IPv6

Frank Kargl

Chaos Computer Club – Ulm

frank.kargl@ulm.ccc.de

Übersicht

- Der Weg von IPv4 zu IPv6
- Das IPv6 Datagramm
- IPv6 Adressen
- Protokolle und Abläufe
- IPv6 Netzwerke
- Linux + IPv6: Der eigene IPv6 Anschluss

Warum ein neues IP?

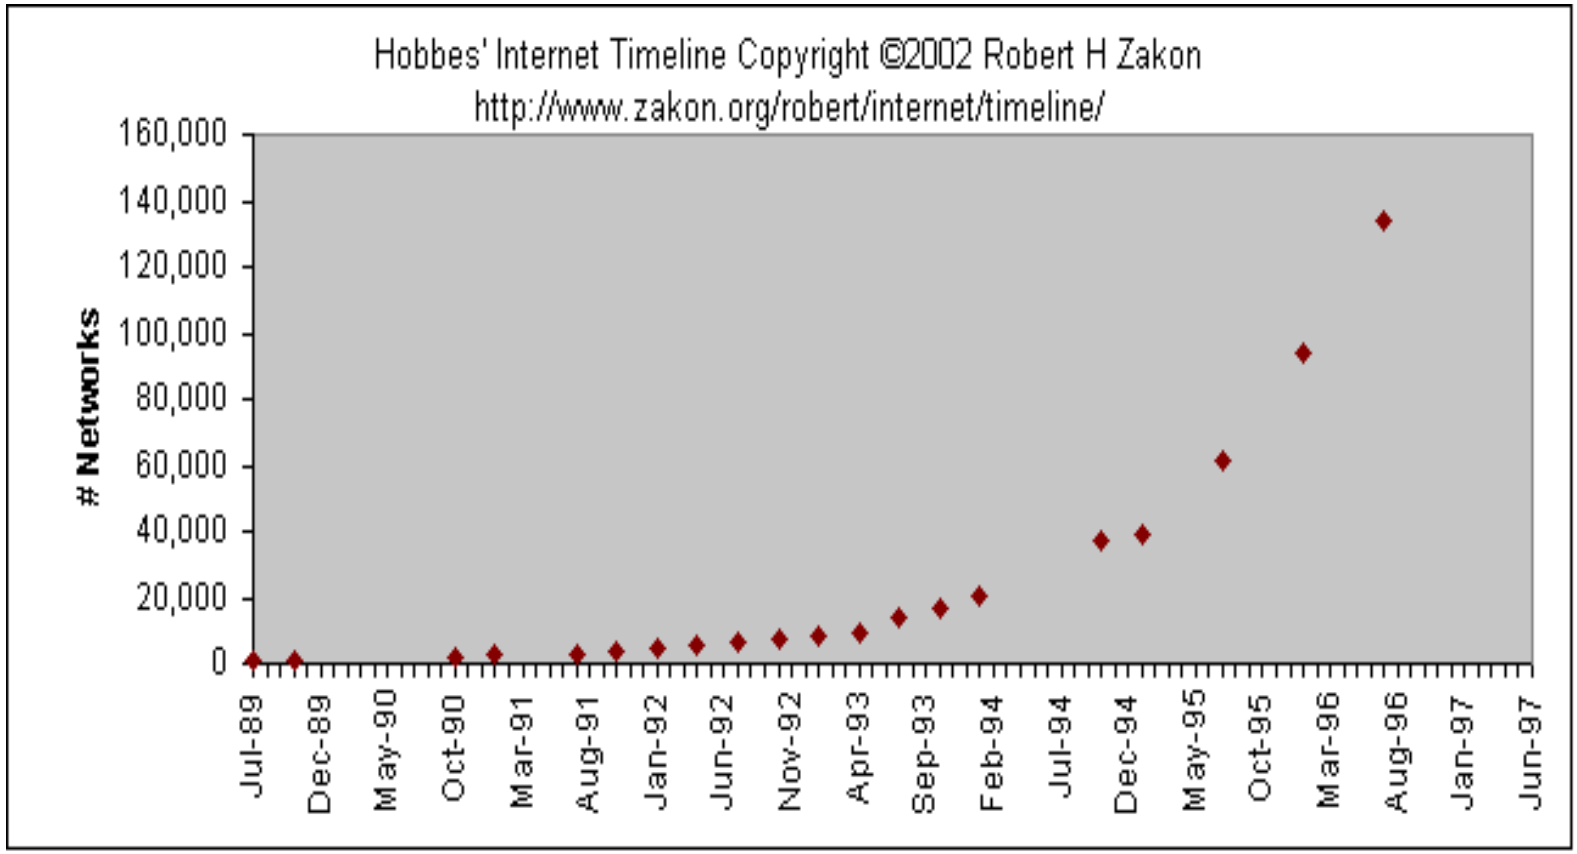

Warum ein neues IP?

