Prospective of IPv6 as a Future Internet Engine

Xing Li

2012-09-13

Outline

- Challenges and requirements
- Design compromise
- Lessons learned
- Rethink of the design principles
- Future

Challenges and requirements

Places, people, things

The majority of traffic is now packet

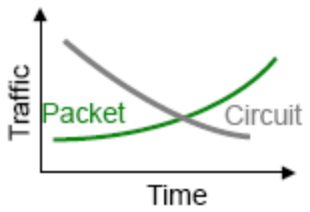

Significant proliferation of mobile broadband

Workforce mobility & increased connectivity

Enterprise IT is shifting to the cloud

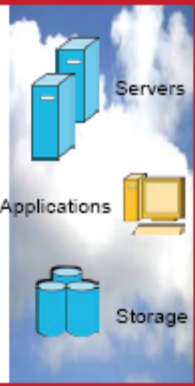

Video, video and more video

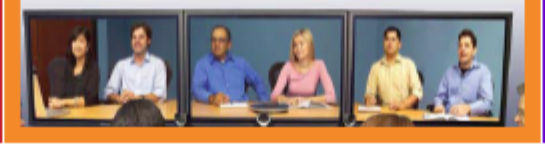

Real-time, low latency applications

High-frequency trading & other financial apps

Medical imaging & other healthcare apps

Native video transport & other video apps

First we Connected Places, Then People, Then Things

RIR IPv4 address run-down

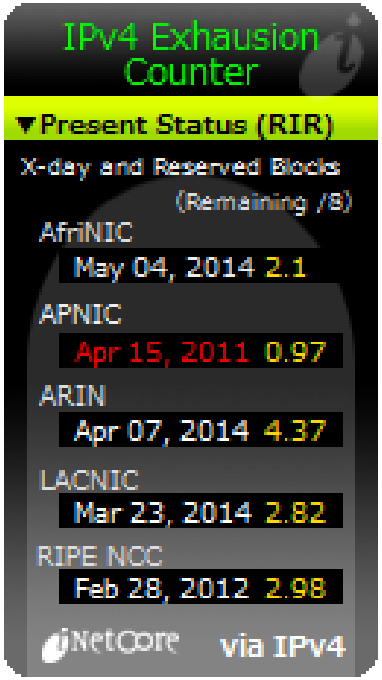

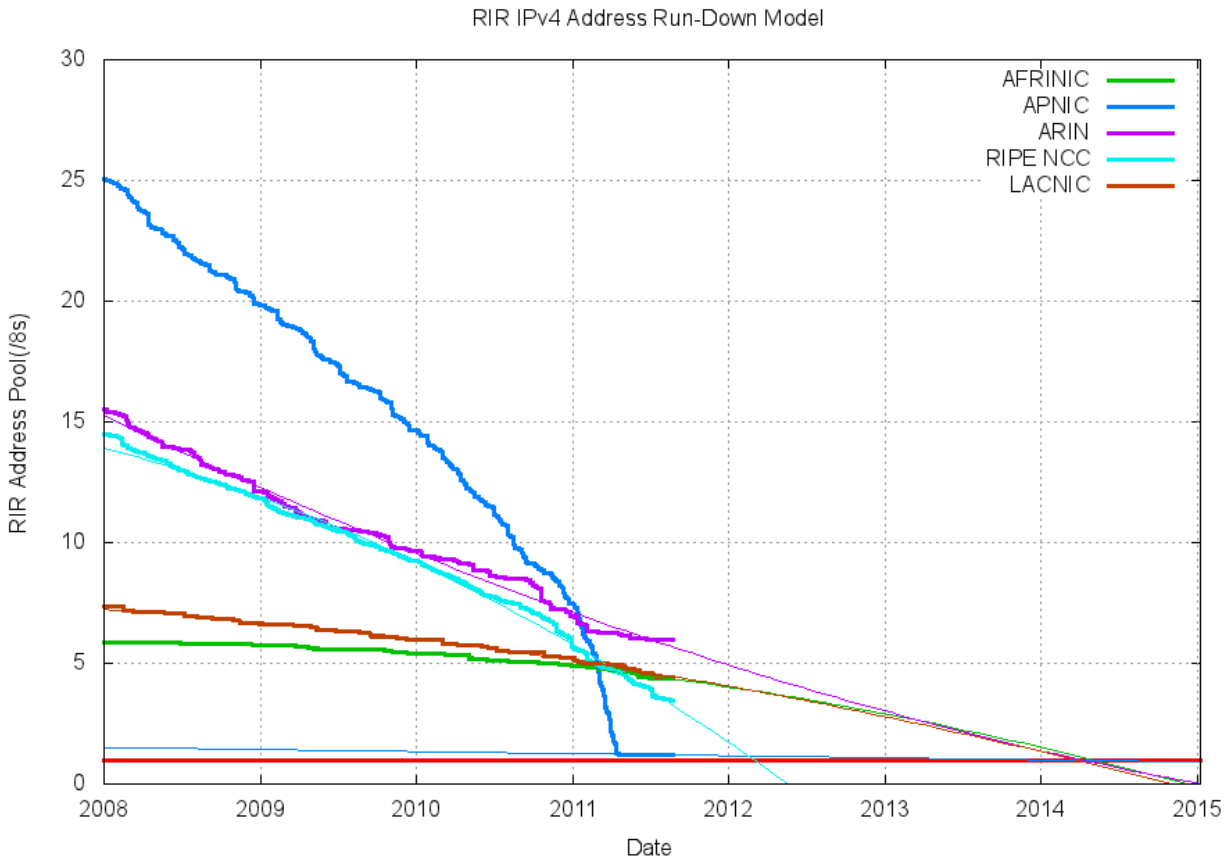

IPv4 routing table explosion

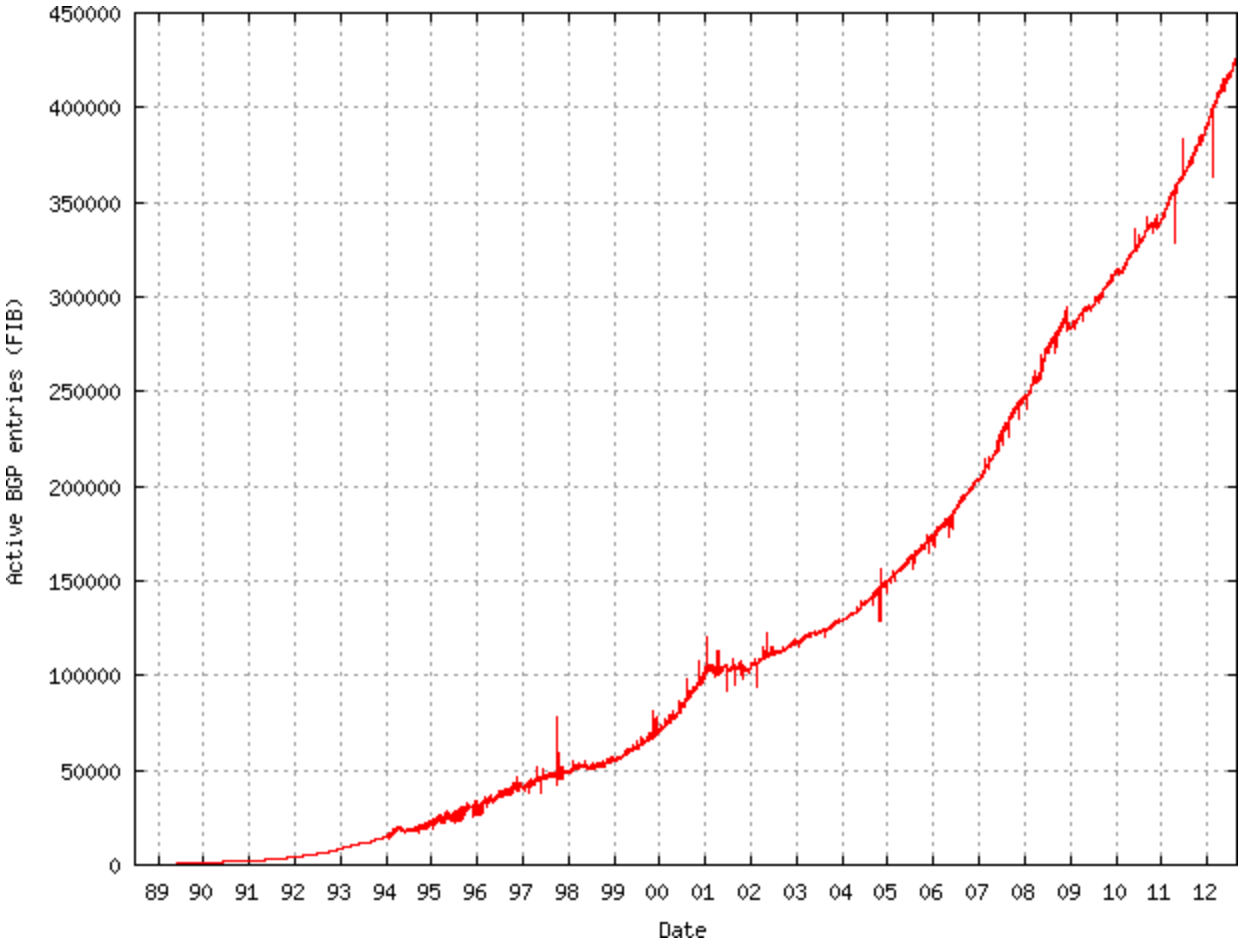

Challenged Business Model

1 Explosive Bandwidth Growth

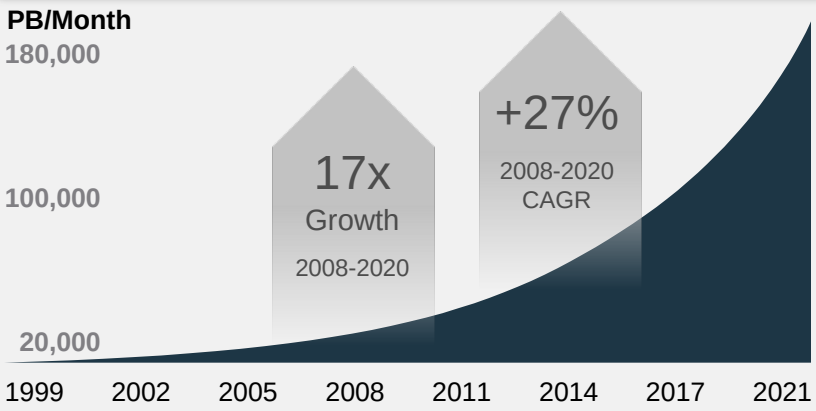

Source: Juniper, Cisco, MINTS

2 Dynamic and Unpredictable

Mobility

IP Apps

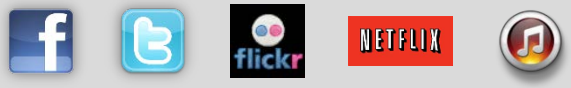

