# IPv4/IPv6 Translation: Framework

Li, Bao, and Baker

# Outcome from the Montreal Interim

- Basically, merging NAT64 and IVI to produce a common translation technology
  - Not to exclude other documents, but these form the basis

- Described in at least four documents:
  - Framework
    - draft-baker-behave-v4v6-framework
  - SIIT Update – basic translation behavior
    - draft-baker-behave-v4v6-translation
  - Extensions for stateful translation
    - draft-bagnulo-behave-nat64
  - DNS Translation gateway
    - draft-bagnulo-behave-dns64
  - Possible future documents
    - FTP ALG etc

# Scenario

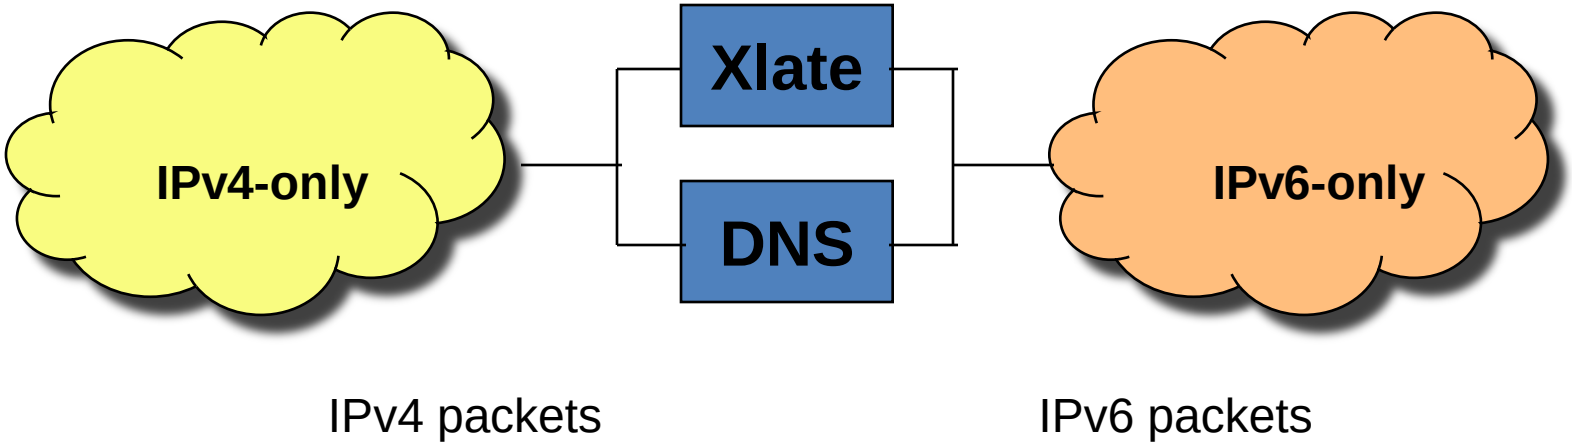


- The IPv4 packets arrived in the IP/ICMP translator will be translated to IPv6 packets.
  - The translator translates the packet headers from IPv4 to IPv6 and translate the addresses in those headers from IPv4 addresses to IPv6 addresses.
- The IPv6 packets arrived in the IP/ICMP translator will be translated to IPv4 packets.
  - The translator translates the packet headers from IPv6 to IPv4 and translate the addresses in those headers from IPv6 addresses to IPv4 addresses.

# Terminology (1)

- State
  - Refers to dynamic per-flow or per-host state
- Stateless translation
  - The translation information is carried in the address itself, permitting both IPv4->IPv6 and IPv6->IPv4 sessions establishment.
- Stateful translation
  - Translation state is maintained between IPv4 address/port pairs and IPv6 address/port pairs, enabling IPv6 systems to open sessions with IPv4 systems.

# Terminology (2)

- IPv4-mapped IPv6 address
  - The IPv4-mapped IPv6 addresses are the IPv6 addresses which have unique relationship to specific IPv4 addresses.
  - This relationship is self described by embedding IPv4 address in the IPv6 address.
- Unmapped IPv6 address
  - The unmapped IPv6 addresses are general IPv6 addresses.
  - There may exist relationship to the IPv4 addresses, but this relationship is maintained as the states (mapping table between IPv4 address/port and IPv6 address/port) in the translator.
  - The states are either manually configured or session initiated.

# Terminology (3)

- IPv4 address pool
  - In the stateful mode, a certain amount of IPv4 addresses are maintained in the translator as the IPv4 address pool.
  - In the stateless mode, there is no IPv4 address pool in the translator. A special block of IPv4 addresses are reserved, embedded in the IPv6 addresses and represented by the IPv6 end systems.