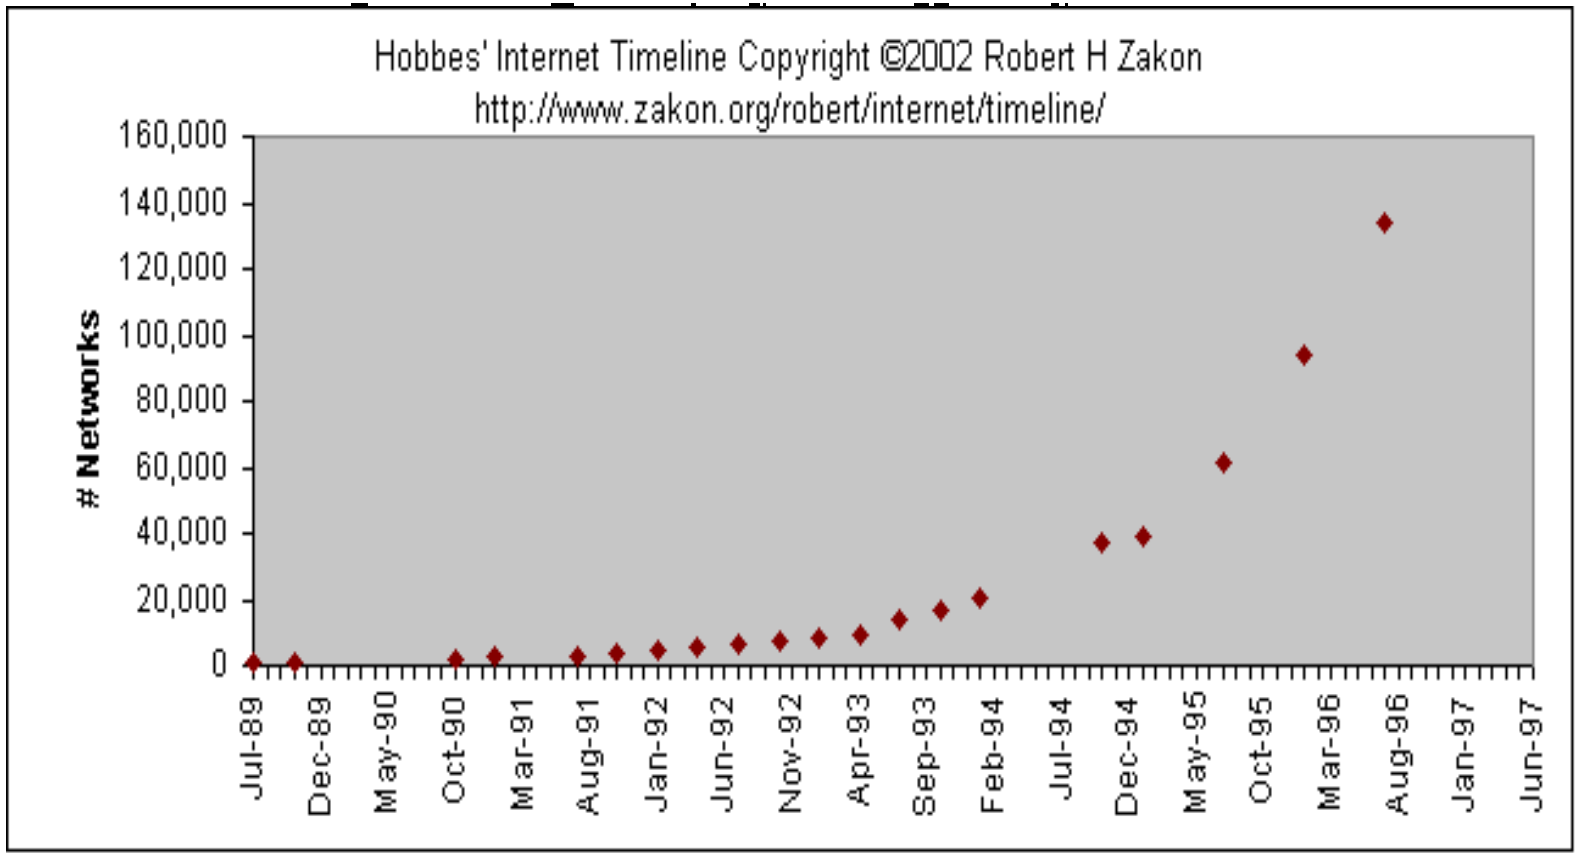

Anforderungen

RFC 1726 "Technical Criteria for Choosing IP, the Next Generation (IPnG)":

- Skalierbarkeit
- Topologische Flexibilität
- Performanz
- Robustheit
- Transitions-Strategie
- Medien-Unabhängigkeit
- "Unreliable Datagram Service"
- Konfiguration, Administration und Betrieb
- Sicherheit

Anforderungen (Forts.)

RFC 1726 "Technical Criteria for Choosing IP, the Next Generation (IPnG)":

- Eindeutige Bezeichner
- Freie Standards
- Multicast
- Erweiterbarkeit
- Dienstegüte
- Mobilität
- Kontroll-Protokoll
- Private Netzwerke

Kandidaten

Jan. 1995: RFC 1752 "The Recommendation for the IP Next Generation Protocol":

- Common Architecture for the Next Generation Internet Protocol (CATNIP)
  Integration von OSI und Novell
- Simple Internet Protocol Plus (SIPP)
  Evolution von IPv4
- TCP/UDP with Bigger Addresses (TUBA)
  Ersetzt IPv4 Netzwerk mit ISO CLNP

Entwicklung

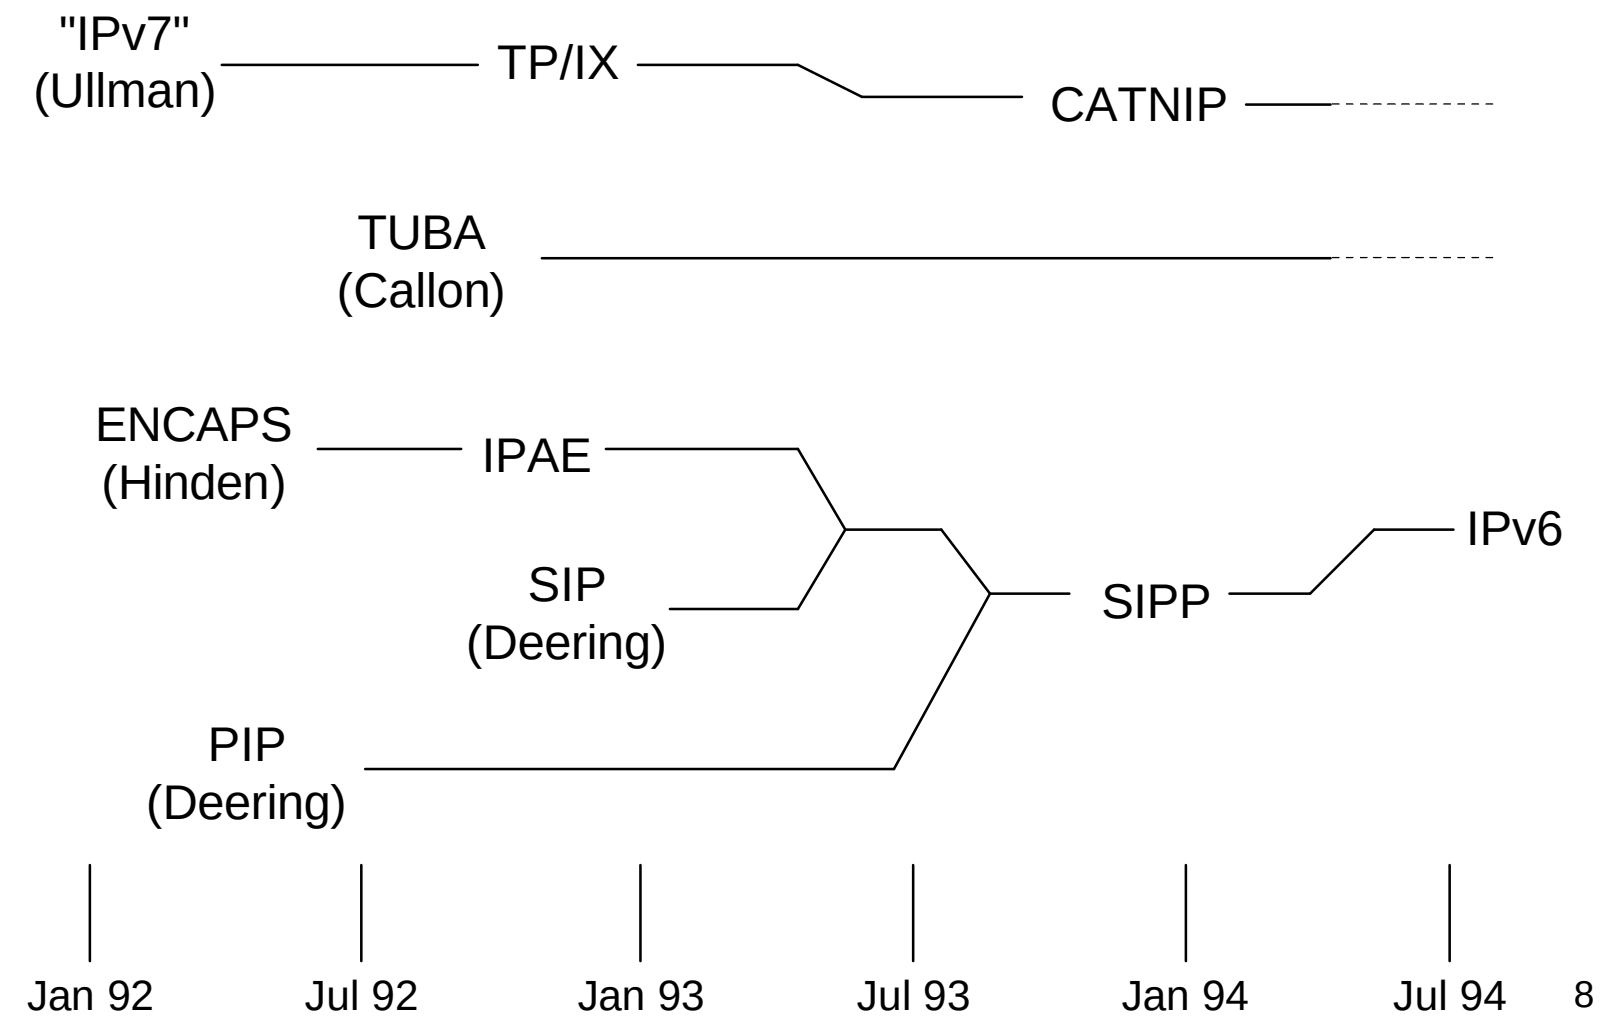

nach: S. Deering, 1997

Eigenschaften

- 128 Bit Adressen
- Vereinfachtes Headerformat
- Erweiterungs-Header
- Authentisierung und Verschlüsselung
- Autokonfiguration
- Source-Routen