Cloud

3 Business Model Challenged

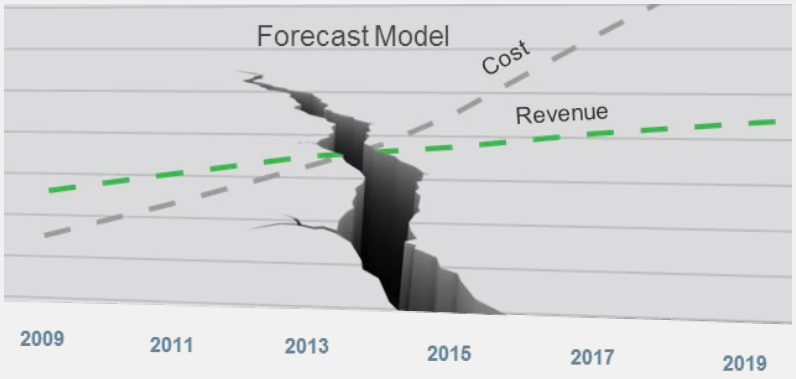

4 Today's Architecture Doesn't Scale

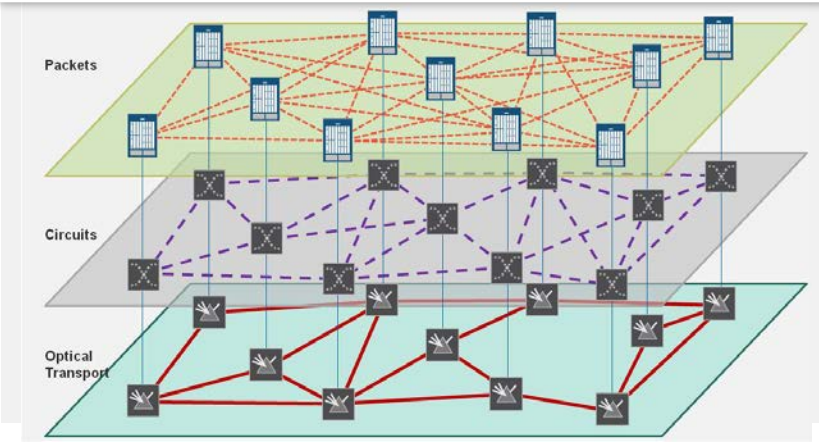

Security incident explosion

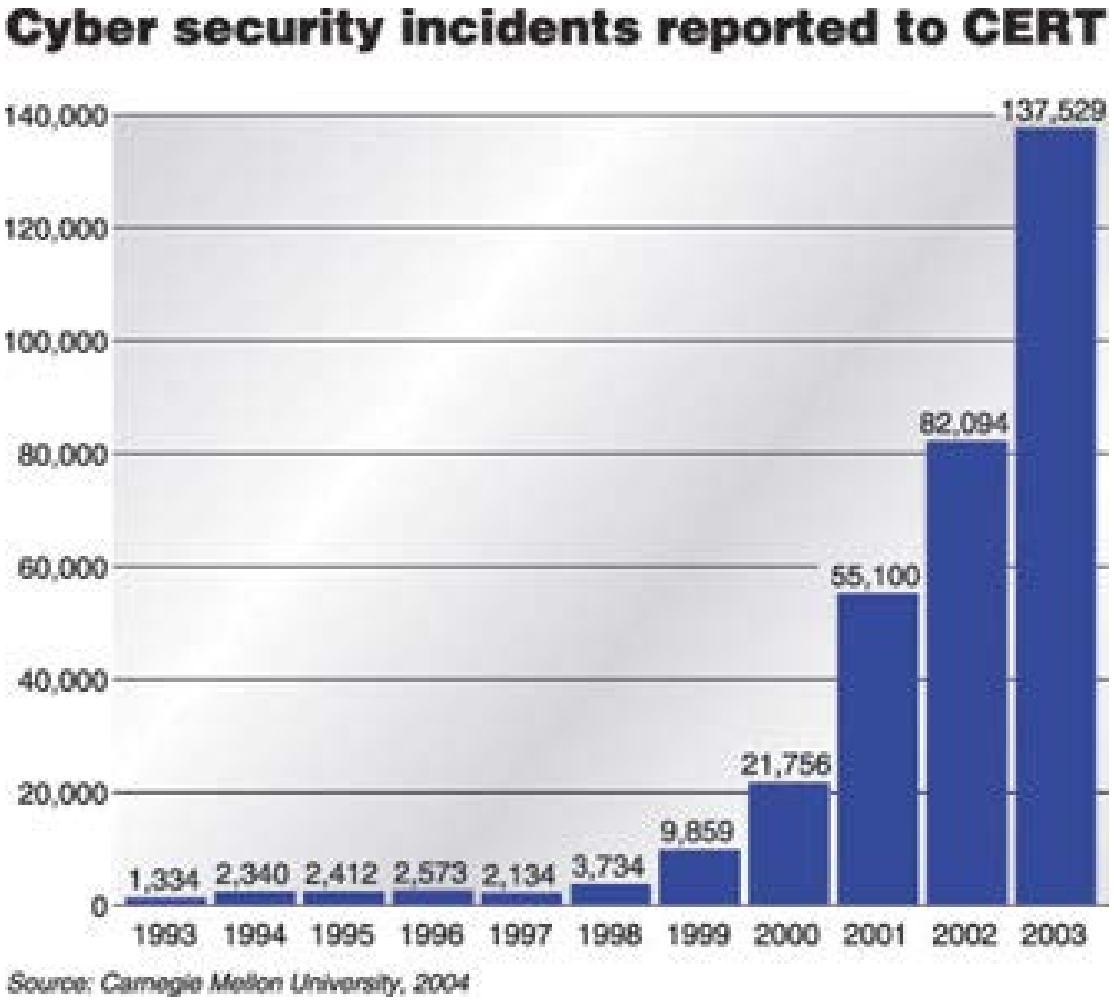

New backbone

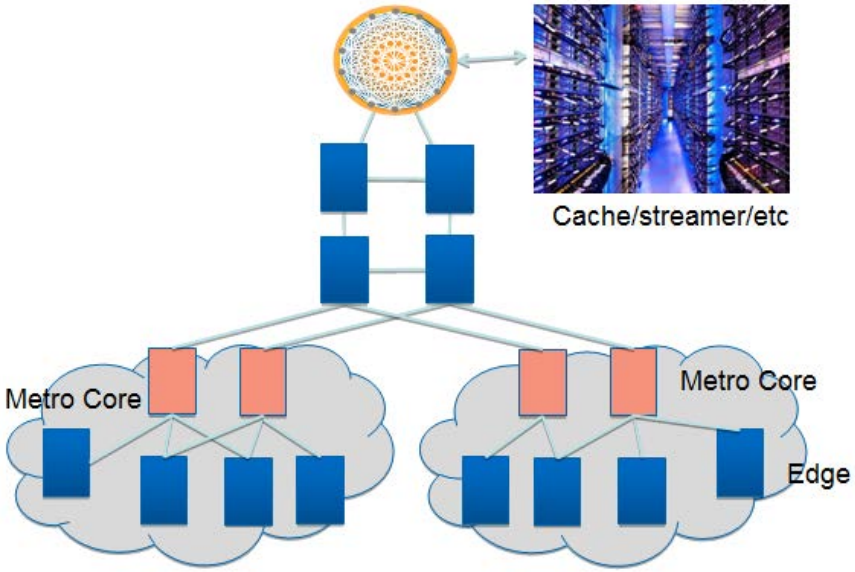

Large Objects in Consolidated DC

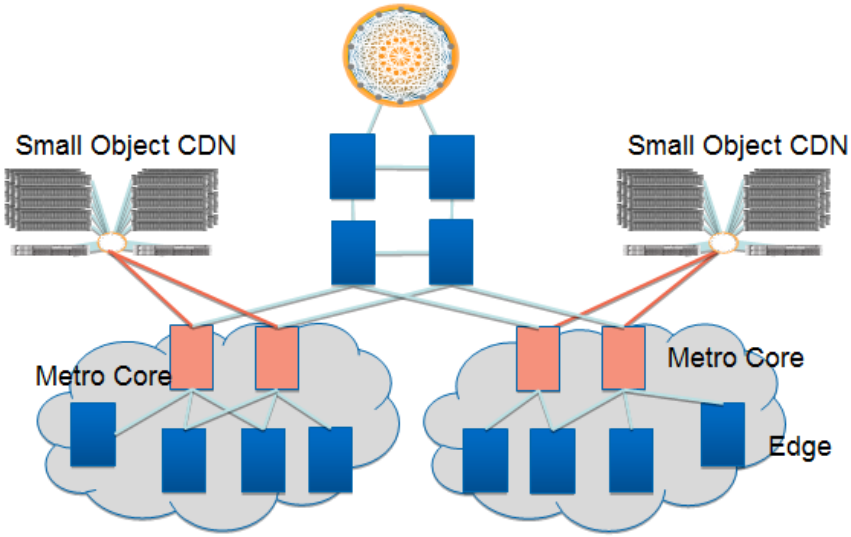

Small Objects In Metro Core

Design compromise

Internet architecture

- Need to interconnect many existing networks
- Hide underlying technology from applications
- Decisions:
  - Network provides minimal functionality
  - “Narrow waist”

