, costs, and priority for CT. Ins. Corp. to move to IPv6 addressing in their network.

As mentioned above, CT. Ins. Corp. has a very extensive and updated lifecycle refresh initiative that is ongoing. During this analysis, any devices that were critical to supporting IPV6 that are slated for lifecycle refresh were not included as an additional cost. Based on this approach all of the critical network and firewall devices in the external/DMZ segment of the CT. Ins. Corp. environment currently support IPv6, will support it with software/IOS upgrades or are planned to be replaced as part of current lifecycle plans.

The following table represents the estimated costs for the upgrades that will need to occur on these devices:

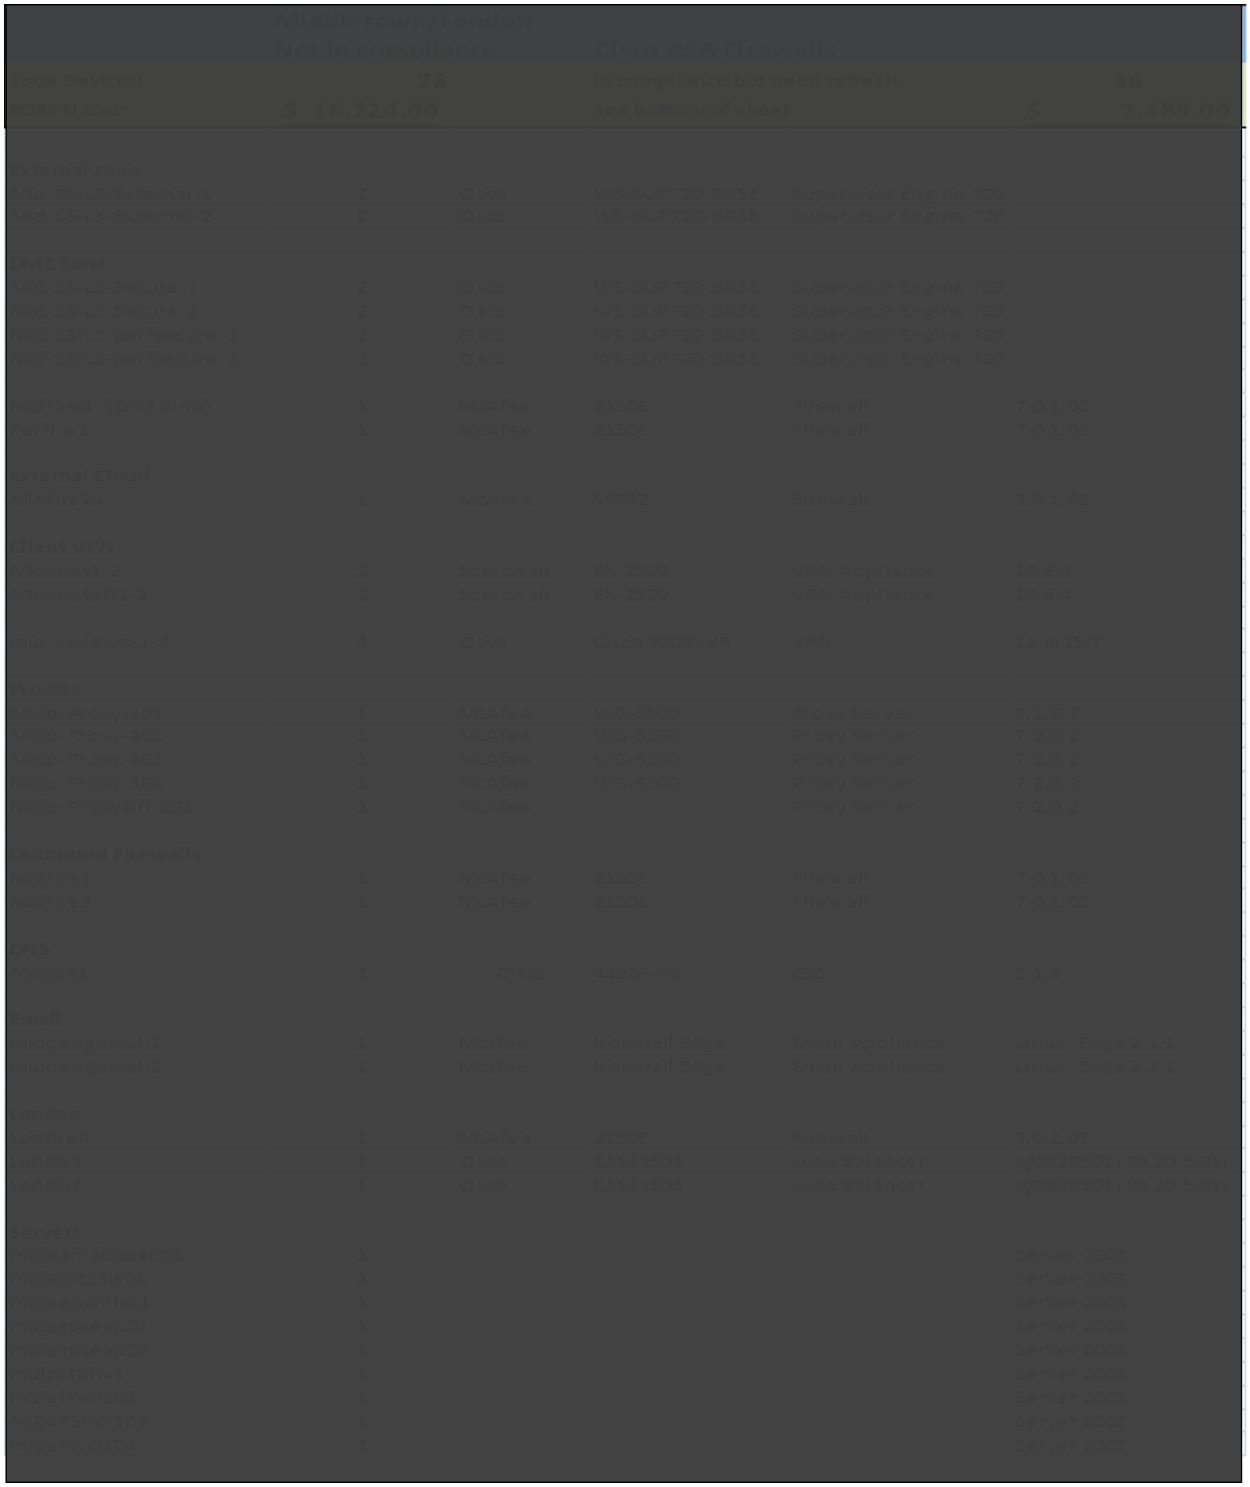

.

IPv6 Architecture

IPv6 Architecture Options

All organizations looking to adopt IPv6 will have two target architectures. The long-term strategic architecture option is dual-stack infrastructure. This will enable services to be available using either IPv4 and/or IPv6 transport. Dual-stack allows IPv4 and IPv6 to co-exist in the same devices and networks.

The short term or tactical architecture option uses tunneling or translation devices to allow IPv6 and IPv4 devices to communicate with each other.

- Tunneling allows IPv6 packets to be transmitted over an IPv4 infrastructure or IPv4 packets over an IPv6 infrastructure and IPv6 becomes more prevalent.
- Translation allows IPv6 only devices to communicate with IPv4 devices when a dual-stack cannot be configured on a device. An example is that the current DMZ mail servers that CT. Ins. Corp. uses do support dual-stack but cannot support simultaneous use of both address types, therefore in order for the mail servers to receive mail externally on an IPv6 address and communicate internally to the corporate mail servers (IPv4), translation must be used.

AMI expects dual-stack to be supported for many years into the future, as there will be legacy applications that cannot be converted to dual-stack. There is a lot of historical evidence for support of legacy protocols, and their need to continue in the face of mounting costs and potential superior technological solutions. The future end-state solution will be to reduce the dual-stack environment to IPv6 only, by removing all IPv4 from the environment. AMI expects that IPv6 only networks will not exist for a number of years (5-10) based on the complexity of migrating all infrastructure and applications to native IPv6.

IPv6 Architecture Road-Map and Recommendations

After completing the discovery, analyzing the collected data, and reviewing the recommended architectural options with CT. Ins. Corp. stakeholders and support contacts the following roadmap is presented.

To address the risk and gaps of global visibility plus position CT. Ins. Corp. for additional technological flexibility with the new protocol, AMI is suggesting the implementation of a hybrid approach combining the strategic direction of Dual-Stack functionality while enabling translation capability to address timing and cost containment to enable IPv6 functionality in the 12-24 month timeframe.

The recommended approach for CT. Ins. Corp. is to utilize a combination of Translation and Dual-Stack technologies at the edge of its network. The immediate benefits are:

- Translation platform refreshes IPv4 SLB/DNS/CSS-GSS platforms
- Provides immediate visibility for IPv6 only customers to CT. Ins. Corp. globally by addressing risk and gaps outlined
- Low impact, cost and effort to execute, eases CT. Ins. Corp. into IPv6
- Provides time and controlled pace for CT. Ins. Corp. to plan and execute its strategic IPv6 goals

Additional benefits and attributes of the tactical approach are outlined in the section ***Migration Strategy and Timelines Phase 5.***

The strategic recommendation is to continue to plan, test and phase in the progression Dual-Stack protocols either from the edge of the network inward towards the enterprise or from the enterprise out towards the tactical translation point, remove translation and complete the Dual-Stack progression.

The dual-stack architecture is preferred, from a scalability and performance standpoint. However based on the requirements to phase in the migration and the current timelines associated with dual-stack support across all infrastructure elements within the target environments, CT. Ins. Corp. will need to focus on the tactical addition of translation capability.

Although tunneling can provide needed functionality to enable migration to IPv6, AMI has found based on customer and lab testing that tunneling can be very problematic in a number of custom and "Off the shelf" applications, has security implications and is not even further considered for CT. Ins. Corp.'s case.

By executing this strategy for IPv6 support, CT. Ins. Corp. will be able to address the following IPv6 use cases:

Case 1: Internet Web Presence.

Risk - Without support for IPv6 CT. Ins. Corp. is in risk of losing additional global visibility to its products and services from IPv6 only site visitors.

Remediated – Translation solution immediately provides CT. Ins. Corp. controlled and phased visibility of web presence for new customers, vendors and members globally.

Additional benefits:

1. Can provide web content and services without touching its current web platform(reduces initial OPEX/CAPEX)
2. CT. Ins. Corp. can control which content goes IPv6 first at translation device. For example start with the top 10 sites and then gradually add more.
3. Can gradually monitor IPv6 traffic using same tools for analytics data(NAT64 IP addresses can be the index). Volume metric data can determine IPv6 demographic usage and possibly for sales use.
4. Provides CT. Ins. Corp. time to plan, test and migrate towards the preferred Dual-Stack architecture.

Case 2: Inbound Employee/Contractor Access

Risk - CT. Ins. Corp.'s extranet VPN/Remote access is a critical productivity tool for its employees and partners. Based on the criticality of this service, the inability for CT. Ins. Corp. to support IPv6 connectivity from broadband endpoints (cable/dsl/cellular) domestically or globally; could have significant productivity and potential financial impact.

Case 3: Inbound B2B Application support

Risk - The domestic risk depends on CT. Ins. Corp.'s agreements with its business partners. The domestic business partners already on IPv4 have time to adopt IPv6 on their end. CT. Ins. Corp. is not forcing existing B2B vendors to embrace IPv6 immediately but new vendors are queried for their IPv6 capabilities. The risk is for international and new global partners that may only communicate via IPv6. CT. Ins. Corp. must provision to, at minimum, have the capability in place when that first IPv6 only request arrives.

Remediated – Translation solution immediately provides IPv6 only CT. Ins. Corp. B2B partners and employees continued ability to utilize existing SSL and IPsec remote access and VPN tools. Translation eliminates IP protocol geographical gap preventing partners and users from accessing CT. Ins. Corp. resources and ensuing continued productivity. The same additional benefits available from user case 1 also apply to this case.

Case 4: Inbound Email

Risk - The global need to process email from native IPv6 packets is important to CT. Ins. Corp.'s ability to remain flexible in its communications stature. External email system today can support Dual-Stack IPv4/6 however if IPv6 is not deployed into the internal zones of the network to talk to the DC mail servers so the external IPv6 paradigm is broken because the email appliances in the DMZ do not conduct translation services. Additional HW/SW solution may need to in place thus adding to the cost and complexity of supporting IPv6 in the External and DMZ Zones.

Remediated – Translation solution immediately provides a global IPv6 email presence for business continuity and maintains CAPEX in existing systems. The same additional benefits available from user case 1 also apply to this case.

A series of summary steps is provided in **the Phase -5 section** on executing just a Dual-Stack progression into the DMZ zones is included for CT. Ins. Corp.'s reference.

AMI recommends the following high-level steps in order to enable IPv6 within the External/DMZ environments:

1. Implement Dual-Stack in the External zone towards the Secure DMZ zone as a first step. Dual-Stack is required for IPv6 to flow from the Internet to the Translation device.

2. Deploy translation capabilities in order to provide IPv6 support for devices which cannot support dual-stack or native IPv6 functionality. This can be implemented either inside and outside the existing external firewalls:

    a. Translation in Zone 1 - Deploy translation capabilities in the External zone towards the DMZ zones to support all the Mid-Int, Mid-Secure(real) zone firewalls, and VPN firewalls, outbound proxies etc. thus enabling IPv6 inbound packets to be translated to IPv4 prior to reaching DMZ zone IPv4 resources. Additional research for outbound will be required. This is referred to Translation Zone 1 deployment.

    b. Translation in Zone2 - Deploy additional translation capabilities within the DMZ zones if required. Dual-Stack will progress to cross the DMZ firewalls as IPv6 packets but will now be translated in their respective zone to IPv4 prior to reaching that zone's IPv4 resources. Additional research for outbound will be required. This is referred to Translation Zone 2 deployment and adds to the OPEX and CAPEX of the project.

3. Further phases will entail CT. Ins. Corp. continuing to work towards deploying Dual-Stack in the DMZ zones as hardware and software is refreshed and validated in the Lab as compliant.

The ***Migration Strategy and Timelines Section*** covers the following phases required to progress through the road map recommended in order to enable IPv6 within the CT. Ins. Corp. external/DMZ environment.

Translation Capability - Tactical

As described in the Architecture recommendations, translation will provide a critical capability enabling CT. Ins. Corp. to support IPv6 endpoints with lower impact from a technical and financial perspective. IPv6 translation is a tactical capability that will allow for a phased migration for supporting IPv6 (native or dual-stack), that will align much closer to the lifecycle/refresh plans that are currently budgeted and ongoing.

Note: Translation capability can be provided by a number of platforms from a variety of vendors. Based on the current and planned expansion of the F5 footprint within the CT. Ins. Corp. environment. AMI does concur with CT. Ins. Corp.'s plans with the F5 product set for providing this functionality within the External/DMZ environments.

Research into the use of BIG IPs F5 LTM or GTM and modules and licensing required for the translation, DNS and SLB package that provides a single consolidated platform that has several advantages.

- Easier to maintain and common
- DNS, SLB, SSL, full IPv6 proxy support in one platform
- Works with existing DNS infrastructure in place or can replace BIND based external component.
- Works with existing GSS/CSMs platforms in place or can replace Cisco based components.
- Used for translation zones plus when Dual-Stack deployment is completed in the DMZ zones translation can be turned off and other v4/v6 services such as DNS, SLB, SSL continue natively.
- 11.2.0 In this release, the BIG-IP IPv6 standards compliance has been improved.

The license received from F5 Networks determines what software modules the BIG-IP system can support. The license ensures that you can activate all software modules you have purchased. An F5 license is applicable for the life of the system, or until you reactivate it, for example, by updating the software version or purchasing additional modules. The modules available for a version of the software include **Local Traffic Manager, Global Traffic Manager,** Link Controller, Application Security Manager, Protocol Security Module, Web Accelerator, WAN Optimization Module, and Access Policy Manager.

An excerpt from F5 regarding its IPv6 capabilities:

> Because a BIG-IP device is a full proxy device, it can view all interactions between client and server. And because it is dual-stacked, it can understand and handle both IPv4 and IPv6 requests, or it can act as a gateway, translating IPv4 addresses to IPv6 and vice versa. Many F5 customers aren't aware that the BIG-IP devices they already own have these capabilities. That's especially significant for organizations like F5 that may need to accommodate website visits from IPv6-only devices.
>
> Customers can take advantage of BIG-IP devices' dual stack capabilities and use their own DNS server to resolve IPv6 requests directly. The IT group used this method for its proof of concept. It configured the BIG-IP 3900 devices for IPv6, assigning them IPv6 virtual addresses that point to www.f5.com web servers—the same physical servers that already host www.f5.com on the IPv4 Internet.
>
> No changes were made to the servers themselves. "We simply added new IPv6 addresses for most of our web properties to our DNS server in BIG-IP GTM and made those addresses publicly available.

Big IP F5 recent additional IPv6 features from TMOS version 11.1.0 and up.

IPv6 to IPv4

In this release, you can configure the BIG-IP Local Traffic Manager (LTM) to load balance IPv6-only client connection requests to IPv4-only servers on your network by returning an AAAA record response to the client.

Route Domains for IPv6

In this release, Route Domains support IPv6, providing the same capabilities of IPv4, including strict isolation between route domains and overlapping IP addresses, as well as support of a single route domain for both IPv4 and IPv6 virtual servers in the same route domain.

Analytics Enhancements

In this release, analytics:

- Support IPv6
- Support VIPRION platforms
- Work in vCMP
- Support statistics in TMSH
- Store AVR statistics in the iStats library
- Support configuration through iControl
- Support iRule hooks

BIG-IP GTM supports next-generation IPv6 networks, resolving AAAA queries without requiring wholesale network and application upgrades. As IPv6 adoption grows, BIG-IP GTM eases the transition to IPv6 by bridging the gap between IPv6/IPv4 DNS. The DNS translation between IPv6 and IPv4 networks is seamless as BIG-IP GTM provides DNS gateway and translation services for hybrid IPv6 and IPv4 solutions, and manages IPv6 and IPv4 DNS servers in DNS64 environments. For AAAA queries from clients, BIG-IP LTM configured with NAT64 transforms IPv6 to IPv4 for those IPv4-only environments. The response data is sent to the client from NAT64 using IPv6.

There is a way to support IPv6 externally while making relatively no changes to the organizational network architecture. An IPv6 gateway can provide the translation necessary to seamlessly support both IPv6 and IPv4. Employing an IPv6 gateway insulates organizations from making changes to internal networks and applications while supporting IPv6 clients and infrastructure externally.

The right IPv6-enabled solution can also help with migration internally. For example, an enabled application delivery controller like F5 BIG-IP LTM (Local Traffic Manager) can act as an IPv4 to IPv6 gateway, and vice-versa, by configuring a virtual server using one IP address version and pool members using the other version. This allows organizations to mix and match IP versions within their application infrastructure as they migrate on their own schedule toward a Dual-Stack then IPv6 only network architecture, internally and externally.

Figure 7- BIG IP F5 Pv6 Translation

Two key technology components of the BIG IP F5 deployed in Translation Zone 1 are NAT64 and DNS64

Address Family Translation(AFT) Using Stateful NAT64

AFT using stateful NAT64 is preferred over the other available IPv6 migration and transition technologies. It facilitates communication using User Datagram Protocol (UDP), Transmission Control Protocol (TCP), or Internet Control Message Protocol (ICMP) between IPv6-only and IPv4-only hosts and networks by performing:

- IP header translation between the two address families using an algorithm defined in RFC 6145/46 (IP/ICMP Translation Algorithm)
- IP address translation between the two address families using an algorithm defined in RFC 6052 (IPv6Addressing of IPv4/IPv6 Translators)

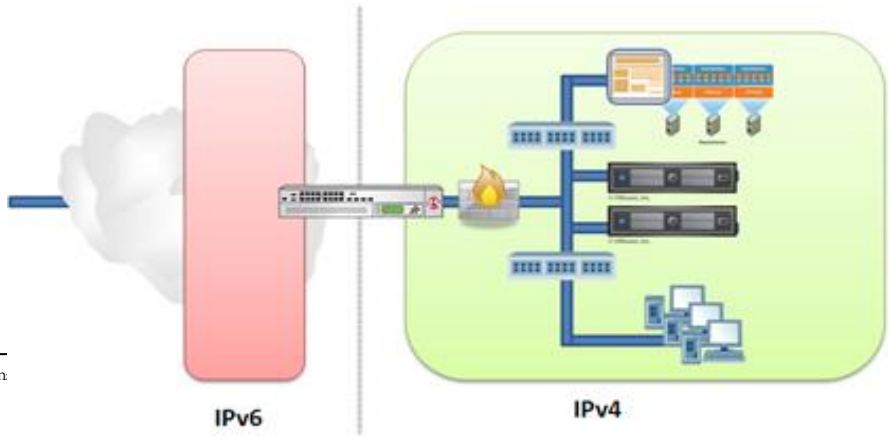

BIG-IP LTM with IPv6 Gateway Module

NAT64 (From RFC 6146)

A mechanism for IPv4-IPv6 transition and IPv4-IPv6 coexistence. Together with DNS64 [RFC6147], these two mechanisms allow an IPv6-only client to initiate communications to an IPv4-only server. They also enable peer-to-peer communication between an IPv4 and an IPv6 node, where the communication can be initiated when either end uses existing, NAT-traversal, peer-to-peer communication techniques, such as Interactive Connectivity Establishment (ICE) [RFC5245]. Stateful NAT64 also supports IPv4-initiated communications to a subset of the IPv6 hosts through statically configured bindings in the stateful NAT64.

