IPv6 – state of development

04. 02. 2005

Wien

Internet Society German Chapter

IPv6 Task Force Germany

Why a new IP?

- We are running out of addresses (are we really?)
- size of routing tables explodes
- more security is required
- users want Plug & play-installation
- easier renumbering would be nice
- Quality of Service is needed
- Other new Applications (where?)

history

- IETF started several projects 1991
- SIP, CATNIP, PIP and TUBA
  - IPng
  - IPv6
- Address-length 64,128, 256 or variable
- IETF recommendation by
  IPng area directors in 1994 called for setup of
  IPng working group 1995
- features, features and more features
- long discussion (still ongoing)

history of IPv4

- 2004: counting more than 285 millions of hosts visible in the Internet
  - more than 8,2 millions of domains using .de
  - more than 33 millions using .com
  - the Internet is everywhere
  - IP-Addresses are requested for everything:
    - Computer
    - Machines
    - Telephones
    - Portable Devices
    - Cars
    - … … ..

Internet hosts 1995-2004

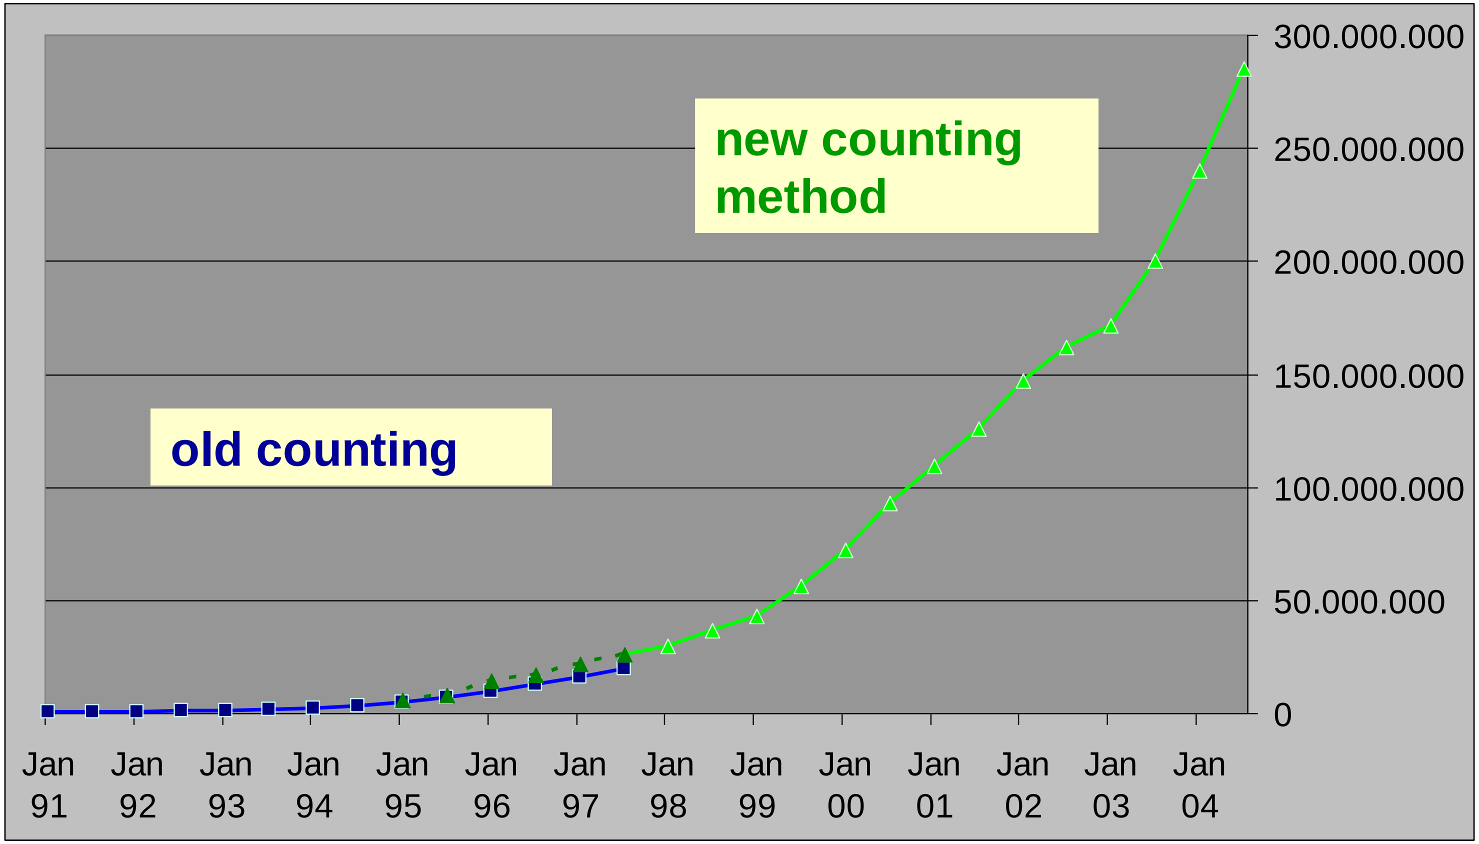

Source: ISC

usage of IPv4 addresses

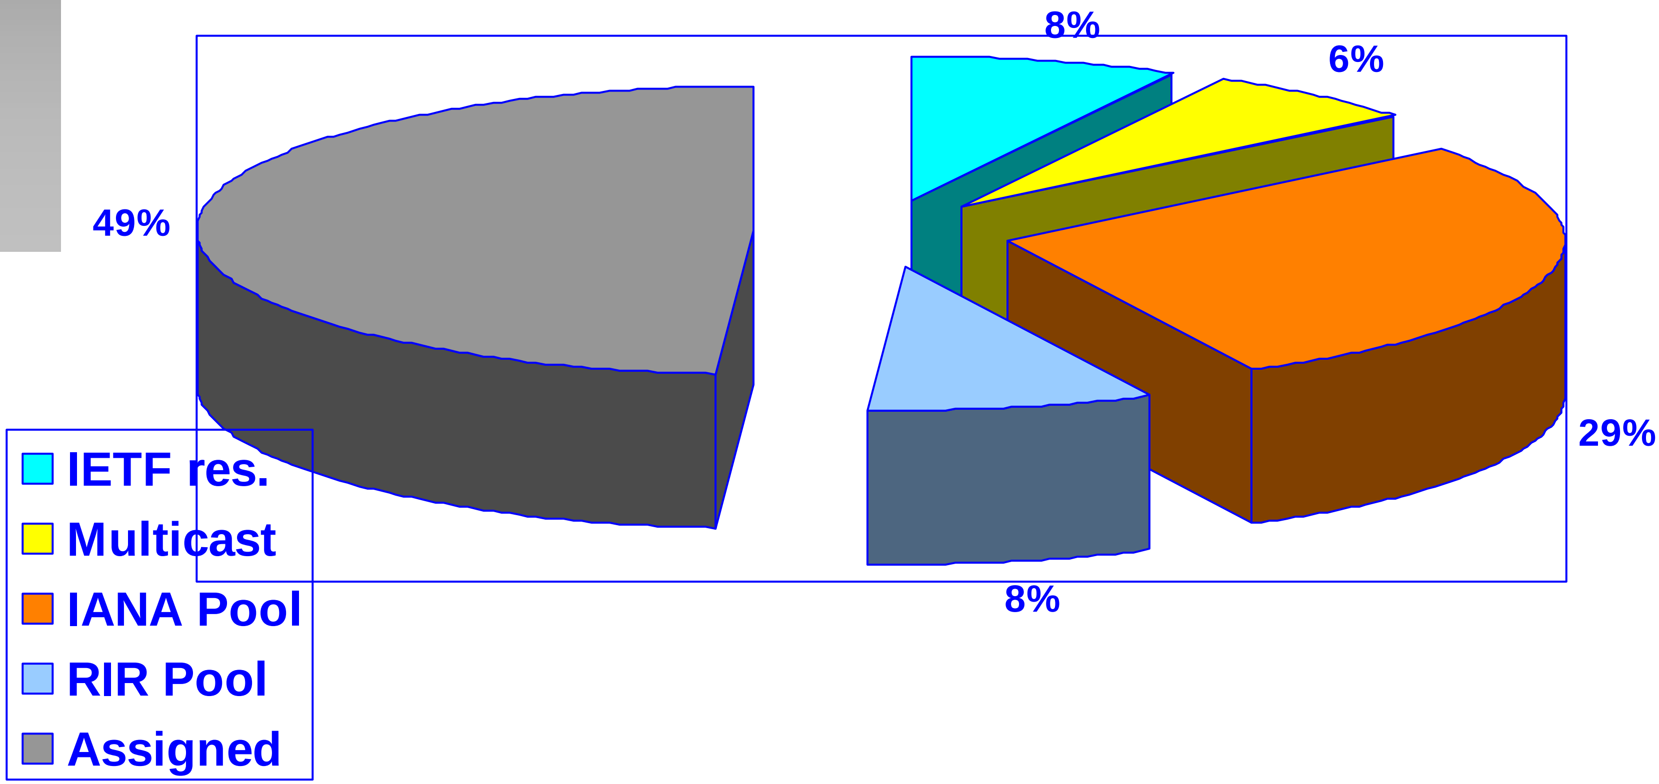

(Source of data: Geoff Houston http://bgp.potaroo.net/)

