Challenge and response: Towards tomorrow's Internet

Brian E Carpenter
Distinguished Engineer
Internet Standards & Technology
IBM

Internet Society

Amsterdam, November 2003

Acknowledgements

❖Material used from
- Ian Foster (Argonne National Lab)
- Steve Deering (Cisco)
- Geoff Huston (Telstra)
- IBM colleagues
- IETF colleagues
- 6NET colleagues

Topics

- The Internet today: as far as Web Services
- The Internet tomorrow: a services platform for computing on demand
- Challenges at the network level (transparency, addressing, routing, multihoming)
- Challenges at the middleware level (service architecture, heterogeneity, security, integrity)
- Challenges outside the technology
- Releasing known potential: beat the challenges
- Summary

The Internet Today

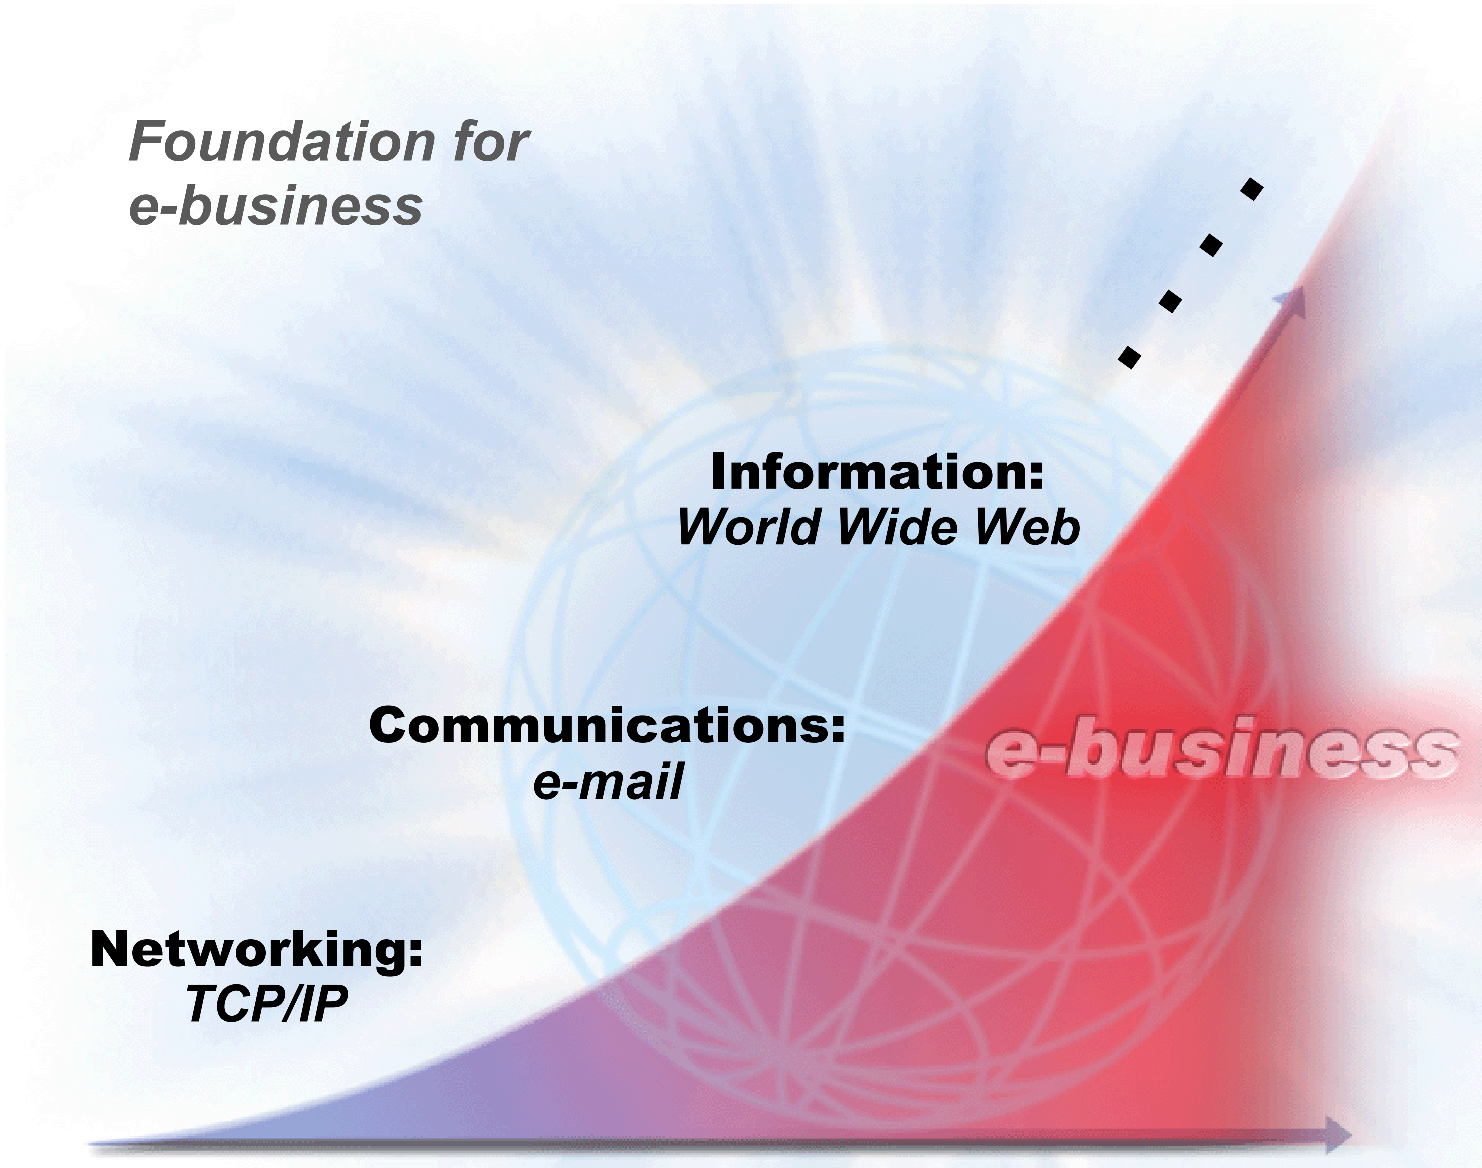

What we really have

- An information web – the normal mode is for clients (users) to suck down bits from a server, like young birds in a nest suck down food from their parents.
  - Using the web to do stuff (buy, sell, play, work) is still somewhat the exception.
  - Using the web on the move is still the exception.
  - Fully trusting the web is still the exception.
- ***Web Services*** are just starting

(Thanks to birds.cornell.edu)

Web Services: multiparty model

www.ibm.com/software/solutions/webservices/

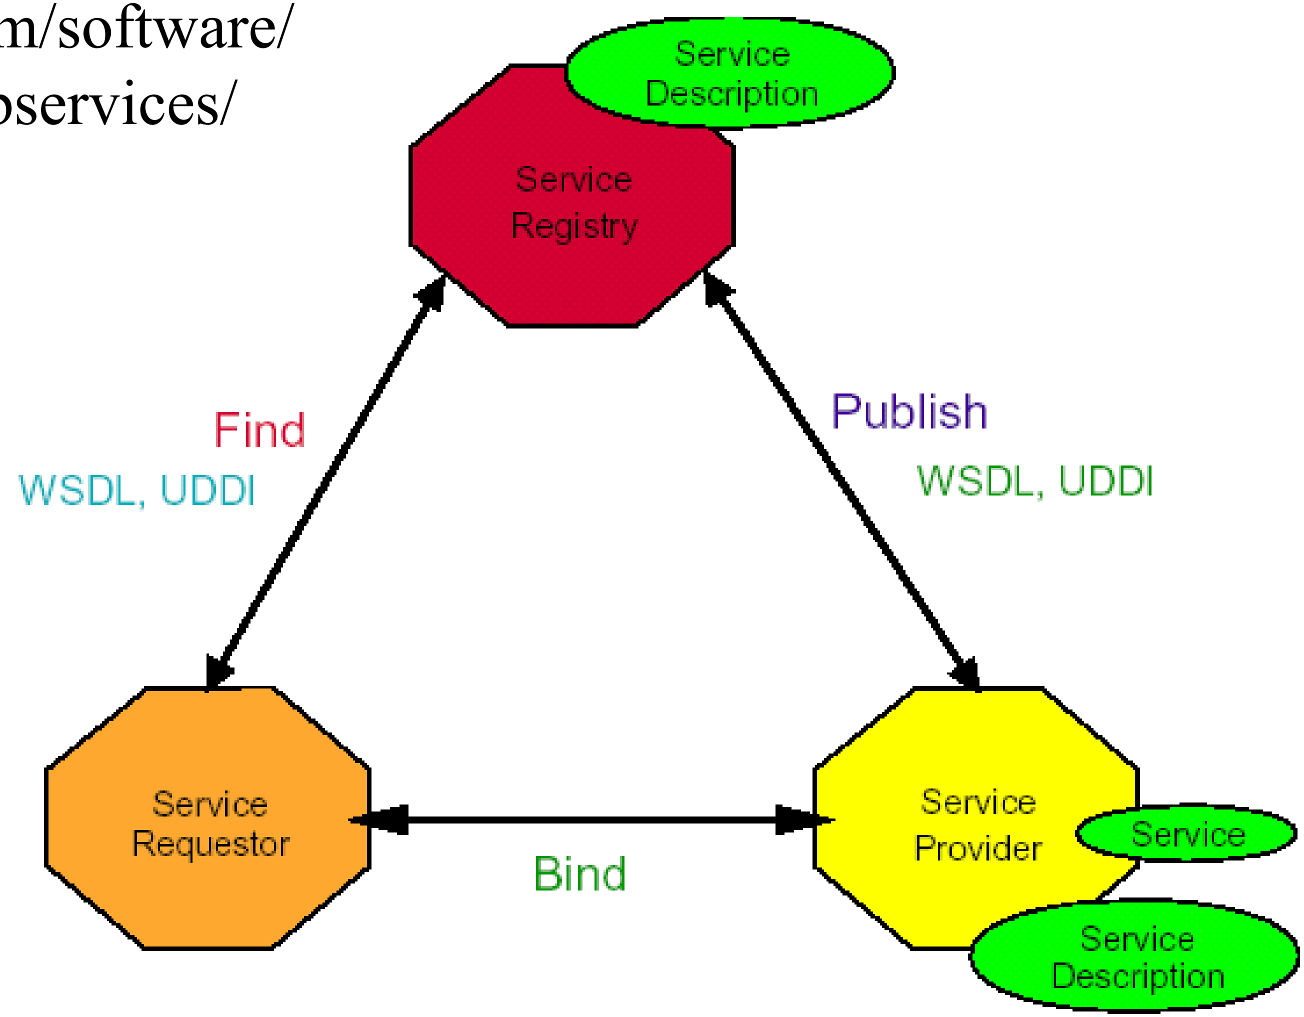

The conceptual Web Services stack

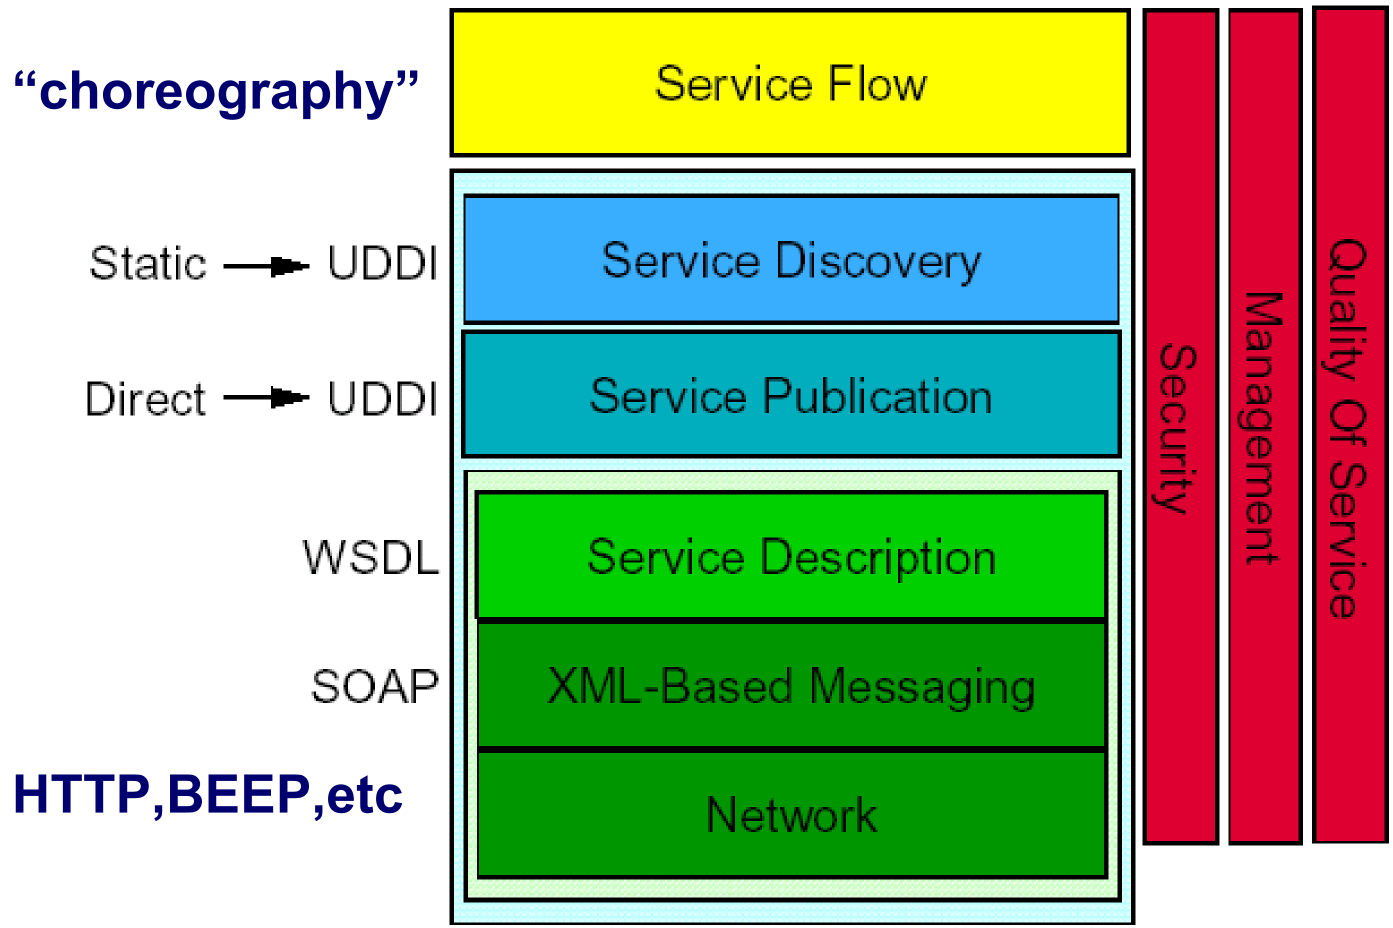

Factors for continued change and growth

- Marketplace requirements
- Technology and the appetite for technology feed on each other
- Internet culture of open standards

Marketplace Requirements

- More efficient use of IT resources
  - Computing, storage, transactions,…
  - Renewed importance of Total Cost of Ownership
  - Chasing out hidden costs
- Industrial strength infrastructure
  - 7x24, secure, robust under attack, disaster recovery
- Integrated, but flexible
  - Distributed, centralized, outsourced..
- Impatient consumers
  - Fast, always on, everywhere, natural, intelligent, easy, and trusted

Enabling IT and Business Value

IT Needs

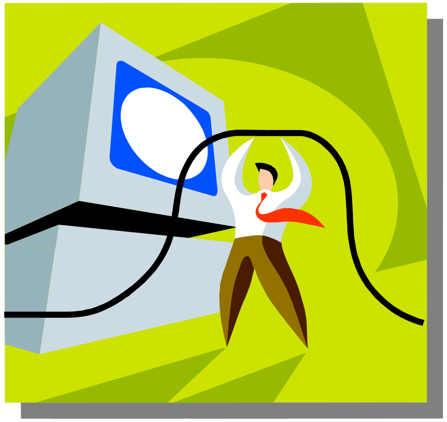

- **Improve Asset Optimization**
- **Integrate Heterogeneous Resources**
- **Enable Data Access, Integration and Collaboration**
- **Strengthen Redundancy and Resiliency**
- **Quickly Respond to Variable Demands**

Business Needs

- **Improve Operating Efficiency/ROI**
- **Reduce Capital Expenses**
- **Accelerate Business Processes**
- **Enhance Enterprise Collaboration**
- **Quickly Adapt to Changing Requirements**

An On Demand Business

An enterprise whose business processes -- integrated end-to-end across the company and with key partners, suppliers and customers -- can respond with speed to any customer demand, market opportunity or external threat

Growth refuses to slow down

- Bandwidth costs can beat Moore's law
- New countries are showing an interest
  - Let's bet on the 10 billion node Internet