Stateful NAT64 is a mechanism for translating IPv6 packets to IPv4 packets and vice versa. The translation is done by translating the packet headers according to the IP/ICMP Translation Algorithm defined in [RFC6145]. The IPv4 addresses of IPv4 hosts are algorithmically translated to and from IPv6 addresses by using the algorithm defined in [RFC6052] and an IPv6 prefix assigned to the stateful NAT64 for this specific purpose. The IPv6 addresses of IPv6 hosts are translated to and from IPv4 addresses by installing mappings in the normal Network Address Port Translation (NAPT) manner [RFC3022]. The current specification only defines how stateful NAT64 translates unicast packets carrying TCP, UDP, and ICMP traffic. Multicast packets and other protocols, including the Stream Control Transmission Protocol (SCTP), the Datagram Congestion Control Protocol (DCCP), and IPsec, are out of the scope of this specification.

DNS64 (From RFC 6147) is a mechanism for synthesizing AAAA resource records (RRs) from A RRs. A synthetic AAAA RR created by the DNS64 from an original A RR contains the same owner name of the original A RR, but it contains an IPv6 address instead of an IPv4 address. The IPv6 address is an IPv6 representation of the IPv4 address contained in the original A RR. The IPv6 representation of the IPv4 address is algorithmically generated from the IPv4 address returned in the A RR and a set of parameters configured in the DNS64 (typically, an IPv6 prefix used by IPv6 representations of IPv4 addresses and, optionally, other parameters).

Together with an IPv6/IPv4 translator, these two mechanisms allow an IPv6-only client to initiate communications to an IPv4-only server using the Fully Qualified Domain Name (FQDN) of the server.

These mechanisms are expected to play a critical role in the IPv4-IPv6 transition and coexistence. Due to IPv4 address depletion, it is likely that in the future, many IPv6-only clients will want to connect to IPv4-only servers. In the typical case, the approach only requires the deployment of IPv6/IPv4 translators that connect an IPv6-only network to an IPv4-only network, along with the deployment of one or more DNS64-enabled name servers. However, some features require performing the DNS64 function directly in the end hosts themselves.

Together, these two mechanisms allow an IPv6-only client (i.e., a host with a networking stack that only implements IPv6, a host with a networking stack that implements both protocols but with only IPv6 connectivity, or a host running an IPv6-only application) to initiate communications to an IPv4-only server (which is analogous to the IPv6-only host above).

An excellent white paper on the details of NAT64

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6553/white_paper_c11-676278.html

Figure 8 depicts the basic mechanics of CT. Ins. Corp.'s use of translation

Figure 8 – Translation Mechanics and Flow

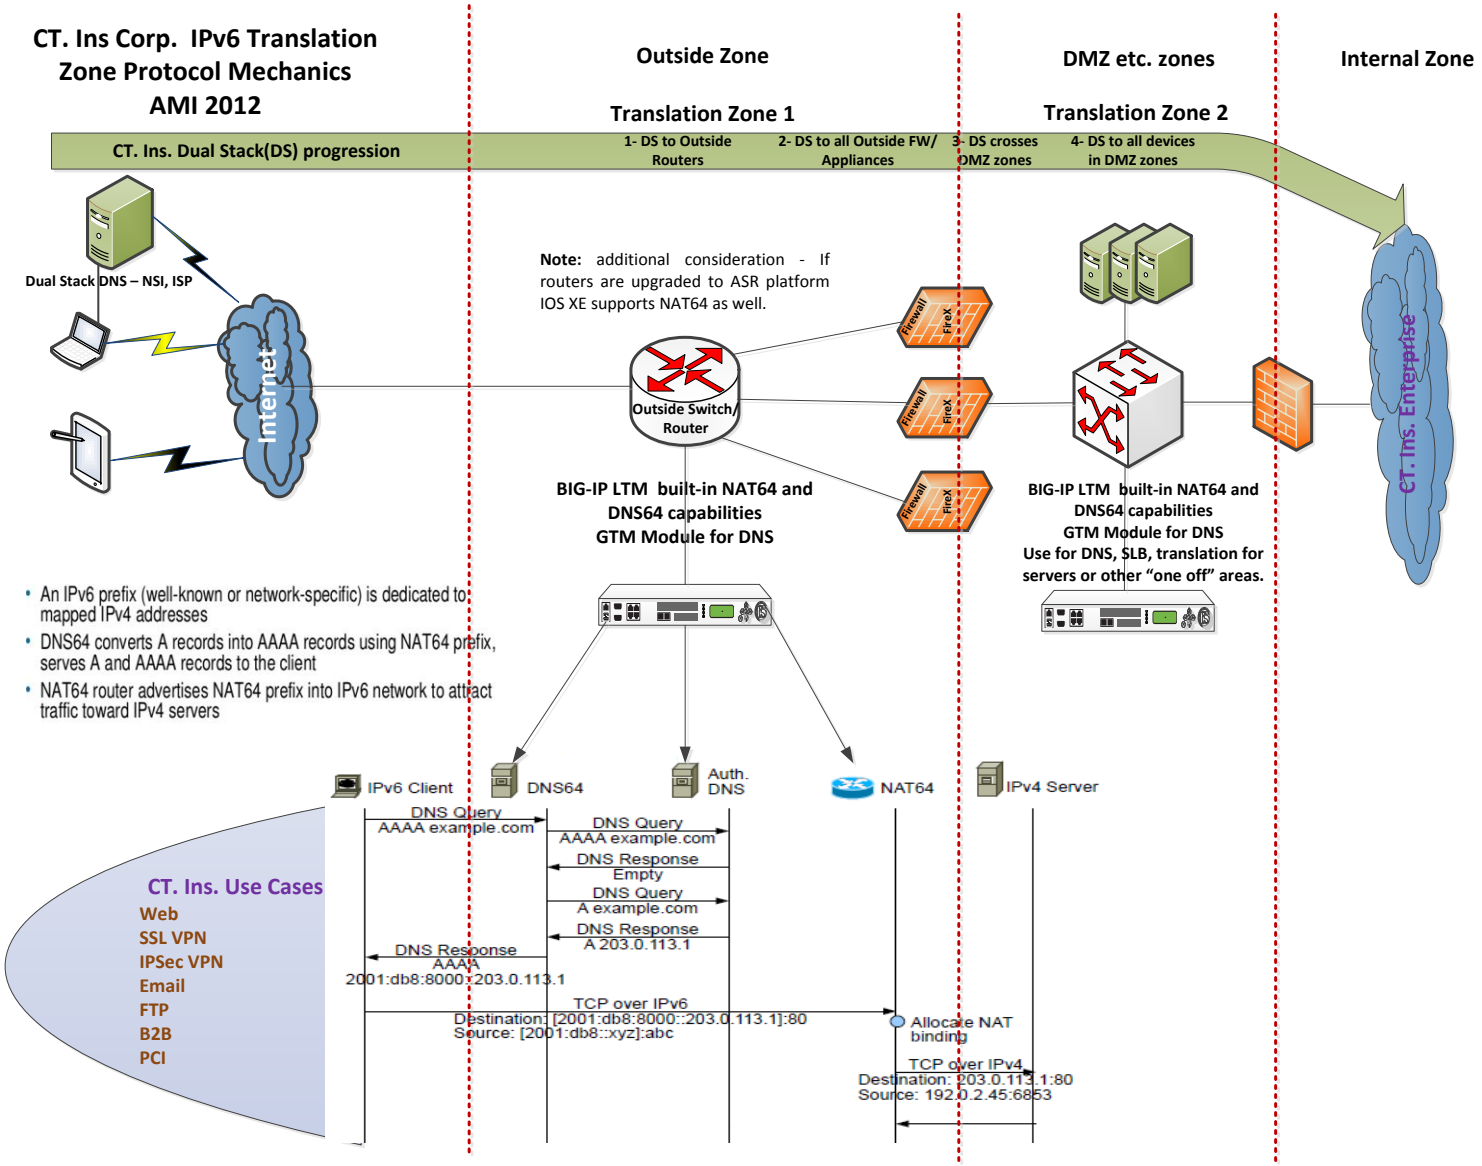

Migration Strategy and Timelines

AS mentioned in the Executive Summary AMI's recommendations are aligned across the primary phases. They provide details on the tasks that must be executed in order to meet the tactical approach's goals. Together these Timelines and Sequences provide a roadmap for addressing the risks associated with the IPv4 address exhaustion and a structured approach to deploying IPv6. Additional work will be required to refine these estimates as AMI and CT. INS. CORP. move forward with the following recommended phases:

- Phase 1 – Detailed Design and Planning - IPv6 Solutions
- Phase 2 – Test and Verify IPv6 Solutions
- Phase 3 – Deploy the Dual-Stack Target Architecture in the External Zone
- Phase 4 – Enable IPv6 Services (DNS/LDAP/NTP/FTP/NMS/VPN) Capability
- Phase 5 - Deploy IPv6 Translation Capability in the External Zone
- Phase 6 – Enable Dual-Stack on Remaining Secure DMZ Infrastructure/Servers

Phase 1 – Detailed Design and Planning - IPv6 Solutions

There are different areas of the network that require different levels of remediation, but all will require engineering and planning resources to build the detailed designs and specific integration plans. The design and planning phase activities include defining priorities, defining detailed functional requirements, developing device configuration templates, establishing standards, and developing migration plans. The activities below provide a summary of the roadmap steps for this phase:

1. Develop an IPv6 Numbering Standard – refer to section ***IPv6 Planning Resources IPv6 Addressing Scheme***
2. Develop IPv6 proof of concept test criteria for sand box and formal lab
3. Design and plan translation support for external DNS
4. Define and prioritize external applications for translation access
5. Design and plan remote user translation VPN access
6. Design and plan external translation electronic mail capabilities
7. Design and plan immediate external translation application capabilities
8. Design and plan network monitoring, management and reporting

Phase 2 – Test and Verify IPv6 Solutions

The testing in this phase focuses on validating the software and hardware in the IPv6 designs to provide the desired functional capabilities in a controlled environment. Testing results will also include development of configuration templates for the various network and application infrastructure elements. This phase also includes building the sand box and formal lab for general and deployment testing activities. Pilot efforts are designed to introduce the new capabilities in a production environment, exposing the solution to a broader set of conditions that cannot be anticipated or easily replicated in a controlled environment.

A summary of the roadmap steps for this phase includes:

1. Build sandbox and formal lab as recommended in section ***IPv6 Testing and Verification (Lab)***
2. Conduct testing in lab and verify translation behavior across all applications and services
3. Plan pilot program for initial deployment
4. Test and verify pilot translation support for external DNS
5. Test and verify pilot remote user translation VPN access
6. Test and verify pilot external translation electronic mail capabilities
7. Test and verify pilot external translation application capabilities
8. Test network monitoring, management and reporting

Phase 3 – Deploy the Dual-Stack Target State Architecture

The tactical approach of the road map calls for the use of Dual-Stack and Translation to be utilized together in the External Zone. Translation cannot work without some amount of the path from the internet to the translation device being completely Dual-Stacked or bi lingual IPv4/v6. This phase stages future phases for translation deployment by first applying Dual-Stack capabilities in the external zone edge routers. The following diagram and summary of steps involved in completing this phase.

Figure 9 – External Zone Dual-Stack

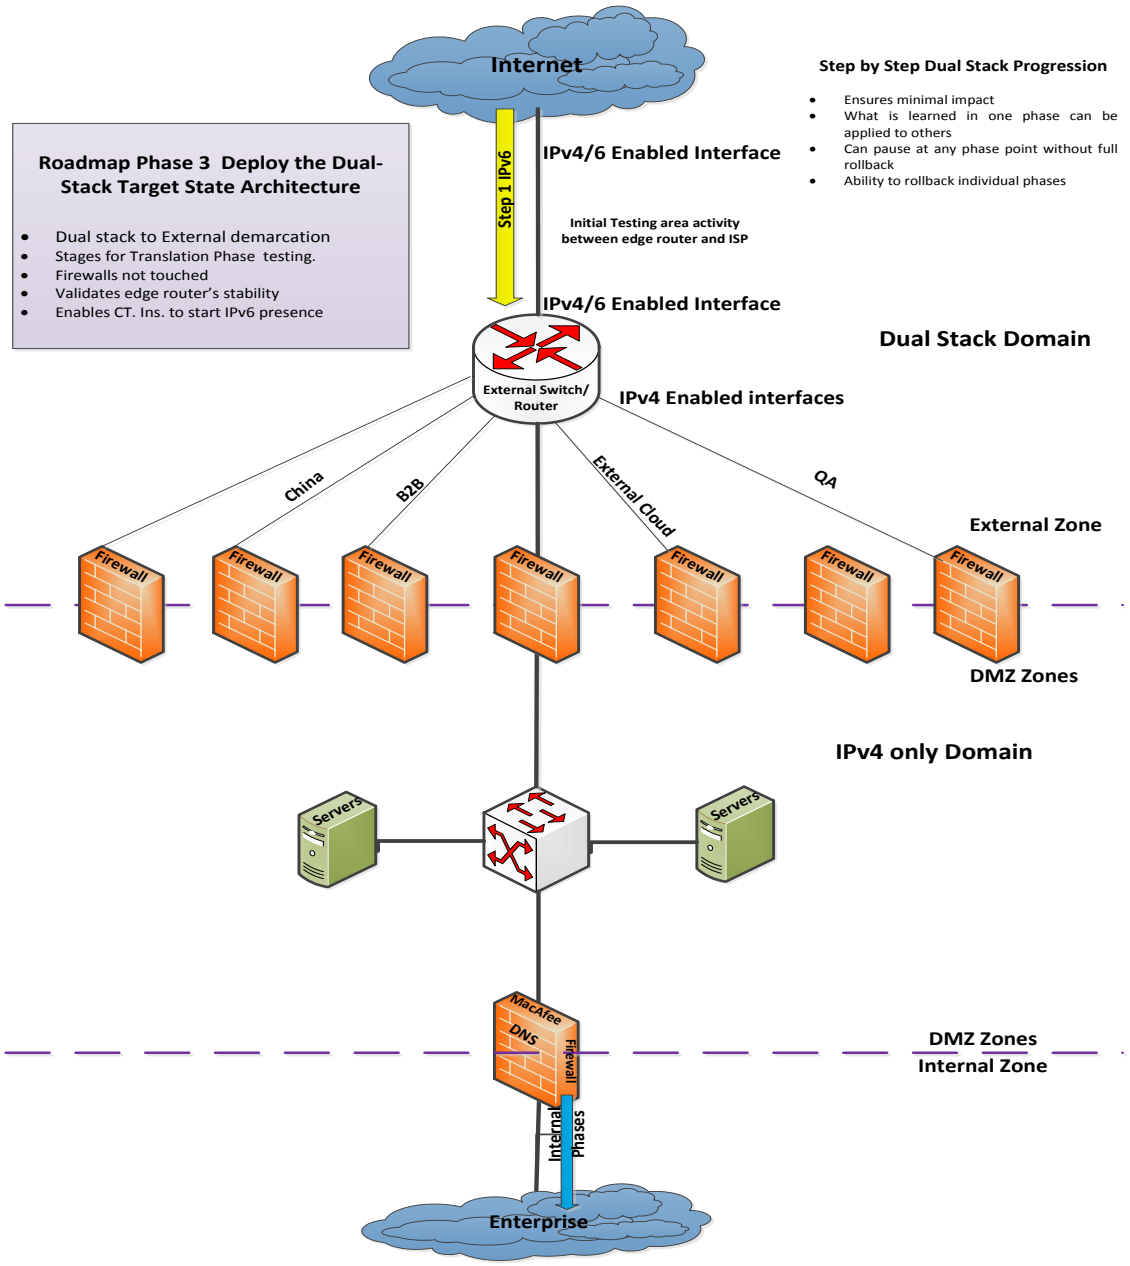

1. Conduct and complete required upgrades or platform refreshes cycles in zone
2. Acquire IPv6 prefix from ARIN if not already completed- see section ***IPv6 Planning Resources - IPv6 Addressing Scheme***
3. Enable IPv6 on just the carrier and external router's IPv6 interface
4. Low impact initial step - CT. Ins. Corp. and ISP can monitor for discovery traffic and check security related statistics
5. CT. Ins. Corp. has PI IPv6 address and ISP is configured on their end and ready for turn up.
6. CT. Ins. Corp. works with ISP for IPv6 interface turn up.
7. Apply IPv6 ACLs to router
8. Review any policies and special configuration requirements ISP may need to have met
9. ISP turns up IPv6 on its end
10. Configure external 6500 interface for IPv6 – pre stage
11. Enable IPv6 on external interface
12. Conduct PtP connectivity and IPv6 related tests, Neighbor discovery/cache etc. with no ACL's active
13. Flip switch on ACLs and rerun connectivity tests - monitor ACLs etc.
14. Provision for default route ::/0
15. Provision and test eBGP AF for IPv6 and advertise prefixes – optional
16. Test IPv4 based NMS and security tools that can glean IPv6 statistics
17. Enable IPv6 on redundant equipment and repeat tests

Once this phase is completed an initial IPv6 POC is established with the ISP and CT. Ins. Corp. becomes connected to the IPv6 Internet.

Phase 4 – Enable IPv6 Services (DNS)

This phase entails planning, testing and staging of DNS to support the tactical approach of utilizing translation. Additional planning and testing is required for the optimal DNS64 architecture is applied for CT. Ins. Corp.'s use. DNS lookup requests are sent to a caching DNS "recursive" that forges (or rather creates) AAAA records for hosts that do not naturally have them configured. That DNS server creates the AAAA records using an IPv6 range which will be configured on the translation device(NAT64) which handles the translation of IPv6 to IPv4. A summary of the process involved is outlined below:

1. The IPv6-only host performs as DNS lookup by triggering a DNS query (AAAA: example.com) to access a service from example.com.
2. The DNS64 translation device receives a DNS AAAA query for resolving example.com.
3. DNS64 triggers an AAAA query to the DNS AAAA authoritative server(McAfee Bind) for the domain being queried. However, because the server has only an A record for example.com an empty AAAA response is returned. Refer to figure 10.

Figure 10 – DNS64 Process

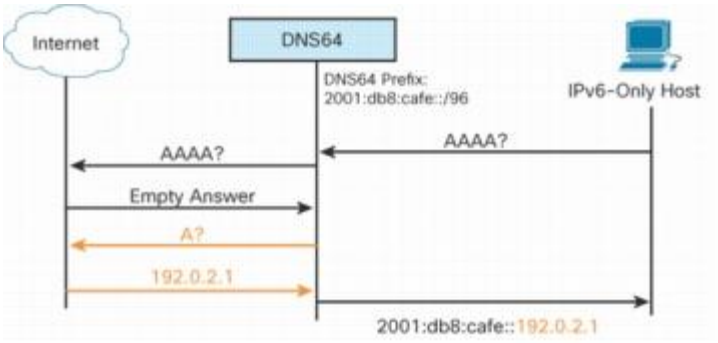

4. On receiving end the empty answer in the response to the AAAA request, DNS64 triggers an A query (A: example.com) to the DNS A authoritative server.
5. DNS64 receives a DNS A record for the service being queried (A: example.com-192.0.2.1).
6. DNS64 synthesizes the AAAA record by prefixing it with the NAT64 prefix, which may be the WKP or an NSP type.

Note: this DNS64 state is maintained in the translations device.

A summary of some of the post planning steps is outlined below.

1. Engage NSI for DNS Staging - Example steps from NSI:

> ***To add an IPv6 AAAA (quad-A) record****, please login to your Network Solutions® account and follow the Advanced User instructions on the Edit DNS page. Your request will be reviewed, authenticated, and processed within 24 to 48 business hours.*
>
> *1. Log into Account Manager.*
> *2. Click on nsWebAddress™ (Domains).*
> *3. Click on Manage Domains.*
> *4. Click on the selected Domain and the Domain Detail page will open.*
> *5. Check the Designated DNS radio button and click on Apply Changes.*
> *6. Click on Move DNS under Change Domain Name Servers.*
> *7. Follow the Advanced User instructions in the box on the Edit DNS page and send an email with the appropriate information to IPV6Req@networksolutions.com*
>
> ***Advanced Users:***
>
> *To specify your IPv6 name server address (IPv6 glue record), e-mail us the domain name, the host name of the name server(s), and their IPv6 address(es).*

2. Engage ISP on IPv6 turn up, prefix delegation and DNS support
3. Setup IPv6 interfaces on BIG-IP LTM
4. Verify IPv4 and IPv6 routing and connectivity
5. BIG-IP LTM Gateway configuration(example from F5)
    a. Setup required IPv6 virtual servers with IPv4 pool members
    b. Test connectivity to new IPv6 addresses
    c. Configure any needed iRules to provide standard application access
    d. Ensure all needed NS AAAA records are configured and recursive IPv6 DNS lookups are functioning
    e. Setup BIG-IP Global Traffic Manager (GTM) Module V6 to V4

6. Conduct connectivity tests up to ISP and beyond from 6500(Mid-SS-L3-External-1)
7. Testing of IPv6 ACL rules
8. Enable IPv6 on redundant equipment and repeat tests
9. GSS Replacement- F5 GTM
10. Enable DNS caching and resolving
11. Enable Application Monitoring and load balancing
12. Enable IPv6 on redundant equipment
13. IPv6 test script is executed to determine impact
14. Monitoring period commences

Below is an example of a simple external workstation side test criteria for Dual-Stack DNS testing in the lab against another vendor's DNS product. This was used to verify what the client sees and also determine if any Happy Eyeballs activity is present. A similar more detailed testing matrix will be required for testing of DNS64 translation behavior.

