the globus project™
www.globus.org

IBM®

IPv6, Grid Computing and Scaling Up the Internet

Brian E Carpenter

IBM and 6NET

Ian Foster

Argonne National Laboratory, University of Chicago, and Globus

6net

June 2003

brian@hursley.ibm.com
foster@mcs.anl.gov

Topics

- General Introduction
- Brief introduction to IPv6
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

General Introduction
Why discuss Grids and IPv6 together?

- Grid computing represents a fundamental shift in the approach to distributed computing, like the fundamental shift in information access introduced by the Web ten years ago.
- IPv6 represents a major step function in the Internet's ability to scale, like the introduction of IPv4 twenty years ago.
- Inevitably there is synergy between these two game changers.
- Let's share a common goal of reaching 10 billion Internet nodes.

Topics

- General Introduction
- Brief introduction to IPv6
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

Why we need IPv6

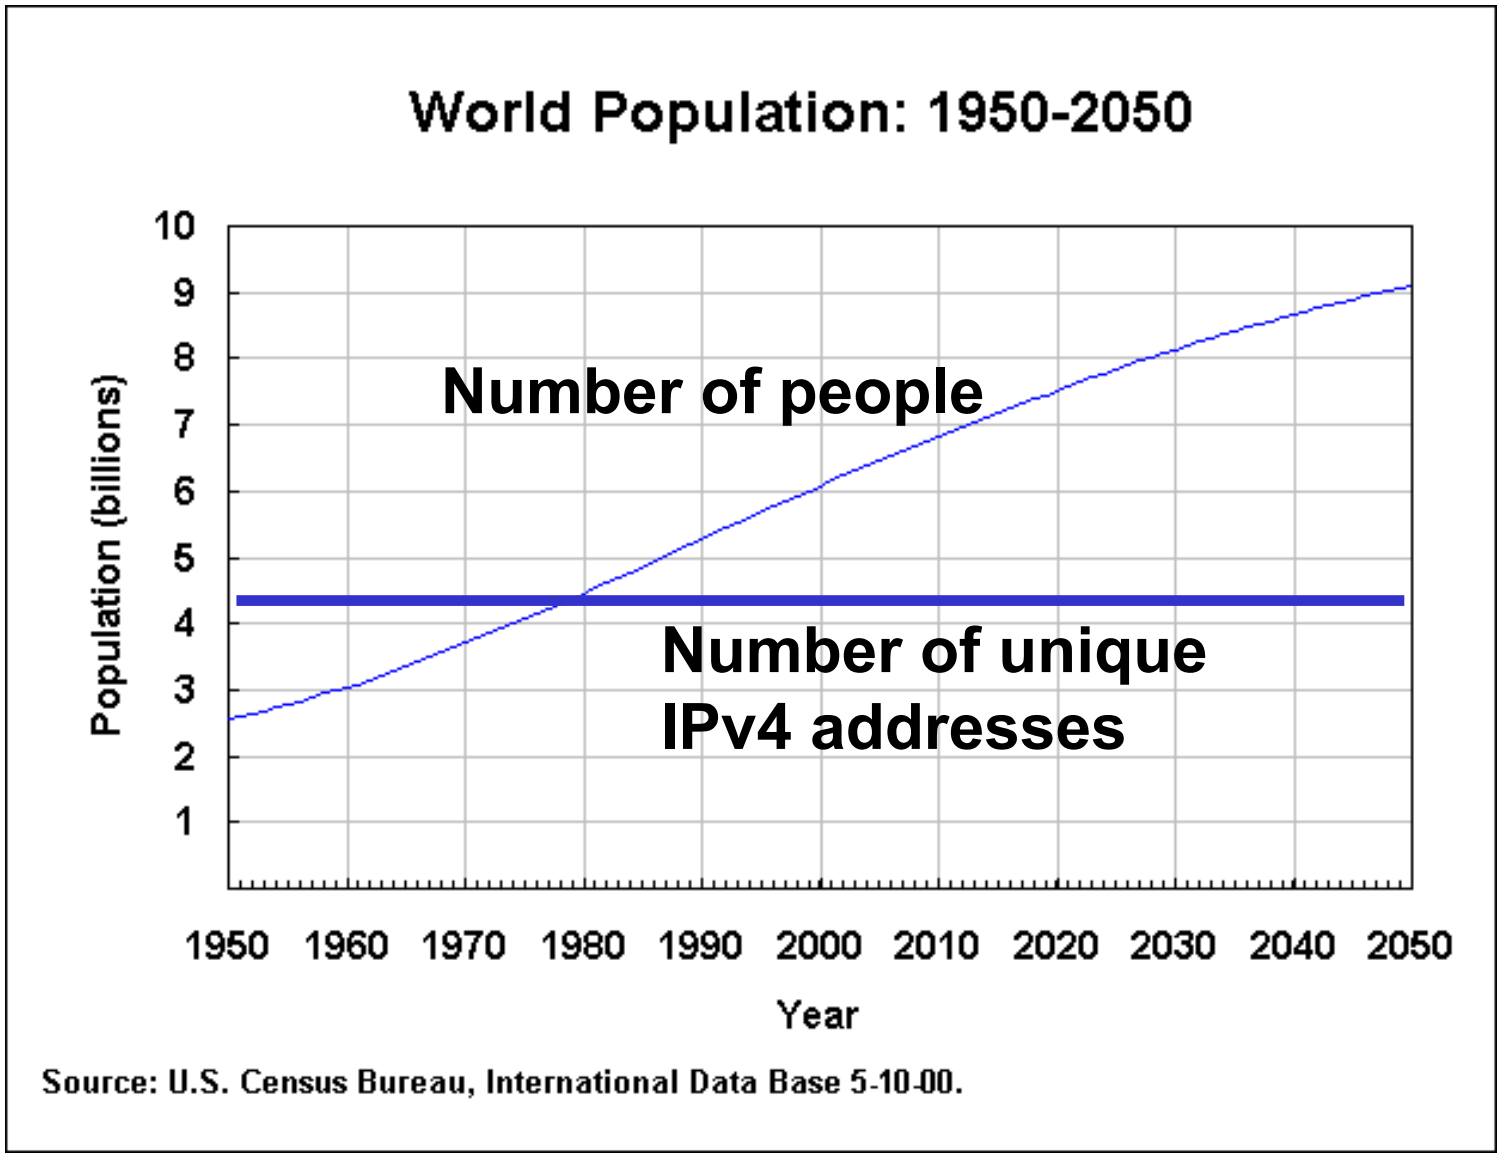

Living with too few addresses

- If we don't have many more addresses than we expect to have devices, we will have a fractured network with artificial internal boundaries.
  - The tense is wrong. Today in the US, there is widespread use of ambiguous (net 10) address space with consequent glitches and hacks.
  - Much more acute problem in (e.g.) China.
- This is a major operational cost and an obstacle to innovative applications.
  - In fact, that is exactly why Cerf and Kahn invented IP, but they didn't go far enough. It's time to fix that bug.

More addresses than people

- Let’s think of ten billion nodes as a modest target; that’s only one device per person.
- The only way out is bigger addresses.
- The IETF picked 128 bits.

Other major benefits of IPv6

- Automatic configuration
  - stateless, for manager-free networks
  - stateful (DHCPv6), for managed networks
  - help for site renumbering
- Better aggregated routing tables than IPv4
- Complete Mobile IP solution
- Global addressability allows IPSEC end to end.
  - mechanisms for secure firewall traversal will come
- Simplified header format with clean extensibility.
  - allows effective header compression
- Provision for a QOS flow label.

