CIDR

The Life Belt of the Internet

Early IP Addressings

- Before 1981 only class A addresses were used
  - Original Internet addresses comprised 32 bits (8 bit net-id = 256 networks)
- In 1981 RFC 790 (IP) was finished and classes were introduced
  - 7 bit class A networks
  - 14 bits class B networks
  - 21 bits class C networks

IP is an old protocol which was born with several design flaws. Of course this happened basically because IP was originally not supposed to run over the whole world.

The classful addressing scheme led to a big waste of the 32 bit address space.

A short address design history:

1980	Classful Addressing	RFC 791
1985	Subnetting	RFC 950
1987	VLSM	RFC 1009
1993	CIDR	RFC 1517 - 1520

Address Classes

- From 1981-1993 the Internet was Classful (!)
- Early 80s: Jon Postel volunteered to maintain assigned network addresses
  - Paper notebook
- Internet Registry (IR) became part of IANA
- Postel passed his task to SRI International
  - Menlo Park, California
  - Called Network Information Center (NIC)

Until 1993 the Internet used classful routing. All organizations were assigned either class A, B, or C network numbers. In the early 1980s, one of the inventors of the Internet, Jon Postel, volunteered to maintain all assigned network addresses—simply using a paper notebook!

Later the Internet Registry (IR) became part of the IANA and Jon Postel's task was passed to the Network Information Center, which is represented by SRI International.

FYI: See http://www.iana.org

Classful – Drawbacks

- "Three sizes *don't* fit all" !!!
  - Demand to assign as little as possible
  - Demand for aggregation as many as possible
- Assigning a whole network number
  - Reduces routing table size
  - But wastes address space

Class B supports 65534 host addresses, while class C supports 254... But typical organizations require 300-1000 !!!

Using the full classes of the addresses it was difficult to match all needs.

Subnetting

- Subnetting introduced in 1984
  - Net + Subnet (=another level)
  - RFC 791
  - Initially only statically configured
- Classes A, B, C still used for global routing !
  - Destination Net might be subnetted
  - Smaller routing tables

By introduction of subnetting (RFC 791) a network number could be divided into several subnets. Thus large organizations who needed multiple network numbers are assigned a single network number which is further subnetted by themselves. This way, subnetting greatly reduced the Internet routing table sizes and saved the total IP address space.

Routing Table Growth (88-92)

MM/YY	ROUTES ADVERTISED	MM/YY	ROUTES ADVERTISED
Feb-92	4775	Apr-90	1525
Jan-92	4526	Mar-90	1038
Dec-91	4305	Feb-90	997
Nov-91	3751	Jan-90	927
Oct-91	3556	Dec-89	897
Sep-91	3389	Nov-89	837
Aug-91	3258	Oct-89	809
Jul-91	3086	Sep-89	745
Jun-91	2982	Aug-89	650
May-91	2763	Jul-89	603
Apr-91	2622	Jun-89	564
Mar-91	2501	May-89	516
Feb-91	2417	Apr-89	467
Jan-91	2338	Mar-89	410
Dec-90	2190	Feb-89	384
Nov-90	2125	Jan-89	346
Oct-90	2063	Dec-88	334
Sep-90	1988	Nov-88	313
Aug-90	1894	Oct-88	291
Jul-90	1727	Sep-88	244
Jun-90	1639	Aug-88	217
May-90	1580	Jul-88	173

Growth in routing table size, total numbers
Source for the routing table size data is MERIT

The list above shows the growth of the routing tables from 1988 until 1992 in total numbers.

Network Number Statistics, April 1992

	Total	Allocated	Allocated %
Class A	126	48	54%
Class B	16383	7006	43%
Class C	2097151	40724	2%

Only 2% of more than 2 million Class C addresses assigned !!!

Source: RFC 1335

The table above shows a statistic for the assignement of IP addresses in April 1992. Obviously, class A and B addresses have been allocated quicker than class C addresses. In the following years the utilization of class C addresses increased rapidly while class A and B addresses were spared.

Especially VLSM and NAT (invented 1994) supported the utilization of class C addresses.