IPv6 lab Infoblox testing of mixed records for production v6 staging		
Client Type	**Infoblox setting**	**What's Returned**
Windows 7		
V4 client only	v6 Listening off	v6 record with v4 record
		v6 record alone for queries aginst a name with just AAAA record
v4/v6 client	v6 Listening off	v6 record with v4 record
		v6 record alone for queries aginst a name with just AAAA record
V4 client only	v6 off Lan interface	v6 record with v4 record
	can not v6 ping device	v6 record alone for queries aginst a name with just AAAA record
v4/v6 client	v6 off Lan interface	v6 portion of WS stack in Windows 7 overrides v4 and thus lookup
	can not v6 ping device	fails. Keeps trying v6 even though the only DNS enabled on
		WS is v4

Note: Happy Eyeballs RFC - http://tools.ietf.org/html/rfc6555

Phase 5 – Enable Translation Capabilities

Translation Zone 1 deployment:

Approach: This is a tactical, low impact, easy to roll back, approach to ensure that disruptions do not occur to IPv4 traffic and business. Keep in mind what is presented here is just a summary of the details involved. The actual planning and steps of the phases will have additional details from lab testing, upgrade, and planning session results.

Note: this sample approach progresses from the internet in towards the **Secure DMZ** demarcation point and all the components in between. This approach validates internet access first.

Refer to figure 11 on the following page.

Figure 11 – Translation Mechanics and Flow

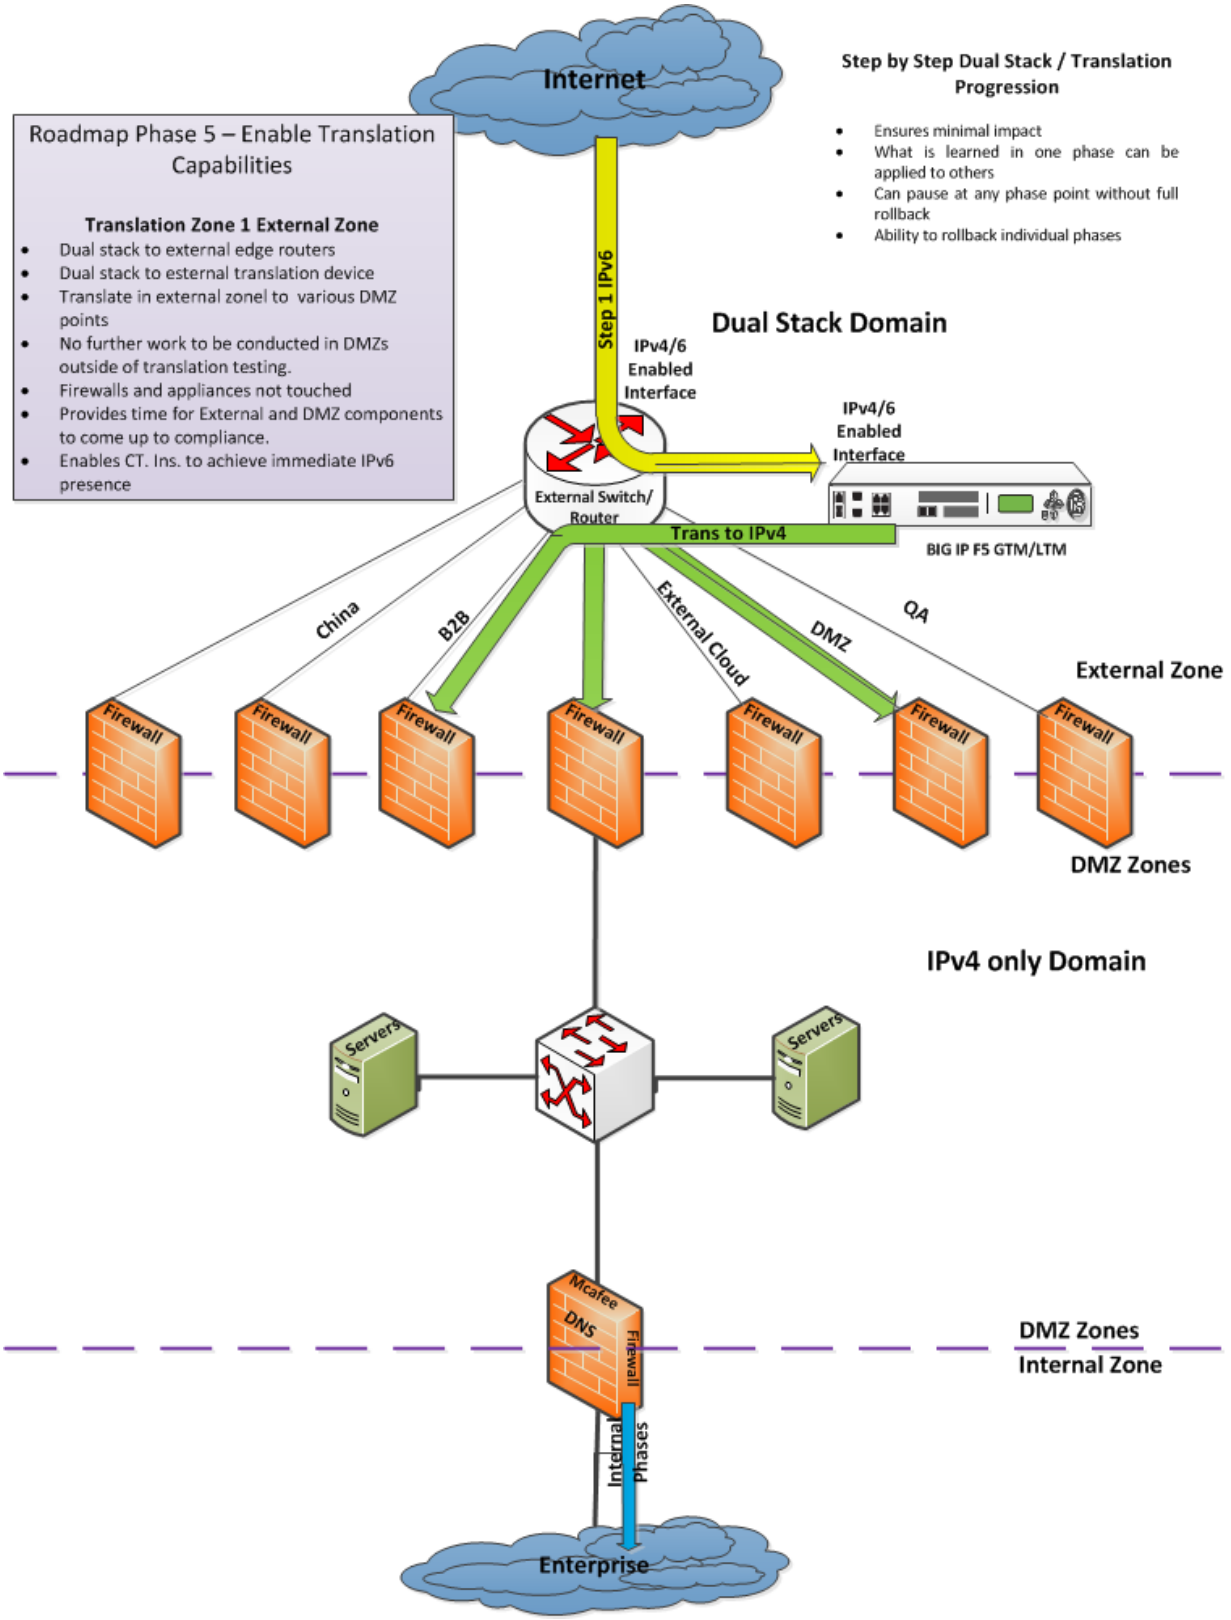

Figure 11 – Translation Zone 1

Translation can be accomplished in one or more points in CT. Ins. Corp. External and DMZ zones. These points are called Translation Zones. Translation Zone 1 encompasses the External Zone of CT. Ins. Corp.'s network. This Zone is the closest to the internet routers and connects directly to the ISP. Deploying a translation device between this zone and the Secure DMZ zones exhibit the following attributes:

1. Simplifies initial Dual-Stack deployment to just the external zone
2. Reduces lab testing for Dual-Stack which can be time consuming and testing for a smaller domain reduces time and cost
3. Reduces impact to enterprise – fewer devices to touch and change
4. Easier to roll back or turn off a change
5. Easier to monitor – IPv6 is only enabled and monitored on a few devices, any anomalies and or security related issues can be quickly identified due to smaller protocol domain
6. No need to touch firewalls and existing equipment in the DMZ zones – saves time and costs
7. Provides CT. Ins. Corp. time to remediate assets in DMZ zones for IPv6 readiness
8. Traffic levels of initial IPv6 for same services will not impact translation device platform capabilities
9. Less secure if completely outside of Firewalls.
10. Translation device sits with Firewalls in external zone, will rely on its own hardened OS and security and edge routers ACLs (Bogons and exploits) for IPv6 protection.
11. Can be separated by dual interfaces or sub interfaces.
12. Provides CT. Ins. Corp. engineers time to test applications and Dual-Stack deployment

However, since the translation devices sits in the External Zone CT. Ins. Corp. must also validate the translation device's security posture.

Example: High Level Steps for Translation 1 deployment

Example: High Level Steps for Translation Zone 1 deployment

- This is the simplest approach for an initial foray into IPv6 to ensure the services and risks identified in the Case Studies section are addressed and supported.
- Enable IPv6 on just the carrier and external router's IPv6 interface
- Low impact initial step - CT. Ins. Corp. and ISP can monitor for discovery traffic and check security related statistics

Assumptions:

1. This is the first contact approach only from one DC, links to other DC, redundancy and routing protocol considerations not fully factored in.
2. All upgrades have been completed in the External zone and its devices are IPv6 compliant
3. Roadmap Phases 3 and 4 are completed
4. Full Lab testing is completed and behavior information is documented
5. IPv6 configuration - addresses, routing, ACLs for Bogons etc. for external routers are already baked – again Phase 3 and 4 completed
6. F5 GTM/LTM integration is in place in DMZ translation area 1
7. F5 initially used just for translation all other services DNS/SLB etc. still use existing IPv4 devices
8. DNS and NSI is staged for IPv6 but not active
9. DMZ zones will not be fully Dual-Stack.
10. For initial Web application testing, VPN, B2B, PerfQa will follow similar process.

11. Steps can be reversed where IPv6 is turned up towards carrier and last step is OC interface turn up.

Configuration Steps

1. Phase 3 is already completed and CT. Ins. Corp. has PI IPv6 address and ready for turn up or already up.
2. Enable DNS caching and resolving
3. Enable Virtual servers and addressing
4. Setup IP sharing
5. Enable Application Monitoring and load balancing
6. Enable F5 for NAT64 services and execute testing script
7. Enable F5 for DNS64 services and execute testing script
1. Monitor
2. Enable on redundant equipment
3. Conduct testing from outside IPv6 users - results should match lab

Translation Zone 2 deployment:

Approach: This is a moderate impact, easy to roll back, approach to ensure that disruptions do not occur to IPv4 traffic and business. Keep in mind what is presented here is just a summary of some of the details. The actual planning and steps of the phases will have additional details from lab testing, upgrade, and planning session results. This option will need to be implemented in each separate DMZ across the CT. Ins. Corp. enterprise since they all utilize separate ASA firewall interfaces. Although the recommended translation device might support multiple interfaces, this would need to be reviewed with the CT. Ins. Corp. security requirements to determine the number of translation points/devices that would be required.

Note: this sample approach progresses from the internet in through the **Secure DMZ** demarcation point. This approach validates internet access first.

Refer to figure 12 on the following page.

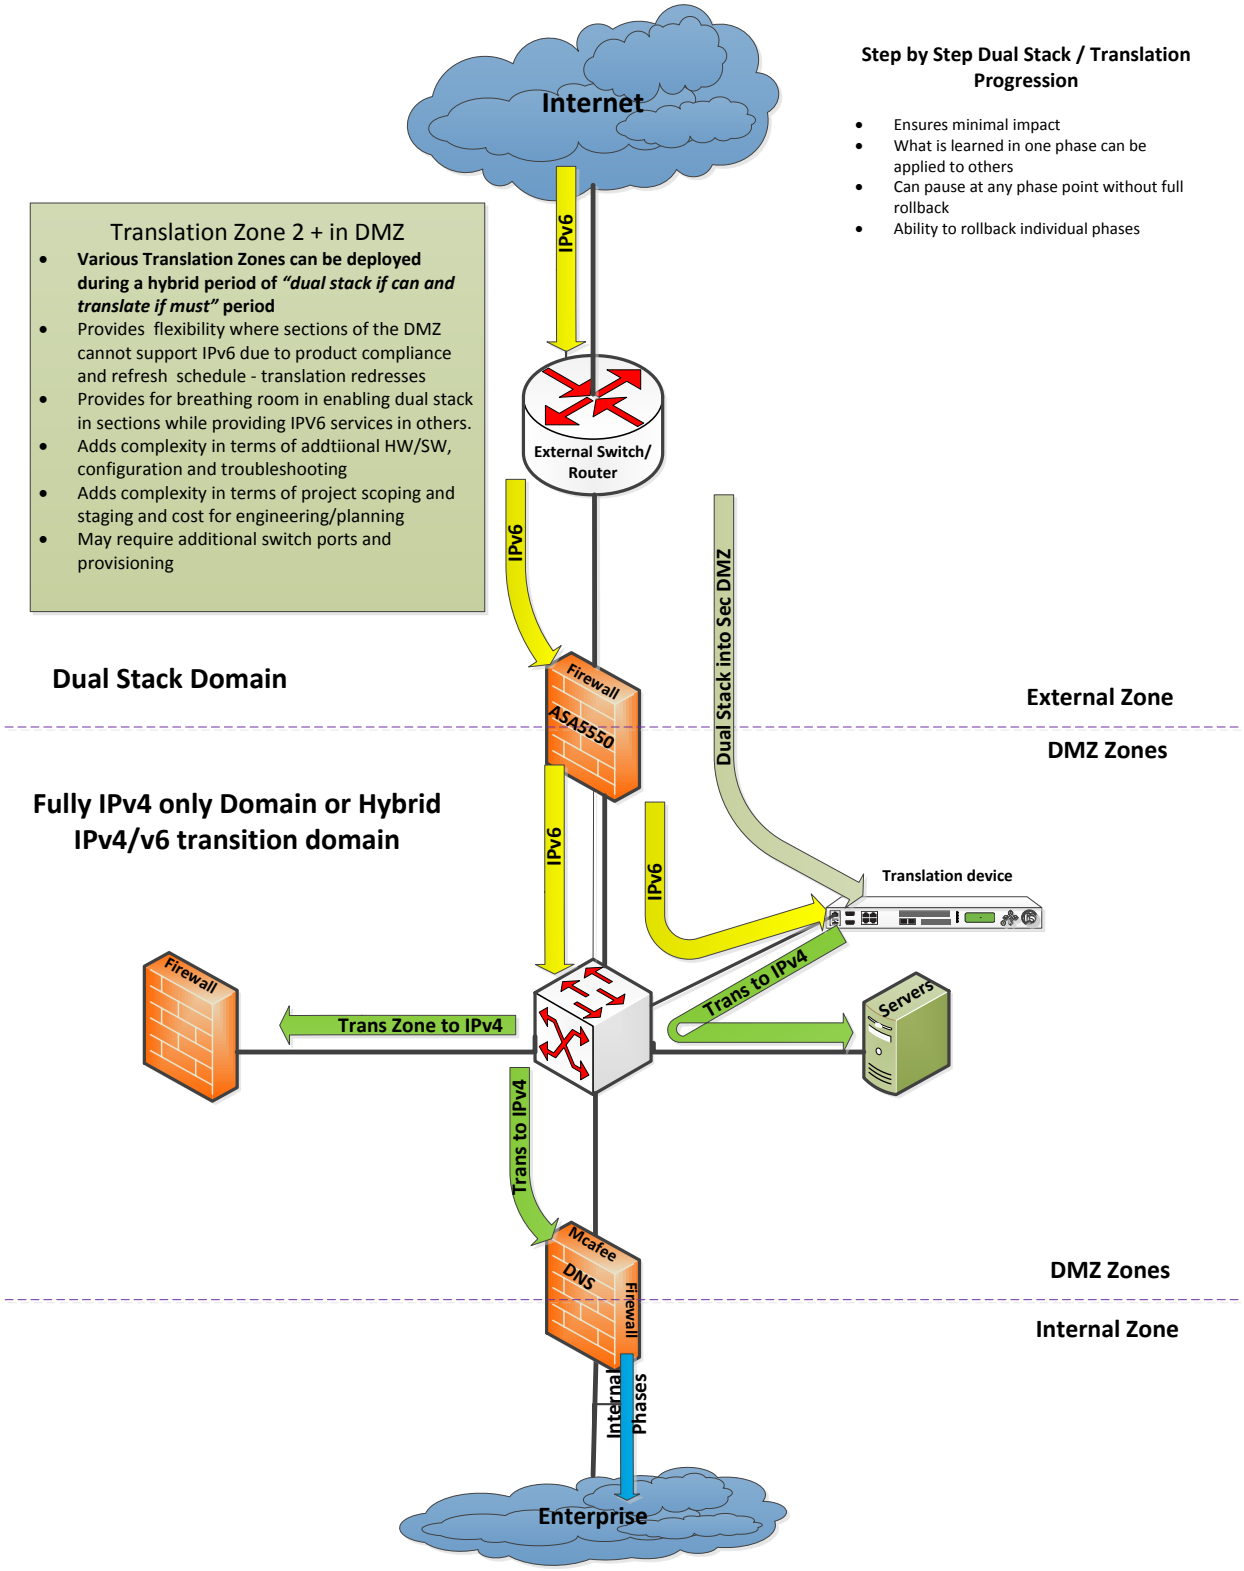

Figure 12 – Translation Zone 2

Translation Zone 2 is in the Secure DMZ zone of CT. Ins. Corp.'s network. This Zone is the closest to servers and resources internet user's access. Deploying a translation device in this zone exhibit the following attributes:

1. Progresses Dual-Stack deployment to into a section of the DMZ
2. Lab testing for further Dual-Stack which can be time consuming and testing for a smaller domain reduces time and cost
3. Increases impact to enterprise – additional devices to touch and change
4. Additional rollback points are available
5. Various Translation Zones can be deployed during a hybrid period of *"Dual-Stack if can and translate if must"* period
6. Provides flexibility where sections of the DMZ cannot support IPv6 due to product compliance and refresh schedule - translation redresses
7. Provides for breathing room in enabling Dual-Stack in sections while providing IPV6 services in others. IPv6 can be enabled on servers heading outward towards translator etc.
8. Adds initial complexity in terms of additional HW/SW, configuration and troubleshooting
9. Adds initial complexity in terms of project scoping and staging and cost for engineering/planning
10. Secure, behind firewalls and additional monitoring is present monitored
11. May require additional switch ports and provisioning
12. A little more to monitor – IPv6 is only enabled and monitored on a few additional devices any anomalies and or security related issues can be identified but will now include External and DMZ zone consideration.
13. Traffic levels of initial IPv6 for same services will not impact translation device platform capabilities
14. Still provides breathing room for CT. Ins. Corp. to progress with Dual-Stack planning and engineering activities in the DMZ
15. Provides ability to reduce translation footprint as sections of DMZ zones become Dual-Stack.

This is another approach for an initial foray into IPv6 to ensure the services and risks identified in the Case Studies are addressed and supported by IPv6. However, since the translation devices sits outside of the DMZ it is CT. Ins. Corp.'s decision as to utilize this approach based on its security posture.

Example: High Level Steps for Translation Zone 2 deployment

Assumptions:

1. This is the first contact approach only from one DC, links to other DC, redundancy and routing protocol considerations not fully factored in.
2. All upgrades have been completed in all zones IPv6 is fully supported in the DMZ zones just not enabled.
3. Lab testing is completed and behavior information is documented
4. IPv6 configuration - addresses, routing, ACLs for Bogons etc. for external routers are already baked
5. F5 GTM/LTM integration is in place in DMZ translation area 2
6. F5 initially used just for translation all other services DNS/SLB etc still use existing IPv4 devices
7. DNS and NSI is staged for IPv6 but not active

8. DMZ zones will not be fully Dual-Stack.
9. For initial Web application testing, VPN, B2B, PerfQa will follow similar process.
10. Steps can be reversed where IPv6 is turned up towards carrier and last step is OC interface turn up.