- Transitionsplan:
  1. Incremental Upgrade
  2. Incremental Deployment
  3. Easy Addressing
  4. Low Start-up Costs
- Quality of Service Unterstützung

IPv4 vs. IPv6

	0	4	8	16	19	24	31
1	VERS	HLEN	SERVICE TYPE	TOTAL LENGTH			
2	IDENTIFICATION			FLAGS	FRAGMENT OFFSET		
3	TIME TO LIVE		PROTOCOL	HEADER CHECKSUM			
4	SOURCE IP ADDRESS						
5	DESTINATION IP ADDRESS						
6	IP OPTIONS (IF ANY)					PADDING	
7	DATA						
8	. . .						
9	. . .						
10	. . .						

IPv4 vs. IPv6

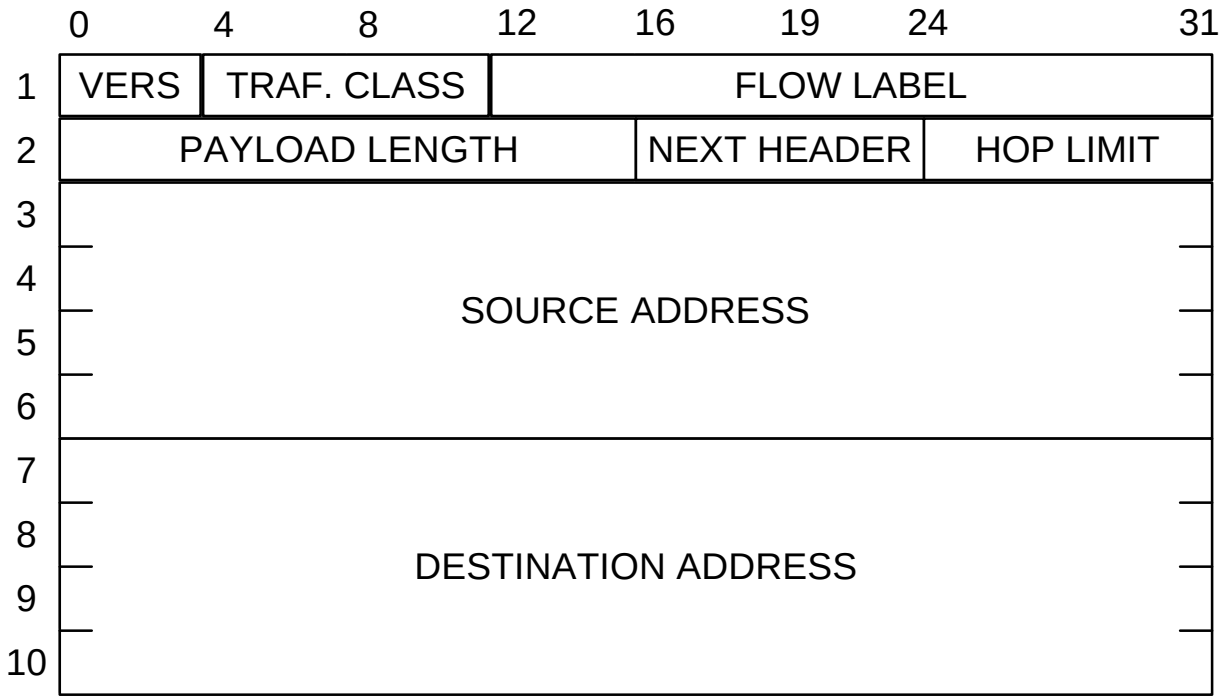

Extension Header

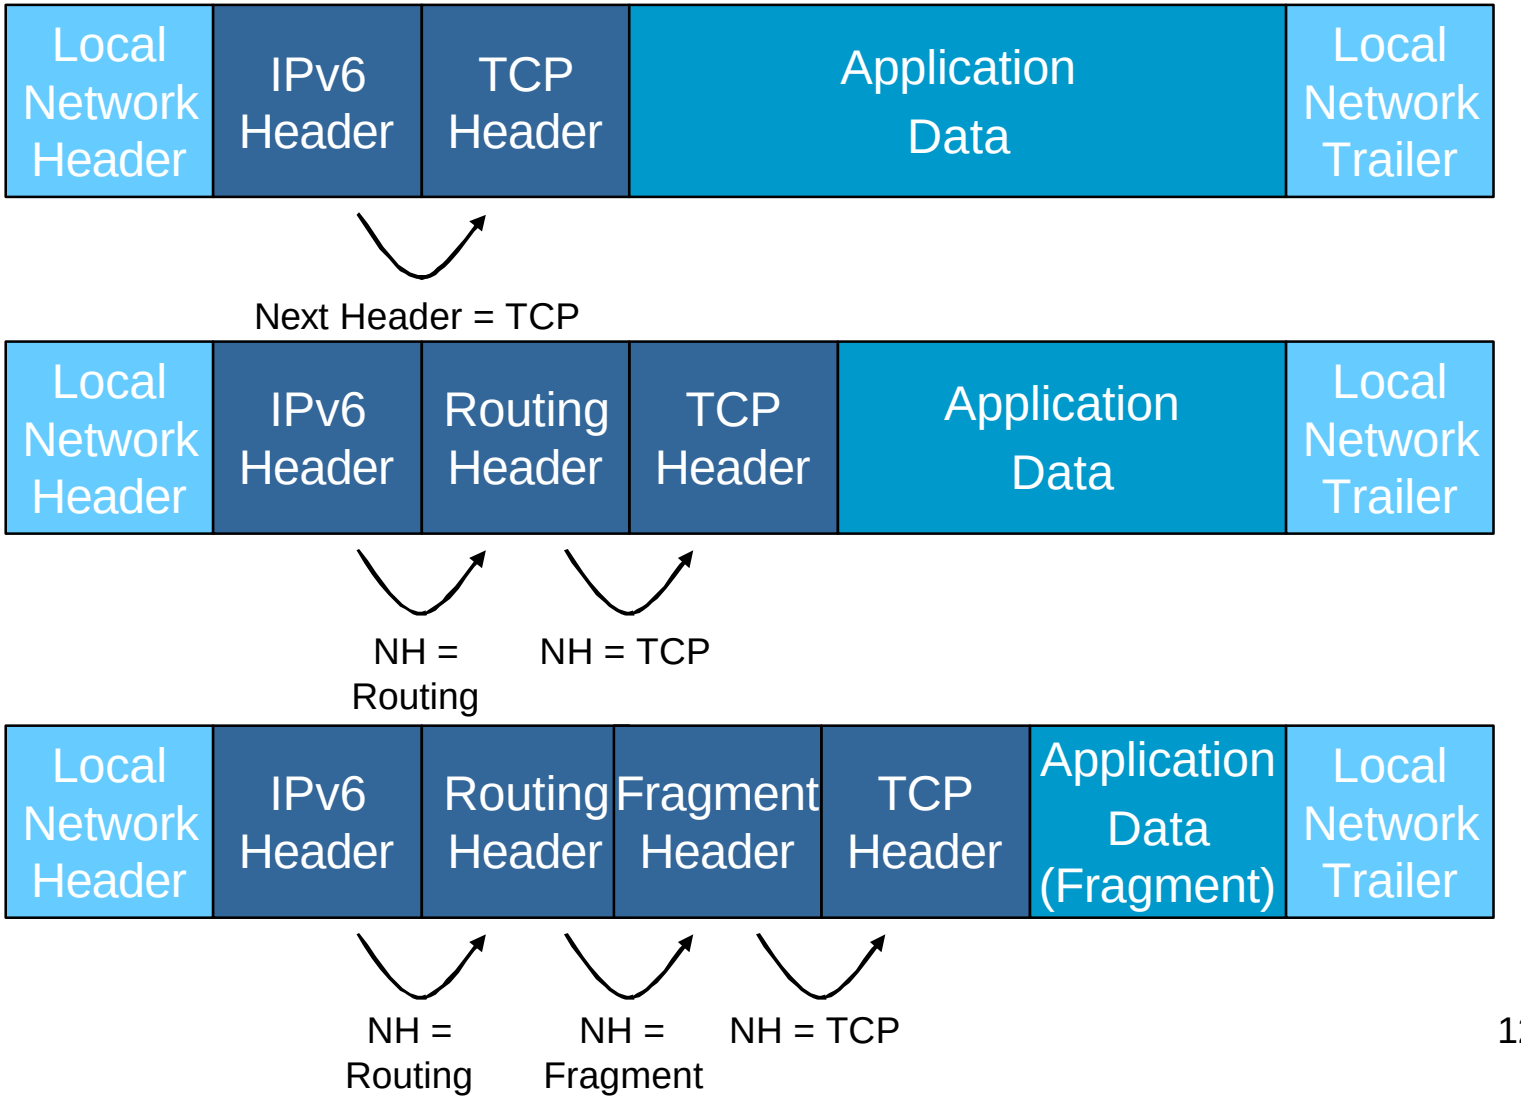

Extension Header

- Hop-By-Hop Options Headers
  - Jumbo Payload Format (64kB < Payload <= 4GB)
- Destination Options Headers
  - Pad1, PadN Options
- Routing Header
  - Routing Type0 Header (Source Routing)
- Fragment Header
- Authentication Header
- Encapsulation Security Payload Header
- No Next Header

IPv4 Adressen

IPv4 Adressen:

1	1	0	0	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

C0	A8	01	01

Dezimal:

192	168	1	1