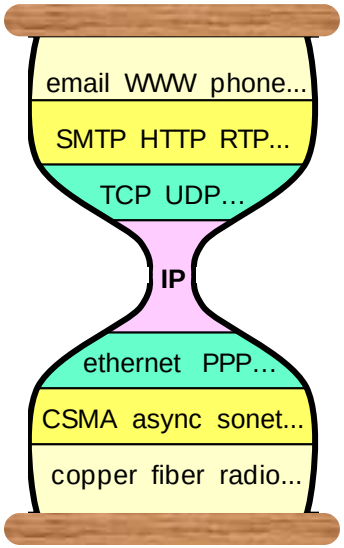

Tradeoff: **No assumptions, no guarantees.**

IPv4 & IPv6 Header Comparison

IPv4 Header

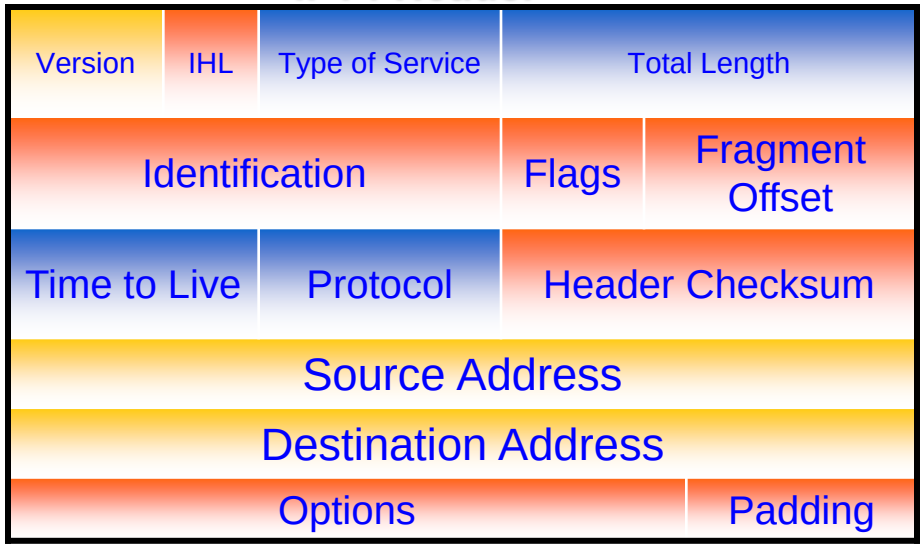

IPv6 Header

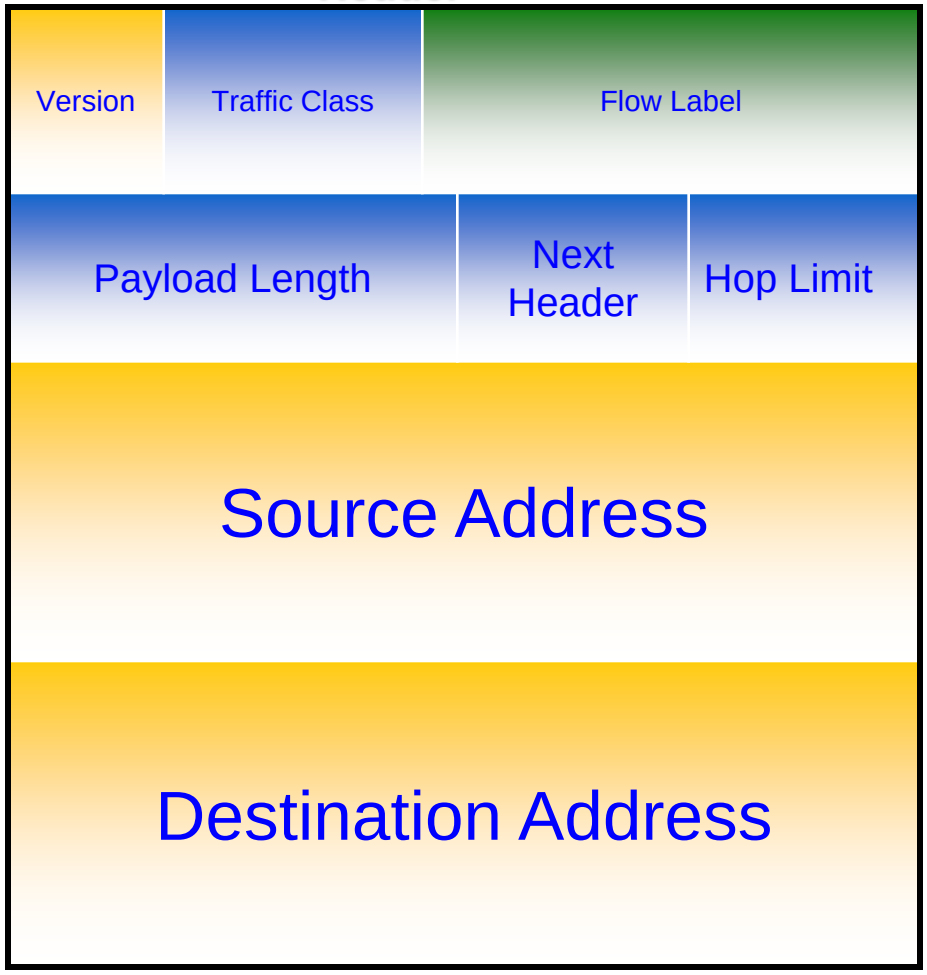

Legend

- **field's name kept from IPv4 to IPv6**
- **fields not kept in IPv6**
- **Name & position changed in IPv6**
- **New field in IPv6**

Controversy on addressing

- The length of addresses
  - 64 bits vs. variable length
- Providers and monopolies
  - keep the network simple
- Flows and services
  - network map
- Variable format and renumbering
  - provider/subscriber

Controversy on protocol

- Do we need more than 255 hops?
  - Large network vs routing loop
- Is the destination address in the right place?
  - 128bit processor
- Should packets be larger than 64K?
  - Jumbogram option
- Can we live without a checksum?
  - Link layer checksum and risk analysis

Controversy on security

- Should we mandate security?
  - export restriction/clipper chip
- Did we choose the correct algorithm?
  - MD5/DES-CBC not secure enough, too slow, or both
- Is this the right layer?
  - New network code
- Do we need additional protection?
  - DOS, traffic pattern

Controversy on flows

- Will flow labels be used?
  - There are more flows than sources
- To reserve or not?
  - Capacity/Adaptive applications
- What about ATM?
  - As simple as possible, in order to scale

Controversy on transitioning

- Dual-stack strategy
  - The IPv6 code, ICMP, neighbor discovery code
  - The handling of IPv6 within TCP and UDP
  - Modifications to the sockets libs
  - The interface with the name service
- Encapsulating IPv6 in IPv4
  - Tunneling
- Translation
  - IPv4 and IPv6 are not compatible

AS count

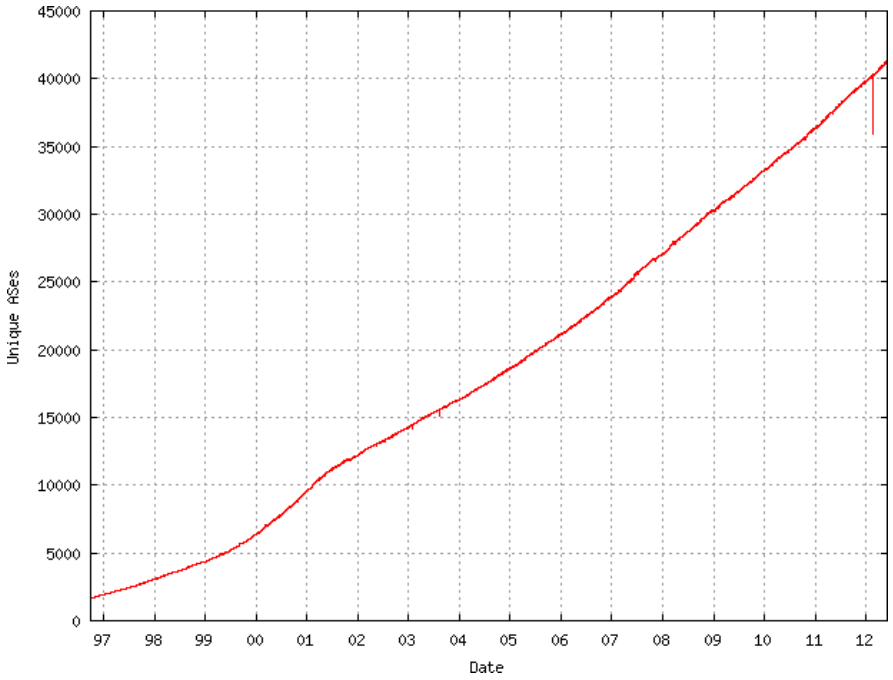

IPv4

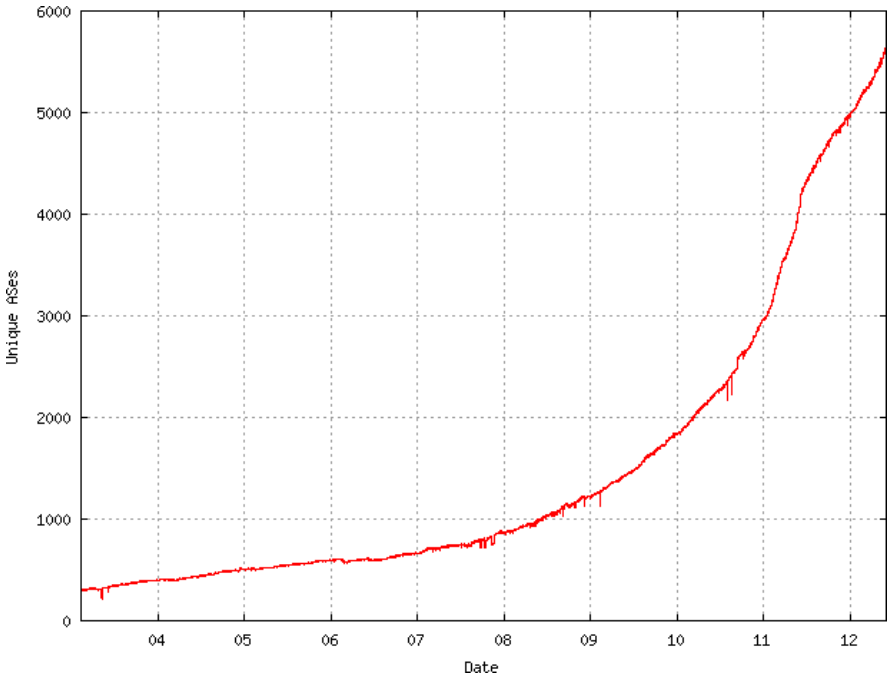

IPv6

Lessons learned

CERNET (IPv4)

