

Getting your hands wet ...

Deploying IPv6

planning, common pitfalls
and security-considerations

About me

- ▶ Stefan Neufeind
- ▶ From Neuss (near Düsseldorf, Germany)
- ▶ Working for SpeedPartner GmbH
(consulting, development, administration)
- ▶ Peering management for / founding member of
OpenCarrier e.G. (ISP-association)

Agenda

- ▶ Levels to consider for deployment
- ▶ Ways to connect
- ▶ Getting IPv6-addresses
- ▶ Address-planning
- ▶ Configuration
- ▶ Aspects to consider when deploying

- ▶ Open part

Levels to consider for deployment

- ▶ End-users / clients (Linux, Windows, ...)
 - ▶ Using services (DNS, Web, Mail, ...)
- ▶ Network equipment
 - ▶ Routers, firewalls, ...
- ▶ Servers
 - ▶ Offering services
 - ▶ Interacting with others

Ways to connect

- ▶ Migration techniques
 - ▶ 6-to-4
 - ▶ Teredo
- ▶ Connectivity in multiple ways
 - ▶ Native (datacenter, fixed line, PPPoE, ...)
 - ▶ Through tunnels (IPv6-in-IPv4, GRE, L2TP / softwires, ...)
 - ▶ Static
 - ▶ Tunnel-brokers

Ways to connect

- ▶ Example for migration technique: 6-to-4

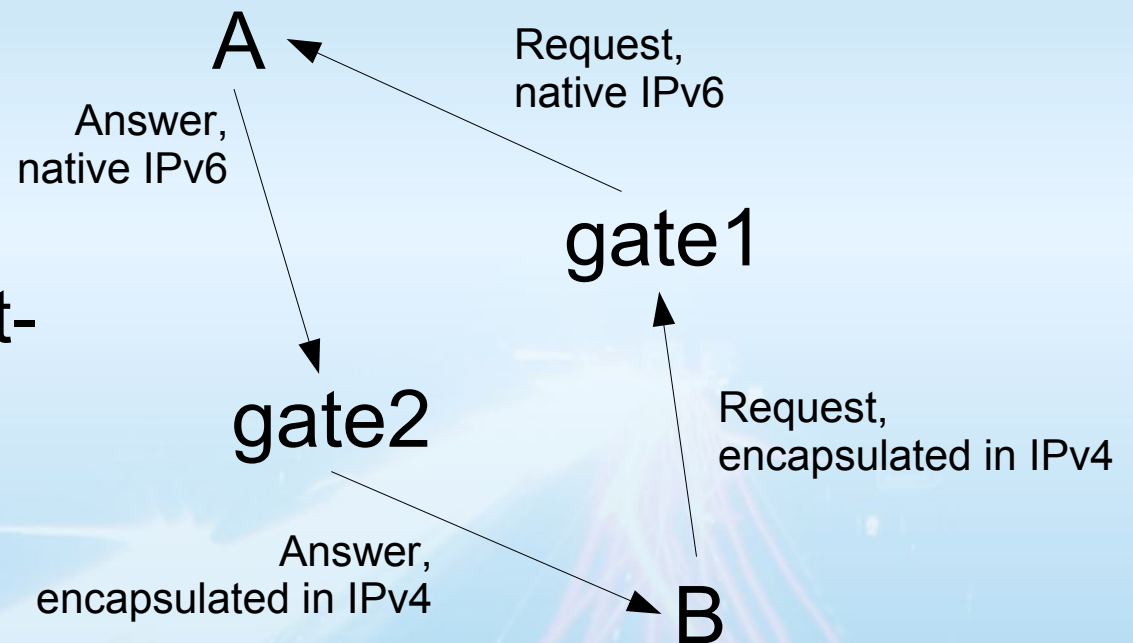
- ▶ Maps a /48 to each IPv4-address

91.184.34.11
↓ ↓ ↓ ↓
2002:5BB8:220B::/48

- ▶ Host encapsulates IPv6-packet in Ipv4
- ▶ Sends with public IPv4-address to a 6-to-4-gateway
- ▶ Either explicitly configured gateway – or anycast:
192.88.99.1 / 2002:c058:6301:: (RFC3068)

Ways to connect

- ▶ Example for migration technique: 6-to-4
- ▶ „Easy“
- ▶ But asynchronous
- ▶ And quality might vary (especially with anycast-chosen gateways)
- ▶ And needs public IPv4-address per host (no NAT)