Dotted decimal notation:

192.168.1.1

IPv4 Adressen

Class A:

Class B:

Class C:

Class D:

Class E:

Address Depletion

- Class-Based IP Adressing:
  Jedes physikalische Netzwerk bekommt eine weltweit eindeutige Netzwerk-ID.
  Physikalische Netze sind über Router verbunden.

→Verschwendung von Netzwerken

Lösungen:

Innerhalb eines LANs: Subnetting

In WANs: Classless Interdomain Routing (CIDR)

Subnetting

Reduktion der Zahl der Netzwerke:

- Transparente Router
- Proxy ARP
- Subnet Addressing

Verwendung einer Netzmaske:

- Dotted Decimal Notation:
  192.168.1.128/255.255.255.127
- Offset:
  192.168.1.128/25

Subnetting

IP Address: 201.1.1.34
Netmask: 255.255.255.240
Wie lautet die Netz-ID und Host-ID?

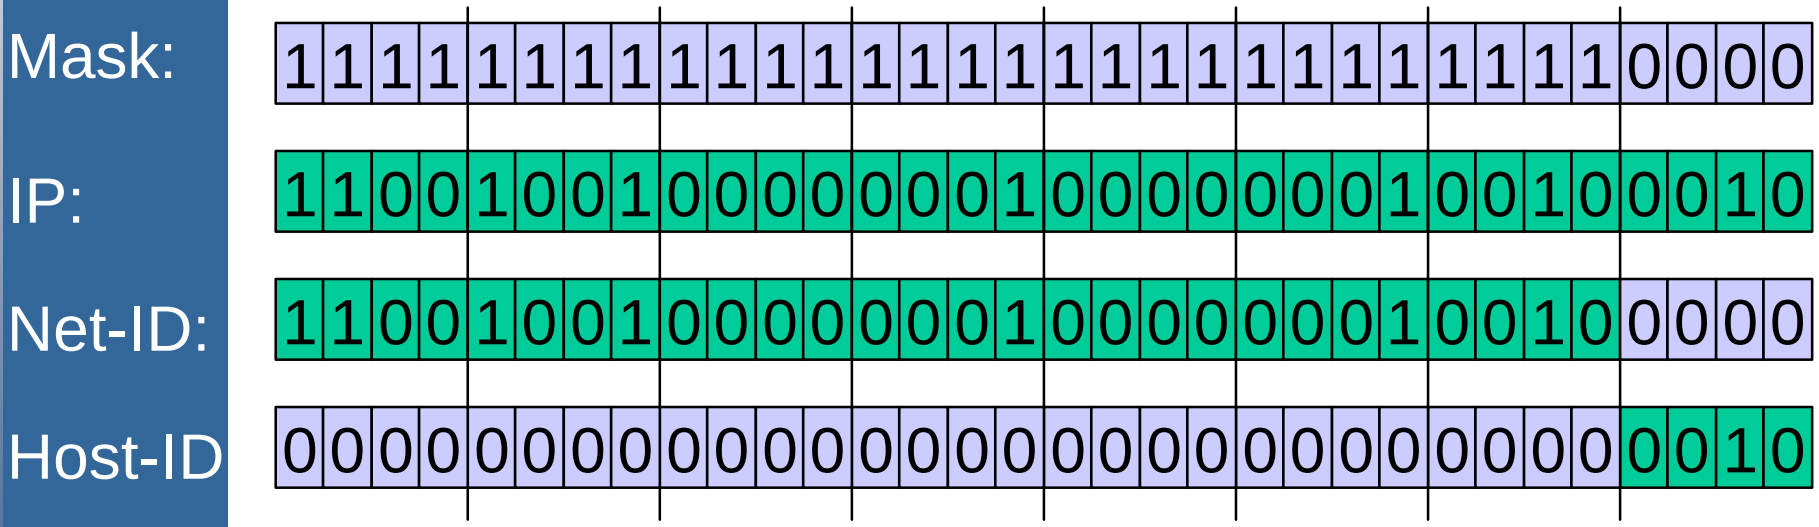

Netz-ID: 201.1.1.32 Host-ID: 0.0.0.2

Classless Interdomain Routing (CIDR)

- Statt eines "Class A" oder "Class B" Netzwerks vergibt man eine Anzahl zusammenhängender "Class C" Netzwerke.
  Problem: Explosion der Routing Tabellen
- Lösung: Route Aggregation