The IPv6 Header

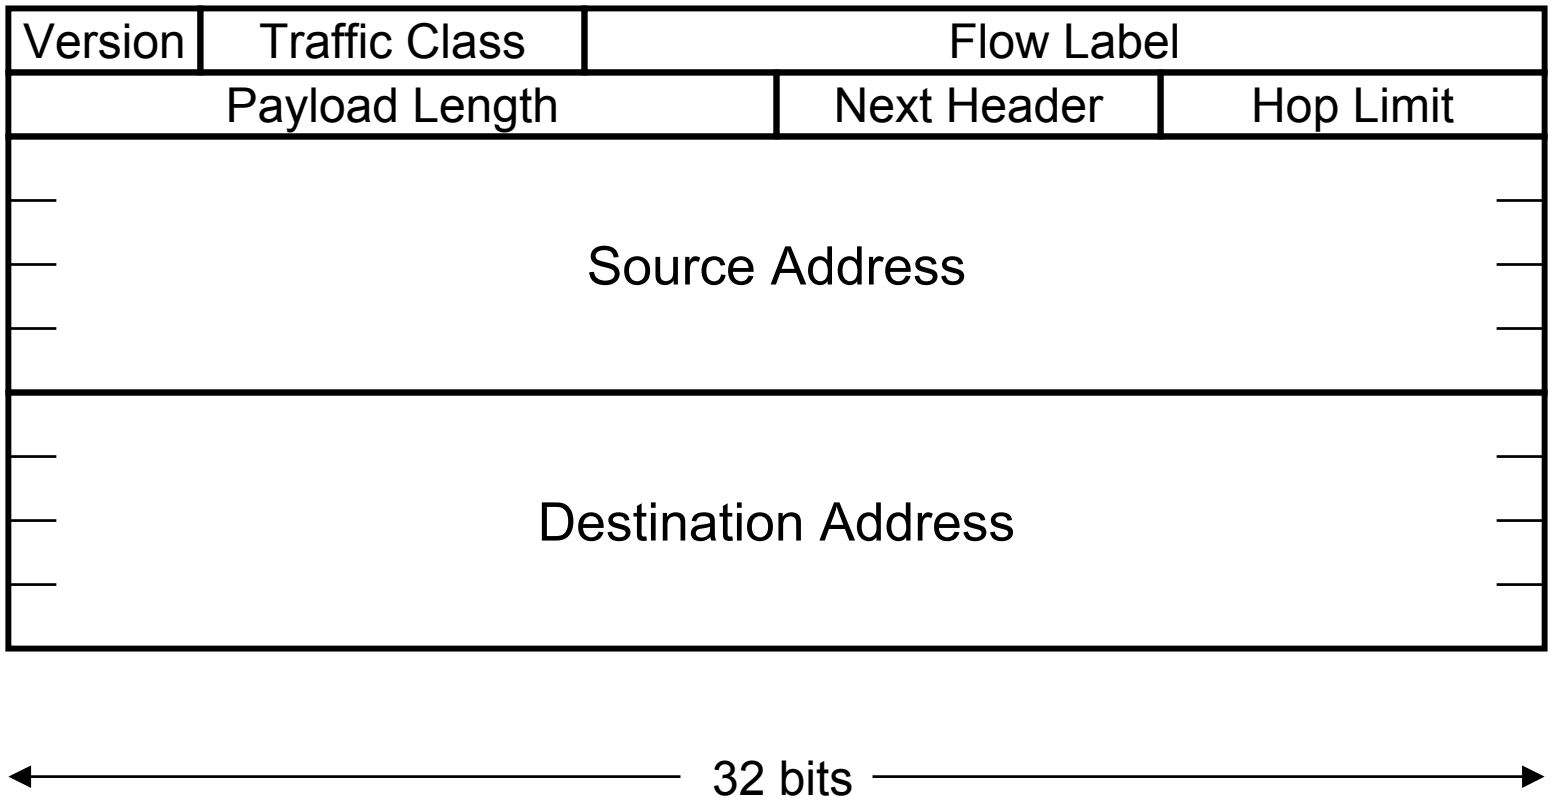

credit: Steve Deering

The IPv4 Header

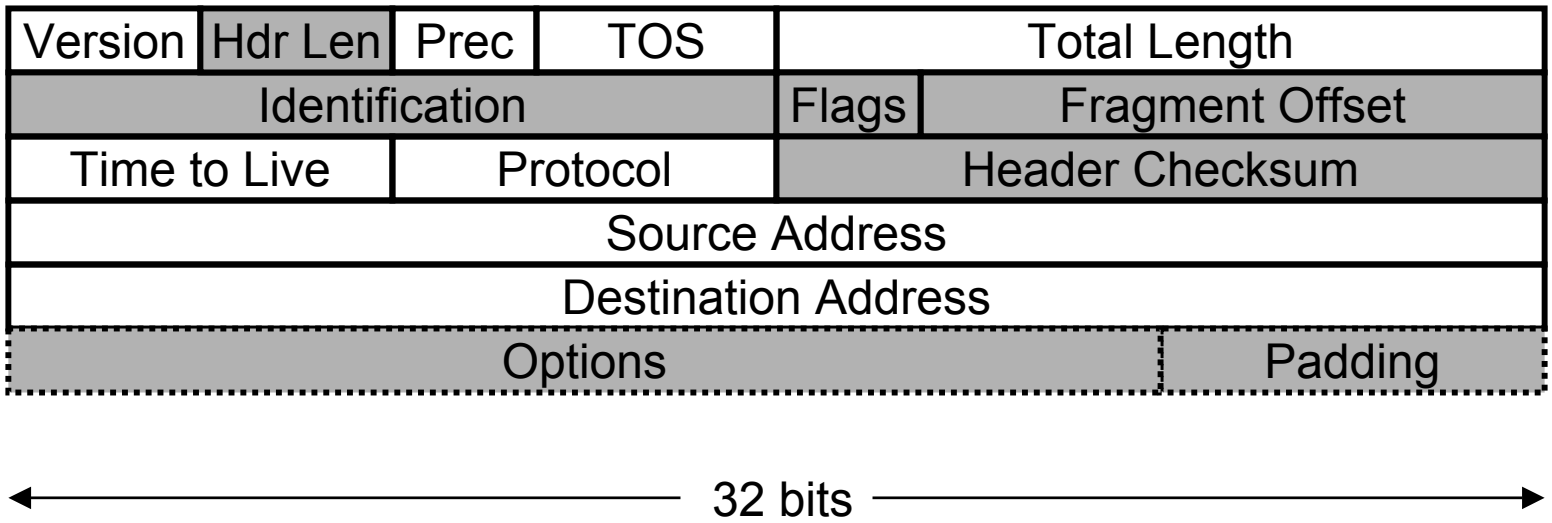

Shaded fields are absent from IPv6 header

credit: Steve Deering

Extension Headers

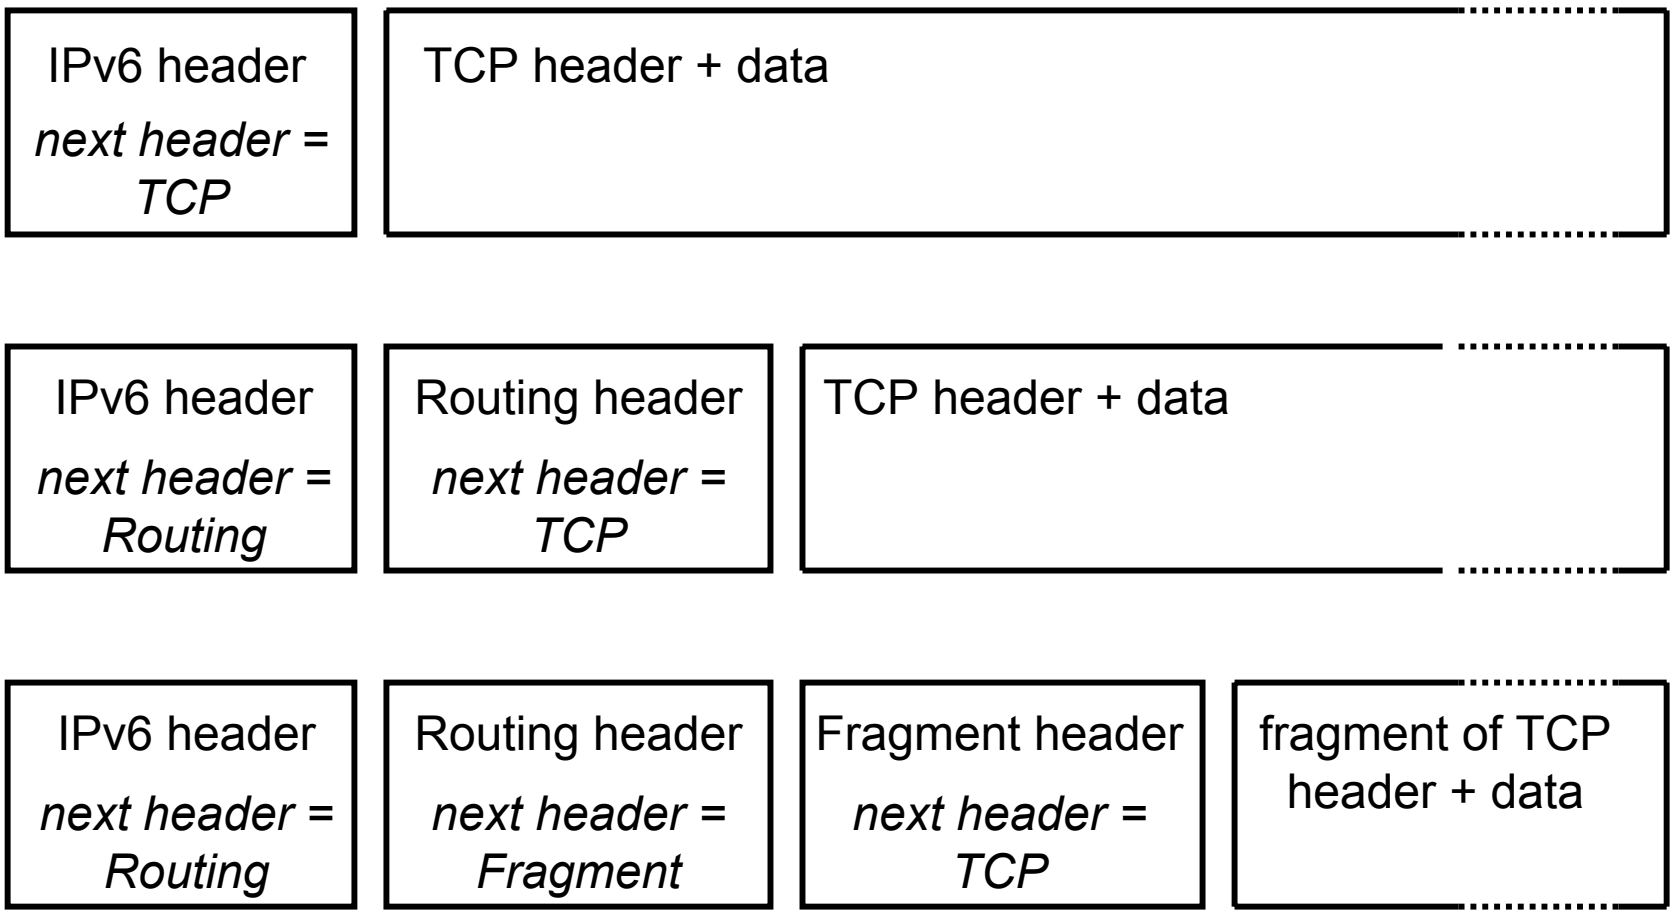

credit: Steve Deering

Global Unicast Addresses

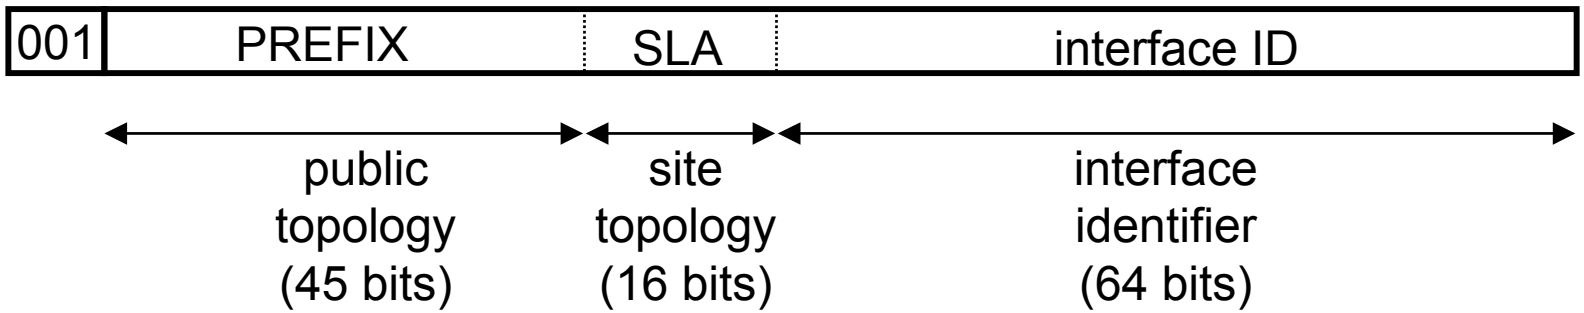

- Prefix ranges may be assigned to providers or exchanges
- Recommended that all sites including homes get 48 bit prefixes (35,184,372,088,832 are available)
- SLA = Site-Level Aggregator (subnet prefix)
- Subfields variable-length, non-self-encoding (cf CIDR) ➔much better route aggregation than legacy IPv4

credit: Steve Deering

Coexistence mechanisms (simple version)