exponential IPv4 projection

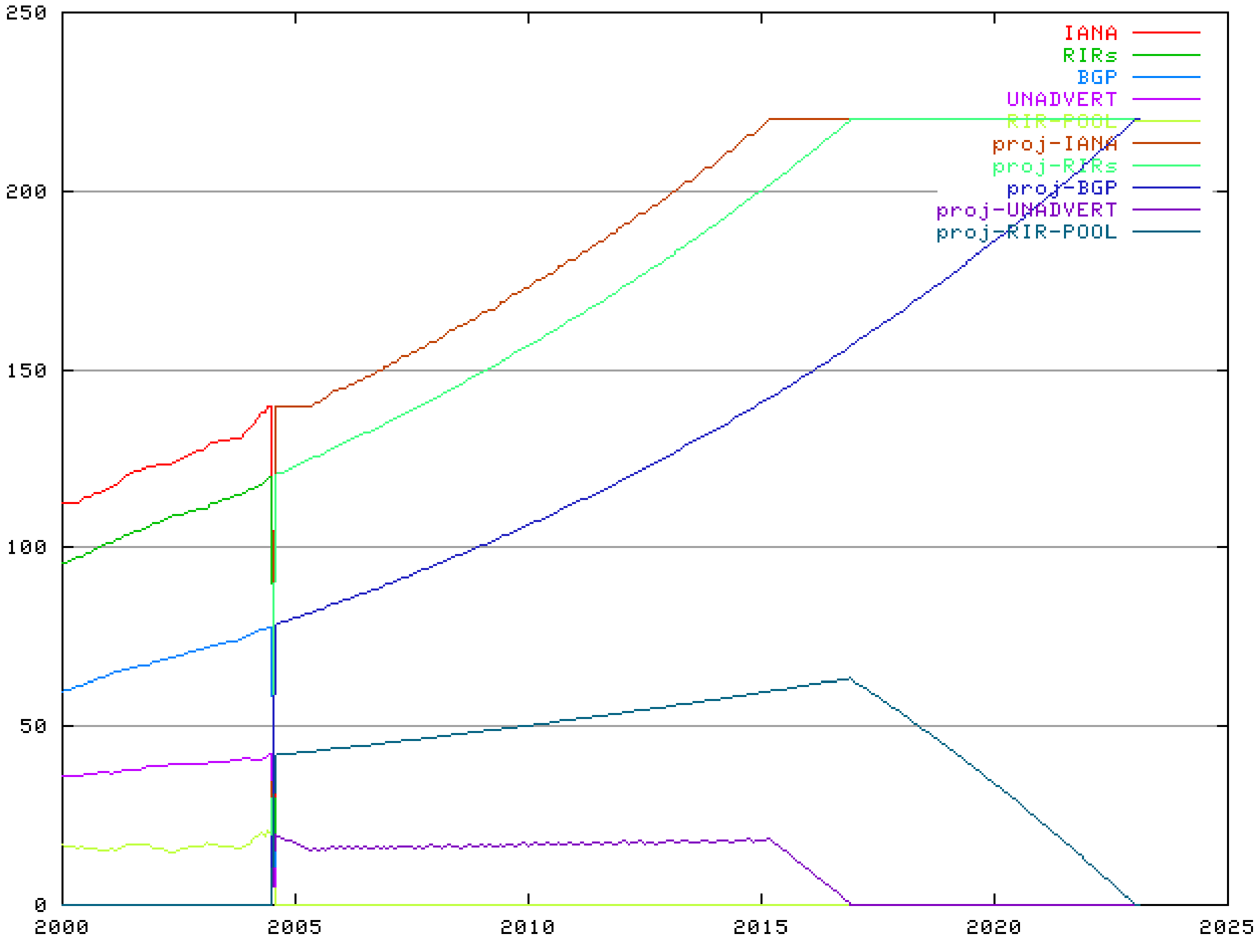

(Geoff Houston http://bgp.potaroo.net/)

increase in BGP-routes

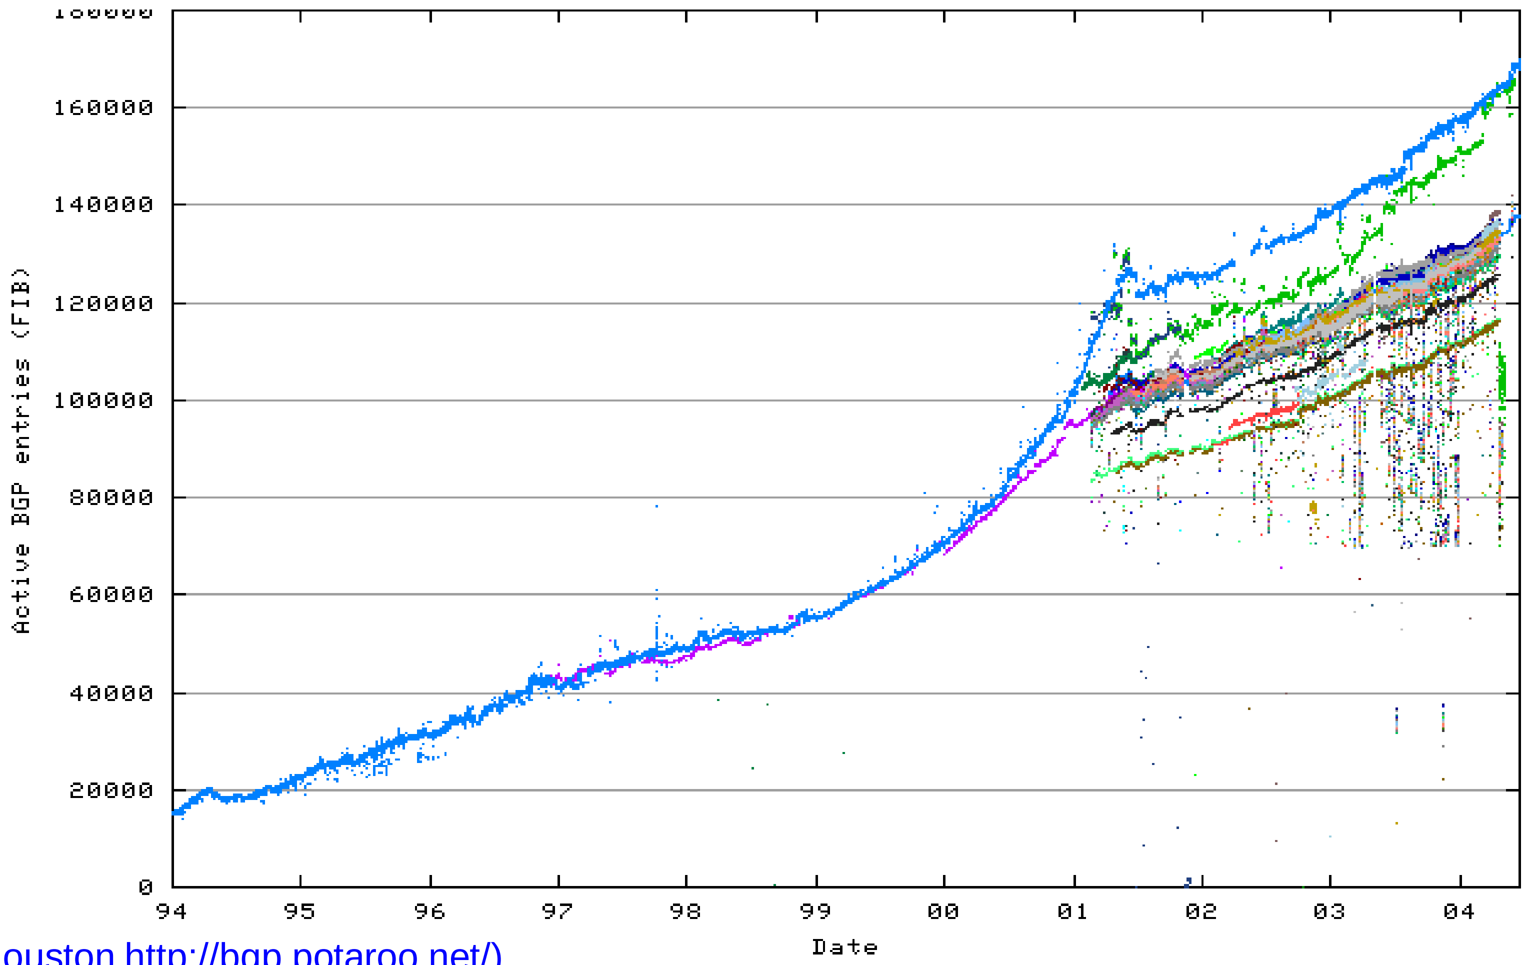

(Geoff Houston http://bgp.potaroo.net/)