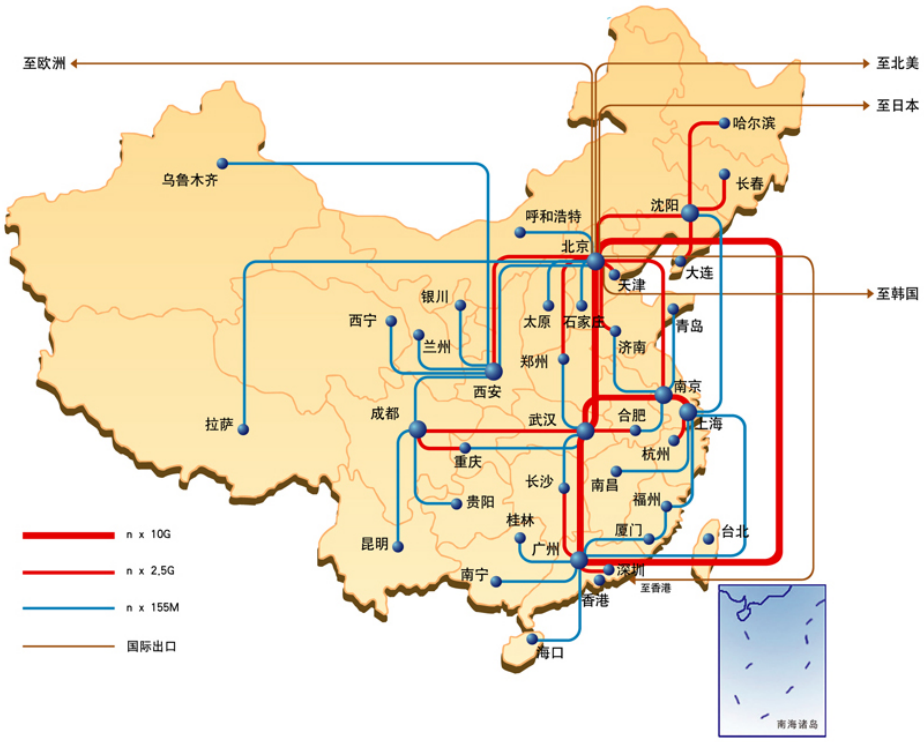

- CERNET is the first (1994) nation wide Internet backbone in China.
- CERNET ranks 23 in global CIDR report.
- Over 2,000 universities on CERNET with about 20M subscribers.

CNGI-CERNET2 design concepts

- Protocol selection
  - IPv6-only
- Complicity
  - Multiple AS's
  - Multiple vendors
- Transition strategy
  - High performance
  - Free
- IETF related works
  - IPv4 over IPv6 (Softwire WG)
  - IPv4/IPv6 translation (Behave WG, softwire WG)
  - Source address validation improvement (SAVI WG)

CERNET2 (IPv6)

- Built in 2004, with national coverage
- CERNET2 is the largest IPv6 backbone in China.
- About 200 universities connected to CERNET2 with about 2M subscribers.

IPv6 applications

- Video
  - Beijing 2008 Olympic website
  - Medical applications
  - Musical performance

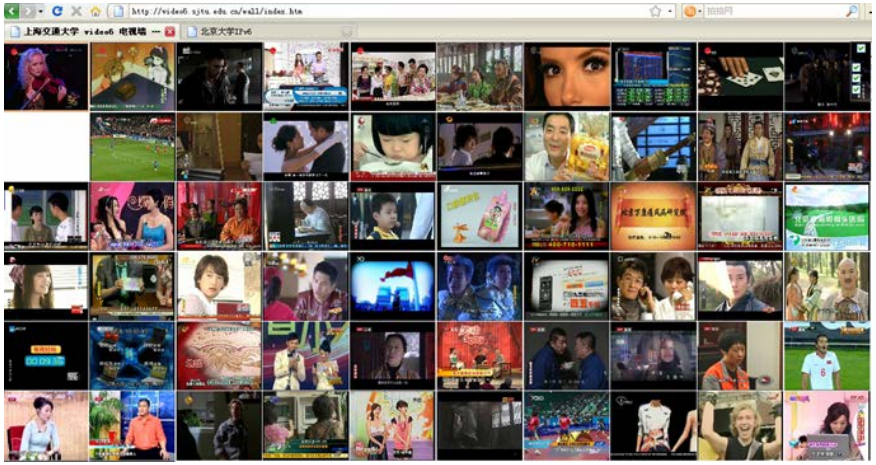

The killer application

- Video?
- P2P?
- Internet of Things?
- The intercommunication with the IPv4 Internet is the killer application of IPv6.