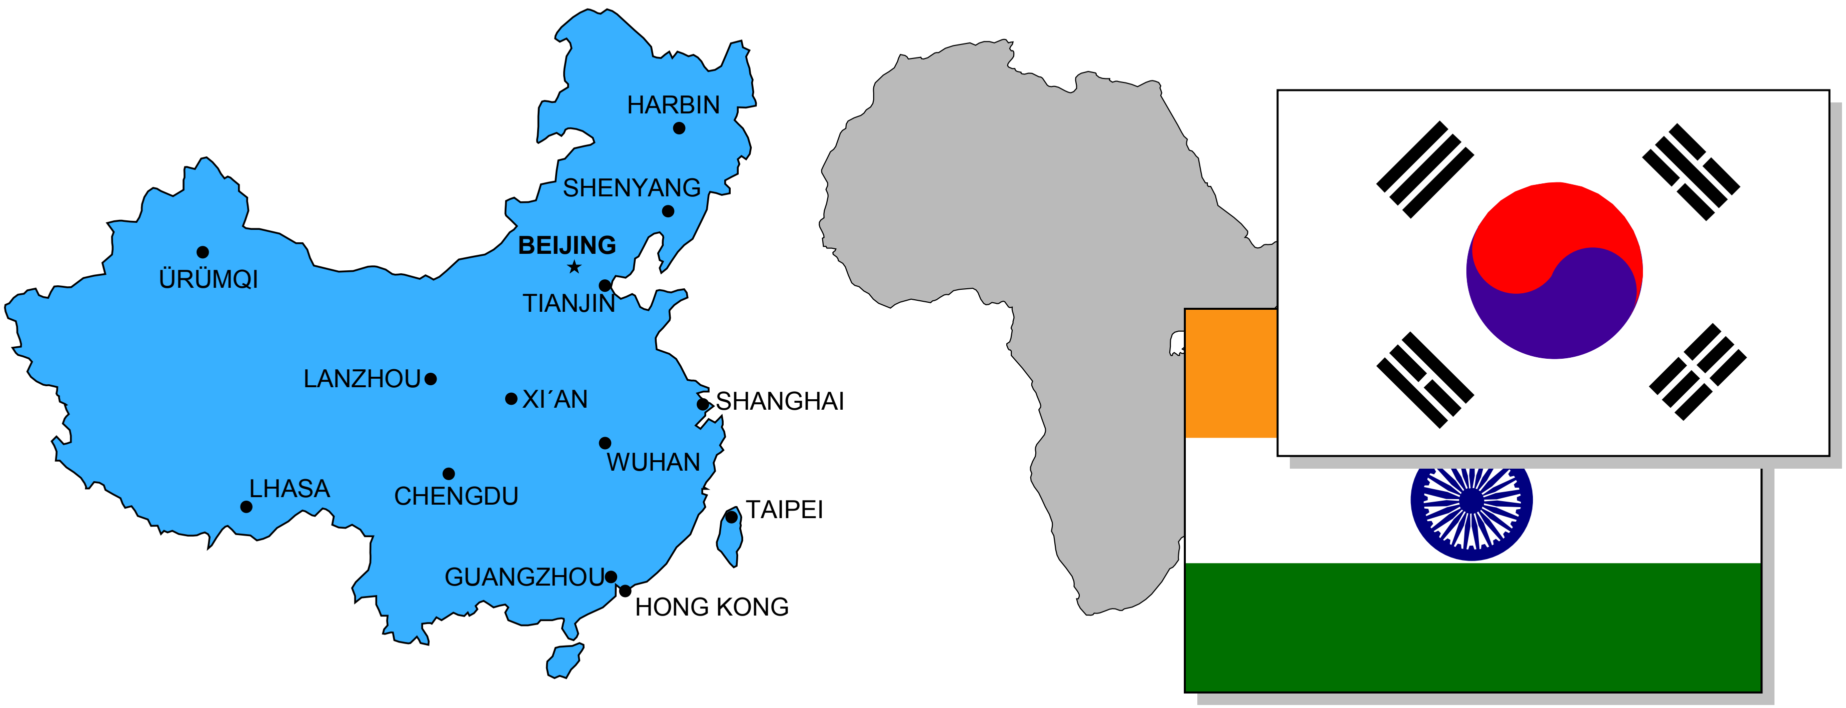

Culture of Standards

Industry trends converge

- Grid computing today is not the same as Web Services, but it was driven in the scientific world by the same forces that drove Web Services for dynamic e-business:
  - *evolving costs*
  - *systems convergence*
  - *resource sharing on the network*
  - *service levels*
  - *security.*

Common eScience/eBusiness Vision

- Link dynamically acquired resources
  - From collaborators, customers, eUtilities, … (members of evolving “virtual organization”)
- Into a “virtual computing system”
  - Dynamic, multi-faceted system spanning institutions and industries
  - Loose coupling of heterogeneous systems
  - Configured ***on demand*** to meet instantaneous needs, for:
- Multi-faceted QoS for demanding workloads
  - Security, performance, reliability, …

Thus: the Internet as a Computing Services Platform

- Building an open infrastructure
  - Web Services + Grid Computing Protocols = ***Open Grid Services Architecture***
- Managing the infrastructure
  - Autonomic Computing, Virtualization
- Using the infrastructure
  - Computing on demand

The Internet Tomorrow

On Demand Operating Environment

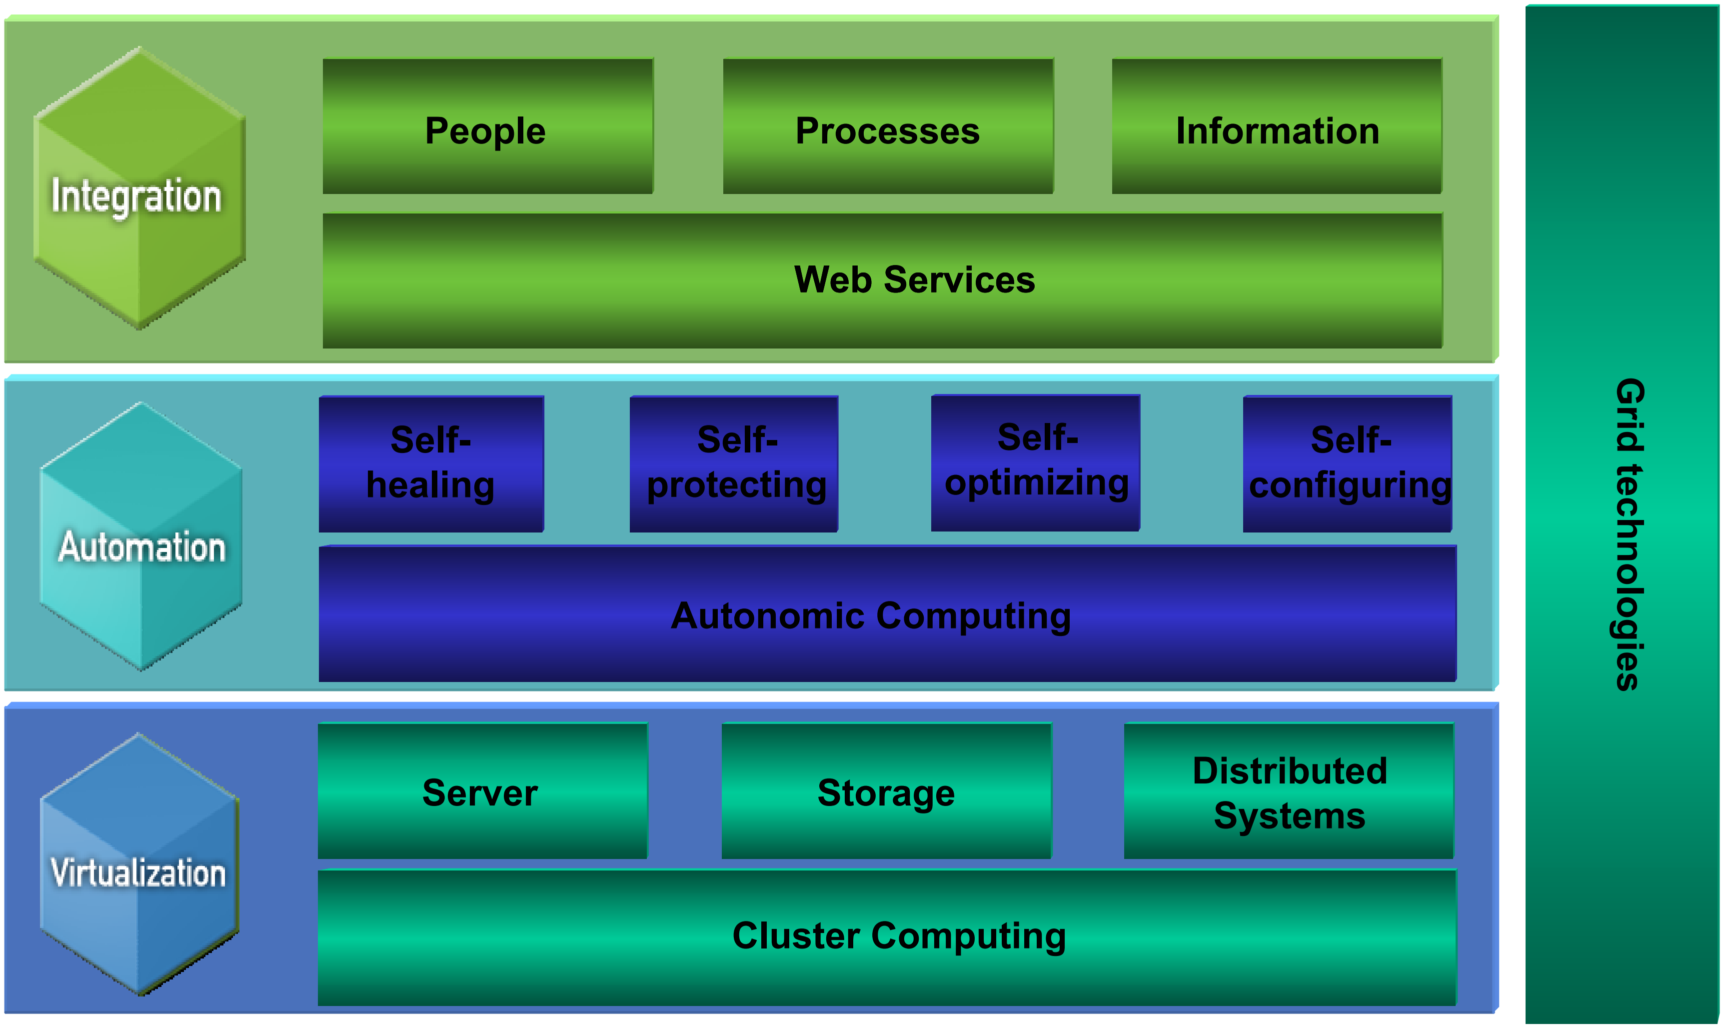

The Grid Is …

a) A collaboration & resource sharing infrastructure for scientific applications
b) A distributed service integration and management technology
c) A disruptive technology that enables a virtualized, collaborative, distributed world
d) An open source technology & community
e) A marketing slogan
f) All of the above

The book…

The Grid model originated with Ian Foster and Carl Kesselman.

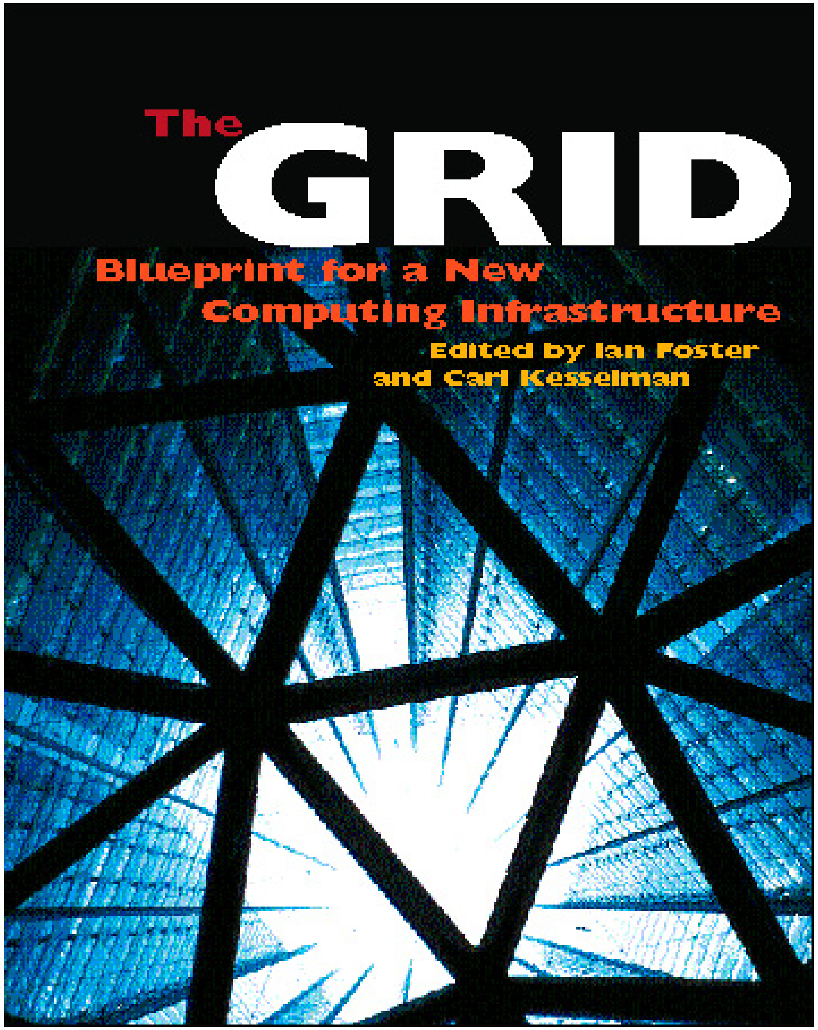

Not Exactly a New Idea …

- “The time-sharing computer system can unite a group of investigators …. one can conceive of such a facility as an … intellectual public utility.”
  - Fernando Corbato and Robert Fano , 1966
- “We will perhaps see the spread of ‘computer utilities’, which, like present electric and telephone utilities, will service individual homes and offices across the country.”
  - Len Kleinrock, 1967

But, Things are Different Now

- Networks are far faster (and cheaper)
  - Faster than computer backplanes
- The Internet has already radically changed the practice of "Computing"
  - Our "computers" have already disintegrated
  - E-commerce increases size of demand peaks
  - Entirely new applications & social structures
- We've learned a few things about software
  - especially that loose coupling and late binding makes for more robust, more flexible distributed systems (at a cost in performance)

But Wait A Minute—Computing isn't Really Like Electricity!

- I import electricity but must export data too
- I can't store unused computing power
- "Computing" is not interchangeable but highly heterogeneous
  - Computers, data, sensors, services, …
- Ok, so the story is more complicated
- But more significantly, the sum can be greater than the parts
  - Real opportunity: Construct new capabilities dynamically from distributed services

⇒ Virtualization & distributed service mgmt

Abstract concept of a computingGrid

- Like public utilities
  - Shared
  - Reliable
  - Running it is someone else's problem
- A computing Grid is a mechanism to "coordinate resource sharing and problem solving in or between physically dispersed virtual organizations (VOs)."
- Assigning resources, users, & applications to VOs is the fundamental Grid technical value proposition.