types of IPv6 addresses

- Addressing modes:
  - unicast — directed to one node
    - global
    - link-local
    - ~~site-local~~ — site-locals are removed from all standards
  - anycast — to the first (nearest) node of a group sharing one prefix (used in MOBILEIP)
  - multicast — to all in a group

  - an interface has always a link-local unicast address
  - an interface has always one or more multicast addresses
  - an interface may have several global addresses

  additional hint:
  - IPv6 has no broadcast-addresses . This function from IPv4 was completely replaced by multicast

IPv6 - addresses

- Scope
  if several addresses may be in conflict (like FF02::1 all nodes on this link on a machine with several links) an additional zone identifier may be added:
  - FF01::1%1 means all nodes on all links with the manually defined zone value 1 and FF01::1%23 means all nodes in zone 23

- private addresses
  - FC00::/7 proposed solution for unique local addresses
  - 7 bit FP
  - FC00::/8 using a 40 bit centrally allocated global identifier
  - FD00::/8 using a 40 bit locally defined identifier
  - 16 bit subnet
  - 64 bit interface ID

  - still under discussion!

interface IDs

- the lower 64 bit of an IPv6-address may be defined using several methods:
  - automatically by auto-configuration computing the 64 bit of the ID from the 48 bit MAC-address as described in EUI-64
  - automatically from a DHCP-server, either using EUI-64 identifier or preset values from the network manager
  - manually defined at the node
  - generated at system-start (or triggered by time, data volume or manual trigger) using a random (pseudo-random) 64 bit value (for privacy reasons in dial-up situations)
  - other methods are up to the wild dreams of system designers

