

draft-anderson-v6ops-siit-eam  
draft-ietf-v6ops-siit-dc  
draft-ietf-v6ops-siit-dc-2xlat

Tore Anderson

Redpill Linpro AS

IETF 92, Dallas, March 2015



SIIT-EAM

Explicit Address Mappings for Stateless  
IP/ICMP Translation

draft-anderson-v6ops-siit-eam-03

# SIIT-EAM in a nutshell


- New draft that updates RFC6145 (SIIT)
- Allows the operator to override the default RFC6052 algorithm on a per-address (or per-prefix) basis
  - RFC6052: 192.0.2.1 <-> 64:ff9b::192.0.2.1
    - × Only IPv4-embedded IPv6 addresses are allowed
  - SIIT-EAM: 192.0.2.1 <-> 2001:db8::f00
    - ✓ Can translate arbitrary IPv6 addresses
- No matching EAM? Fallback on RFC6052

# Problem statement

- RFC6145/RFC6052 requires IPv6 nodes to be numbered using IPv4-translatable addresses
  - Fundamentally incompatible with e.g. SLAAC
- Operator choices:
  - 1) Import the IPv4 address space limitations into IPv6
  - 2) Secondary addresses on participating IPv6 nodes
    - Unaggregatable host routes in the IPv6 network
    - Duplicate ACL entries for native IPv6 and translated IPv4
    - Duplicate monitoring, duplicate config on host, etc.
    - “dual-stack light”

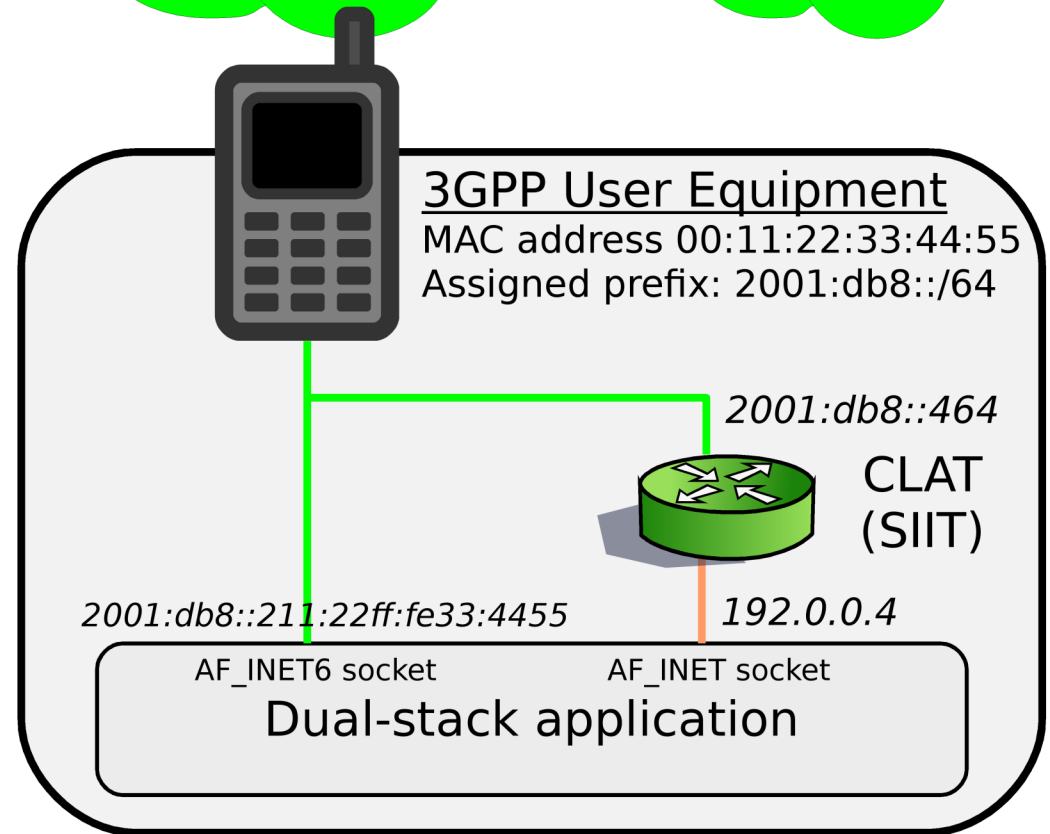
# Use case #1: 464XLAT (RFC6877)

PLAT (Stateful NAT64)



When a dedicated /64 prefix is not available for translation from DHCPv6-PD [RFC3633], the CLAT may perform NAT44 for all IPv4 LAN packets so that all the LAN-originated IPv4 packets appear from a single IPv4 address and are then statelessly translated to one interface IPv6 address that is claimed by the CLAT via the Neighbor Discovery Protocol (NDP) and defended with Duplicate Address Detection (DAD).