# IPv4/IPv6 Translation: temporary tool to help coexistence/transition

- IPv4 addresses
  - Embedded in an IPv6 prefix in the IPv6 domain
  - Stateless and stateful translation
- Connectivity provided:
  - IPv4 <-> IPv4
  - IPv6 <-> IPv6
  - 1:N IPv6 -> IPv4 (unmapped)
  - 1:1 IPv6 <-> IPv4 (mapped)
- Attributes:
  - Enables services in both domains
  - Stateless translation works in multiple providers, multiple translators
- Experience:
  - IVI 2 years in CERNET
  - NAT-PT/SIIT commercially deployed

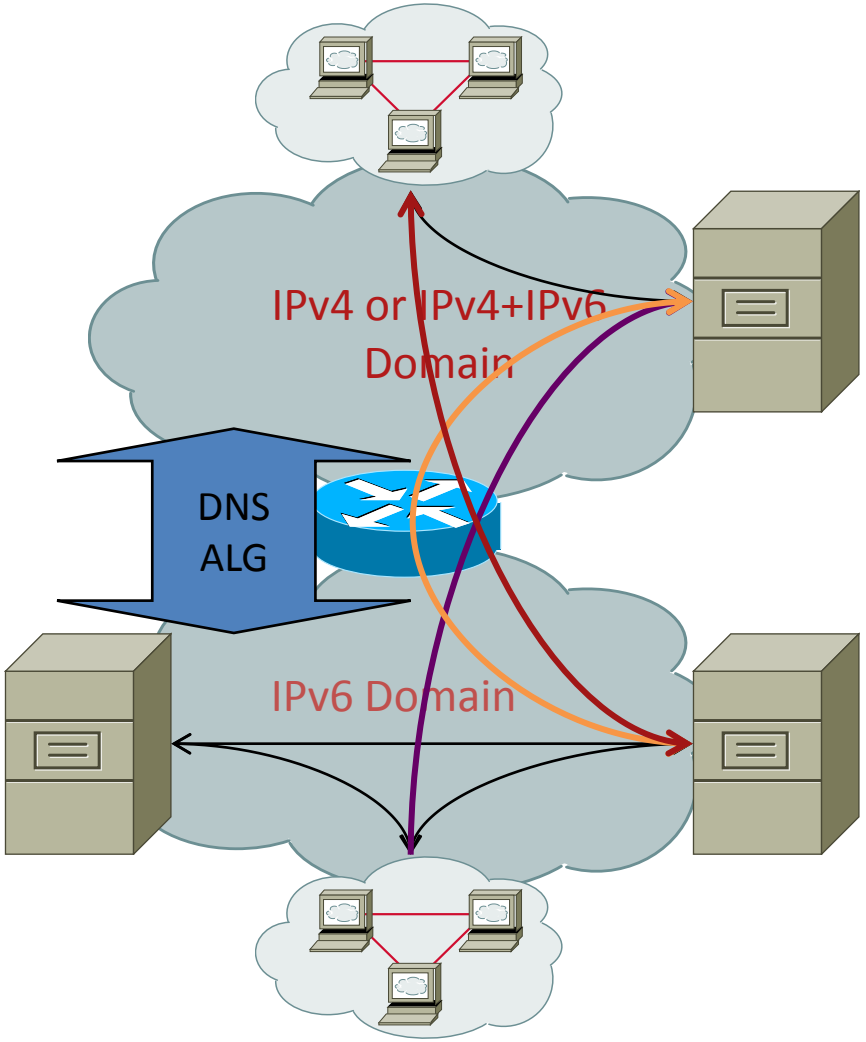

# The address format chosen

- Basic format:
  - IPv4 address embedded in IPv6 address
- Prefix: provided by the network administration
  - 0::0/3 format has routing issues with multiple translators and with multiple IPv4 domains
  - 0::0/3 format partially deprecated in RFC 4291
- Placement of IPv4 address:
  - Cook's choice: IPv4 bit 0 in IPv6 bit 33..63 or 96
  - Prefix64::/96 format appropriate for CPE and for stub IPv4 networks
  - Putting upper part of prefix in routing locator appropriate for ISP usage

# ISP usage #1

- Carrier Grade NAT, if you will
  - Designed to facilitate carrier transition with customers in various phases of transition
- Enables service:
  - IPv6 /48 or longer general prefix to customer
  - Equivalent of IPv4 /24 or longer to customer in IPv6 form for access by remote IPv4-only hosts with 1:1 stateless translation
    - Requires advertisement of /64 by edge network for IPv4-mapped IPv6 addresses
  - IPv6-only service with
    - remote IPv4 hosts accessing local mapped IPv6-only servers and
    - local IPv6 hosts accessing remote IPv4-only servers