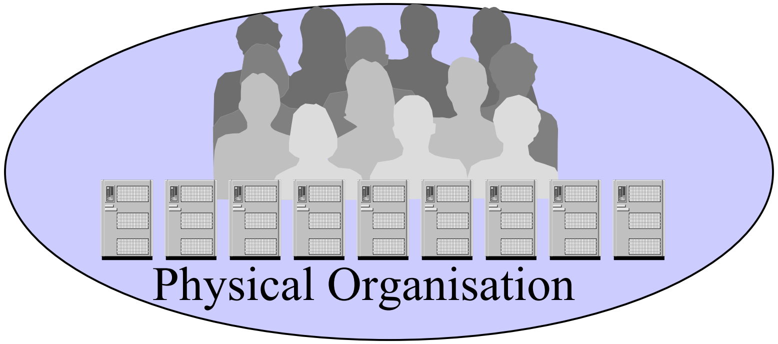
Physical Organisation

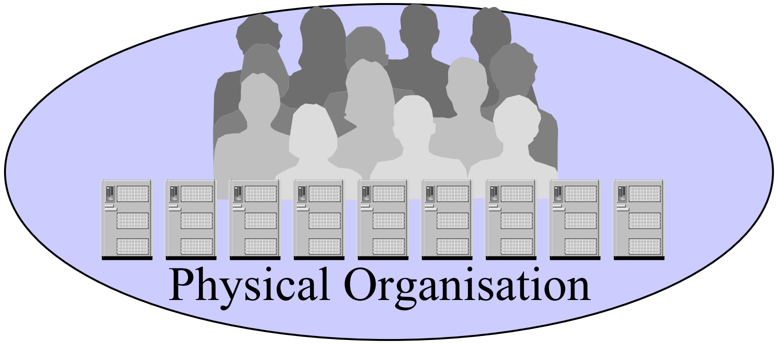
Physical Organisation

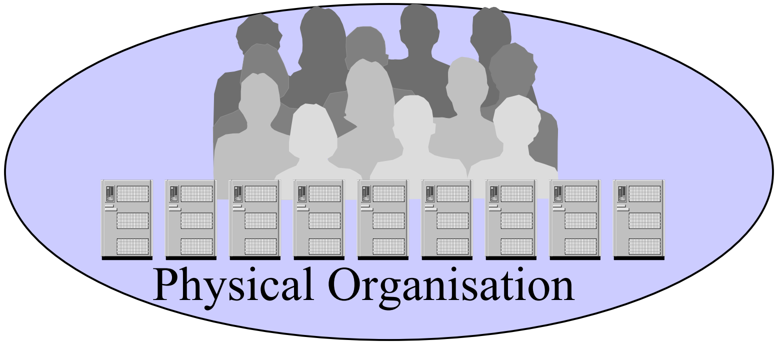
Physical Organisation

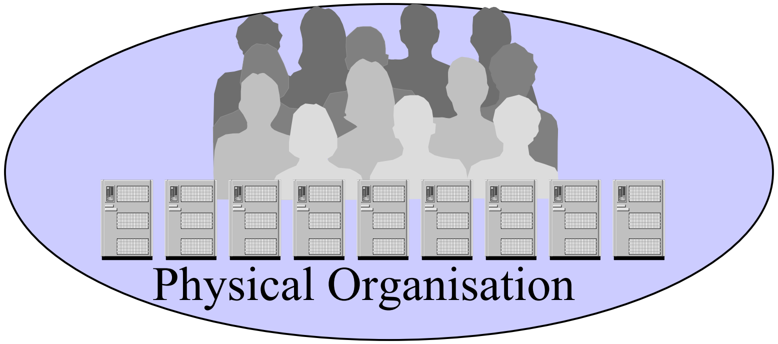
Physical Organisation

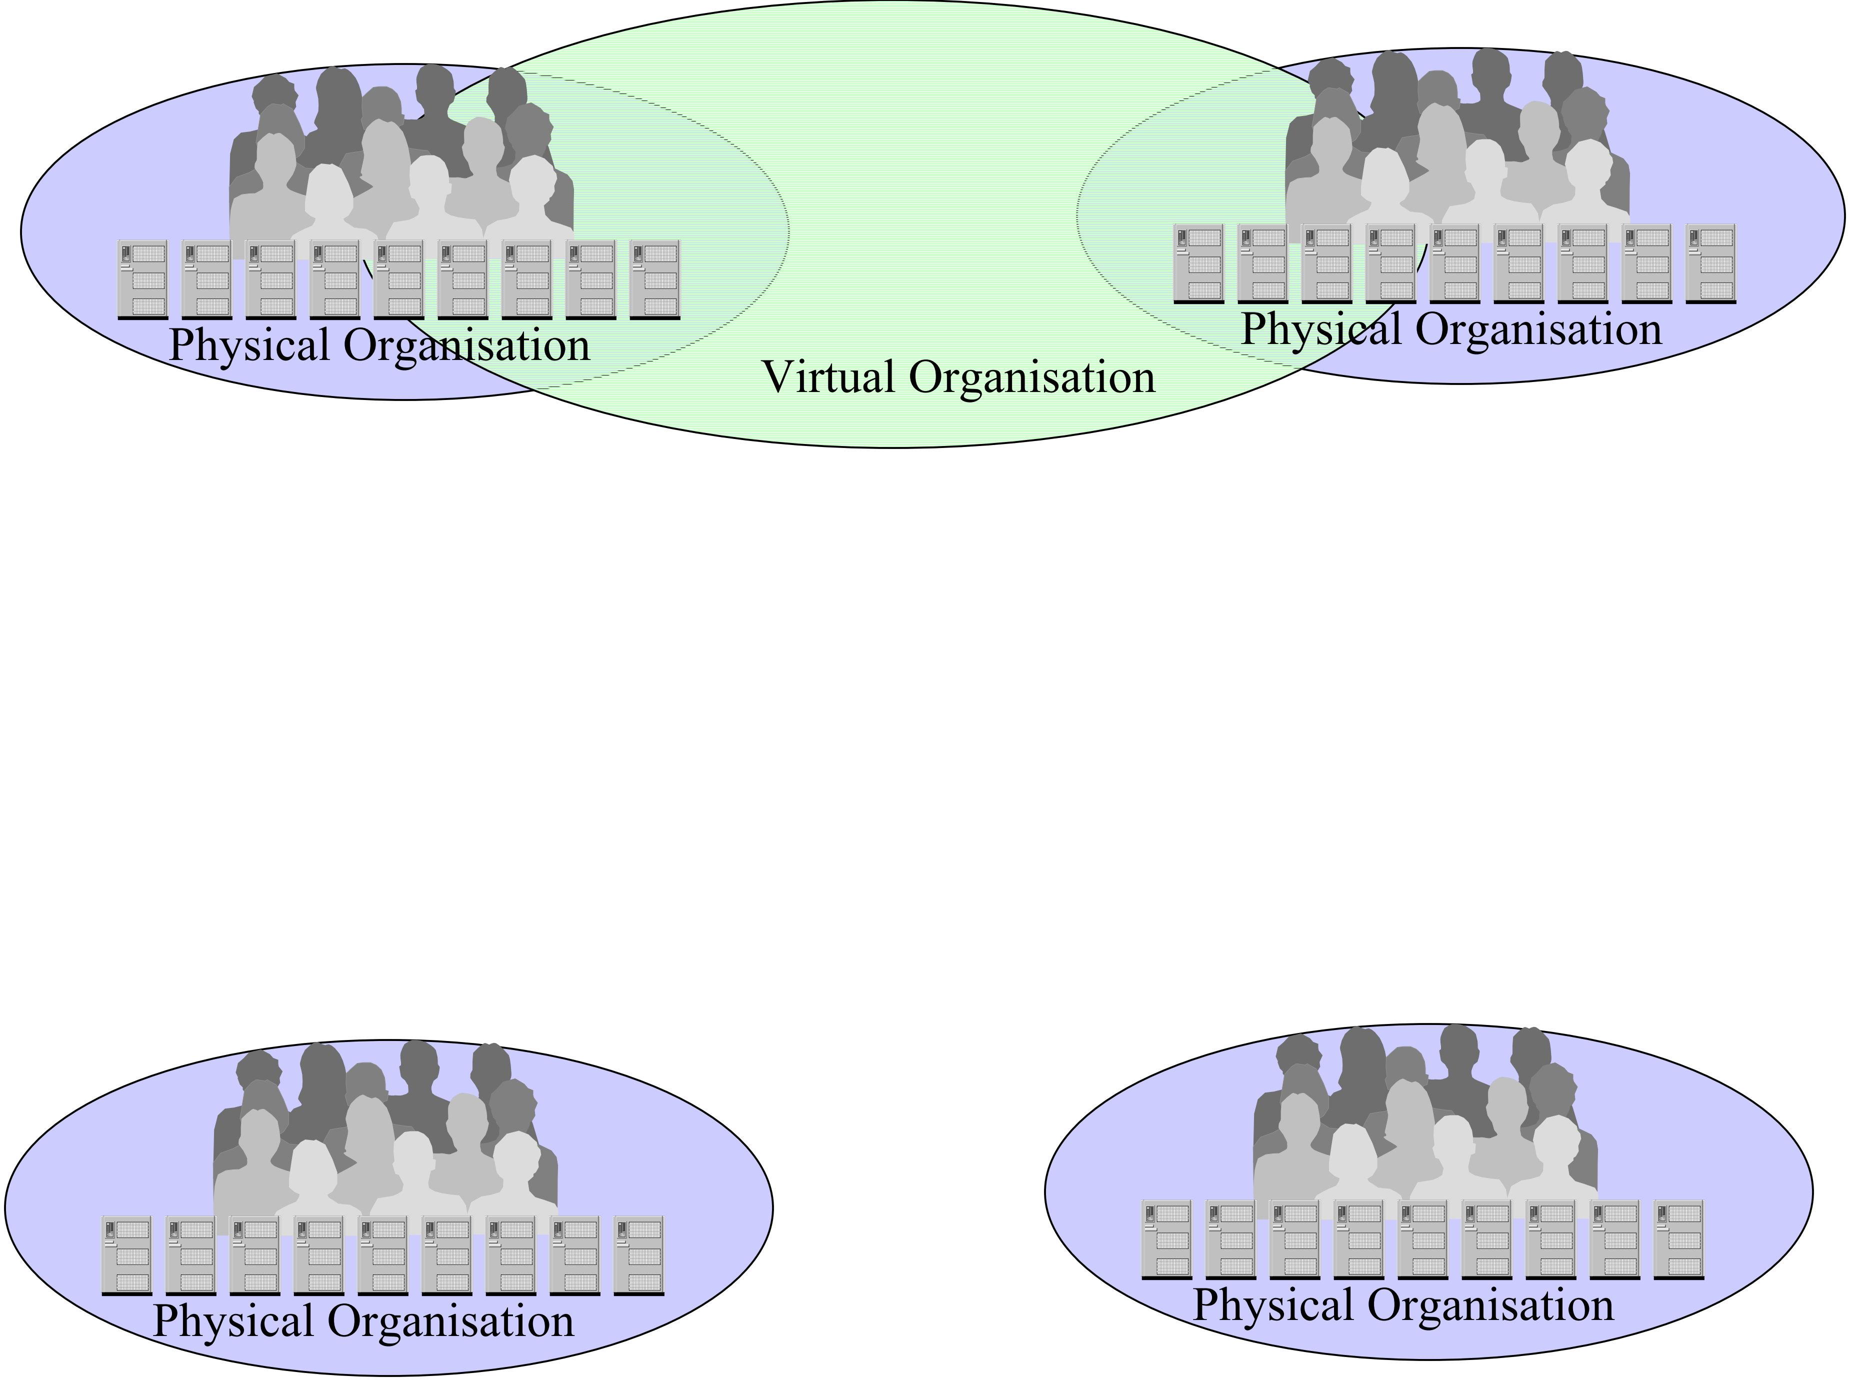
Physical Organisation
Virtual Organisation
Physical Organisation
Physical Organisation
Physical Organisation

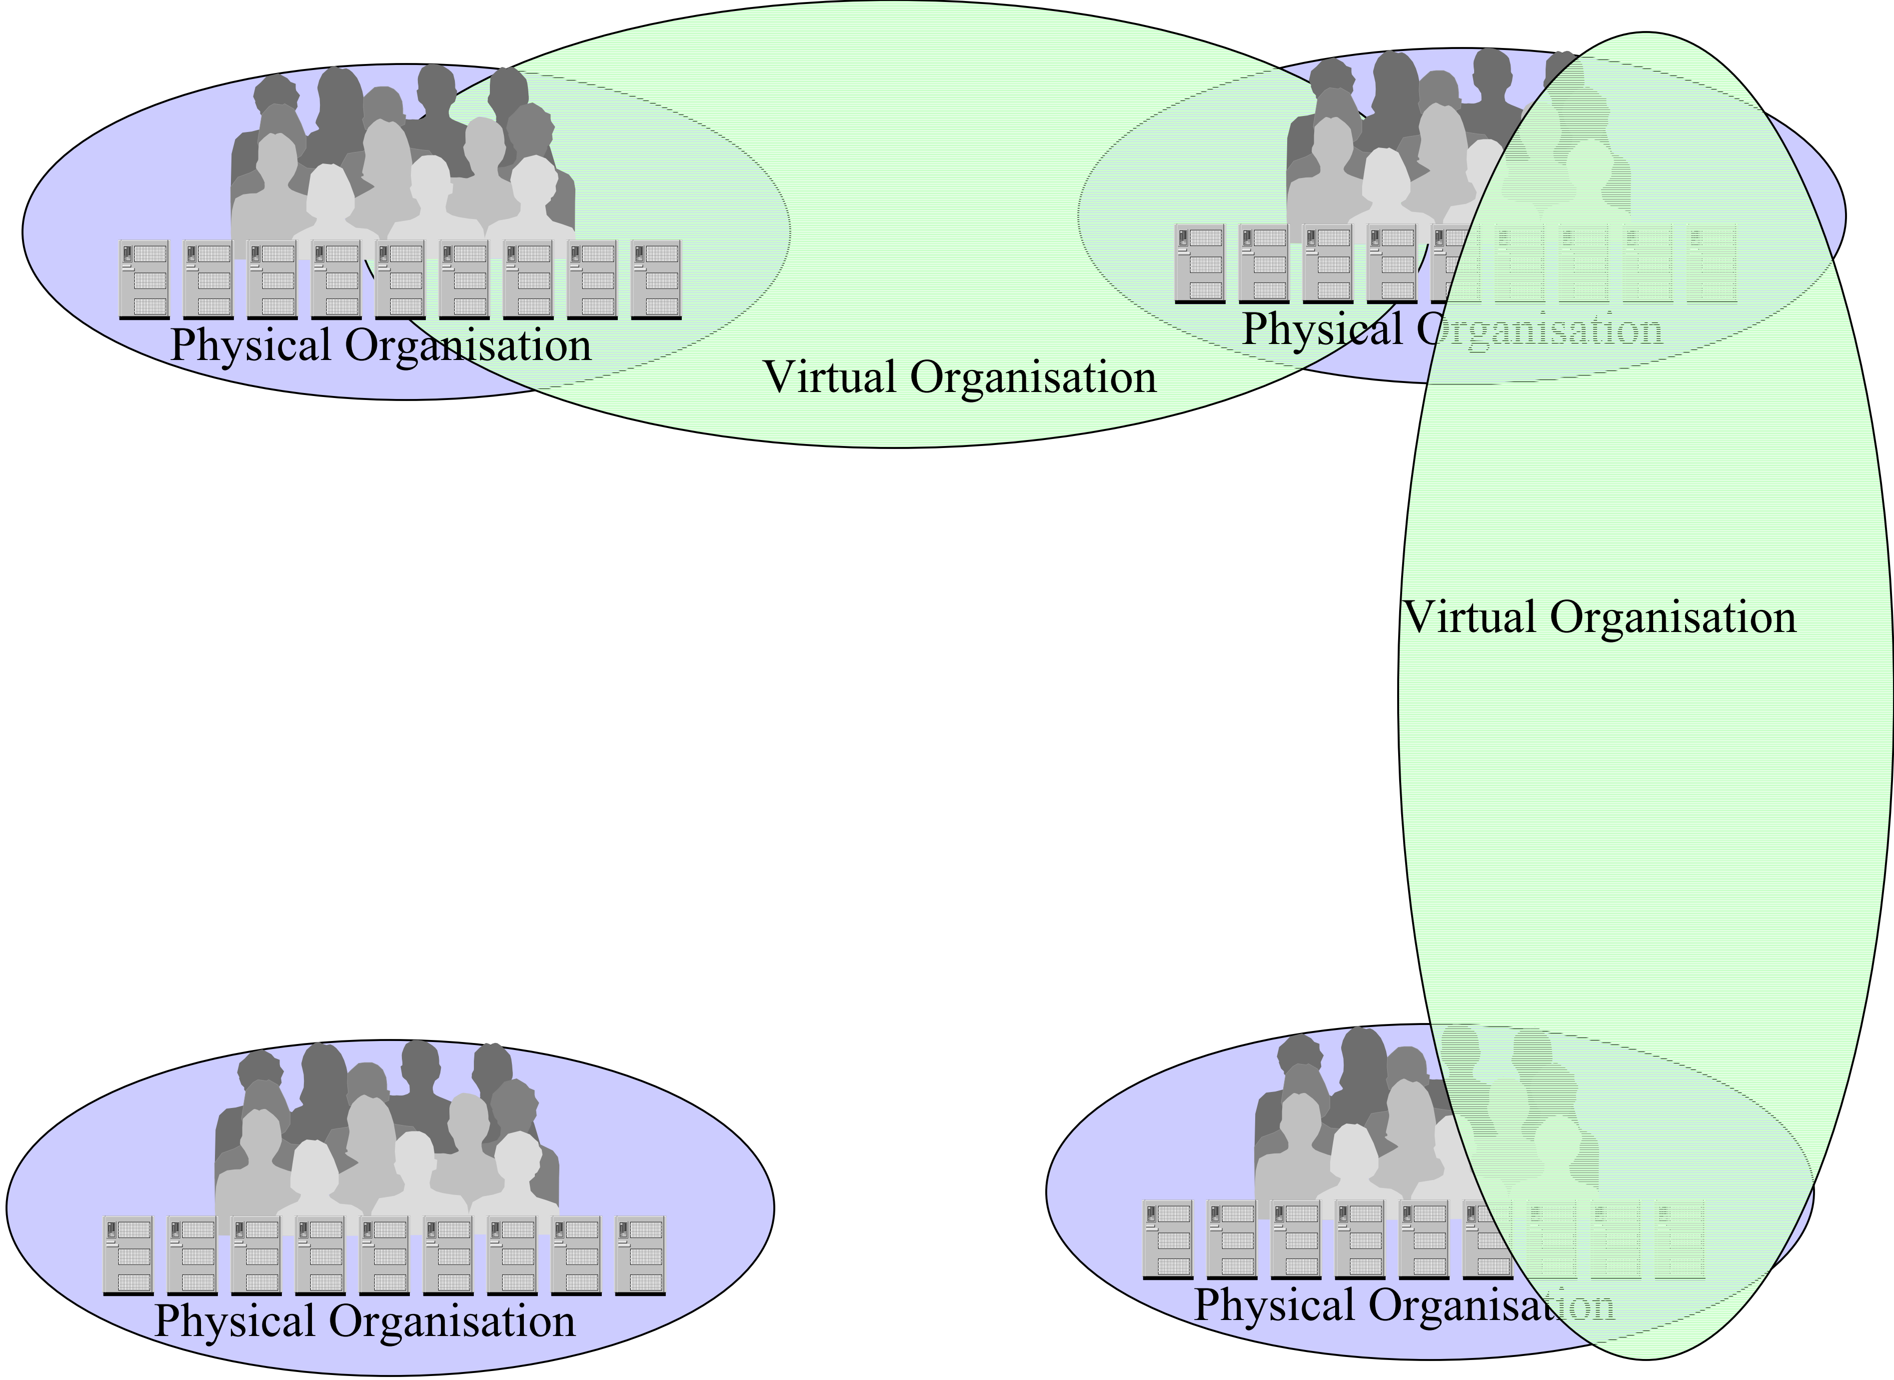
Physical Organisation
Virtual Organisation
Physical Organisation
Virtual Organisation
Physical Organisation
Physical Organisation