-- RFC6877 section 6.3



3GPP User Equipment

MAC address 00:11:22:33:44:55

Assigned prefix: 2001:db8::/64

EAM entry required in CLAT:  
192.0.0.4 <-> 2001:db8::464

Example entry for tethering:  
192.168.0.0/16 <-> 2001:db8::464:0/112

# Use case #2: IVI (RFC6219)

EAM entry in IVI translator:

0.0.0.0/0 <-> 2001:db8:ff00::/40

2001:db8:ffc0:0002:0b00::  
[2001:db8:ff{192.0:2.11}00::]

2001:db8:ffc0:0002:fe00::  
[2001:db8:ff{192.0:2.254}00::]

The IPv4 Internet



IVI XLAT  
IVI4(i): 192.0.2.0/24

SP's IPv6 network




The IPv6 Internet



2001:db8:ffc0:0002:0100::  
[2001:db8:ff{192.0:2.1}00::]

2001:db8:ffc0:0002:0a00::  
[2001:db8:ff{192.0:2.10}00::]

The IVI address format is defined based on an individual ISP's IPv6 prefix, as shown in the following figure

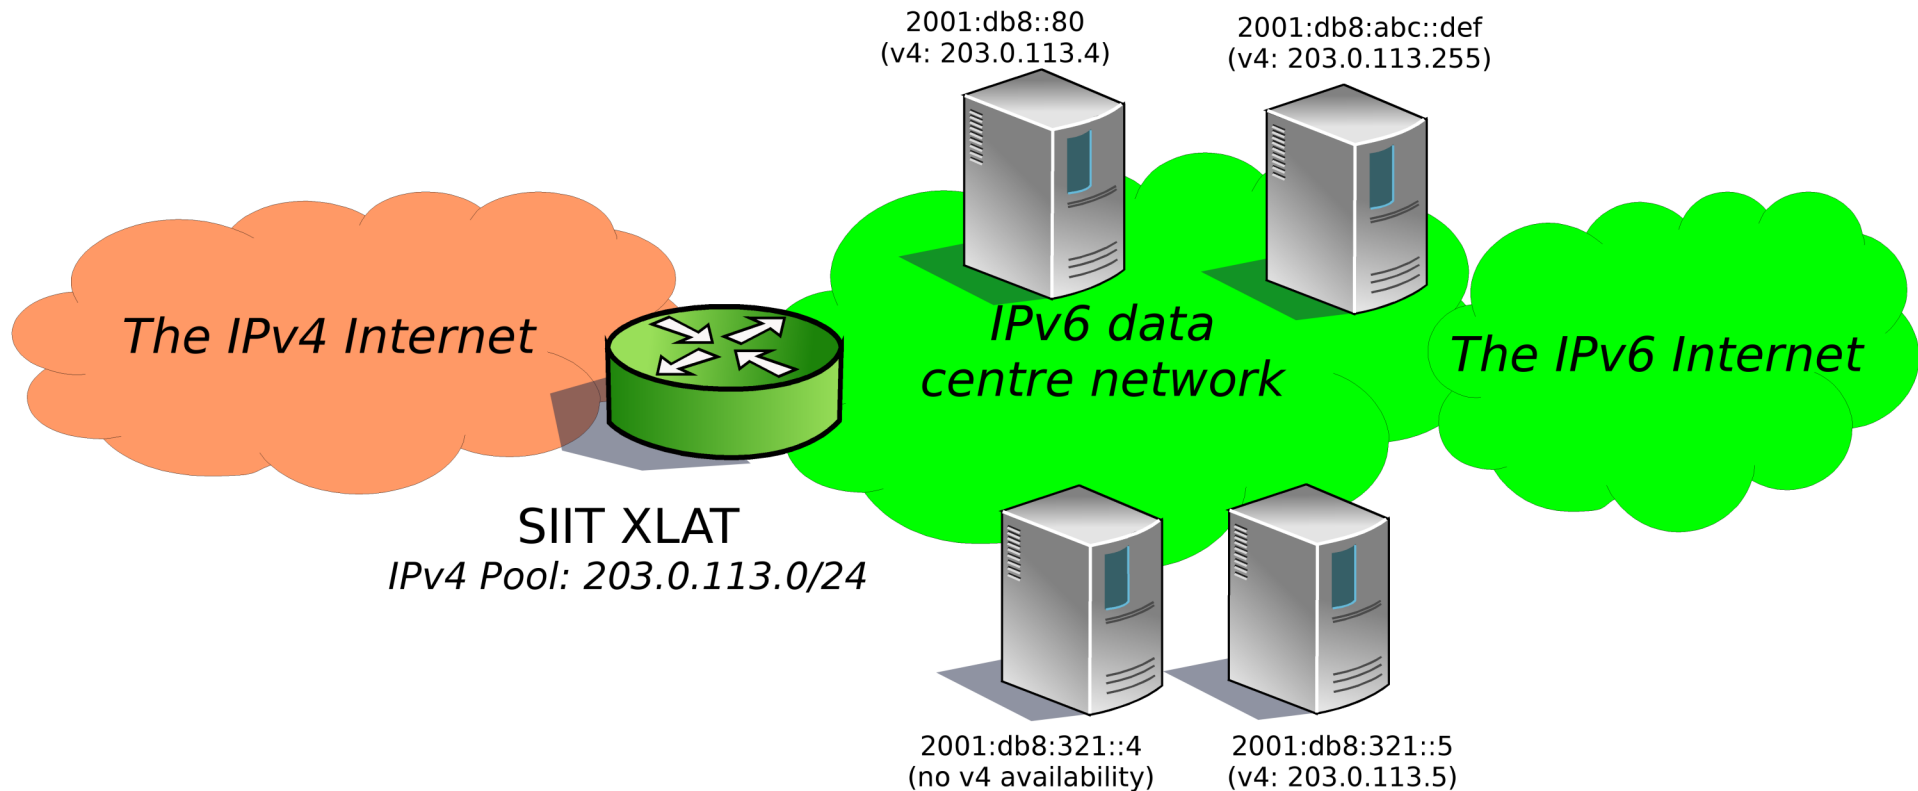


where bit 0 to bit 31 are the prefix of ISP(i)'s /32 (e.g., using document IPv6 address IPS6=2001:db8::/32) in the CERNET implementation, bit 32 to bit 39 are all ones as the identifier of the IVI addresses, and bit 40 to bit 71 are embedded global IPv4 space (IVIG4), presented in hexadecimal format (e.g., 2001:db8:ff00::/40)

-- RFC6219 section 3.1

**(Note:** IVI translation prefixes differ from 40-bit RFC6052 prefixes due to embedding the 32 IPv4 address bits in a different location of the IPv6 address.)

# Use case #3: SIIT-DC (v6ops draft)



## EAM entries in translator:

203.0.113.4 <-> 2001:db8::80  
203.0.113.5 <-> 2001:db8:321::5  