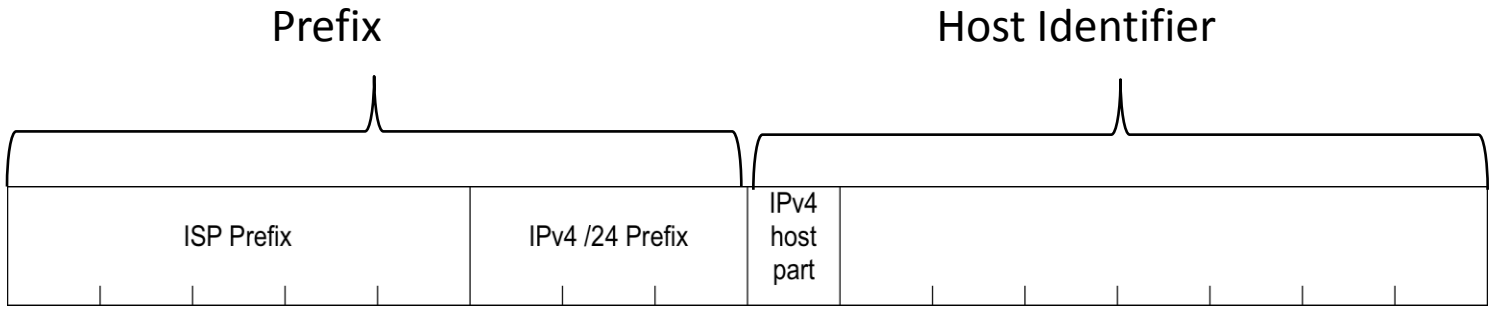

# ISP usage #2 (residential/SOHO/SMB)

- Dual stack customers around IPv6-only network
- /64..48 to customer results in
  - One /64 translated to IPv4
  - $2^n$-1 /64 IPv6 subnets
  - No IPv4-accessible servers

