TROOPERS

ERNW
providing security.

Remote Access & Business Partner Connections

Enno Rey, erey@ernw.de
@enno_insinuator

Agenda

- IPsec in IPv6 Networks
  - Some very quick remarks on current state
- Remote Access
  - The "use case" & its implications
- Business Partner Connections

IPsec in IPv6 Networks

Overall state of affairs

- IKEv2 is less complex than IKEv1
  - → better interoperability between different devices/OSs expected (& observed in real life).
    - https://blogs.technet.microsoft.com/networking/2014/12/26/vpn-interoperability-guide-for-windows-server-2012-r2/
    - https://wiki.strongswan.org/projects/strongswan/wiki/Interoperability
- RFC 6434 *IPv6 Node Requirements* strongly recommends ("SHOULD") that IPv6 nodes do have IPsec.
  - Once IPv6 is there, it will have IKEv2.
  - However, we expect some stacks to lack IPsec support, namely in IoT devices.

IPsec in IPv6 Networks

What does this mean?

- In the long run deploying IPsec might be a reasonable strategy for some parts of an organization's network.
  - Mainly those where COTS operating systems are used (e.g. data centers).
  - Identify threat/risk model and carefully evaluate sec benefit vs. operational effort.
- Key management will still be crucial!
  - Probably X.509v3 certs only viable way, once "those are on the systems anyway".
- We recommend to gain extensive operational experience with core IPv6 first.

Remote Access

Common IPv6 "initial use case" given DS-lite deployments (in cable networks)

- Questions to be clarified in advance
  - Addressing approach
  - Route propagation strategy
  - Road warriors (only) or business partners (S2S VPNs) also?
    - For the latter see below.
  - Exact VPN setup/"what does IPv6 mean?"
  - How VPN devices get their default route.

"Multiple Address Space"

Pros & Cons

- Pros
  - Consistent with initial mindset.
  - *Could* be helpful in the long-term
    - → Core of debate/speculation
- Cons
  - Creation of respective `route6` objects in different RIRs can be cumbersome/tricky.
    - In particular once outsourcing involved.
  - In the long-term potentially fragmented address space *within* global network.

"Cohesive Address Space" Approach

Pros & Cons

- Pros
  - Easier to handle wrt `route6` objects.
  - Unified address space in the long-term (as desired goal).
- Cons
  - Leads to *out-of-region* announcements
    - Good, bad, sth else?
  - Needs renumbering if probs turn up later
    - DNS is your friend.
  - Geo IP Location !?
    - Might be solvable, but considered significant issue by quite some global organizations.

Addressing Approach

Case Study / Decision Actually Taken

- For the moment go with "cohesive approach" and monitor situation/global (route) availability.
- Much easier handling with $SERVICE_PROVIDER expected.
- Allows to gain experience with
  - Out-of-region announcements
  - Provider capabilities
- We can always revert to use "multiple address space" approach.

Routing Propagation Strategy

Variants

Overall long term strategy (in case study): null-route specific prefixes which are supposed not to be reachable from untrusted networks.

- Implement long term strategy from the beginning
- For the moment go with *selective annoucements*, and monitor situation
  - As of today propagate only /48s

Start with *Selective Announcements* Strategy

Pros & Cons

See also:
https://www.troopers.de/media/filer_public/8a/6c/8a6c1e42-f486-46d7-8161-9cfef4101ecc/tr15_ipv6secsummit_langner_rey_schaetzle_slash48_considered_harmful_update.pdf

- Pro
  - one can gain experience with the approach and find out if "strict IPv6 prefix filtering" is (still) really a problem.
    One might note that currently ~45% of the IPv6 routes in the DFZ are /48s and the majority of those is without *covering aggregate*.
  - One doesn't get all the "usual noise" (network traffic from bots and the like) for a full /32 from the very beginning.
- Con
  - Potentially not aligned with long term strategy (which still might change though).

Routing Propagation Strategy

Case Study / Decision Actually Taken

- For the moment go with selective announcements (specific /48s only, see below).
  - Gain experience (not least as for $PROVIDER's maturity when it comes to route filtering & propagation).
  - Avoid noise.