CIDR / Aggregation

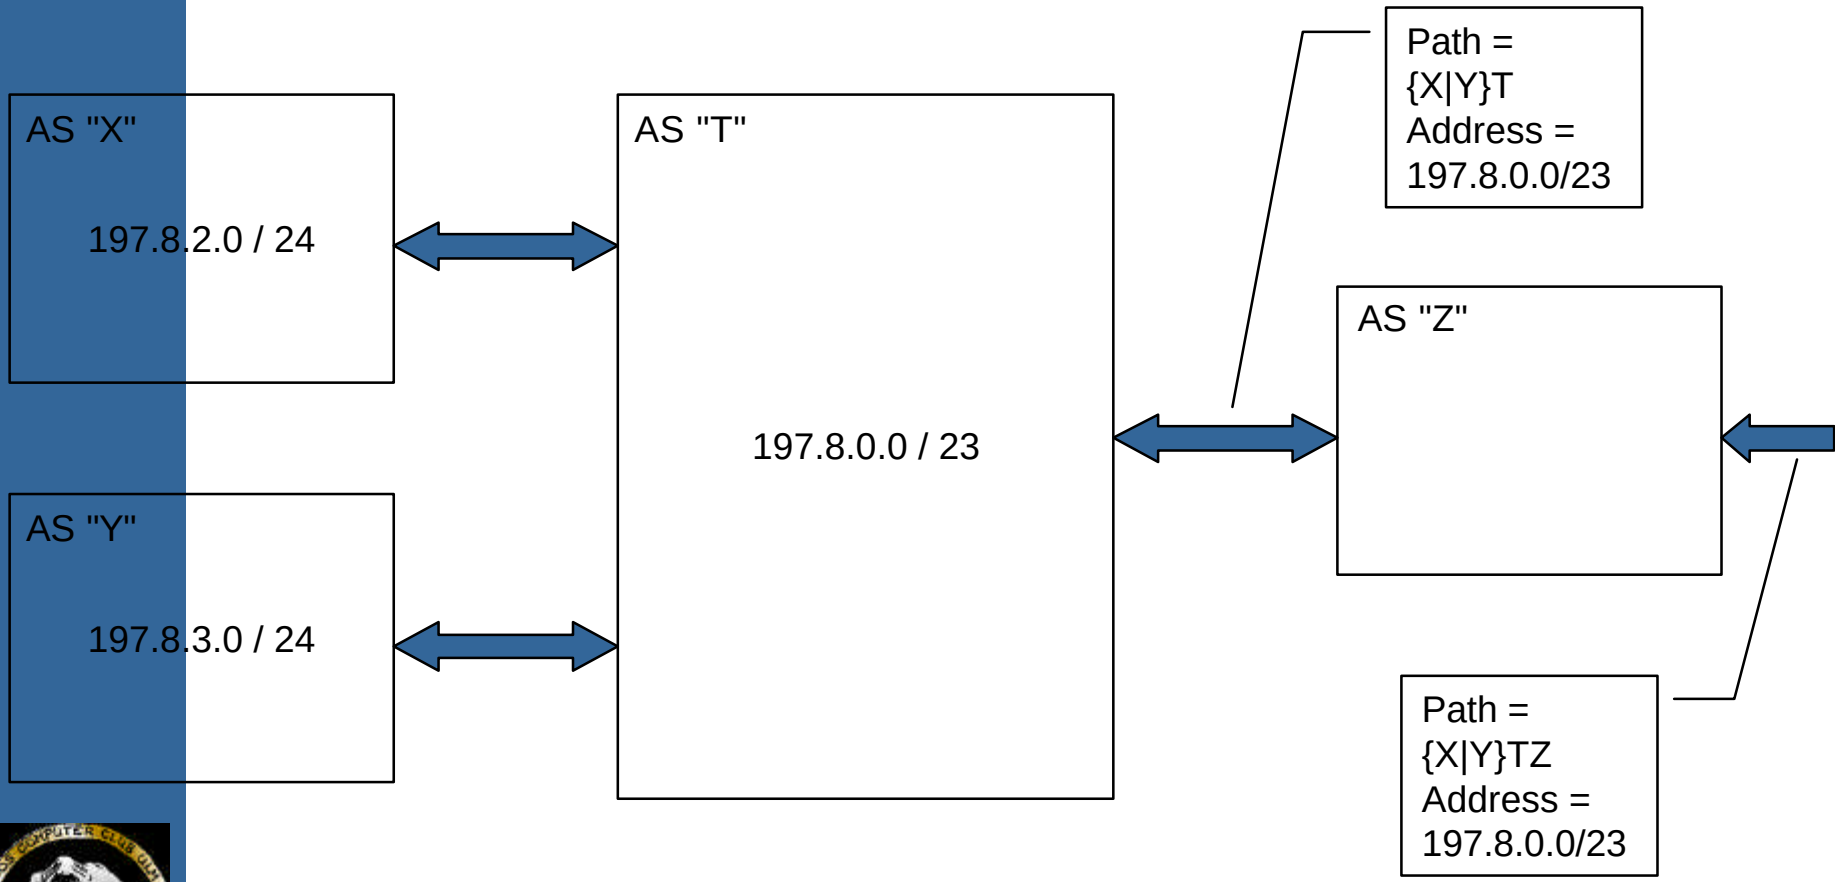

Class C Verteilung

- Multiregional: 192.0.0.0 - 193.255.255.255
- Europa: 194.0.0.0 - 195.255.255.255
- Others: 196.0.0.0 - 197.255.255.255
- North America: 198.0.0.0 - 199.255.255.255
- Central/S.America: 200.0.0.0 - 201.255.255.255
- Pacific Rim: 202.0.0.0 - 203.255.255.255
- Others: 204.0.0.0 - 205.255.255.255
- Others: 206.0.0.0 - 207.255.255.255
- (planed)
- Geographical vs. Provider-based ?
- Weitere Lösung: Network Address Translation (NAT)