IVI

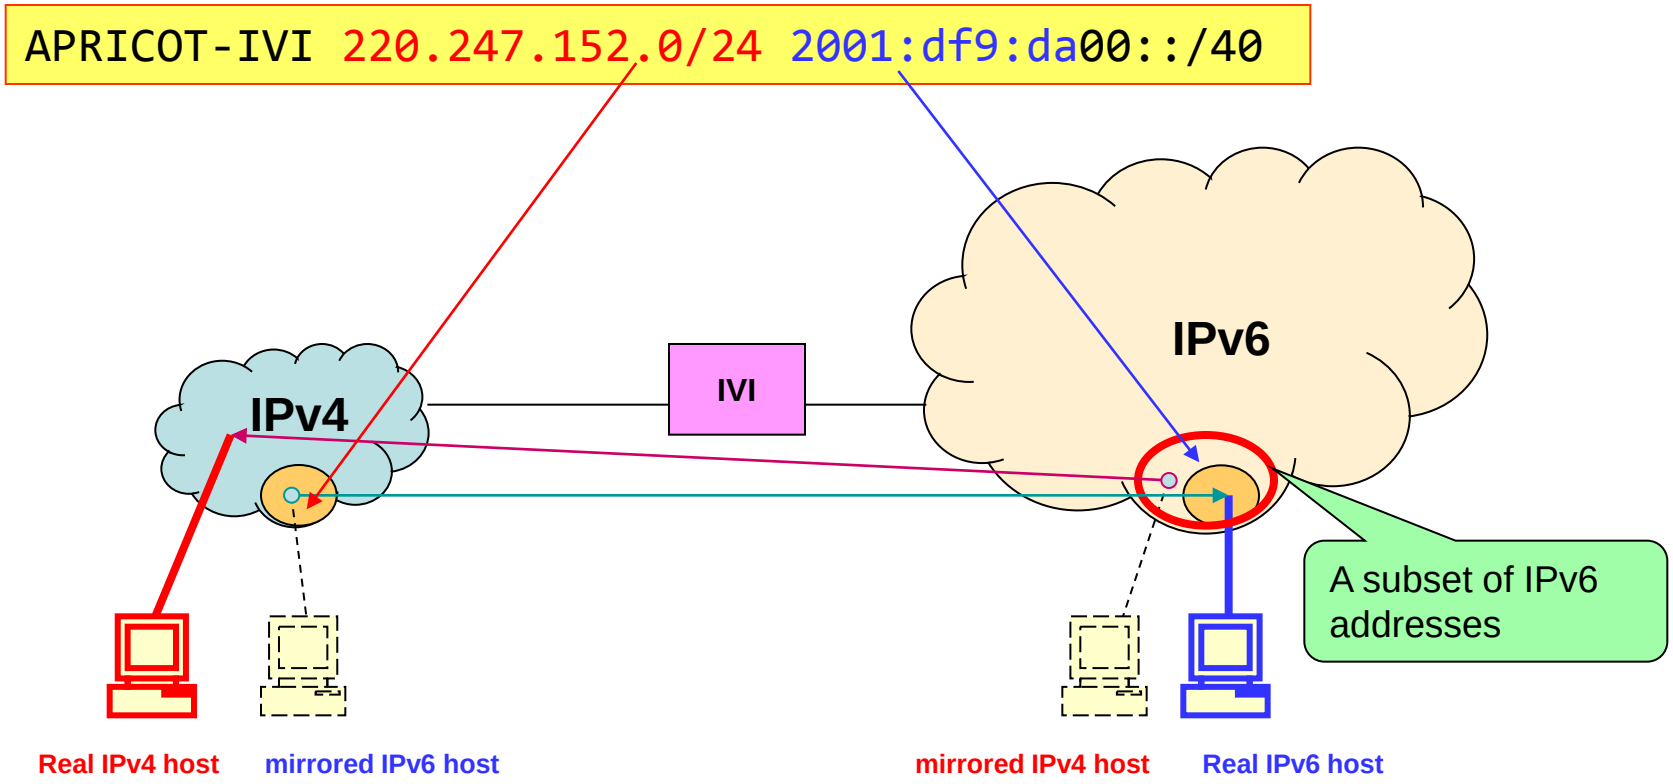

A subset of IPv6 addresses

CERNET/CNGI-CERNET2

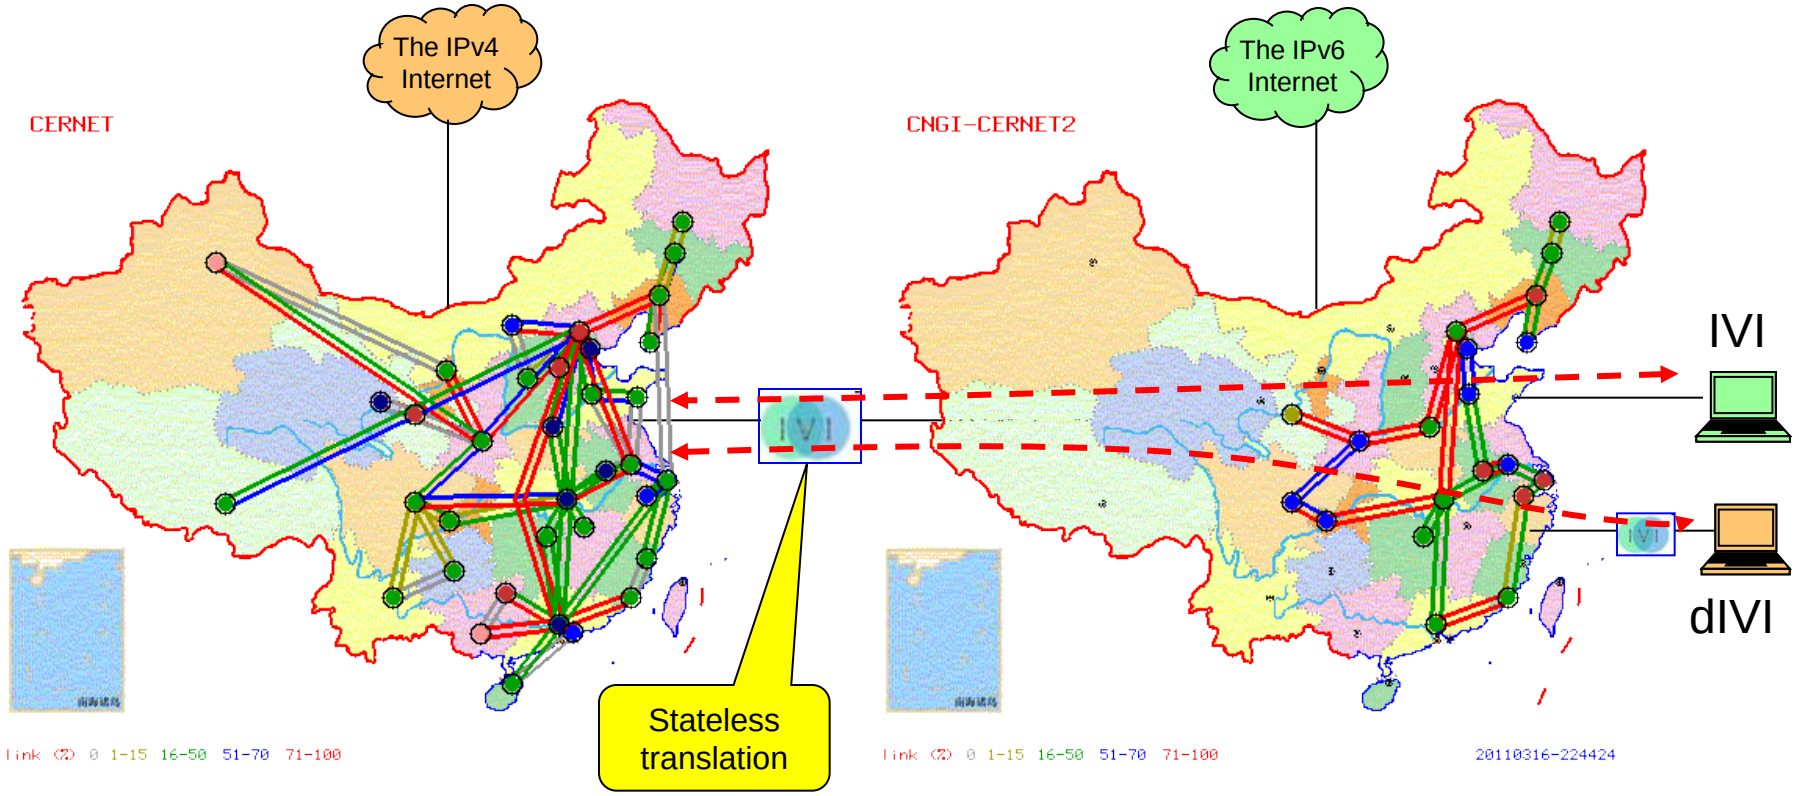

Traffic comparisons

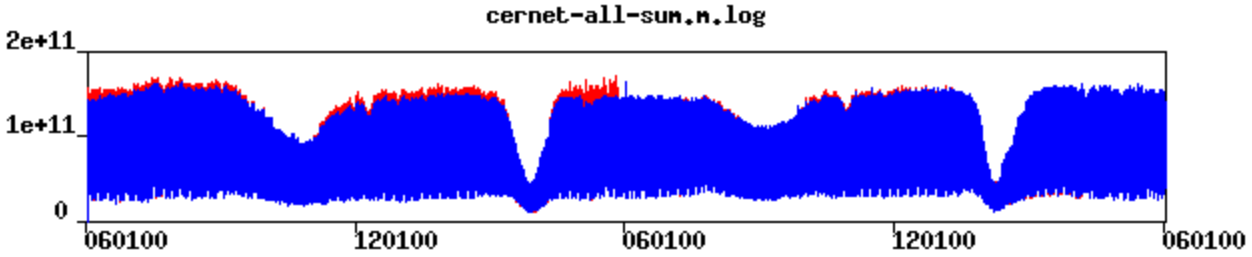

2008 2010 2012

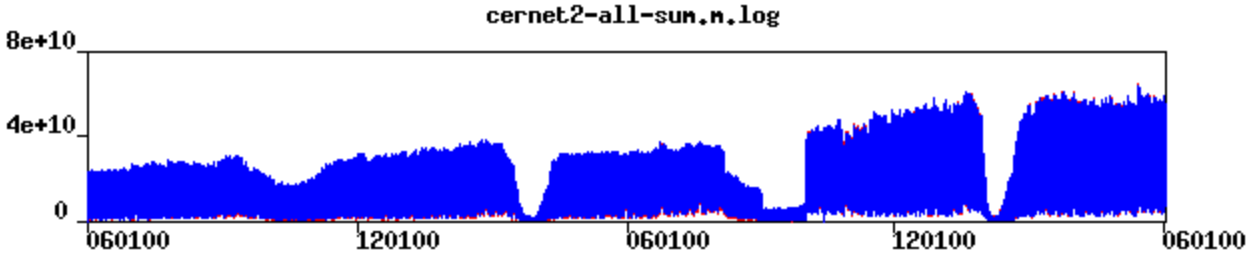

2008 2010 2012

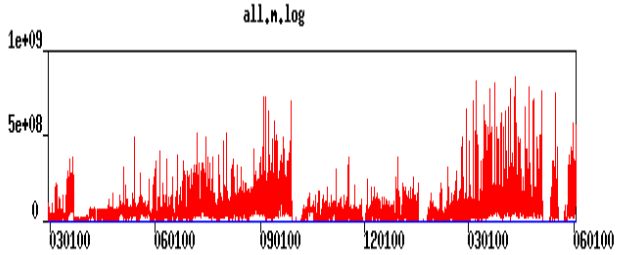

- **CERNET IPv6' traffic is about 20% of IPv4**

CERNET IPv6 transition experience

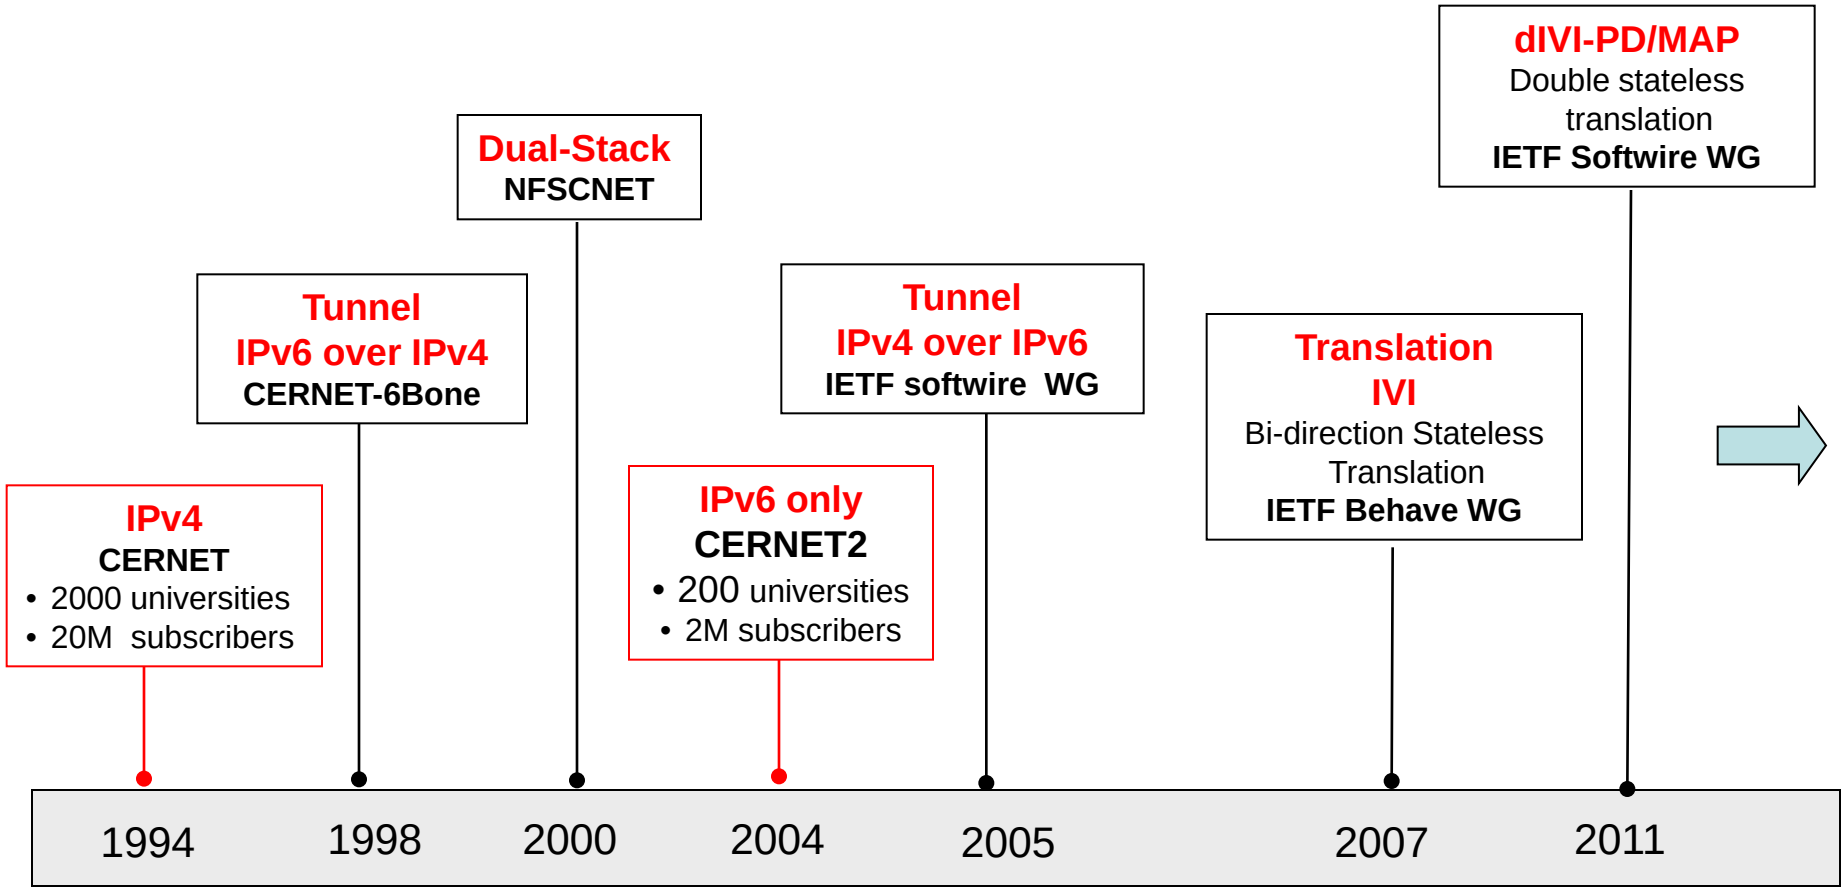

Transition evolution

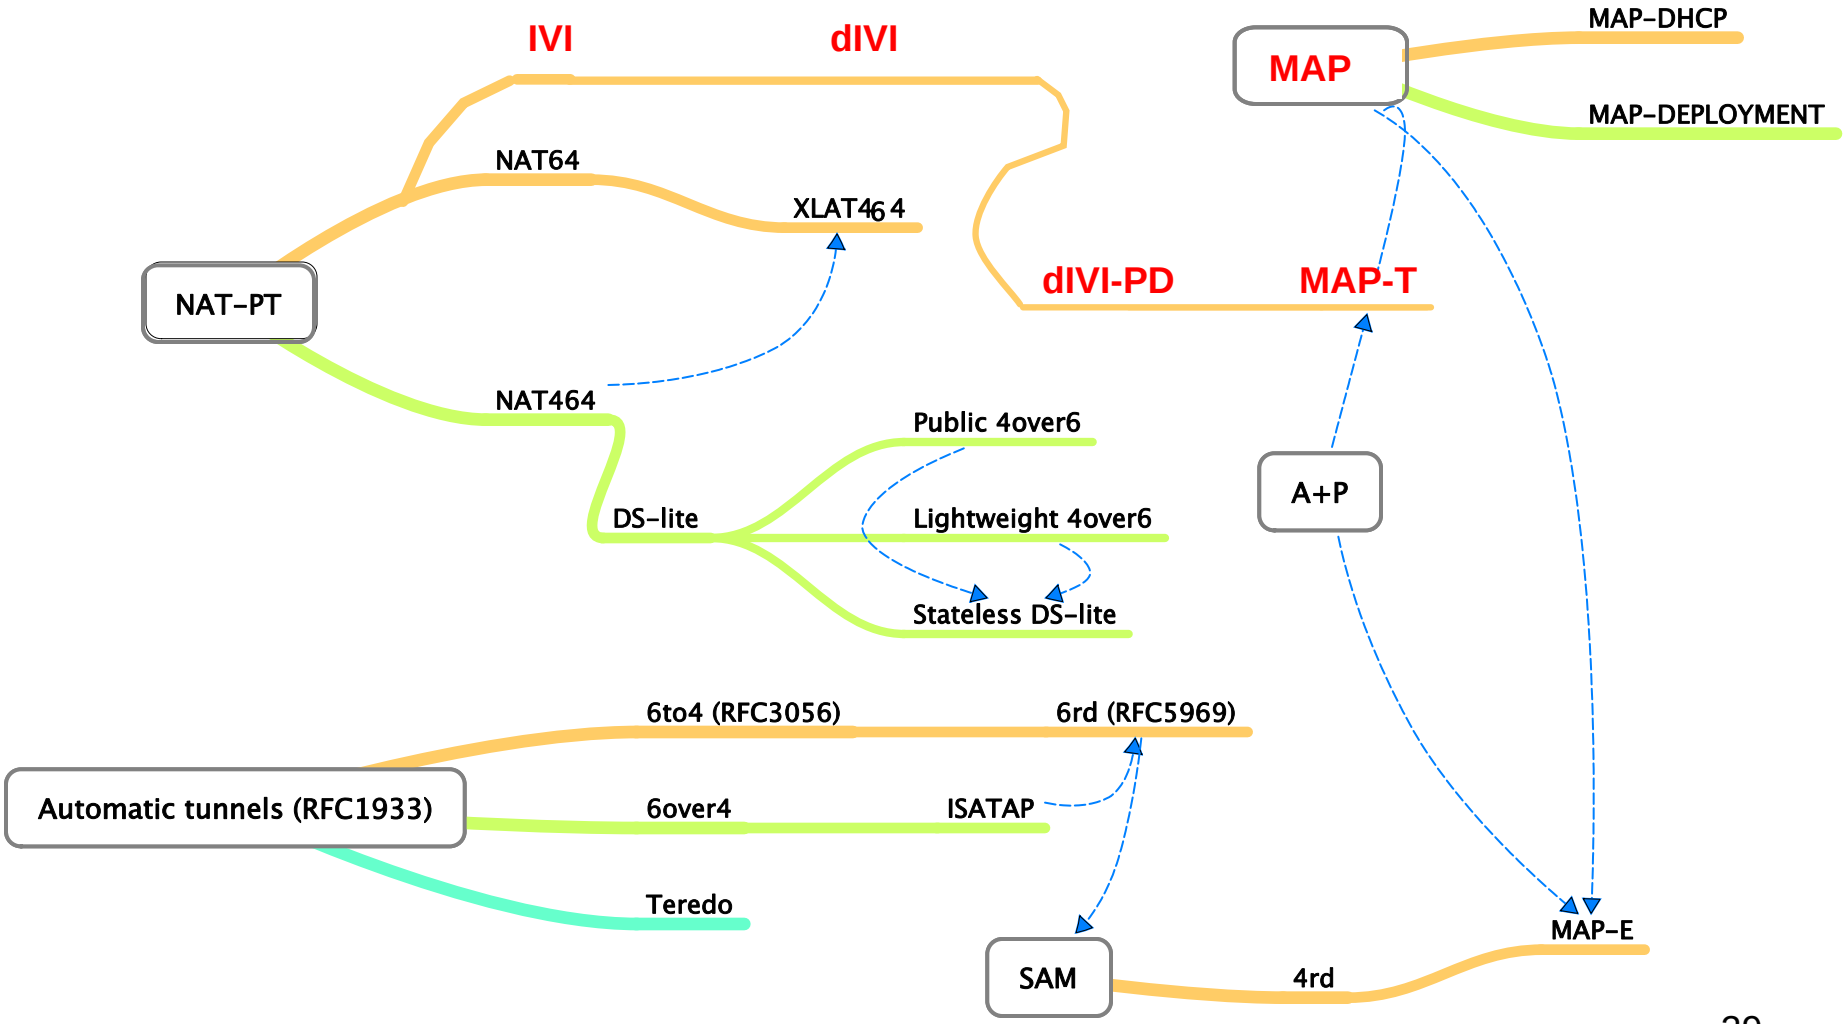

Remarks

- Single and double translation can be mixed
- Double translation and encapsulation can be unified
- If stateless works, the stateful should work

Principles of IPv4/IPv6 transition

- Native IPv6 (both ends are in IPv6)
- Single translation (the other side is in IPv4)
- Double translation (native IPv4 app and ALG)
- Encapsulation (IPv4 header transparency)

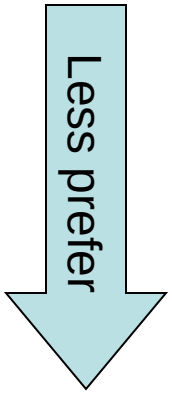

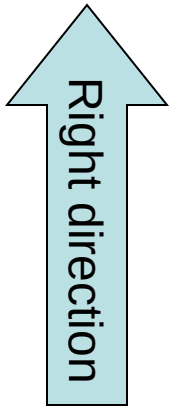

Rethink of the design principles

General Design Issues

- Heterogeneity is inevitable and must be supported by design.
- If there are several ways of doing the same thing, choose one.