VPN Use Case(s) / Setup

- Road warrior only or incl. S2S VPNs (business partners)?
  - For the latter see below.
- Keep in mind that, in context of a remote access solution, "IP connectivity" can actually mean two things:
  - Reach VPN gateways over IPv6
  - Be able to use IPv6 over/within tunnel

VPN Setup

Case Study / Decision Actually Taken

- Devices will be accessible over IPv6 but *no* IPv6 will be available within the tunnel.
  - No config of IPv6 address pools.
  - Else huge implications as for IPv6 addressing/routing in corp intranet.

Cisco AnyConnect

Bug CSCur82067

$SOMECUSTOMER runs 3.1.05187 and they don't see any problems.

Bug Search › CSCur82067

No DNS query is seen if IPv6 DNS server is configured on PHY interface

CSCur82067

Description

Symptom:
DNS query is not sent out from all of interface on the PC while connecting VPN.

Conditions:
- split-tunel is configured on the tunnel interface
- group-policy assigns DNS server for the tunnel interface (server address is outside of split tunnel)
- IPv6 is enabled and DNS server is configured on the PHY interface
- DNS server for IPv4 is also configured on the PHY interface (address is outside of split tunnel)

Workaround:
None

Further Problem Description:

Was the description about this Bug Helpful? ☆ ☆ ☆ ☆ ☆ *(0)*

Details

Last Modified: Dec 17,2015

Status: Fixed

Severity: 3 Moderate

Product: Cisco AnyConnect Secure Mobility Client

Support Cases: 7

Known Affected Releases: *(1)*
3.1(5152)

Known Fixed Releases: *(4)*
3.1(11004)
3.1(12020)
4.1(6020)
4.2(96)

Download software for Cisco AnyConnect Secure Mobility Client

How Do VPN Gateways Get Their Default Route?

Assuming they sit in $SOME_DMZ

- Perform full static configuration incl. address and default gateway
  - (Multi-) HSRP could come into play

  or

- Configure static address but learn default gateway from *Router Advertisements*
  - Clear PIO

Default Route of VPN Devices

Fully static configuration

- Pros
  - Allows highest level of control
  - Does not interfere with other devices in segment.
- Cons
  - It's against core IPv6 principles.
  - Potentially requires tedious configuration.
  - Viable long-term strategy?
  - What happens once RAs get enabled anyway?
  - Does not allow for dynamics.

Default Route of VPN Devices

Border gateways emit RAs on internal interface

- Interference with (non IPv6-enabled) devices?
  - How do they react?
  - Logging?

RAs or Not

Proposed approach in case study

- Use *router advertisements* for configuration of default routing.
  - Monitor/adapt behavior & logging on other devices in segment (firewalls).
- Reflect on interaction with (Multi-) HSRP
  - → Lab!

Assignment of Static Addresses

General procedure in "infrastructure"/ DMZ-like segments without "traditional hosts offering services"

- ::1 - ::f
  - Router interfaces (incl. HSRP), mainly default route[s].
- ::10 - :: 99
  - VPN gateways, proxies, mail gateways etc.
- ::100 -
  - all others incl. dynamically generated addresses of various systems

Assignment of Static Addresses

Additional notes/rules

- Align with bit boundaries (e.g. for ACLs) and leave some space so that at any time an additional device can be added for redundancy incl. (possibly several) VIP addresses (e.g. VRRP):
  - ::10 first functionality/device "A"
  - ::11 backup device of "A"
  - ::12 VRRP address of "A" pair/cluster
  - ::13 other VRRP address of "A"
  - ::20 second functionality/device "B"

ToDo

From case study organization

- Create `route6` objects for the involved /48 prefixes
  - Include $PROVIDER as `mnt-routes`?
- Announce routes via $PROVIDER, leading to respective DCs/site(s)
  - Monitor propagation
  - Try going with /40 once affected by *strict filtering* (keep `route6` objects in mind!)
- Configure border gateways
  - Addresses on external/internal IFs
  - Proper (w/out PIO) *router advertisements* on inside IF
- Configure VPN gateways
  - Address(es) only, default route tb learned

Business Partner Connections

Business Partner Connections

What this section is about

- In industry sectors with a deep & complex supply chain structures (e.g. *automotive*) it's quite common to have a lot of business partner connections.
  - These are not necessarily implemented by means of centralized "BP gateways" (ofc some traffic filtering still happens).
  - Direct/somewhat closed "interconnectivity" networks (e.g. ENX/ANX) might be involved.
  - In the vast majority of cases there's some N:1 NAT/masquerading involved, somewhere.
    - Else all those 10.0.0.0/8 networks would clash.