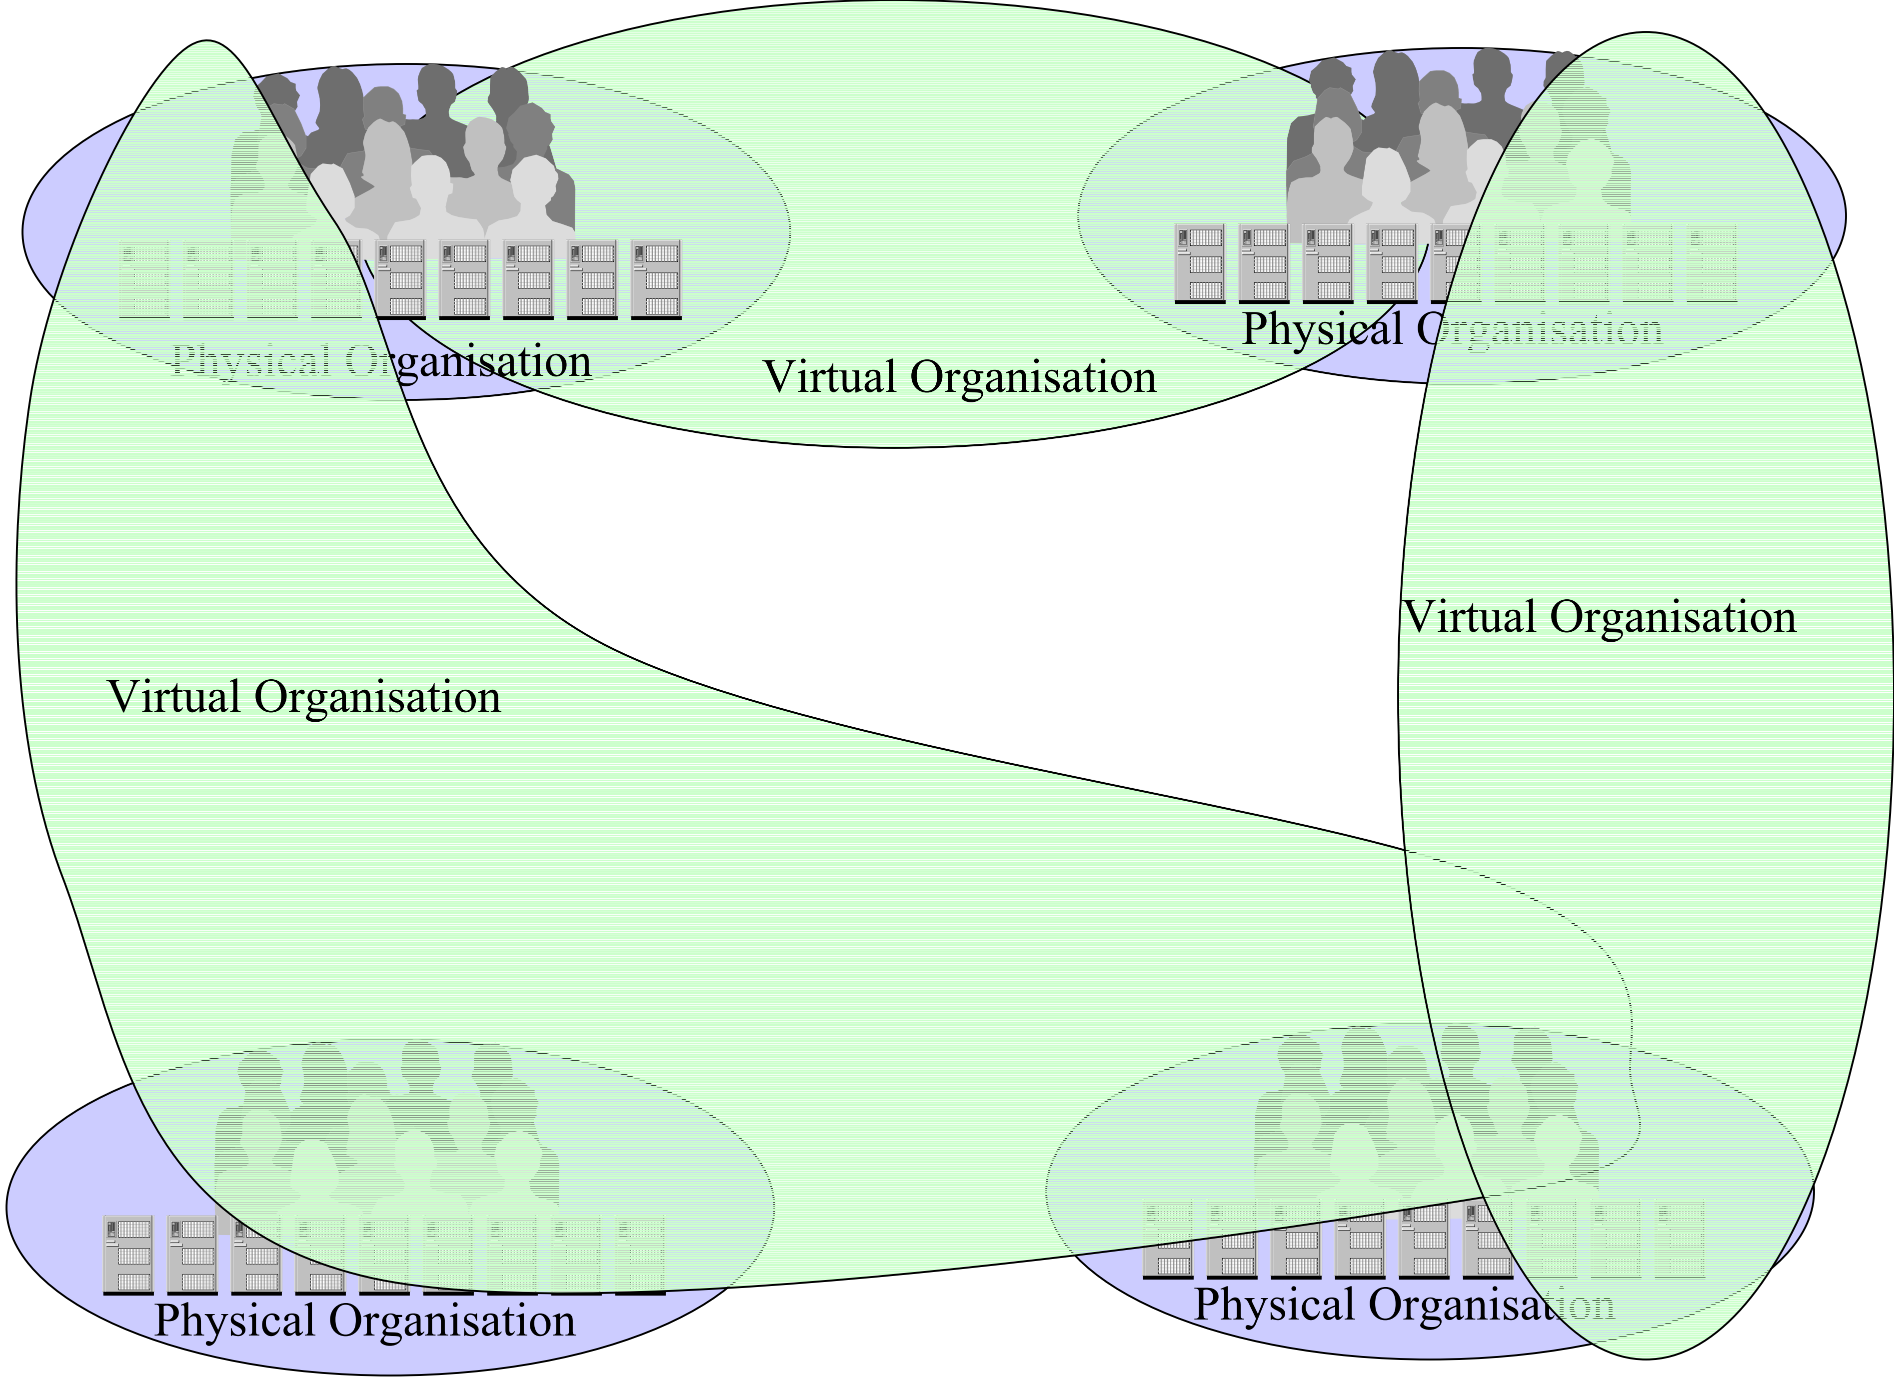
Physical Organisation
Virtual Organisation
Physical Organisation
Virtual Organisation
Virtual Organisation
Physical Organisation
Physical Organisation

Why it isn't trivial to do

- "Lack of central control, omniscience, trust"
- "Primary challenge: to *enable*, *maintain,* and *control* the sharing of resources to achieve a common goal"
- Technical challenges
  - Heterogeneity, WANness (latency and disconnects), scale, autonomy, dynamic nature, unpredictability, privacy and security
- Match or exceed the resilience and self-healing of the Internet itself

Need for management and open standards

- Grids are much more than peer-to-peer computing
  - Grids must create & manage VOs
  - Therefore, Grids require strong resource & security management
- Grid computing cannot be proprietary
  - Grids must run on heterogeneous platforms
  - Therefore, Grids require open standards and APIs.
  - The dominant open solution is GLOBUS.

GLOBUS overview

- GLOBUS is an open source toolkit developed initially by the “big science” computing community in the US (Argonne National Lab, USC, etc.)
- Freely available for various platforms under its own open source licence at **http://www.globus.org**

GLOBUS toolkit (v2) components

- Grid Security Infrastructure
- Grid Resource Allocation Management
- Uses resource brokers (e.g., Load Leveler, Condor Matchmaker)
- Grid Resource Information Service
- GridFTP
- Etc.

GLOBUS and standards

- GLOBUS v2 uses a snapshot of older Internet standards such as
  - LDAP for information services
  - SSL, X.509 for security
- GLOBUS v3 will move to Open Grid Services Architecture (OGSA) using Web Services standards such as
  - WSDL
  - SOAP
  - WS-Security
- Open standards work in Global Grid Forum (GGF), **http://www.ggf.org**

First Grid usage: Revolution in Science

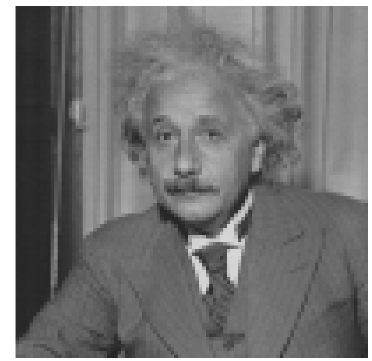

- Pre-Internet
  - Theorize &/or experiment, alone or in small teams; publish paper
- Post-Internet
  - Construct and mine large databases of observational or simulation data
  - Develop simulations & analyses
  - Access specialized devices remotely
  - Exchange information within distributed multidisciplinary teams

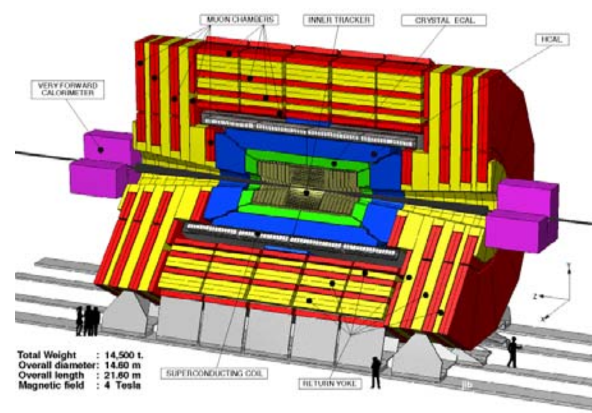

What applications are suitable for a Computational or Data Grid?

- Many traditional High Performance Computing applications, e.g.
  - Big Physics
  - Seismology
  - Fluid dynamics
  - Protein analysis
  - Bioinformatics & Medical imaging
- Large-scale engineering design
  - Automobile & aerospace design
- Financial systems
  - Market modelling

The Internet Tomorrow

Science Grids

Building the National Virtual Collaboratory for Earthquake Engineering Research

NEESgrid

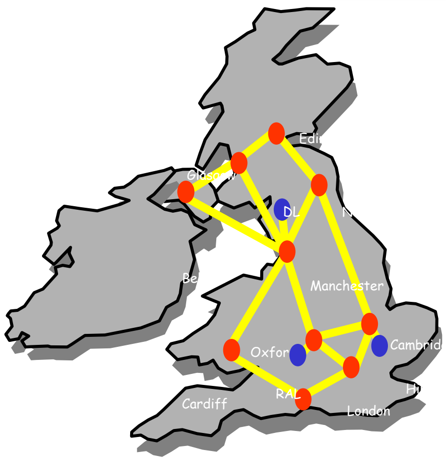

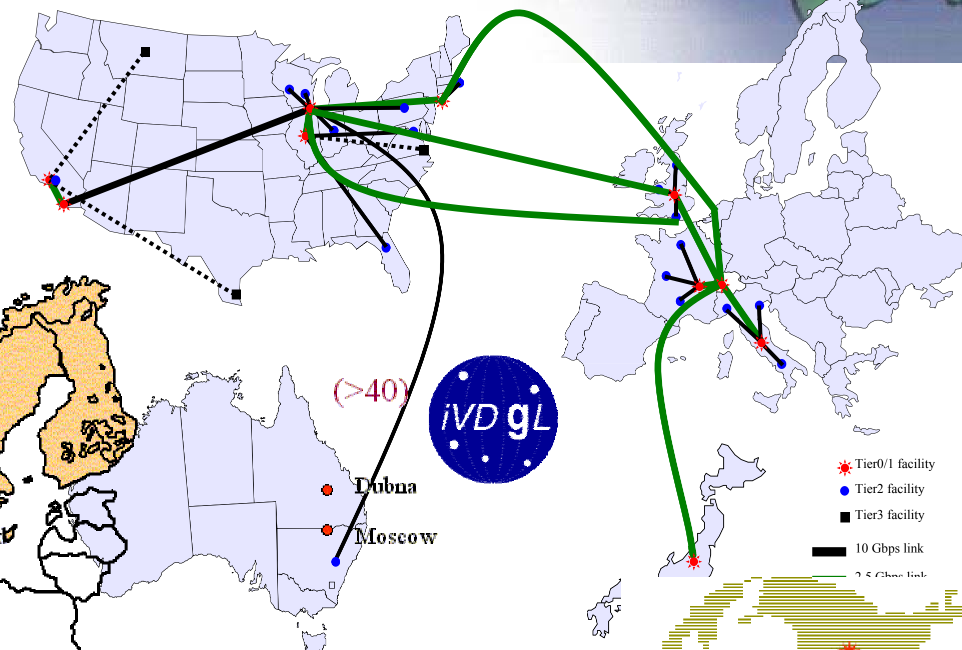

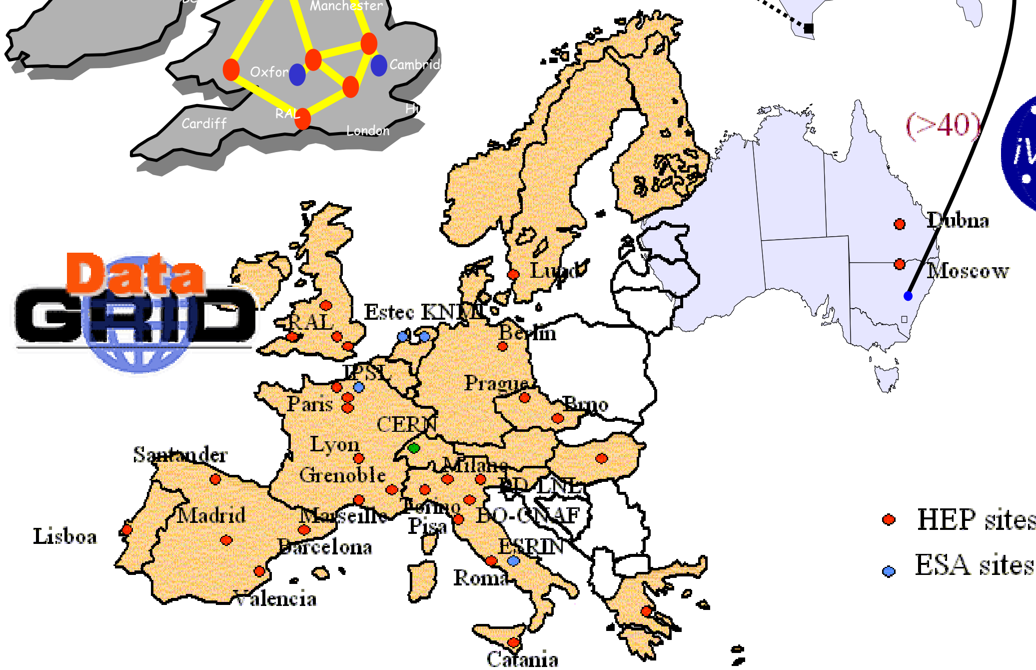

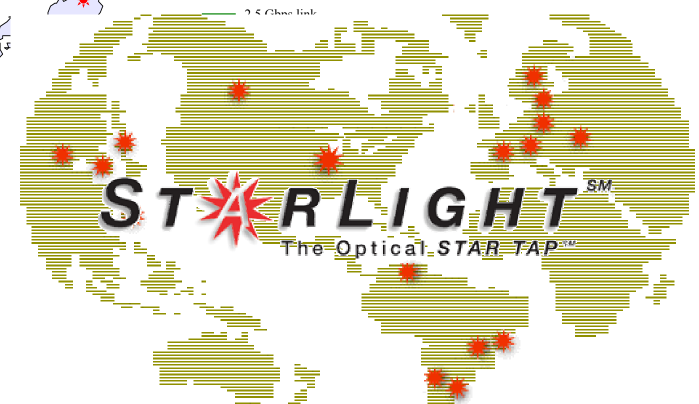

Grid concepts apply to commercial computing

- On-demand access to transaction processing & Web Services power *requires* resources virtualized across clusters and sites.
- Business models *require*
  - Flexible server-to-server interactions with standard protocols and late binding
  - Service location & resource management
  - QOS: guaranteed availability, utilisation
  - Security: management domains, authentication, privacy
  - Central monitoring, reporting, accounting