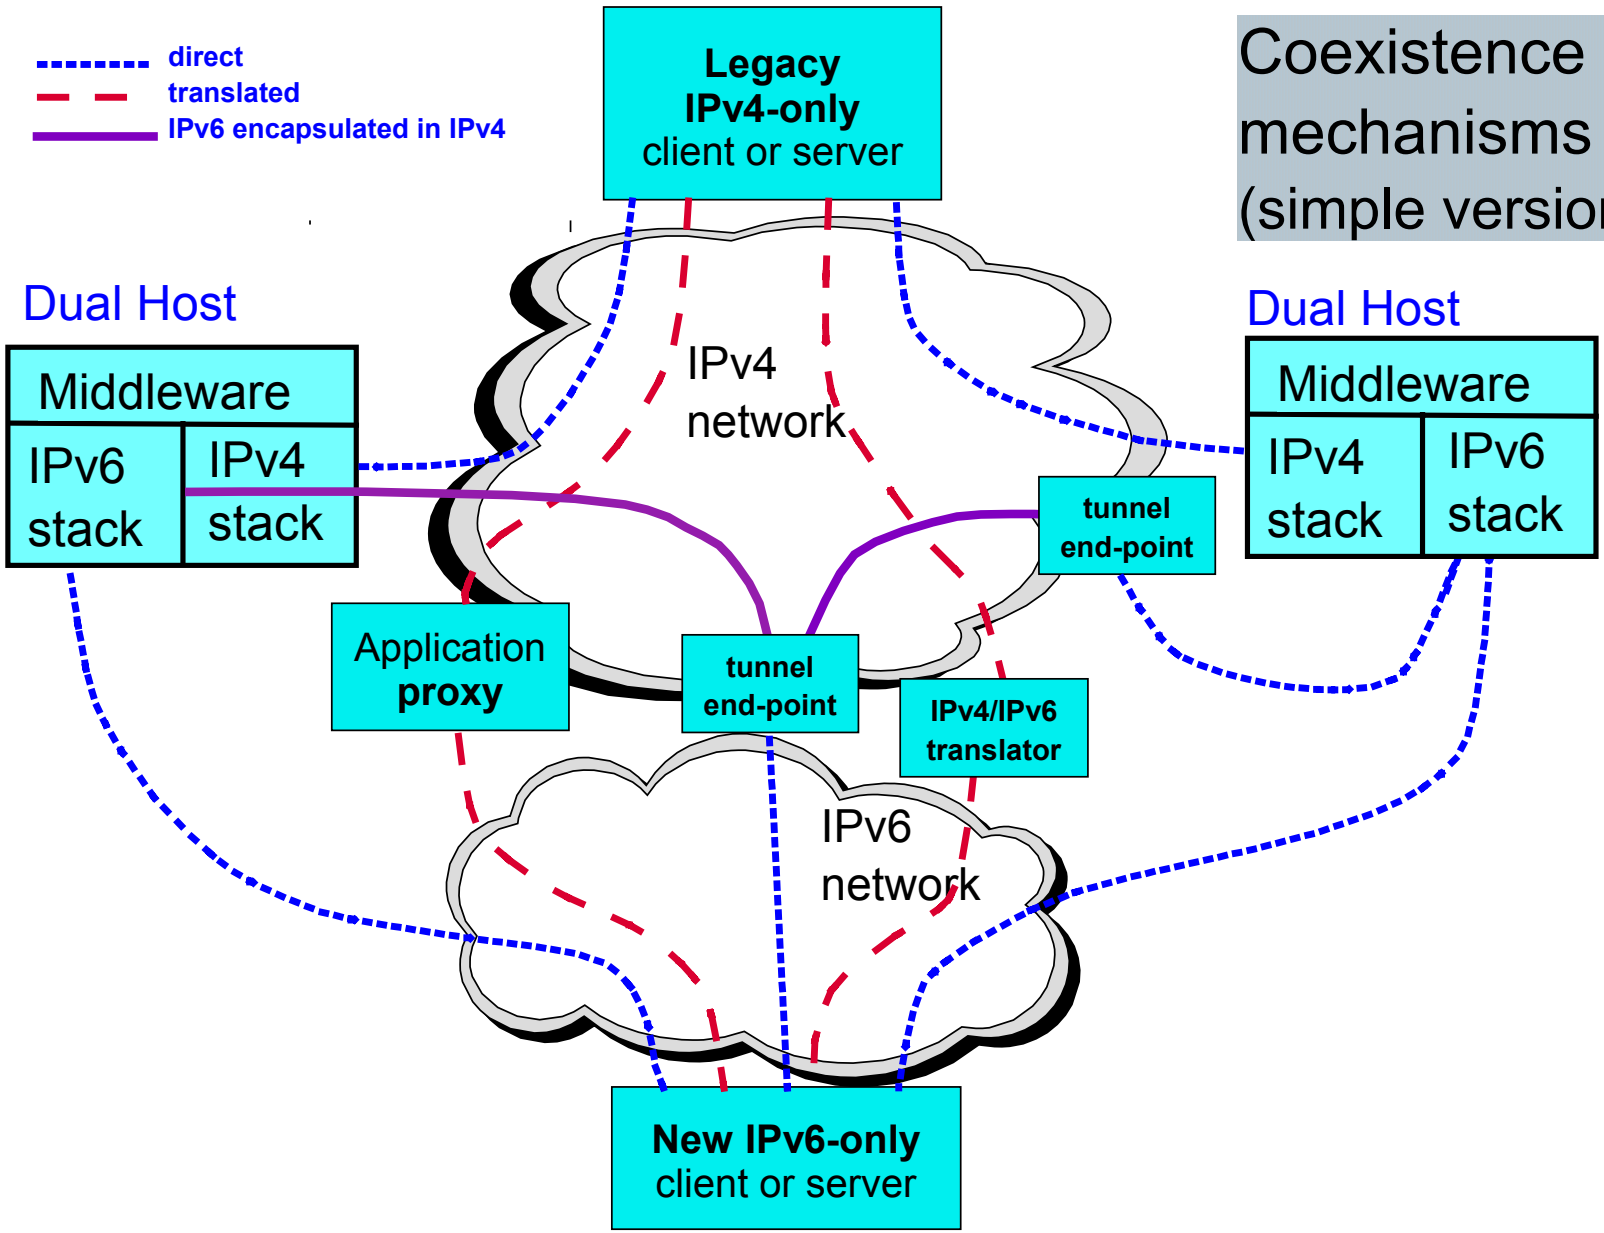

Example coexistence cases from 3G phone world

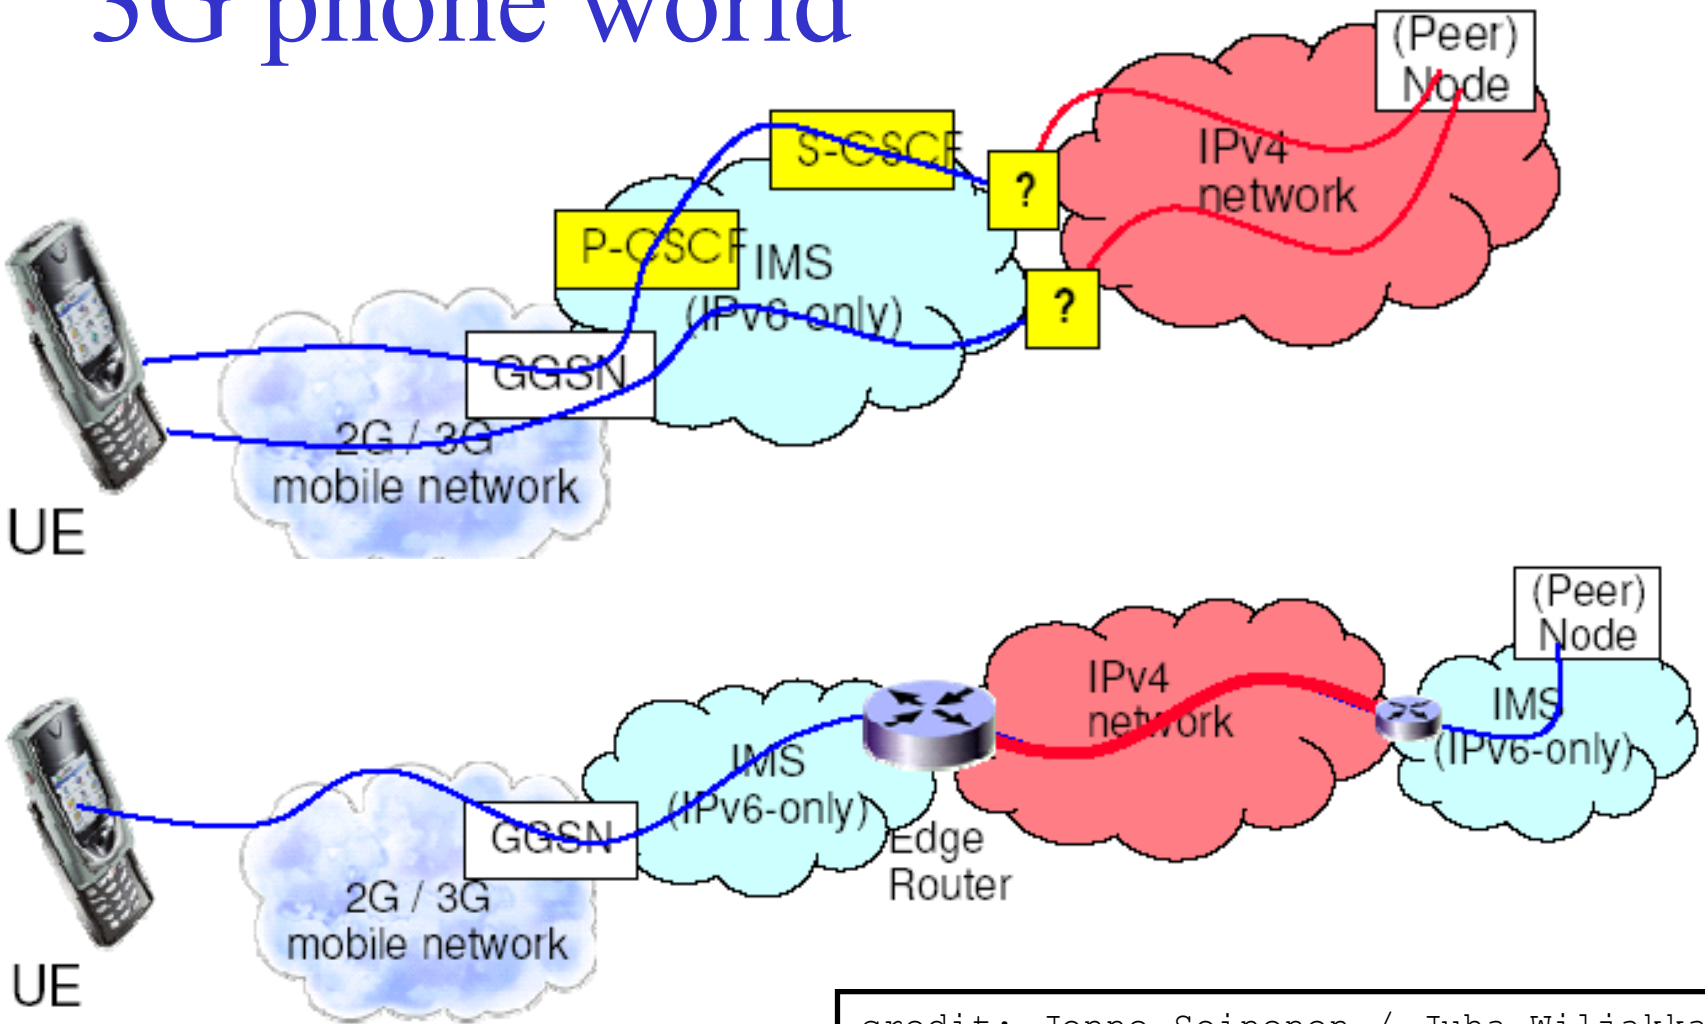

credit: Jonne Soinenen / Juha Wiljakka

A few words about DNS

- Dual-stack DNS needs careful thought.
- Need to resolve IPv6 queries over IPv4, and vice versa.
- If a host has an IPv4 address and a few IPv6 addresses, a DNS query should return several answers.
- Which one should we try?
- Getting this right remains tricky

Standards status

- Basic standards for the protocol, auto-configuration, mobility, socket API, DNS, and coexistence mechanisms are done.
- IETF work continues on
  - site multihoming
  - address space for disconnected sites
  - coexistence scenarios
  - dependencies within other IETF protocols
  - endless refinements
- IPv6 is required by 3GPP standards

Implementation status

- All significant operating systems and router vendors now support dual IPv4/IPv6 stacks and socket APIs