Configuration Steps

1. Phases 3 and 4 completed
2. Configure Midfire5 FW inside interface for IPv6
3. Enable DNS caching and resolving
4. Enable Virtual servers and addressing
5. Setup IP sharing
6. Enable Application Monitoring and load balancing
7. Enable F5 for NAT64 services and execute testing script
8. Enable F5 for DNS64 services and execute testing script
9. Monitor
10. Enable on redundant equipment
11. Conduct testing from outside IPv6 users - results should match lab
12. Monitor IPv6 traffic across Mid 5 or other firewalls.
13. Enable on redundant equipment
14. Conduct testing from outside IPv6 users - results should match lab

Note: At this point the tactical approach has extended into a DMZ zone but so has Dual-Stack. CT. Ins. Corp. can continue to Dual-Stack devices in these zones to fulfill its strategic goals.

Translation Zone Flexibility:

Translation HUB
If the security requirement is to keep translation behind all firewalls then multiple Translation Zones of type 2 may be required. This adds to the CAPEX cost of additional translation devices possibly required. Additional research should be conducted to determine if a "translation zone consolidation point" or "hub" can be engineered for IPv6 traffic among the different DMZ zones to a common pair of translation devices.

Translation Firewall
One other option for **Translation Zone 1** is to utilize just an IPv6 only Firewall. For example, from the external router carve out the IPv6 traffic either using a sub or actual physical interface and pass it through its own IPv6 only FW then onward towards Translation device. This approach only adds one other Dual-Stack device in the External Zone, the IPv6 only FW. The traffic is separated for easier support and operations perspective uses and is secure. Any immediate IPv6 related attacks or exploit will be easier to identify and police or shut down without affecting the IPv4 traffic. The IPv6 only FW can have an IPv4 out of band connection for management.

After recent discussion with CT. Ins. Corp. Sr. Network Engineer, CT. Ins. Corp. is interested in utilizing the Translation Firewall concept with the F5. The firewall considered are the new Cisco ASA 55000-X series. The protocol translations will still occur at the F5 due to its history and features supporting IPv6. The ASA will operate in possibly a "carved out" IPv6

only connection to the F5 as noted above. The ASAs are part of the refresh cycle thus assisting CT. Ins. Corp.'s CAPEX goals.

Dual-Stack Progression

Dual-Stack Only – AMI provided these summary steps to show CT. Ins. Corp. a glimpse of the steps involved for planning considerations.

Approach example: This is a moderate to high impact, easy to roll back, approach to ensure that disruptions do not occur to IPv4 traffic and business. Keep in mind what is presented here is just a summary of the details involved. The actual planning and steps of the phases will have additional details from lab testing, upgrade, and planning session results.

A second even more conservative approach is to just execute the phases in reverse. However, some of the testing and execution dependencies will change. This approach tests those appliances, applications (web servers) etc. all work using IPv6 before as the phase progress up towards the internet first contact. This approach ensures confidence that as CT. Ins. Corp. gets to the point of turning on the first IPv6 "spigot" it knows its web sites speak native IPv6 for troubleshooting purposes.

Refer to figure 13 on the following page.

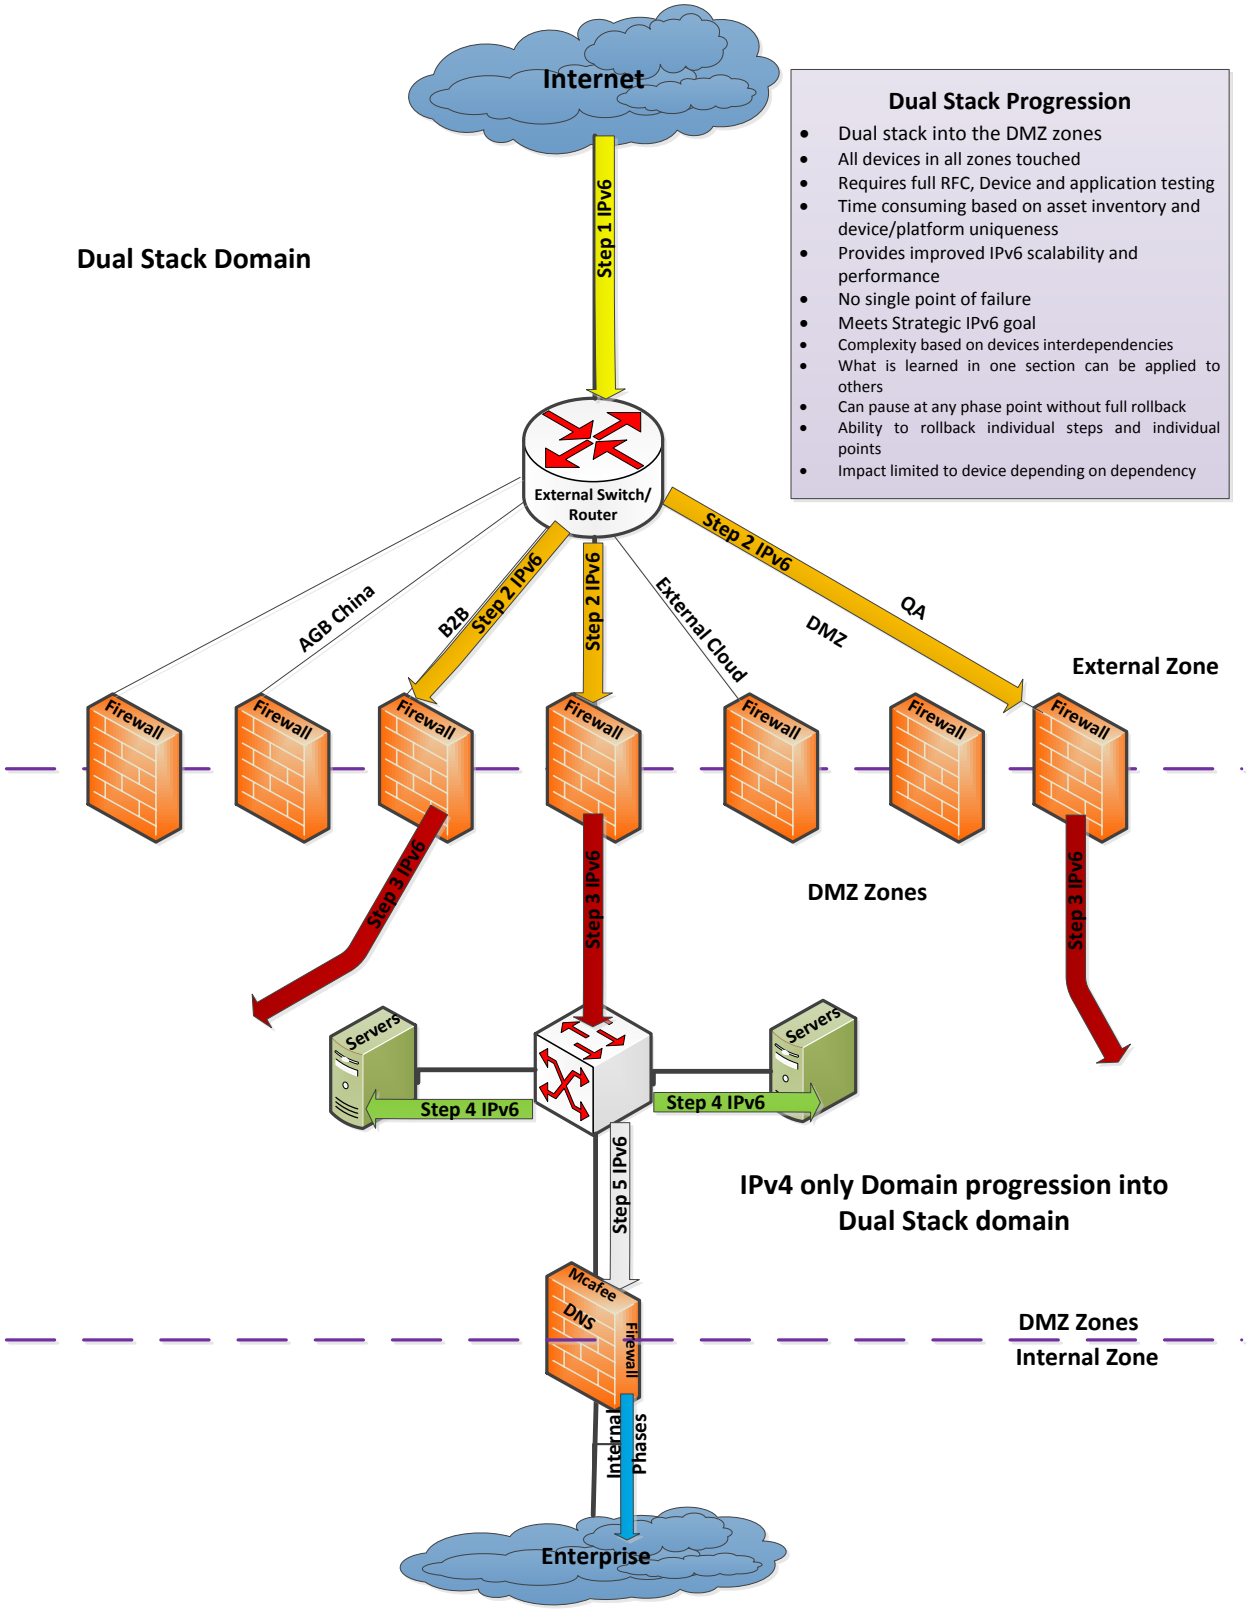

Figure 13 – Dual-Stack Progression

Assumptions:

1. CT. Ins. Corp. has PI IPv6 address and ISP is configured on their end and ready for turn up.
2. All upgrades have been completed in all zones IPv6 is fully supported
3. Lab testing is completed and behavior information is documented
4. IPv6 configuration - addresses, routing, ACLs for Bogons etc. for external routers are already baked
5. DNS and NSI is staged for IPv6 but not active
6. Email will not be supported unless IPv6 is brought all the way into the enterprise
7. No translation unless just for Email

Step 1 First Contact

1. Enabling IPv6 on just the carrier and external router's IPv6 interface
2. Low impact initial step. CT. Ins. Corp. and ISP can monitor for discovery traffic and check security related statistics
3. CT. Ins. Corp. works with ISP for IPv6 interface turn up.
4. Review any policies and special configuration requirements ISP may need to have met
5. ISP turns up IPv6 on its end
6. Configure external 6500(Mid-SS-L3-External-1) interface for IPv6 – pre stage
7. Enable IPv6 on external interface
8. Conduct tests and note neighbor discovery
9. Monitors ACLs etc.
10. Enable IPv6 on redundant equipment

Step 2 Heading in towards DMZ Zones

1. Configure external 6500(Mid-SS-L3-External-1) interfaces or SVI towards Mid x Firewalls for IPv6
2. Configures Midfire x outside interfaces for IPv6
3. Configure IPv6 routing
4. Conduct connectivity tests up to ISP and beyond from 6500
5. Testing of FW IPv6 rules
6. Conduct tests to other IPv6 devices and note neighbor discovery
7. Monitors ACLs etc.
8. Enable IPv6 on redundant equipment

Step 3 Enabling IPv6 inside DMZ zones

1. DNS IPv6 staging and execution tasks on FW
2. Configure Midfire x FW inside interfaces
3. Configure Midfire5 FW inside interface
4. Configure Secure DMZ 6500(Mid-SS-L3-Secure-1) interfaces or SVI towards Mid 5 Firewall for IPv6
5. Conduct tests and note neighbor discovery
6. Monitors ACLs etc.
7. Enable IPv6 on redundant equipment

Step 4 Enabling IPv6 for all devices in the Secure DMZ zones

1. Enable IPv6 on outside firewall interfaces facing external zone
2. Enable IPv6 on inside firewall interfaces facing DMZ zones
3. Enable IPv6 on individual appliances, components and servers in the respective DMZ Zones(VPN, PerfQa, B2B, Cloud etc.)
4. Conduct tests and note neighbor discovery
5. Monitors ACLs etc.
6. Enable IPv6 on redundant equipment

Step 5 Enabling IPv6 at demarcation Surf Zone point.

1. Enable IPv6 on firewall interfaces connecting Surf to DMZ zones only
2. At this point IPv6 is Dual-Stacked across External Zone and both sides of Secure DMZ firewalls and other DMZ zones
3. Configure additional Secure DMZ devices one at a time for IPv6 based on lab tested script and sequencing.
4. Conduct testing of Secure DMZ and its zones for IPv6 any to any connectivity between devices for basic IPv6 address only connectivity.

Step 6 Enabling DNS and utility services

1. DNS is enabled on Midfire5 for IPv6
2. IPv6 test script is executed to determine impact
3. Monitoring period commences

IPv6 Planning Resources

IPv6 Addressing Scheme

IPv6 provides a tremendous amount of address space. One of the IETF's original goals for IPv6 was to achieve what did not happen in IPv4. True open end to end connectivity between any devices around the globe. This means there is enough address space for end user devices to communicate directly between each other using globally routed addresses, securely (IPsec), without any intermediary devices in-between them(NAT).

This section describes the IPv6 address scheme reviewed with CT. Ins. Corp.. CT. Ins. Corp. currently does not own an IPv6 address. The initial scope of the schema is to cover addressing requirements for the external and DMZ zones of CT. Ins. Corp.'s US network with the flexibility to address the internal, national remote sites and global regions outside of its RIR in the future.

CT. INS. CORP. is interested in allocating a globally unique address (GLA) IPv6 address from the North American Region RIR (Regional Internet Registry) The 2600::/24(Provider Independent) range is allocated to Enterprises that require multi-homing of their address space from multiple carriers.

PI address space refers to IP addresses a customer receives directly from a RIR for their own use; it is not allocated to an ISP or NSP. Customers may advertise that IP address space to their ISP, but ISPs may only announce the aggregate /48 IPv6 block to its customers and or Peers. A /48 IPv6 address block refers to a range of IPv6 addresses where the first 48 bits of the IPv6 address block are masked or fixed. The customer must advertise an aggregate /48 or a less specific address block (e.g., /47, /46, /32, etc.).

This address space differs from the 2001::/ Provider Aggregately (PA) IPv6 Address Space: PA address space refers to IP addresses a customer is allocated from a specific ISP or NSP. A customer location connected to two ISPs or NSPs is often referred to as multi-homed. Multi-homed customers using ISP provided addresses may advertise PA IP address space to other ISPs, but the other ISP may only announce the aggregate /48 or less specific (e.g., /47, /46) IPv6 block to its customers or Peers. A customer may advertise more specific (smaller) PA IPv6 address blocks to the ISP but again the ISP may not advertise the PA IPv6 address blocks smaller than /48 to its customer or peers. Also, if the customer decides to change ISPs there may be a need to renumber.

Note: The addressing schema's presented here are based with the strategic Dual-Stack architecture in mind but also apply and are adaptable to the tactical translation architecture.

CT. Ins. Corp. is first considering to apply for a /32.

All IP addresses (IPv4 and IPv6) are allocated in large blocks to a regional registry (ARIN, RIPE, etc…) these registries then in turn allocate addresses down to ISPs, ISPs to customers and so on. By default the registry will allocate a /32 to an ISP and the ISP is supposed to allocate a /48 to a customer. The IPv6 equivalent of an IPv4 /24 subnet is a /64 which is the space a customer should break their network segments into.

IPv6 /32s are assigned to ISPs and large institutions that provide the proper justification as per ARIN's ***ARIN Number Resource Policy Manual sections 6.6.1.1 and 6.6.1.2*** . They typically break them in 2^16 (65,526) /48 subnets, which they in turn assign. Each /48 subnet can then be split in 2^16 (65,536) /64s. Each /64 has 2^64 possible end addresses.

6.5.8.2.2. Extra-large sites

In rare cases, an organization may request more than a /48 for an extra-large site which requires more than 16,384 /64 subnets. In such a case, a detailed subnet plan must be submitted for each extra-large site in an organization's network. An extra-large site qualifies for the next larger prefix when the total subnet utilization exceeds 25%. Each extra-large site will be counted as an equivalent number of /48 standard sites.

The numbers are unbelievably large. CT. Ins. Corp. can properly subnet their /32 space to sites such as the following as an example only – future stages of the deployment will provide the time and resources to plan and design a detailed schema which should include considerations of an IPAM system to be integrated as well.

Figure 14 – CT. Ins. Corp. ARIN /32 example

Example: 2601:db8::/32

/64 networks = 2^{32} = 4,294,967,296

Hosts = 2^{96} = 79,228,162,514,264,337,593,543,950,336

Hosts per /32 network - 18,446,744,073,709,551,616

ARIN - /32	Subnets	Remaining /64 for hosts
2601:db8:	0000:0000:	0000:0000:0000:0001 or ::1
2601:db8:	0000:0000	Entire /32 or carved further to /48 for loopbacks or utility use – tremendous size.
2601:db8:	0000:0001	Loc 1 Data Center overlay
2601:db8:	0000:0002	Loc 1 Data Center overlay
2601:db8:	0000:0003	Loc 2 Data Center overlay

The /32 can then be further assigned for additional granularity and use.

The possible benefits with a /32 are a tremendous address space for CT. Ins. Corp. to utilize essentially providing CT. Ins. Corp. ISP like addressing capabilities to handle growth across a spectrum uses:

- Mobile IPv6
- 802.11 Wireless granularity for nodes, location services, IDS/IPs, asset tracking of any and all CT. Ins. Corp. assets(we are talking RFID box level)
- Multiple addresses per device
- Large Flat Networks for DC Fabrics
- VPN/Partner Extranet segmentation
- M&A use allocation for example for all M&As they get a /48 from CT. Ins. Corp. as if they were assigned a /48 by ARIN.
- Aggregation options for BGP – Advertise just the /32 or several top/48s(as if an ISP) to the world only to reach internal /48s. Keeps external edge policies simple and routing tables smaller. Conceals amount of networks and visibility to outside world. This is also dependent upon how much of CT. Ins. Corp.'s network will be open for GLA for true IPv6 end to end communication. But the option is there in the future.

Some possible issues with a /32 prefix:

- Ability to obtain one initially, CT. Ins. Corp. must provide ample justification to receive a /32 and the process may take longer than a typical /48 request. Refer to 6.5.8. Direct assignments from ARIN to end-user organizations

  6.5.8.1. Initial Assignment Criteria

  Sub section 6.5.8.2.2. Extra-large sites

- In rare cases, an organization may request more than a /48 for an extra-large site which requires more than 16,384 /64 subnets. In such a case, a detailed subnet plan must be submitted for each extra-large site in an organization's network. An extra-large site qualifies for the next larger prefix when the total subnet utilization exceeds 25%. Each extra-large site will be counted as an equivalent number of /48 standard sites.
- Issues with ISP peers accepting the /32 and aggregating to /48. Route visibility between peers and RIRs may require periodic oversight and policy changes as other ISP and RIR work towards supporting handling of PI based assignments.
- The management of /48s can become unyielding and if /40 and other nibble boundary or non-nibble based boundary used up to /64 it can become difficult to manage, understand and troubleshoot.

/36 /40 and /44 prefixes and options for their use can be reviewed in future stages of the IPv6 deployment project.

However, if CT. Ins. Corp. cannot garner a /32 they do wish to acquire 2 /48s if possible.

A single /48 still provides a tremendous amount of addressing space and the same benefits achieved with a /32 outlined earlier can be accomplished with /48 outside of further aggregation.

The numbers are just as unbelievably large for the /48. CT. Ins. Corp. can properly subnet their /48 space to sites such as the following as an example only – future stages of the deployment will provide the time and resources to plan and design a detailed schema which would include an IPAM system to be considered as well.

Figure 15 – CT. Ins. Corp. ARIN /48 example

Example 2601:db8:1000::/48

/64 networks = 2^{16} = 65,536

Hosts = 2^{80} = 1,208,925,819,614,629,174,706,176 Total

Hosts per /48 network - 18,446,744,073,709,551,616

ARIN - /48	Subnets	Remaining /64 for hosts
2601:db8:1000:	0000:	0000:0000:0000:0001 or ::1
2601:db8:1000:	0000	Entire /32 or carved further to /48 for loopbacks or utility use – tremendous size.
2601:db8:1000:	0001	Loc 1 Data Center overlay
2601:db8:1000:	0002	Loc 1 Data Center overlay
2601:db8:1000:	0003	Loc 2 Data Center overlay

It is easier to obtain and is the standard accepted prefix by ISPs around the globe. A /48 is considered basically RIR independent. ARIN section 6.5.8.2. Initial assignment size will grant a /48.

However, ISPs may accept a longer prefix such as a /56 and Carrier will accept a shorter (/40, /32) but that prefix may be changed or dropped when it is handed off to another carrier. Due to this uncertainty, **AMI recommends** an IPv6 addressing plan based on the known policies in place today as well as the recommend practice from ARIN **which is /48.**

Additional sizing for subnets for building, sites and P2P links can be further allocated utilizing

/50 /54 and /60 bit prefixes and these options can be reviewed in future stages of the IPv6 deployment project.