203.0.113.255 <-> 2001:db8:abc::def

*The implementation MUST allow the operator to configure an arbitrary number of Static Address Mappings which override the default [RFC6052] algorithm. It SHOULD be possible to specify a single bi-directional mapping that will be used in both the IPv4=>IPv6 and IPv6=>IPv4 directions, but it MAY additionally (or alternatively) support unidirectional mappings.*

*An example of such a bidirectional Static Address Mapping would be:*

*o 192.0.2.1 <=> 2001:db8:12:34::1*

*-- draft-ietf-v6ops-siit-dc-00 section 5.2*

# Why a separate draft for SIIT-EAM

- Dave Thaler's comment on SIIT-DC at IETF91:
  - «*This isn't a new protocol, just good use cases*»
  - Not accurate, as “vanilla” RFC6145 cannot be used with non-IPv4-embedded IPv6 addresses
- But: SIIT-EAM seems useful beyond SIIT-DC
  - For example 464XLAT and IVI
  - Therefore I believe the SIIT-DC draft never was the appropriate place for an RFC6145 update
- Adopt as WG document?
  - Allows me to remove RFC6145 update from SIIT-DC



SIIT-DC

Stateless IP/ICMP Translation for IPv6 Data  
Centre Environments

draft-ietf-v6ops-siit-dc-00

# Status

- WG adoption in IETF91; re-uploaded with appropriate new file name
  - Just minor changes since the previous version
- Future work
  - If SIIT-EAM was adopted: Remove RFC6145 protocol update from SIIT-DC and drop its status to BCP or Informational
  - If SIIT-EAM was rejected: Continue as-is with RFC6145 protocol update inside SIIT-DC draft