distribution of addresses

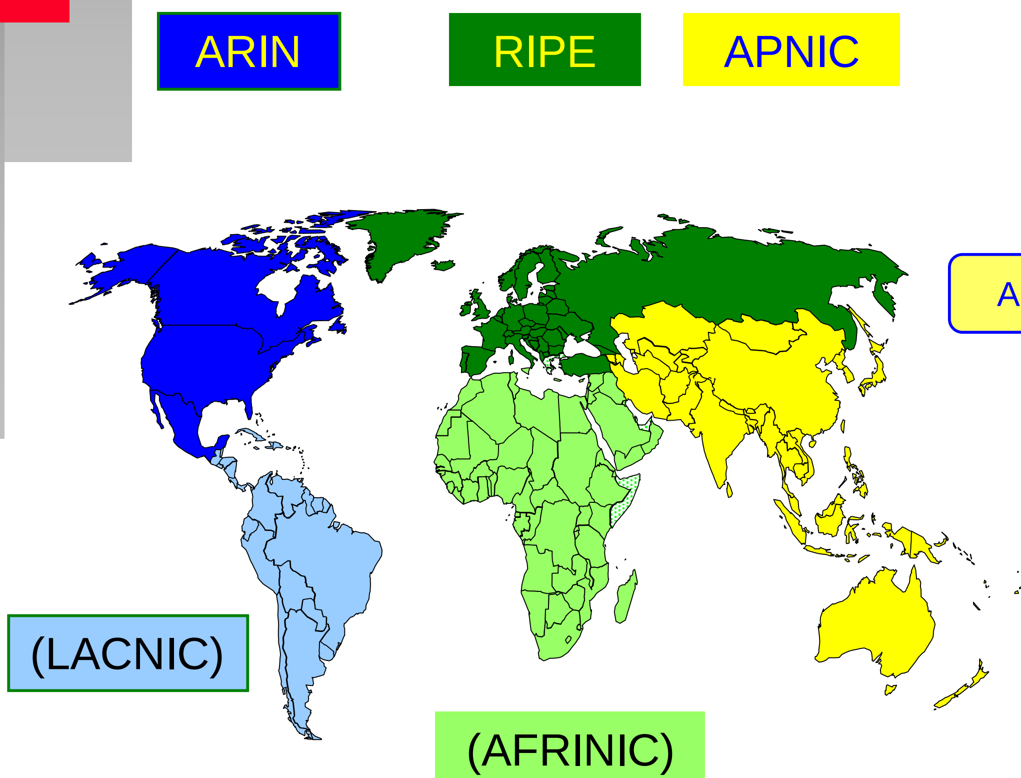

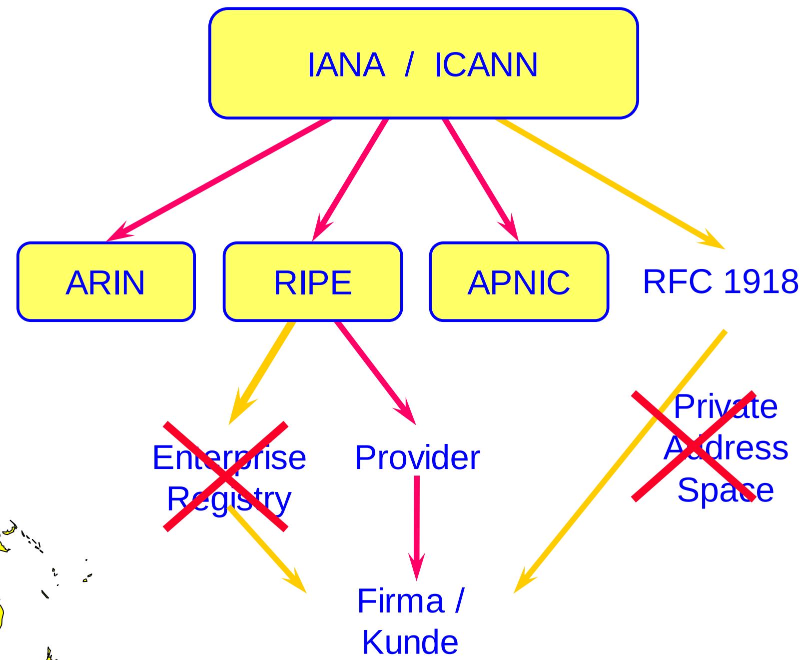

address block policy

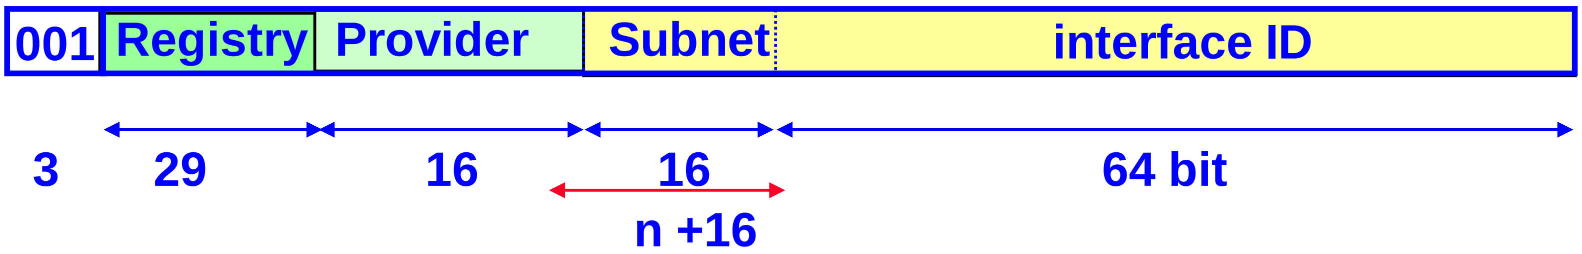

- a Local Internet Registry (LIR) (= ISP or provider) starts with a /32
  - RIPE requires a plan to attach at least 200 customers to accept you as provider
- if the number of customers increases additional neighboring /32 may be allocated
- for large LIRs larger chunks are possible
- a normal end-user gets a /48 network
- a /64 may be assigned, if only one subnet will be used by design
- a /128 may be assigned if its absolutely known that only one device may be attached at any time in future
- if needed a provider may assign several /48 to one customer resulting in a /47, /46 or bigger

start into new territory

- first usable standards 1995 - 1996
- 6bone started 1996
  - universities and research institutes
- commercials ISPs added from 1999
- AMS-IX in Amsterdam started IPv6 1999
- 9/2001 DECIX with first test
  - end of 2004:
    25 % of all DECIX connections carry also IPv6
- address-distribution from RIPE since 2000 available on firm rules

vendor support

- 6bone started with Solaris and Linux based routers mostly on tunnels
- CISCO started beta in 1997 and finally turned 2003 into production release
- Juniper started 2001 with IPv6
- most other vendors on same state
  – it is possible to use IPv6
  but still lacking optimal integration
  (ACL-implementation or ASIC-integration)

- first exploitable bug in IPv6-Software reported by US-CERT January 2005 for CISCO IOS

software support

- Solaris and NT4 started 1997 with very unstable patches or betas
- KAME and USAGI projects in Japan did a lot of implementations for BSD and Linux
- now nearly all operating systems include a more or less fully developed support for IPv6
  - Windows XP and Windows 2003
  - MAC-OS 10.2
  - AIX > 4.3, Solaris > 2.7, HP-UX > 11i, Irix > 6.5 ….
  - FreeBSD > 4.0, NetBSD > 1.5, RedHat > 7.2, SuSE > 8.1…
- still only few applications support IPv6
- still lacking support in “glue” software
- still very few support in management software

open points

- many different versions of tunneling available
- cleanup is needed
- routing still a little chaotic
  - tunnels tend to end in unforeseen places
  - BGP announcements of IPv6 connectivity is often chaotic
  - clean hierarchical peerings and route-announcements must still be developed
  - often “beta-equipment” and “alpha-operators” inject wrong routes
  - tunneling leads to wrong hop counts
    - a tunneled route with only two hops each way across the Atlantic seems to be shorter than the direct route with 6 native hops from Germany to Italy

open points

- NAT is evil
  - true --- but necessary for IPv6 deployment
  - we need PNAT to cross from IPv6 islands to existing IPv4-services
  - tunneling solves the transport problems but the access issues still exist – not everybody speaks IPv6
- better solution is dual-stack
  - available for some (but not all) central services
    - mail (smtp, pop, imap ….)
    - DNS
    - WWW
  - BIND, sendmail, postfix, qmail, exim, qpopper, apache, IIS, Proftp, ssh, telnetd and many more support IPv6
  - but here are still several missing peaces – but you have to start somewhere

problem areas

- legacy equipment
  - webcams, USV-interfaces and many other embedded systems have no IPv6
  - old routers are often not supported by new software
- legacy software on legacy systems
  - legacy printers don’t speak IPv6
- low-cost systems
  - all actual cheap SOHO-routers lack IPv6
- missing links
  - integration in large DHCP-systems
  - security in Firewalls – just starting – still beta
  - Accounting – not too much available
  - IDS – doesn’t know about IPv6
  - diagnostic tools (SLA-tools)
  - all the configuration tools (YAST, CONFIXX, PLEX …….)

