IPv6 Transition and NAT

Building Bridges

Carsten Strotmann

CREATED: 2025-01-30 THU 09:43

IPv6 Address Translation (NAT)

IPv6 NAT

- IPv4 Network Address Translation (Cone-NAT) has broken the Internet "end-to-end" connectivity paradigm
- "end-to-end" communication allows more resilient, decentralized networks
- IPv6 aims to re-establish "end-to-end" communication between hosts

IPv6 NAT

- For various reasons, network address translation might still be required in IPv6
  - For IPv4 to IPv6 transition technologies
  - For policy reasons
  - For network topology reasons

NAT and IPv6

- IPv6 tries to avoid NAT (Network Address Translation) if possible
  - NAT has many negative side effects on Internet protocols
- Sometimes, NAT can be useful even with IPv6
  - But be careful - NAT should be an exception in IPv6
- See ↪RFC 5902 (Informal) "IAB Thoughts on IPv6 Network Address Translation"

NAT and IPv6

- Use cases for IPv6 NAT:
  - Avoiding renumbering
  - Facilitating multihoming
  - Making configurations homogeneous
  - Hiding internal network details
  - Providing simple security

NAT and IPv6

- ↪RFC 6052 "IPv6 Addressing of IPv4/IPv6 Translators" (Standards Track) defines a well known IPv6 prefix for (IPv6 to IPv4) address translation: `64:ff9b::/96`
  - This prefix is *checksum neutral*.
  - The sum of the hexadecimal numbers `0064` and `ff9b` is `ffff`, i.e., a value equal to zero in one's complement arithmetic.
  - An IPv4-embedded IPv6 address constructed with this prefix will have the same one's complement checksum as the embedded IPv4 address.

IPv6 NAT

- We will cover
  - MAP - Mapping of Address and Port
  - IPv6-to-IPv6 Network Prefix Translation
  - 464XLAT
  - NAT64/DNS64
  - DS-Lite/lightweight DS-Lite
  - Carrier Grade NAT (CGN) / challenges with CGN NAT
  - deprecated IPv6 NAT technologies

Network Prefix Translation

Network Prefix Translation

- ↪RFC 6296 "IPv6-to-IPv6 Network Prefix Translation" (experimental) defines NPTv6
  - End-to-end reachability is preserved, the prefix is translated, the interface-id is kept
  - Firewalls are required to control the flow of traffic

- Prefix translators do not store or need to share state, multiple independent translators possible
- Translation is 1:1 at the network layer, there is no need to modify port numbers or other transport parameters (but IP addresses in upper layer protocols create issues)

IPv6 NPT (1/8)

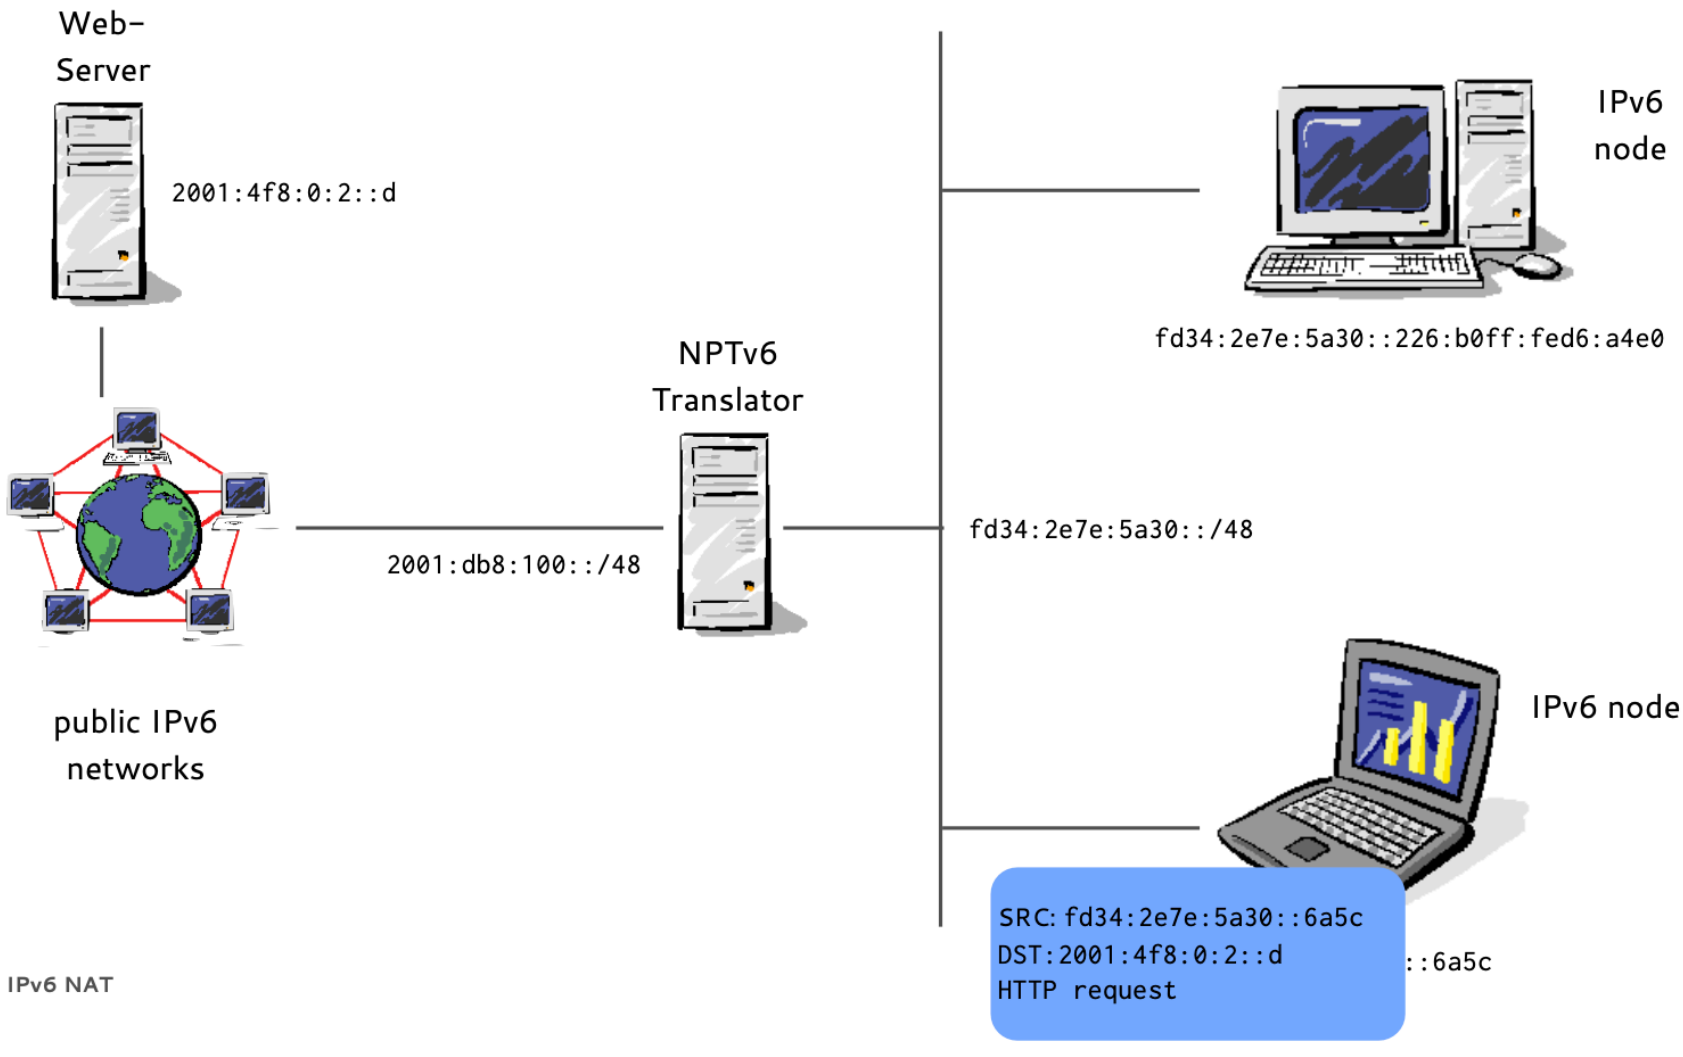
Web-
Server
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
2001:db8:100::/48
fd34:2e7e:5a30::/48
public IPv6
networks
IPv6 node
SRC: fd34:2e7e:5a30::6a5c
DST:2001:4f8:0:2::d
HTTP request
::6a5c

IPv6 NAT

IPv6 NPT (2/8)

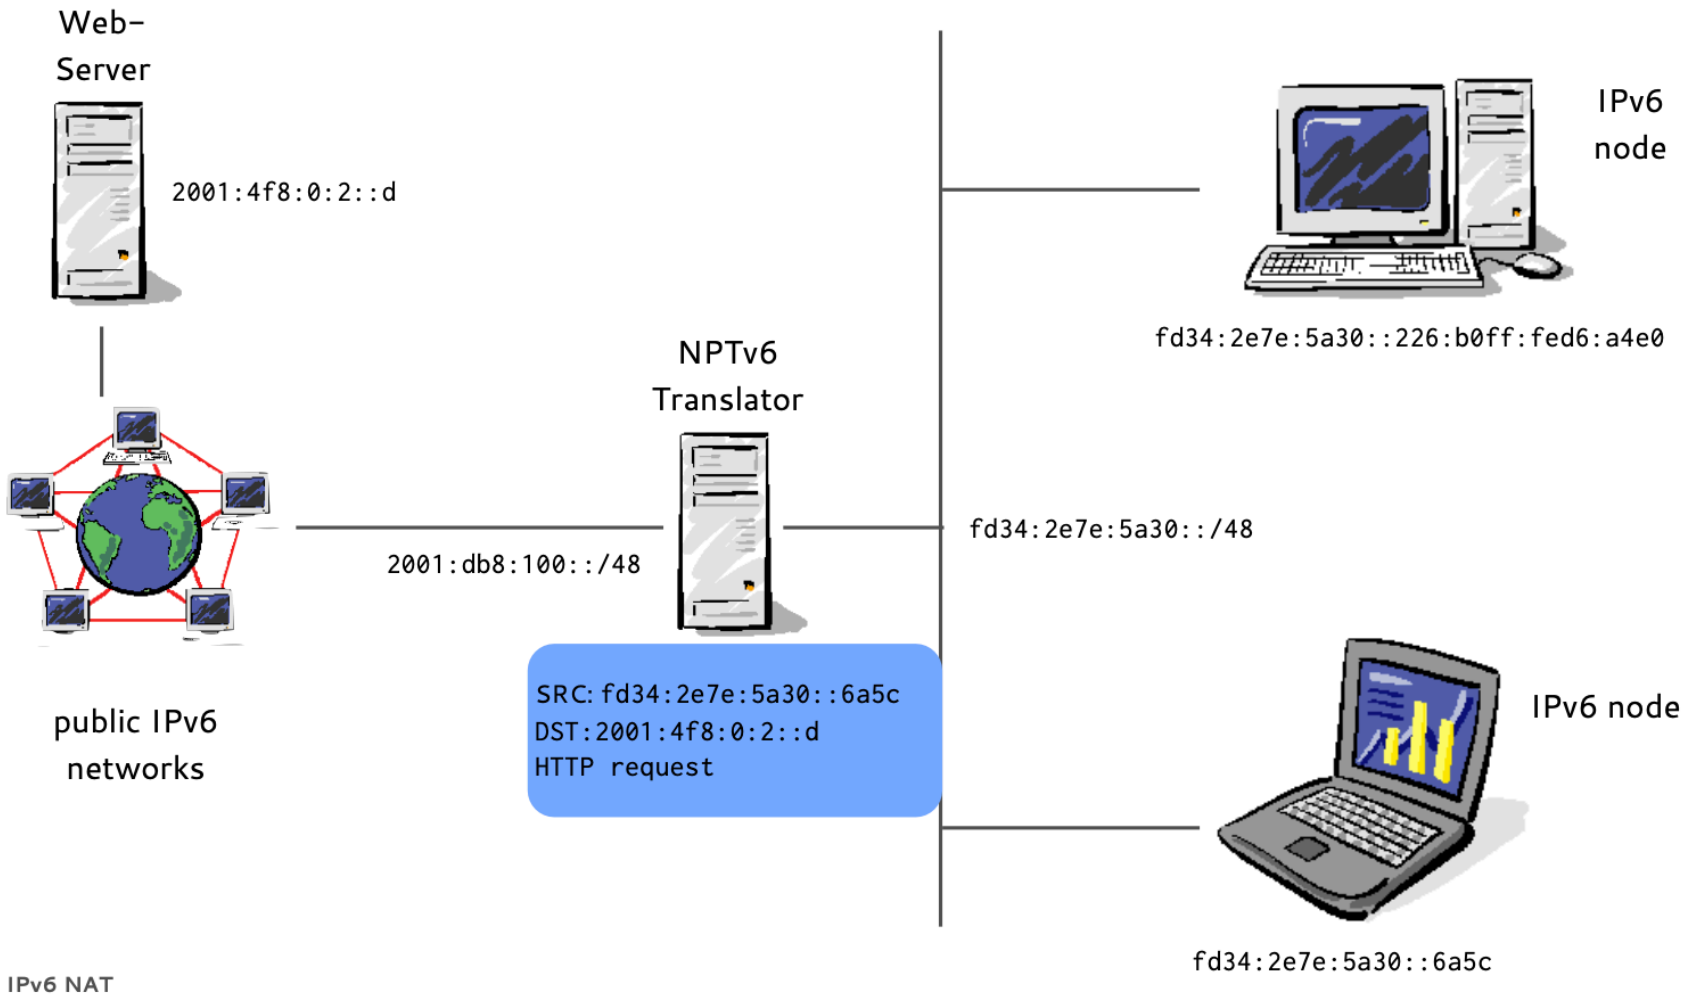
Web-
Server
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
public IPv6
networks
2001:db8:100::/48
fd34:2e7e:5a30::/48
SRC: fd34:2e7e:5a30::6a5c
DST:2001:4f8:0:2::d
HTTP request
IPv6 node
fd34:2e7e:5a30::6a5c

IPv6 NAT

IPv6 NPT (3/8)

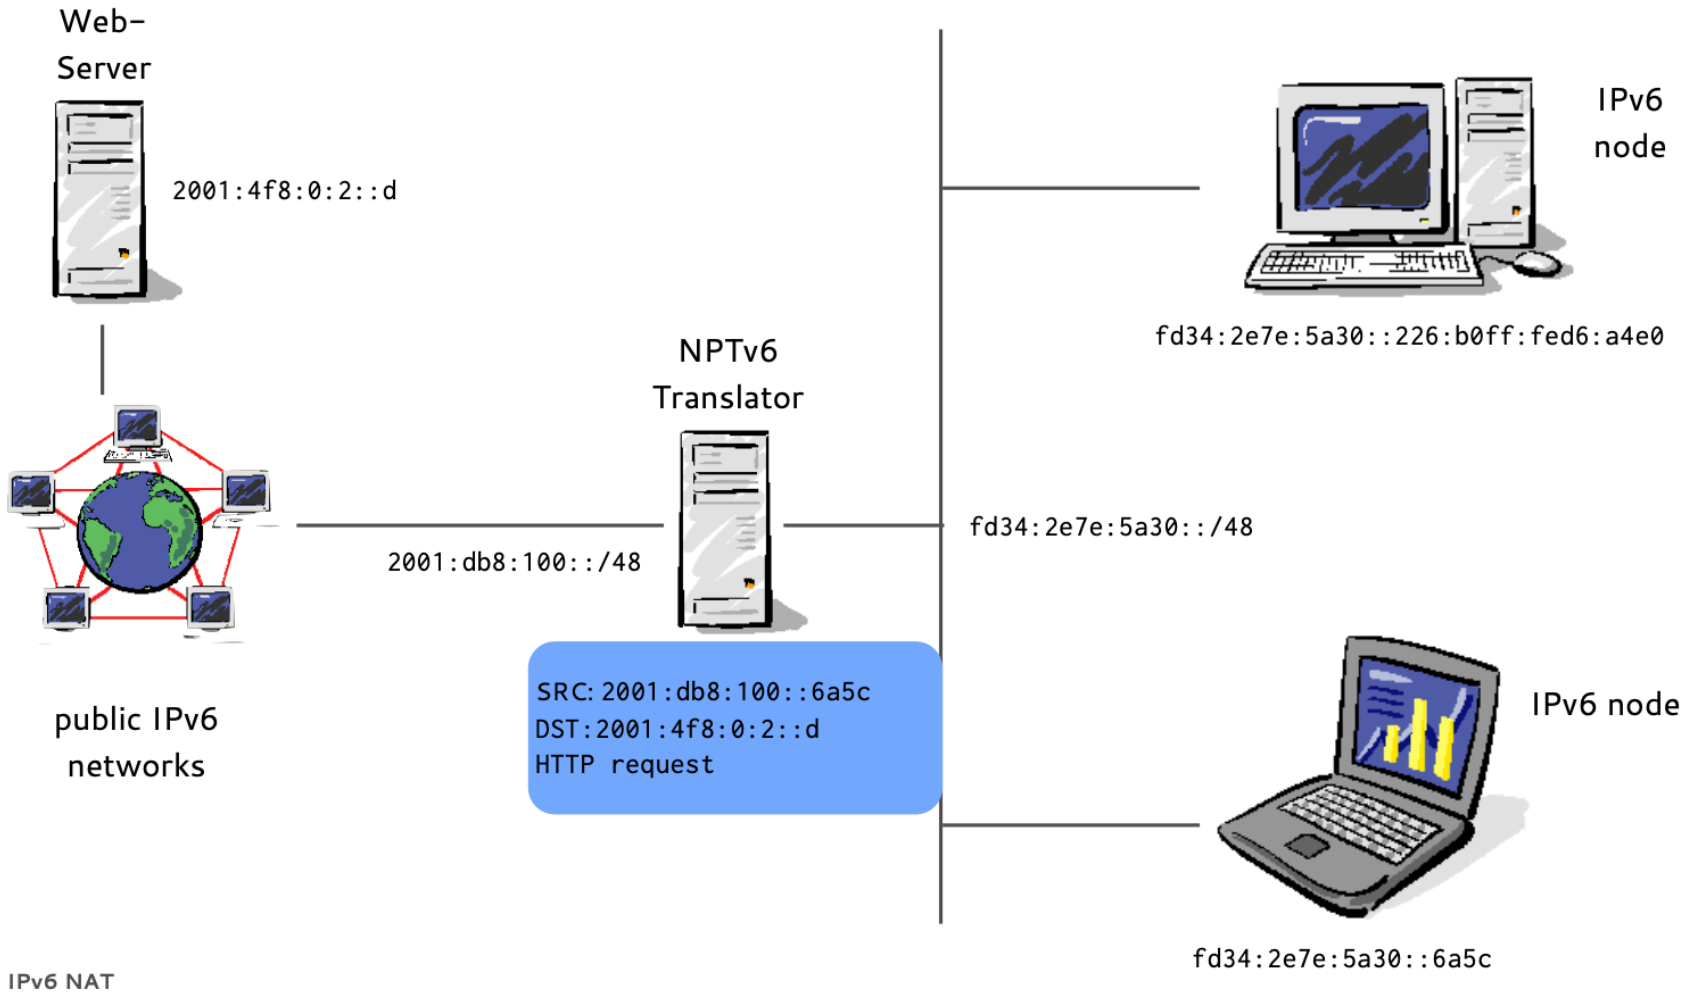
Web-
Server
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
2001:db8:100::/48
fd34:2e7e:5a30::/48
public IPv6
networks
SRC: 2001:db8:100::6a5c
DST:2001:4f8:0:2::d
HTTP request
IPv6 node
fd34:2e7e:5a30::6a5c

IPv6 NAT

IPv6 NPT (4/8)

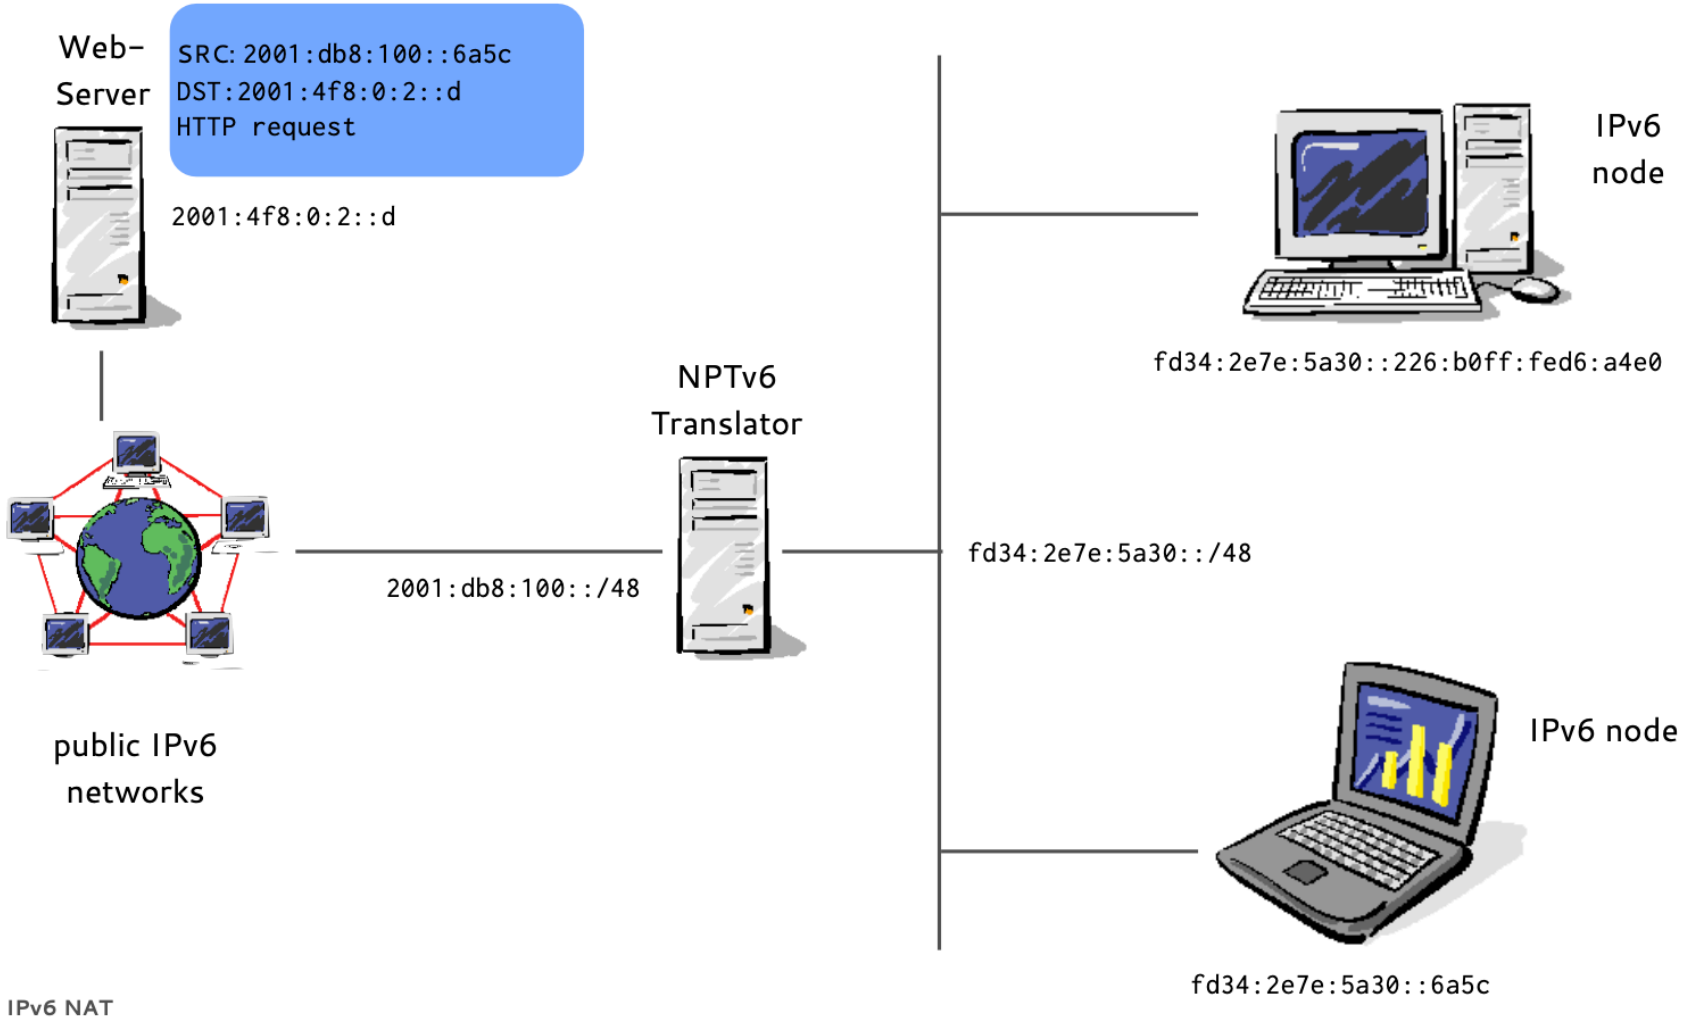
Web-
Server
SRC: 2001:db8:100::6a5c
DST:2001:4f8:0:2::d
HTTP request
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
2001:db8:100::/48
fd34:2e7e:5a30::/48
public IPv6
networks
IPv6 node
fd34:2e7e:5a30::6a5c