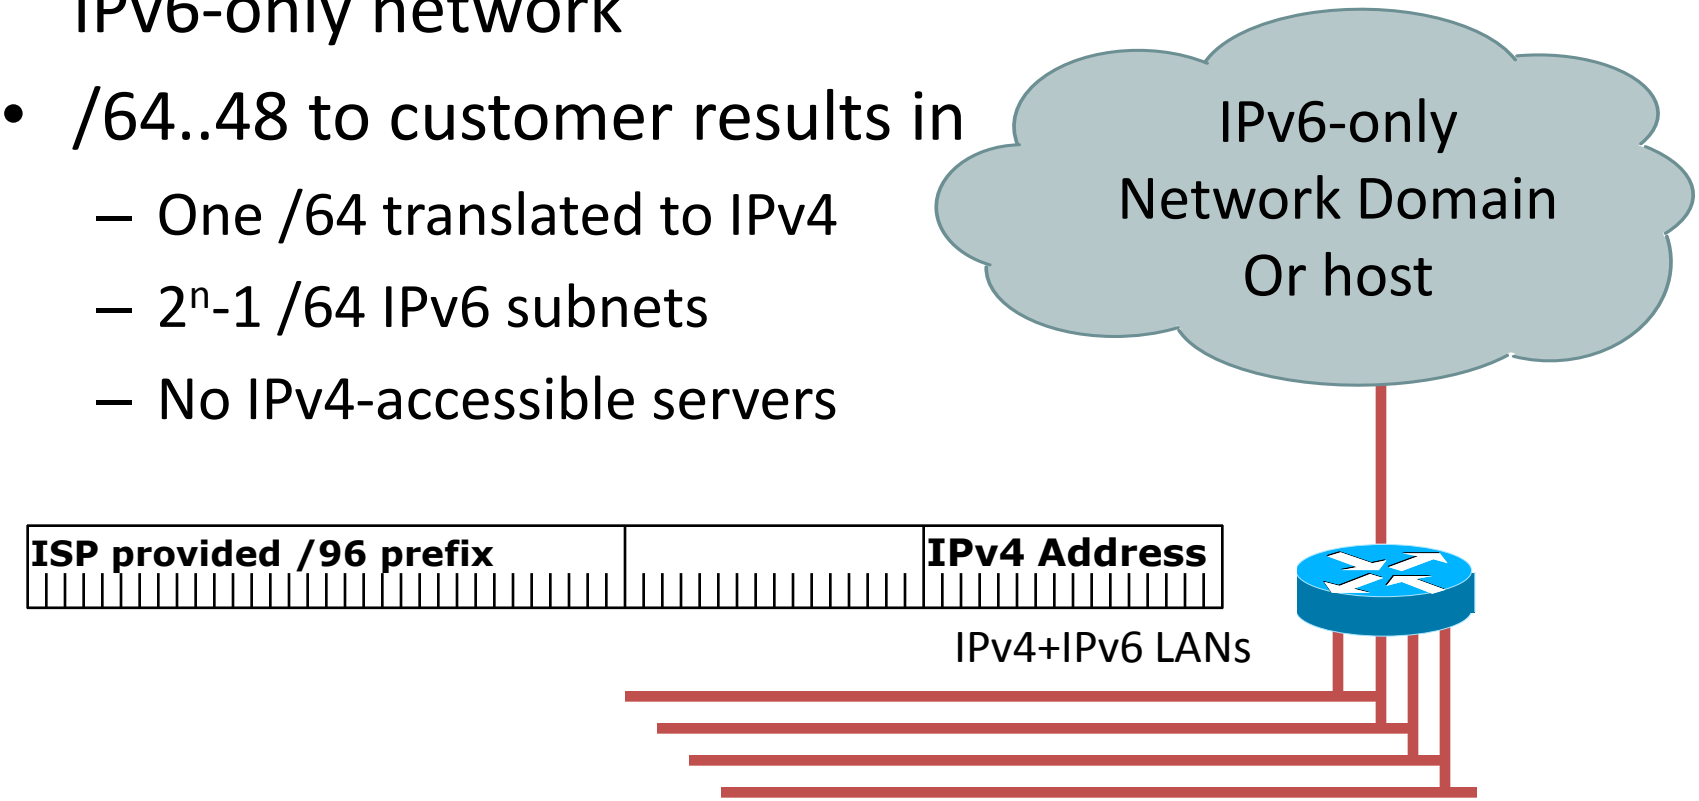

# Stub network usage: Access to legacy equipment

- IPv6-only network, IPv4-only equipment (could be dual stack but network chooses not to)
- /64 prefix to RFC 1918 space with 1:1 stateless translation