- All designs must scale readily to very many nodes per site and to many millions of sites.
- Performance and cost must be considered as well as functionality.
- Keep it simple. When in doubt during design, choose the simplest solution.
- Modularity is good. If you can keep things separate, do so.
- In many cases it is better to adopt an almost complete solution now, rather than to wait until a perfect solution can be found.
- Avoid options and parameters whenever possible.
- Be strict when sending and tolerant when receiving.
- Be parsimonious with unsolicited packets, especially multicasts and broadcasts.
- Circular dependencies must be avoided.
- Objects should be self decribing (include type and size), within reasonable limits.
- All specifications should use the same terminology and notation, and the same bit- and byte-order convention.
- And perhaps most important: Nothing gets standardised until there are multiple instances of running code.

RFC1958

The principles of design (1)

- *Principles such as simplicity and modularity are the stuff of software engineering*
- *Decentralisation and tolerance are the life and breath of Internet*

Tim Berners-Lee

The principles of design (2)

- *Avoid needless complexity*
- *Support existing content*
- *Be conservative in what you send; be liberal in what you accept*
- *Solve real problems*
- *Pave the cowpaths*
- *Degrade gracefully*
- *The value of a network is proportional to the square of the number of connected users of the system (n2)*
- *Design for humans first, machines second*
- *Rough consensus and running code*

What IPv6 can do for FI?

- Principles
  - End-to-end → hop-to-hop
  - Stateless → stateful
- Resources
  - 128 bits addresses
  - Traffic class
  - Flow label
- Reality
  - Knowledge
  - Equipment
  - Networks

Future

Future Internet

- Should the current Internet principle be kept?
  - Connectionless
  - End-to-end
  - Best effort
- What is the transition scheme?
  - Multiple stack
  - Encapsulation
  - Translation

Remarks


THE INTERNET
WISHING TREE
BY THE INTERNET SOCIETY
Add your wish
Use the keyboard left / right arrows to navigate
I wish there is ONE Internet in our globe , not two Internets (IPv4 and IPv6).
Xing