IPv6 NAT

IPv6 NPT (5/8)

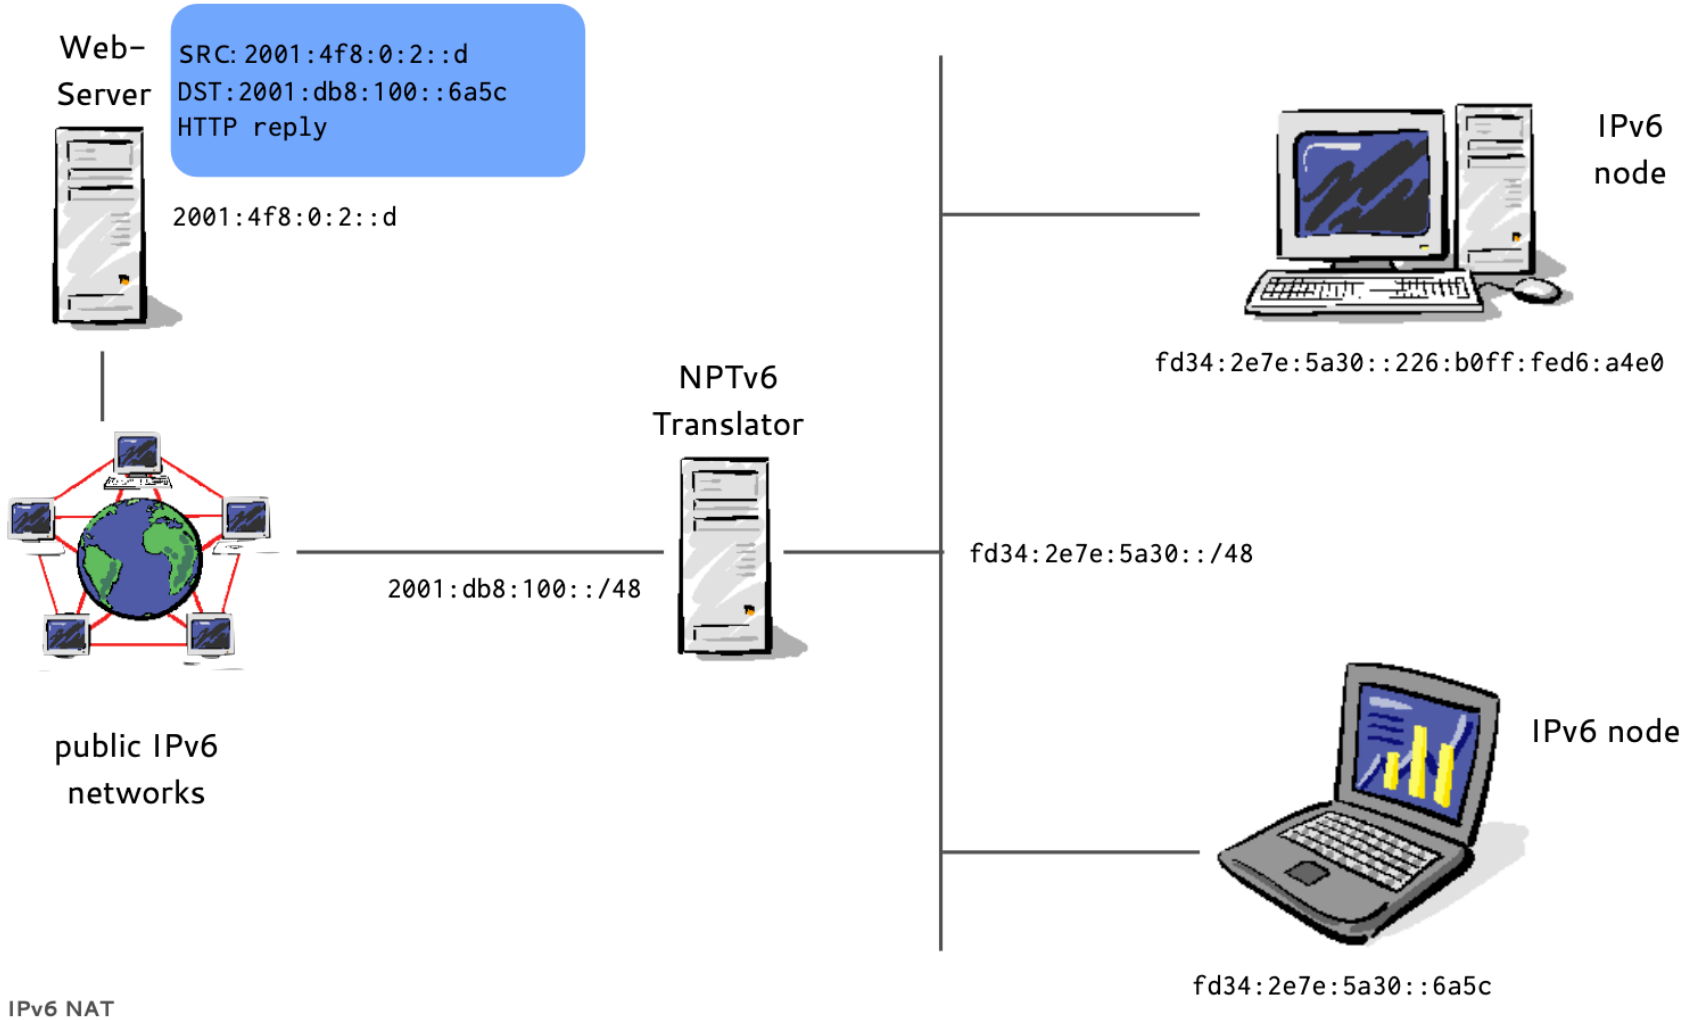
Web-
Server
SRC: 2001:4f8:0:2::d
DST:2001:db8:100::6a5c
HTTP reply
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
2001:db8:100::/48
fd34:2e7e:5a30::/48
public IPv6
networks
IPv6 node
fd34:2e7e:5a30::6a5c

IPv6 NAT

IPv6 NPT (6/8)

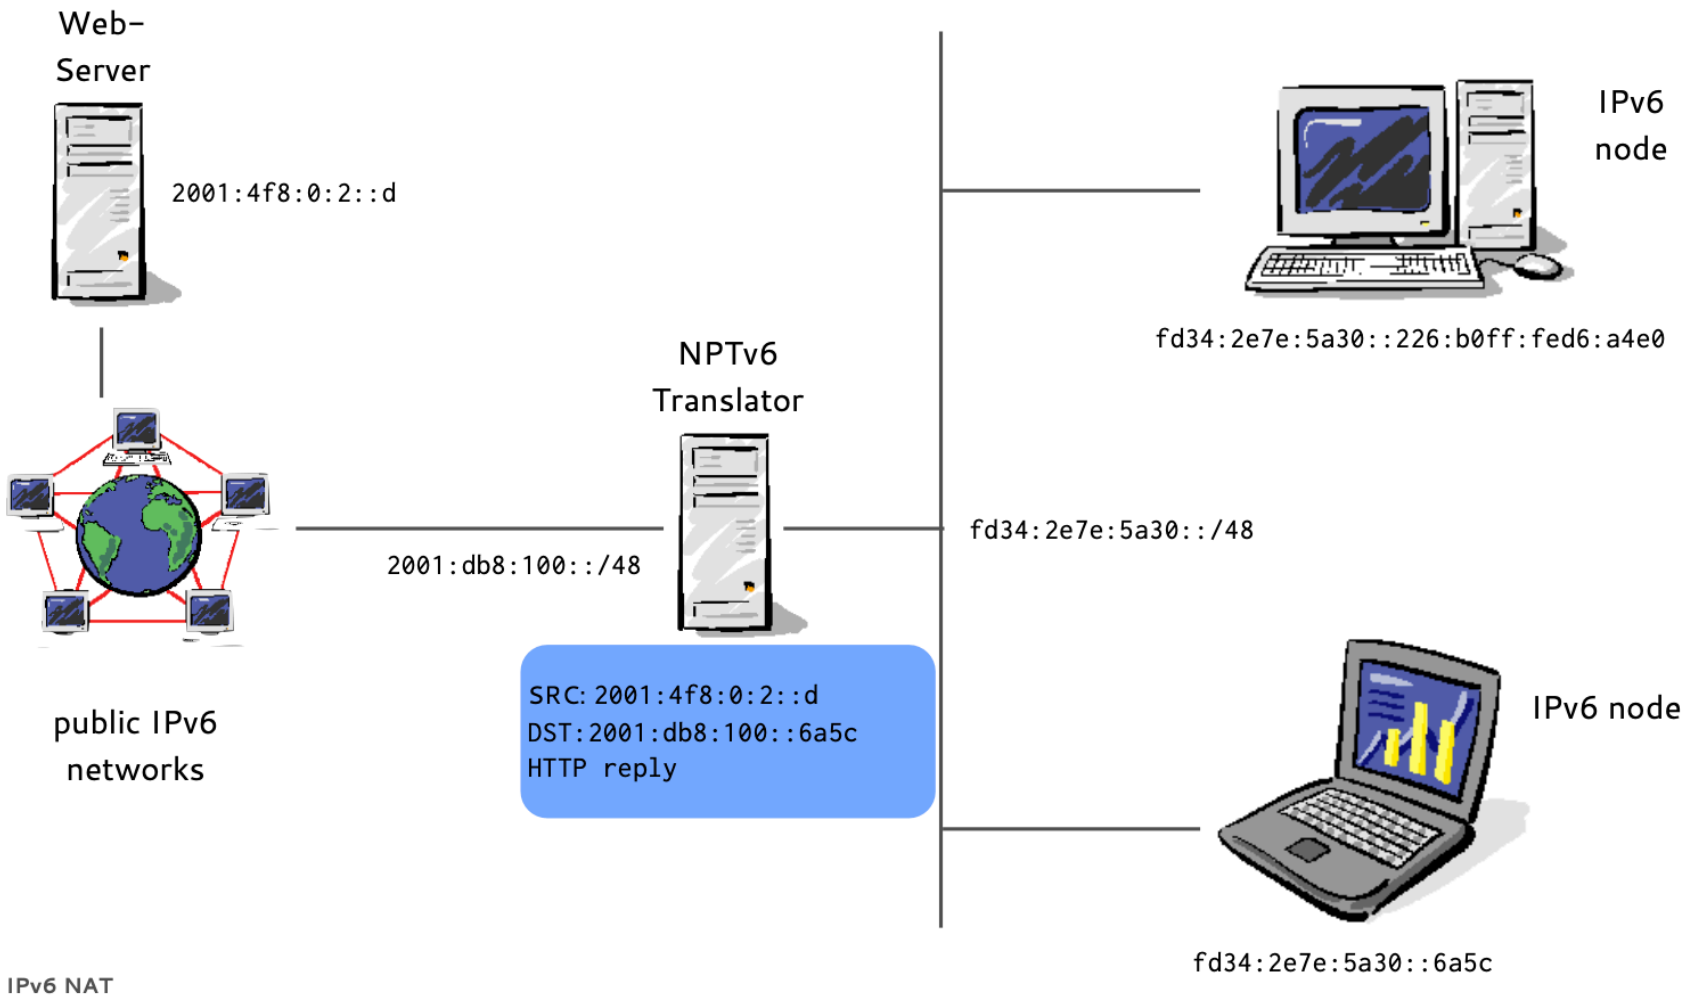
Web-
Server
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
2001:db8:100::/48
fd34:2e7e:5a30::/48
public IPv6
networks
SRC: 2001:4f8:0:2::d
DST:2001:db8:100::6a5c
HTTP reply
IPv6 node
fd34:2e7e:5a30::6a5c

IPv6 NAT

IPv6 NPT (7/8)

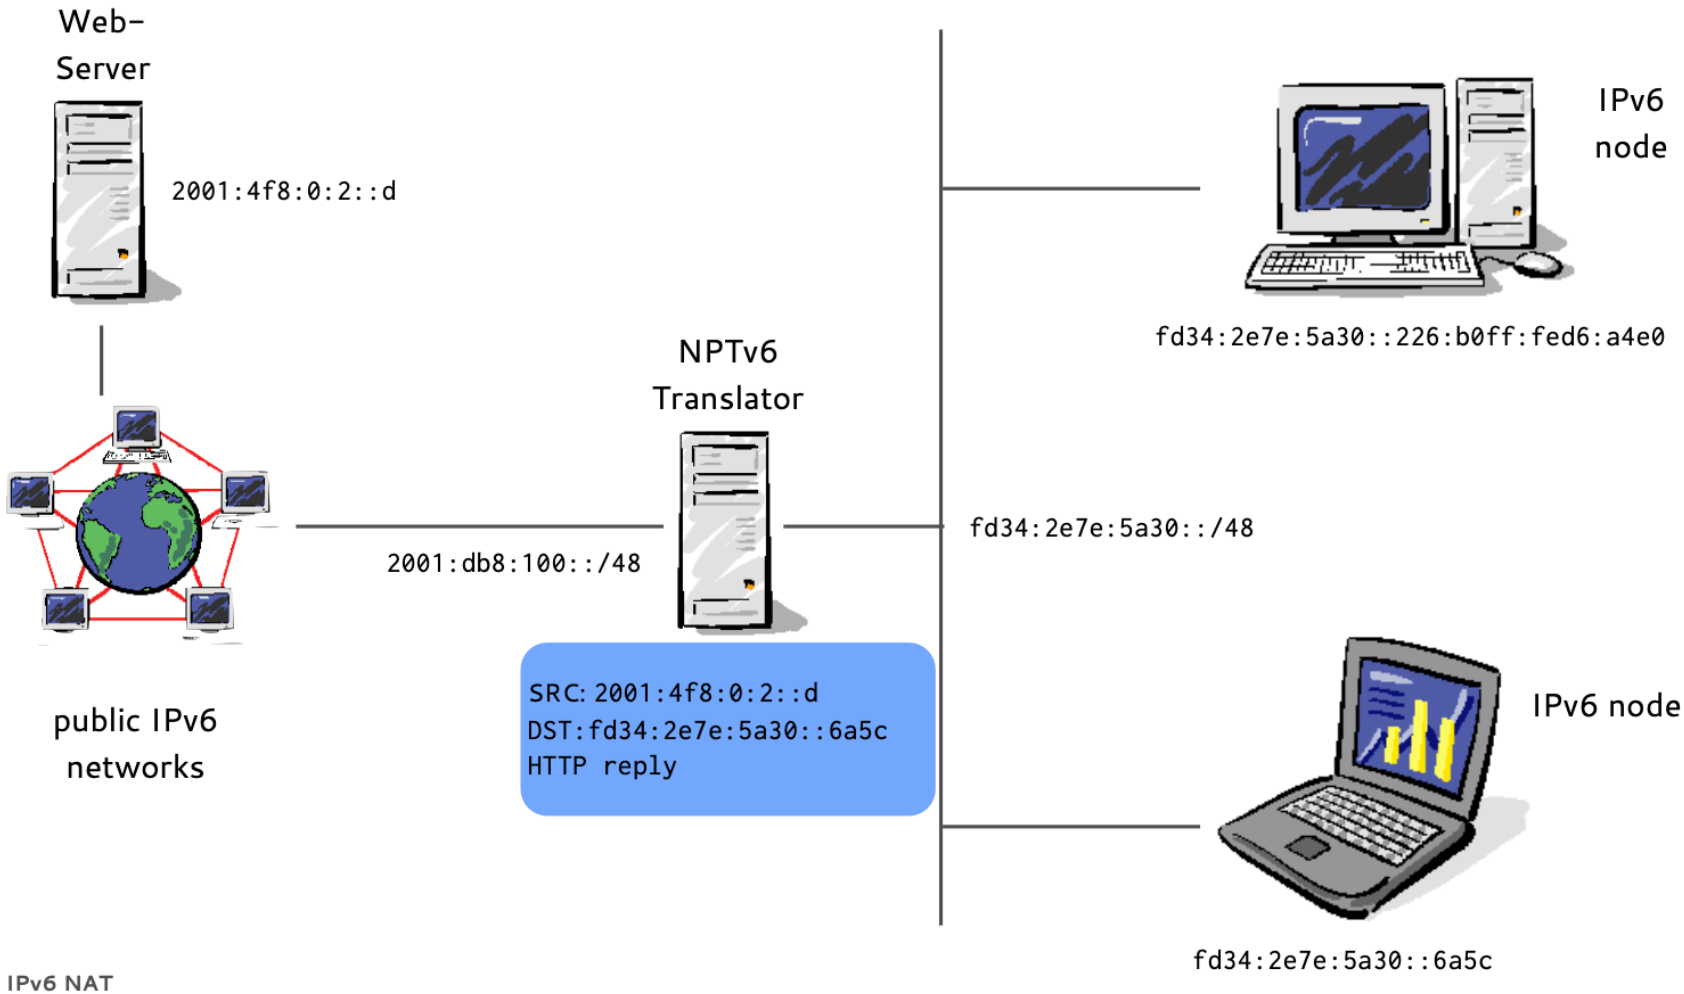
Web-
Server
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
2001:db8:100::/48
fd34:2e7e:5a30::/48
public IPv6
networks
SRC: 2001:4f8:0:2::d
DST:fd34:2e7e:5a30::6a5c
HTTP reply
IPv6 node
fd34:2e7e:5a30::6a5c

IPv6 NAT

IPv6 NPT (8/8)

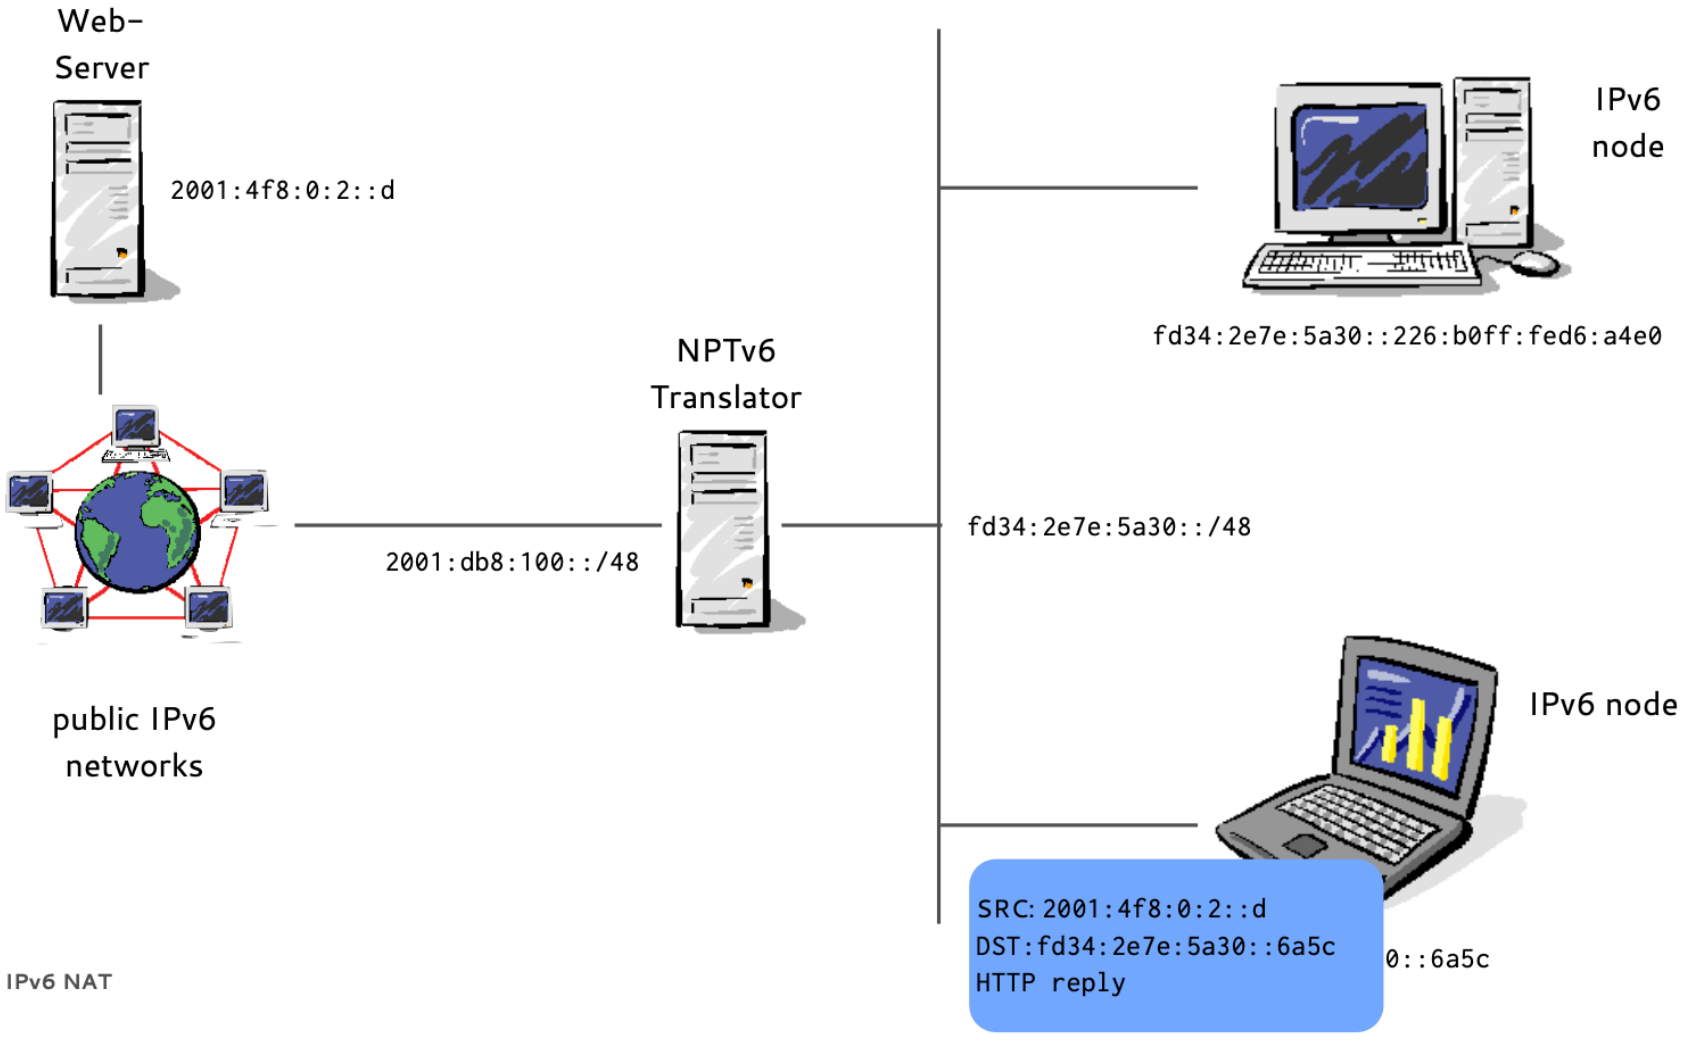
Web-
Server
2001:4f8:0:2::d
IPv6
node
fd34:2e7e:5a30::226:b0ff:fed6:a4e0
NPTv6
Translator
public IPv6
networks
2001:db8:100::/48
fd34:2e7e:5a30::/48
IPv6 node
SRC: 2001:4f8:0:2::d
DST:fd34:2e7e:5a30::6a5c
HTTP reply
0::6a5c

IPv6 NAT

IPv6 Network Prefix Translation

- NPTv6 is available for
  - Linux ip6tables / nftables
  - xBSD "pf" Firewall
  - Commercial Router OS (Juniper JunOS ...)