Keine Lösung

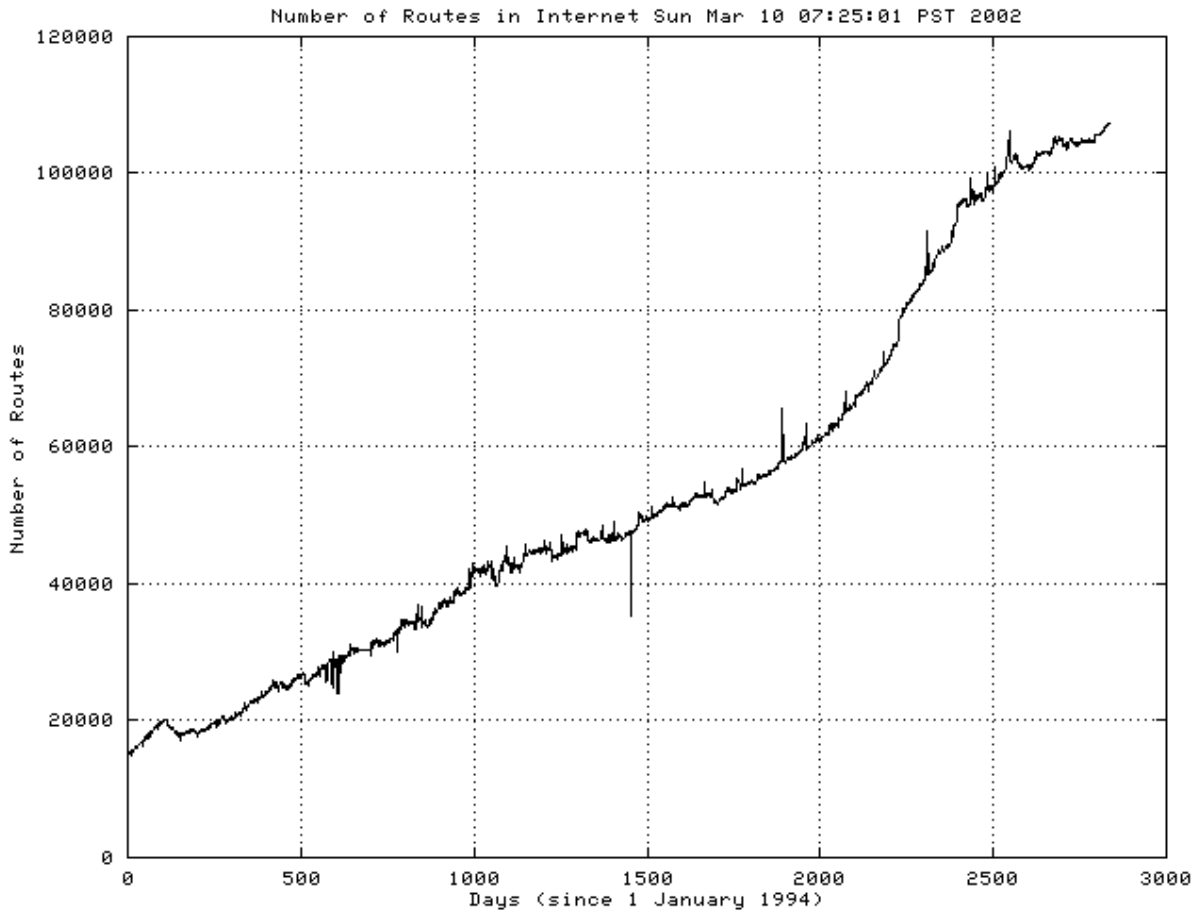

Quelle: http://www.employees.org/~tbates/cidr.hist.plot.html

IPv6 Adressformen

Unicast

Multicast

Anycast

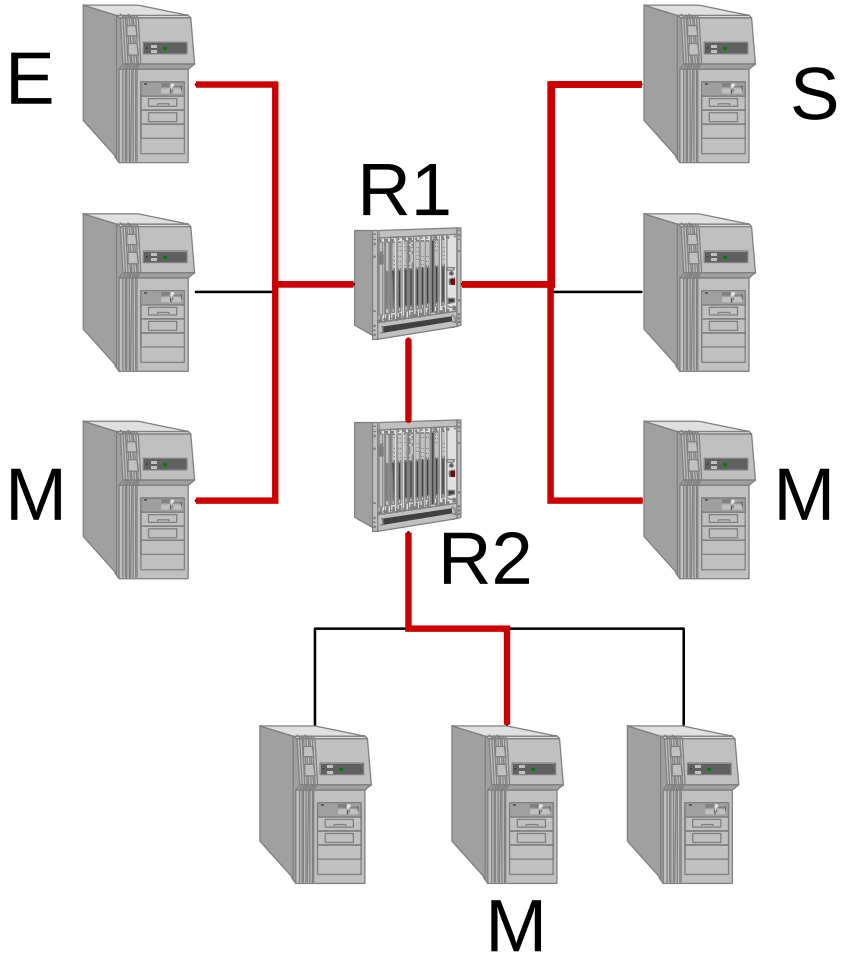

IPv6 Adressen

- Hexadezimale Schreibweise:
  `3ffe:0400:0060:004d:0250:04ff:fe44:b099`
- Ohne führende Nullen:
  `3ffe:400:60:4d:250:4ff:fe44:b099`
- Abgekürzte Schreibweise:
  `3ffe:0:0:4d:250:4ff:fe44:b099`
  `3ffe::4d:250:4ff:fe44:b099`
- Mit Präfix:
  `3ffe:400:60:4d:250:4ff:fe44:b099/64`

Adressaufteilung

IPv6 Adressaufteilung nach RFC 2373:

Prefix	Belegung
0000 0000	**Reserved**
0000 001	**Reserved for NSAP Allocation**
0000 010	**Reserved for IPX Allocation**
001	**Aggregatable Global Unicast Addresses**
1111 1110 10	**Link Local Unicast Addresses**
1111 1110 11	**Site Local Unicast Addresses**
1111 1111	**Multicast Addresses**
Rest	**Unassigned**

Unicast Adressen

- Loopback ::1
- IPv4-compatible IPv6 ::8612:7601
- IPv4-mapped IPv6 ::ffff:8612:7601
- NSAP + IPX Adressen
- Aggregatable Global Unicast Addresses (RFC 2374)
- Testing Addresses (RFC 2471)
  3ffe:400:60:4d:250:4ff:fe44:b099

Aggr. Global Unicast Address

FP: Format Prefix (001)

TLA: Top-Level Aggregation Identifier

RES: Reserved

NLA: Next-Level Aggregation Identifier

SLA: Site-Level Aggregation Identifier

IID: Interface Identifier IEEE EUI 64 Format

6bone Adressen

FP: 001

TLA: 0x1FFE

RES: 0x00

NLA: 0x0400: JOIN (pseudo-TLA)
0x0060: UNI-ULM NLA

SLA: 0x004d: "Subnetz 70"

IID: Interface Identifier IEEE EUI 64 Format
00:50:04:44:B0:99 → 0250:04ff:fe44:b099

Adresse: 3ffe:400:60:4d:250:4ff:fe44:b099/64

Local Use Adressen

10	54	64
1111111010	0	EUI-64 IID

Unicast Link-Local Adresse

fe80::250:4ff:fe44:b099

10	38	16	64
1111111011	0	Subnet	EUI-64 IID

Unicast Site-Local Adresse

fec0:7d:250:4ff:fe44:b099

Multicast Adressen

Flags:	0000	Well-Known Multicast Group
	0001	Transient Multicast Group
Scope:	0 reserved	
	1 node-local	
	2 link-local	
	5 site-local	
	8 organization-local	
	E global	
Beispiele:	ff01::1	all nodes
	ff01::2	all routers
	ff05::102	SGI Dogfight (site-local)

IPv6 Node Adressen

- Link-Local Adresse (pro Interface)
- Assigned Unicast Adresse
- Loopback Adresse
- All-Nodes Multicast Adresse
- Solicited-Node Multicast Adresse (pro Unicast/Anycast Adresse)

ICMPv6

ICMPv6 Error Messages:

- Destination Unreachable
- Packet too big
- Time exceeded
- Parameter problem

ICMPv6 Informational Messages:

- Echo Request/Reply
- Multicast Verwaltung
- Router Solicitation/Advertisement
- Neighbor Solicitation/Advertisenent
- Redirect

Neighbor Discovery Protocol

Aufgaben (RFC 2461):

- Router Discovery
- Prefix Discovery
- Parameter Discovery
- Address Autoconfiguration
- Address Resolution
- Next-Hop Determination
- Neighbor Unreachability Detection
- Duplicate Address Detection
- Redirect

Neighbor Solicitation

1. Send Neighbor Solicitation for IPv6 address; set source to link-local
2. Wait for Neighbor Advertisement containing MAC address