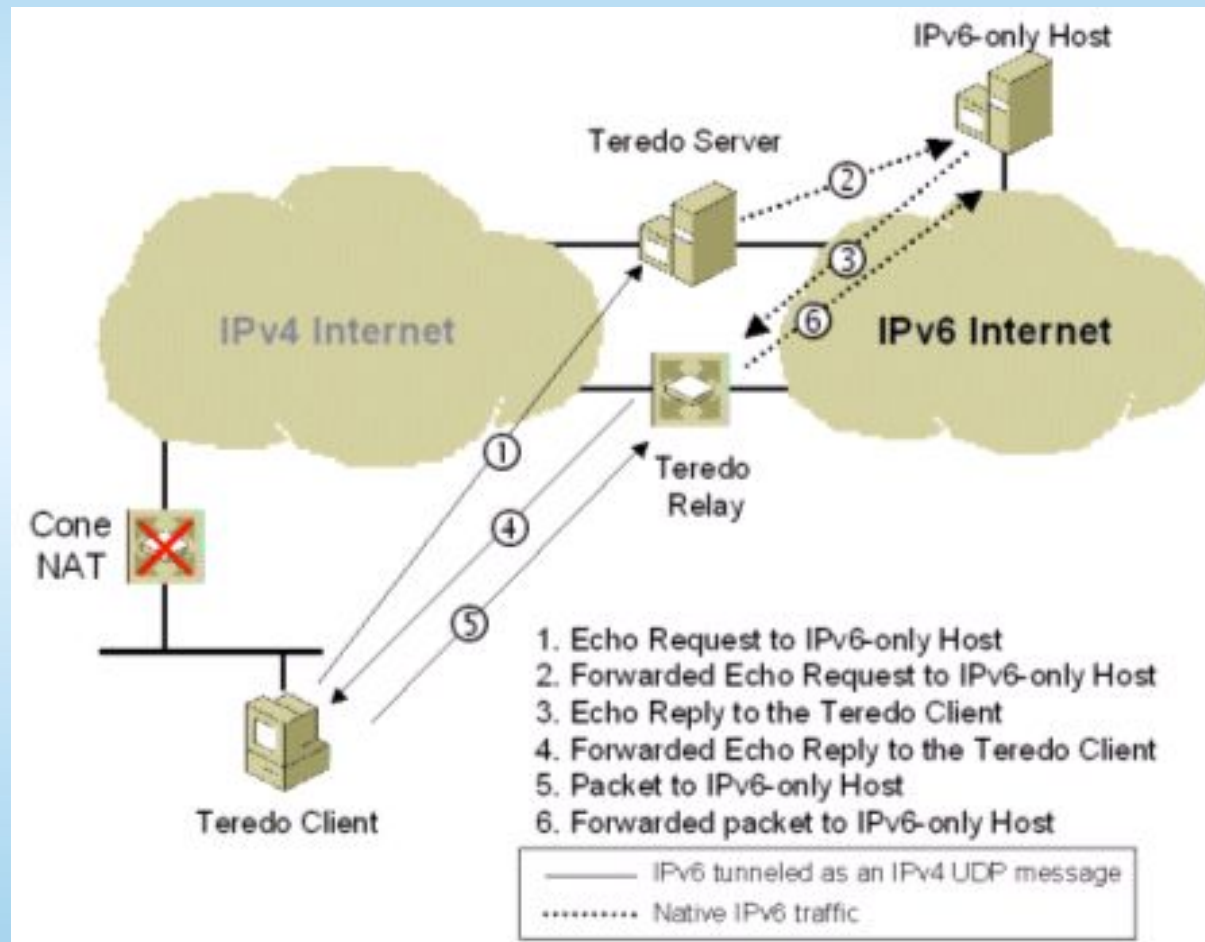


Ways to connect

- ▶ Example for migration technique: Teredo (only briefly!)
- ▶ Works over UDP-IPv4, also with NAT
- ▶ Teredo client: Only IPv4-connectivity
- ▶ Teredo server: For holepunching (with NAT)
- ▶ Teredo relay: For routing traffic
- ▶ Client-address constructed from Teredo-prefix (2001:0000::/32), NAT-type, UDP-port, NAT public IPv4-address