Supernetting (RFC 1338)

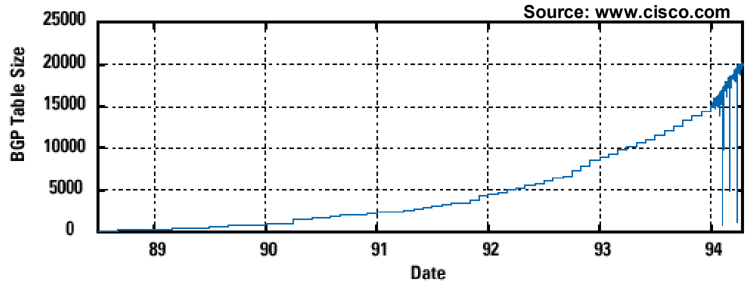

- In 1992: RFC 1338 stated scaling problem:
  - Class B exhaustion
  - No class for typical organizations available
  - Unbearable growth of routing table
- Use subnetting technique also in the Internet !
  - Do hierarchical IP address assignment !
  - Aggregation = "Supernetting"
    (Smaller netmask than natural netmask)

RFC 1338 introduced Supernetting: an Address Assignment and Aggregation Strategy, now obsoleted by RFC 1519.

Classful Routing Update

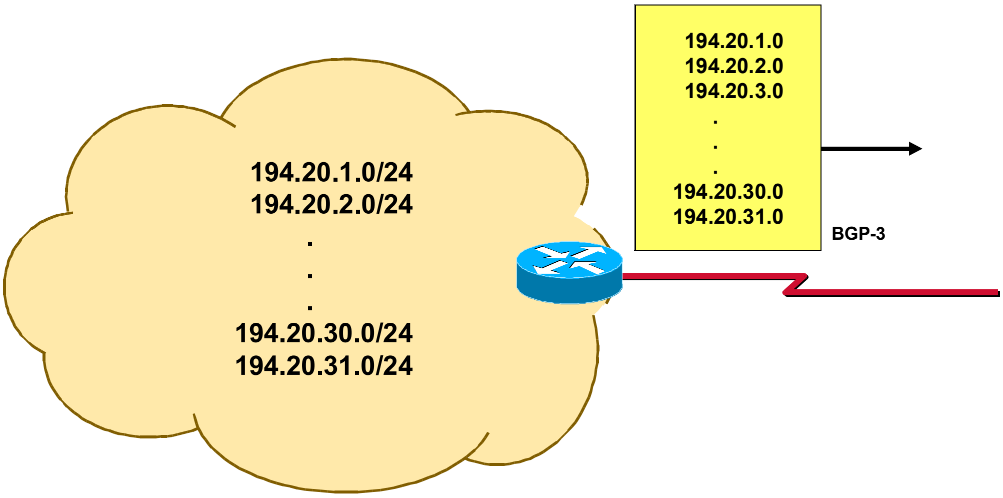

BGP-3 was a classfull routing protocol, sending the information about major class A, B, and C networks only.

Now Classless and Supernetting

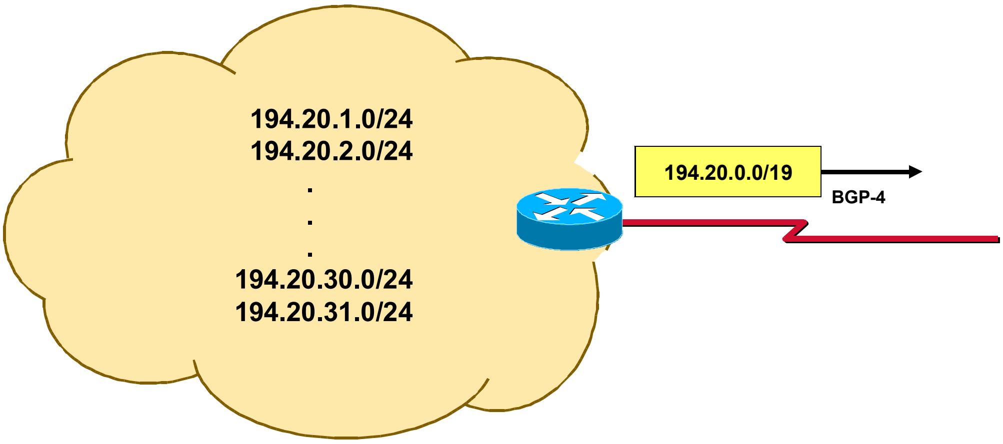

BGP-4 is classless, it can aggregate a range of class C network in one supernet.