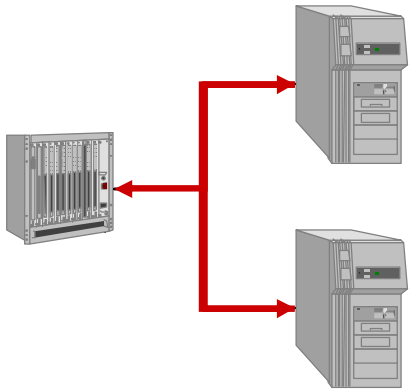

Automatische Konfiguration

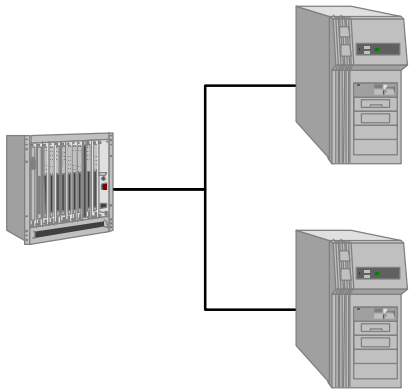

Stateless-Address-Autoconfiguration:

- Calculate Link Local Adresse
- Duplicate Detection:
  Neighbor Solicitation/
  Neighbor Advertisement
- Join All-Nodes Multicast Group ff02::1
- Send Router-Solicitation to All-Routers Multicast Group ff02::2
- Receive Router-Advertisement incl. Prefix(es), Router-Addresses
- Configure additional addresses

Alternativ:

Stateful-Address-Autoconfiguration

DNS mit IPv6

- Forward: Record Type AAAA
- Reverse: Domain ip6.int

```
ip6.hosts:
nathan  IN      AAAA    3ffe:0400:0060:004d:0250:04ff:fe44:b099
6bone   IN      AAAA    3ffe:0400:0060:004d:0250:04ff:fe44:b099
6bone   IN      A       134.60.70.253
sixbone IN      CNAME   6bone.ip6.uni-ulm.de.
scuba   IN      AAAA    3ffe:0400:0060:004d:0250:04ff:fe44:b09b
wave2   IN      AAAA    3ffe:0400:0060:004d:0290:96ff:fe07:6915

004d-ip6.rev:
9.9.0.b.4.4.e.f.f.f.4.0.0.5.2.0 IN PTR nathan.ip6.uni-ulm.de.
b.9.0.b.4.4.e.f.f.f.4.0.0.5.2.0 IN PTR scuba.ip6.uni-ulm.de.
5.1.9.6.7.0.e.f.f.f.6.9.0.9.2.0 IN PTR wave2.ip6.uni-ulm.de.
```


DNS mit IPv6

Alternativ

- Forward: Record Type A6
- Reverse: Record Type DNAME (ip6.arpa.)

```
ip6.hosts:
nathan  IN      A6      64 ::0250:04ff:fe44:b099 sub70
sub70   IN      A6      48 ::004f:0:0:0:0 uni-ulm
uni-ulm IN      A6      0 3ffe:0400:0060::

004d-ip6.rev:
$ORIGIN sub70-rev.example.com
\[x025004fffe44b099]   IN PTR nathan.ip6.uni-ulm.de.
\[x004d/16]            IN DNAME sub70-rev.ip6.uni-ulm.de.
\[x3ffe04000060/48]    IN DNAME uni-ulm-rev.ip6.uni-ulm.de.
```


Weitere Protokolle

- DHCPv6
- IPv6 Path MTU Discovery
- RIPv6
- OSPFv6
- BGP4+
- Mobile IPv6

6Bone

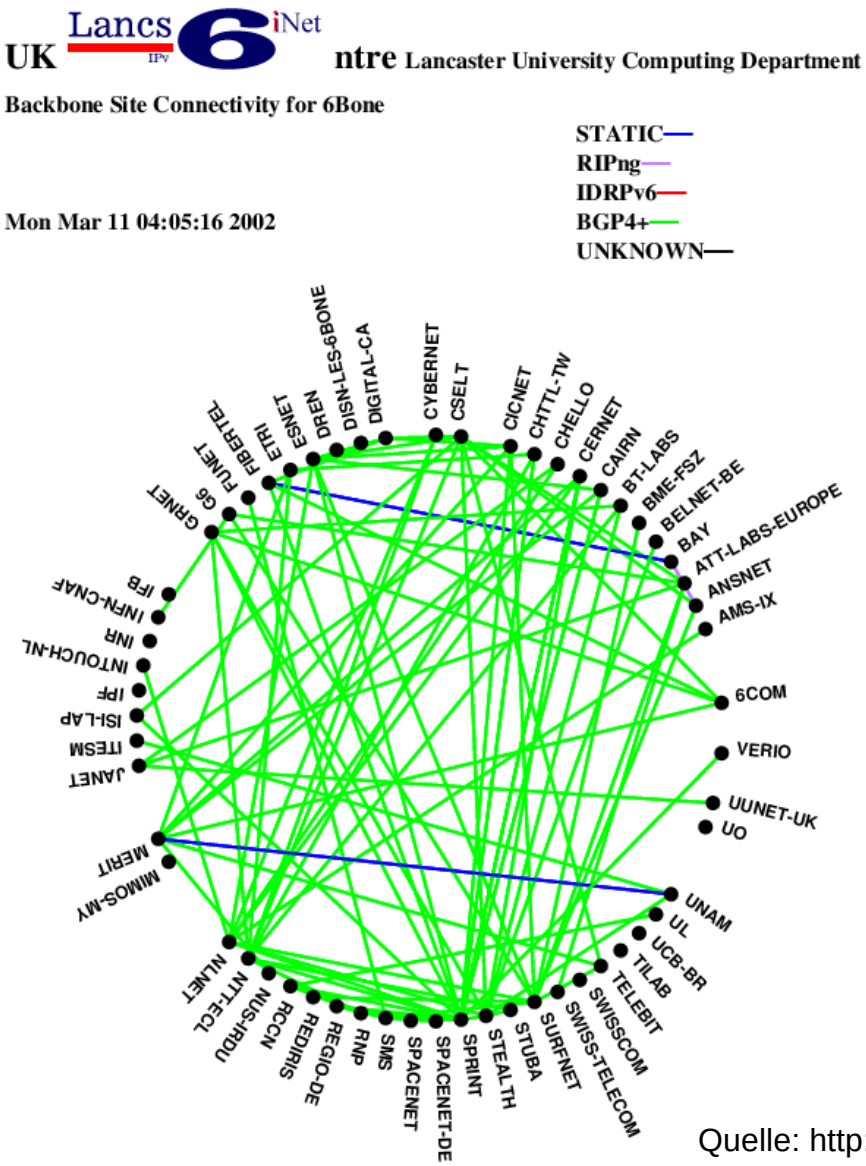

- Real existierendes IPv6 Testnetz
- Besteht vorwiegend aus 6-over-4 Tunneln
- Backbone