selection of packet size

- the smallest MTU-value (maximum transport unit) for a given link is defined with 576 bytes
  - old value with IPv4 = 68 bytes
  - packets may be smaller, it must only be guaranteed that packets with fewer than 576 bytes can be transported without undergoing fragmentation
  - if a link (example ATM) has smaller transport units, either the link-layer or a layer between link layer and IP must do the fragmentation and reassembly hiding it from IPv6
- all IPv6 nodes must be able to do MTU-discovery to find out the MTU of a given link
- nodes which are absolutely sure that all their data fits into packets smaller than 576 bytes may skip MTU-detection
  - may be suitable for embedded devices sending only alarms
  - may be usable in voice applications
  - not optimal for normal server or client usage

cascaded MTU-discovery

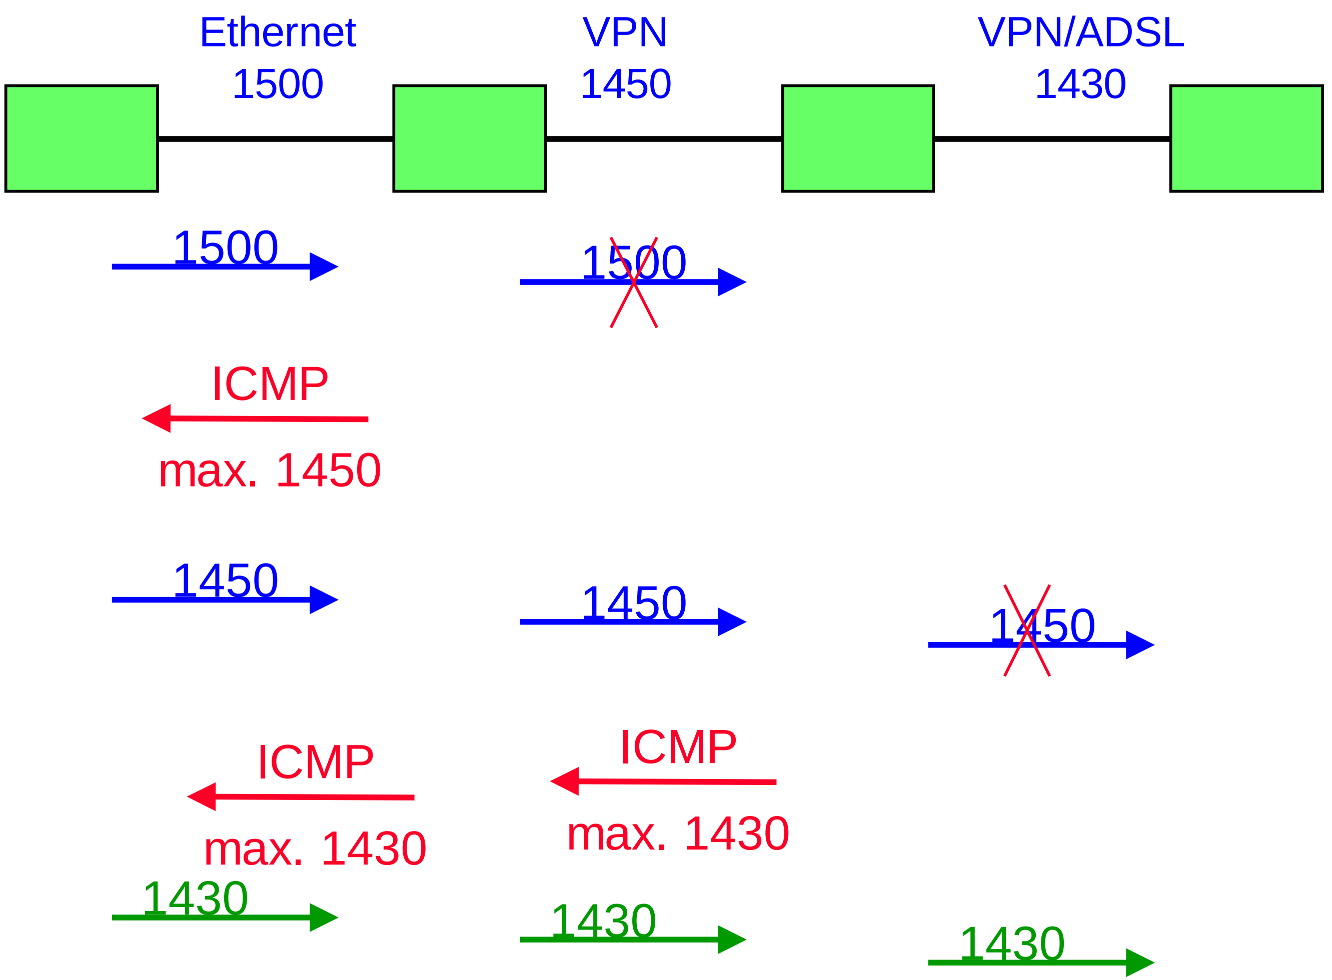

how to configure

- IPV6 gives Router-Announcements
  - what to announce
  - when to re-announce changes if routes are broken
- DHCPv6 gives DNS and other data
  - when to switch-over
  - who reigns?
- Other methods still under discussion
  - Network attachment detection
    - allows quick changes and hand-over in 3GPP
  - NETCONF
    - Mini-DHCP for DNS

selection of source-address

- A node may have several (many) addresses on every interface
- still under discussion, current proposal:
  - using the destination-address all prefixes of possible source-addresses will be checked for a longest match and this address shall be used.
  - are several candidates available, a manually configurable routing-weight should decide. Other proposals favor rules using local addresses first
- There are still drafts on the table, which define fixed ordering of priorities on address-types

multi-homing

- Multi-Homing is easy with IPv6
  - but
  - move decision-point from router to end-node
  - how to detect right path
  - when to switch-over
  - how to test and select right address
- possible solutions in HIP or MULTI6
  - new address-independent identifier
  - inserting shim-layer between application and IP
  - adding persistent connections

multicast

- Multicast has a much more central role with IPv6
- Multicast (and Anycast) in LAN-environments replace Broadcasts used by IPv4
- Multicast over WAN or the global Internet are nearly identical from IPv4 to IPv6:
  - Multicast-protocols were adapted (PIM)
  - additional ICMPv6-protocol elements are used for better control of multicast-groups

routing

- All existing routing-protocols must be adopted for IPv6
  - RIP - done
  - BGP - done
  - OSPF - done