Figure 16 – CT. Ins. Corp. /48 to IPv4 manual mapping example

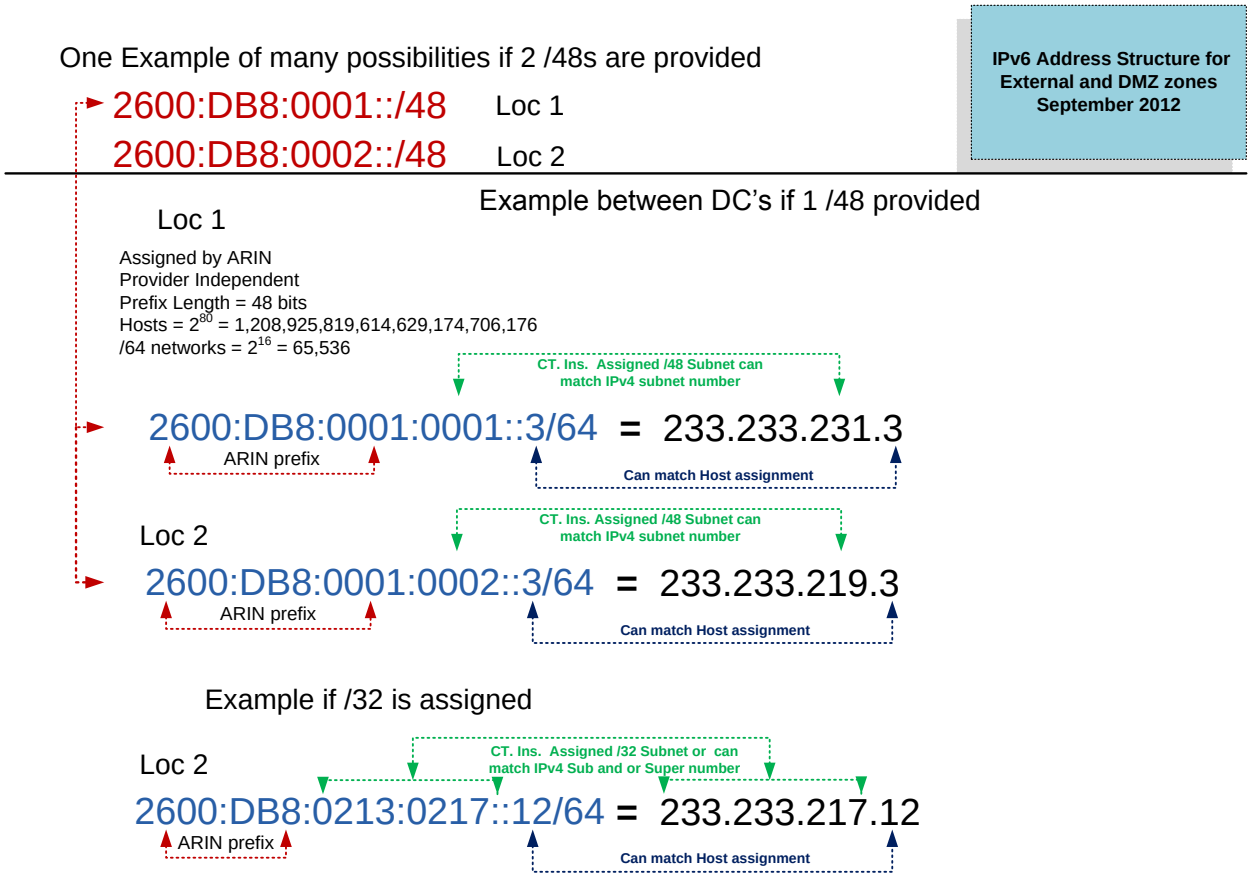

Translation Solutions and IPv6 addressing

The Addressing schemas and options presented here also work with Address Family Translation(AFT) and Using Stateful NAT64. The **Network-specific prefix (NSP)** is the IPv6 prefix assigned by an organization for use in algorithmic mapping between address families; it is usually carved out of the organization prefix and can be globally routable: for example, 2001:db8:cafe::/96 carved out of organization prefix 2001:db8:cafe::/48

Additional planning will be required to map out the ::/96 pools for the translation zones.

IPv6 Link Local Addressing Policy

AMI CT. Ins. Corp. IPv6 Link Local use policy for External and DMZ internet facing network zones

CT. INS. CORP. can have Link Local addresses statically assigned with their own addresses in External and DMZ areas for easier identification of CT. INS. CORP. IPv6 systems for audit and management purposes. This assignment is akin to the classic LAA type addresses assigned to network adapters. This approach helps CT. Ins. Corp. identify new IPv6 systems that may appear on any zone with a Link Local address automatically configured and not using the CT. Ins. Corp. static configuration and for troubleshooting purposes. Reading esoteric link locals to device in troubleshooting sessions can be difficult plus using some easily identifiable addresses improves the productivity of the personnel. There is no security risk since the addresses are not routable. The identifiers that will make up the CT. Ins. Corp. Link Local address will occur in the last 64 bits of the address.

Link-Local addresses are for use on a single link. Link-Local addresses have the following format:

Figure 167– Link Local Address Structure

bits	10	54	64
field	*prefix*	*zeroes*	*interface identifier*

The *prefix* field contains the binary value 1111111010 or FE80 hex. The 54 zeroes that follow make the total network prefix the same for all link-local addresses, rendering them non-routable.

As per RFC 4291 the Link-Local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration, neighbor discovery, or when no routers are present.
Routers must not forward any packets with Link-Local source or destination addresses to other links.
The interface identifier of the Link Local will be used and split into sections for identification use.

For example **FE80::AAA:BB:3**

FE80:: prefix and 54 leading zeros – per RFC 4291

The **:AAA** in interface identifier section will be akin to the OUI and denotes the data center

The **:BB** will denote the Zone for example :AA = External - :BB can be DMZ et. Al.

The **:3** will denote the node number in the DMZ zone and can match the v4/v6 Global address if need be – still not routable.

To maintain such an addressing nomenclature for the benefit noted above CT. Ins. Corp. must ensure that this policy is included in the same process as that of assigning an IPv4 address for a new or changed device in the DMZ or Screened network segments.

Please refer to Figure 18 on the following page.

Figure 18 – CT. Ins. Corp. Link Local Assignent Example

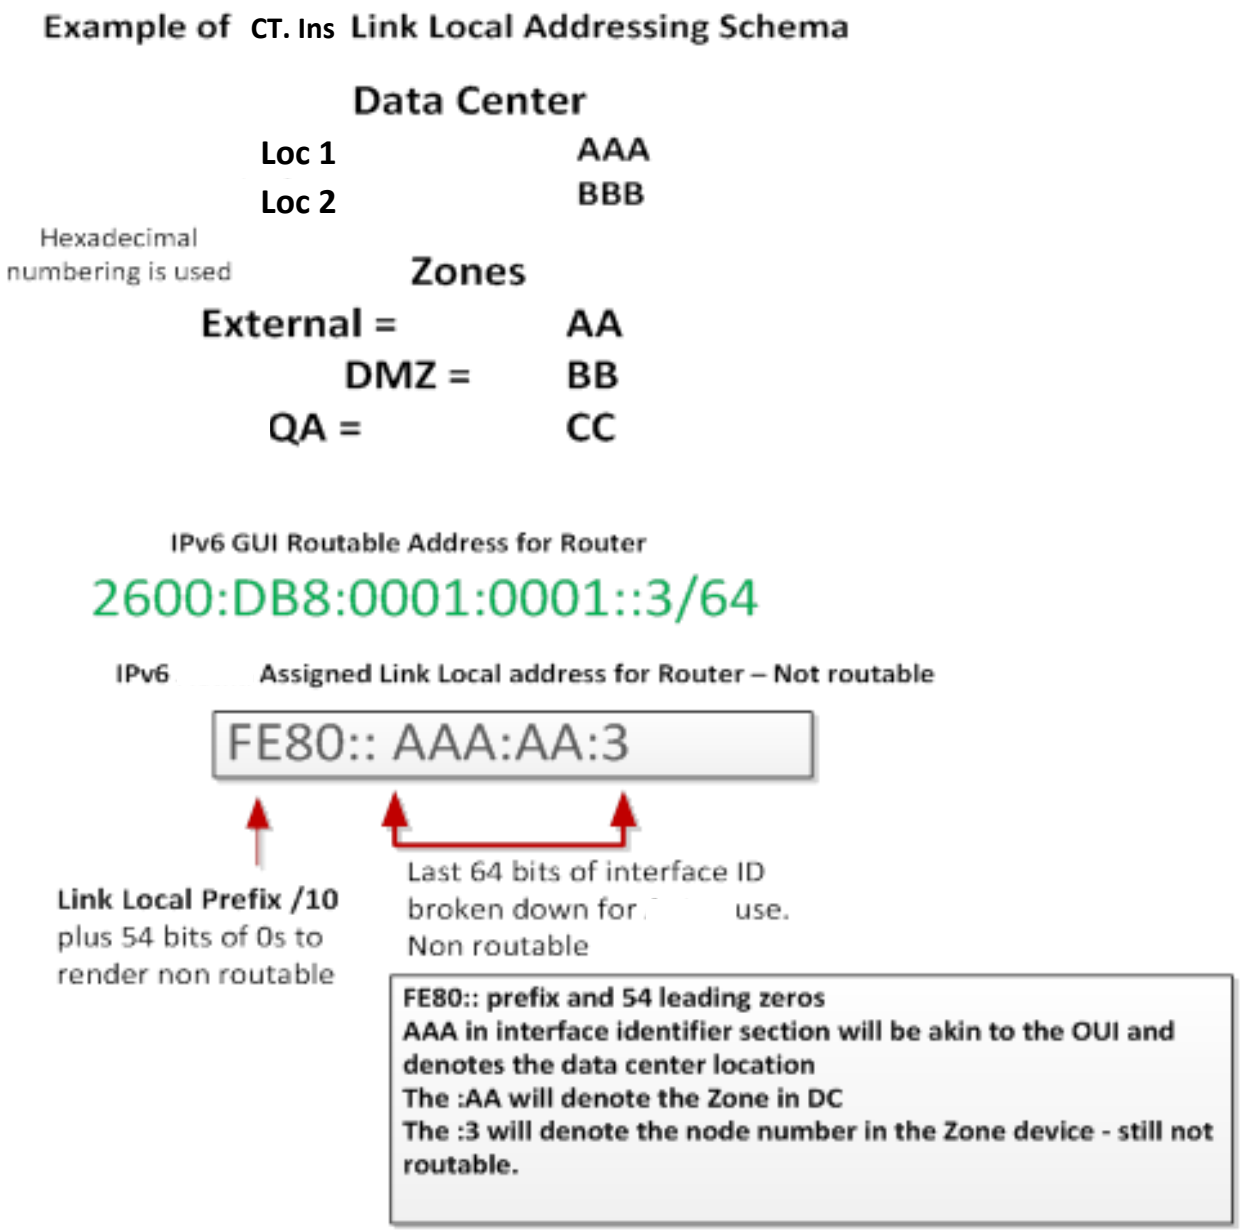

From the readiness analysis Carrier can handle the following Prefix assignments options:

Table 13 – Carrier IPv6 Prefix options

What is the minimum IPv6 prefix announcement that is allowed via the Internet peering? (/30, 32, /36, /40, /44 or /48)		Carrier will advertise subnets to peers only if they are /48 or larger blocks (/48, /40, etc.).
Do you accept prefix announcements larger than a /48 such as a /49, /50 etc.?		Carrier will accept customer route announcements for use in our core between customer locations of variable subnet length up to /64, however, Carrier will advertise subnets to peers only if they are /48 or larger blocks (/48, /40, etc.).
Do you filter out prefix announcements larger than a /48 from other carriers?		No. We will accept blocks larger than /48. (/40, /36, etc.)
Will you accept an IPv6 prefix assigned from any RIR (such as ARIN or RIPE) or is it restricted to the RIR within the region where the internet point of presence is located?		Carrier will accept PI prefixes from any RIR in all of the global regions we operate. However the RIRs are encouraging the use of regional PI prefixes and not across regions. Therefore when this policy becomes more mature, Carrier may change its policy to reflect the new standard.
Do your downstream peering partners accept provider independent IPv6 announcements from all RIR's?		Yes, at this time.
Are there known filters for prefix sizes?		Minimum of /48
What is the minimum prefix announcement allowed?		/48

From the readiness analysis Verizon can handle the following Prefix assignments options:

Table 14 – Verizon IPv6 Prefix Options

What is the minimum IPv6 prefix announcement that is allowed via the Internet peering? (/30, 32, /36, /40, /44 or /48)		**Awaiting outstanding Answer from Verizon** **Assumed industry practice of /48**	
Do you accept prefix announcements larger than a /48 such as a /49, /50 etc.?		Verizon will accept IPv6 announcements all the way to /128. Prefixes longer than /48 will not be advertised to other regions, customers, or Peers.	
Do you filter out prefix announcements larger than a /48 from other carriers?		Yes. This is the current industry best practice.	
Will you accept an IPv6 prefix assigned from any RIR (such as ARIN or RIPE) or is it restricted to the RIR within the region where the internet point of presence is located?		**Yes. We have no advertising restriction based on which RIR the IP space originated from.**	
Do your downstream peering partners accept provider independent IPv6 announcements from all RIR's?		**Yes, to the best of our knowledge.**	
Are there known filters for prefix sizes?		**Yes.**	
What is the minimum prefix announcement allowed?		**The current industry best practice is /48.**	

A complete list of IPv6 options and services the carriers provide can be found in section ***Carriers and Third Party Vendors.***

Transport/Circuits the Asset database.

From the initial review of what the carriers supplied to AMI the parity of services is adequate for use with /48 prefix assignments. Verizon filters out prefixes larger than /48 thus in this case the /32, if CT. Ins. Corp. obtains one, will have to follow-up with Verizon on any additional routing policies that the carrier will need to put into effect to support the /32.

The requirement for summarization of IP addresses and advertisement of the address space will be addressed by CT. Ins. Corp. in future planning stages of the IPv6 rollout.

AMI can further work with CT. Ins. Corp. on the detail addressing plans first for the lab and then for the external deployment phases.

An excellent Best Current Operational Practice (BCOP) paper from Chris Grundemann should be reviewed and followed when planning for expansion. http://www.ipbcop.org/wp-content/uploads/2012/02/BCOP-IPv6_Subnetting.pdf

IPv6 Testing and Verification (Lab)

An IPv6 lab is a critical component to ensuring that IPv6 is integrated in a non-disruptive fashion. Outside of a lab, the only option would be to pilot IPv6 capabilities on the production network. CT. Ins. Corp. currently has a staging lab located in their 151 Farmington Ave. Atrium facility. The lab equipment at this facility is expressly used to support production related testing and the lab environment needs to remain as close to production as possible for testing code upgrades and security vulnerabilities. For IPv6 testing, equipment will be borrowed to build an IPv6 only lab that mimics the production environment.

Approach and Recommendations

CT. Ins. Corp. expressed that its initial approach to IPv6 testing is to use any equipment they have available and just build a basic lab for learning and understanding of the protocol's behavior and operation.

This lab is considered a ***sandbox*** where different than production equipment can be connected, that is all that is available, and IPv6 protocol and configuration testing can occur in a less formal atmosphere against the HW and SW in the lab.

If the project schedule is at a point for this lab to be built then formal class and reference training should commence.

What the engineers learn in class and with the educational references provided in this paper they can safely experiment with different protocols and configuration options in the lab without overstepping other configurations or plans

Once CT. Ins. Corp. reaches a point where the various support teams has had some time to learn and understand the protocol and is comfortable from a design, configuration and support standpoint they would look to get hardware similar to production from spare inventory.

This would constitute the ***formal*** IPv6 Lab.

In this lab, the components would match the production environment as closely as possible. It should have the following attributes

1. Mock ISP IPv6 presence for – IPv6 Inbound Internet connectivity Internet Presence (mock DC 1 is sufficient at first but a mock Windsor for BGP testing should be planned)
2. Mock External zone with similar equipment to test for Dual-Stack and translation
3. Mock Secure DMZ and SLB DMZ – with similar equipment to test for Dual-Stack and translation

Once the External and DMZ zones are built for IPv6 testing CT. Ins. Corp. can continue testing:

1. DMZ support services, Firewalls, load balancers, DNS etc. testing
2. Internet Front Facing Server
3. Employee Remote IPSec/SSL VPN – Translation and Dual-Stack remote site Internet connectivity

It is recommended that CT. Ins. Corp. start a process to acquire any equipment possible, power and rack space to build the generic ***sandbox*** lab first. This lab should include routers, switches, firewalls, servers and some laptops for end clients and traffic generator stimuli use. The equipment should be as close to production as possible if possible. This lab should not be connected to the production lab or business environment. The lab's initial use of the flowchart below indicates should be for IPv6 first then test for Dual-Stack with IPv4.

The formal lab will match the production environment as close as possible for hardware but must be a match for software. This ensures that the results from protocol and operational behavioral testing conducted will match that in production. This lab is not to be confused with the PerfQA lab or any IPv4 production lab currently in place. This lab is deemed ***formal*** as opposed to the sandbox for it will be used for pre-production testing and deployment phase process testing.

A lab testing flowchart is provided as an example of the process involved. AMI can assist CT. Ins. Corp. in defining a detailed custom testing approach to achieve their goals.

F5 Related testing considerations:

This section briefly outlines some of the considerations CT. Ins. Corp. must keep in mind when testing in the sandbox or formal lab.

The F5 offers several modes for load balancing:

Static modes – Round Robin and Ratio
Dynamic Modes – Least Connections, Fastest, Observed, Predictive, and Dynamic Ration
Also, each can load balance by Pool Member or by Node.

What this means is any testing for IPv4 for the load balancing DNS and content traffic must also be tested with IPv6 individually and then combined to determine if any issues arise due to performance, memory allocation and any other nuances that may appear.

For example a virtual server and pool relation created for XWEB Member web instances for IPv4 will be duplicated for an IPv6 virtual server and pool(if not shared) and the load balancing may react differently statistically since the translated has a different delay involved. We may have a state that for native IPv4 the load balancing mode selected meets CT. Ins. Corp.'s needs but for translated IPv6 the same nodes in the pools may require a different method. This also applies to Priority Group activation methods tool. CT. Ins. Corp. must fully test all load balancing features in IPv4 and translation mode to gauge the behavior.

For VPN traffic the ASA's currently load balance their traffic so the consideration is flagged here to ensure that if translated VPN traffic is to be translated and load balanced at the F5 just to reach the ASA pool to be load balanced a second time may or may not introduce any type of asymmetrical state pattern between the appliances. CT. Ins. Corp. should test for this behavior as well. It could provide CT. Ins. Corp. additional insight as to possibly achieving a higher efficiency utilizing the F5's load balancing algorithms for IPv4/V6 over the ASA's, thus turning off the ASA feature.

Additionally DNS Express for the GTM, GTM CMP also play a role in the testing for not only the IPv4 deployment but for the translated IPv6 in relation to performance tuning.

DNS64 should be tested between their authoritative servers internally and NSI.
The use of Anycast is another testing consideration if used for IPv4 the translated requests be tested as well.

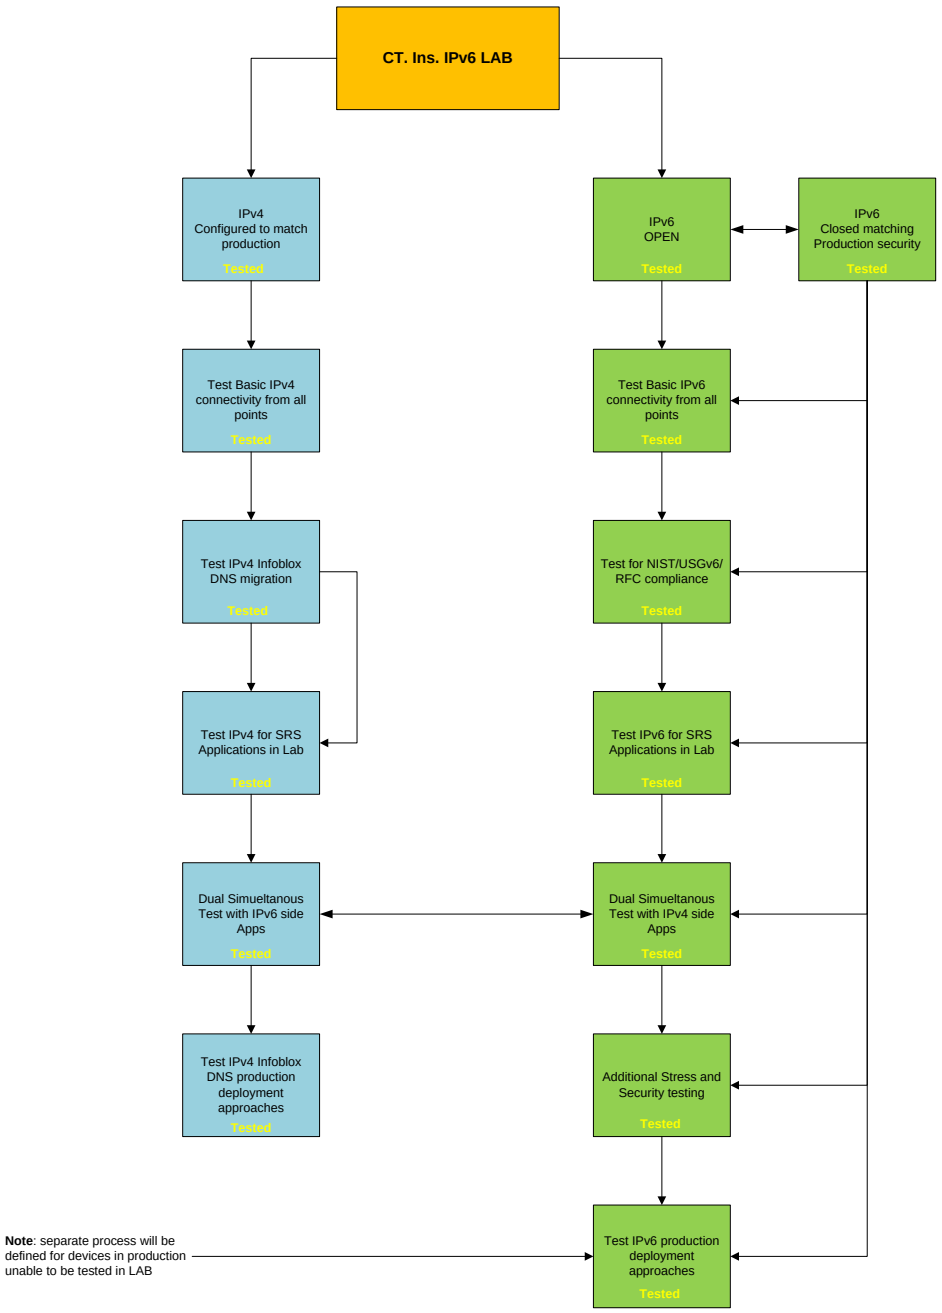

Figure 19 – IPv6 Lab Testing Flowchart

The flowchart is an overall high-level summary of the major steps for testing IPv6 in the formal lab. In short, the testing would comprise of first just enabling IPv6 only and executing a test script against the equipment.