Quelle: http://www.cs-ipv6.lancs.ac.uk/

6Bone

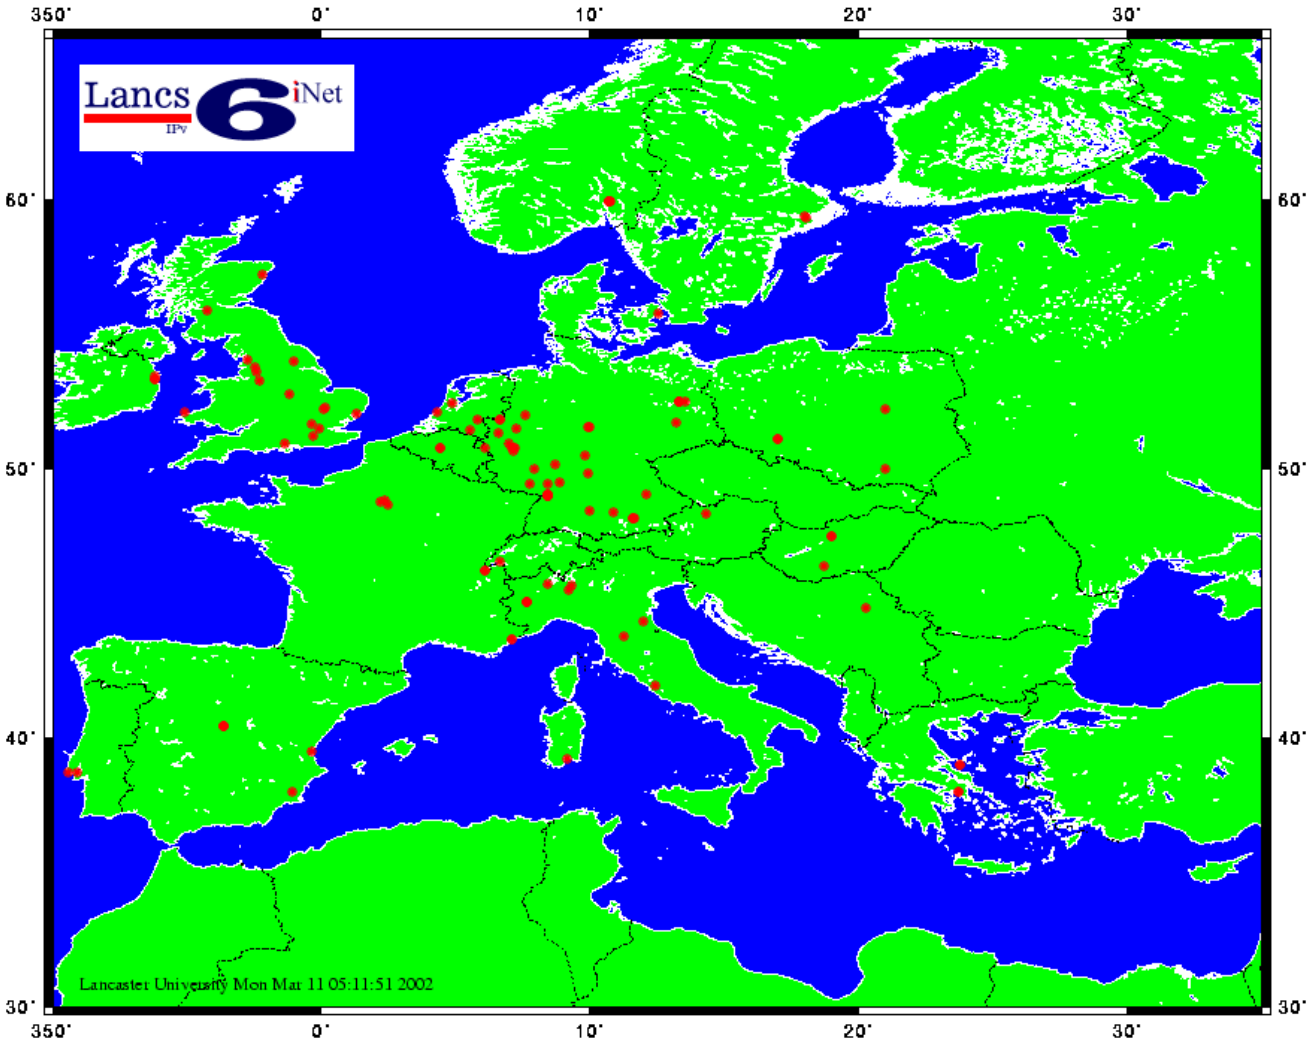

Quelle: http://www.cs-ipv6.lancs.ac.uk/

Teilnahme am 6Bone

- Informationen auf Join-Homepage:
  http://www.join.uni-muenster.de/6bone/6bone-teilnahme.html
- Join: Nur Netze mit fester IP
- Bei Einwahl mit dynamischer IP:
  Tunnelbroker Freenet, Uni Leipzig

Produktionsnetze

- APNIC, ARIN, RIPE-NCC vergeben Adressen
- Diverse Provider (WIN, Spacenet, ...) bieten ausgewählten Kunden IPv6 Dienst an
- Verschiedene Dienste laufen parallel zu IPv4 auch unter IPv6
  (WWW, IRC Uni-Erlangen, ...)

IPv6 Implementierungen

- Verfügbar für praktisch alle Betriebssysteme:
  http://www.join.uni-muenster.de/software/hersteller.html
- Bei vielen Routerherstellern im Release Code (z.B. Cisco IOS 12)
- Neuere Linux Distributionen IPv6 ready (z.B. Suse seit Version 7.0)

IPv6 unter Linux

- Demonstration am Gerät
  - ifconfig
  - route –Ainet6
  - ping6 + traceroute6
  - dig / host
  - /etc/rc.config.d + /etc/init.d/inet6
  - /etc/radvd.conf
  - ssh
  - IRC (irc.irc6.net)

IPv6 API

- BSD Socket API mit kleinen Erweiterungen

Beispiel-Programm

IPv6 - Ausblick

- Wann wird IPv6 IPv4 ablösen?
  - Niemals ;-)
  - Laufende Projekte
    - Bundeswehr
    - Telcos, insb. Mobilfunkcarrier
    - Uni Ulm ;-)
    - ...
  - Schätzungen: etwa 2005
    (wird von Zeit zu Zeit korrigiert)

Material

- Linksammlung JOIN
  http://www.join.uni-muenster.de/info/wwwlinks.html
- Bücher:
  Miller, Implementing IPv6, 2nd ed.
  Huitema, Routing in the Internet, 2nd ed.
  Stevens, Unix Network Programming, 2nd ed.

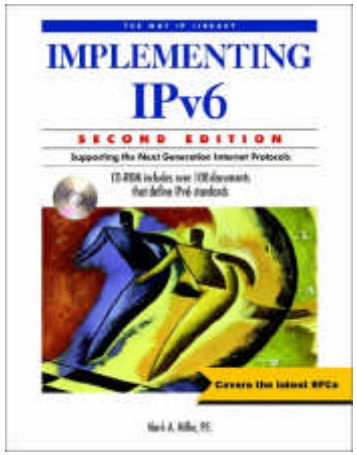

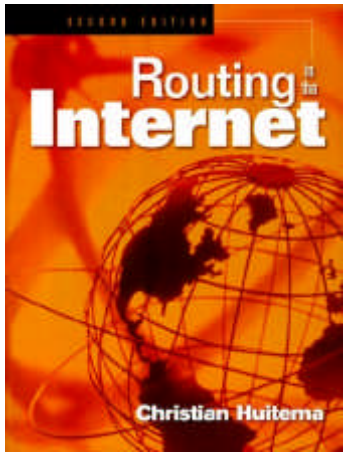

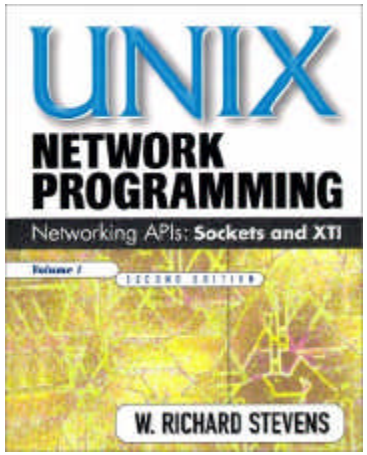

The End

Fragen, Kommentare, Anmerkungen?