- Alternatives to NPTv6 are documented in ↪RFC 7157 "IPv6 Multihoming without Network Address Translation"

464XLAT

464XLAT - IPv4 over IPv6-only networks

- 464XLAT is defined in ↪RFC 6877 - 464XLAT: Combination of Stateful and Stateless Translation (Informational)
- 464XLAT "tunnels" IPv4 traffic over IPv6-only networks between *CLAT* and *PLAT* (more on these in a bit)
- 464XLAT has been originally invented for mobile carrier networks, but has seen popularity in other types on networks (enterprise networks) in the last years

464XLAT - CLAT and PLAT

- In 464XLAT, the *CLAT* encapsulates the IPv4 address into an IPv6 address, the IPv6 prefix is routed towards the *PLAT*
  - *CLAT*: client/customer side translator. Can be a separate box, but is usually build into an operating systems network stack (Android, iOS, macOS, Linux)
  - The *CLAT* is stateless

464XLAT - CLAT and PLAT

- The encapsulated traffic is send over IPv6 towards the *PLAT*
  - *PLAT*: provider translator
  - The *PLAT* has IPv4 connectivity and creates a new IPv4 packet from the encapsulated IPv4 address
  - The *PLAT* needs to keep *state* (like a traditional NAT44 device)
  - The *PLAT* is located at the edge of the IPv6-only network (despite the name, it does not need to be operated by a "provider")

464XLAT

- Like NAT44, 464XLAT allows applications to reach IPv4 destinations, but there is no end-to-end connectivity from the Internet back towards the client
- IPv6-only hosts find the PLAT/NAT64 prefix by sending an `AAAA` (IPv6 address) DNS query for the name `ipv4only.arpa` (↪RFC 7050 and ↪RFC 8880)

464XLAT operation

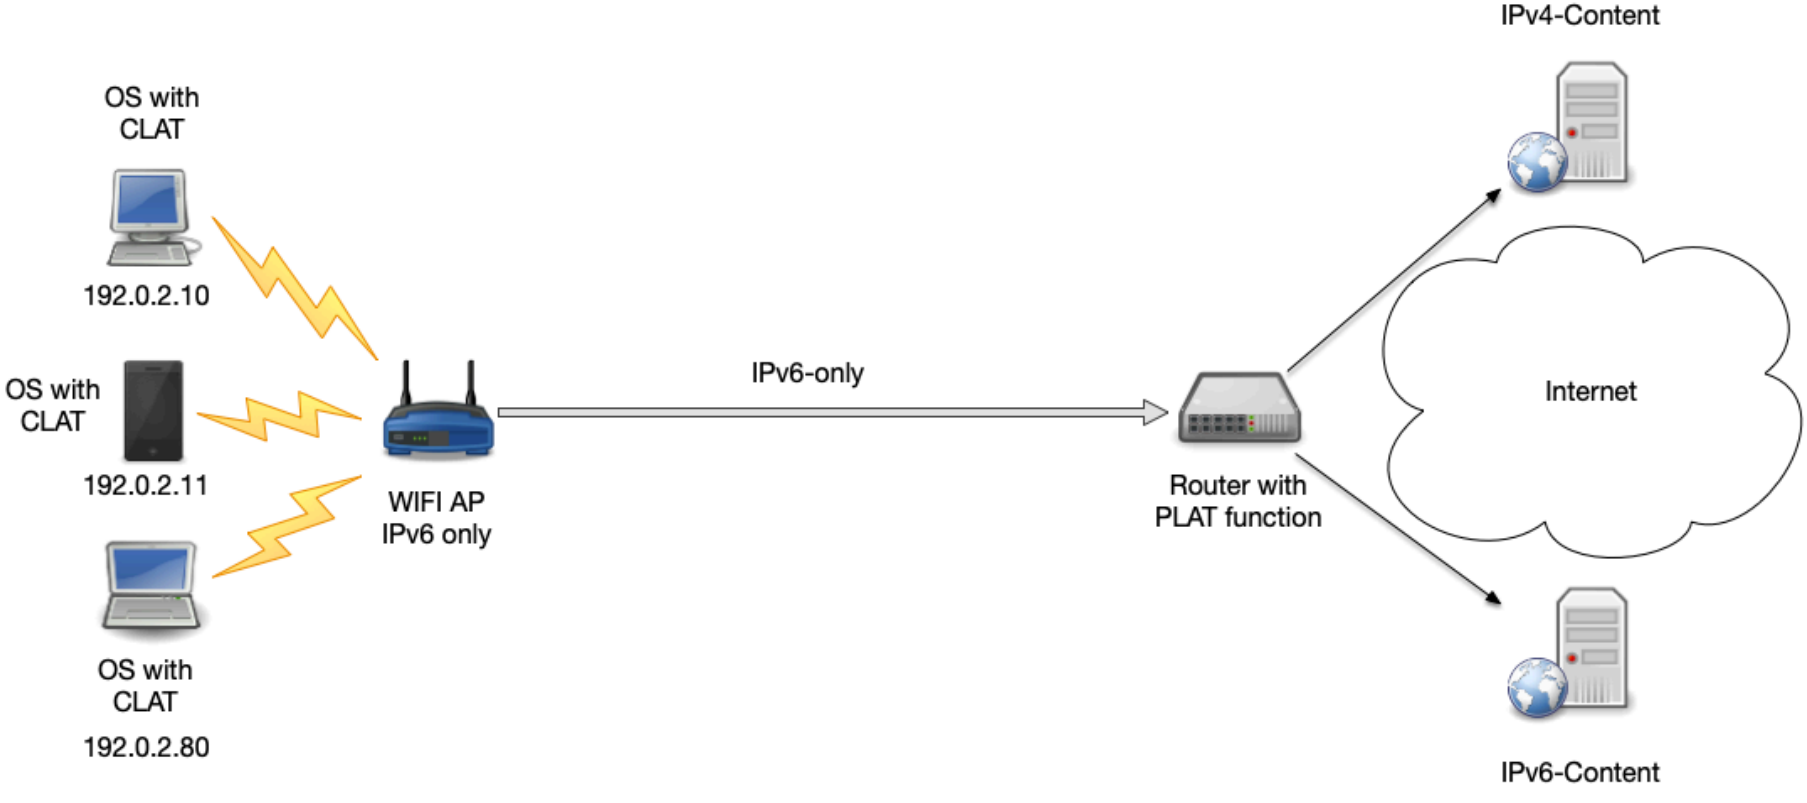

NAT64/DNS64

NAT64/DNS64

- NAT64/DNS64 builds on the functions for IPv4/IPv6 translation as seen before
- It uses a special DNS function (DNS64) to translate between IPv4 and IPv6 addresses

DNS64 / NAT64

- DNS64 and NAT64 are used to connect IPv6 only hosts to IPv4 only systems on the Internet
  - DNS64 translates DNS request
  - NAT64 does Network Address Translation

DNS64 / NAT64

- If a host has no IPv6 AAAA-Record in DNS
  - Instead of returning "NODATA", a DNS64 DNS Server translates IPv4 A-Records to IPv6 AAAA-Records
  - A NAT64 Gateway translates between the 'synthesized' IPv6 Addresses and the real IPv4 host between IPv6 and IPv4

DNS64/NAT64 example (1/19)

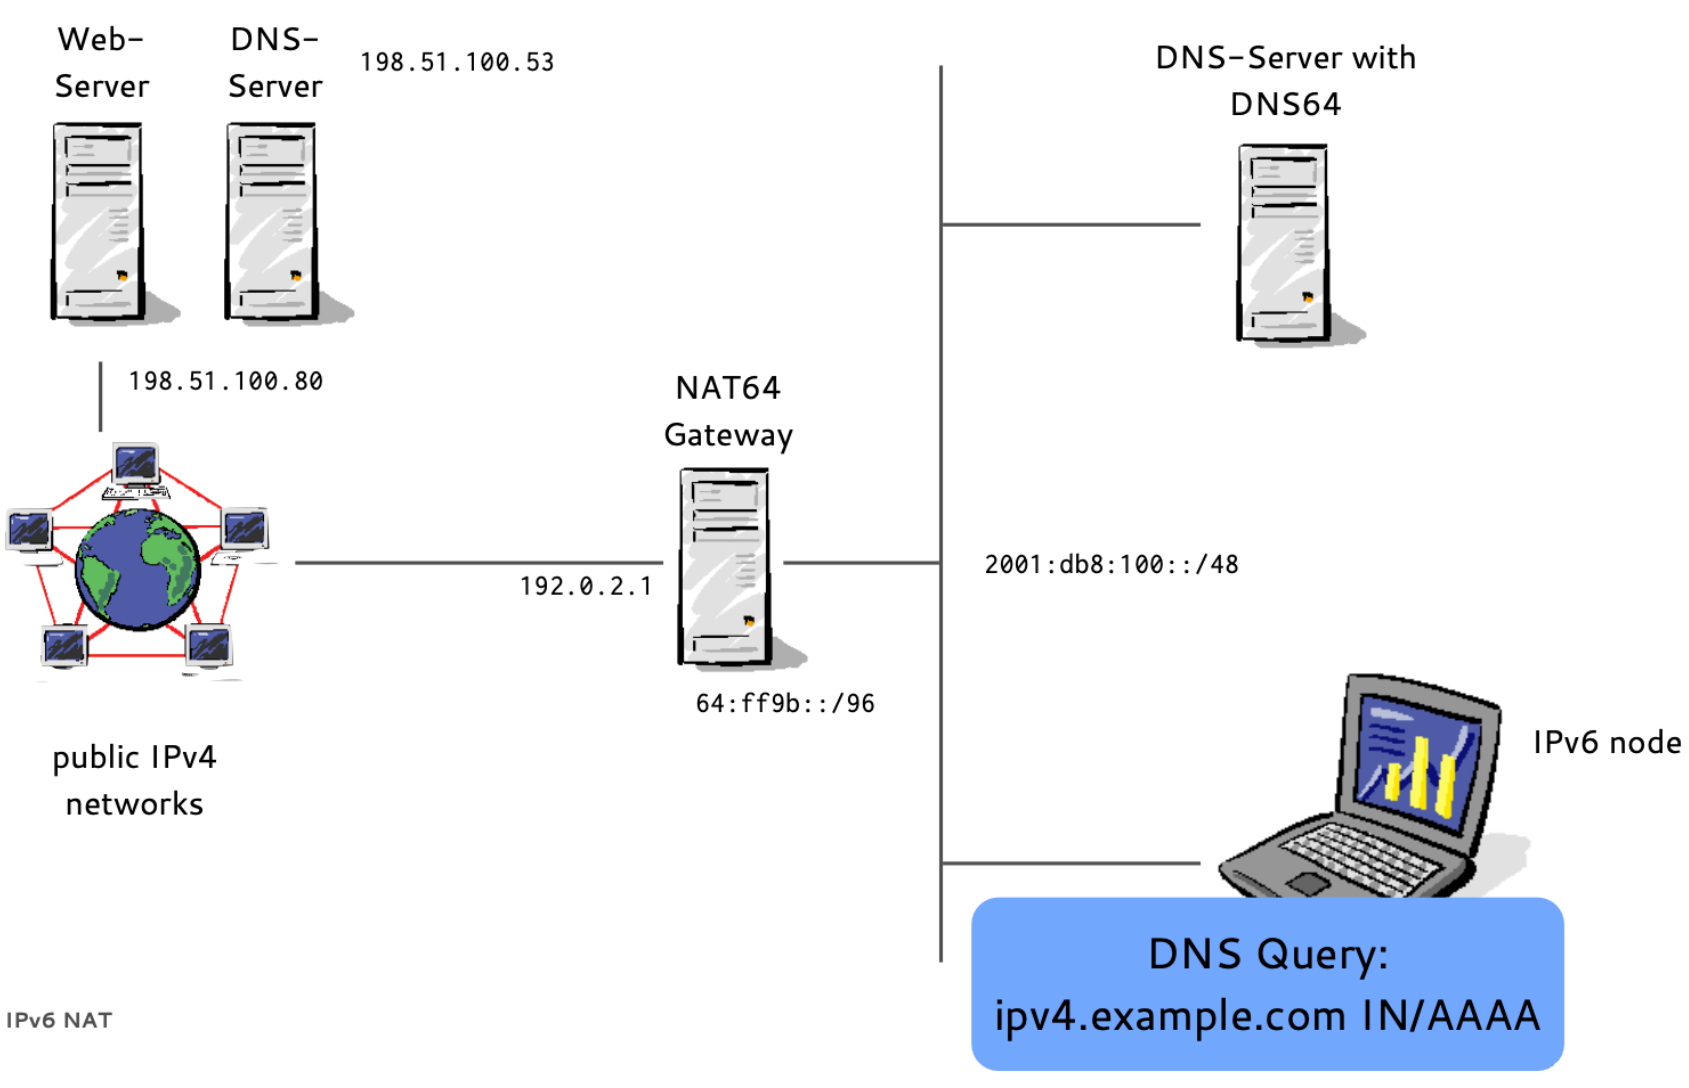
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
DNS Query:
ipv4.example.com IN/AAAA

IPv6 NAT

DNS64/NAT64 example (2/19)

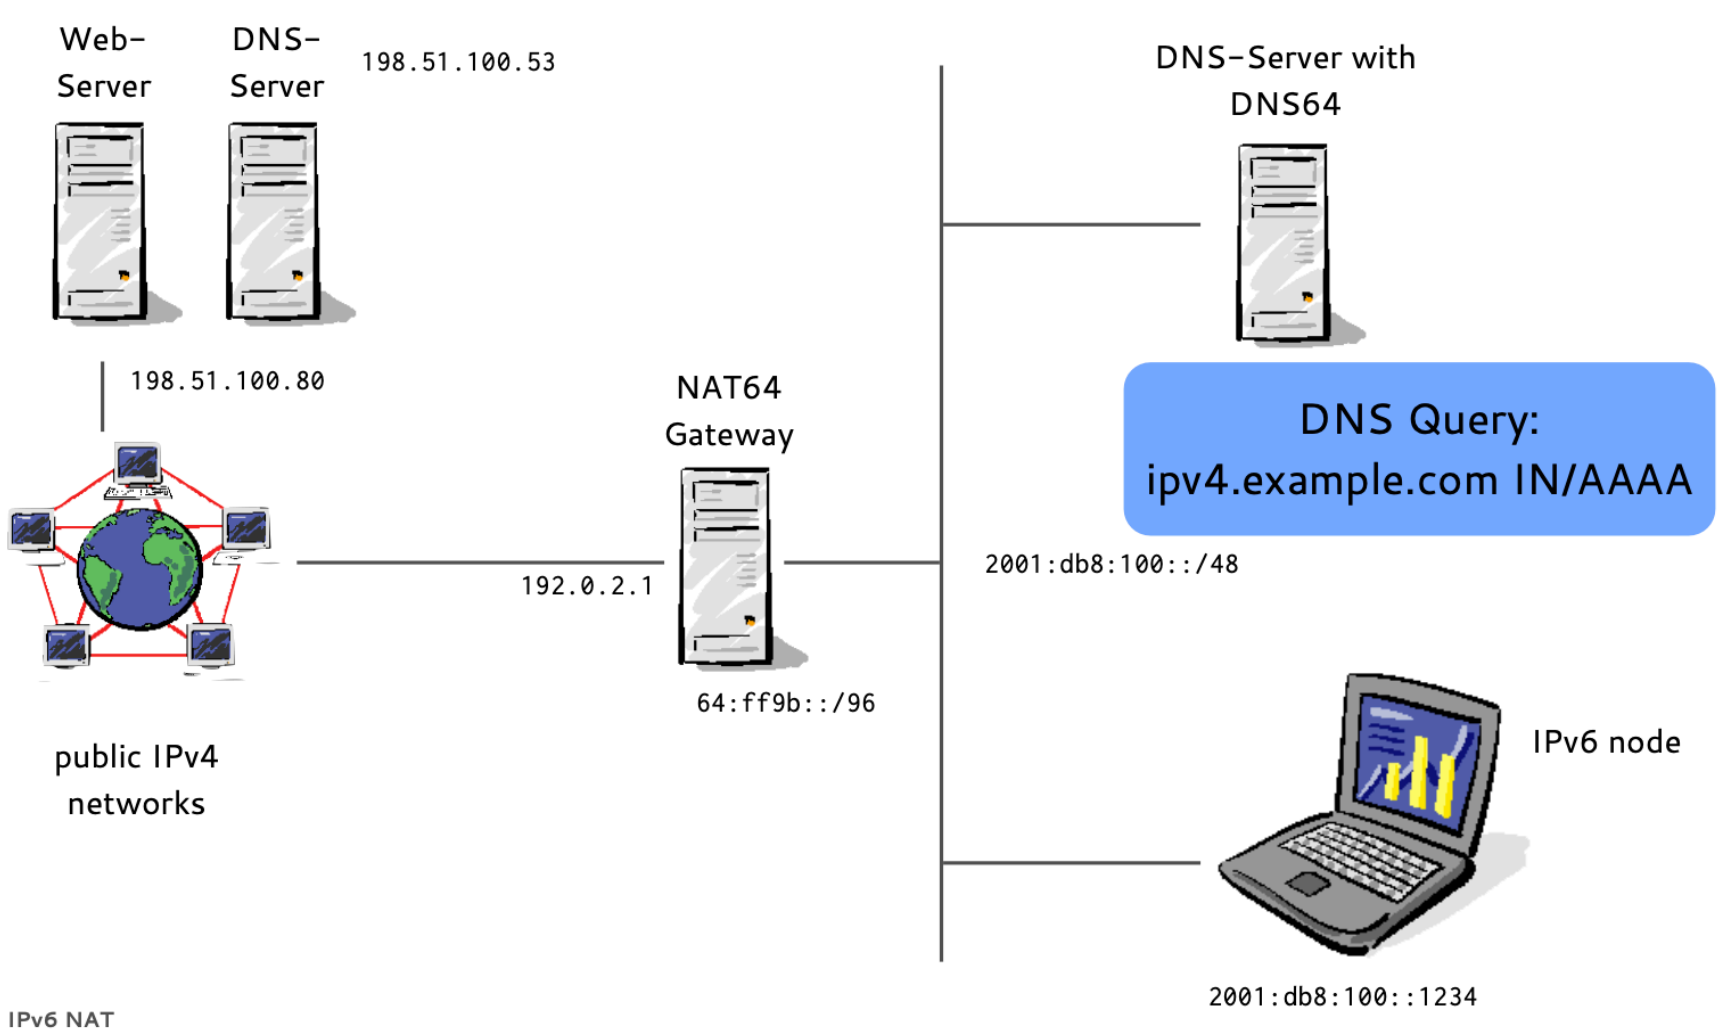
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
DNS Query:
ipv4.example.com IN/AAAA
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (3/19)

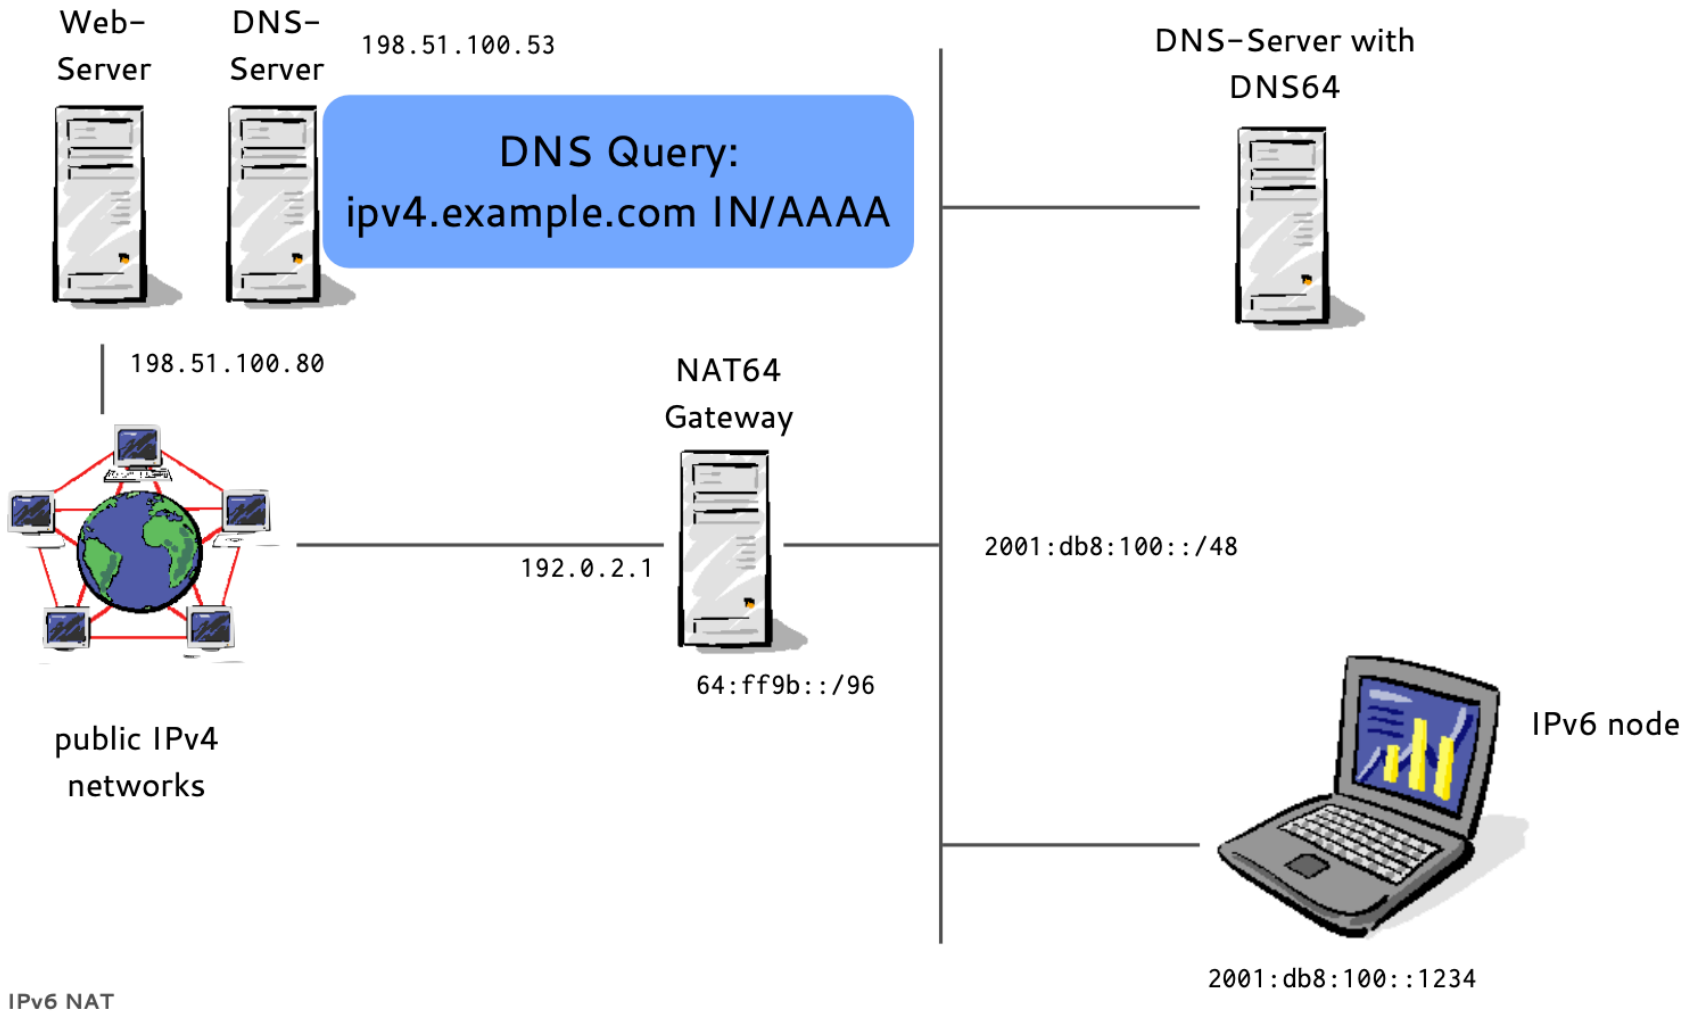
Web-
Server
DNS-
Server
198.51.100.53
DNS Query:
ipv4.example.com IN/AAAA
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (4/19)