- BIND DNS, PowerDNS, djbdns support IPv6
- Java 1.4 supports IPv6
- Many public domain applications support IPv6
- The conversion of commercial applications is beginning

Deployment status (1)

- Multiple R&D IPv6 testbeds running around the world
- Numerous commercial IPv6 services on offer, but we have a classical chicken/egg deadlock.
- National and EU IPv6 Task Forces starting up.
- Required by 3GPP
- Emerging requirement in RFPs; US DoD requirement from 10/03.

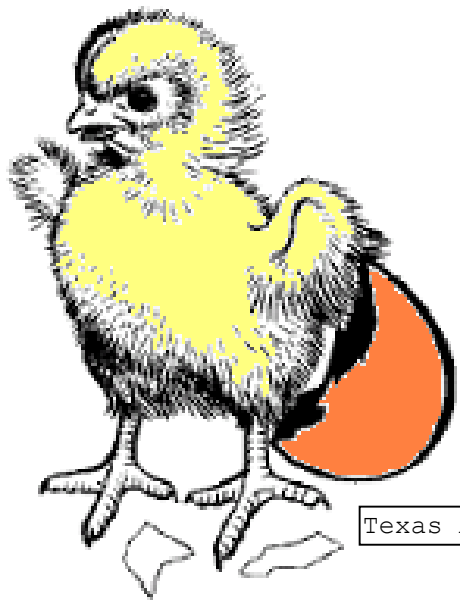

Deployment status (2)

- About 350 “production” IPv6 prefixes allocated, which mainly belong to ISPs.
  - Hard to know how many offer commercial IPv6 (certainly at least 25, of which ~10 in Japan)
  - Remember that customer prefixes are mainly aggregated behind ISP prefixes
  - Connectivity is real, see http://net-stats.ipv6.tilab.com/bgp/bgp-page-otherIANA.html

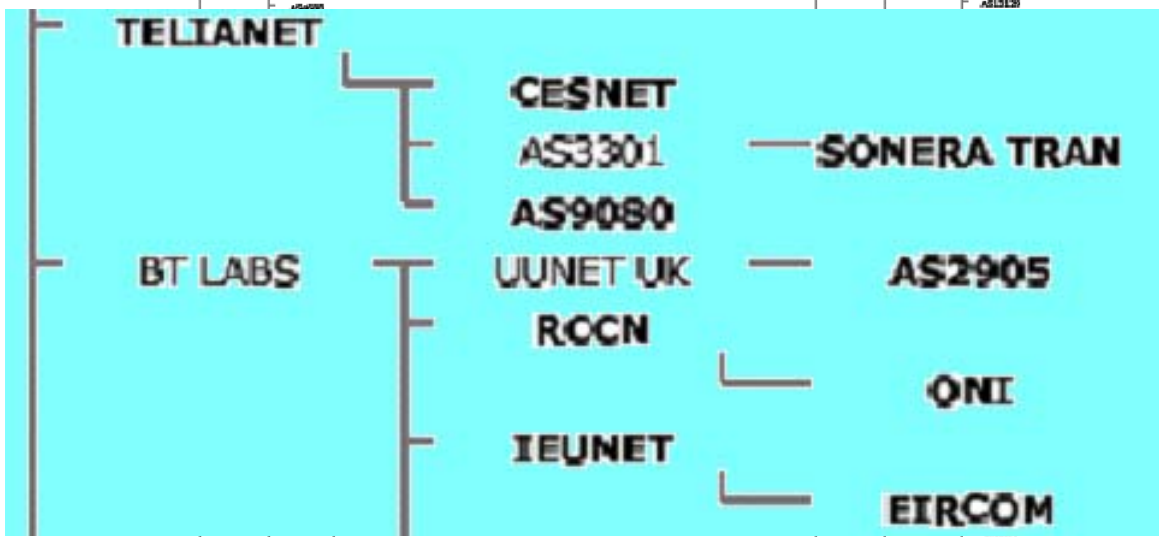

Active IPv6 topology, 2003-04-09.