Grid Deployment Scenarios

- “intraGrid” to flexibly share resources within a distributed organisation
- “extraGrid” to share resources with business partners
- “interGrid” or “Service Grid” to share resources among a variety of customers

Grid Deployment Options

A function of business need, technology and organizational flexibility

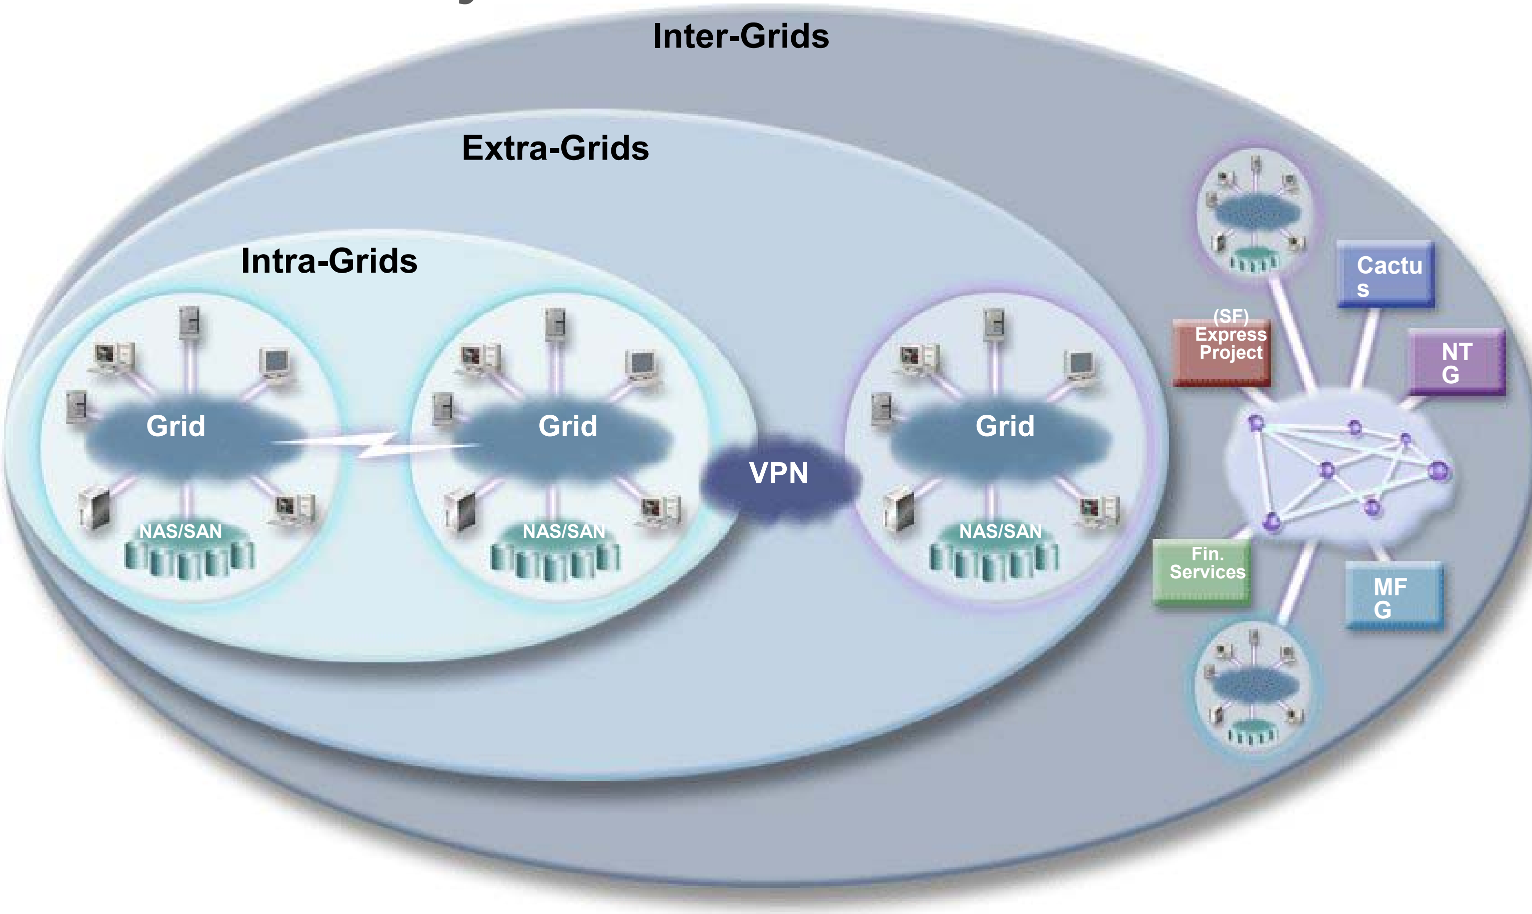

Initial Grid Focus Areas

Research & Development	Engineering & Design	Business Analytics	Enterprise Optimization	Government Development
Accelerate and enhance the R&D process by enabling the sharing of data and computing power seamlessly for research intensive applications	Share data and computing power, for computing intensive engineering and scientific applications, to accelerate product design	Enable faster and more thorough business planning and analysis through the sharing of data and computing power	Optimize computing and data assets to improve utilization, efficiency and business continuity	Create large-scale IT infrastructure to drive economic development and/or enable new collaborative government services

Industry support

- 500..1000 attendees at Global Grid Forum meetings
- Some of the companies supporting Globus, GGF, or both:
  - Avaki, BAE, Boeing, BT, Cisco, Cray, Entropia, Ford, Fujitsu, Hitachi, HP (was Compaq), IBM, InSORS, Intel, Johnson & Johnson, Juniper, Level 3, Microsoft, NEC, Platform Computing, Qwest, SGI, Sun Microsystems, United Devices, Veridian, Veritas.

Open Grid Services Architecture (OGSA)

- An architecture originated by IBM and the Globus team, being developed in the Global Grid Forum.
  - Defines WSDL interfaces and behaviors that define a Grid Service, plus extensibility elements → a component model
  - Uses SOAP as the binding protocol →loose coupling
  - Globus v3 is OGSA based
  - Unifies Grid & Web Services in one framework for the definition of composable, interoperable services
- Uses Web Services Security Architecture published by IBM, Microsoft and Verisign and being standardised in OASIS.

http://www.globus.org/ogsa/

OGSA Interactions

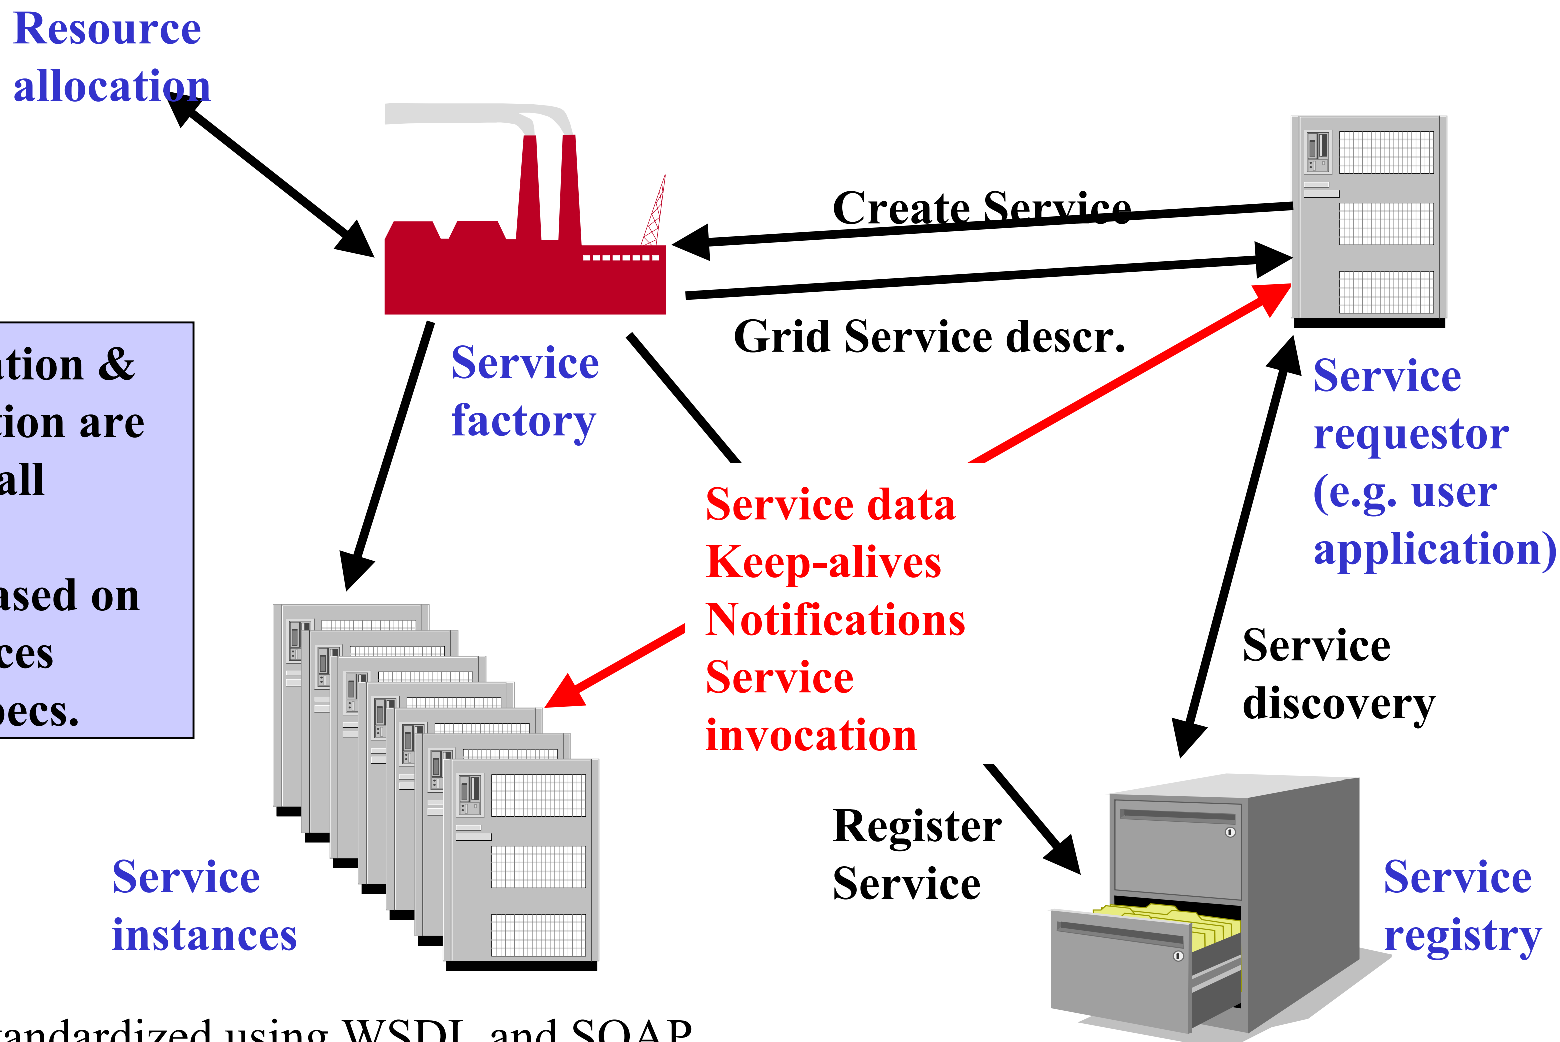

Authentication & Authorization are applied to all requests.
Security based on Web Services Security specs.

Interactions standardized using WSDL and SOAP

The Internet Tomorrow

OGSA builds on Web Services

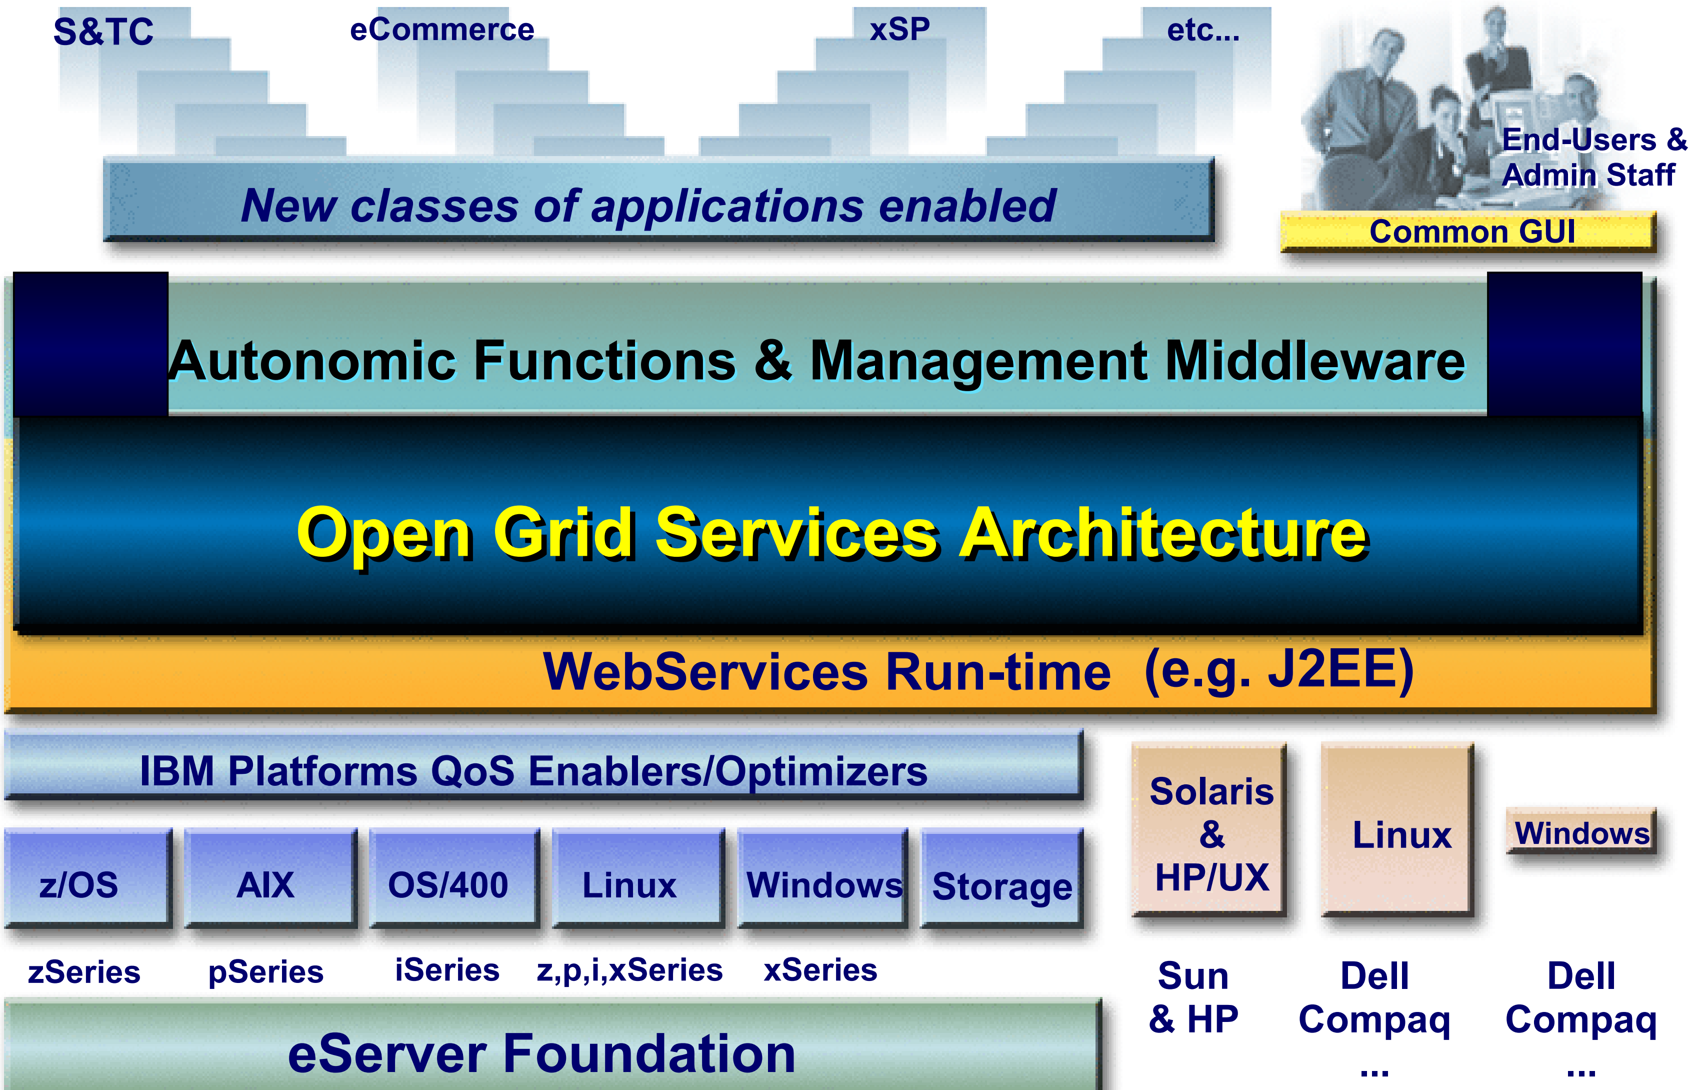

OGSA Security Requirements

- **Naming users, attributes, targets & operations across real and VOs**
- **Managing VO membership and access policies**
- **Mapping identities & policies across VO and local realms for single sign-on & local access enforcement**
- **Trust delegation and credential propagation**
- **Secure logging, audit & accountability**
- **Cross-VO intrusion detection and anti-virus measures**
- **E2E communication / session security (confidentiality & integrity)**