  - new versions carry longer address-fields
  - implementation and availability still varies from manufacturer to manufacturer

DNS and IPv6

- IPv6 in DNS is a simple extension to IPv4-rules:

myv4host	IN	A	1.2.3.4
myv6host	IN	AAAA	4321:0:1:2:3:4:56:789a
mypc	IN	AAAA	4321::4567:89ab

second DNS-format

- proposed version using recursion

my6host	IN	AAAA	4321:0:1:2:3:4:56:789a
➡ -			
mynet	IN	A6	4321:0:1:2::
my6host	IN	A6	::3:4:56:789a 64 mynet
mypc2	IN	A6	::1:9:7654:fedc 64 mynet

This solution is now depreciated and removed from the standards track because of implementation problems and load problems caused by complex recursion especially in caching environments.

reverse DNS

- To find names from addresses the IPv4-PTR was only extended:
  - who has: 4321::12:3:4:567:89ab

    DNS entry:

b.a.9.8.7.6.5.0.4.0.0.0.3.0.0.0.2.1.0.0.0.0.0.0.0.0.0.0.1.2.3.4.ip6.arpa IN PTR my6host

- Hint:
  at the start of the standardization (still found in some machines) the DNS-tree .int instead of .arpa was proposed for this functionality

problems with DNS

- DNS is a hierarchical tree
  - not all combinations of IPv4-servers and IPv4-forwarders are working seamlessly with IPv6-based resolvers and clients
  - a standard for gateways is still under discussion
- ICANN has approved IPv6-Service for DNS-root-servers
  - this has not yet been fully rolled out to all servers
  - not all TLD-servers are running IPv6-DNS
  - not all registries are accepting IPv6-addresses
- Software Registry and Registrars
  - not too much support in existing tools and WEB-interfaces
  - all internal workflow must be adapted and extended
  - not all registrars and ISPs are supportive for IPv6

conclusions

- ☺ standards – enough
- ☹ standards – not yet completely stable
- ☺ hardware – available
- ☺ network software – available
- 😐 application software – some
- ☹ management software – few
- ☹ operational software – very few
- ☺ DNS – supported in TLDs
- ☹ new applications – not yet
- ☹ user interest – not enough

current RFC-list

- RFC 1887 An Architecture for IPv6 Unicast Address Allocation
- RFC 1883 Internet Protocol, Version 6 (IPv6) Specification obsoleted by RFC 2460
- RFC 1884 IP Version 6 Addressing Architecture obsoleted by RFC 2373
- RFC 1885 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) obsoleted by RFC 2463
- RFC 1886 DNS Extensions to support IP version 6
- RFC 1888 OSI NSAPs and IPv6
- RFC 1897 IPv6 Testing Address Allocation obsoleted by RFC 2471
- RFC 1970 Neighbor Discovery for IP Version 6 (IPv6) obsoleted by RFC 2461
- RFC 1972 A Method for the Transmission of IPv6 Packets over Ethernet Networks obsoleted by RFC 2464
- RFC 1981 Path MTU Discovery for IP version 6
- RFC 2019 Transmission of IPv6 Packets Over FDDI obsoleted by RFC 2467
- RFC 2023 IP Version 6 over PPP obsoleted by RFC 2472
- RFC 2073 An IPv6 Provider-Based Unicast Address Format obsoleted by RFC 2374
- RFC 2133 Basic Socket Interface Extensions for IPv6 obsoleted by RFC 2553
- RFC 2147 TCP and UDP over IPv6 Jumbograms obsoleted by RFC 2675
- RFC 2292 Advanced Sockets API for IPv6 obsoleted by RFC 3542
- RFC 2373 IP Version 6 Addressing Architecture obsoleted by RFC 3513
- RFC 2374 An IPv6 Aggregatable Global Unicast Address Format obsoleted by RFC 3587
- RFC 2375 IPv6 Multicast Address Assignments
- RFC 2450 Proposed TLA and NLA Assignment Rules

current RFC-list

- RFC 2452 IP Version 6 Management Information Base for the Transmission Control Protocol
- RFC 2454 IP Version 6 Management Information Base for the User Datagram Protocol
- RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
- RFC 2461 Neighbor Discovery for IP Version 6 (IPv6)
- RFC 2462 IPv6 Stateless Address Autoconfiguration
- RFC 2463 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
- RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
- RFC 2465 Management Information Base for IP Version 6: Textual Conventions and General Group
- RFC 2466 Management Information Base for IP Version 6: ICMPv6 Group
- RFC 2467 Transmission of IPv6 Packets over FDDI Networks
- RFC 2470 Transmission of IPv6 Packets over Token Ring Networks
- RFC 2471 IPv6 Testing Address Allocation obsoleted by RFC 3701
- RFC 2472 IP Version 6 over PPP
- RFC 2473 Generic Packet Tunneling in IPv6 Specification
- RFC 2491 IPv6 over Non-Broadcast Multiple Access (NBMA) networks
- RFC 2492 IPv6 over ATM Networks
- RFC 2497 Transmission of IPv6 Packets over ARCnet Networks
- RFC 2507 IP Header Compression
- RFC 2526 Reserved IPv6 Subnet Anycast Addresses
- RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels

current RFC-list

- RFC 2545 Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
- RFC 2546 6Bone Routing Practice
- RFC 2553 Basic Socket Interface Extensions for IPv6 obsoleted by RFC 3493
- RFC 2590 Transmission of IPv6 Packets over Frame Relay Networks Specification
- RFC 2675 IPv6 Jumbograms RFC 2710 Multicast Listener Discovery (MLD) for IPv6
- RFC 2711 IPv6 Router Alert Option
- RFC 2732 Format for Literal IPv6 Addresses in URL's
- RFC 2740 OSPF for IPv6
- RFC 2765 Stateless IP/ICMP Translation Algorithm (SIIT)
- RFC 2766 Network Address Translation - Protocol Translation (NAT-PT)
- RFC 2767 Dual Stack Hosts using the Bump-In-the-Stack Technique (BIS)
- RFC 2874 DNS Extensions to Support IPv6 Address Aggregation and Renumbering
- RFC 2893 Transition Mechanisms for IPv6 Hosts and Routers
- RFC 2894 Router Renumbering for IPv6
- RFC 2921 6BONE pTLA and pNLA Formats (pTLA)
- RFC 2928 Initial IPv6 Sub-TLA ID Assignments
- RFC 3019 IP Version 6 Management Information Base for the Multicast Listener Discovery Protocol
- RFC 3041 Privacy Extensions for Stateless Address Autoconfiguration in IPv6
- RFC 3053 IPv6 Tunnel Broker
- RFC 3056 Connection of IPv6 Domains via IPv4 Clouds
- RFC 3089 A SOCKS-based IPv6/IPv4 Gateway Mechanism

current RFC-list

■	RFC 3111	Service Location Protocol Modifications for IPv6
■	RFC 3122	Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification
■	RFC 3142	An IPv6-to-IPv4 Transport Relay Translator
■	RFC 3146	Transmission of IPv6 Packets over IEEE 1394 Networks
■	RFC 3162	RADIUS and IPv6
■	RFC 3175	Aggregation of RSVP for IPv4 and IPv6 Reservations
■	RFC 3177	IAB/IESG Recommendations on IPv6 Address Allocations to Sites
■	RFC 3178	IPv6 multihoming support at site exit routers
■	RFC 3226	DNSSEC and IPv6 A6 aware server/resolver message size requirements
■	RFC 3266	Support for IPv6 in Session Description Protocol (SDP)
■	RFC 3306	Unicast-Prefix-based IPv6 Multicast Addresses
■	RFC 3307	Allocation Guidelines for IPv6 Multicast Addresses
■	RFC 3314	Recommendations for IPv6 in 3GPP Standards
■	RFC 3315	Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
■	RFC 3316	Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts
■	RFC 3363	Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)
■	RFC 3364	Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)
■	RFC 3484	Default Address Selection for Internet Protocol version 6 (IPv6)

current RFC-list

RFC 3493	Basic Socket Interface Extensions for IPv6
RFC 3513	IP Version 6 Addressing Architecture
RFC 3531	A Flexible Method for Managing the Assignment of Bites of an IPv6 Address Block
RFC 3542	Advanced Sockets Application Program Interface (API) for IPv6
RFC 3572	Internet Protocol Version 6 over MAPOS (Multiple Access Protocol Over SONET/SDH)
RFC 3574	Transition Scenarios for 3GPP Networks
RFC 3582	Goals for IPv6 Site-Multihoming Architectures
RFC 3587	IPv6 Global Unicast Address Format
RFC 3590	Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
RFC 3595	Textual Conventions for IPv6 Flow Label
RFC 3633	IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6
RFC 3646	DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
RFC 3697	IPv6 Flow Label Specification
RFC 3701	6bone (IPv6 Testing Address Allocation) Phaseout
RFC 3736	Stateless Dynamic Host Configuration Protocol (DHCP) Service for Pv6
RFC 3750	Unmanaged Networks IPv6 Transition Scenarios
RFC 3756	IPv6 Neighbor Discovery (ND) Trust Models and Threats

current RFC-list

■	RFC 3769	Requirements for IPv6 Prefix Delegation
■	RFC 3775	Mobility Support in IPv6
■	RFC 3776	Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents
■	RFC 3810	Multicast Listener Discovery Version 2 (MLDv2) for IPv6
■	RFC 3831	Transmission of IPv6 Packets over Fibre Channel
■	RFC 3849	IPv6 Address Prefix Reserved for Documentation
■	RFC 3879	Deprecating Site Local Addresses
■	RFC 3898	Network Information Service (NIS) Configuration Options for Dynamic Host Con-figuration Protocol for IPv6 (DHCPv6)
■	RFC 3901	DNS IPv6 Transport Operational Guidelines
■	RFC 3904	Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks
■	RFC 3919	Remote Network Monitoring (RMON) Protocol Identifiers for IPv6 and Multi Pro-tocol Label Switching (MPLS)
■	RFC 3956	Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address
■	RFC3974	SMTP Operational Experience in Mixed IPv4/v6 Environments

sources for RFCs and Drafts

- www.ietf.org
- www.ripe.net
- playground.sun.com

URLs for IPv6

- playground.sun.com
- www.ipv6forum.com
- www.join.uni-muenster.de
- www.ipv6-net.de
- www.6bone.net
- www.kame.net
- www.ipv6-tf.de

Dipl. Inform. Hans Peter Dittler

- 72 - 77 Computer Science at University of Karlsruhe
- 77 - 79 Research Fellow at University of Karlsruhe
- 80 - 89 Engineer for Data Communication Products at Conware Computer Consulting
- 90 - 94 CEO of Conware
- 95 - 96 BRAINTEC Consultant
- since 97 CEO of BRAINTEC Netzwerk-Consulting GmbH
- since 86 active in various groups at IEEE and IETF
- Vice-Chair of ISOC.DE
- member of German IPv6-Taskforce

BRAINTEC Netzwerk-Consulting GmbH

Hans Peter Dittler

www.braintec-consult.de
hpdittler@braintec-consult.de
dittler@isoc.de

Herstellerunabhängige Beratung für
Vernetzung und Kommunikation
Karlsruhe