Ways to connect

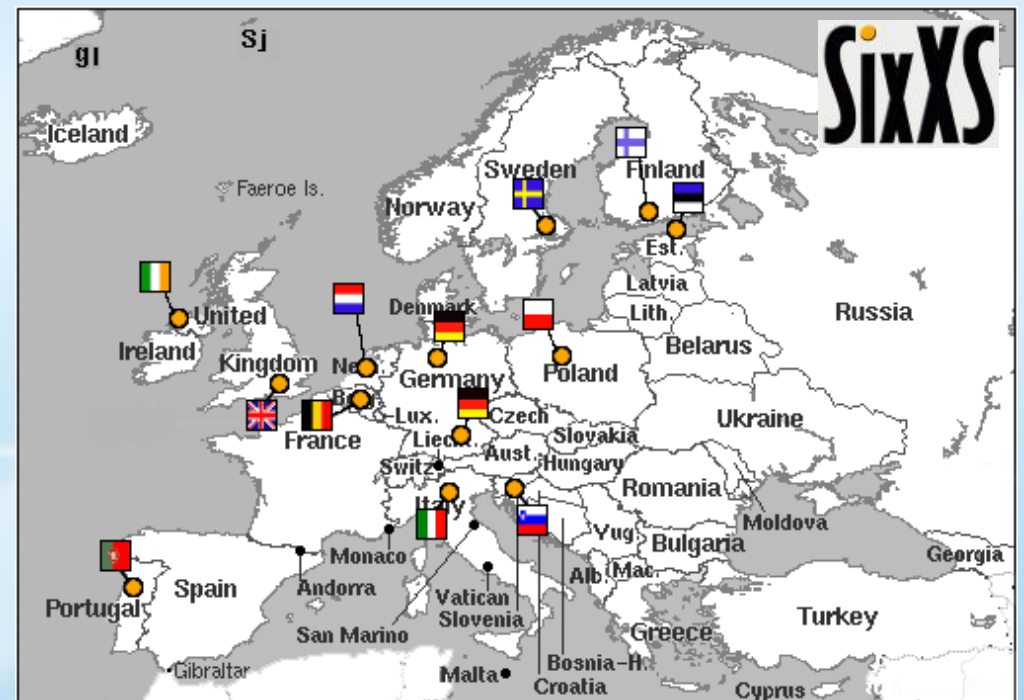
- Example for migration technique: Teredo (only briefly!)



Source:
Microsoft TechNet

Ways to connect

- ▶ Example for tunneled connectivity: SixXS
- ▶ Besides other things: IPv6-in-IPv4 with tunnel broker
- ▶ Handle e.g. registration and tunnel discovery
- ▶ Several POPs available
- ▶ Works with „aiccu“-client (available on several platforms including e.g. OpenWRT)



Ways to connect

- ▶ Example for tunneled connectivity: SixXS
- ▶ Statically assigned addresses (/64 and optionally /48) per tunnel
- ▶ Can work with changing Ipv4-addresses
- ▶ „Heartbeat protocol“ allows detection of stale tunnels
- ▶ But requires public Ipv4-address on host

Getting IPv6 addresses

- ▶ For single client / „small network“:
usually /64 prefix from provider
 - ▶ assigned
 - ▶ dynamically through tunnel/autodiscovery
- ▶ For network: prefix from provider (usually /48 prefix)
- ▶ PI-address-space?
 - ▶ Not yet in RIPE-region
 - ▶ Proposals were under discussion:
<http://ripe.net/ripe/meetings/ripe-52/presentations/ripe52-plenary-ipv6-pi.pdf> (by Jordi Palet)

Address-planning

- ▶ Interesting finding: „There are the same number of /32s in IPv4 space as in IPv6 space.“ :-)
(by david.conrad at icann.org on the global-v6-mailinglist)
- ▶ Fixed length of 16 bytes (128bit)
- ▶ Usually noted in 8 groups of each 4 hexadecimal digits
- ▶ Possible to leave out one or more groups of zeros with two colons (::) instead; e.g. 2001:0db8::1

Address-planning

- ▶ Subnetting usually done on /16-boundary (/32, /48, /64, ...)
- ▶ Minimum /64 if you want to use auto-configuration
- ▶ Subnetting/numbering of hosts up to you. Examples:
 - ▶ Sequentially
 - ▶ By building/POP, city, country, ...
 - ▶ Including VLAN-ID
 - ▶ AS-numbers (there are also 32-bit AS-numbers now!)
 - ▶ For fun: Why „A5“ in e.g. „2001:7F8:1::A504:1692:1“?
Looks like „AS“ ... serious! :-)