- **Grid-specific twists**
  - **VOs**
  - **Dynamism**
  - **Openness**
  - **Scale**
- **WS-security will already provide**
  - **Communication security**
  - **E2E Conversation security**
  - **Security policy exchange**
  - **Trust management**
  - **Federation management**
  - **Authorization management**
  - **Privacy management**

http://www.cs.virginia.edu/~humphrey/ogsa-sec-wg/

OGSA Security Directions

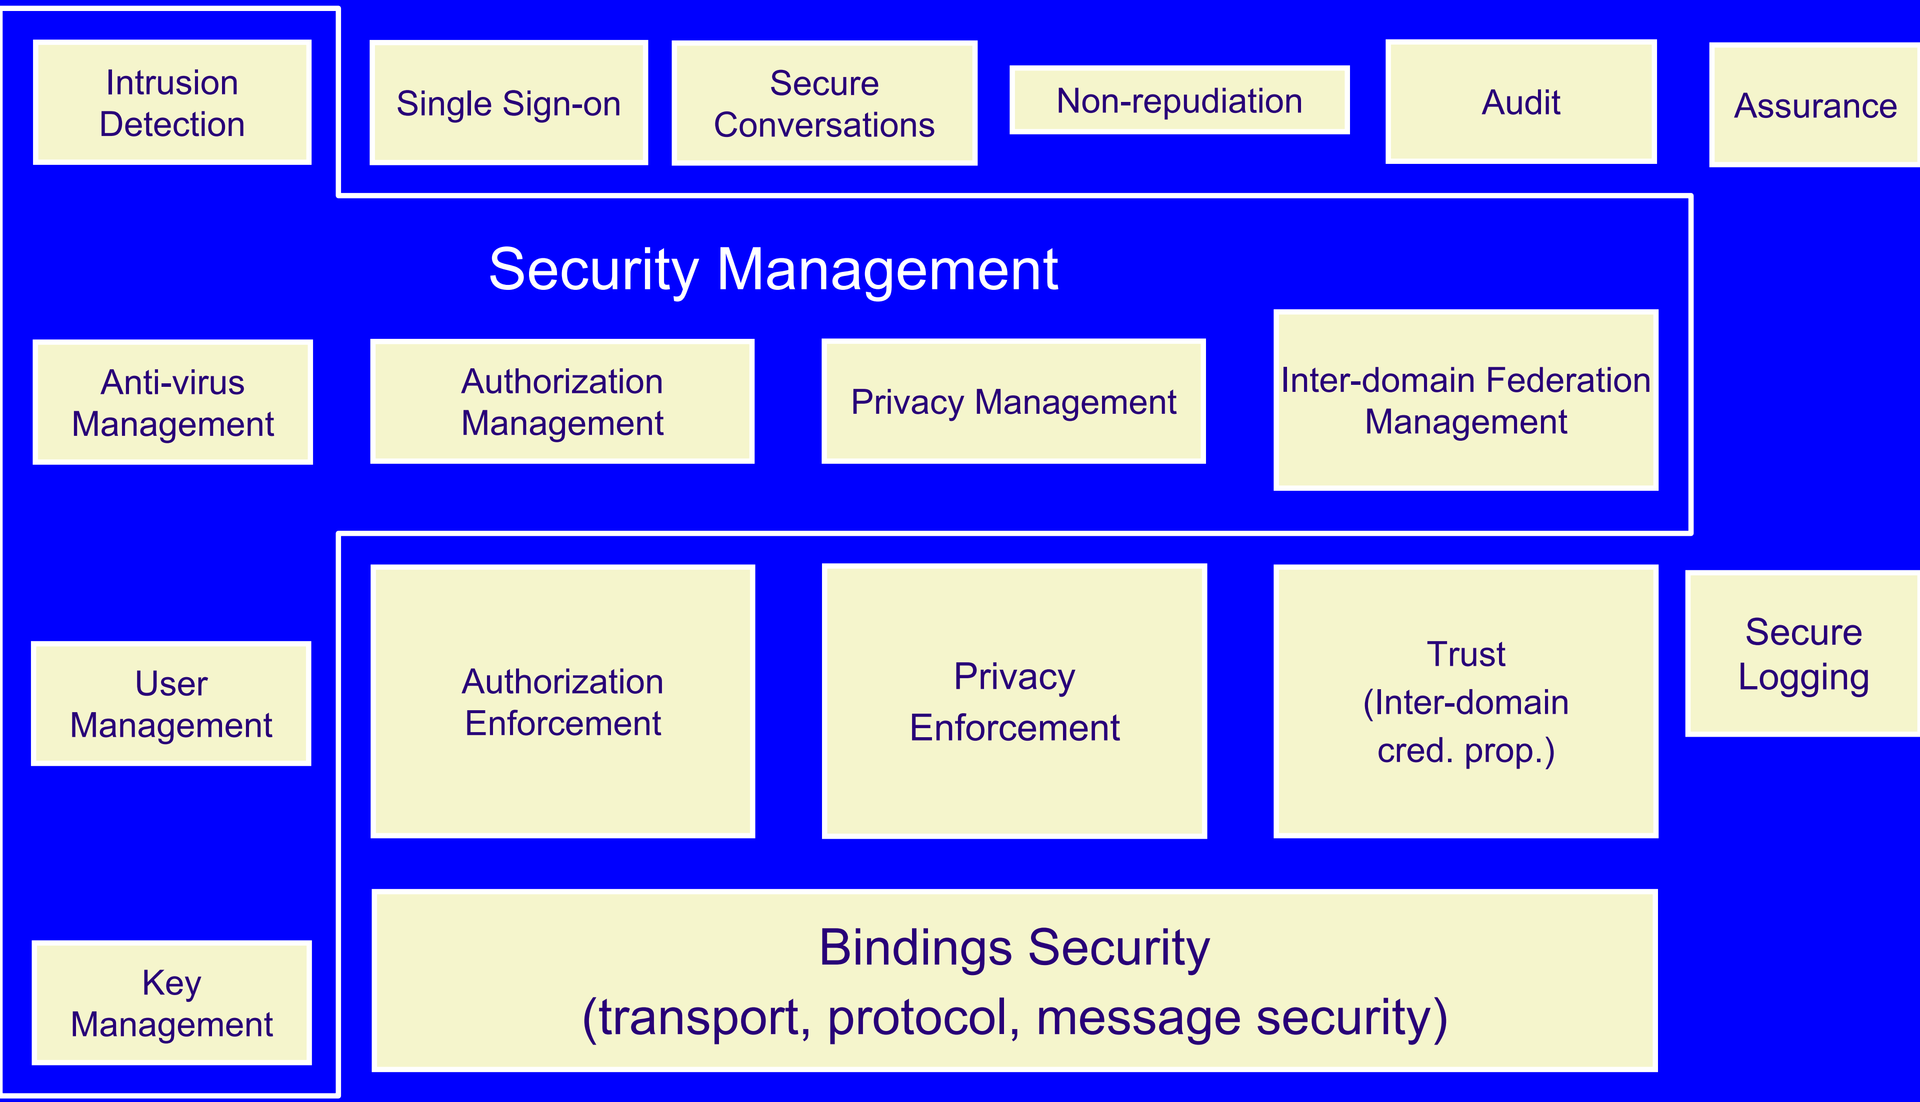

Open Grid Services Infrastructure (OGSI)

- The first detailed technical specification for OGSI
- Defines the WSDL structures needed in an OGSA service interface
- With Web Services tooling and run-time support, provides an OGSA environment

http://www.ggf.org/ogsi-wg/

Globus Toolkit v3 (GT3) Open Source OGSA Technology

- ➢ Snapshot downloads from globus.org
- 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, …
- Growing adoption and contributions
  - "Linux for the Grid"

A note on performance

- GT3 is not finalised and not optimised
- First reports suggest the the OGSI mechanisms do cause a significant performance penalty, especially for trivial (“hello world”) applications
- Caused by the overhead of interpreting WSDL text on the critical path
- Optimisation techniques will improve this but will never eliminate it completely – this is the price of loose coupling and late binding – i.e. the price of robustness and flexibility

Topics

- The Internet today: as far as Web Services
- The Internet tomorrow: a services platform for computing on demand
- Challenges at the network level (transparency, addressing, routing, multihoming)
- Challenges at the middleware level (service architecture, heterogeneity, security, integrity)
- Challenges outside the technology
- Releasing known potential: beat the challenges
- Summary

Scaling the address space

- Known problem since 1992
- Solution chosen in 1994
- IPv6 products since 1997
- Stable IPv6 standards since <2000
- So why is it so slow to start?
  - Operational costs of conversion; operational conservatism
  - Lack of strategic incentives in a fundamentally short-term industry
  - Pain from NAT is spread too thinly and not applied to the decision makers

The backbone routing system

- Another problem known since 1992, but far harder in principle than scaling the address space.
  - See RFC 3221
  - See http://bgp.potaroo.net/ for the curve
  - BGP4+ is not adequate for much longer
    - maybe 5 years to go?
  - Still a research topic, see http://www.irtf.org/charters/routing.html draft-irtf-routing-reqs-01.txt

Geoff Huston's BGP graph (89-02)

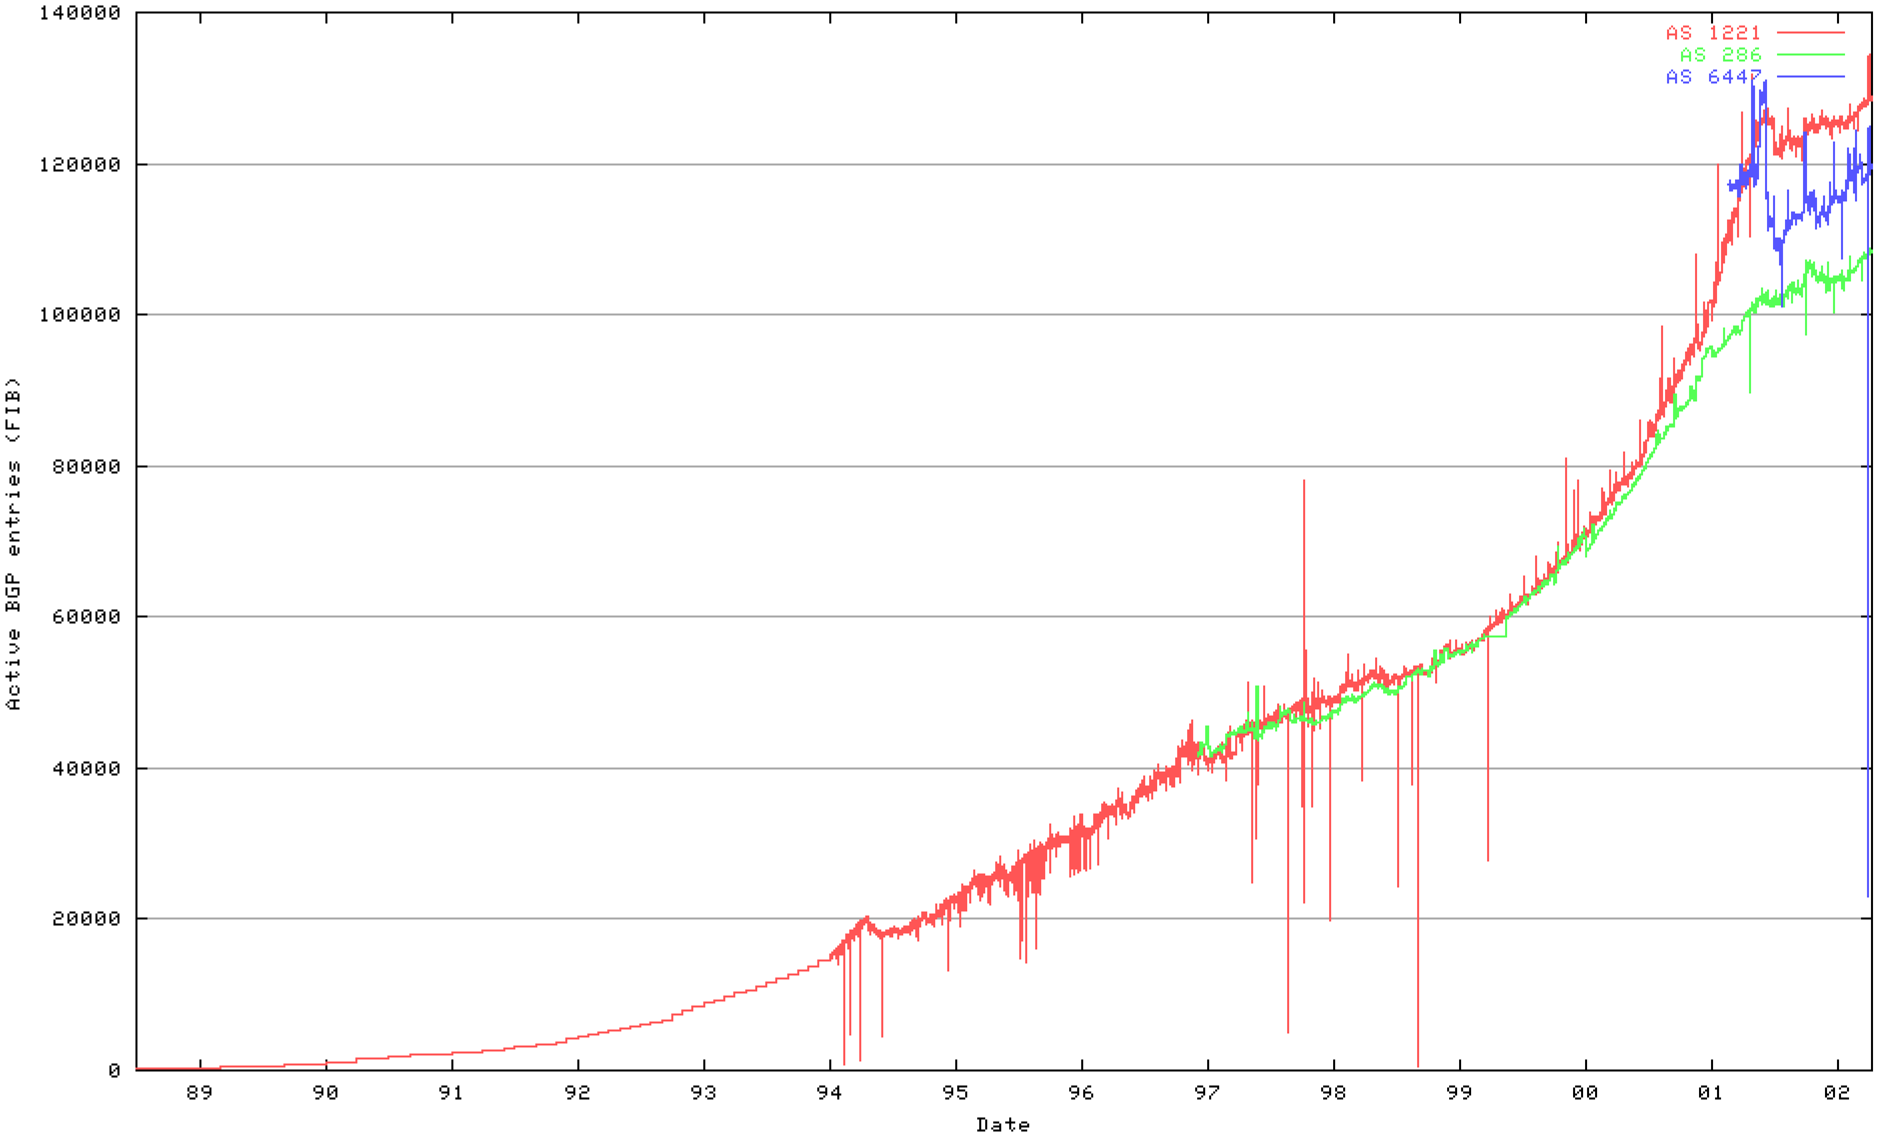

Geoff Huston's BGP graph (94-03)