SIIT-DC-2XLAT


SIIT-DC: Dual Translation Mode

draft-ietf-v6ops-siit-dc-2xlat-00

# Quick recap

- Reverse SIIT-DC translation before passing IP packets to application
  - Support IPv4-only software, devices, or protocols
  - Support protocols that need end-to-end address transparency / cannot tolerate NAT (e.g., FTP)
- Inspired by 464XLAT. Notable differences:
  - “PLAT” is stateless (464XLAT: use Stateful NAT64)
  - “CLAT” is configured with public IPv4 address on the inside (464XLAT: Use private RFC 7335 IPv4)

# Host Centric translation



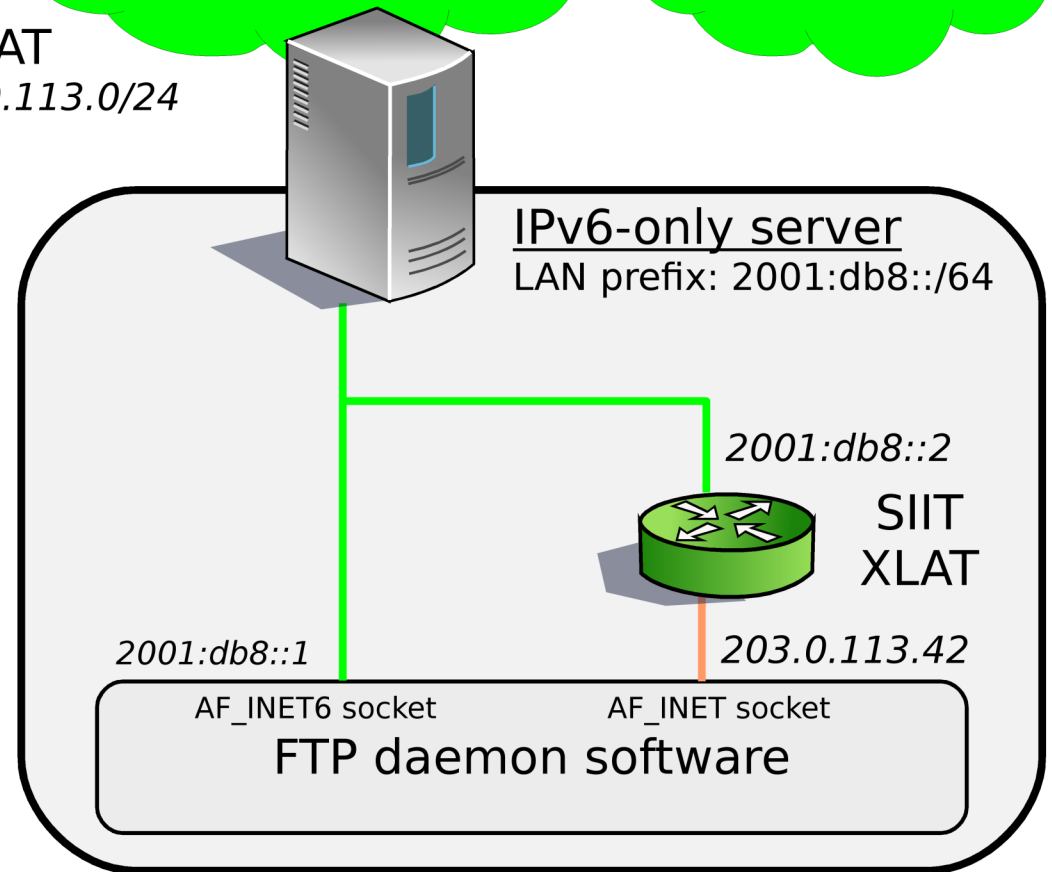
SIIT XLAT  
IPv4 Pool: 203.0.113.0/24

Use case: Server OS supports IPv6-only operation, but needs to run IPv4-only software, or an application protocol that does not support NAT (e.g., carries IP literals in application layer-7 payload)