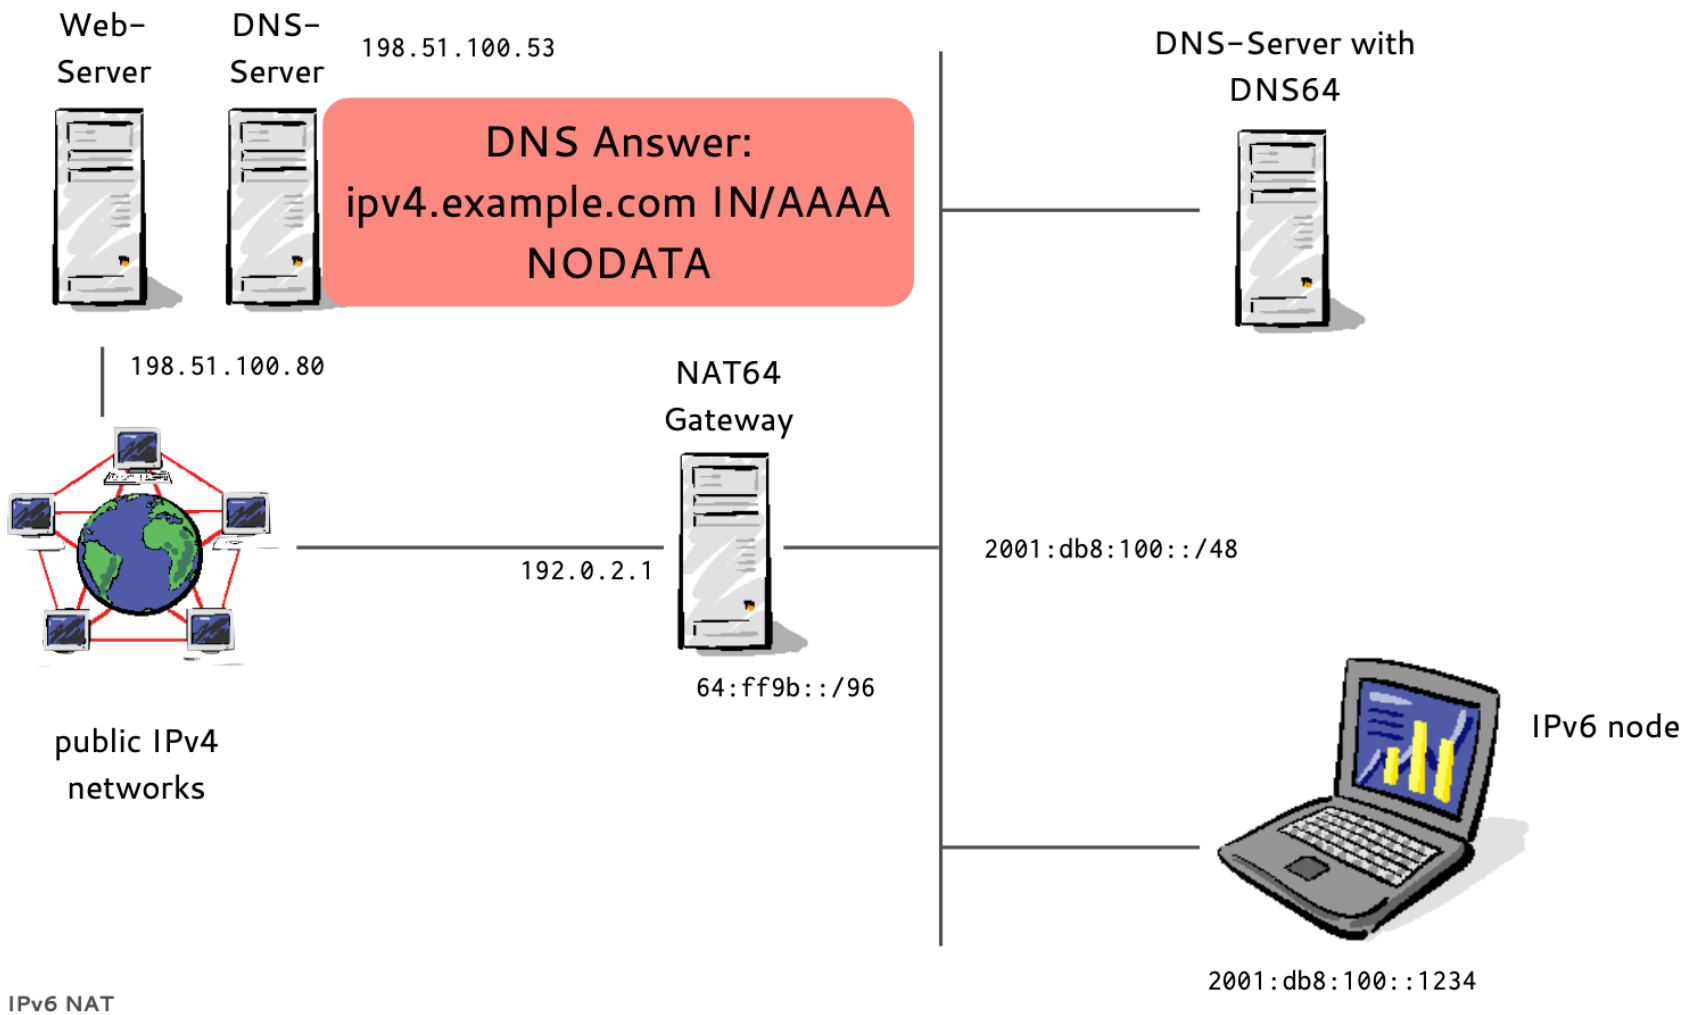
Web-
Server
DNS-
Server
198.51.100.53
DNS Answer:
ipv4.example.com IN/AAAA
NODATA
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (5/19)

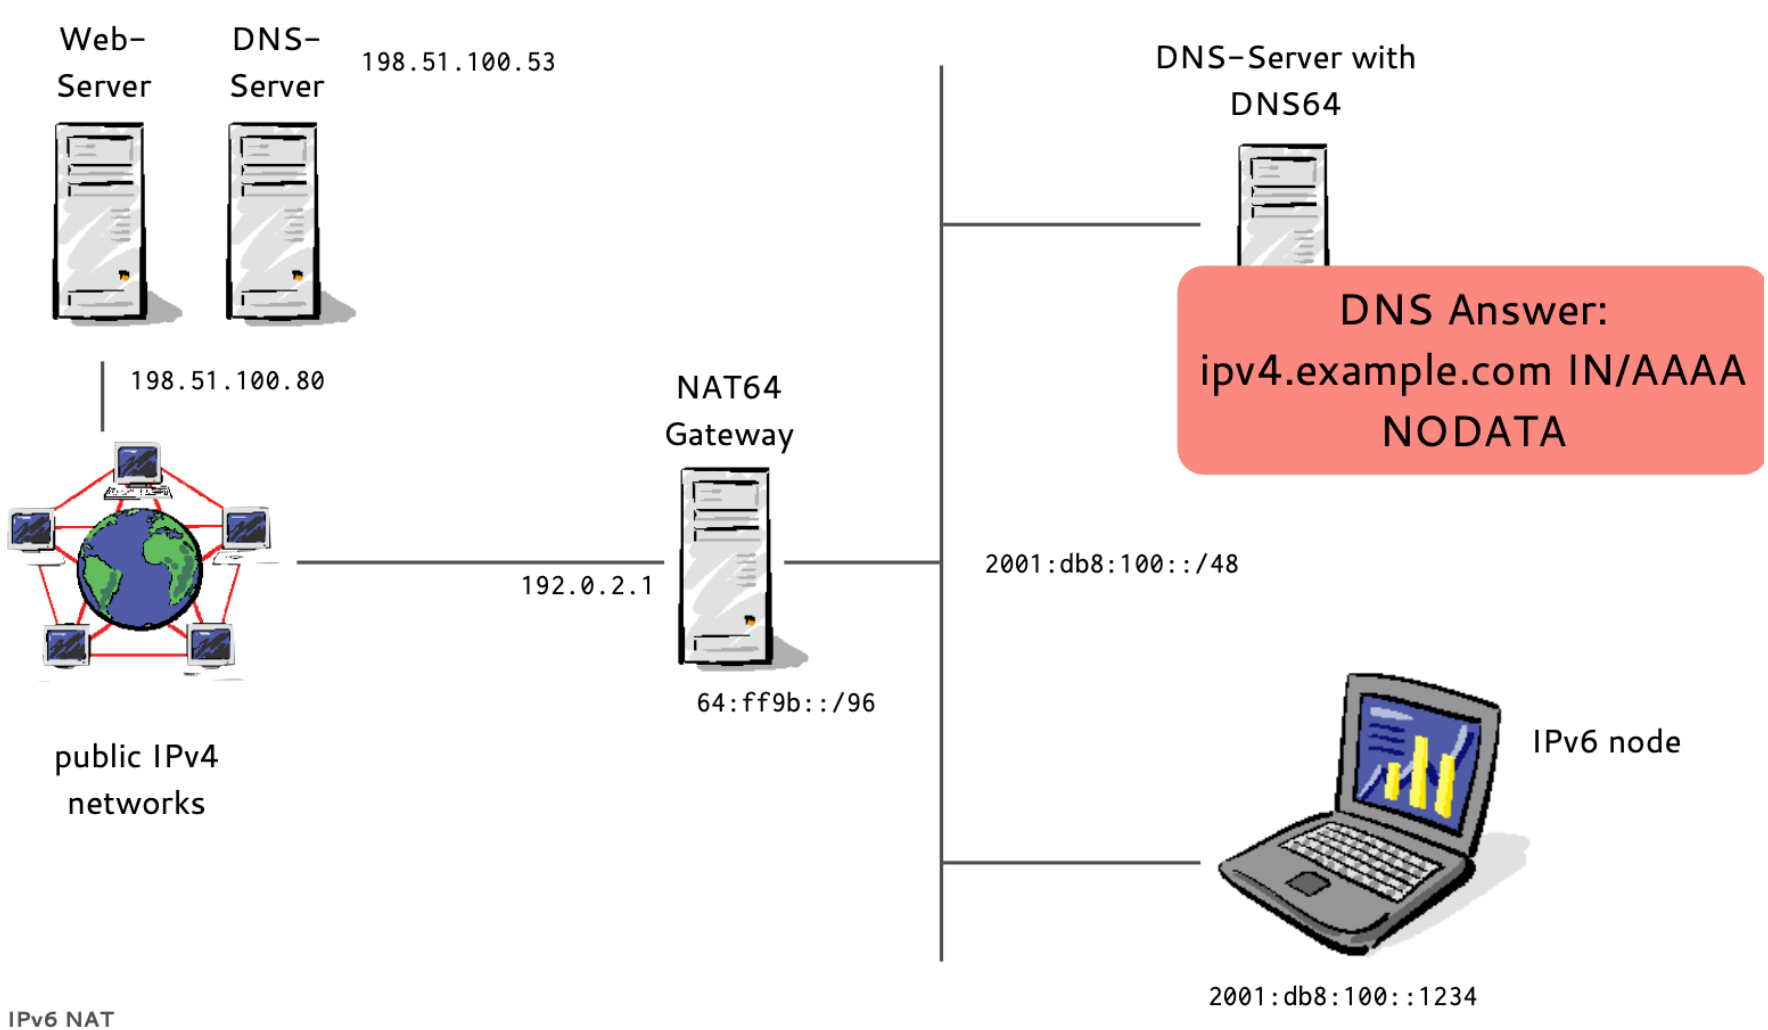
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
DNS Answer:
ipv4.example.com IN/AAAA
NODATA
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (6/19)

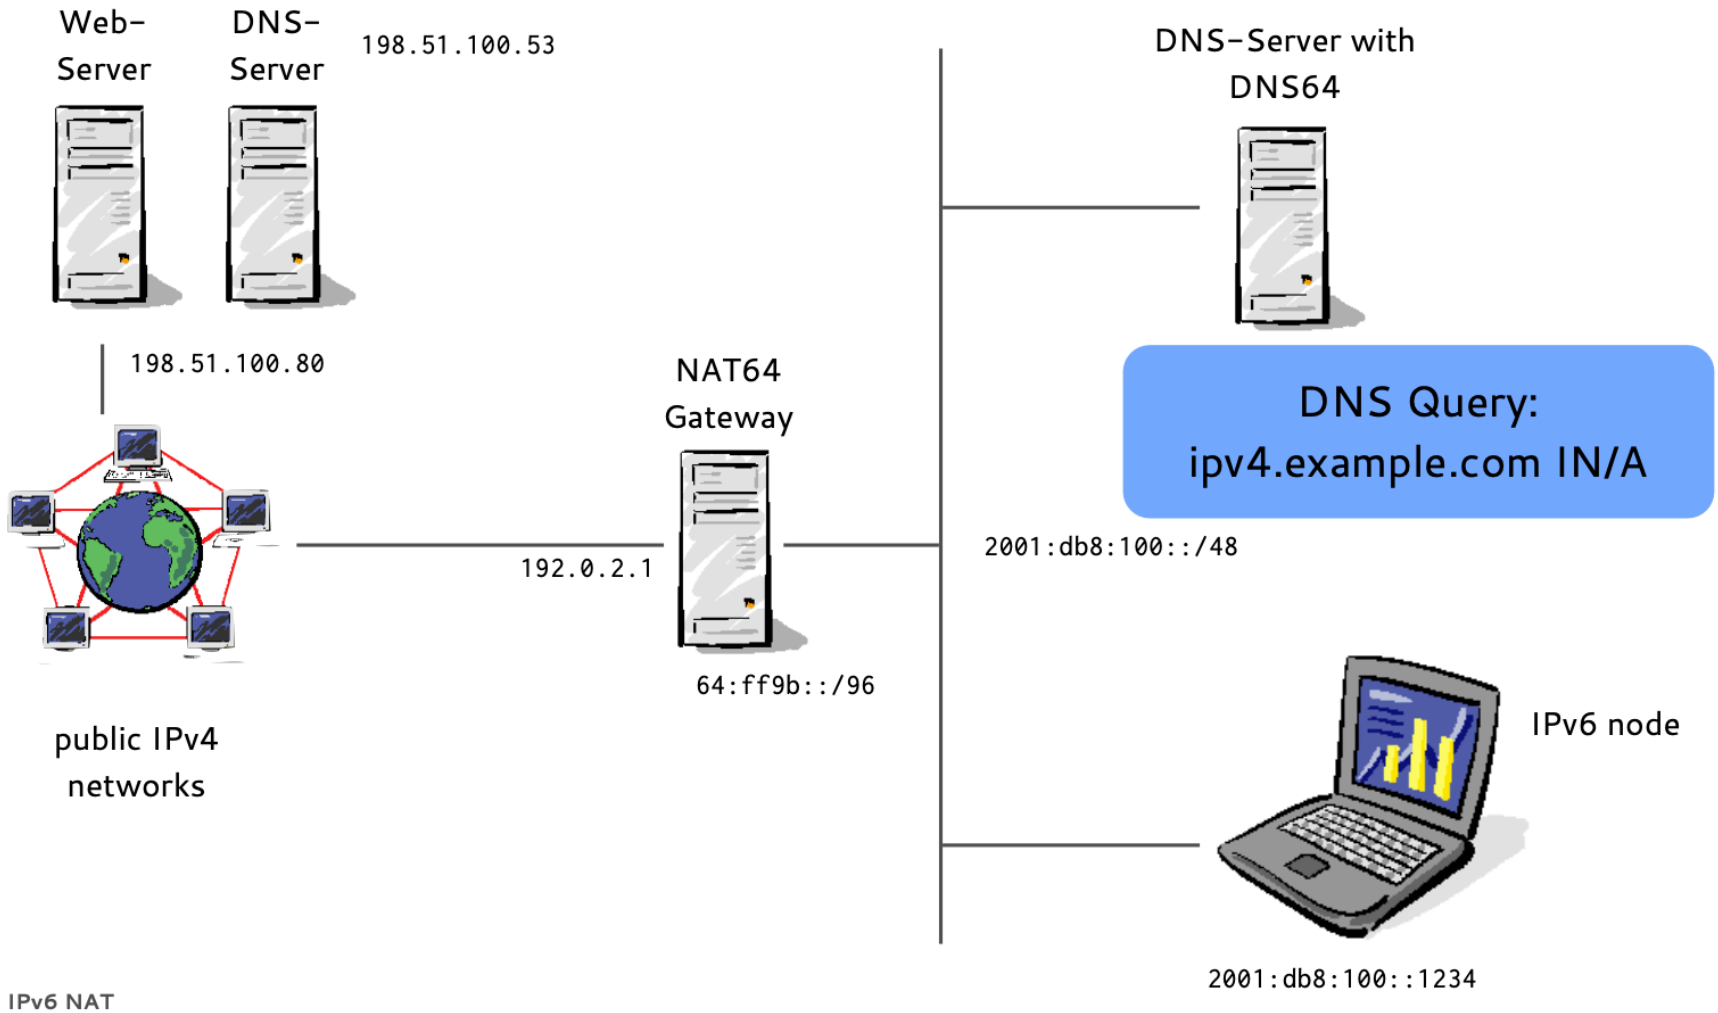
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
DNS Query:
ipv4.example.com IN/A
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

DNS64/NAT64 example (7/19)

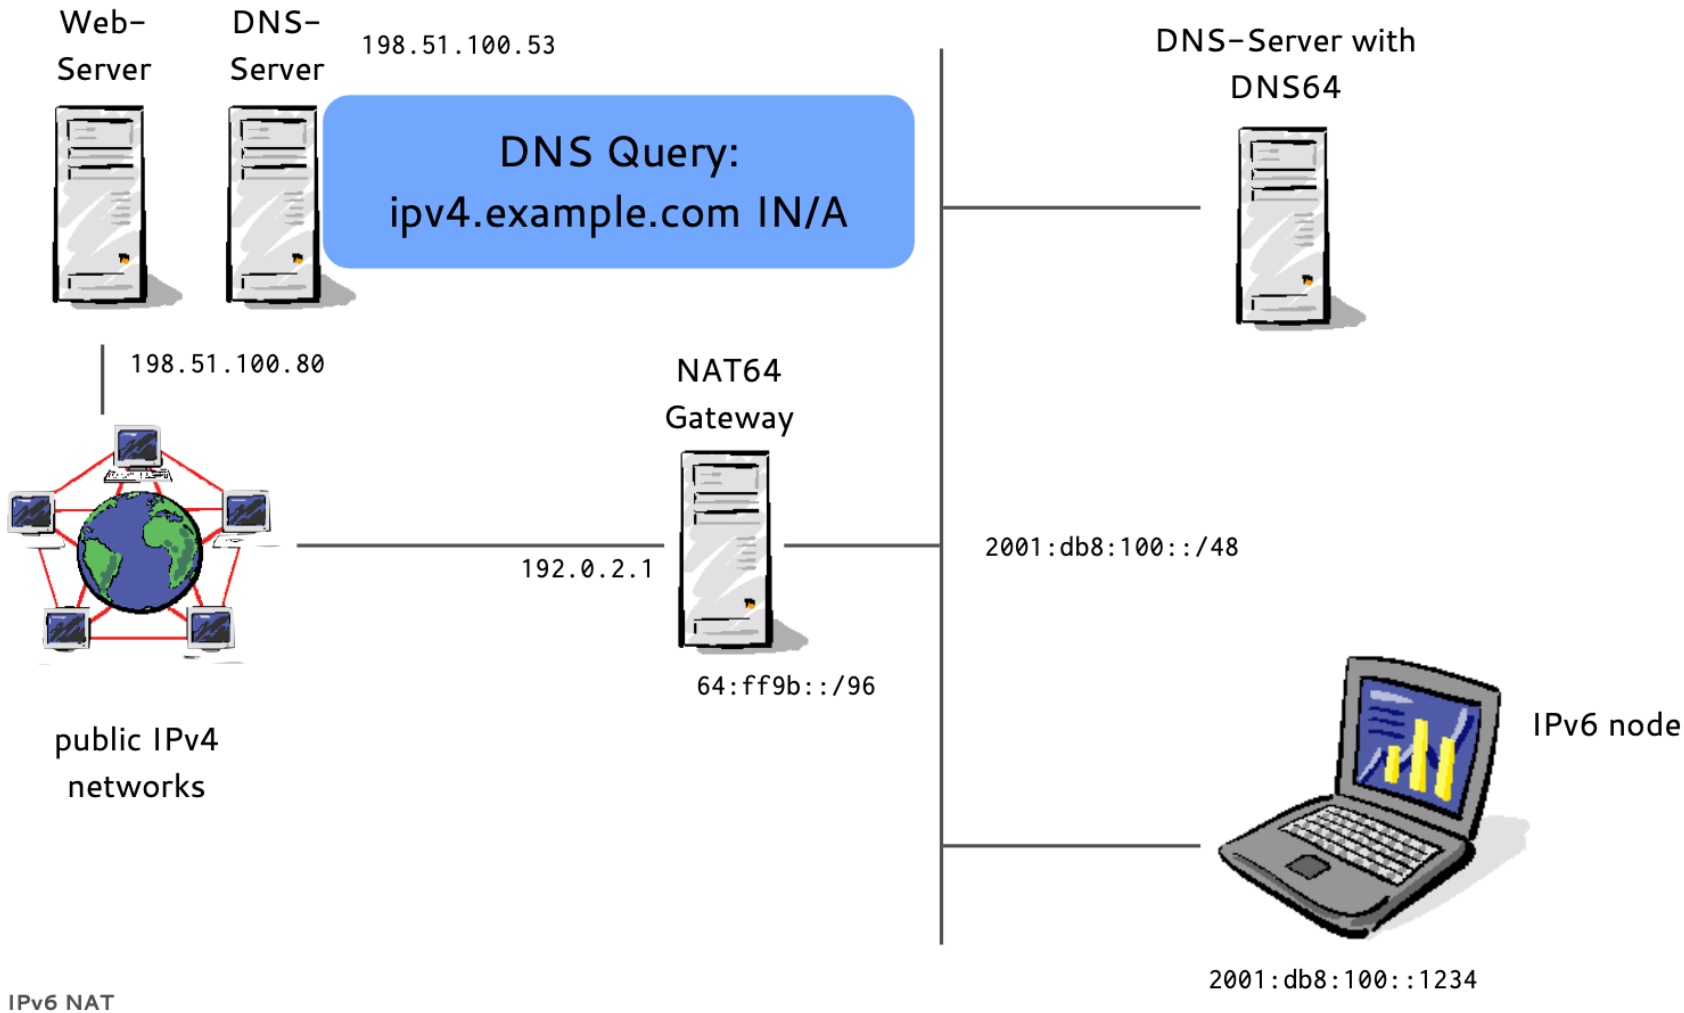
Web-
Server
DNS-
Server
198.51.100.53
DNS Query:
ipv4.example.com IN/A
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (8/19)