CIDR

- September 1993, RFC 1519: Classless Inter-Domain Routing (CIDR)
- Requires classless routing protocols
  - BGP-3 upgraded to BGP-4
  - New BGP-4 capabilities were drawn on a napkin, with all implementors of significant routing protocols present (legend)
  - RFC 1654

RFC 1519 introduced **Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy**

RFC 1654 a draft standard for BGP – 4

RFC 1771 a standard for BGP - 4

Address Management

- ISPs assign
  contiguous blocks of
  contiguous blocks of
  contiguous blocks ...
  of addresses to their customers
- Aggregation at borders possible !
- Tier I providers filter routes with prefix lengths larger than /19
  - But more and more exceptions today...

To minimize the sizes of the routing tables ISPs use agregation, giving the customers the contiguous blocks of networks or subnets. Most of the ISPs would not accept routes from other ISP if the prefix is longer than /19.

International Address Assignment

- August 1990, RFC 1174 (by IAB) proposed regionally distributed registry model
  - Regionally means continental ;-)
- Regional Internet Registries (RIRs)
  - RIPE NCC
  - APNIC
  - ARIN

RFC 1174 IAB Recommended Policy on Distributing Internet Identifier Assignment.

This RFC represents the official view of the Internet Activities Board (IAB), and describes the recommended policies and procedures on distributing Internet identifier assignments and dropping the connected status requirement.

RIRs

- RIPE NCC (1992)
  - Réseaux IP Européens (RIPE) founded the Network Coordination Centre (NCC)
- APNIC (1993)
  - Asia Pacific Information Centre
- ARIN (1997)
  - American Registry for Internet Numbers
- AfriNIC
  - Africa
- LACNIC
  - Latin America and Caribbean

RIPE NCC is located in Amsterdam and serves 109 countries including Europe, Middle-East, Central Asia, and African countries located north of the equator. The RIPE NCC currently consists of more than 2700 members.

APNIC was relocated to Brisbane (Australia) in 1998. Currently there are 700 member organizations. Witin the APNIC there are also five National Internet Registries (NIRs) in Japan, China, Korea, Indonesia, and Taiwan, representing more than 500 additional organizations.

AfriNIC and LACNIC are relatively new RIRs (2002?).