EAM entry required in both XLATs:  
203.0.113.42 <-> 2001:db8::2

DNS configuration:

ftp.example.com. IN A 203.0.113.42  
ftp.example.com. IN AAAA 2001:db8::1



IPv6-only server  
LAN prefix: 2001:db8::/64

2001:db8::2

SIIT  
XLAT

2001:db8::1

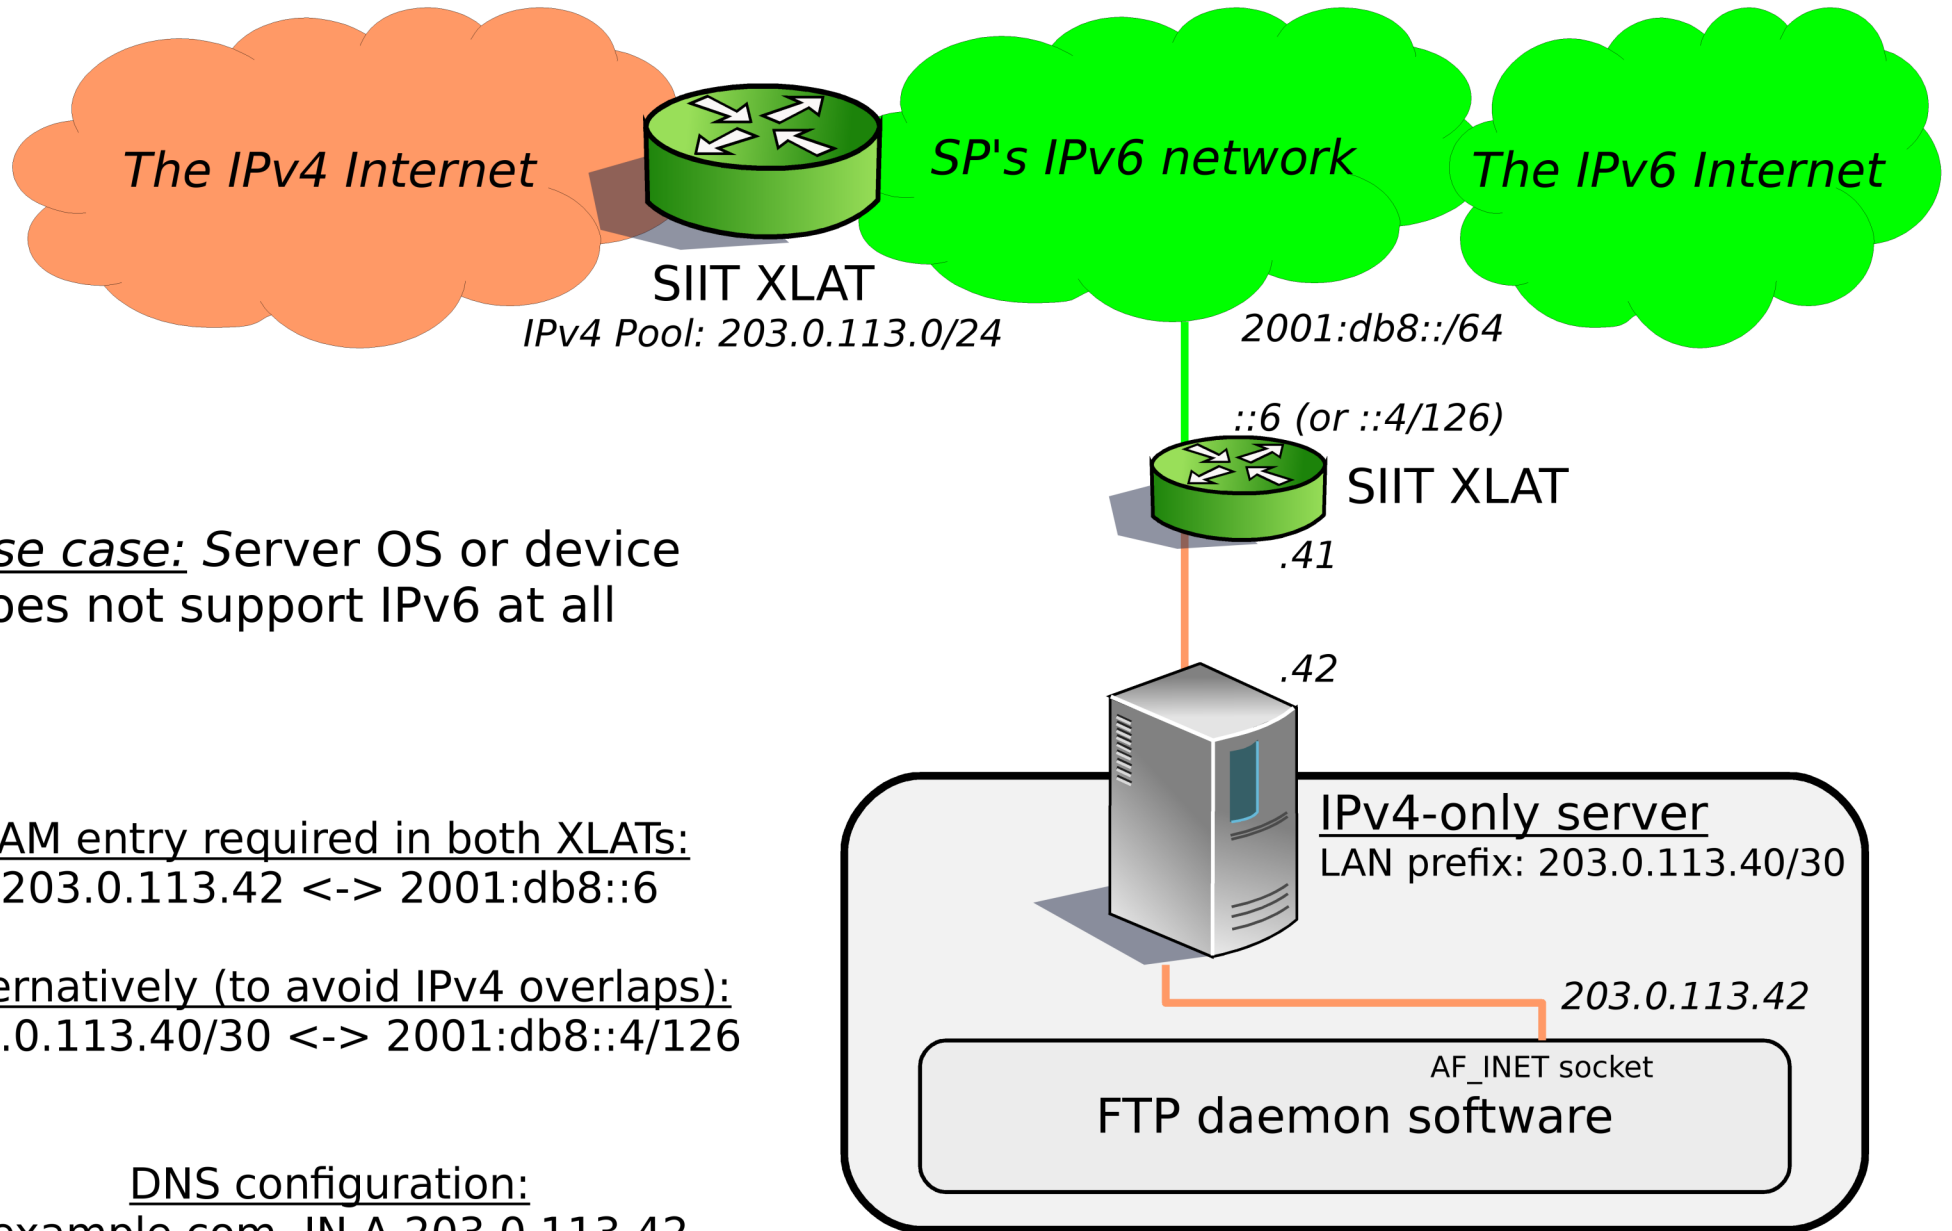
203.0.113.42

AF\_INET6 socket

AF\_INET socket

FTP daemon software