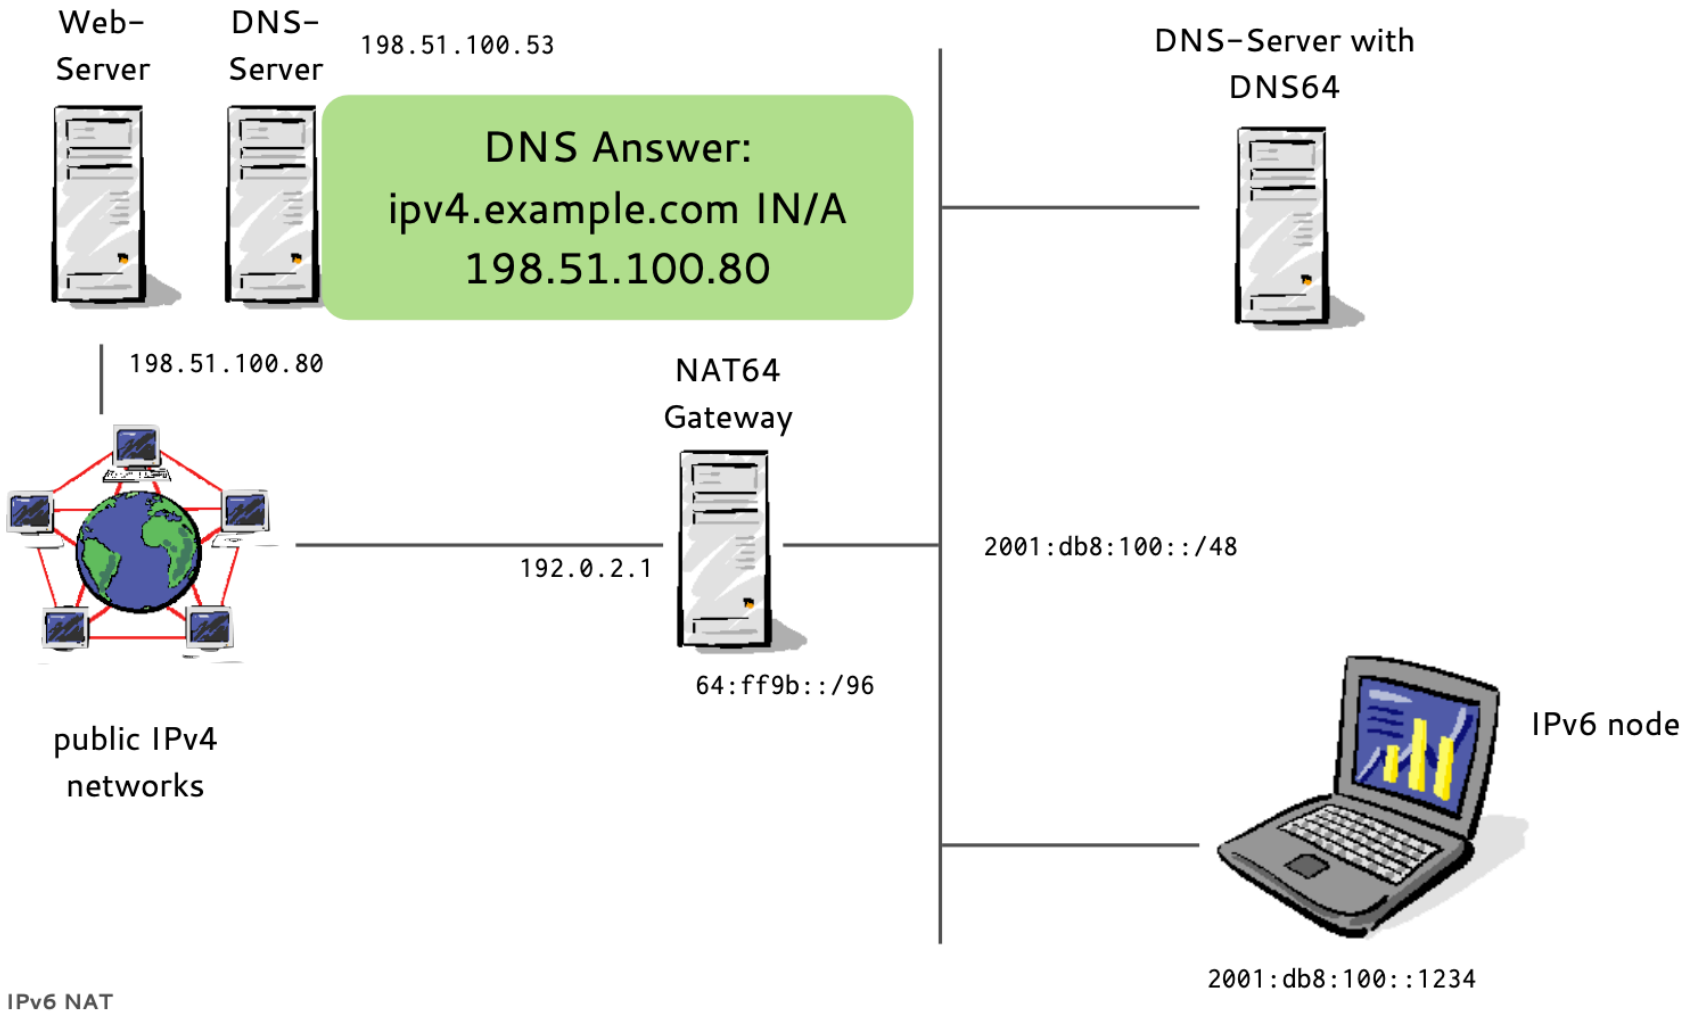
Web-
Server
DNS-
Server
198.51.100.53
DNS Answer:
ipv4.example.com IN/A
198.51.100.80
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

DNS64/NAT64 example (9/19)

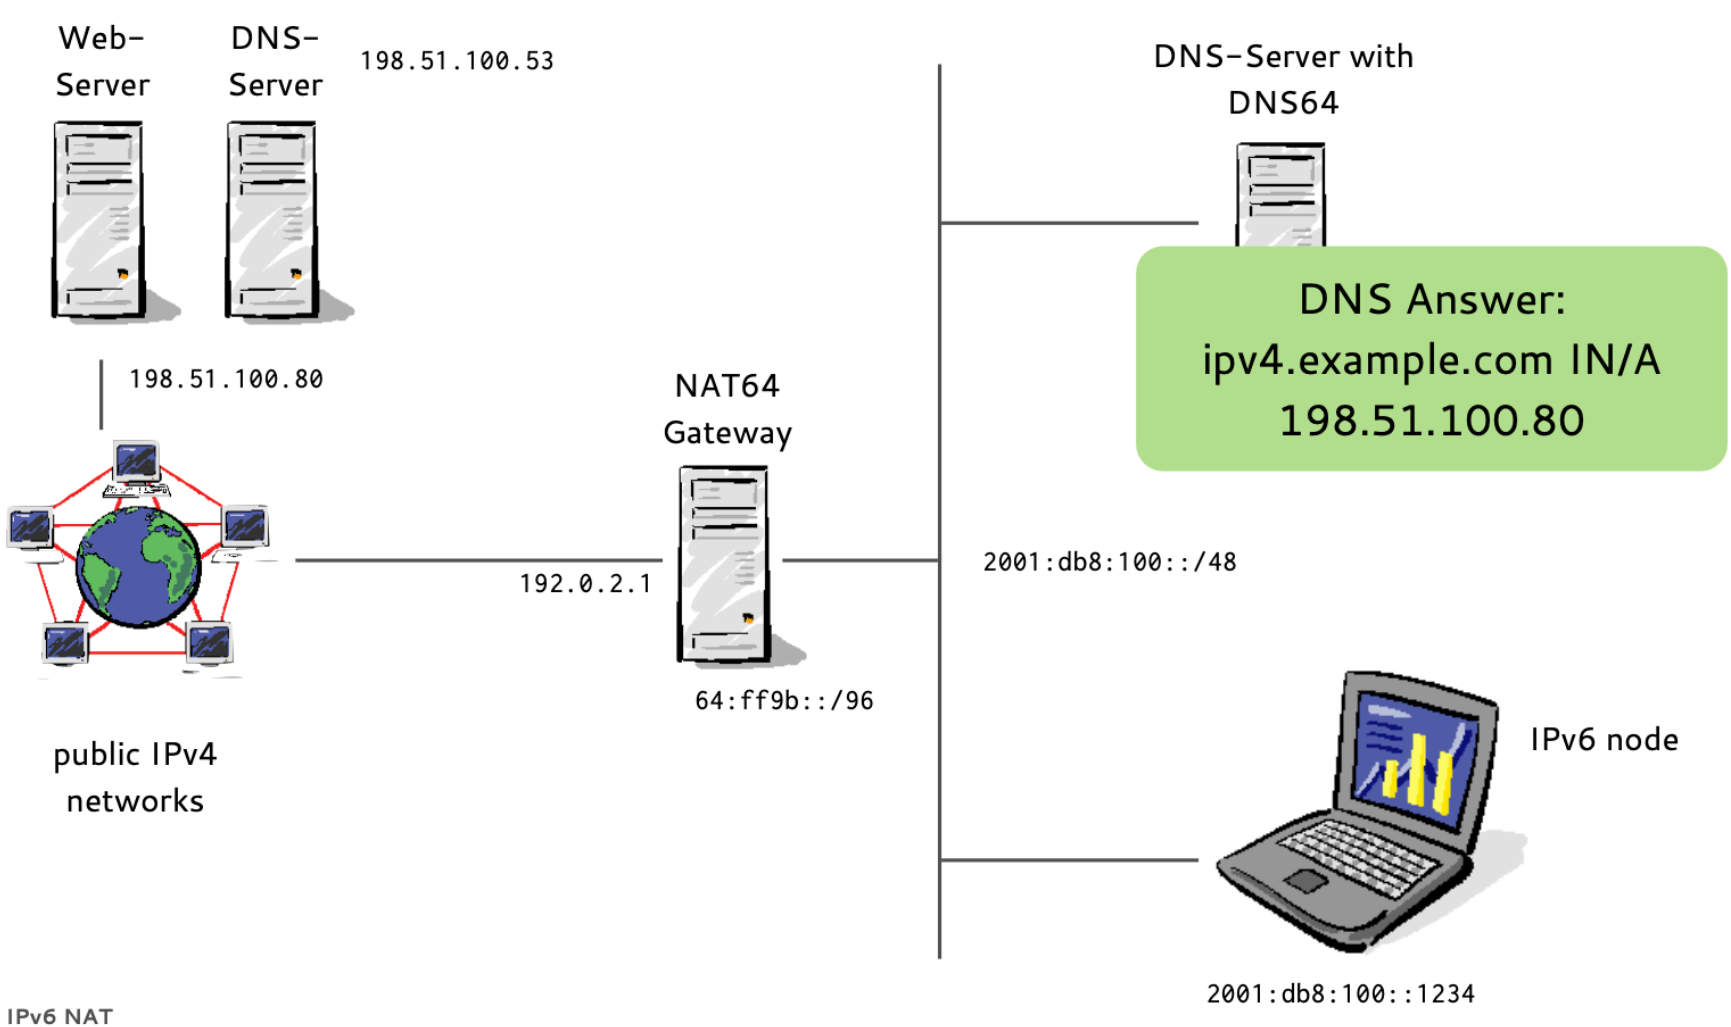
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
DNS Answer:
ipv4.example.com IN/A
198.51.100.80
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (10/19)

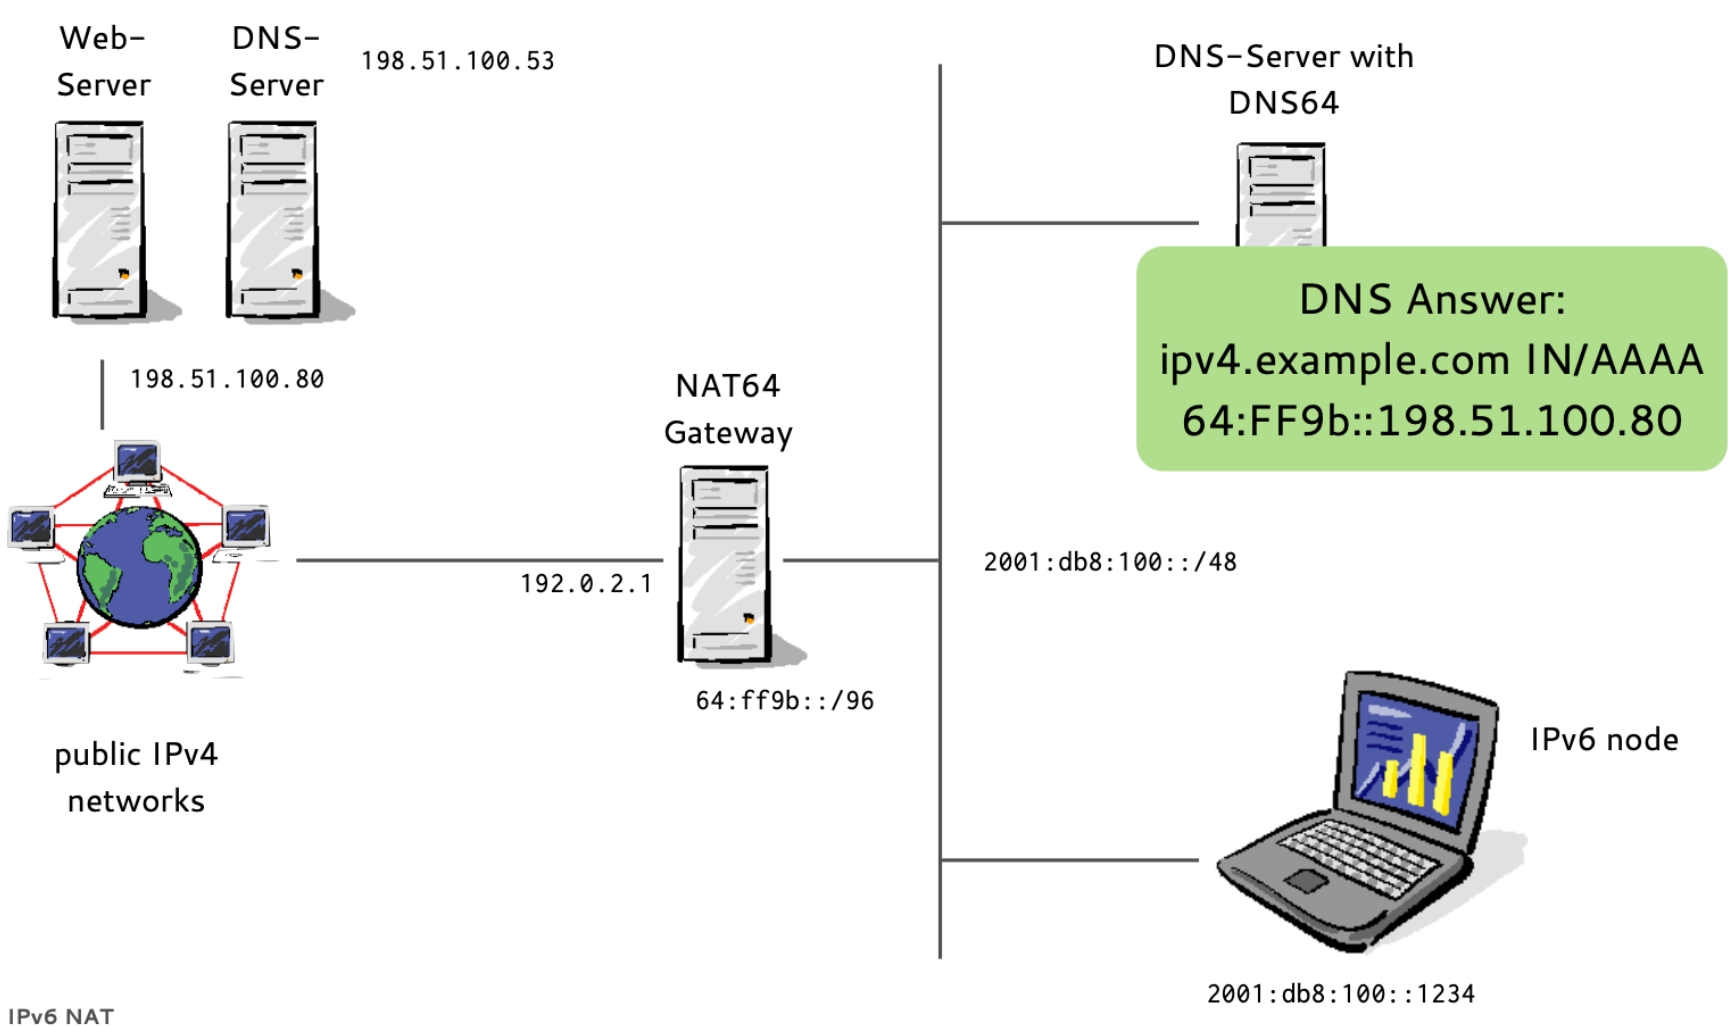
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
DNS Answer:
ipv4.example.com IN/AAAA
64:FF9b::198.51.100.80
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (11/19)

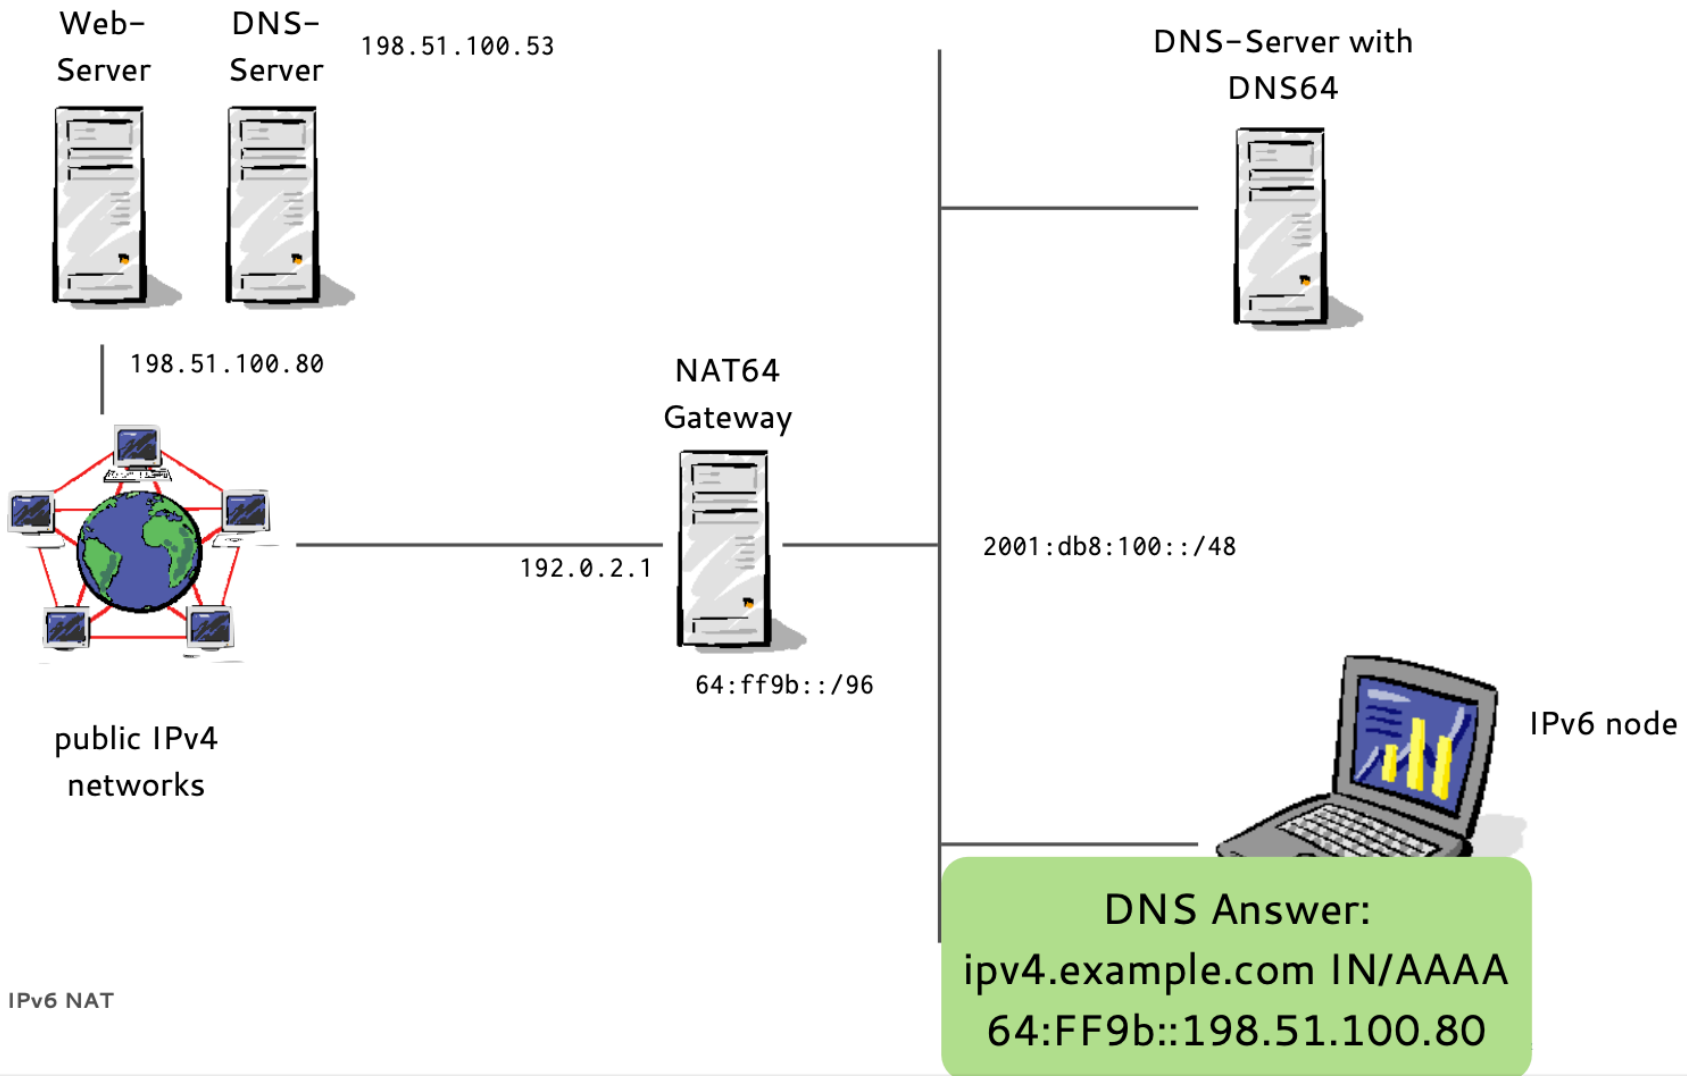
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
DNS Answer:
ipv4.example.com IN/AAAA
64:FF9b::198.51.100.80

IPv6 NAT

DNS64/NAT64 example (12/19)

