

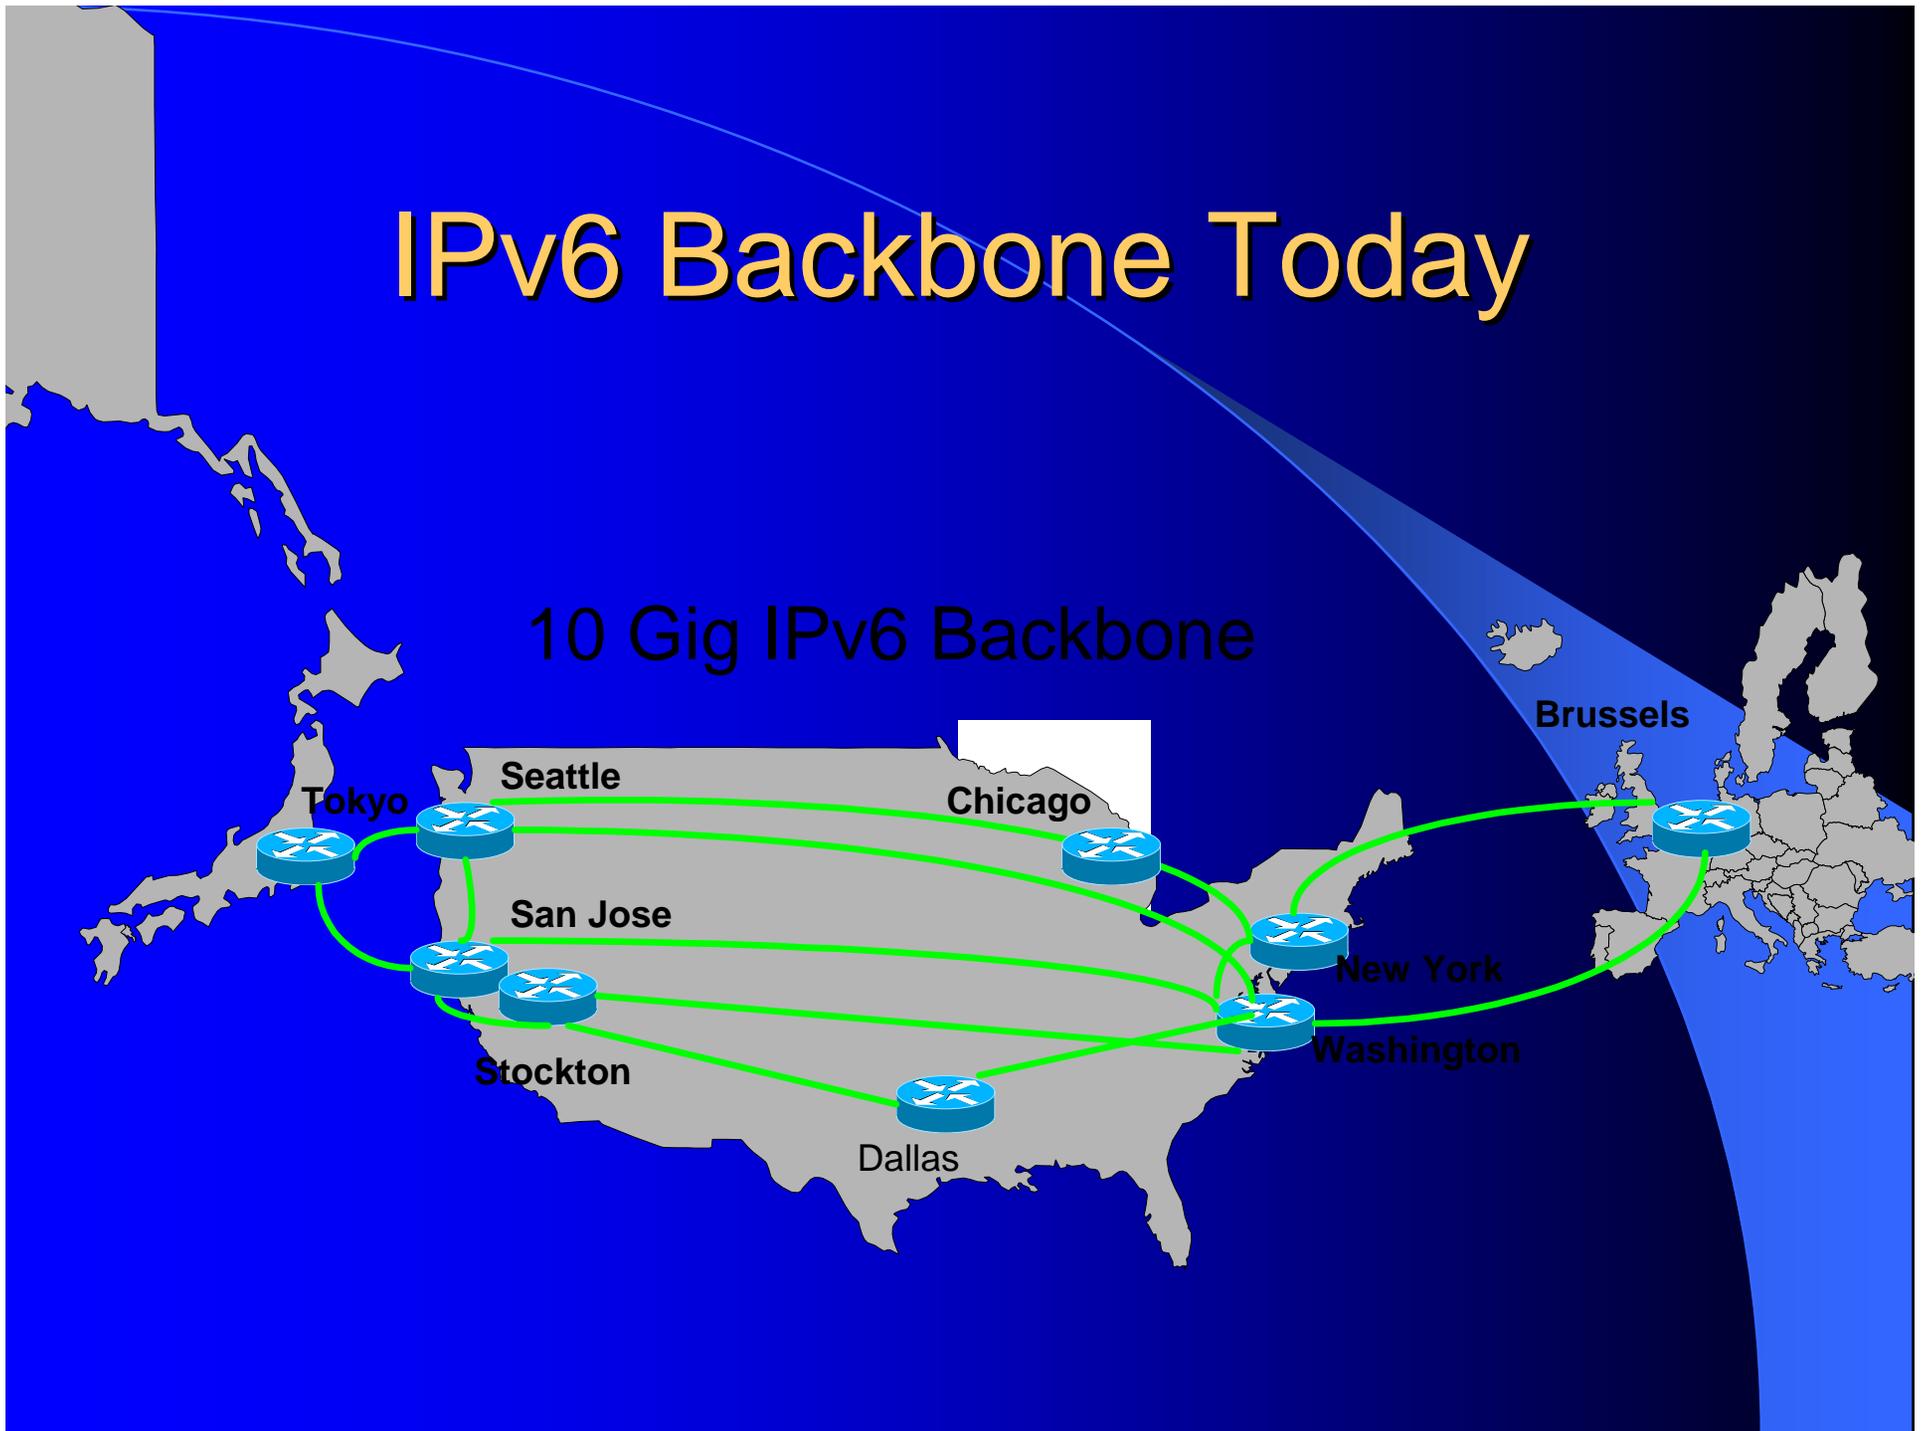
<marketing Logo here>

# IPv6 at Sprint

rrockell@sprint.net

# IPv6 Backbone Today

10 Gig IPv6 Backbone



# SprintLink IPv6 history

- 1997: Obtained 6bone address space (3ffe:2900::/24)
  - Original router under my desk ☺
- 1998: Totaling 15 customers using tunnels to 6bone
- 1999: Totaling 40 customers using tunnels to 6bone
  - Move router out to the network...
- 2000: Obtained ARIN space (2001:440::/35 → /32)
  - EOY 2000 Total: 110 customer using tunnels to 6bone.
- 2001-2002: Added 4 more IPv6 capable PoP's
  - Brussels, Washington DC, San Jose, New York
  - Member of the NY6IX exchange
  - Turning up customers at 2-3 per week.
- 26-May-2004: 300 Ipv6 Tunneled Connections; 2 Native
  - Request frequency has slowed considerable (1 every week)

# Sprint IPv6 offering

- Routers are IPv6 stand-alone boxes
  - No dynamic protocol-level interaction with the IPv4 network (SprintLink; AS1239).
- GRE tunneling is used between IPv6 routers
  - Over SprintLink (IPv4) infrastructure.
- iBGP full-mesh between AS6175
- ISIS runs as IGP
  - Looks eerily similar to SprintLink backbone, but with Hex addresses.

# Why we did it this way

- Router OS Dependancies:
  - Either features/code we need to run IPv4 do not support IPv6, or visa versa.
- Customer requirements:
  - Customer: “Yes, we require IPv6, IMMEDIATELY!”