# Network Centric translation (1/2)



Use case: Server OS or device does not support IPv6 at all

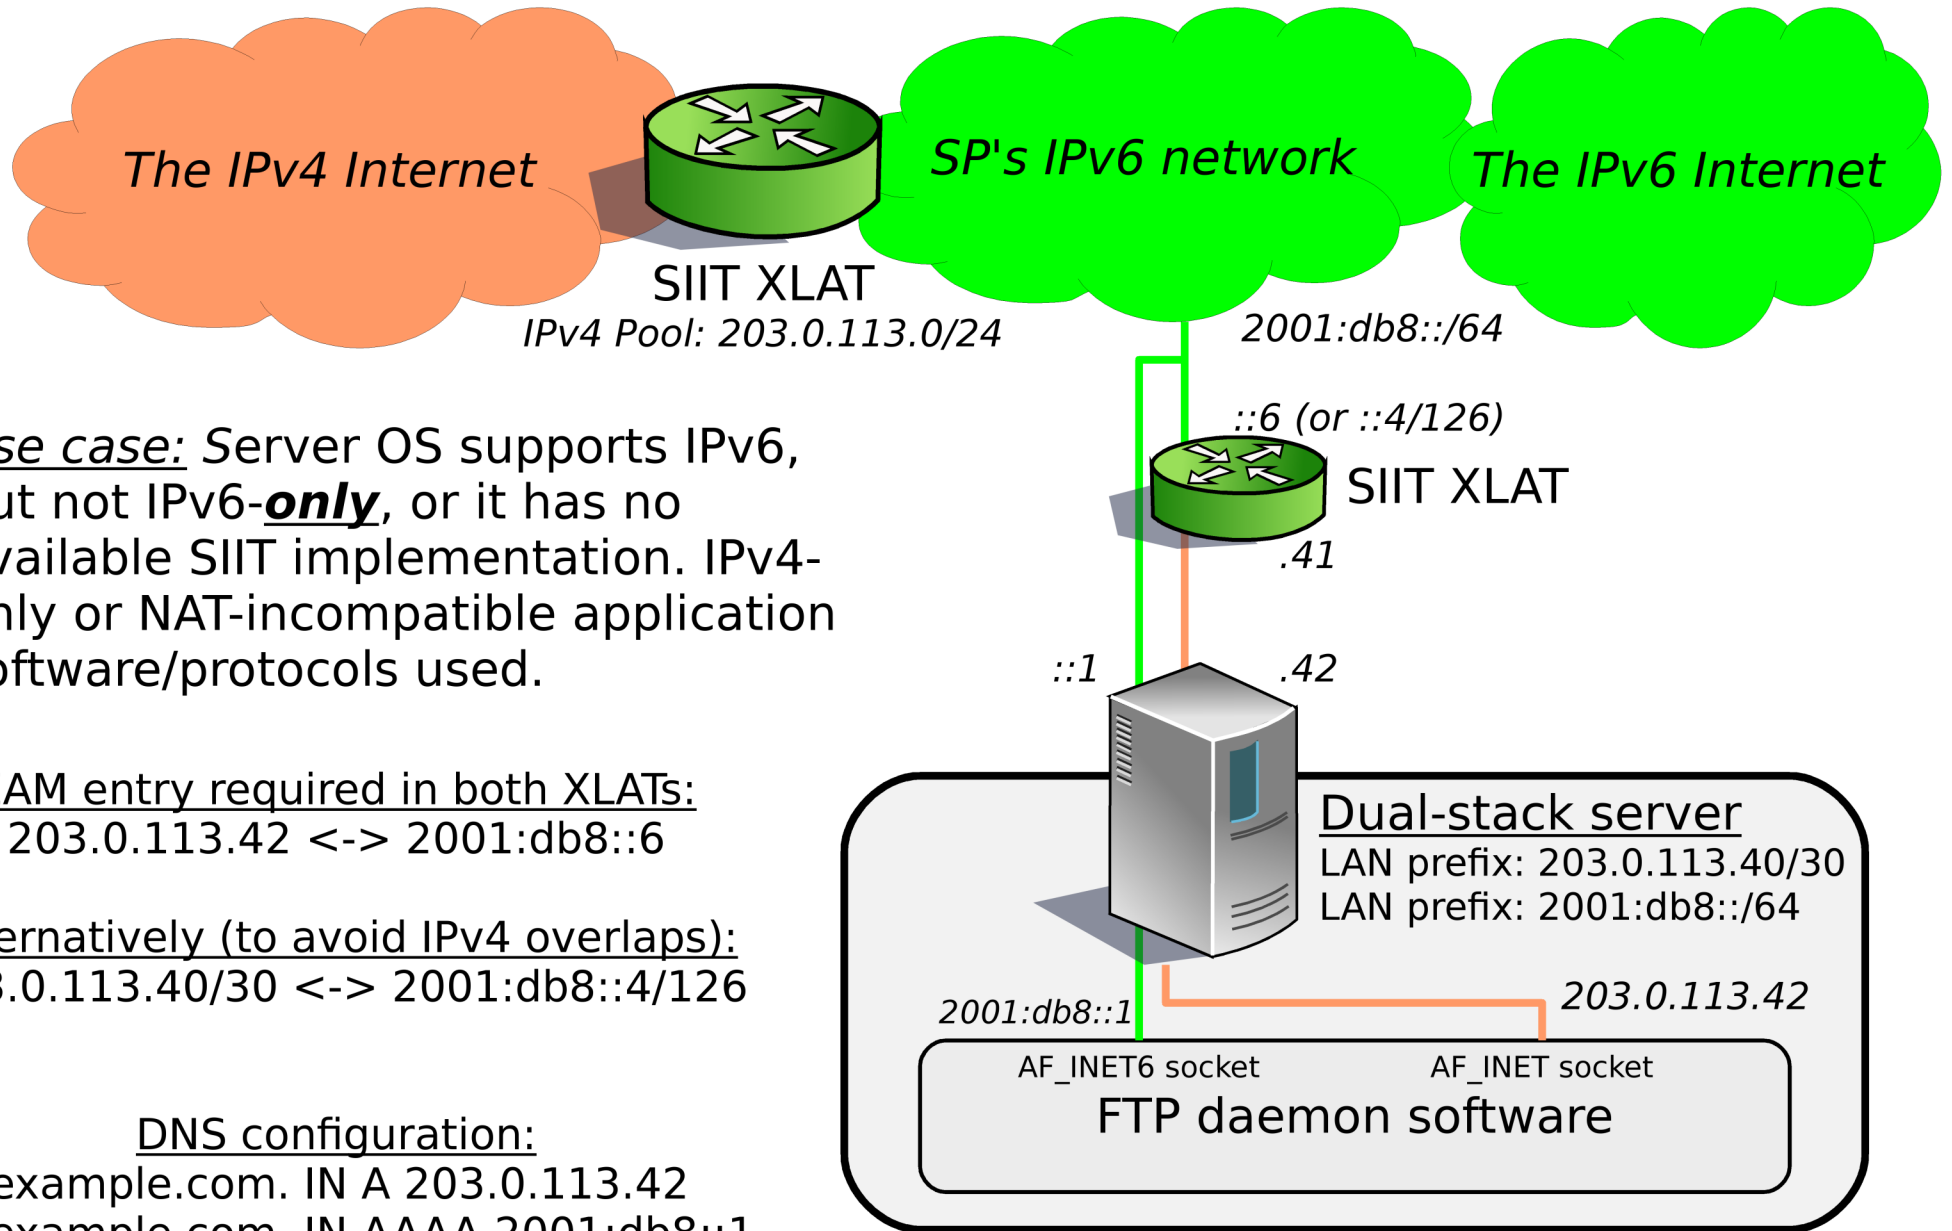
EAM entry required in both XLATs:  
203.0.113.42 <-> 2001:db8::6