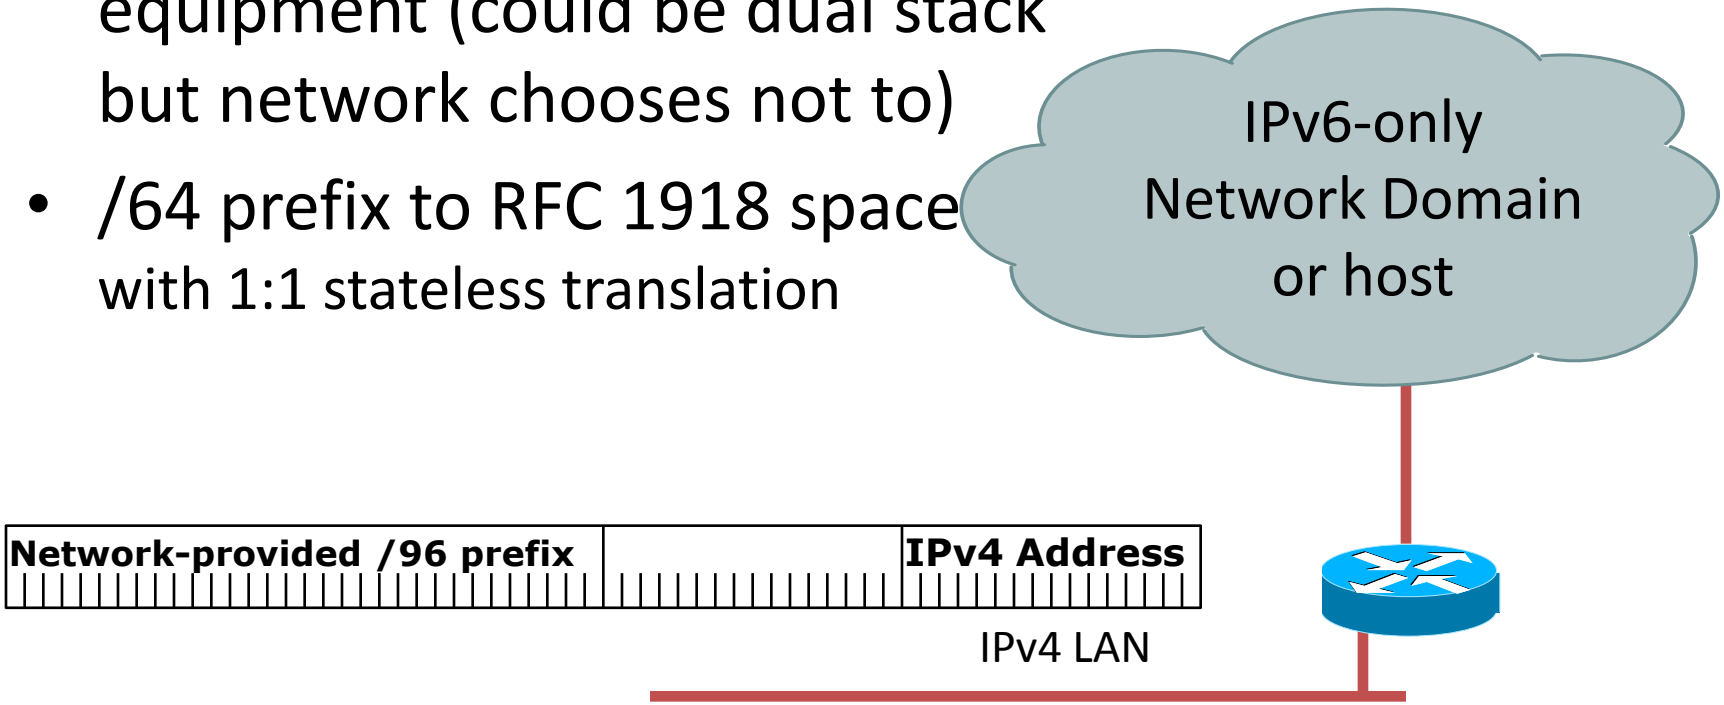

# Routing advertisements by translator

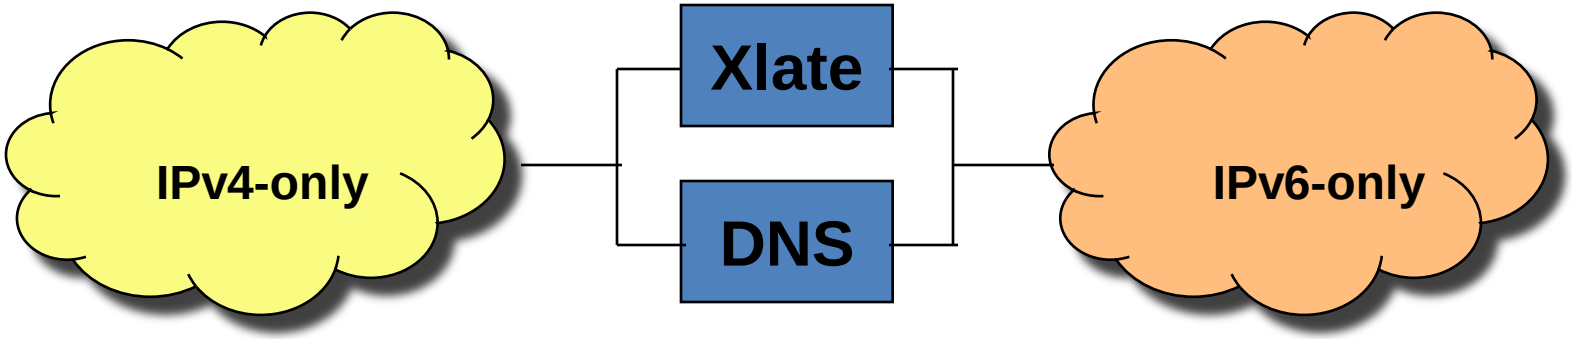


- In the IPv4 network
  - Translator advertises an IPv4 prefix for stateless translation in ISP#1 case
  - Translator advertises an IPv4 prefix for the stateful translation address pool
  - Attracts traffic destined for translation to IPv6
- In the IPv6 network
  - Translator advertises an IPv6 prefix for entire IPv4 address space
  - Attracts traffic destined for translation to IPv4

# Usage of 1:n translation

- Primarily to let IPv6-only hosts with general format addresses access IPv4-only servers/peers
- IPv4 access to general IPv6 hosts excluded due to complexity

# Usage of DNS translator

- Client/Server and Peer/Peer
  - Enable IPv6 hosts with mapped addresses to be accessible to IPv4 clients/peers
  - Enable IPv4 hosts to be accessed by IPv6 clients/peers
- Designed for simplicity and maintainability
  - Simplest case is static configuration of records
  - Capable of dynamic translation A<->AAAA
  - Capable of multiple DNS servers with predictable results and no state other than DNS caches