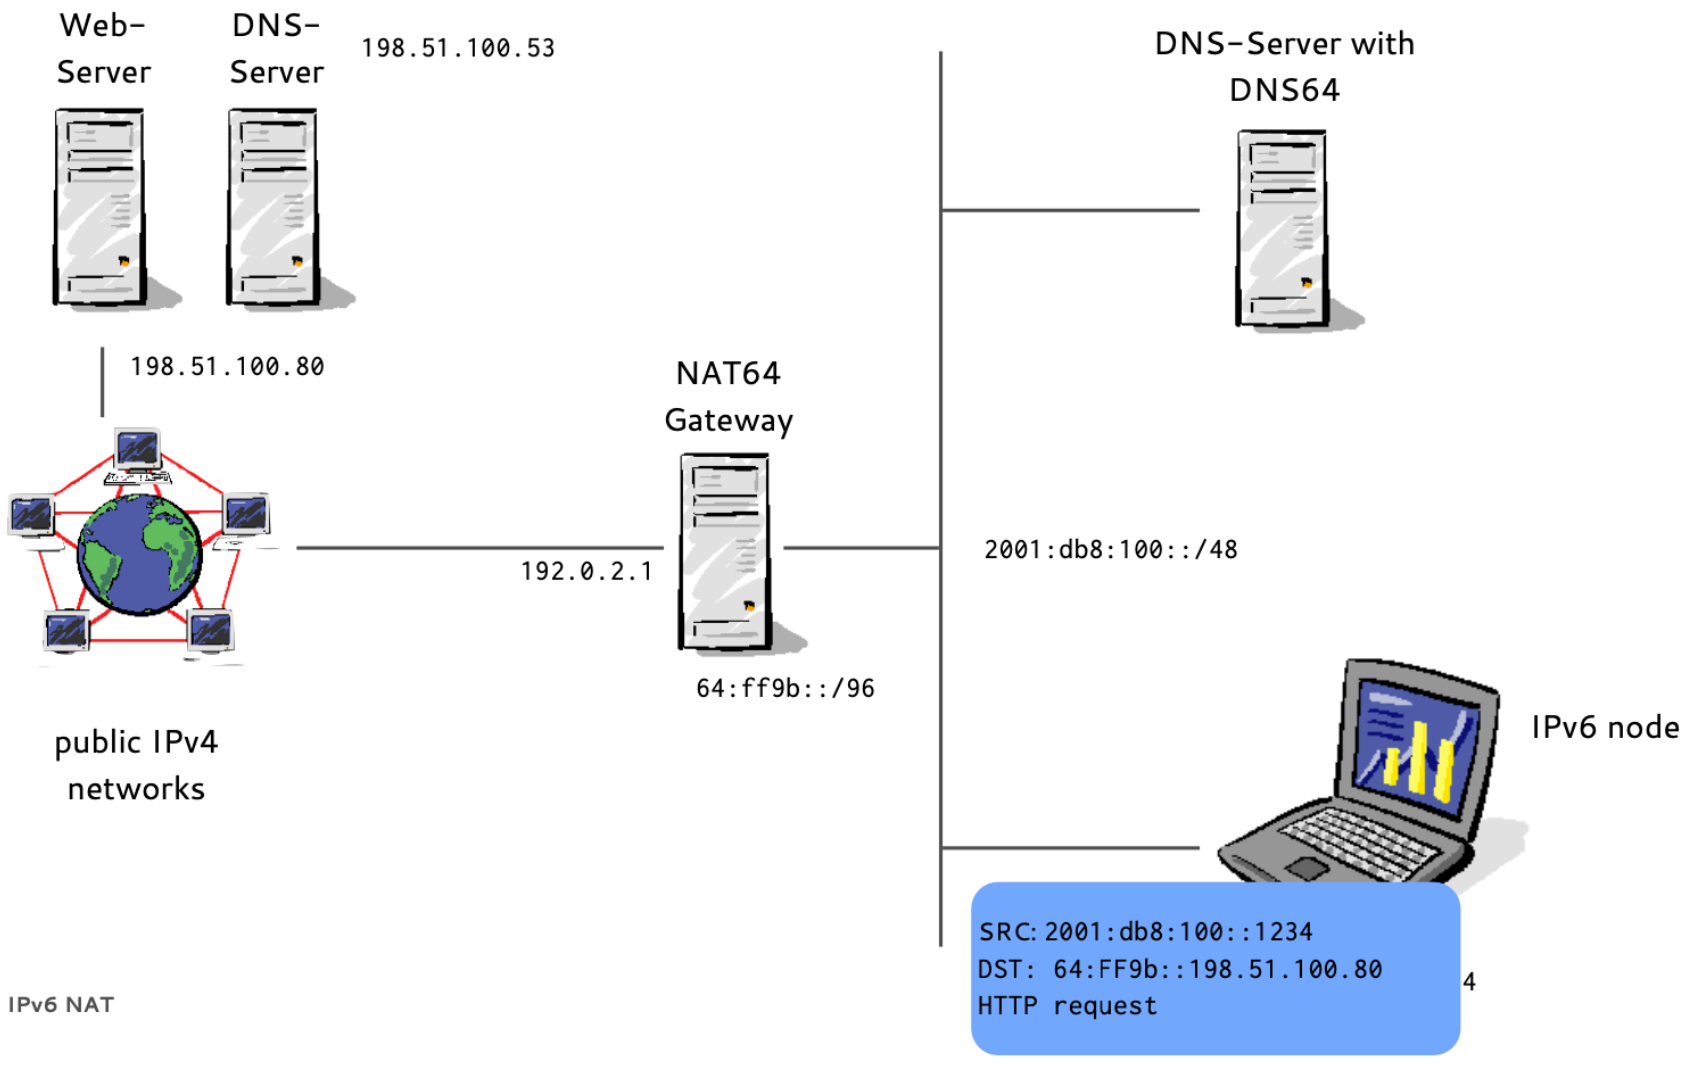
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
SRC: 2001:db8:100::1234
DST: 64:FF9b::198.51.100.80
HTTP request
4

IPv6 NAT

DNS64/NAT64 example (13/19)

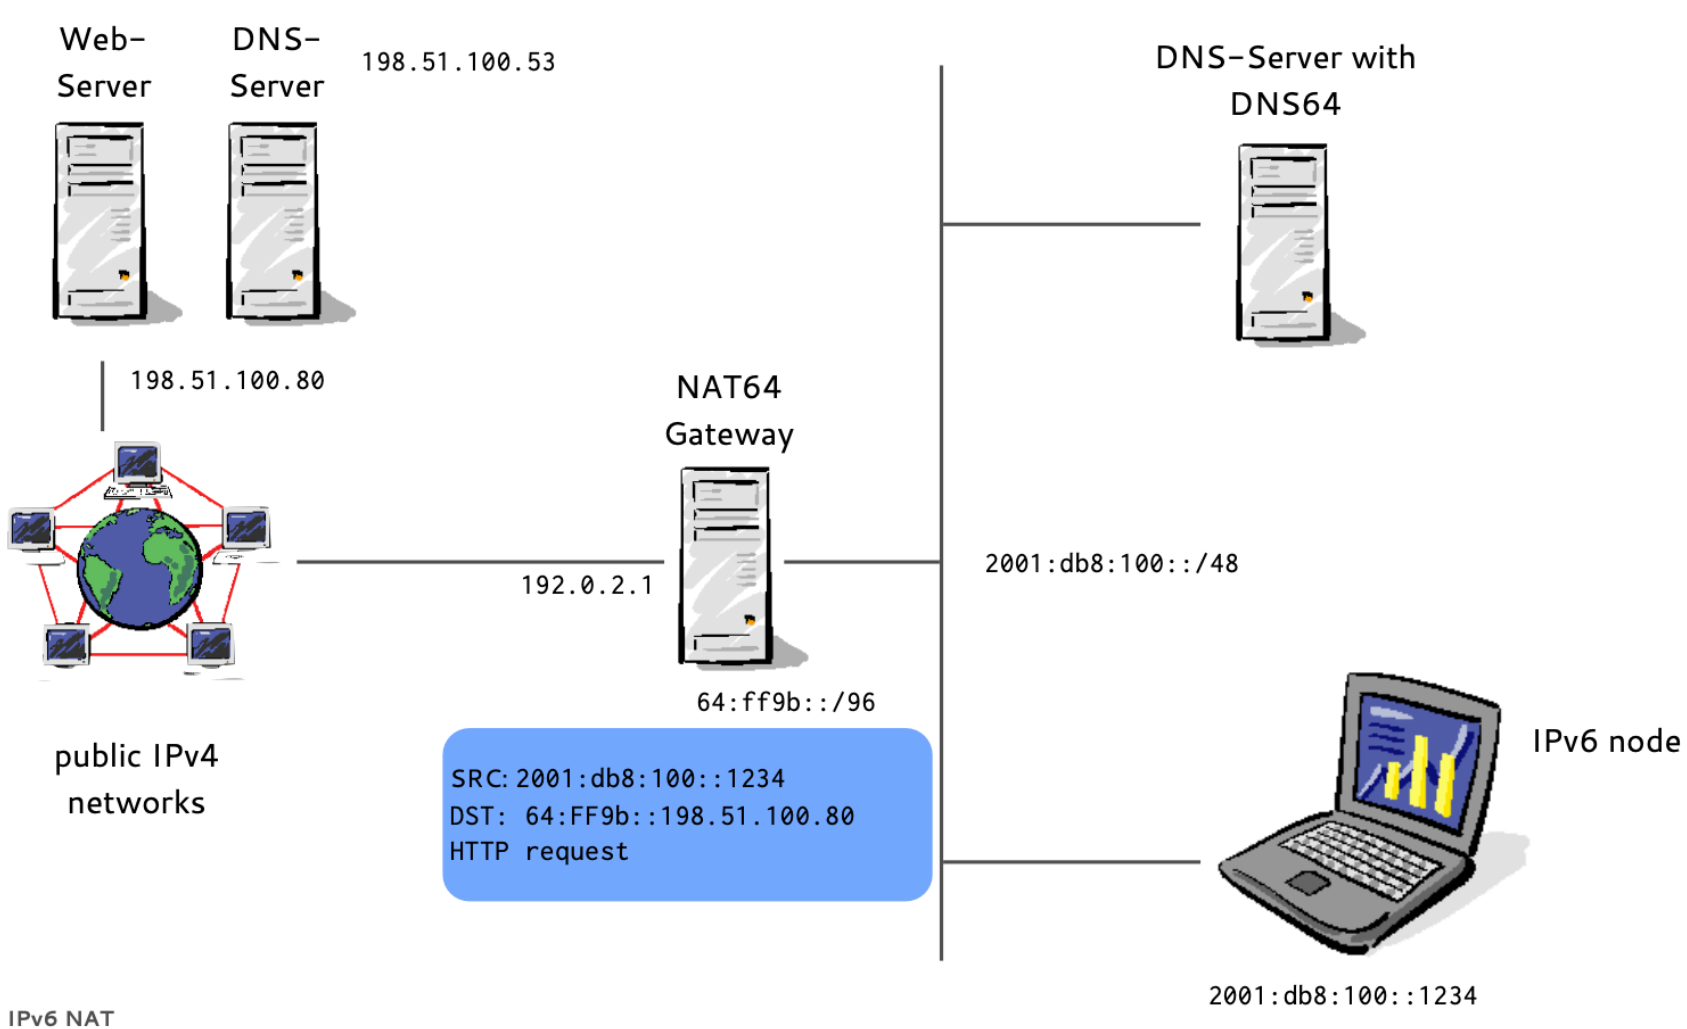
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
SRC: 2001:db8:100::1234
DST: 64:FF9b::198.51.100.80
HTTP request
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (14/19)

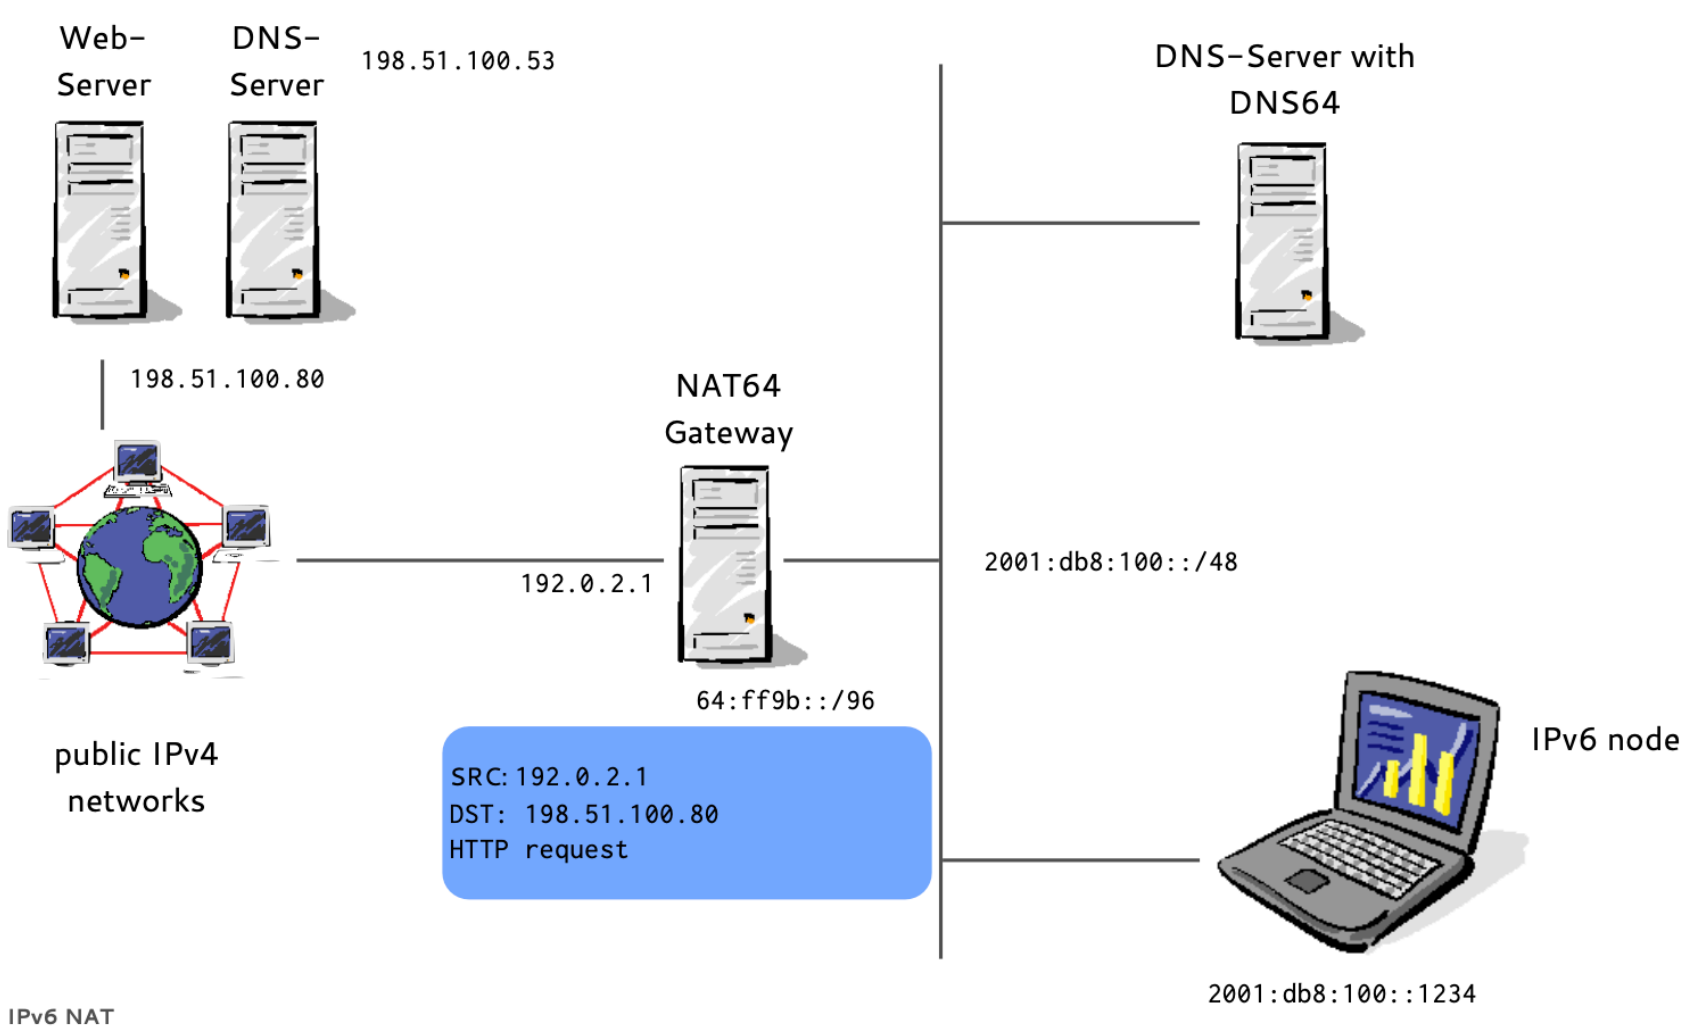
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
SRC: 192.0.2.1
DST: 198.51.100.80
HTTP request
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (15/19)

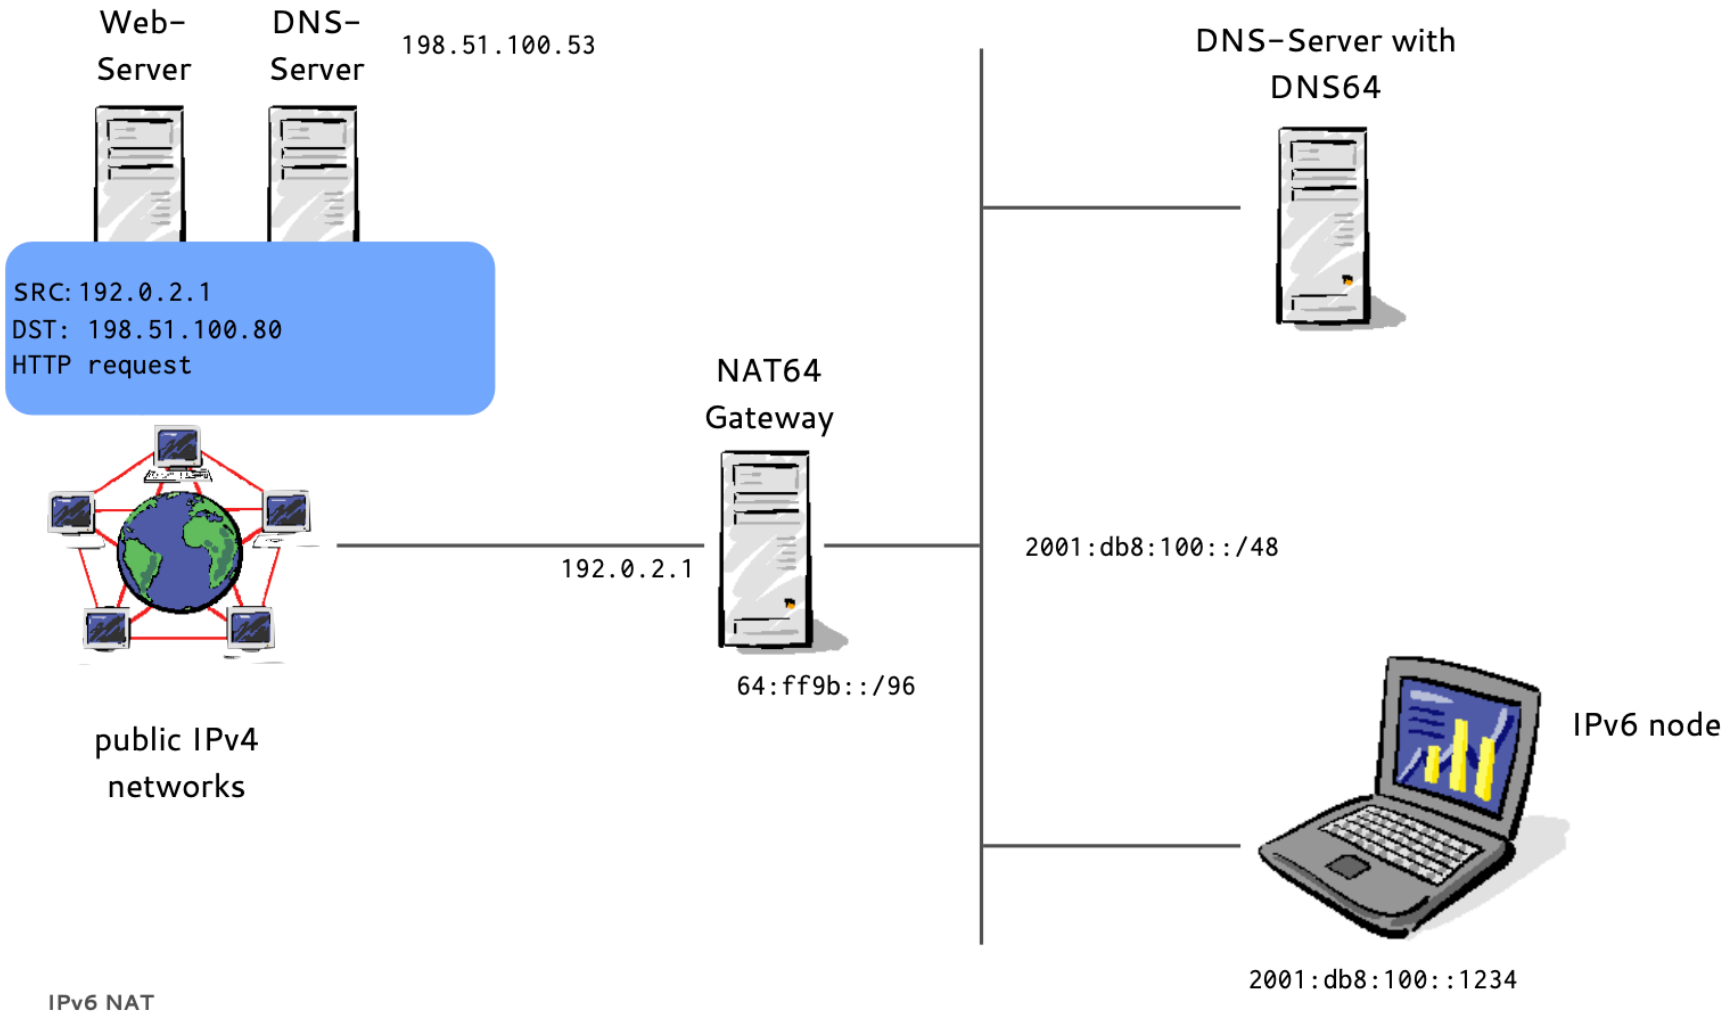
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
SRC: 192.0.2.1
DST: 198.51.100.80
HTTP request
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (16/19)