Alternatively (to avoid IPv4 overlaps):  
203.0.113.40/30 <-> 2001:db8::4/126

DNS configuration:

ftp.example.com. IN A 203.0.113.42

# Network Centric translation (2/2)



Use case: Server OS supports IPv6, but not IPv6-**only**, or it has no available SIIT implementation. IPv4-only or NAT-incompatible application software/protocols used.

EAM entry required in both XLATs:  
203.0.113.42 <-> 2001:db8::6

Alternatively (to avoid IPv4 overlaps):  
203.0.113.40/30 <-> 2001:db8::4/126

DNS configuration:

```
ftp.example.com. IN A 203.0.113.42  
ftp.example.com. IN AAAA 2001:db8::1
```

# Status

- WG adoption in IETF91; re-uploaded with appropriate new file name. Major changes:
  - New Network Centric mode. Thanks to Ray Hunter!
    - “Host Agent” therefore renamed to “Edge Translator”
  - Discuss IPv4<->IPv4 intra-DC communication.  
Thanks to Andrew Yourtchenko and Shucheng LIU!
- Future work
  - Merge with SIIT-DC document or keep them separate? (compare Stateful NAT64 vs. 464XLAT)



# Running code

- Commercial appliances:
  - Brocade ServerIron ADX
  - Cisco ASR/CSR
  - F5 BIG-IP LTM
- Open source:
  - clatd (<https://github.com/toreanderson/clatd>)
  - Jool (<http://jool.mx>)
  - nat46 (<https://github.com/ayourtch/nat46>)
  - TAYGA (<http://www.litech.org/tayga>)

# Wrapping it up

- Need more reviews on the v6ops list!
  - Everyone who sends reviews are hereby invited to join me for a free beer or equivalent in Prague... :-)
  - I would also welcome co-authors or editors
- Questions, comments, feedback?