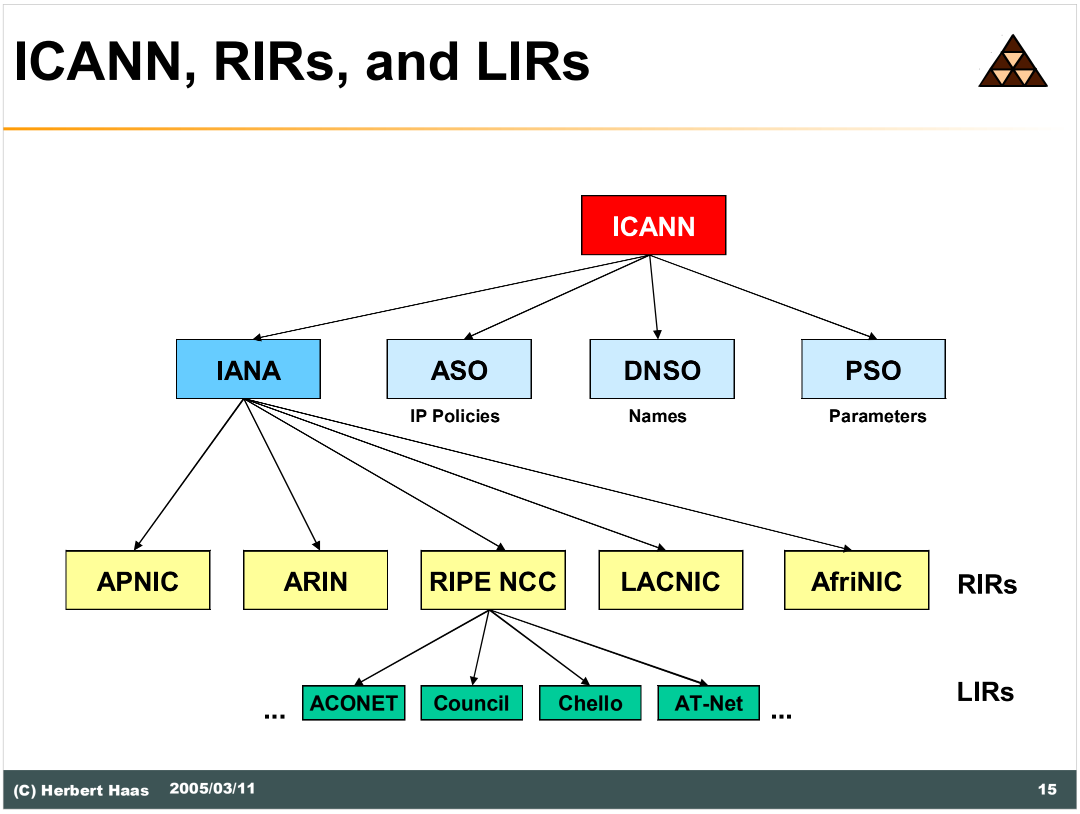

After foundation of the **ICANN**, the **Internet Assignment Numbers Authority (IANA)** is only responsible for IP address allocation to RIRs.

Other sub-organizations of the ICANN:

> **Address Supporting Organization (ASO)**, which was founded by APNIC, ARIN, and RIPE NCC, and should oversee the recommendations of IP policies
>
> **Domain Name Supporting Organization (DNSO)** is responsible for maintaining the DNS
>
> **Protocol Supporting Organization (PSO)** is responsible for registration of various protocol numbers and parameters used by RFC protocols

Originally, all tasks of these sub-organizations were performed by the IANA only. Today the IANA only cares for address assignment to the RIRs.

The slide above shows a few of the long list of LIRs in Austria. These LIRs are those who are widely known by Internet users as "Internet Service Providers".

CIDR Concepts Summary

- Coordinated address allocation
- Classless routing
- Supernetting

RFC 1366 Address Blocks

- 192.0.0.0 - 193.255.255.255 ... Multiregional
- 194.0.0.0 - 195.255.255.255 ... Europe
- 198.0.0.0 - 199.255.255.255 ... North America
- 200.0.0.0 - 201.255.255.255 ... Central/South America
- 202.0.0.0 - 203.255.255.255 ... Pacific Rim

RFC 1366 Guidelines for Management of IP Address Space, was obsoleted by 1466 in 1993, in 1996 an RFC 2050 came out.

Class A Assignment

- **IANA responsibility**
  - **RFC 1366 states:** ***"There are only approximately 77 Class A network numbers which are unassigned, and these 77 network numbers represent about 30% of the total network number space."***
- **64.0.0.0 – 127.0.0.0 were reserved for the end of (IPv4) days ?**
  - **Recent assignments (check IANA website)**

The Class A addresses assignment is controled by the IANA.

Class B Assignment

- IANA and RIRs requirements
  - Subnetting plan which documents more than 32 subnets within its organizational network
  - More than 4096 hosts
- RFC 1366 recommends to use multiple Class Cs wherever possible

In order to receive a class B address, an organization must fulfill strict requirements such as employing more than 4096 hosts and more than 32 subnets.

Class C Assignment

- If an organization requires more than a single Class C, it will be assigned a bit-wise contiguous block from the Class C space
- Up to 16 contiguous Class C networks per subscriber (= one prefix, 12 bit length)

```
       Organization                               Assignment

  1) requires fewer than 256 addresses     1 class C network
  2) requires fewer than 512 addresses     2 contiguous class C networks
  3) requires fewer than 1024 addresses    4 contiguous class C networks
  4) requires fewer than 2048 addresses    8 contiguous class C networks
  5) requires fewer than 4096 addresses   16 contiguous class C networks
```


Example (RFC 1366) for Class C assignment:

For instance, an European organization which requires fewer than 2048 unique IP addresses and more than 1024 would be assigned 8 contiguous class C network numbers from the number space reserved for European networks, 194.0.0.0 - 195.255.255.255. If an organization from Central America required fewer than 512 unique IP addresses and more than 256, it would receive 2 contiguous class C network numbers from the number space reserved for Central/South American networks, 200.0.0.0 - 01.255.255.255.

RFC 1918 – Private Addresses

- In order to prevent address space depletion, RFC 1918 defined three private address blocks
  - 10.0.0.0 - 10.255.255.255 (prefix: 10/8)
  - 172.16.0.0 - 172.31.255.255 (prefix: 172.16/12)
  - 192.168.0.0 - 192.168.255.255 (prefix: 192.168/16)
- Connectivity to global space via Network Address Translation (NAT)

RFC 1918 defines an "Address Allocation for Private Internets", that is three address spaces, which should only be used in private networks.