Much better route aggregation than IPv4.

IPv6 is real!

Topics

- General Introduction
- Brief introduction to IPv6
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

Grid Deployment Options

A function of business need, technology and organizational flexibility

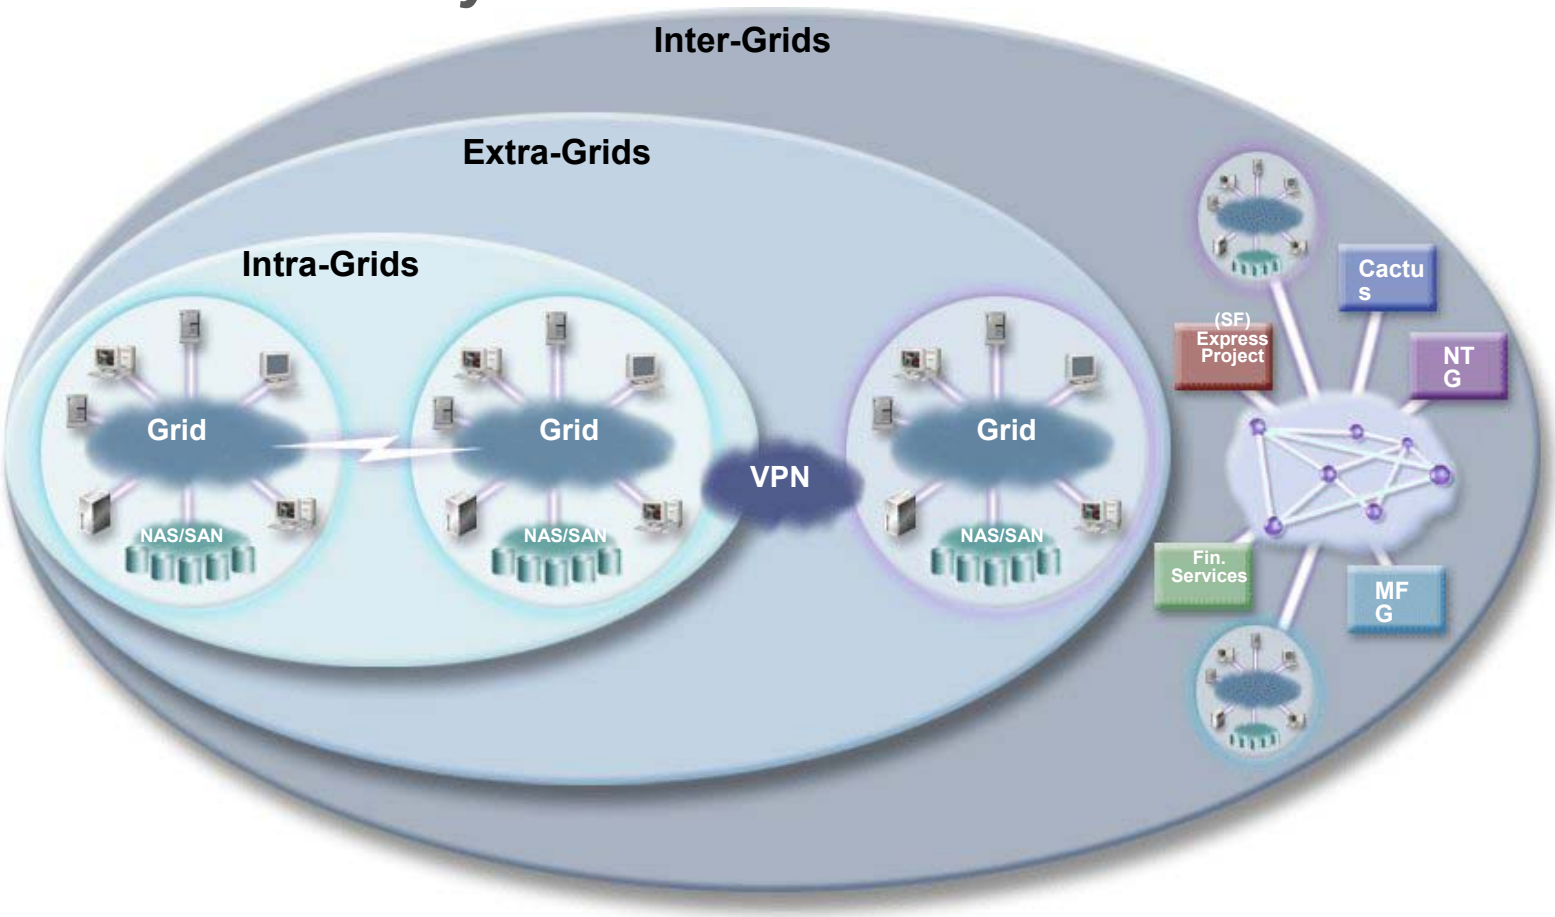

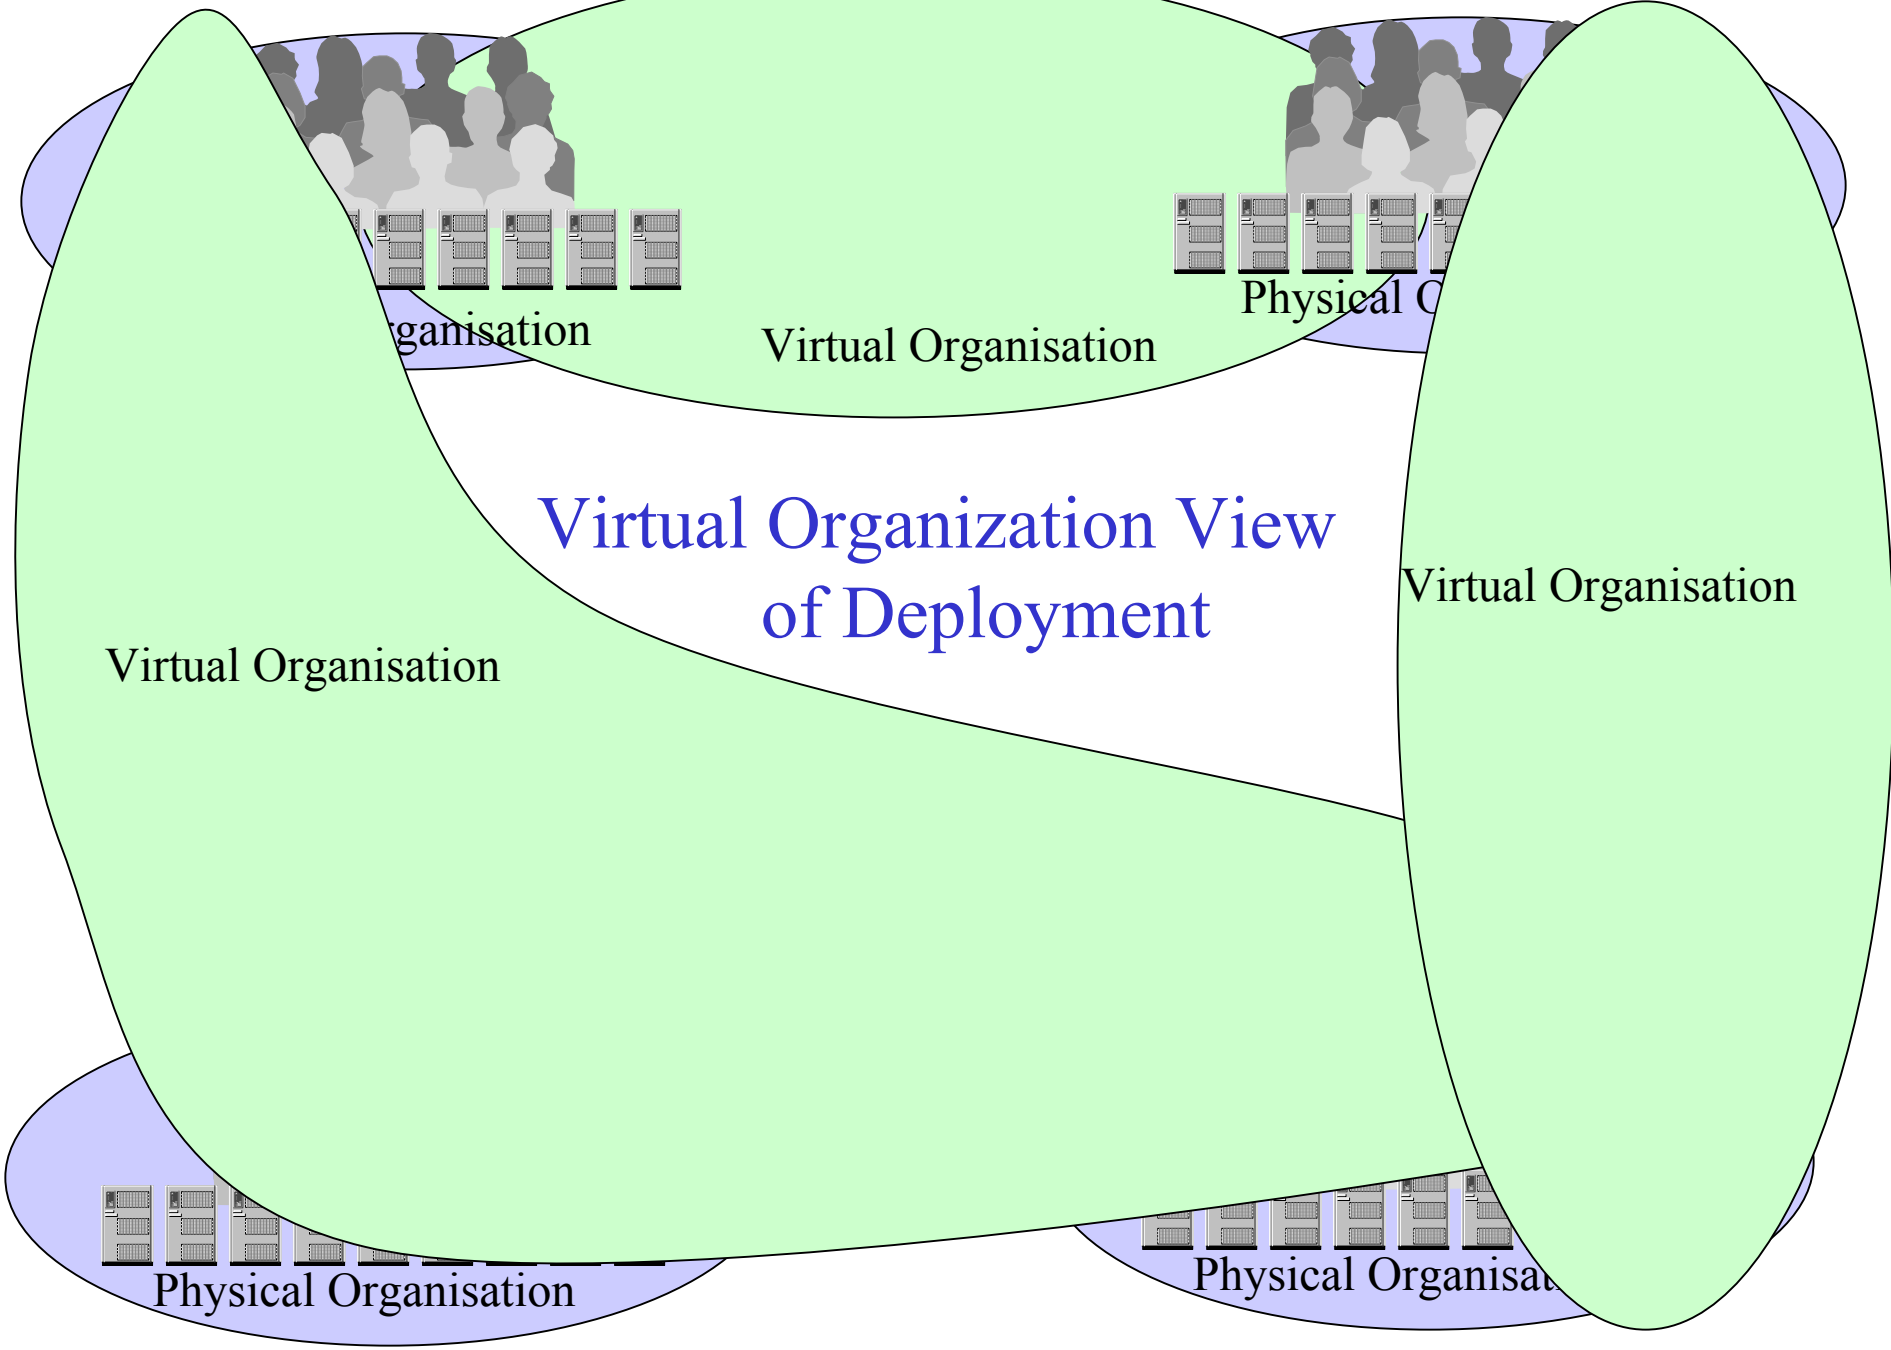
ganisation
Virtual Organisation
Physical C
Virtual Organization View of Deployment
Virtual Organisation
Virtual Organisation
Physical Organisation
Physical Organisat

Virtual organizations look like dynamic mergers & acquisitions

- The effect of a Grid VO on networks is like a temporary partial merger of the organizations.
- Merging two networks is very painful today:
  - "private" IPv4 address space becomes ambiguous
  - worst case: forced to renumber both networks
- Temporary partial mergers of an arbitrary number of IPv4 networks are unthinkable.

➢ IPv4 based Grids are forced to rely on HTTP proxying between organizations: inefficient, and cannot exploit network level security.