This type of IPv6 only testing helps the engineer capture the following:

- Validating configuration and commands required for each device
- Capturing known expected positive behavioral results and save to compare for production deployment
- Capture any anomalies, errors and bugs discovered
- Testing basic and advance features supplied for IPv6 based on the HW/SW point tested
- Certifying that the HW/SW really works with IPv6 and no future upgrades are required
- IF upgrades are required then this is the point to test and validate
- Conduct stress and traffic generation testing against lab for performance and buffer/queue behavior
- Testing of security components – stress testing as well
- Testing of translation components if Dual-Stack is not extended in the lab.
- Testing of translation integration first then removal and Dual-Stack progression in lab

Once the IPv6 only tests are completed the following should occur regardless if translation and Dual-Stack approach are used:

- Apply production IPv4 configuration to all devices in lab. The lab must look and act similar to IPv4 production.
- Test for IPv4 production behavior, accessing web sites etc.
- Validate that IPv4 is working in lab against all services as tested earlier for IPv6 only
- Conduct side by side parity testing with v4/v6- includes stress testing
- Test for IPv4/v6 configuration option removal and addition options.
- Remove IPv6 altogether and lab only reflect production IPv4
- Start testing IPv6 integration stage process testing, zone at a time, and document the process and procedures for production migration script.
- Tune and adjust process as needed
- Ratify migration process and schedule for change

Application Testing Methodology

As discussed earlier, there are many unknowns surrounding application interoperability over IPv6. This fact coupled with the extensive leveraging of the Internet for CT. Ins. Corp. use cases connectivity and application service, puts a large percentage of CT. Ins. Corp.'s applications at risk of needing to support IPv6.

Translation technologies such as NAT64 will provide CT. Ins. Corp. with the ability to mitigate the number and types of applications that would need to support dual-stack day one post exhaust. However, translation technologies such as NAT64 are emerging technologies and the ability of these technologies to support CT. Ins. Corp. applications is unknown. In addition, it must be stated that NAT64 is a transition mechanism that is an interim step to integrating IPv6 and should not be considered as a steady state solution.

AMIis recommending in-depth testing of translation capabilities in order to minimize the impact the IPv4 exhaust dates will have on the CT. Ins. Corp.'s application environment. This is a priority and should be completed as soon as possible. However, regardless of the results of the NAT64

capabilities, it is assumed that planning for application testing will need to commence immediately and that an approach and resource allocations must be defined in the near term.

CT. Ins. Corp. will need to characterize their applications to identifying the importance of the application as well as the capabilities and functions within the application that will need to be tested. Once the applications are characterized, CT. Ins. Corp. can prioritize what applications to address first. It will not be feasible to test all applications prior the exhaust date, therefore developing a "priority list" will allow CT. Ins. Corp. to address the applications that represent the most risk to CT. Ins. Corp.. A first "sweep" of applications, their importance and priority to the business is included in the report in section ***Applications - External Facing.*** This list can be used as an initial short list for test candidate consideration.

AMIrecommends that CT. Ins. Corp. focus on the applications that introduce the most risk to CT. Ins. Corp. moving forward. Applications that affect CT. Ins. Corp. revenue, market reputation or customer SLA's/satisfaction should be priority. The following areas are considered primary and should be executed in phase 1 of the transition effort.

1. Identify and test mission critical applications that are traversing the Internet from remote sites. These applications run over IPSEC VPN tunnels and SSL connections today. Validate ability of applications to work over tunneling and SSL technologies.
2. Identify and test applications that are delivered over the Internet outside of the remote site applications. Validate ability of these applications to work over translation technologies.
3. All other applications that are not impacted by the IPv4 address exhaust can remain on IPv4 indefinitely. Focusing on the Internet services will provide CT. Ins. Corp. with the ability to learn IPv6 transition skills for this subset of CT. Ins. Corp.'s application space. Once the IT teams have been through the public facing applications, it will then be feasible to address all applications.

Regardless of which applications are tested first or last, CT. Ins. Corp. may need to modify and enhance their application testing processes and develop a plan for testing applications prior to enabling them for IPv6. The application team will need to work with the network teams to provide requirements necessary to test applications. This will include plans for testing applications in the lab or via pilots.

Below is a framework for the development of an Application Transition Plan. This framework can be used as a guide for testing the CT. Ins. Corp. applications.

1. Application functionality requirements
2. Use of standardized APIs
3. IPv6 capability requirements
4. Dual use (IPv4 & IPv6) or single use (IPv4 or IPv6)
5. IPv6-capable transition requirements
6. Transition approach
7. One set of applications that support both IPv4 & IPv6
8. Separate applications running in native IPv4 or IPv6 mode
9. Timing of application transition with network transition
10. Application audit & analysis
11. Identify all applications in use today
12. Determine if they are impacted
13. Identify method of transition
14. Application Transition Resources
15. Who will modify the application
16. COTS product (vendor responsibility)
17. Contractor developed/modified
18. Internal developed/modified
19. Budget considerations
20. Prioritization versus other upgrades and patches
21. Rolling in new versions of software
22. Support for legacy applications
23. Length of time they will be supported
24. Transition mechanisms to be used to extend life

Regardless of the testing approach, significant coordination and collaboration is required from the network, application, and server teams in order to facilitate the testing of these applications.

IPv6 Skills and Training

IPv6 skills are required across all IT groups. While very similar to IPv4 in functionality, IPv6 and IPv4 are separate protocols. IPv6 will be enabled in parallel with IPv4 in the External and DMZ zones and then eventually internally thus CT. INS. CORP. resources will need to understand the specifics how to enable IPv6 as a separate protocol on all the devices and systems under their control. Fortunately the fundamentals of the new technology is IP so there is an inherent knowledge foundation already in place, there just needs to be a ramp up to the IPv6 version. IPv6 skills will be initially required of CT. Ins. Corp. Internet Engineers and Network Engineers for the continuing planning, testing and deploying of IPv6 in the phases outlined in this paper.

Training is a critical part of any IT organization. Keeping staff up to date in technical knowledge and skills ensures the organization will receive efficient results from its IT staff.

CT. Ins. Corp.'s approach to training is when the business is deploying a new technology CT. Ins. Corp. will follow through with the proper training to ensure its engineers can support the new technology in a timely manner. For the purposes of this paper the focus of training is for the Internet and Network Engineering groups.

- Pursuing certifications is encouraged within the groups and supported by management
- Global Knowledge is used as their main training vendor.
- The staff experience and certification levels range from CCN/DA- CCN/DP – CCIP - CCIE
- A two day custom basic Introduction to IPv6 was provided to the groups in September
- CT. Ins. Corp. is considering a five day Global Knowledge course which covers IPv6 in detail with hands on labs included which is Cisco technology based.
- When training is to be provided to the entire group the training is brought in house
- When training is to be provided to an individual for a specific need or for a certification/boot camp the individual is sent to the training provider's facilities
- CBTs and having one engineer train the others are not considered, formal classes are preferred

CT. Ins. Corp.'s IPv6 training plans are as follows:

- Budget of $40k over 2 years of IPv6 training for 20 engineers
- The 2nd and 3rd quarter of 2013 are planned for training to cover the expected period when IPv6 is first deployed in the External and DMZ zones
- The 1st and 2nd quarters of 2014 are planned for training to cover the internal related phases of IPv6 deployment.

CT. INS. CORP. will need to focus the initial training thrust on the groups that will be impacted the greatest by the IPv6 integration effort. This includes the following teams responsible for the External and DMZ zones.

1. Internet and Network Engineering teams
2. Operations Management Teams
3. Server Teams
4. Security
5. Application Teams

Each IT group within CT. INS. CORP. will need to develop basic intermediate and advanced skillsets required to deploy IPv6. Some groups will be impacted more than others, but all IT groups including network, server, security, and operations will need to ramp up on IPv6 to the same level of proficiency as IPv4. Each IT organization will need to ensure that the training, though different per discipline is consistent in terms of when it is conducted in relation to the overall IPv6 deployment schedule.

Network Administration and Operations Management Teams

The greatest impact is on the network design/engineering and operations teams in terms of breadth of skills required. Turning up IPv6 on a desktop device or server only requires a basic understanding of IPv6. However, design, testing, deployment, and management of the vast array of network equipment require significantly advanced skills.

The core network skills required for IPv6 are the same for as for IPv4, however while IPv6 provides the same functionality as IPv4, all of the commands and configurations used to enable IPv6 are different. This in itself is a major challenge with IPv6 skills adoption. In addition to the configuration skills, IPv6 introduces new addressing structure, tunneling technologies and translation capabilities. This requires expertise in the design of IPv6 technologies that are not evident in IPv4. Finally, understanding how IPv6 works and how to troubleshoot IPv6 network issues require advanced understanding of IPv6. Development of basic, intermediate, and advanced IPv6 skills by the network and operations teams is required. These individuals should embark upon a training regimen that includes leveraging of instructor led training (where possible) combined with self-paced study and support of lab and pilot deployments. The following is a breakdown of skills for each level:

1. Basic Understanding – Knowledge of IPv6 addressing and skills required to enable dual-stack interfaces on a device. Knowledge of IP routing ability to troubleshoot device connectivity
2. Intermediate Understanding – Knowledge of specific device functionality and features required to enable IPv6 functionality. Knowledge of tunneling technology and translation functionality and the ability to design, engineer, enable and test basic IPv6 functionality capabilities on devices under span of responsibility. In addition, the ability to troubleshoot IPv6 functionality over the network will be required.
3. Advanced Understanding – Architecture level knowledge of vendor products is required and the ability to work with vendors to develop testing scenarios, and validate transport and IPv6 communications over the vendor products.

he network and operations team.

Table *15 - Network and Operations Training Recommendations* provides a mapping of responsibilities and skillsets for the operations, CCNA level staff and CCNP level staff on the network and operations team.

Table 15 - Network and Operations Training Recommendations

	IPv6 Responsibilities	IPv6 Skills

	IPv6 Responsibilities	IPv6 Skills
Operations	• First-level problem detection, triage and trouble ticketing • Provides status on network, system and service affecting issues • Capture of IPv6 alerts, traps and logs • Configure equipment interfaces and basic routing for IPv6 connectivity • IPv6 troubleshooting to validate network connectivity • Understanding of the processes and policies for escalation and coordinate with other IPv6 resources	• IPv6 Addressing & Subnetting • IPv6 Addressing types including ULA, GUA, and LLA • Assigning IPv6 addresses to interfaces • IPv6 tunneling • Dynamic routing in IPv6 • Static routes in IPv6 • Floating static routes in IPv6 • IPv6 ping
Tech Support CCNA Level staff	• First-level problem detection, isolation, and trouble ticketing • Follows NOC policies, processes, and procedures • Performs regular network, system, and service surveillance • Keeps NOC team lead and manager informed of any perceived trends or real operational trends • Ability to work with service providers in the resolution of network, system and service issues • Develops equipment configurations for enabling IPv6 connectivity • Develops equipment configurations for enabling IPv6 and IPv4 feature parity • Develops integration plans • Works with level 3 architects in construction of IPv6 test plans • Works with testing teams to set up labs and validate IPv6 functionality • Updates and corrects to engineering, process, and procedural flows as necessary • Ensures problems are satisfactorily resolved	• Understanding of IPv6 addressing techniques (SLAAC, DHCP) • Understanding of IPv6 subnetting • Understanding of IPv6 routing • Understanding of IPv6 QoS • Understanding of IPv6 multicast • Understanding of IPv6 tunneling technologies • Understanding of IPv6 translation capabilities • Understanding of IPv6 traffic engineering
Engineering CCNP Level staff	• Provides technical leadership regarding IPv6 integration • Interfaces with other technical teams to discuss impact of IPv6 • Develops and maintains architecture standards for IPv6 • Develops configuration standards for IPv6 equipment • Provides product analysis and validation of IPv6 functionality • Coordinates with vendors and leads discussions around product IPv6 support • Leads or provides direct input to the IPv6 lab design and testing approach	• Proficient in IPv6 routing, addressing, dual-stack and tunneling technologies • Able to develop configuration details surrounding IPv6 integration, functionality and testing for IPv6 infrastructure assets • Able to evaluate and recommend product strategy and direction • Able to troubleshoot IPv6 connectivity, routing and service delivery issues over the production network

Training for the Server, Security and Application teams should be reviewed during the planning phases of the IPv6 deployment. This ensures that training budgets and timelines can coincide with

the overall IPv6 project plan and approach utilized. For example if the initial phases of the IPv6 deployment is to utilize a translation solution and the servers will not be converted to a Dual-Stack configuration then server training for server administrators would not be needed until just before the period of adding IPv6 to the servers. There is no benefit to training server administrators on a protocol they may not get to implement for another year.

IPv6 Skills and Training Recommendations

CT. Ins. Corp. already recognizes the important of training and has plans in place to train all its engineers in IPv6 skills. This is ensures skill parity across team members and flexibility in support coverage.

Some additional recommendations:

a. Scheduling of training must coincide with the IPv6 projects overall schedule. This means that training should commence when the informal IPv6 lab or sandbox is planned and built. This ensures that the engineers receiving the training can put it to immediate use and this help reinforces the knowledge and skills acquired for future IPv6 deployment phases.
b. Additional training must commence at a minimum during the formal lab phases and not before or during production deployment. The reason for this is in training there may be tools, tips other deployment case studies from other companies disclosed from the instructor or another student that may benefit CT. Ins. Corp.'s deployment and CT. Ins. Corp. would need time to test in the lab.
c. Ensure any additional training coves in-depth packet level and forensic troubleshooting for Internet/Network Engineering and Security teams. Introductory classes may provide a brief lab for this this but it may be too broad. Security based classes provide focused material and help ensure the engineer has the knowledge and skills to troubleshoot security related IPv6 issues.
d. Engineers should familiarize themselves with some critical aspects of IPv6 after formal training via additional studying of RFCs and online materials. Some aspects considered should be ICMPv6, Extension Headers, Fragmentation, DNS extensions, Neighbor Discovery, Default Address Selection. A list of RFCs is provided in **Appendix A**
e. Check with Network Management and IDS/IPS vendor's to see what training offerings they provide for their products related to IPv6. For example, is there "add on" to a product that requires training?
f. The Internet and Network Engineers should also attain the IPv6's Forums Silver or Gold Certification status. The IPv6 Forum is a worldwide consortium of leading Internet vendors, Industry Subject Matter Experts, Research & Education Networks, with the mission to advocate the adoption of IPv6 technology.IPv6 Forum offers the IPv6 Education Certification Logo Program. This program encourages and accelerates education on IPv6 and promotes swifter adoption of IPv6 by certifying courses, engineers and trainers.

   By taking the CCNA, CCNP, CCIE Routing and Switching, CCDA, CCDP, or CCDE certification exams - certified by the IPv6 Education Certification Program - IT

professionals will be able to demonstrate that they have attained IPv6 knowledge and skills at the associate, professional or expert level.

Cisco also offers additional training with the **"IPv6 Fundamentals, Design and Deployment Course (IP6FD)".** This course has been Gold certified by IPv6 Forum. After successfully completing the course, candidates can include on their resumes that they have attended the IPv6 Forum Gold Certified course.

IPv6 Forum Logo Guidelines

The prime objective of this Program is to encourage, accelerate the education on IPv6, and promote thereby swifter adoption of IPv6 in the education curriculum and programs of the universities, research institutes, vendors and training specialists. The IPv6 Education Program is a program intended to increase practical engineering expertise and hands-on knowledge to tackle this large undertaking ahead of us extending thereby user confidence by demonstrating that qualified engineers will can IPv6.

Table 16 – Cisco and IPv6 Forum Certification Levels

If candidate has passed	>	CCNA or CCDA	CCNP or CCDP	CCIE Routing and Switching or CCDE
Candidate falls under this IPv6 Forum Certification Category	>	Silver	Gold	Gold
Candidate can obtain this logo from IPv6 Forum	>			

Additional information about the IPv6 Forum's Certification can be found here:
http://www.ipv6forum.com/ipv6_education/

Note: U.S. vendors have the most IT products that have been approved by the IPv6 Forum's IPv6 Ready program, which runs conformance and interoperability tests. U.S. companies including Cisco, HP and Juniper have run 425 networking products such as routers and hosts through the IPv6 Ready process. This compares to 350 IPv6 Ready products from Japanese vendors and 250 from Taiwanese vendors.

CT. INS. CORP. utilizes Cisco equipment as the core enabler of IPv6 communications. Fortunately, there are multiple avenues available for Cisco or general IPv6 training including web based, instructor led and lab based offerings. All resources will need to understand the fundamentals of IPv6 and it is recommended that self-study and web based delivery be utilized for this effort during any Lab or pre deployment phases of an IPv6 project. There is a tremendous amount of free and subscription online related training and materials available and some very good books as well. It is recommended that CT. Ins. Corp. takes some time to review these sites and books listed and choose what will best serve their needs.

- http://www.6deploy.org/index.php?page=tutorials - *excellent repository of tutorials, reference and case studies of deployments.*
- http://www.slideshare.net/feb_989/cisco-i-pv6-laband-techtorial-workshop-v0
- *Cisco IPv6 Lab and Techtorial by Hinwoto*
- http://www.6diss.org/e-learning/ - *web based IPv6 class*
- http://www.ipv6tf.org/ - *IPv6 Portal another repository of IPv6 information*
- http://ipv6.he.net/ - *Hurricane Electric IPv6 tools and education*
- http://www.ipspace.net/Webinars - IpSpace IPv6 - *IPv6 Webinars*

The following documents, web based and instructor led training is recommended.

Documents/Books/Whitepapers

- Introduction to IPv6 – *Wikipedia IPv6*
- *IPv6 for Enterprise Networks*, Shannon McFarland, Muninder Sambi, Nikhil Sharma, Sanjay Hooda; Copyright 2011 Cisco Systems, Inc., Cisco Press, Indianapolis, IN; ISBN: 1587142279
- *IPv6 Security IPv6 Security*, Scott Hogg and Eric Vyncke, Copyright© 2009 Cisco Systems, Inc, ISBN-10: 1-58705-594-5
- *Understanding IPv6: Your Essential Guide to IPv6 on Windows Networks*, Joseph Davies, Microsoft Press; Third Edition (June 27, 2012), ISBN-13: 978-0735659148
- Cisco IPv6 Configuration Guide - *Cisco IPv6 Configuration Guide*
- *Deploying IPv6 in the Internet Edge* – SBA Slide Deck
- *Deploying IPv6 Networks*, Ciprian Popoviciu, Eric Levy-Abegnoli, Patrick Grossetete, Copyright 2006 Cisco Systems, Inc., Cisco Press, Indianapolis, IN; ISBN: 1587052105
- *Global IPv6 Strategies: From Business Analysis to Operational Planning*, Patrick Grossetete; Ciprian Popoviciu; Fred Wettling; Copyright 2008 Cisco Systems, Inc., Cisco Press, Indianapolis, IN; ISBN 1587053438

1. Training Courses
    a. IPv6 Fundamentals, Design, and Deployment (IP6FD) – Cisco Learning Partners
    b. Introduction to IPv6 Features and Services – Online Delivery
    c. IPv6 Deployment Fundamentals – Online Delivery
    d. Introduction to IPv6 Transition Mechanisms and Deployment – Online Delivery

Server & Desktop Team Training

The server team will need to be able to enable IPv6 on each server platform that requires IPv6 support. CT. Ins. Corp. will need to develop the IPv6 skills necessary to deliver basic, intermediate, and advanced support for all servers and desktops under their span of control. From a server standpoint, CT. Ins. Corp. primarily has servers running on Windows Server 2008 R2, Linux, VMWare, IBM AIX, HP-UX, and Sun platforms. These servers support IPv6.