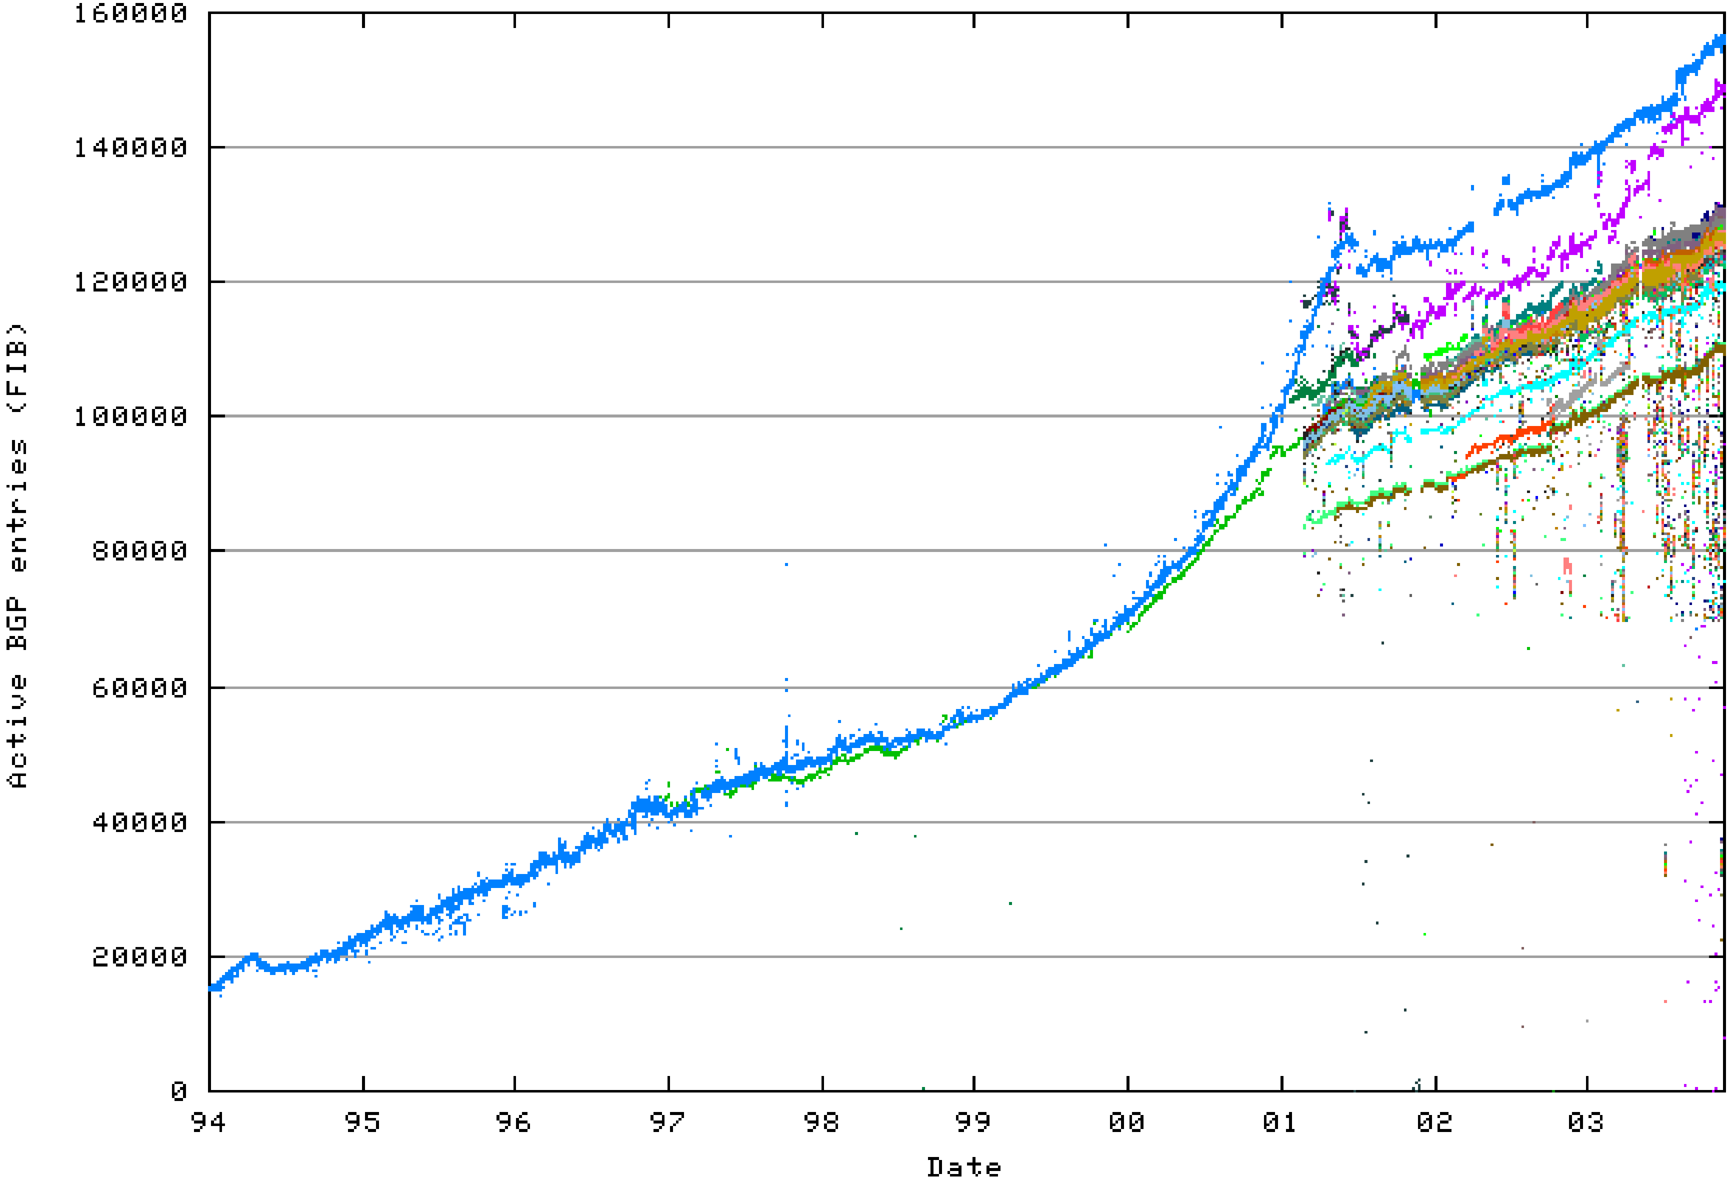

Multihoming

- An important requirement for enterprises and local ISPs is the ability to be connected to multiple upstream ISPs with automatic switch-over from one to another when needed.
  - Today that causes further explosion of the routing table (one extra entry per multihomed customer)
  - Work continues on how to avoid this scaling problem for IPv6

Quality of Service

- We've invented session-oriented (intserv) and stateless (diffserv) models for Internet QOS.
  - MPLS (layer 2.5) also supports diffserv
  - IETF is designing a new signalling system to replace RSVP
- Both technologies are available in widely used products. Neither has swept the world.
- Like IPv6: how can we get a new technology into the current practice of every network operator?
  - See RFC 2990

Network Address Translation

- It was such a tempting quick fix…
- It could even be marketed as a security system (by pre-configuring it to allow nothing)
- And it breaks many non-client-server applications as well as network level security
  - See RFC 2993

Layer Violation Boxes ("Level 4 switches" etc.)

- Let's just peek into application layer headers…
- Let's just send this packet to a different server…
- Let's just proxy this request without being asked...
- Let's just rewrite this little piece here…
- They were all such tempting quick fixes
- Result: unpredictable, inexplicable glitches & failures
  - See RFC 3234
  - Middleboxes should be architected, not thrown together

Let's just put it in the DNS

- The DNS was narrowly designed, as a replacement for */etc/hosts* with distributed update and distributed lookup
- It was also designed to be extensible
- But it wasn't designed as a directory
- It is abused as a directory (pimples.com)
- It still isn't secured
  - See RFC 3467

Crunchy outside, soft inside

- Corporate firewalls attempt to divide the world into a trusted inside and a mistrusted outside (usually with a half trusted DMZ)
- Very vulnerable
  - to dishonest employees
  - to tricks with “safe” protocols
- Don’t meet new requirements
  - compartmentalized & dynamic trust relationships
  - end-to-end, any-to-any trust relationships across administrative boundaries

Internationalisation

- We thought it was straightforward: rely on ISO 10646/Unicode (RFC 2277). But…
- Some uses of text are hidden entirely in protocol elements and need only be read by machines, while other uses are intended entirely for human consumption (presentation). Many uses lie between these two extremes, which leads to conflicting implementation requirements.
  - Humans can handle ambiguity, protocol engines can't
  - Humans care about cultural aspects, protocol engines are allergic to them
  - Thus, matching & folding requirements are different in the two cases
- Some good news: Internationalised DNS is here

Let's just run it over HTTP

- HTTP was narrowly designed, to carry HTML requests and responses
- It was also designed to be easy to use
- Firewall operators are bound to let it through
- But it wasn't designed as a transport protocol
- It is abused as a transport protocol & firewall penetration technique
  - See RFC 3205

The mythical PKI

- It was thoughtless to imagine that by creating technology capable of supporting a universal public key infrastructure, such an infrastructure would come into existence.
- As a result, we have a big challenge in actually deploying public key based solutions except within closed worlds.

Some standards organisations

- IETF (Internet Engineering Task Force)
- W3C (World Wide Web Consortium)
- GGF (Global Grid Forum)
- ISO JTC1 (Specific WGs of SC 2, 6, 25, 27, 29, 32, 34)
- ITU-T (various subcommittees)
- GSC (Global Standards Collaboration)
- ETSI (European Telecommunications Standards Institute)
- ECMA (formerly European Computer Manufacturers Association)
- ICTSB(European ICT Standards Board)
- CEn/ISSS(European IT standards portal)
- Telcordia
- Web Services Interoperability
- Eclipse
- OASIS
- P2P WG
- WAP Forum
- DVB (Digital Video Broadcasting project)
- IEEE
- ATM Forum
- Frame Relay Forum
- BlueTooth SIG
- Universal Plug and Play
- jini
- Salutation
- Home Audio Video Interoperability
- UMTS Forum
- 3GPP
- 3GPP2
- Network Processing Forum
- Mobile Wireless Internet Forum
- The Open Group
- New Productivity Initiative (NPi)
- OMG (Object Management Group, CORBA)
- OSGI(Open Services Gateway Initiative)
- Unicode Consortium
- JavaSoft
- IPv6 Forum
- MPLS Forum
- Internet Software (DNS BIND)Consortium
- MINC (Multilingual Internet Names Consortium)
- IMTC (International Multimedia Telecommunications Consortium)
- Telemanagement Forum (formerly Network Management Forum)
- DMTF (Distributed Management Task Force)
- WfMC (Workflow Management Coalition)
-

Hubris

Function: noun

Etymology: Greek

Date: 1884

: **exaggerated pride or self-confidence**

(Merriam-Webster on line)

- Those who created the Internet have reason to be proud, but
  - should not lose sight of the real problems
  - should not ignore the impact of success on the original design principles of the network.

Gold diggers – guess the year

(http://www.webcom.com/~walsh/)

The Commercial domain grew by over 10,000 in the first two weeks of Aug. Kraft Foods registered 133 product names … In the second two weeks the companies switched tactics. … Procter & Gamble started registering ailments, afflictions and body parts (e.g. diarrhea.com,pimples.com and underarms.com, etc.) 36 more.

ICANN

- Administer protocol parameters
- Coordinate allocation of address blocks to the regional registries
- Coordinate allocation of TLD names to TLD registries
- Coordinate root server operations
- *How can this **possibly** cost $6M/year?*

Regulators & politicians

- National & international telecomms regulators find the Internet very tempting, but hard to get hold of. However, they are persistent.
  - When in doubt, make a regulation!
- Politicians also find it very tempting, and threatening (free speech? unwelcome material? tax free?). Also, they are unpredictable.
  - When in doubt, pass a law!
  - Never mind these geeks who say “that’s technically impossible.” Pass the law anyway.

WSIS

- World Summit on the Information Society
  - Geneva, 12/2003 and Tunis, 2005
- A mixture of industrial, NGO and developing country interests
  - multiple sources of conflict
  - strong risk of international regulation, but in whose interests?
  - strong risk of unintended consequences
  - surely better to stick to self-regulation
  - don’t hand the Internet over to the ITU!

Artificial Barriers

- Some of these challenges created artificial barriers to progress beyond the "information web" stage
  - NATs, firewalls, & thoughtless middleboxes inhibit deployment of any2any solutions (vs. client/server)
  - The firewall/intranet model & the PKI problem inhibit deployment of any2any trust & fine-grain security
- Solutions will exist
  - IPv6 is ready to roll
  - Architected middleboxes (Web Services, MIDCOM, OPES, etc)
  - Any2any trust models will emerge (VOs, intergrids)

Topics

- The Internet today: as far as Web Services
- The Internet tomorrow: a services platform for computing on demand
- Challenges at the network level (transparency, addressing, routing, multihoming)
- Challenges at the middleware level (service architecture, heterogeneity, security, integrity)
- Challenges outside the technology
- Releasing known potential: beat the challenges
  - IPv6 update
- Summary

Why it isn’t trivial to do

- It’s hard to imagine deploying OGSA (or any other generic any-to-any services architecture) across a 10 billion node network without removing the barriers identified earlier.

Why the Internet as a Computing Services Platform needs IPv6

- 10 billion nodes squeezed into 4 billion IPv4 addresses –why would we do that to ourselves?
- Immediate benefit for applications that are being actively hurt by NAT today
  - release the known potential
- Strategic benefit for the next 50 years at least
  - the opportunity cost of staying with IPv4

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 is unthinkable.

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

Overlapping virtual organizations

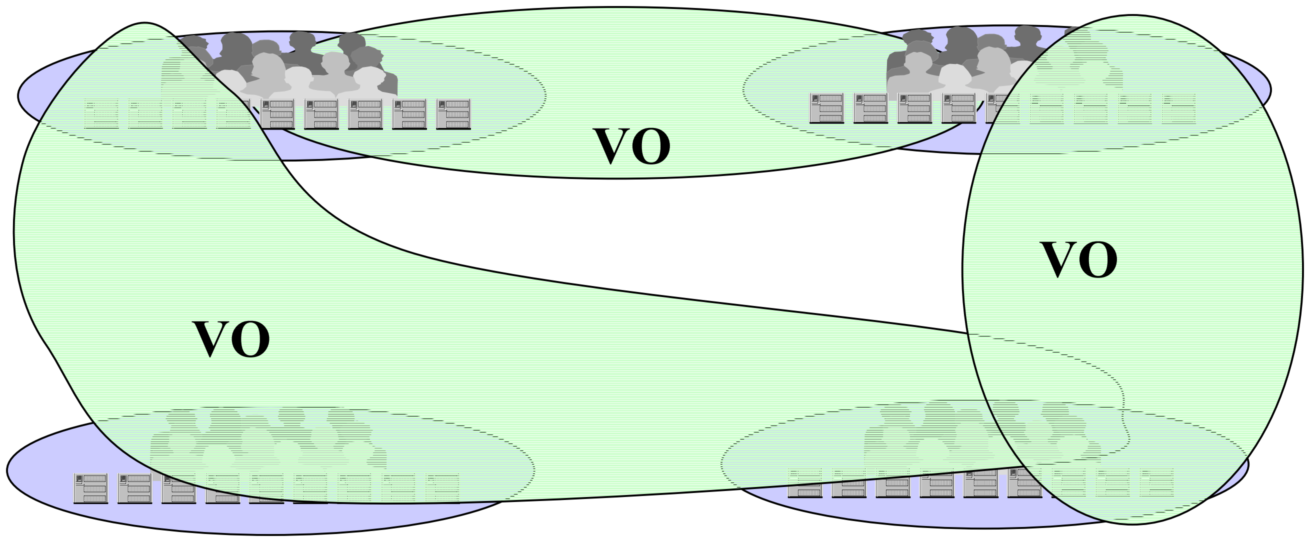

- Any system can be in any number of VOs with any number of other systems
  - needs uniform address space to avoid proxies & allow IPSEC
  - addressing ambiguities unacceptable
  - security boundaries ≠ organization boundaries
  - **can’t meet these constraints at massive scale with IPv4**

Critical advantages of IPv6 for a services architecture such as OGSA

- Potential for massive scaling
- Avoid the NAT handicap
- Autoconfiguration is a big plus for infrastructure configuration
- (Since Grids use transport and application level security, we can't claim a security advantage for IPv6)