Business Partner Connections

Preliminary observation

- Pretty much all organizations having such networks/connections heavily struggle with transferring this into the IPv6 world.
  - All of them plan to use GUAs instead of RFC 1918 addresses in their internal networks.
  - N:1 NAT/masquerading is not foreseen in IPv6 anyway.
    - Some devices can actually do it, but it's not standardized (in contrast to v4/RFC 1631).

Business Partner Connections with IPv6

Potential Objectives

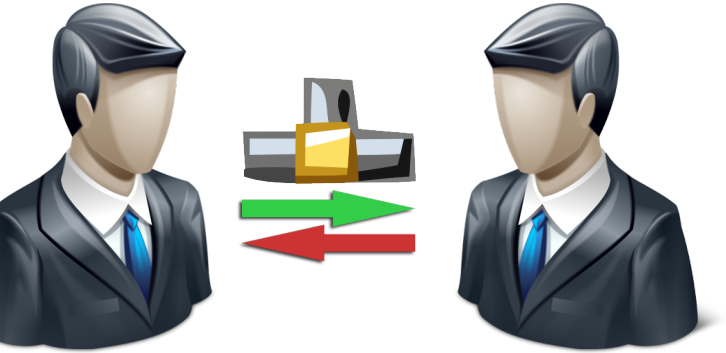

- Manageability of routing (protocols) within corporate network
- Stability of overall routing system
- Support of *routing layer security*
- (Ease of) Filtering / ACLs
- Traceability
- Play nice with your peers.

Objectives

Manageability of routing (protocols) within corporate network

- The main question here is:
  - Do we want to carry (how many?) external routes on devices/platforms within our corporate network.
- What is the potential impact on
  - memory → probably negligible.
  - CPU (when recalculating) → depends.
  - operations → depends.
- Most people we know *don't* like the idea too much.
  - But it might be less painful than alternative approaches ;-)

Manageability of Routing (Protocols)

Another aspect to be discussed

- In case you accept routes from business partners, how exactly do you get those?
  - Dynamically (= by means of $ROUTING_PROTOCOL + redistribution)
  - Static configuration on interconnection points (+ redistribution)

Objectives

Stability of overall routing system

- The main question here is:
  - "How to avoid undesired routing (protocol) interaction?"
- Undesired interaction can include:
  - $ORG inadvertently becomes transit network for BP's Internet traffic.
  - $ORG inadvertently becomes transit network for intra-site traffic of BP.
  - Route leakage.

Objectives

Support of *routing layer security*

See also:

https://www.insinuator.net/2015/12/developing-an-enterprise-ipv6-security-strategy-part-2-network-isolation-on-the-routing-layer/

- Main question:
  - What is the impact of the chosen approach on a potential 'routing layer security' strategy?
  - Does it support this/not?

Objectives

(Ease of) Filtering / ACLs

- Main questions:
  - Is traffic filtering (e.g. by means of ACLs) performed on intersection points within $ORG's network?
  - If so, what's the (operational) impact of the chosen (BP connection) approach?
  - E.g. does the presence of external networks (routes) help here, or not?
    - Can be identified more easily.
    - Grouping/aggregation probably more difficult though.

Objectives

Traceability

- Main questions:
  - When it comes to logging & log analysis (e.g. in order to identify attacks from supply chain networks), does the chosen approach support this?
- Any approach involving translation would require inventory (of mappings) which should be accessible in real-time (e.g. for CSIRT).
  - How well would that work process-wise? ;-)

Objectives

Play nice with your peers

- In one customer environment it was discussed to force business partners to use prefixes from $ORG's GUA space for (their) systems that establish connections to $ORG.
  - Remember that in IPv6 you can use multiple (global) addresses in parallel.
  - Adress selection as of RFC 3484/6724 expected to take care of...
- We didn't think this
  - would technically work very well.
  - is the right approach dealing with BPs.
    - and it could induce delays making mgrs uneasy ;-)

Objectives

Operational Feasibility

- Overall number of interconnection points or "route aggregators" might play a role here.