Any route to this network must be filtered in the Internet! Any router in the Internet must not keep any RFC 1918 address in its routing table!

Together with these addresses, Network Address Translation (NAT) is needed if private networks should be connected to the Internet.

This solution greatly reduces the number of allocated IP addresses and also the routing table size because now class C networks can be assigned very efficiently, using a prefix up to /30.

NAT Example

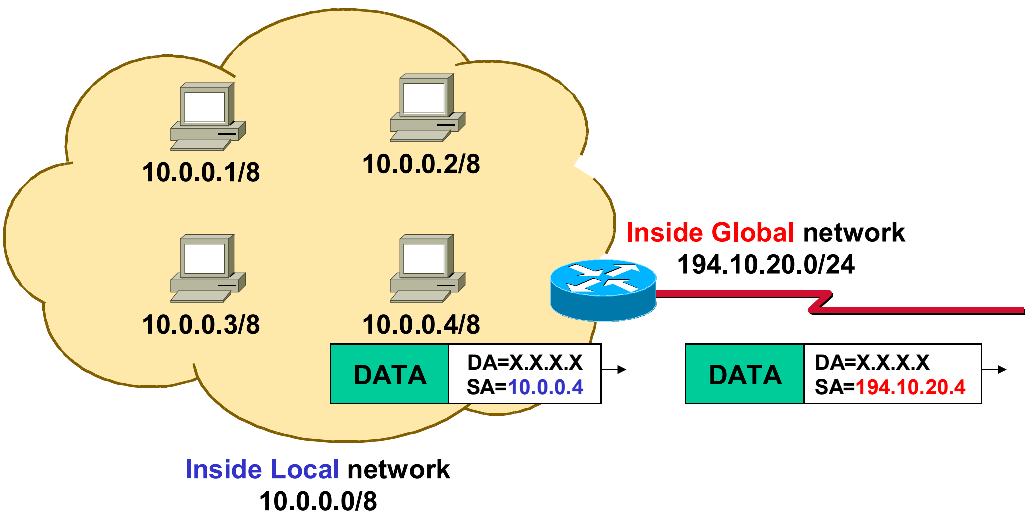

Network Address Translation (NAT) In order to be able to comunicate with internet we have to translate private addresses (inside local) into oficial, assigned by an ISP (inside global).

But...

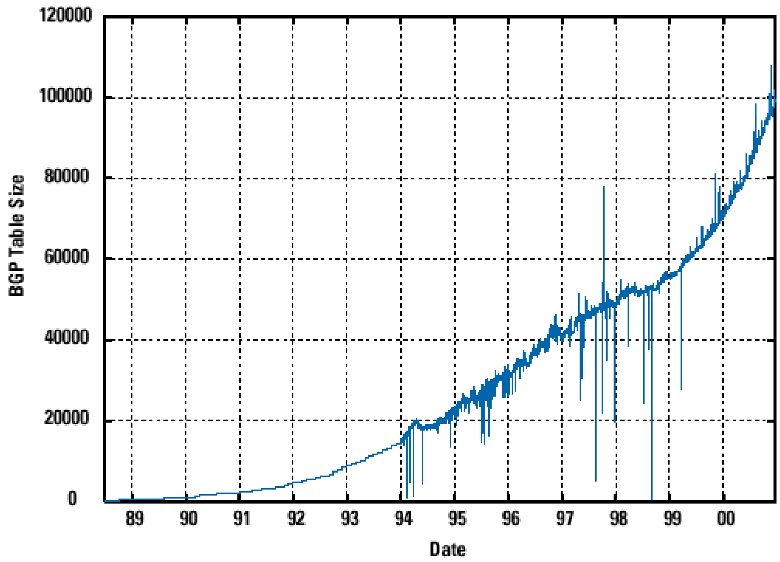

Source: www.cisco.com

But this is not really the end of the story. The growth rate of the Internet was and is generally exponential, that is exp(k*x). Soon after the introduction of CIDR the progressive factor k increased dramatically, thus even CIDR could only reduce k, but not the general exponential character.

It is interesting to question how long the (also exponential) growth rate of silicon memory and processing power together with CIDR and NAT can mitigate the effects of the Internet growth.

As for today, the only solution to deal with this problem in the long run is to introduce IPv6 and a more hierarchical routing strategy.