Configuration

- Routers

Example on Cisco, only new things to configure:

- Basics

```
ipv6 unicast-routing  
ipv6 cef
```

- Interfaces

```
interface GigabitEthernet 0/1  
  ipv6 address 2001:DB8::1/64  
  no ipv6 redirects  
  ipv6 nd suppress-ra  
  ipv6 ospf 1 area 0
```

- OSPF v3

```
ipv6 router ospf 1  
  redistribute connected
```

- BGP4-MP

```
router bgp 65000  
  address-family ipv6  
  neighbor 2001:DB8::2 remote-as 65001  
  neighbor 2001:DB8::2 activate
```


Configuration

- ▶ Routers
- ▶ Interfaces

Example on Juniper, only new things to configure:

```
interfaces {  
  ge-0/0/0 {  
    unit 1 {  
      family inet6 {  
        address 2001:db8::1/64;  
      }  
    }  
  }  
}
```

Configuration

- ▶ Routers
 - ▶ BGP4-MP
 - ▶ OSPF v3

Example on Juniper, only new things to configure:

```
protocols {
  bgp {
    group peers-kleyrex-v6 {
      type external;
      family inet6 {
        unicast;
      }
      neighbor 2001:db8::2 {
        peer-as 65001;
      }
    }
  }
  ospf3 {
    export OSPF3;
    area 0.0.0.1 {
      interface ge-0/0/0.791;
      interface ge-0/0/0.792;
    }
  }
}
```

Configuration

- ▶ End-user systems
- ▶ Servers
- ▶ Services
 - ▶ DNS / rDNS (as well as RIR-objects)
 - ▶ Web- / mailsystems etc.
 - ▶ Monitoring / accounting
(awstats, netflow/sflow, home-grown tools, ...)

Configuration

- Documenting network-relationships in RPSL:

```
mp-import: afi ipv6.unicast from AS34225 accept AS-SPEEDPARTNER  
mp-export: afi ipv6.unicast to AS34225 announce ANY
```

```
mp-import: afi ipv6.unicast from AS41692 accept AS-OPENCARRIER6  
mp-export: afi ipv6.unicast to AS41692 announce ANY
```

- Consider separate AS-sets for IPv4 / IPv6
in case of large difference of downstream-size IPv4 / IPv6

Configuration

➤ DNS in bind:

In a normal zone:

```
r1.example.com IN AAAA 2001:db8::1
```

For reverse-DNS:

```
$ORIGIN 0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
```

```
@ 86400 in SOA ns1.example.com. hostmaster.example.com. (
```

```
[...]
```

```
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 1D IN PTR r1.example.com.
```

➤ Reverse-delegation at RIPE:

```
Domain: 8.b.d.0.1.0.0.2.ip6.arpa
admin-c: ABCD-RIPE
tech-c:  ABCD-RIPE
zone-c:  ABCD-RIPE
nserver: ns1.example.com
nserver: ns2.example.com
mnt-by:  ABCD-MNT
```

Aspects to consider when deploying

- ▶ Does your gear support IPv6?
 - ▶ In software or hardware?
 - ▶ Since what firmware?
 - ▶ How stable?
- ▶ Running dual-stack?
- ▶ Performance: Latency/jitter in sub-optimal IPv6-scenarios
- ▶ Export-protocol for monitoring / accounting
 - ▶ Sflow? Netflow?
- ▶ Security-aspects (firewalls, ...)

Links

- ▶ 6to4 with Linux: <http://www.wlug.org.nz/6to4>
- ▶ Teredo:
<http://www.microsoft.com/technet/network/ipv6/teredo.msp>
- ▶ Public 6to4 routers: <http://www.kfu.com/~nsayer/6to4/>
- ▶ Collection of IPv6-content:
<http://www.sixxs.net/misc/coolstuff/>
- ▶ Audio/video, IPv6-IPv4 gateways, IPv6 newsservers
- ▶ IPv6 filter recommendations:
<http://www.space.net/~gert/RIPe/ipv6-filters.html>
<http://www.cymru.com/Bogons/ipv6.txt>

Thank you!

Slides available at:

<http://talks.speedpartner.de/>

Questions?

neufeind (at) speedpartner.de