  - Sprint: “Yes Sir! How much will you pay for IPv6?”
  - Customer: “Pay? No, we just \*WANT\* it. We don’t want to \*PAY\* for it”
- Overlay model removes Router software dependencies
  - Allows for more experimentation on the IPv6 side of things.
  - Allows us to deploy minimum capital boxes (or depreciated hardware) to support IPv6 for the price-point that customers require (to wit; \$0.00/meg).
- Protocol is not ‘fully-cooked’ yet.

# What this allows us to do.

- Offering is free to any IPv4 customer of Sprint
  - Or any entity with a static tunnel endpoint (if you are nice to us).
    - This is what customers expect to pay for IPv6 today.
- Goal is to promote the usage of IPv6, within the confines of the current abilities of the protocol
  - We are very careful not to bend the rules (rfc2772)
    - This brings the “issues” with the current rules to light, and forces attention upon them
    - Multi-Homing, Micro-mobility, renumbering, etc...
- Today, Sprint uses 6Bone space for all customer numbering
  - We assume that ‘real’ ipv6 numbering schemas will evolve with the protocol, and do not want to get ourselves stuck

# Why is the IPv6 network Tunneled??

```
sl-bb1v6-rly#sho int pos 0/0/0 | inc rate
```

1 minute input rate 181000 bits/sec, 64 packets/sec

1 minute output rate 189000 bits/sec, 65 packets/sec