The initial training approach should focus on the skills necessary to support IPv6 on the Server OS's and the Windows7 operating system for Desktops. This is the direction moving forward and the investment in training should align with the strategic direction for the server and desktop environment. IPv6 should be incorporated into a pilot and a subsequent Windows7 desktop OS load to ensure that the skills are being developed for production deployment of IPv6. For desktop operating systems, CT. Ins. Corp. will need to develop the IPv6 skills necessary to deliver basic, intermediate, and advanced support for all of the servers and desktops under their span of control. All desktops today have IPv6 "Disabled" at the hardware level. These desktops will require a push to enable IPv6. For Win7, the standard for desktop operating systems today and beyond will require a push to complete IPv6 enablement. For older models such as Windows XP and Windows Vista, the assumption is in an IPv6 environment, these machines would be removed.

Training should focus on the Microsoft aspects of the environment. There will need to be expertise for non-Microsoft platforms, but the immediate need is in the Microsoft space. Microsoft provides a wealth of information for enabling IPv6 on their platforms. Below are some key documents and tutorials for enabling Microsoft services.

Online Content:

a. *Microsoft TCP/IP Technical Reference* – Windows Server 2008
b. *Support for IPv6 in Windows Server 2008 R2 and Windows 7*
c. IP Configuration Migration Guide – Windows Server 2008 R2
d. *IPv6 for Microsoft Windows: Frequently Asked Questions*

Courses:

a. 6419 - Configuring, Managing and Maintaining Windows Server® 2008-based Servers
b. 6421 - Configuring and Troubleshooting a Windows Server® 2008 Network Infrastructure
c. 10215 - Implementing and Managing Microsoft® Server Virtualization

Application Team Training

The application teams will have to ensure that IPv6 is enabled on the current applications and incorporate IPv6 into all future application development activities. This will require porting of current IPv4 capabilities to the IPv6 protocol stack. AMI did not assess the application team current skills; however, it is assumed that there are basic, intermediate, and advanced skillsets evident in these groups. These resources will need to understand how to port applications to IPv6 and how IPv4 and IPv6 will operate in parallel on a particular application.

There is little training available on IPv6 application enablement. Microsoft has provided guidance on how to port applications that can be found via the links below:

a. *How to Convert an Application from IPv4 to IPv4/IPv6*
b. *Testing your Application or Device in an IPv6 Environment*
c. *IPv6 Guide for Windows Sockets Applications*
d. *Running the Checkv4 Utility* - The Checkv4 utility documentation includes tips about the order in which to address IPv6-enabling issues. This is the order in which to approach the conversion process in your applications.

AMI's own experience with IPv6 training has shown that there are few external offerings for IPv6 training. There are offerings in the market for instructor led training on Cisco equipment but no similar offerings for server, desktop and application training. This puts a heavy reliance on these CT. Ins. Corp. teams to learn these skills through self-study and actual hands on experience with the technologies. This requires investment in an IPv6 lab or pilot rollouts on the production network.

Europe

The Europe DC was not originally scoped to be included into this readiness analysis, however since the data was captured, AMI has added this DC as a separate section in this report. The same analysis methodologies and review applied to the US sites also applied to this DC as well. The asset inventories and costing applicable to this DC is included. CT. Ins. Corp. can include this DC in their major US IPv6 plans or carve out a separate project to handle this DC.

CT. Ins. Corp. Maintains a Europe Data center and that site handles most of the web traffic for that region in Europe. The data center utilizes a managed service provider which provides the internet POC into the DC. SunGard is the provider and the connections are Ethernet based. The Europe DC utilizes similar Hardware and Software as the US DC 1 and Windsor DCs. SunGard provides two Ethernet circuits for redundancy. The Internet traffic flows inbound only. All outbound traffic flows through the CT. Ins. Corp. WAN to the US DCs. Europe employs a Cisco CSS Appliance for server load balancing.

Refer to figure 20 on the following page.

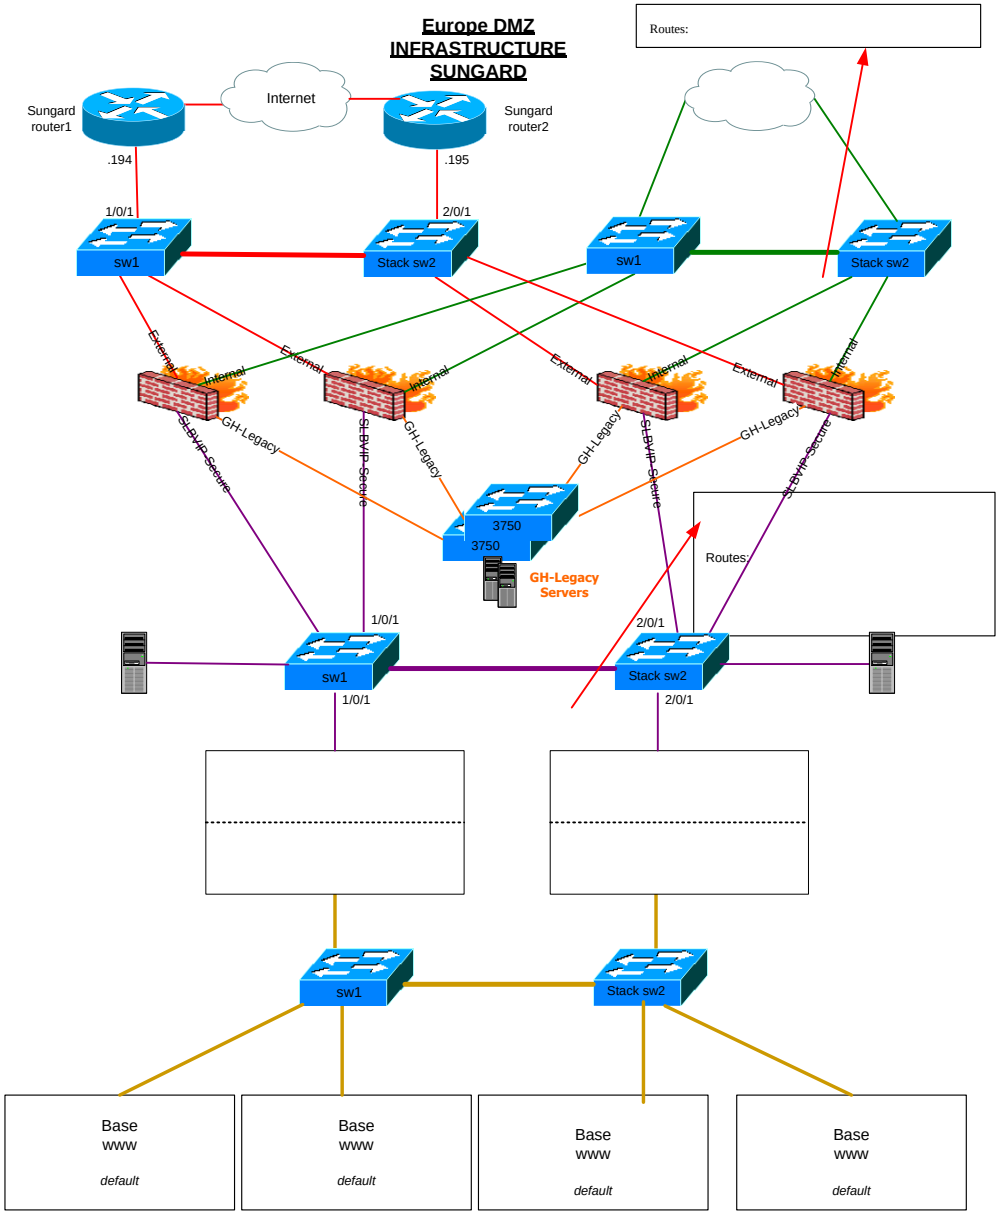

Figure 20 – London DC External Infrastructure

There are two approach options present for CT. Ins. Corp.:

1. Use London as a small pilot. It has few devices to upgrade and if the F5 is used for translation services CT. Ins. Corp. can take a small low impact step into IPv6. From what is learned from the London Pilot can be applied to the US DCs.
2. In contrast CT. Ins. Corp. can continue to deploy first in the US and what is learned from the US deployment is applied to London.

AMI recommends the following:

1. Conduct upgrades on assets required to meet IPv6 compliance based on their refresh cycle or priority
2. Deploy Dual-Stack in the External zone towards the Secure DMZ zone as a first step.
3. Deploy Translation capabilities in the External zone towards the DMZ zone to support all the, Cisco and McAfee firewalls, for VPN and Web applications thus enabling IPv6 inbound packets to be translated to IPv4 prior to reaching DMZ zone IPv4 resources.

Refer to figure 21 on page 94.

Note: Additional research for outbound sourced IPv6 will be required since current outbound traffic does not leave London via SunGard's services but through the CT. Ins. Corp. WAN towards Ct. this asymmetric flow is acceptable for IPv4 but requires additional planning for IPv6.

Additionally there remains the possibility of an IPv6 translated accessed application transaction breaking due to the asymmetric IPv4 path flow. For example, if an IPv6 client in Europe via SunGard to access an CT. Ins. Corp. Website/services and that transaction/flow is translated to IPv4.

However, on the internal side's IPv4 server there is another thread to the transaction that goes towards the internal enterprise in Ct., will the transaction state(s) along the path hold and return to be translated towards SunGard via IPv6 or will it break at some point? Additional research must be conducted against the applications and servers employed in London prior any deployment.

Being consistent with the US DC the use of Big IPs F5 GTM/LTM can provided the translation capabilities. Refer to section ***IPv6 Architecture - Translation Capability - Tactical*** for details about the US recommendations use of the F5.

Step 1 Enabling IPv6 in External Zone

SunGard enables IPv6 on its managed internet services interfaces and routes.

Assumptions:

- All components are upgraded and are IPv6 compliant if required
- SunGard has supplied CT. Ins. Corp. and AMI indication that it will support IPv6 services by the end of 2012
- Determine if SunGard provides IPv6 address or CT. Ins. Corp. uses it's PI
- SunGard routers will have their own managed IPv6 address similar to IPv4 space.
- F5 translation appliance will either receive a SunGard IPv6 address or use their ARIN assigned
- If ARIN assigned CT. Ins. Corp. PI address is used check for RIR prefix submission issues
- If BGP policies are present this is SunGard's responsibility

Low impact initial step. CT. Ins. Corp. and ISP can monitor for discovery traffic and check security related statistics. CT. Ins. Corp. works with ISP for IPv6 interface turn up.

1. Engage SunGard and determine IPv6 addressing stance
2. SunGard configures IPv6 on its routers interfaces
3. CT. Ins. Corp. deploys F5 in External Zone
4. CT. Ins. Corp. deploys F5 in DMZ for CSS if applicable
5. CT. Ins. Corp. turns up IPv6 on F5 and tests basic IPv6 communications with SunGard
6. SunGard monitors IPv6
7. CT. Ins. Corp. progress with sequence of deploying and testing of translation for network services, DNS, FTP, NTP etc.
8. CT. Ins. Corp. commences with enabling translation service for web traffic and other applications
9. Conduct validation testing
10. Enable IPv6 on redundant equipment – if applicable

Dual-Stack progression

Once the Translation is in place CT. Ins. Corp. can proceed with deploying IPv6 on the remainder of its assets either from the External Zone towards the DMZ or vice versa. IPv6 can be configured on these devices but not enabled since translation is still working. A sequenced transition plan must be created that covers approaches such as surgically enable IPv6 up through the DMZ towards the External zone F5 and then disabling translation per service. In contrast a big bang approach of turning off translation altogether and enabling IPv6 Dual-Stack during a maintenance window can be employed.

There are some unique differences with using the managed provider vs. the approach in the US when it comes to further progressing into the DMZ.

If CT. Ins. Corp. utilizes its own PI addresses in the External Zone F5, it can continue to deploy IPv6 towards the firewalls and into the DMZ for full Dual-Stack capabilities. However, if SunGard provides the address for the F5 either the CT. Ins. Corp. External Firewall outside interfaces will require SunGard IPv6 address and then use CT. Ins. Corp. PI address on their inside interfaces or the F5 will have to route IPv6 between the SunGard IPv6 space and CT. Ins. Corp.'s PI space so CT. Ins. Corp. can apply its addresses on outside firewall interfaces and inward.

The Cisco CSS CSS11503 appliance does support IPv6 with a software upgrade. If for budgetary reasons this hardware is to continue to be used then the upgrade path is recommended as an intermediary Dual-Stack state since the F5 is in place. The CSS can work in a Dual-Stack environment.
However, there has been some discussion around Big IP F5 being deployed. If that is the case the expectation is that the F5 will replace the CSS in the DC.

Refer to figure 21 on the following page.

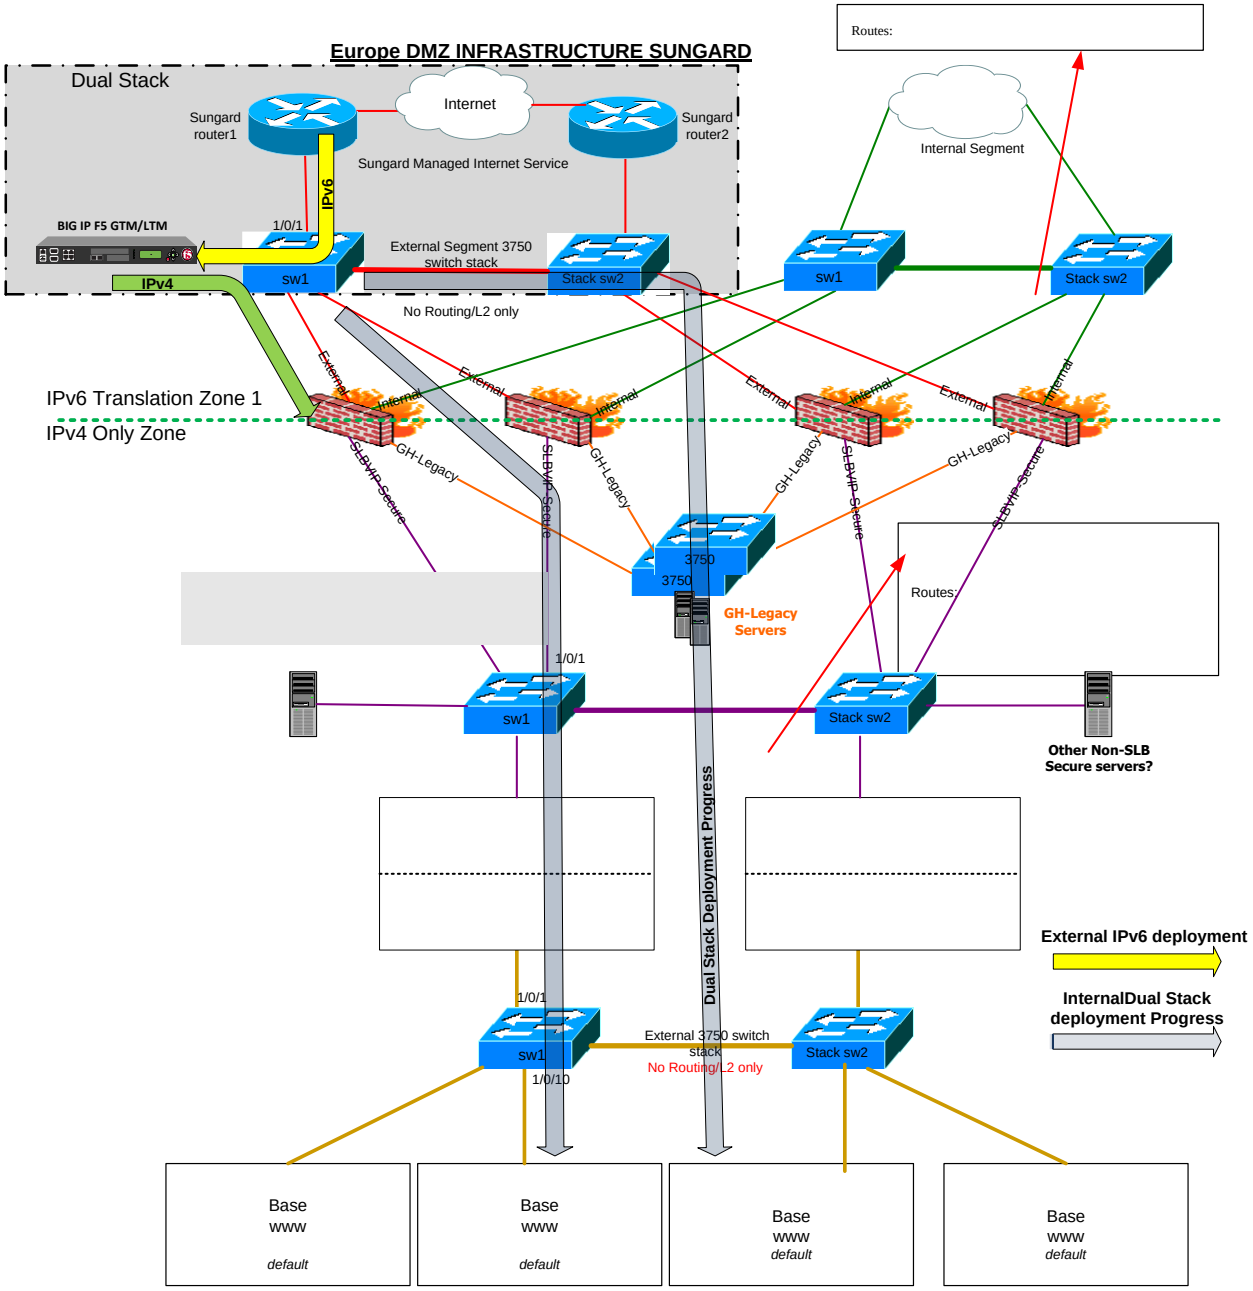

Figure 21 – Europe DC Translation Zone

Appendix A – IPv6 Related Links

#!/bin/the hacker's choice - THC
6DEPLOY - IPv6 Deployment and Support - Tutorials
6DISS > IPv6 Dissemination and Exploitation
6DISS.org - IPv6
7. next stop i pv6. how do we finally get there james wong
ARIN Number Resource Policy Manual
AskF5 | Informational: SOL9279 - BIG-IP automatically translates addresses between IPv4 and IPv6 when necessary
Author Expert: Enterprise IPv6
BCOP-IPv6_Subnetting.pdf (application/pdf Object)
BIND | Internet Systems Consortium
Brian McGehee IPv6 Address Utility
CAIDA: The Cooperative Association for Internet Data Analysis
Cisco Blog » Blog Archive » IPv6 Myths
Cisco Catalyst 6500 Architecture White Paper [Cisco Catalyst 6500 Series Switches] - Cisco Systems
Cisco Catalyst 6500: Building IPv6-Ready Campus Networks [Cisco Catalyst 6500 Series Switches] - Cisco Systems
Cisco Feature Navigator - Cisco Systems
Cisco IOS IPv6 Configuration Guide, Release 12.2SR - Implementing EIGRP for IPv6 [Cisco IOS Software Releases 12.2 SR] - Cisco Systems
Cisco IOS IPv6 Configuration Guide, Release 12.4 - Implementing IPv6 over MPLS [Cisco IOS Software Releases 12.4 Mainline] - Cisco Systems
Cisco IOS IPv6 Configuration Guide, Release 12.4 - Implementing Tunneling for IPv6 [Cisco IOS Software Releases 12.4 Mainline] - Cisco Systems
Cisco IOS IPv6 Provider Edge Router (6PE) over MPLS [Cisco IOS Software Releases 12.2 Mainline] - Cisco Systems
Cisco Network Registrar - Products & Services - Cisco Systems
Cisco Security Advisory: Cisco IOS Software IPv6 Denial of Service Vulnerability - Cisco Systems
Cisco.press.ipv6.Security
Core Networking and Security: IPv6 Subnet Calculators
Deploy360 Programme | Providing real-world deployment info for IPv6, DNSSEC and more…
Deploying IPv6 in Campus Networks [Design Zone for Campus] - Cisco Systems
Deploying IPv6 in Unified Communications Networks with Cisco Unified Communications Manager - Cisco Systems
Design Zone for IPv6 - Cisco Systems
dns_v6.pdf (application/pdf Object)
draft-ietf-v6ops-3gpp-eps-08 - IPv6 in 3GPP Evolved Packet System
draft-kohno-ipv6-prefixlen-p2p-03 - Using 127-bit IPv6 Prefixes on Inter-Router Links
draft-townsley-troan-ipv6-ce-transitioning-00 - Basic Requirements for IPv6 Customer Edge Routers - multihoming and transition
End-of-Sale and End-of-Life Announcement for the Cisco IOS Software Release