Business Partner Connections with IPv6

Possible Approaches

Main differentiator is IPv6 source address of business partner connection.

- Inbound connection has source address from $ORGANIZATION's GUA prefix.
  - As a native address. and/or
  - Translated through NPTv6.
- Inbound connection has source address from $PARTNER's prefix.
  - Could potentially be GUA or ULA prefix.
- Inbound connection has source address from some other prefix.
  - E.g. from trusted 3rd party network (like, in automotive, the ENX network) or some mutually agreed upon prefix.

NPTv6

RFC 6296 IPv6-to-IPv6 Network Prefix Translation

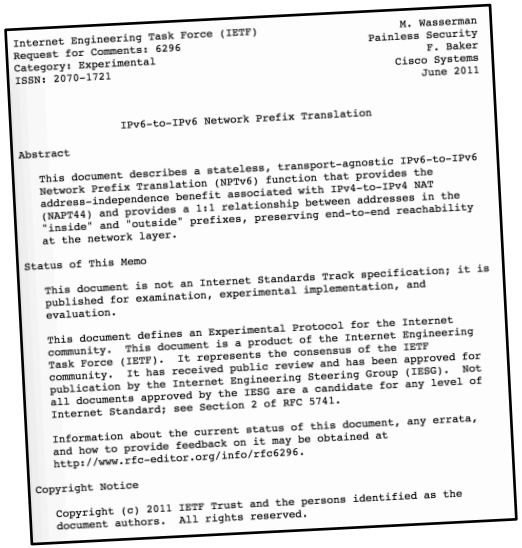

Internet Engineering Task Force (IETF)
Request for Comments: 6296
Category: Experimental
ISSN: 2070-1721

M. Wasserman
Painless Security
F. Baker
Cisco Systems
June 2011

IPv6-to-IPv6 Network Prefix Translation

Abstract

This document describes a stateless, transport-agnostic IPv6-to-IPv6 Network Prefix Translation (NPTv6) function that provides the address-independence benefit associated with IPv4-to-IPv4 NAT (NAPT44) and provides a 1:1 relationship between addresses in the "inside" and "outside" prefixes, preserving end-to-end reachability at the network layer.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6296.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

- Stateless, algorithmic translation of prefixes, performed on L3 devices.
- BP will (have to) use dedicated /48 for their side → translated to dedicated, BP-specific /48 on corp network side
  - Appropriate process to track mappings will be needed for traceability.
- Much less disruptive than NAT44 but will still impact services using IP addresses in ULP payload, e.g. FTP.
  - Gateways/ALGs necessary for those.

NPTv6

Support on commercial gear

- Most commercial firewalls (e.g. Cisco ASA, Palo Alto) support NPTv6.

- Some Cisco guy wrote to me in Oct 2015:

  "The target release is XE 3.17 which is due on CCO 11/30. This will deliver on ASR1K / ISR4K and CSR1Kv.
  For the ISR family we would be looking at a release in mid 2016 for these platforms."

  As of today (Mar 2016) it's not yet in IOS XE, but apparently it's in IOS since 15.5.3M.

Business Partner Connections over IPv6

Case study

- That's a tough one ;-)
- Allowing "external" prefixes in corp network's routing tables is not an option.
- Forcing BPs to use $ORGANIZATION's GUAs will not work out easily/nicely.
- → Approach including NPTv6 pretty much only option left.
  - We know that's a kind-of ugly one.