Overlapping virtual organizations

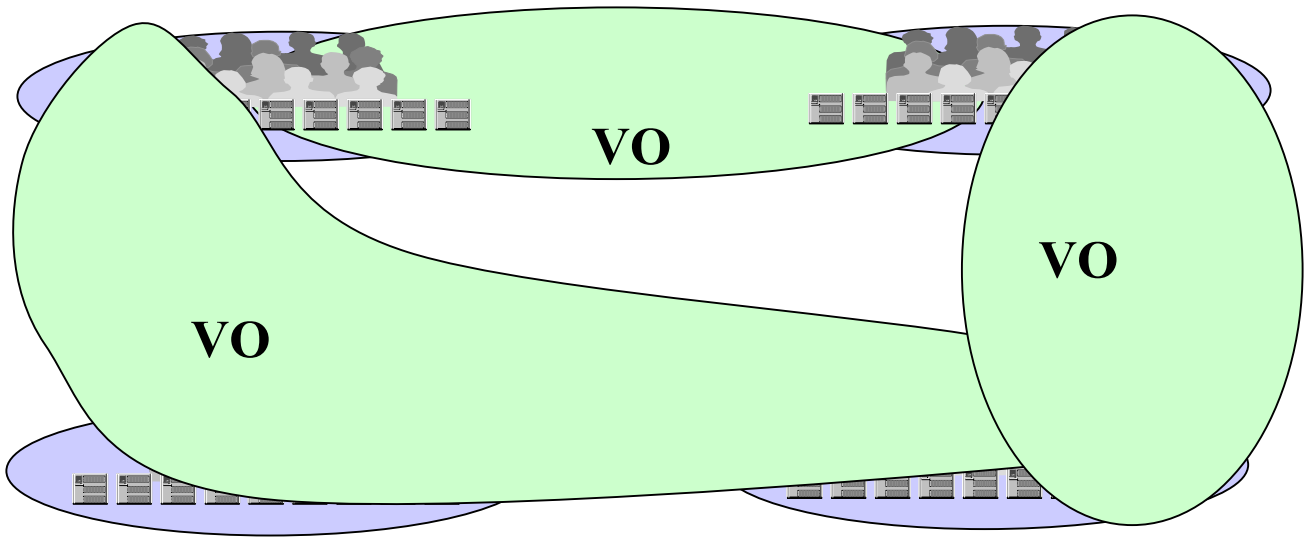

- Any system can be in any number of VOs with any number of other systems
  - needs uniform address space to avoid proxies & allow IPSEC
  - addressing ambiguities unacceptable
  - security boundaries ≠ organization boundaries
  - **can't meet these constraints at massive scale with IPv4**

Critical advantages of IPv6 for OGSA

- Potential for massive scaling
- Uniform global address space eliminates the problem of ambiguous “private” addresses and network address translation.
  - Wasteful proxies can be avoided
  - Network level security can be used
- Autoconfiguration is a big plus for infrastructure configuration

There's no such thing as an IPv6 killer app, but...

- It would be nice to find the killer app that only works on IPv6.
- OGSA won't be that, but there is a good chance that it will be the first major middleware suite to be IPv6-capable out of the box almost from Day One.
- The Grid community should make the most of it.

Topics

- General Introduction
- Brief introduction to IPv6
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

Globus Toolkit v3 (GT3) Open Source OGSA Technology

- Implements OGSI interfaces
- Supports primary GT2 interfaces
  - High degree of backward compatibility
- Multiple platforms & hosting environments
  - J2EE, Java, C, .NET, Python
- New services
  - SLA negotiation, service registry, community authorization, data management, …
- Rapidly growing adoption and contributions: “Linux for the Grid”

6net

Testing, testing

- 6NET is a three-year European Union funded project to demonstrate that continued growth of the Internet can be met using IPv6.
- It includes a work package for ***IPv6 Middleware and User Application Trials*** (led by IBM).
- Globus will be the subject of a trial (lead site: UCL)
  - Target is Globus Toolkit 3, i.e. OGSA
  - GT3 (OGSA) alpha code is now available and being tested on IPv6/Linux at UCL
  - *Credits: Sheng Jiang, Piers O'Hanlon, Peter Kirstein*

Status (early May, evolving daily)

- Background: GT3 runs mainly over Java; Java2 Development Kit 1.4 supports IPv6, although GT3 alpha ships with JDK 1.3.
- In principle, all components of GT3 implemented in Java will "just work" with IPv6 by switching to JDK 1.4
  - In practice, it is not so simple.
  - For C code, need new Globus XIO module
- UCL has a GT3 testbed and has begun testing with JDK 1.4.
  - The GT3 Master Host Environment listens on both IPv6 and IPv4. An elementary test initiated from IPv6 completes OK, but some IPv4 packets are observed.
- Exact environment is GT3 alpha code with Java SDK 1.4.1. on Redhat Linux 7.3 and 8.0.

Status (update May 6th)

- Tested the postgresql IPv6 patch; found one bug there: after removing IPv6 items from the configuration file, the IPv6 address was still enabled.
- Working to deploy GT3 core on Apache Tomcat Java servlet container (IPv6 enabled).
- Waiting for IBM Websphere to do the same.
- Starting to port the OGSA stand-alone web container to be IPv6-enabled
  - *Sheng Jiang*

Status (update June 4th)

- Testing on an IPv6-only network started
- Still see some IPv4 traffic *within* the hosts, related to JDBC and Gatekeeper
- Will report details to Globus Bugzilla
  - (June 25th) 1 fix and 3 issues reported so far

– *Sheng Jiang*

GridFTP (evolving daily)

- GridFTP (striped FTP) as shipped with GT3 alpha is C code that does not support IPv6 sockets.
- Globus is developing a new generic I/O module called XIO that does support IPv6 sockets.
- GridFTP is being rewritten, still in C, to exploit XIO and IPv6 extensions to basic FTP.
  - UCL will test XIO alpha code shortly.
- Any other components of GT3 that remain in C can also support IPv6 via XIO.

6net

Further plans (evolving daily)

- Plan is to make more extensive tests with successive GT3 alphas, with about 10 nodes
  - Issues with IPv6 will be reported into the Globus bug-tracking system
  - Good relations established between 6NET and Globus teams
- Also need to consider what is required to operate GT3 in the cases of:
  - IPv6 only
  - IPv6 and IPv4 coexistence
- Final goal is a realistic systematic trial between 6NET sites

Topics

- General Introduction
- Brief introduction to IPv6
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

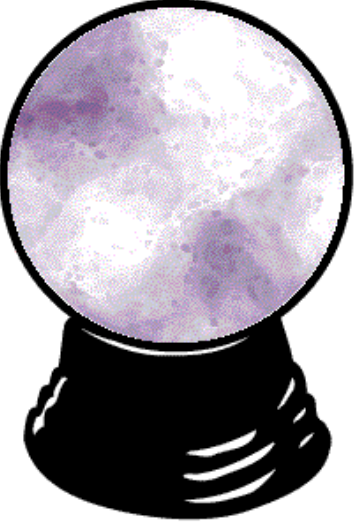

Future Directions

- GT3 will become fully functional for IPv6.
- Global Grid Forum needs to chase down any IPv4 dependencies in its standards.
- Grid computing will become a key enabler of new applications based on resource virtualization and loosely coupled distributed service integration.
- IPv6 will enable Grid Virtual Organizations to span existing network boundaries smoothly and securely.

IBM®

Summary

- OGSA is the key model for sharing networked IT resources efficiently & securely.
  - Transforming the Internet into a computing platform for e-business on demand
- Deployment of IPv6 is the key technology for continued growth of Internet coverage outside its traditional range.
- A marriage between OGSA and IPv6 is the key to massive scaling in a fully connected but secure network environment.

6net

the globus project™
www.globus.org

IBM

Pointers

www.globus.org

www.gridforum.org

www.ipv6forum.org

6net

brian@hursley.ibm.com
foster@mcs.anl.gov