12.2(33)SXH [Cisco Catalyst 6500 Series Switches] - Cisco Systems
Episode 94: TechWiseTV: The IPv6 Implementation Action Plan - YouTube
F5 Friday: 'IPv4 and IPv6 Can Coexist' or 'How to eat your cake and have it too'
Global Traffic Manager Overview | F5 Networks
How to configure NetFlow on Cisco routers for IPv6 - NetFlow & sFlow Network Monitoring - Systrax
http://www.iana.org/assignments/icmpv6-parameters
http://www.ietf.org/rfc/rfc3315.txt
http://www.ietf.org/rfc/rfc3627.txt
Internet Protocol version 6 (IPv6) - Citrix eDocs
Internet Protocol Version 6 (IPv6) Parameters
InterOperability Laboratory: Services: Testing: IPv6
IP Version 6 (IPv6) - Cisco Systems
IPv4 Address Report
IPv6
IPv6 - Cisco Systems
IPv6 - Ubuntu Wiki
IPv6 | Keeping It Classless
IPv6 address - Wikipedia, the free encyclopedia
IPv6 at home, Part 1: Overview, Teredo « Thorsten on (mostly) Tech
IPv6 Configuration Guide, Cisco IOS Release 15.1M&T - Implementing IPv6 Addressing and Basic Connectivity [Cisco IOS 15.1M&T] - Cisco Systems
IPv6 Configuration Guide, Cisco IOS Release 15.1M&T - Start Here: Cisco IOS Software Release Specifics for IPv6 Features [Cisco IOS 15.1M&T] - Cisco Systems
IPv6 Default Router Preference simple example « Cisconinja's Blog
IPv6 Extension Headers Review and Considerations [IP Version 6 (IPv6)] - Cisco Systems
IPv6 Forum :: Driving IPv6 Deployment
IPv6 Migration Related Products | F5 Networks
IPv6 Migration: Implementation Checklist
IPv6 Ready Logo Program Approved List
IPv6 Ready Logo Site | Home
IPv6 Support in Microsoft Products and Services
IPv6 test - IPv6/4 connectivity and speed test
IPv6 Test and Dual-Stack Test For Network Connectivity
IPv6 to Standard
IPv6: - Verizon Enterprise Solutions
IPv6—A Service Provider View in Advancing MPLS Networks - The Internet Protocol Journal - Volume 8, Number 2 - Cisco Systems
IPv6hackingTools - Campus IPv6 Wiki
Local Traffic Manager Features | F5 Networks
Main Page - ARIN IPv6 Wiki
McAfee Product and Technology Support Lifecycle - Appliances | McAfee Support
NAT64 Technology: Connecting IPv6 and IPv4 Networks [IPv6] - Cisco Systems
NetScaler 10: A complete IPv6 powerhouse | Citrix Blogs
OpenDNS > Prepare for IPv6 with the OpenDNS IPv6 Sandbox

ostinato - Packet/Traffic Generator and Analyzer - Google Project Hosting
Peer Name Resolution Protocol
Re: Is there any version of IOS that supports IPv6 for Cisco 4700 routers?
RFC 1886 - DNS Extensions to support IP version 6
RFC 2460 - Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 - Neighbor Discovery for IP Version 6 (IPv6)
RFC 2462 - IPv6 Stateless Address Autoconfiguration
RFC 2710 - Multicast Listener Discovery (MLD) for IPv6
RFC 3041 - Privacy Extensions for Stateless Address Autoconfiguration in IPv6
RFC 3056 - Connection of IPv6 Domains via IPv4 Clouds
RFC 3168 - The Addition of Explicit Congestion Notification (ECN) to IP
RFC 3177 - IAB/IESG Recommendations on IPv6 Address Allocations to Sites
RFC 3484 - Default Address Selection for Internet Protocol version 6 (IPv6)
RFC 3493 - Basic Socket Interface Extensions for IPv6
RFC 3542 - Advanced Sockets Application Program Interface (API) for IPv6
RFC 3627 - Use of /127 Prefix Length Between Routers Considered Harmful
RFC 3646 - DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
RFC 3697 - IPv6 Flow Label Specification
RFC 3736 - Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6
RFC 3756 - IPv6 Neighbor Discovery (ND) Trust Models and Threats
RFC 3810 - Multicast Listener Discovery Version 2 (MLDv2) for IPv6
RFC 3971 - SEcure Neighbor Discovery (SEND)
RFC 3986 - Uniform Resource Identifier (URI): Generic Syntax
RFC 4191 - Default Router Preferences and More-Specific Routes
RFC 4193 - Unique Local IPv6 Unicast Addresses
RFC 4213 - Basic Transition Mechanisms for IPv6 Hosts and Routers
RFC 4291 - IP Version 6 Addressing Architecture
RFC 4361 - Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)
RFC 4380 - Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)
RFC 4443 - Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
RFC 4659 - BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
RFC 4795 - Link-local Multicast Name Resolution (LLMNR)
RFC 4861 - Neighbor Discovery for IP version 6 (IPv6)
RFC 4862 - IPv6 Stateless Address Autoconfiguration
RFC 4864 - Local Network Protection for IPv6
RFC 4941 - Privacy Extensions for Stateless Address Autoconfiguration in IPv6
RFC 5006 - IPv6 Router Advertisement Option for DNS Configuration
RFC 5072 - IP Version 6 over PPP
RFC 5213 - Proxy Mobile IPv6
RFC 5214 - Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)
RFC 5340 - OSPF for IPv6
RFC 5569 - IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)

RFC 5952 - A Recommendation for IPv6 Address Text Representation
RFC 6104 - Rogue IPv6 Router Advertisement Problem Statement
RFC 6144 - Framework for IPv4/IPv6 Translation
RFC 6147 - DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers
RFC 6296 - IPv6-to-IPv6 Network Prefix Translation
RFC 6346 - The Address plus Port (A+P) Approach to the IPv4 Address Shortage
RFC 6555 - Happy Eyeballs: Success with Dual-Stack Hosts
Routing IPv6 in 2011 — RIPE Labs
Search Results | Tools | F5 Networks
Setting Up an IPv6 Teredo Tunnel in Windows 7
SixXS - IPv6 Deployment & Tunnel Broker :: Main
SLAAC Attack – 0day Windows Network Interception Configuration Vulnerability | InfoSec Institute – IT Training and Information Security Resources
Test your IPv6.
The IPv6 Portal
Welcome to Cat Karat Packet Builder
Welcome to Cat Karat Packet Builder
Will the sky fall if you don't deploy IPv6?
http://www.team-cymru.org/

Appendix B – External Router IOS Management.

External Edge Router AMI options:
The external router CT. Ins. Corp. will utilize is the first point of performance, defense and visibility into the IPv6 world. This router is critical to CT. Ins. Corp.'s use and support of IPv6 for the enterprise so it is important it has the capabilities and features to provide as much visibility as possible for the CT. Ins. Corp. engineers. Cisco IOS provides many additional features that provide tools for supporting IPv6 right at the edge without the need to deploy additional HW/SW.

AMI has identified and listed some of them for CT. Ins. Corp. consider. Keep in mind these options are IOS platform dependent so not all features may be available to the compliant code base CT. Ins. Corp. utilizes.

1. Netflow – for statistics
2. Control plan policing
3. General IPv6 DoS packet policer to throttle IPv6 traffic(malicious or benign)
4. Kill switch
5. Packet tagging of IPv6 flow labels etc. for IDS use
6. Extended ACLs for Bogons and logging
7. Dynamic or reflected ACL for customer traffic
8. CEF for performance
9. NBAR for traffic classification and identification
10. Embedded capture manager – sniffer on router
11. Embedded Event manager – simple action scripting to control IPv6 events
12. Auto BGP Bogon service by Team-Cymru
13. Custom IOS command set using EEM and Alias to facilitate IPv6 management
14. Custom IOS menus to facilitate IPv6 management
15. SNMP 3 and MIBS for additional control.

A custom menu was created for a previous client's personnel to utilize in facilitating IPv6 administration

The example Menu was created on an Cisco ASR 1006 perimeter router running IOS-XE series code. Most of the commands are familiar versions of the ones used to administer IPv4 networks. Feature parity of commands varies with OS version and output will vary for IPv6

```
Server "rtr"    Line 0    Terminal-type (unknown)
 ****** IPv6 Router Management Menu ******
   1        Display ISP Side IPv6 Interface Details
   2        Display Screened Side IPv6 Interface Details
   3        Display IPv6 Routers
   4        Display IPv6 Neighbors
   5        Display IPv6 Traffic
   6        Display IPv6 Routes
   7        Display IPv6 CPU Processes
   8        Display IPv6 CEF
   9        Display IPv6 Policy To Mark Packets
   10       Enable NBAR on Interfaces
   11       Disable NBAR on Interfaces
   12       Display NBAR statistics for IPv6
   13       Enable IPv6 DoS Policing
   14       Disable IPv6 DoS Policing
   15       Display IPv6 Policing Statistics
   16       Display IPv6 Flex Netflow v9 Statistics
   17       Apply IPv6 ACL to interface ***ON BY Default**
   18       REMOVE IPv6 ACL from interface ***OFF BY Default**
   19       Display IPv6 ACL statistics
   20       Display Router LOG
   21       CLEAR all IPv6 related Counters - EEM Applet
   22       SHUT DOWN IPv6 Protocol on ISP side Interface(Emergency use)
   23       ENABLE IPv6 Protocol on ISP side Interface(Recovered)
   24       Start IPv6 Embedded Packet Capture Trace
   25       Stop IPv6 Embedded Packet Capture Trace
   26       Display IPv6 Embedded Packet Capture Session
   27       Clear Embedded Packet Capture Session Buffer
   28       Export Embedded Packet Capture File to flash
   29       SETUP Embedded Packet Capture(Already completed by default)
   30       Exit CT. Ins. Corp. Menu for CLI

 >>
```

Alias and EEM commands to aid in configuration and troubleshooting in the future: Custom IPv6 command set for IPv6 routers:

CLI Command	Action
aclon	Applies IPv6 ACL to Screened and ISP side Interfaces
acloff	Removes IPv6 ACL Screened and ISP side Interfaces
v6on	Enables Sub Interface running IPv6 towards ISP
v6off	Disables Sub Interface running IPv6 towards ISP – This is your “Kill Switch”
nbaron	Turns on Network Based Application Recognition for IPv6 on Interfaces
nbaroff	Turns off Network Based Application Recognition for IPv6 on Interfaces
doson	Turns on Denial of Service Traffic Policer on Outside ISP side interface
dosoff	Turns off Denial of Service Traffic Policer on Outside ISP side interface
flow	Displays IPv6 Flexible Netflow v9 Records and statistics
nbar	Displays Network Based Application Recognition for IPv6 statistics
police	Displays IPv6 DOS Policer Statistics when enabled
clearipv6	EEM script to clear all IPv6 related counters (does not clear log)
tagon	Turns on IPv6 DSCP tagging of frames entering our network
tagoff	Turns off IPv6 DSCP tagging of frames entering our network
cleartrace	Clears Embedded Packet Capture buffer
exporttrace	Exports completed Packet capture buffer to bootflash file ipv6capture
menuon	Adds Custom IPv6 Menu to IOS configuration
menuoff	Removes Menu from IOS configuration
m	Displays Custom IPv6 Menu
tracesetup	EEM script to setup Embedded Packet capture environment
acl	Displays IPv6 ACL and what is matched
traceon	Starts an Embedded Packet Capture session on router
traceoff	Stops an Embedded Packet Capture session on router
nocapture	EEM script to remove Embedded Packet capture environment on router
qfp	Shows Quantum Flow Processor Statistics
eemapson	Adds EEM scripts to the router
eemapsoff	Removes EEM scripts from router
loadv6acl	Loads entire IPv6 ACL set into router’s running configuration
remv6acl	Removes entire IPv6 ACL set from router’s running configuration/ run acloff first
flowon	Enables FlexNetflow v9 for IPv6 on AMI side interface
flowoff	Removes FlexNetflow v9 for IPv6 from AMI side interface

During the lab phases AMI can assist CT. Ins. Corp. engineers in creating a custom command set for their own use.

Appendix C – CT. Ins. Corp. Supporting Documentation

Project documentation was stored on the "AMI Business Direct" website, which is a secured knowledge management and document-sharing website. This site provided access to both AMI and CT. Ins. Corp. Employees during the project. AMI considered this documentation shared by CT. Ins. Corp. to be confidential and shared any documentation with only need-to-know individuals for research purposes. Once this project is completed, AMI will provide to CT. Ins. Corp. the full contents of the website and decommission.

This project collected and used the following list of documents for this engagement:

Filename	Description

These links below can be used to view additional documentation that was used by AMI that exists on CT. Ins. Corp.'s SharePoint sites.

Appendix D - Reference Documents

Supporting documents can be found under the Deliverable directory in a subdirectory called Supporting Documentation. Key documents in this subdirectory are listed below along with links to vendor documents are included below:

- **CT. Ins. Corp. IPv6 Costs by Phases**

This MS Excel spreadsheet contains data files delivered from CT. Ins. Corp. but manipulated into pivot tables that are used to complete the IPv6 Compliance and Cost sections of the deliverable.

- **CT. Ins. Corp. IPv6 S&R Deliverable Drawings**

The MS Visio file contains many of the most common drawings used in the deliverable. Most of the drawings come from CT. Ins. Corp. directly while others were created by AMI Consulting. The drawings are already included in the document but may have been converted to Meta files of JPG files.

- **AMI/ CT. Ins. Corp. IPv6 Strategy and Roadmap Project Notebook**

This MS Excel file contains key project related material like CT. Ins. Corp. SME Contacts, Project Contact List, Proposed Timelines, Assumptions, and Project Management related information.

IPv6 Compliance Reference Material

For the analysis of IPv6 hardware and software compatibility and compliance, the following files and web sites were used to define the readiness and impact to CT. Ins. Corp..

Cisco IOS Software Release Specifics for IPv6 Features

http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-roadmap.html#GUID-BF112A82-4C40-4C08-9573-C4A46BE7B108

http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-roadmap.pdf

Cisco feature navigator

http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp

Cisco CCO site end of life and end of sale site

http://www.cisco.com/en/US/products/prod_end_of_life.html

IPv6 Support in Microsoft Products and Services

http://technet.microsoft.com/en-us/network/hh994905.aspx

Appendix E - Project Contacts

Jeff Sicuranza
Sr. Consultant, AMI Consulting

Appendix F - Acronyms and Abbreviations

Acronym	Meaning
/8	Eight bit network mask
A	An IPv4 Resource Record
AAA	Authentication, Authorization, and Accounting
AAAA	An IPv6 DNS Resource Record
ACE	Application Content Engine
ACL	Access Control List
ACS	Access Control Server
AD	Active Directory
AES	Advanced Encryption Standard
ALG	Application Layer Gateway
ALS	Access Layer Switch
AMPO	Architecture Management Program Office
AMS	Application Management System (WEB Site)
ANX	
APA	Asia Pacific Africa (Geographical Area in CT. Ins. Corp.'s network)
ARP	Address Resolution Protocol
AS	Autonomous System
ASA	Adaptive Security Appliance
ASN	Autonomous System Number
AVPN	AMI Virtual Private Network
BA	Business Alliance
BGP	Border Gateway Protocol
BIA	Bump in the API
BIS	Bump in the Stack
CDS ID	Corporate Directory Services Identification
CERT	Computer Emergency Readiness Team
CGN	Carrier Grade NAT
CIDR	Classless Inter-Domain Routing
CMSA	Canada, Mexico, and South America
CMDB	Configuration Management Database
CMM	Communication Media Module
CMS	Content Management Server
CPN	Control Production Network
CSM	Cisco Services Module or Content Switching Module
CSS	Content Services Switch
DAD	Duplicate Address Detection
DC	Data Center or Delivery Content Site
DDOS	Distributed Denial of Service
HCP	Dynamic Host Control Protocol
DHCPv6	Dynamic Host Control Protocol version 6

DNS	Domain Name Service
DNSSEC	DNS Security Extensions
DPI	Deep Packet Inspection
EAA	Enterprise Architecture Assurance
EAG	Enterprise Architecture Governance
ECC	Engineering Computer Center (1 of 2 major data centers in CT. Ins. Corp._CITY)
EDC	Enterprise Data Center (combination of DC1 and ECC)
EGP	Exterior Gateway Protocol
EH	Extended Header
EoL	End of Life (Cisco Announcement of EoSale and LDoS Dates)
EoRFA	End of Routine Failure Analysis (Cisco)
EoSale	End of Sale (Cisco)
EoSCR	End of Service Contract Renew (Cisco)
EoSWM	End of SW Maintenance
ERP	Enterprise Resource Planning
EU	Europe (Geographical Area in CT. Ins. Corp.'s network)
EUI-64	Extended Unique Identifier 64 bit Address
EVPN	Enhanced Virtual Private Network
SITE1	CT. Ins. Corp.
DC1	1) CT. Ins. Corp. Computer Center (1 of 2 major data centers in CT. Ins. Corp._CITY) 2) CT. Ins. Corp. Credit Corporation
CNA	CT. Ins. Corp. North America (Geographical Region)
FQDN	Fully Qualified Domain Name
FSA	CT. Ins. Corp. South America (Geographical Region)
GRE	Generic Routing Encapsulation
GVA	Global Voice Architecture
HA	Home Agent
HIPAA	Health Insurance Portability and Accountability Act
HIPS	Host Intrusion Protection System (HIPS
HMAC	Hashed Message Authentication Code
IAG	Intelligent Application Gateway
IANA	Internet Assigned Number Authority
ICMP	Internet Control Message Protocol
ICT	Intra-Campus Transport
IDC	Internet Data Center
IDF	Intermediate Distribution Frame - connects to MDF, next level of distribution
IFP+	Infrastructure Fabric Plan +
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IIS	Internet Information Services
IKE	Internet Key Exchange
IOS	Internetwork Operating System

IP	Internet Protocol
IPSec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IRT	Intra-Router Transport
ISA	Internet Security Acceleration
ISAKMP	Internet Security Association and Key Management Protocol
ISATAP	Intra-site Automatic Tunnel Addressing Protocol
ISP	Internet Service Provider
IST	Intra-Site Transport
IT	Information Technology
ITII	Information Technology Infrastructure Improvement
ITO	Information Technology Operations team (the portion of IT under Paul Landray)
LAN	Local Area Network
LDOS	Last Day of Support (one of Cisco's lifecycle milestones)
LEAF	Intermediate site in the network
LIR	Local Internet Registry
LISP	Locator/ID Separation Protocol
LLx	Various levels of CT. Ins. Corp. Management (LL6 report to LL5, who report to LL4, etc.)
LOE	Level of Effort
LWAP	Lightweight Wireless Access Point
MAC	Message Authentication Code; Media Access Control
MACD	Moves, Adds, Changes, Disconnects
MAD	Merger Acquisition and Divestiture
MD5	A CT. Ins. Corp. Domain (Maintenance Domain 5)
MD5	Message-Digest Algorithm 5
MDF	Master Distribution Frame (hub of the entire network)
MDR	Main Distribution Router
MIS	Managed Internet Service
MLD	Multicast Listener Discovery
MPLS	Multi-Protocol Labeling System
MTU	Maximum Transmission Unit
NA	North America (Geographical Area in CT. Ins. Corp.'s network)
NA	Network Advertisement
NAT	Network Address Translation
ND	Neighbor Discovery
NDA	Non-Disclosure Agreement
NEOS	Network End-of-Service (CT. Ins. Corp. Term)
NetOps	Network Operations Team
NH	Next Header (Used in IPv6 IP Header)
NIR	National Internet Registry
NS	Neighbor Solicitation; Name Server

NTP	Network Time Protocol
OS	Operating System
OSC	Operational Service Center
OSPF	Open Shortest Path First
OTC	One Time Cost
PA	Provider Aggregate (IPv6 Address Type)
PDA	Personal Digital Assistant
PDT	Personal Data Terminals
PGA	Profitable Growth For All
PI	Provider Independent (IPv6 Address Type)
PIM	Protocol Independent Multicast
PIX	Private Internet eXchange
PLC	Programmable Logic Controller
PMTU	Path Maximum Transmission Unit
PMTUD	Path Maximum Transmission Unit Discovery
PoC	Proof of Concept
POC	Point of Contact or Point of Connection
QoS	Quality of Service
RA	Router Advertisement
RADIUS	Remote Authentication Dial-In User Service
RDR	Remote Distribution Router
RFC	Request for Comment
RIR	Regional Internet Registrar
RR	Resource Record
RS	Router Solicitation
RSVP	Resource Reservation Protocol
SA	South America (Geographical Area in CT. Ins. Corp.'s network)
SA	Security Association
SAR	Special Attention Review
SDR	Secondary Distribution Router
SEND	Secure Neighbor Discovery
SHA	Secure Hash Algorithm
SLA	Service Level Agreement
SLAAC	Stateless Address Auto-Configuration
SMTP	Simple Mail Transport Protocol
SNMP	Simple Network Management Protocol
SNTP	Simple Network Time Protocol
SRV	Service Location
S&R	Strategy and Roadmap
SSH	Secure Shell
SSID	Service Set Identifiers
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TLS	Transport Layer Security
TMG	Threat Management Gateway
TOS	Type of Service
TSO	Technology Process

UA	Unified Access (WLAN)
UAG	Unified Access Gateway
UC	Unified Communication
UDP	User Datagram Protocol
ULA	Unique Local Address
URL	Universal Resource Locator
VDI	Virtual Desktop Infrastructure
VLAN	Virtual LAN
VoIP	Voice of IP
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
WAN	Wide Area Network
WCS	Wireless Control System
WiFi	Wireless Fidelity
WLAN	Wireless Local Area Network
WLC	Wireless LAN Controller