Business Partner Connections over IPv6

Recommendation

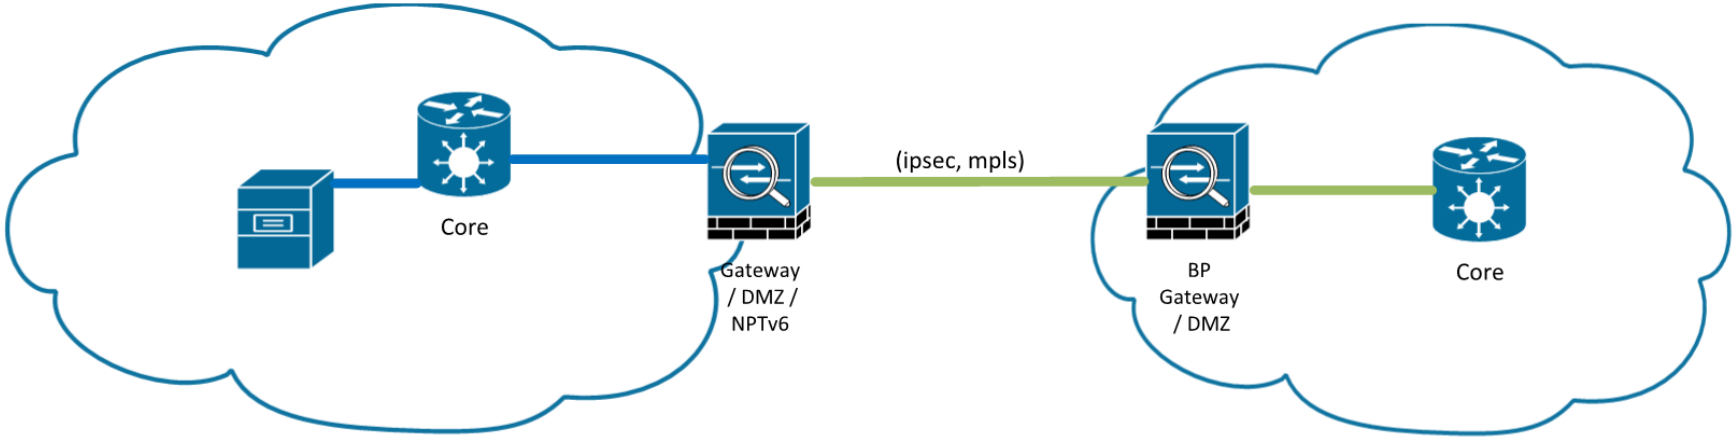

Evaluation of Objectives

Case Study Organization

Objective	BP uses own prefixes, no translation, BP routes are redistributed	BP uses own prefixes, no translation, static routes at GWs (+ redist.)	BP uses own prefixes, those are translated via NPTv6	BP uses $ORG's prefixes for segments which establish connections	BP (& $ORG) use well-known/3rd party pref. (e.g. dedicated or ENX)
Manageability of routing	2 (low)	2 (low)	5 (very high)	5 (very high)	4 (high)
Feasibility to apply filtering/ACLs within DCN	2 (low)	2 (low)	3 (medium)	3 (medium)	4 (high)
Traceability ("Nachvollziehbarkeit")	5 (very high)	5 (very high)	2 (low)	3 (medium)	3 (medium)
Support of "isolation on routing layer"	1 (very low)	1 (very low)	3 (medium)	3 (medium)	3 (medium)
Stability of overall routing system	2 (low)	4 (high)	4 (high)	4 (high)	3 (medium)
Maintaining a cooperative relationship with BPs	4 (high)	4 (high)	3 (medium)	1 (very low)	3 (medium)
Overall operational feasibility	2 (low)	1 (very low)	3 (medium)	2 (low)	2 (low)
Sum of factors (equal weight assumed)	18	19	23	21	22

Business Partner Connections over IPv6

NPTv6 based approach / Additional notes

- Technical
  - Use names for everything!
  - Ensure infrastructure supports NPTv6
  - Integration with IPAM (?)
- Processes & Politics
  - Documentation!
  - Periodic review of mappings

Conclusions & Summary

There's never enough time...

THANK YOU...

...for yours!

@Enno_Insinuator

erey@ernw.de

Slides & further information:
https://www.troopers.de
https://www.insinuator.net
(..soon)

Questions?

Image Credits

- Icons made by Freepik from www.flaticon.com are licensed by CC 3.0 BY.