```
sl-bb20-nyc#sho int pos 2/0 | inc rate
```

1 minute input rate 3001474000 bits/sec, 1057309 packets/sec

1 minute output rate 5028579000 bits/sec, 1303165 packets/sec

Time: 1930 EST, 10-Dec-2003

# Next steps

- Build the native transition plan:
  - Somewhat done.
- Use Layer 2 protocol ID field to dynamically tunnel IPv6 packets to overlay network.
  - Gives customers the native ‘feel’ without the native ‘pain’.

# Why does Sprint not lead the charge to Migrate?

- 3 Main Reasons:
  - We do what brings in money
    - IPv6 is a definite “check-box” on RFPs, but it is not a stand-alone profit generator for Sprint (yet)
  - The Router Vendors do not compel us to do so.
    - Router Vendor Software is not there yet. It has IPv6 forwarding, but misses the things that would entice enterprises
      - Next-headers, IPSec, Mobility extensions, etc..
  - Multi-Homing is not solved yet.
    - If you do IPv6 now you might either:
      - Run out of Router memory (at scale)
      - Run out of Addresses (eventually)
    - Something has to give.
    - Solution MIGHT leave IPv6 looking VERY different than it does today.

# Vendors: the vicious circle

Enterprises don't deploy  
it (for real)

There isn't software  
support for IPv6

Carriers aren't required  
to have it (for real)

Hardware Vendors are  
not driven to develop it  
(for real)