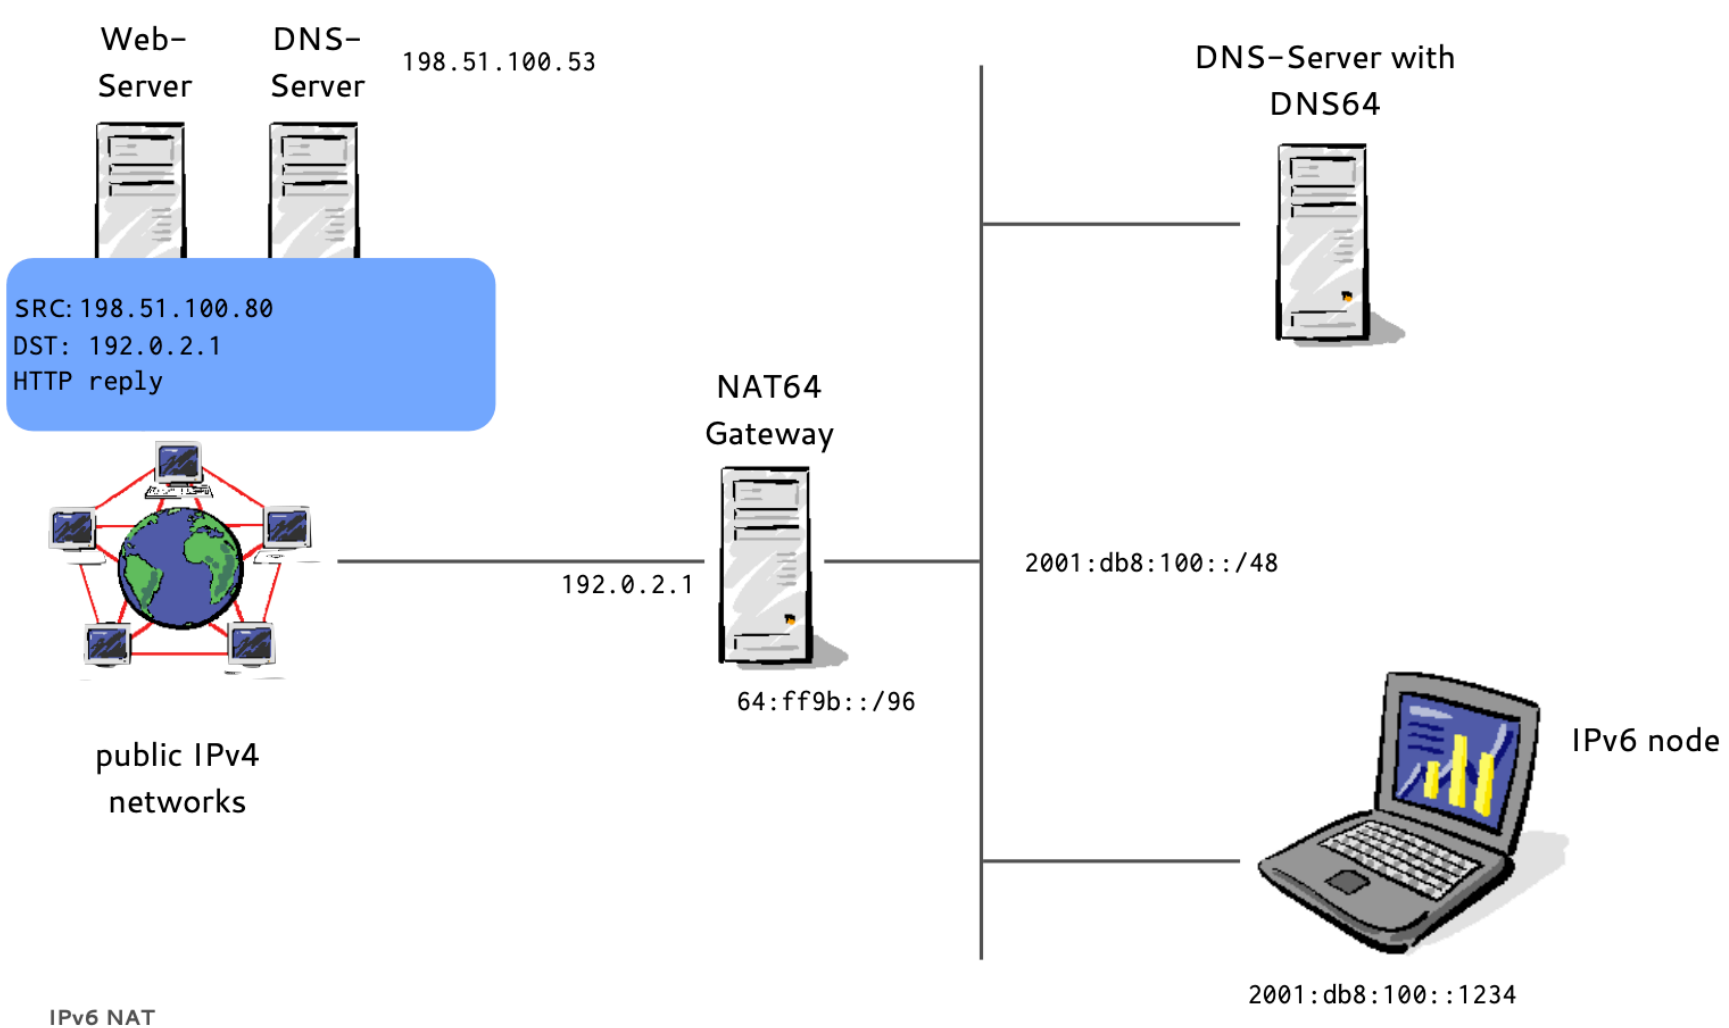
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
SRC: 198.51.100.80
DST: 192.0.2.1
HTTP reply
NAT64
Gateway
192.0.2.1
2001:db8:100::/48
64:ff9b::/96
public IPv4
networks
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (17/19)

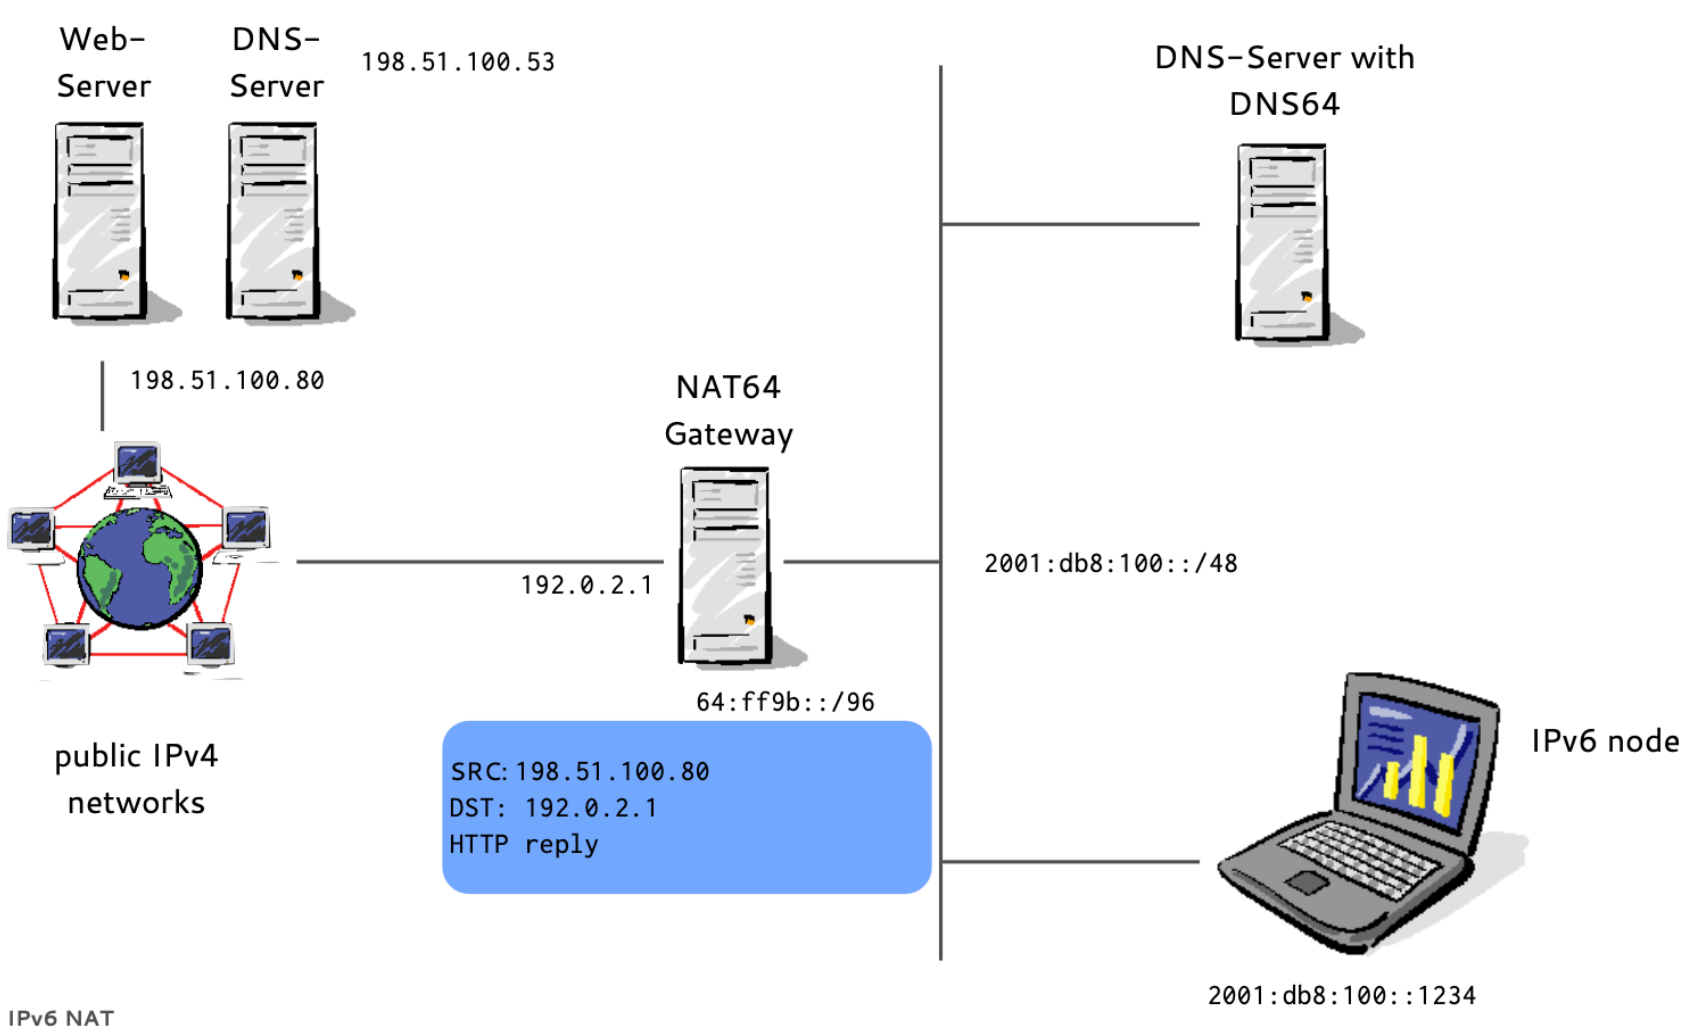
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
SRC: 198.51.100.80
DST: 192.0.2.1
HTTP reply
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (18/19)

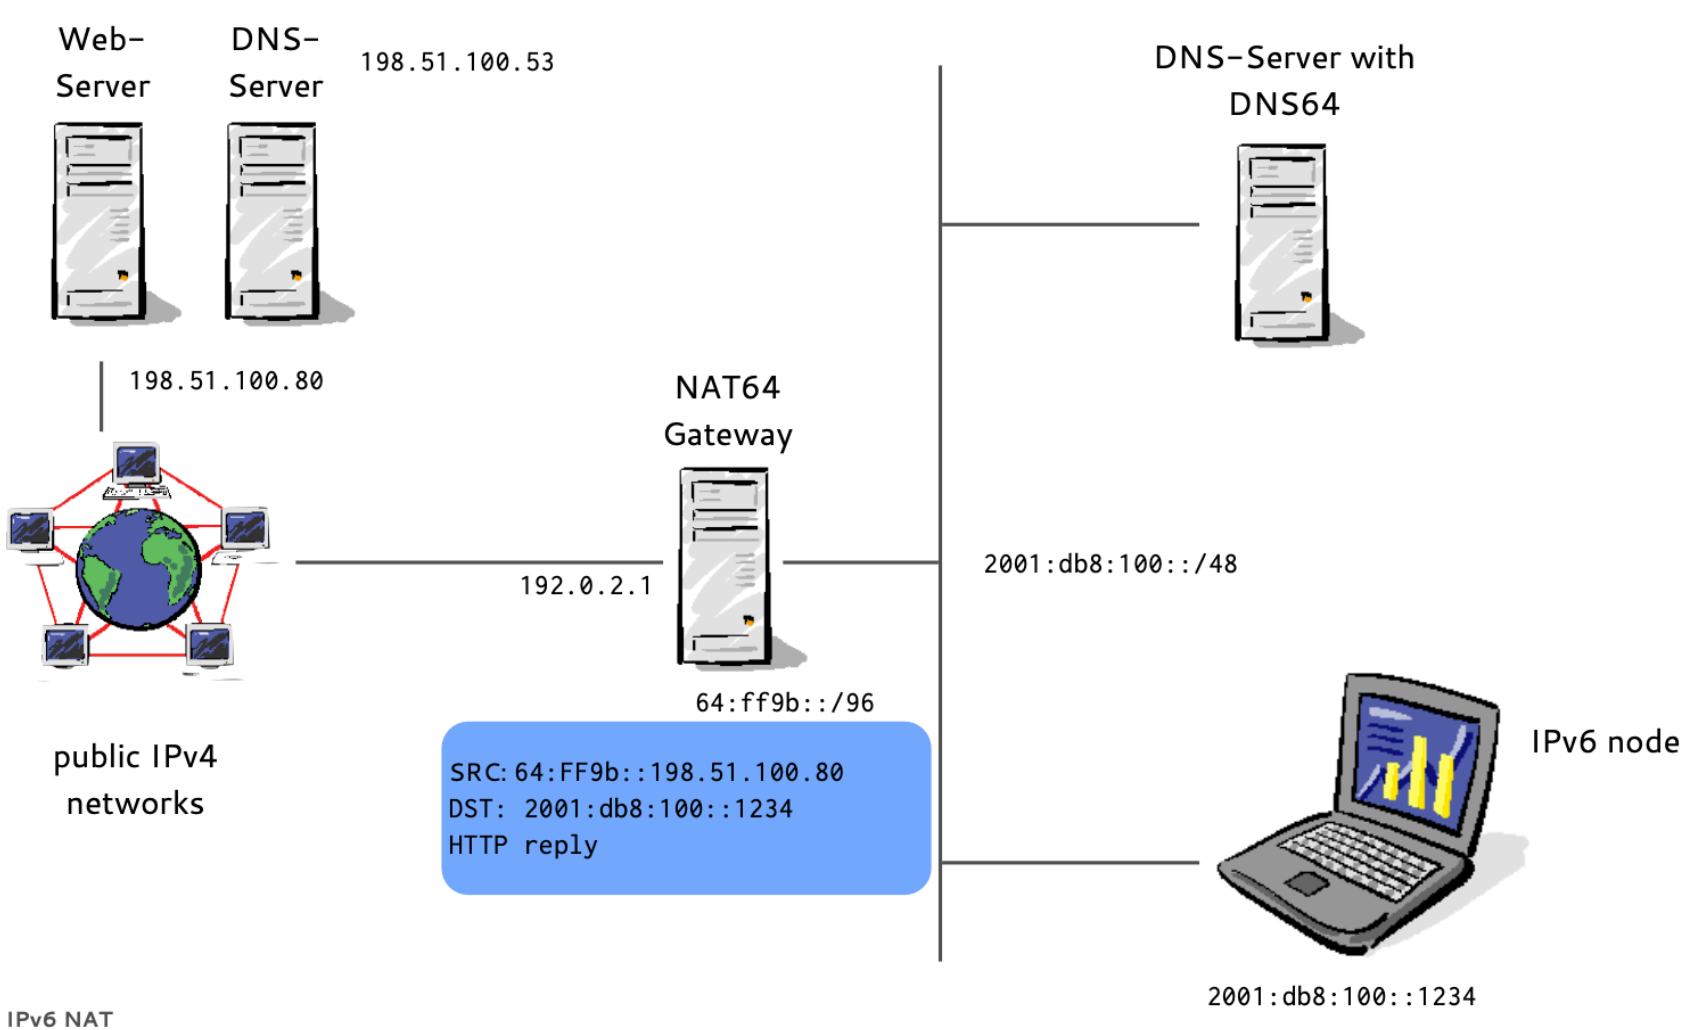
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
SRC: 64:FF9b::198.51.100.80
DST: 2001:db8:100::1234
HTTP reply
IPv6 node
2001:db8:100::1234

IPv6 NAT

DNS64/NAT64 example (19/19)

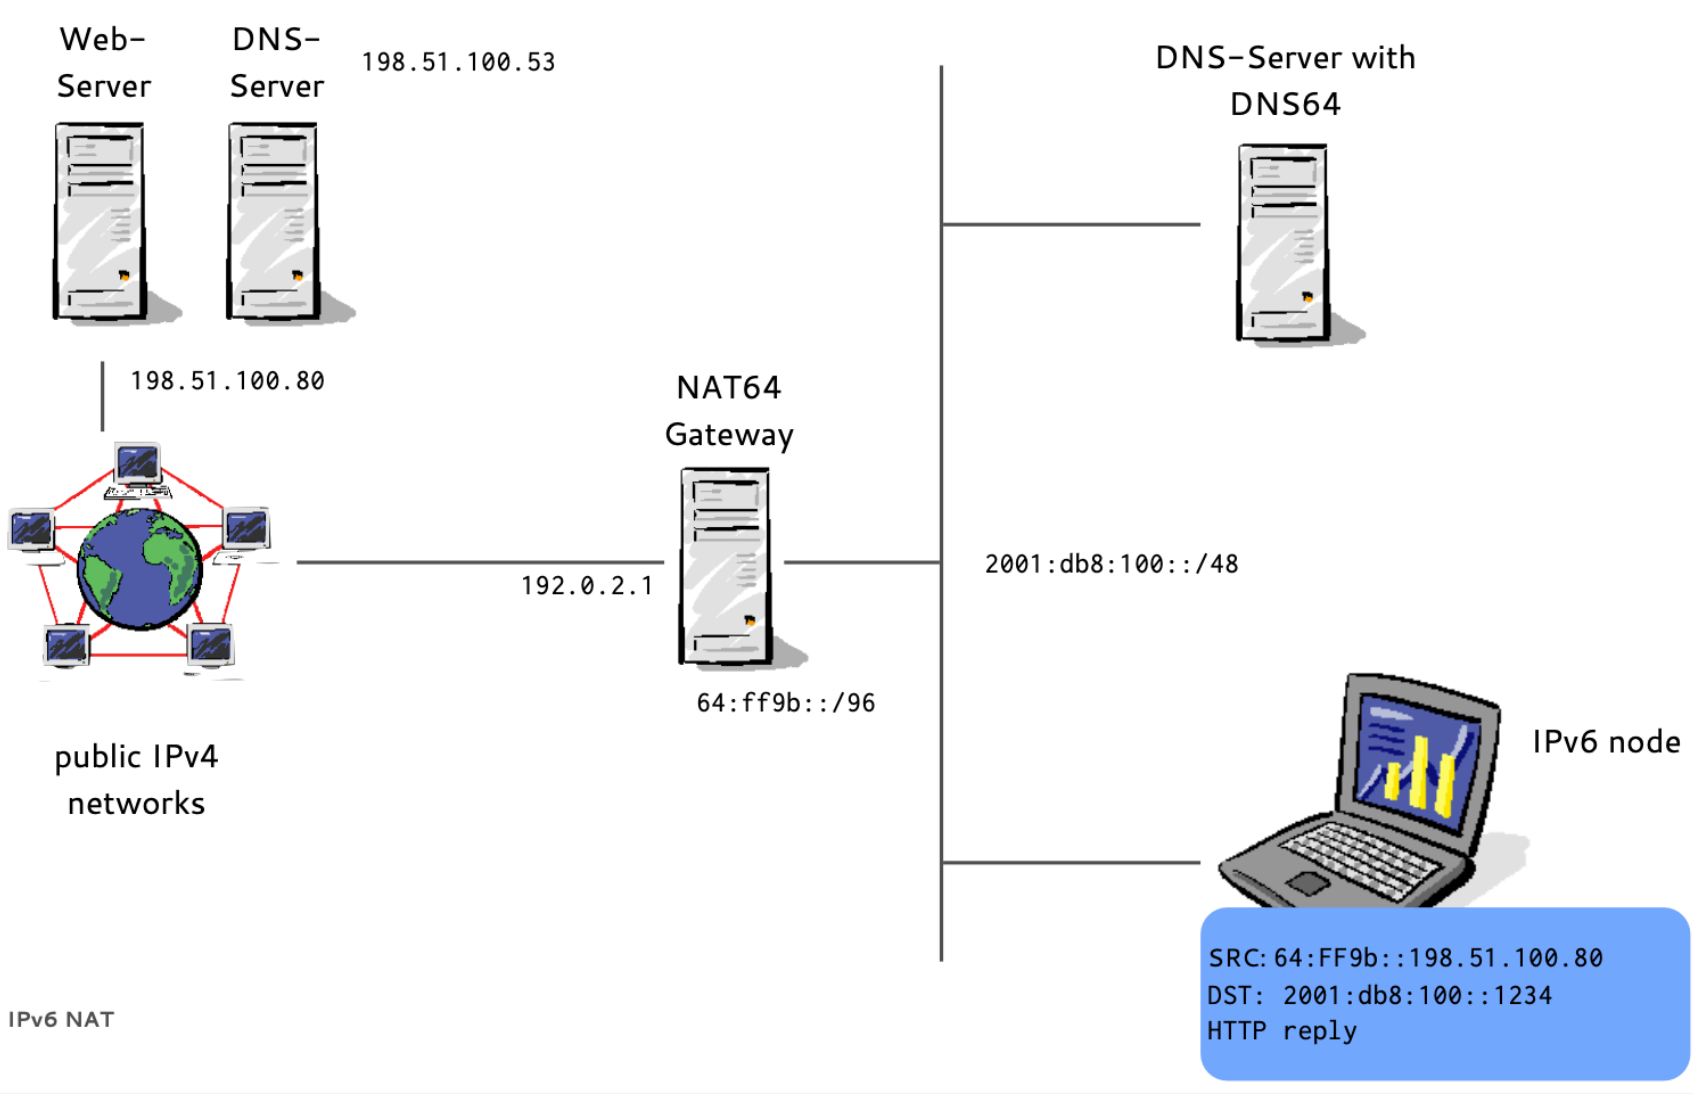
Web-
Server
DNS-
Server
198.51.100.53
DNS-Server with
DNS64
198.51.100.80
NAT64
Gateway
2001:db8:100::/48
192.0.2.1
64:ff9b::/96
public IPv4
networks
IPv6 node
SRC: 64:FF9b::198.51.100.80
DST: 2001:db8:100::1234
HTTP reply

IPv6 NAT

DNS64/NAT64

- NAT64 and DNS64 do not share any state
- Bridges the gap between "IPv6 only" and "IPv4 only" networks

The IP protocols

IPv4 world

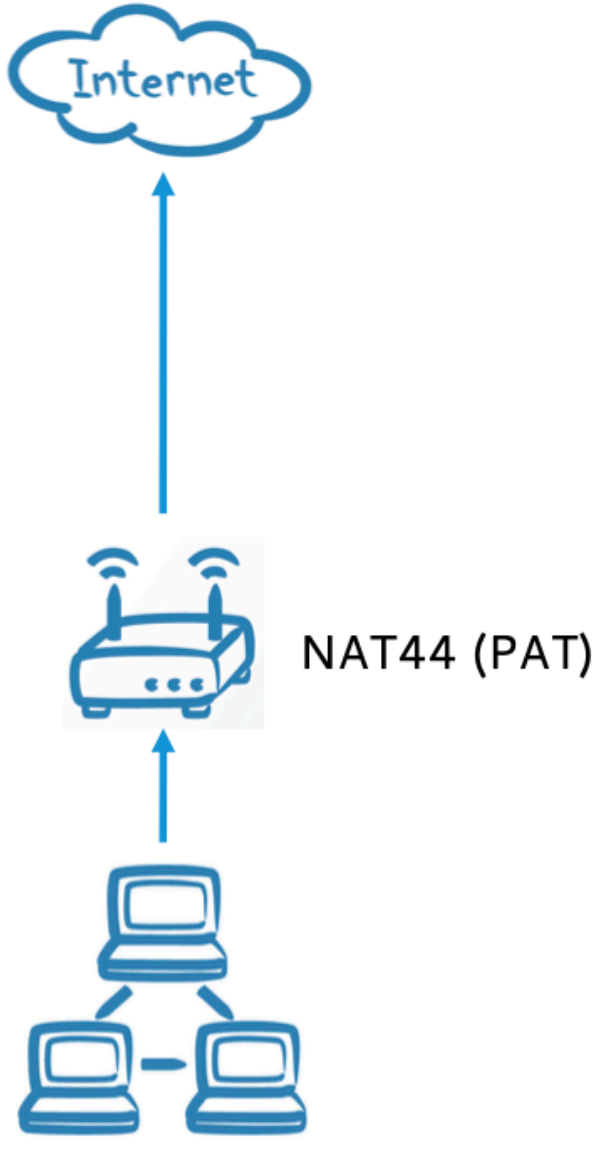
Internet
NAT44 (PAT)

IPv4 world today as envisioned by the IETF in 1996

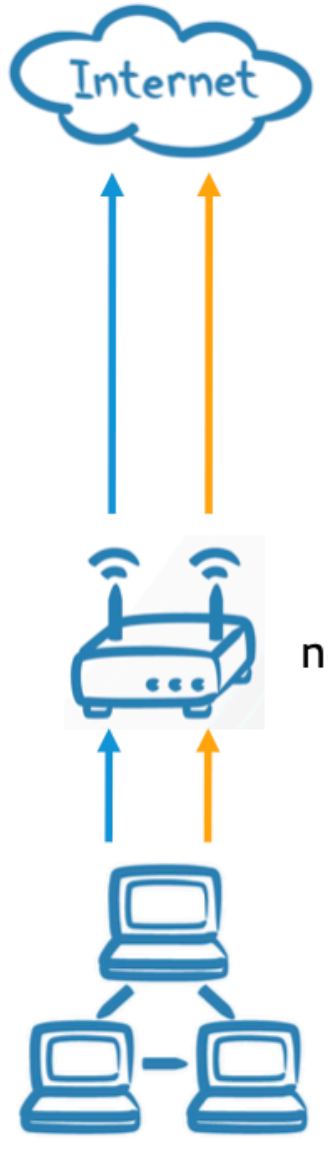

no NAT, transparent 1-to-1

Dual-stack issues

- The Internet is still growing
- New customers need to be connected
- Some "content" is still available on IPv4
- But no new IPv4 addresses are available

IPv4 for access networks

- Service provider try to find solutions to deliver IPv4 access over IPv6 only access networks
- Several solutions have been (and are developed) in the IETF:
  - DS-Lite/lightweight DS-Lite
  - 464XLAT
  - 4rd
  - MAP-E and MAP-T

Dual-Stack-Lite (DS-Lite)

DS-Lite

- DS-Lite is specified in ↪RFC 6333 "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion"
- DS-Lite provides IPv4 access over IPv6 using NAT and IPv4-in-IPv6 tunnel
    - Works with unmodified, legacy IPv4 applications and devices

DS-Lite

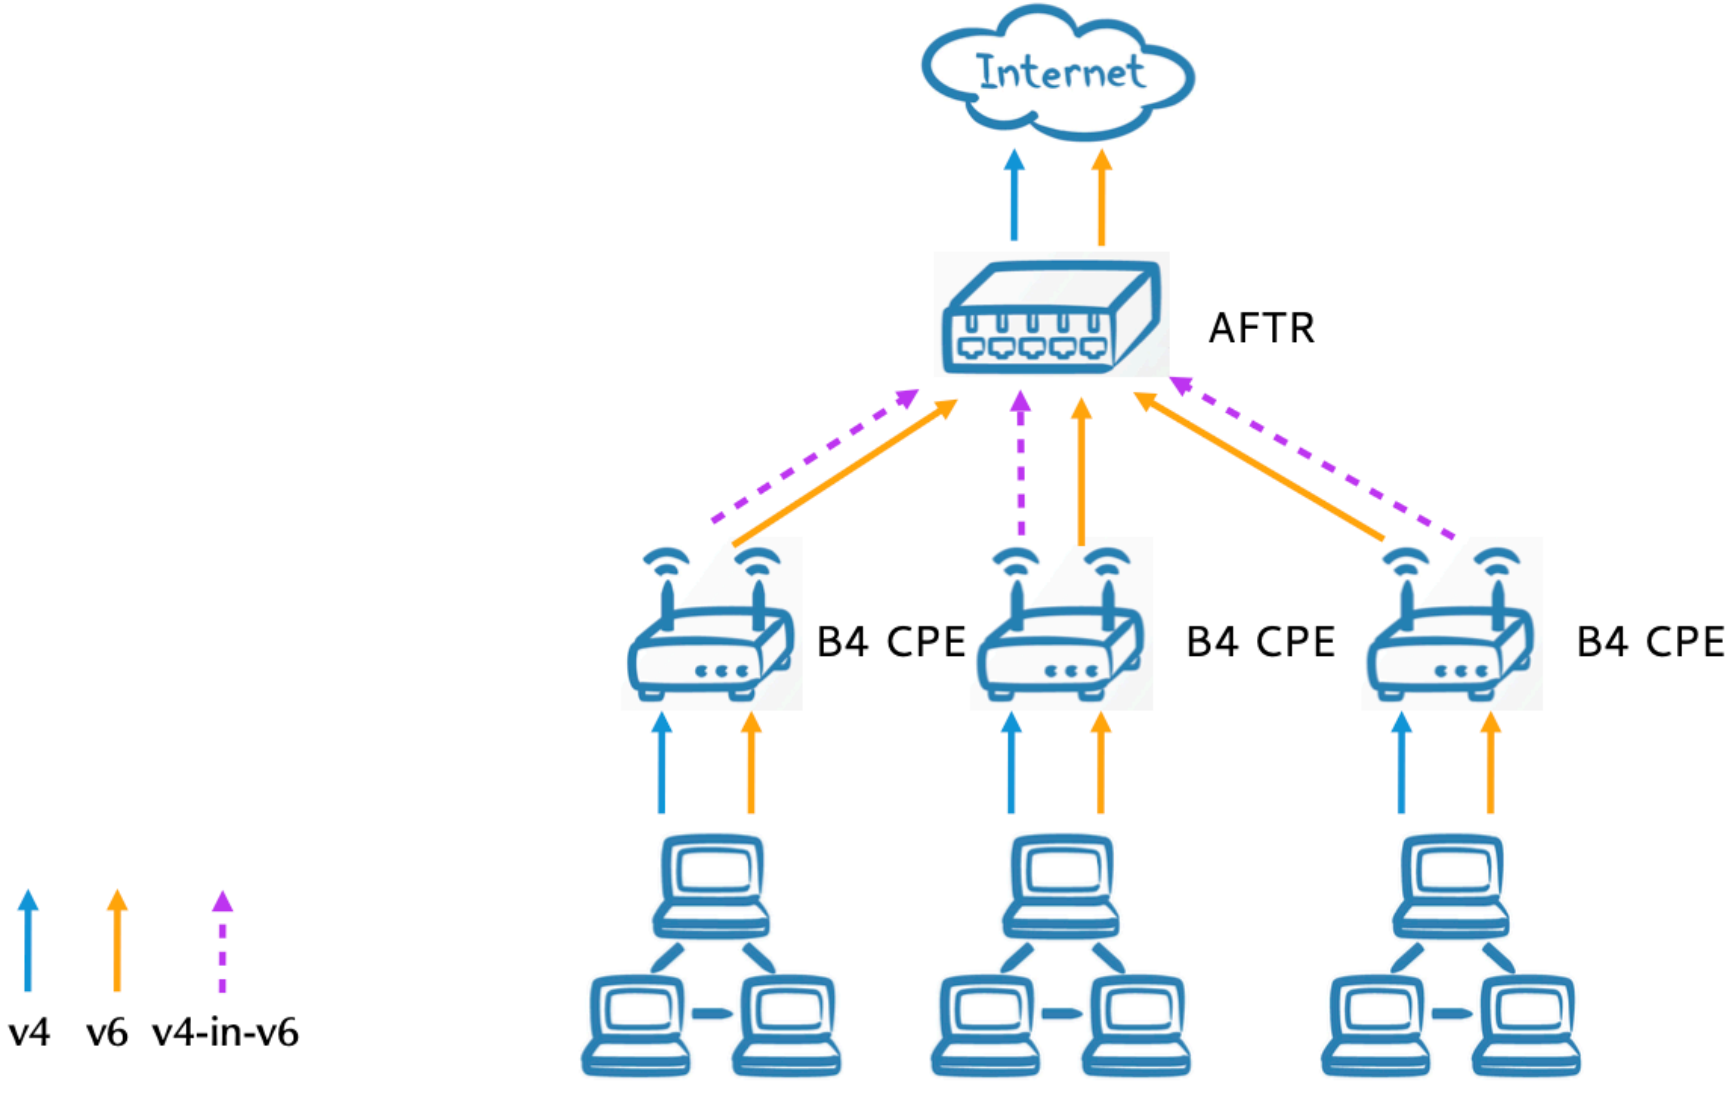
Internet
AFTR
B4 CPE
B4 CPE
B4 CPE
v4
v6
v4-in-v6

Dual-Stack Lite modules

- B4 ("Basic Bridging BroadBand"")
  - Usually build into CPE devices (CPE: Customer Premises Equipment)
  - Can also run on standard PC hardware

- Purpose:
  - Find the other end of the tunnel
  - Encapsulate IPv4 in IPv6

- No NAT on the B4 element!

Dual-Stack Lite modules

- AFTR (Address Family Transition Router)
  - Establish tunnel with many B4 elements
  - Accept and decapsulate IPv6 tunnel traffic
  - NAT between customer IPv4 and the IPv4 Internet

MAP-E

What is MAP-E

- MAP-E is defined in ↪RFC 7597 "Mapping of Address and Port with Encapsulation (MAP-E)" (Standards Track)
- MAP-E (Mapping of Address and Port with Encapsulation) is a protocol that allows IPv4 packets to be transported across an IPv6 network using encapsulation.
- MAP-E helps alleviate the issue of IPv4 address exhaustion by allowing multiple customer premises equipment (CPE) to share the same public IPv4 address through Carrier-Grade NAT.

MAP-E

- In a MAP-E deployment, a MAP-E CE device (such as a router) uses a Basic Mapping Rule (BMR) to configure itself with an IPv4 address, prefix, or shared IPv4 address from an IPv6 prefix.
- The BMR can also be used for forwarding packets in scenarios where an IPv4 source address and source port are mapped into an IPv6 address/prefix.

MAP-E

- MAP-E is an alternative to Carrier-grade NAT and DS-Lite, pushing the IPv4 IP address/port translation function into the existing customer premises equipment IPv4 NAT implementation, thus avoiding the NAT444 and statefulness problems of carrier-grade NAT.

MAP-E

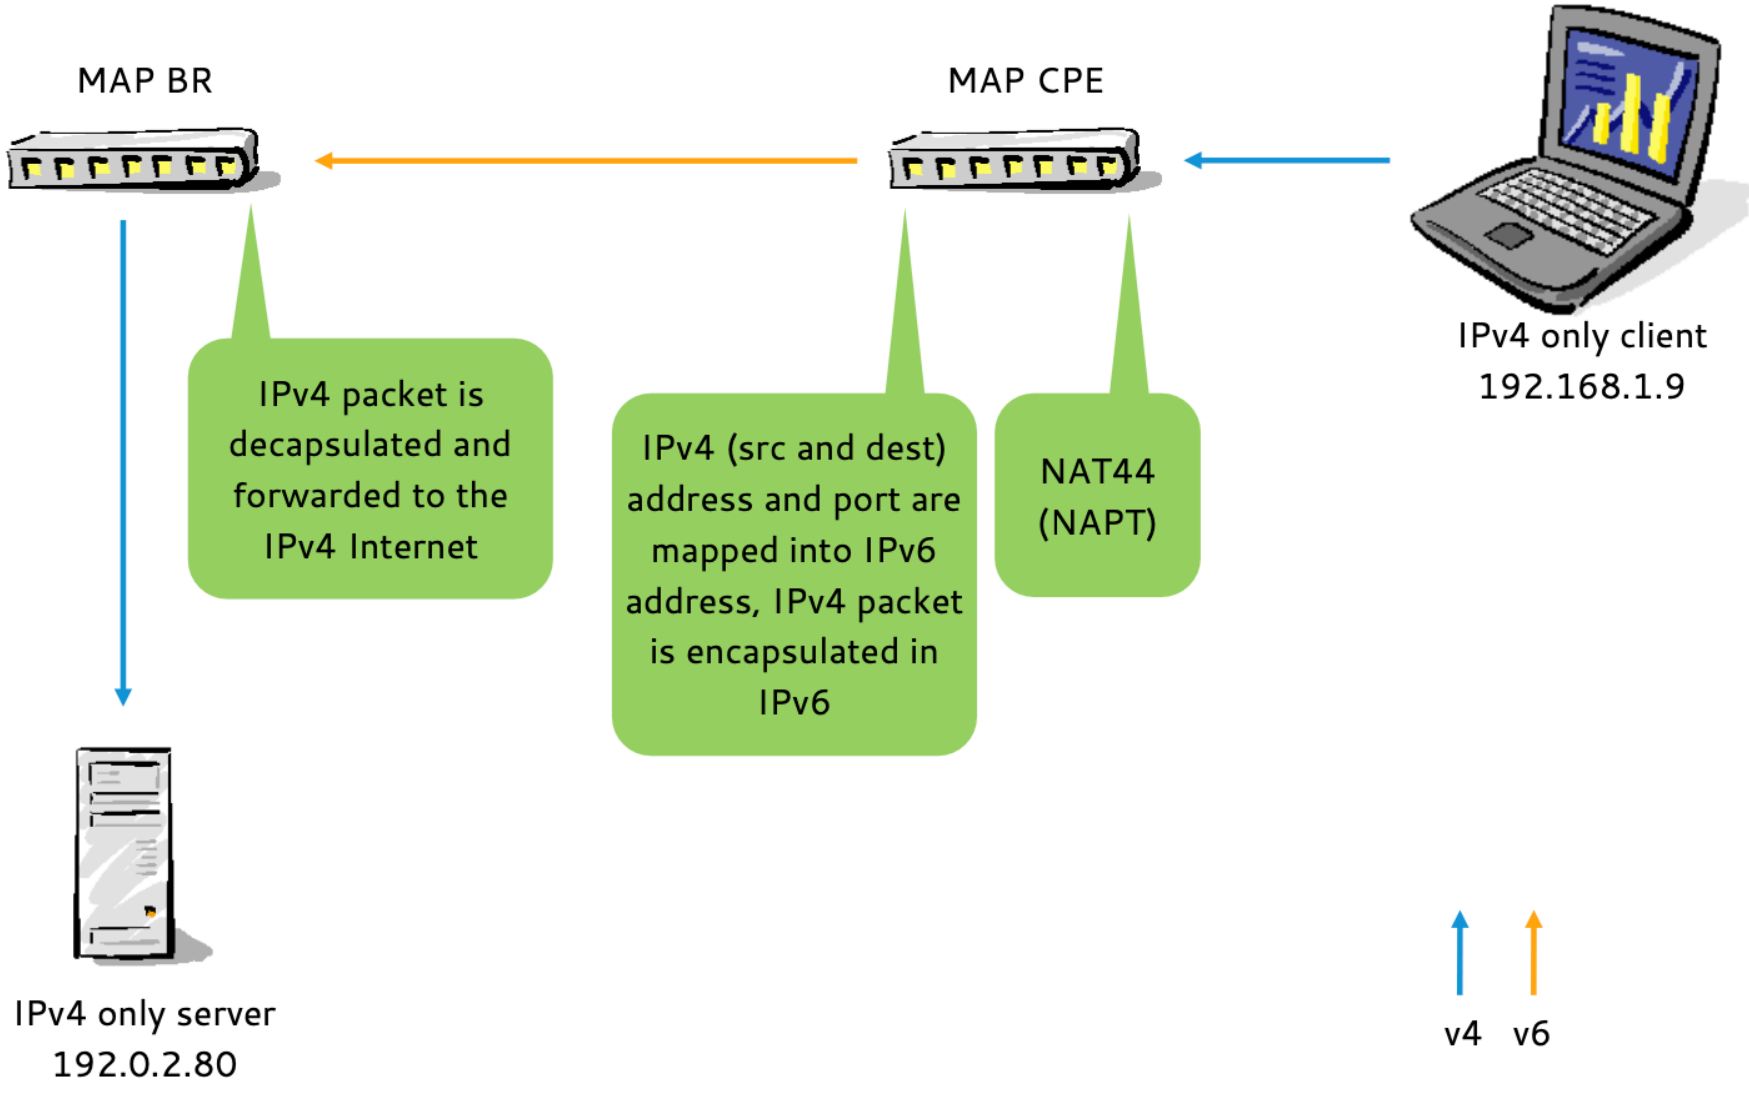
MAP BR
MAP CPE
IPv4 only client
192.168.1.9
IPv4 packet is decapsulated and forwarded to the IPv4 Internet
IPv4 (src and dest) address and port are mapped into IPv6 address, IPv4 packet is encapsulated in IPv6
NAT44 (NAPT)
IPv4 only server
192.0.2.80
v4
v6

MAP-E

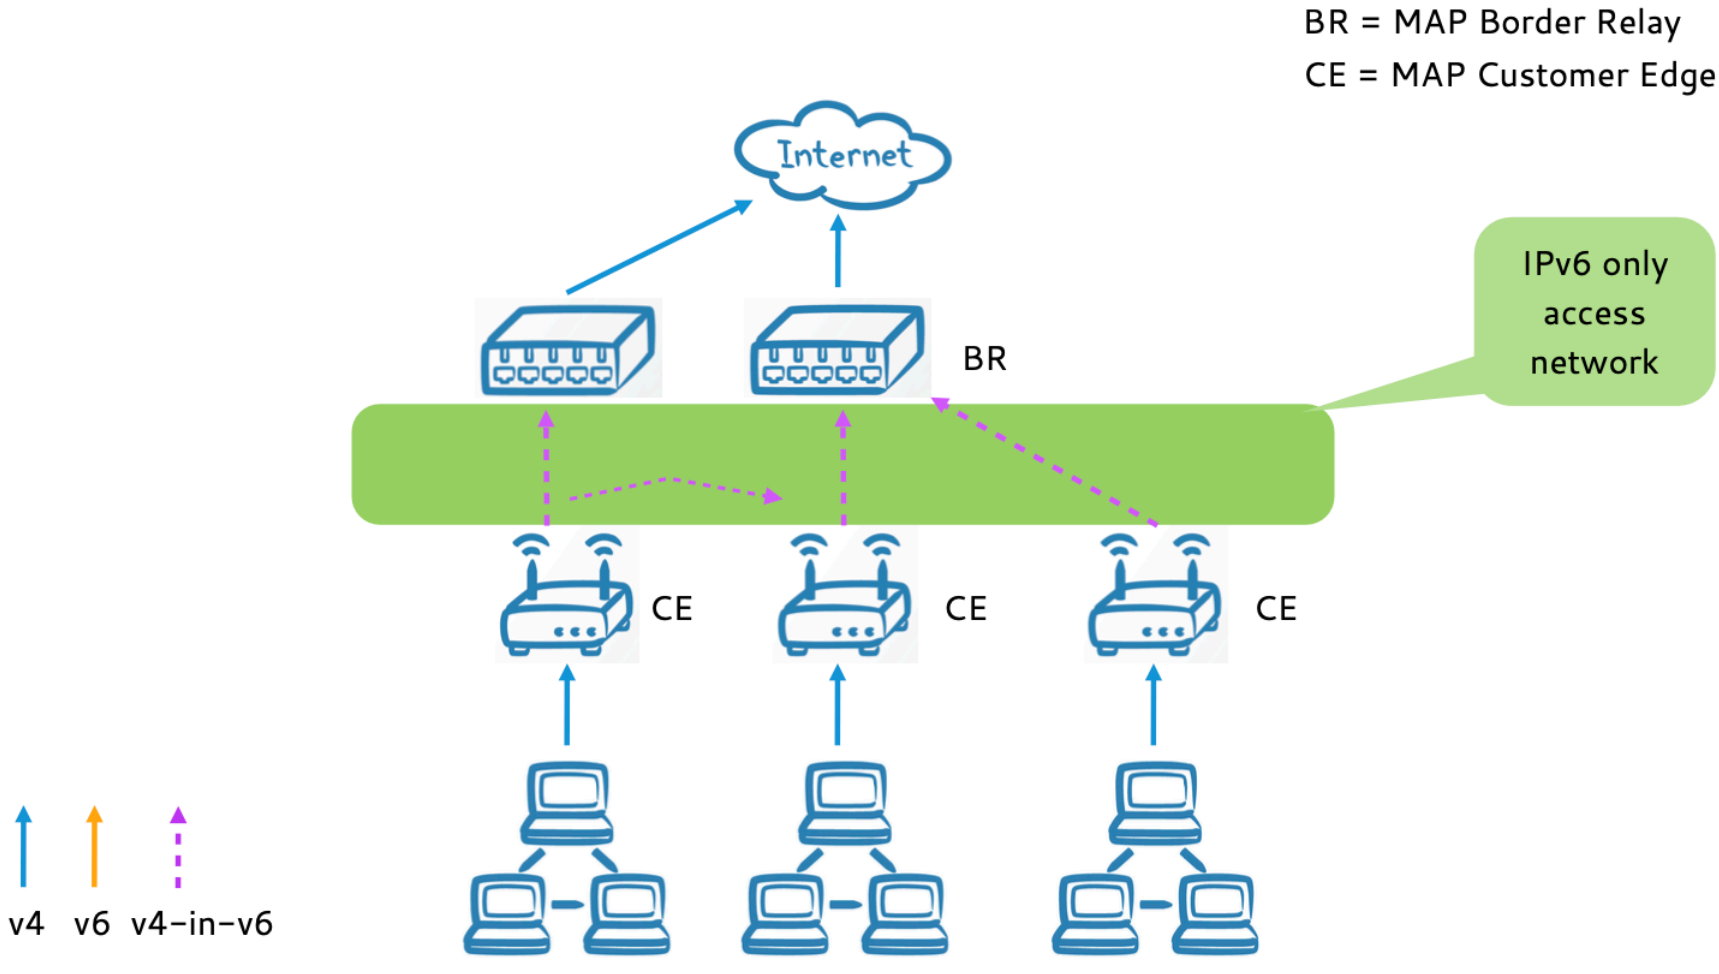
BR = MAP Border Relay
CE = MAP Customer Edge
Internet
BR
IPv6 only access network
CE
CE
CE
v4
v6
v4-in-v6

MAP-E

- MAP-E supports
  - Shared IPv4 addresses (multiple CEs sharing one IPv4 address) on the BR (Border Relay)
    - The customer edge device (CE) performs NAT44
    - Each CE negotiates with the BR a UDP/TCP port range assigned it can use

- A single IPv4 address per CE configured on the BR
- An IPv4 prefix per CE configured on the BR

MAP-E

- The default MAP mode of operation is a "hub-and-spoke" mode (BMR - basic mapping rule)
  - All traffic from a CE must pass a BR for forwarding
- Additionally, there can be one or more "forwarding mapping rules (FMR)" installed on the CEs to enable a "mesh" mode
  - Traffic can be forwarded between CEs between passing through a (central) BR

MAP-E

- The MAP parameters and IPv4 address(es) for the CE are be provisioned via
  - DHCPv6
  - "TR-69" Broadband Forum's Residential Gateway management interface
  - XML-RPC over IPv6

Carrier Grade NAT (CGN)

CGN

- Carrier Grade NAT is a label given to IPv4 depletion mitigation technologies that require network address translation at the provider end of an access network
- Examples:
  - NAT444 (NAT IPv4-to-IPv4 at the CPE, another NAT IPv4-to-IPv4 at the Border Router)
  - (original) DS-Lite
  - 4rd

Challenges with CGN

- CGN creates service regression for network functions or services
    - Many network functions can break or have non-optimal performance in CGN environments

- ↪RFC 6269 "Issues with IP Address Sharing" gives additional details

Challenges with CGN

- Restricted allocations of outgoing ports will impact performance for end-users
- Incoming port negotiation mechanisms may fail
- Incoming connections to Assigned Ports will not work
- Port discovery mechanisms will not work
- Assumptions about parallel/serial connections may fail
- Reverse DNS will be affected
- Inbound ICMP will fail in many cases

Challenges with CGN

- Amplification of security issues will occur
- Fragmentation will require special handling
- Single points of failure and increased network instability may occur
- Port randomization will be affected (DNS cache poisoning)
- Frequent keep-alives will reduce mobile device battery life

Challenges with CGN

- Service usage monitoring and abuse logging will be impacted for all elements in the chain between service provider and content provider
  - Penalty boxes will no longer work
  - Spam blacklisting will be affected
  - Geo-location services will be impacted
  - Geo-proximity mechanisms will be impacted
  - Traceability of network usage and abusage will be affected
- IPv6 transition mechanisms will be affected

Deprecated IPv6 NAT technologies

Deprecated IPv6 NAT technologies

- Network Address Translator - Protocol Translator (NAT-PT) and Network Address Port Translator - Protocol Translator (NAPT-PT) has been defined in the beginning of IPv6 NAT design work (↪RFC 2766, February 2000)
  - NAT-PT and NAPT-PT have several security and scalability issues, which lead to deprecation of these NAT solutions
  - NAT-PT and NAPT-PT have been moved to historical status in ↪RFC 4966 (July 2007)

Questions?