Why we need IPv6

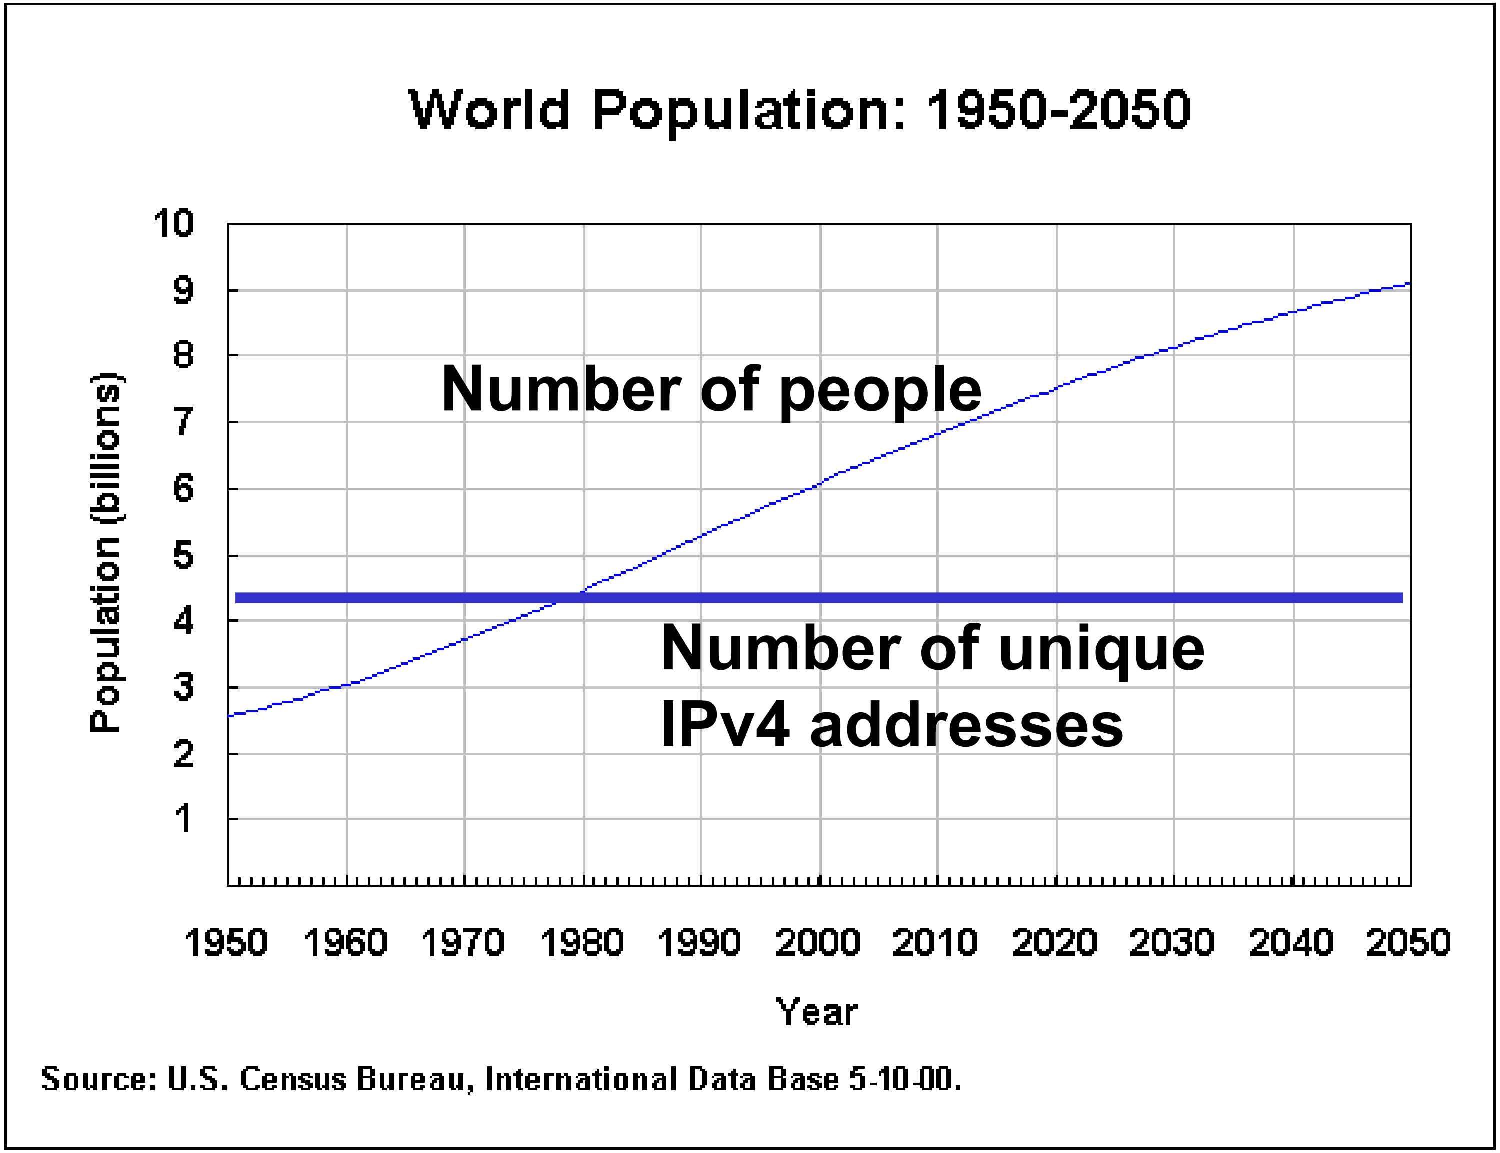

Living with too few addresses

- 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

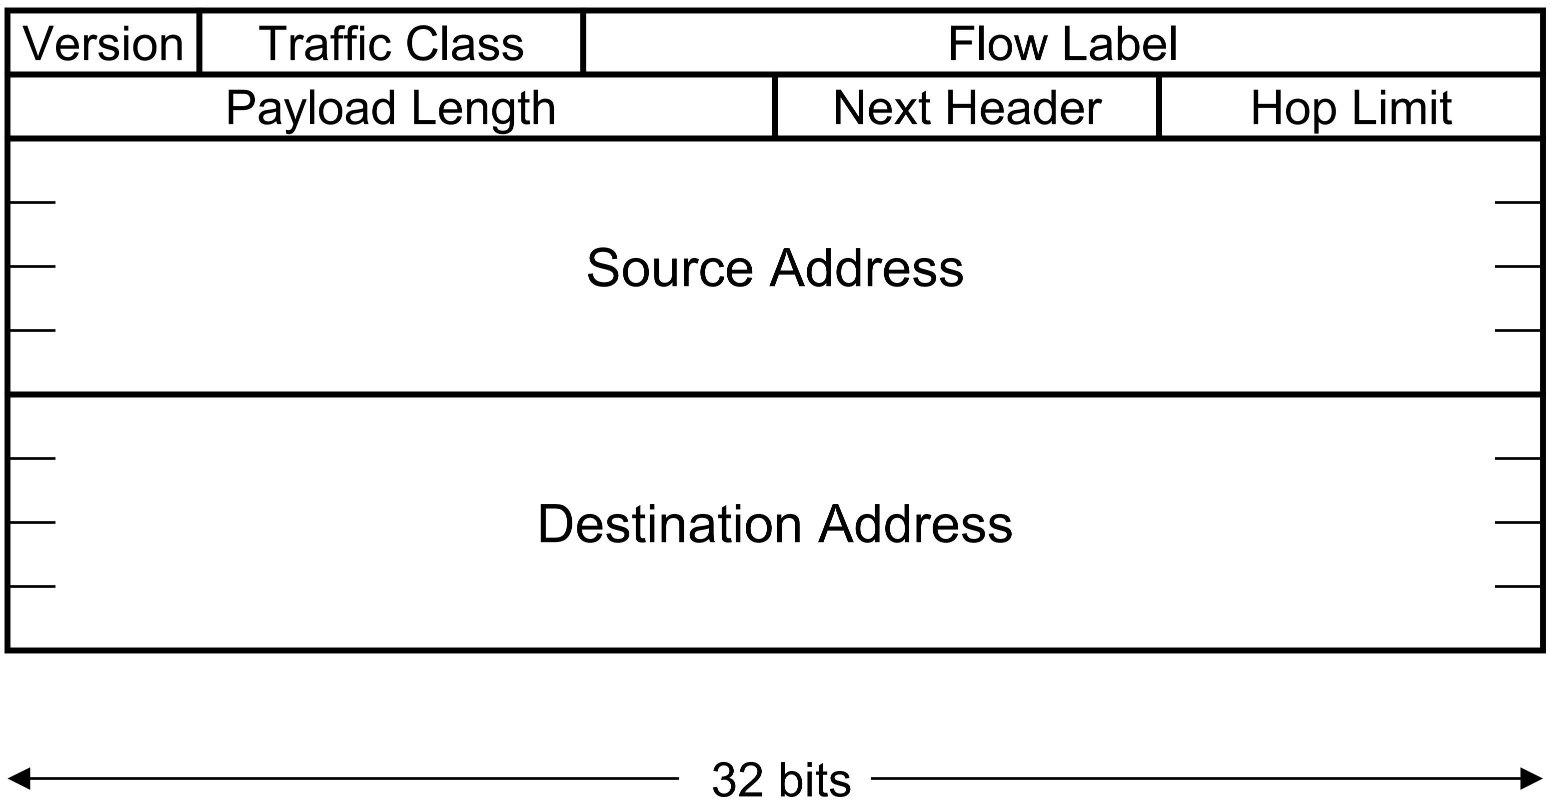

credit: Steve Deering

The IPv4 Header

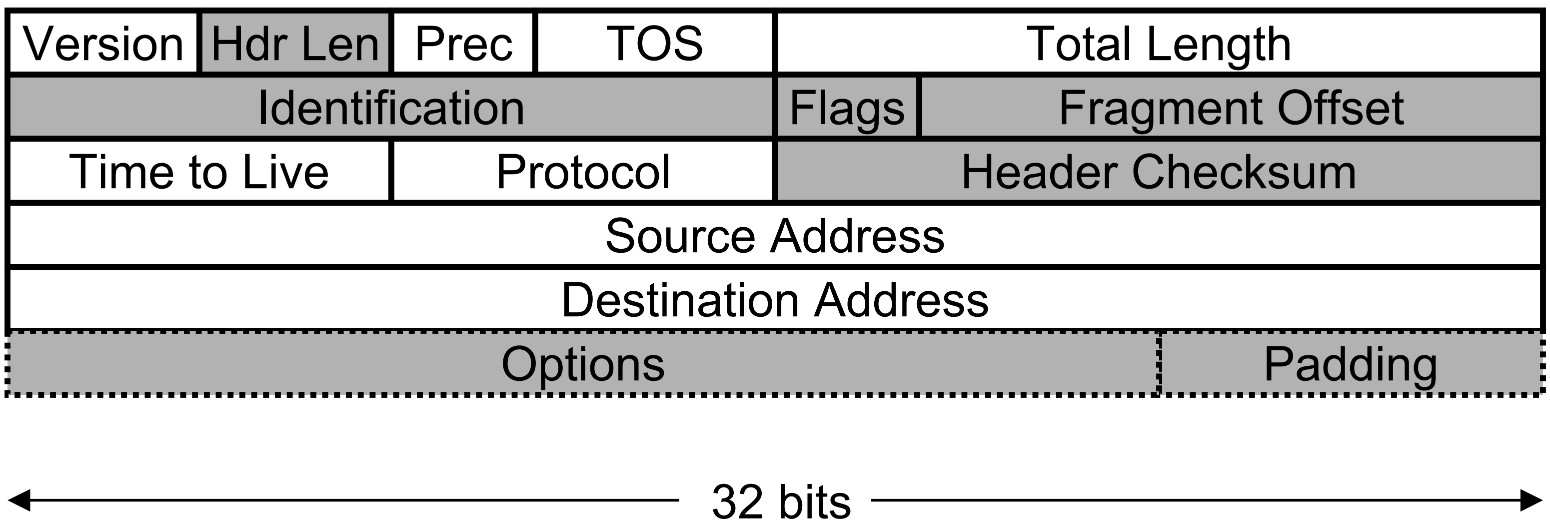

Shaded fields are absent from IPv6 header

credit: Steve Deering

Extension Headers

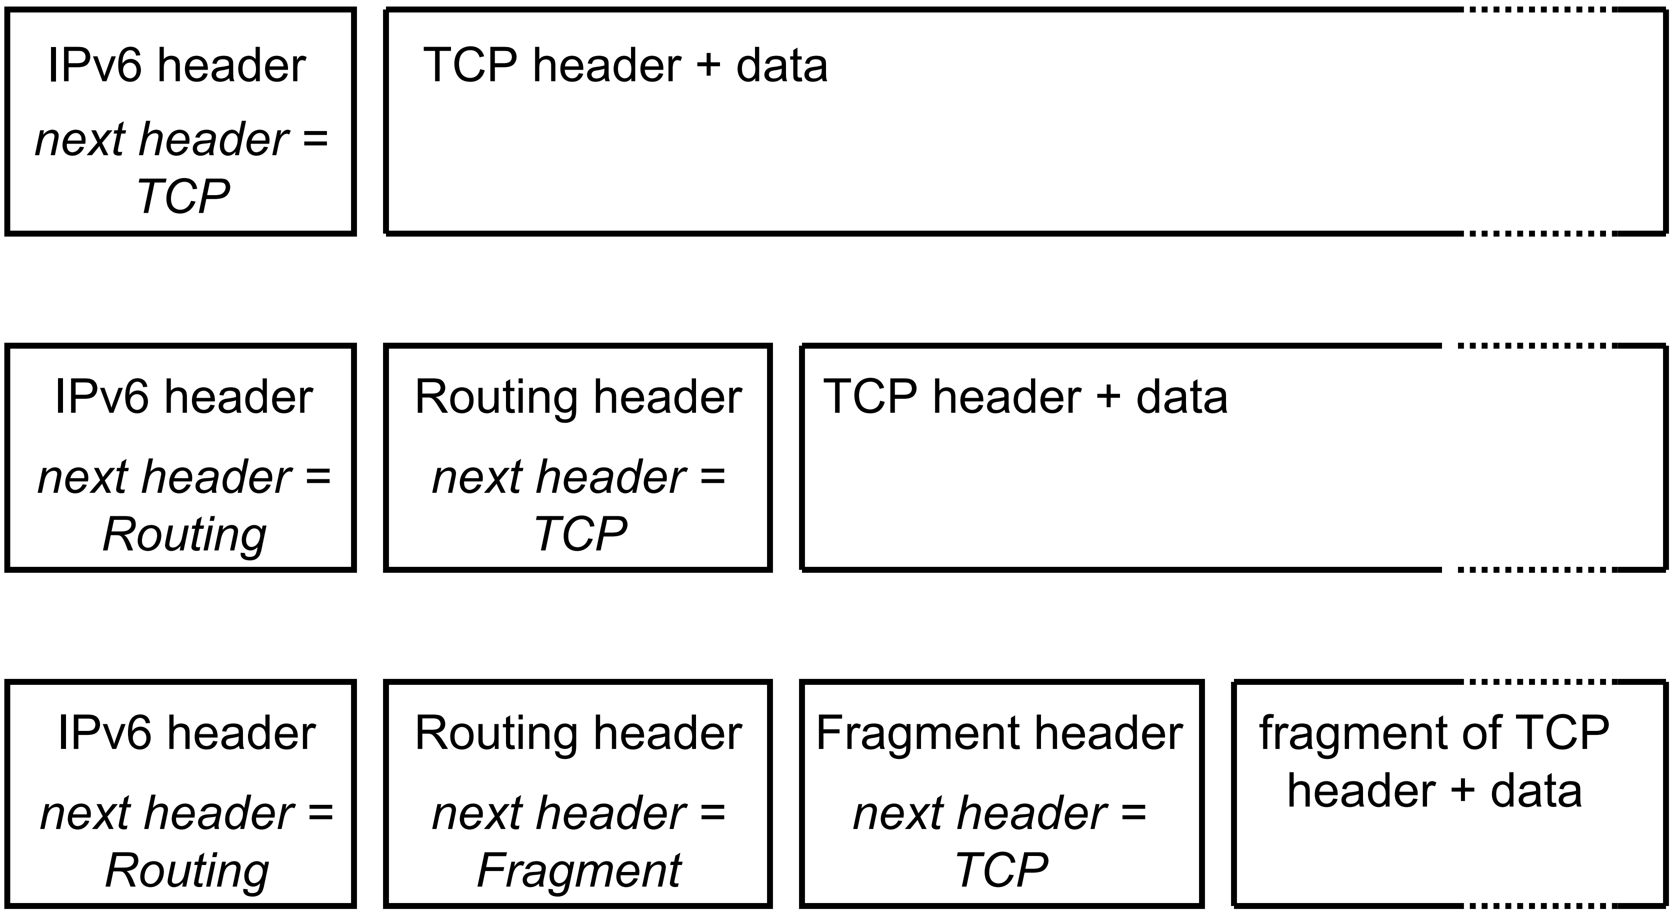

credit: Steve Deering

Global Unicast Addresses

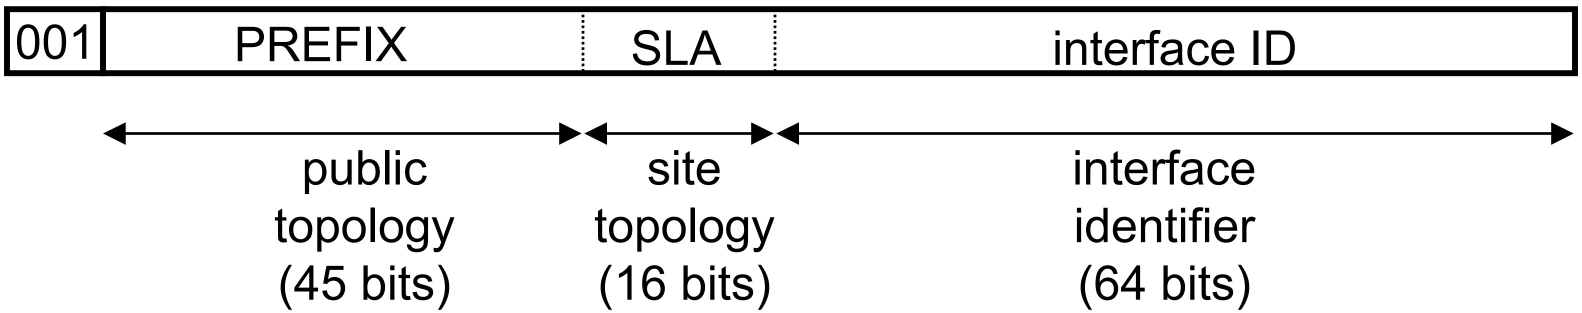

- Prefix ranges may be assigned to providers or exchanges
- Recommended that all sites including homes get 48 bit prefixes (35,184,372,088,832 are available)
- SLA = Site-Level Aggregator (subnet prefix)
- Subfields variable-length, non-self-encoding (cf CIDR)
  ➔much better route aggregation than legacy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

credit: Steve Deering

IPv6 update

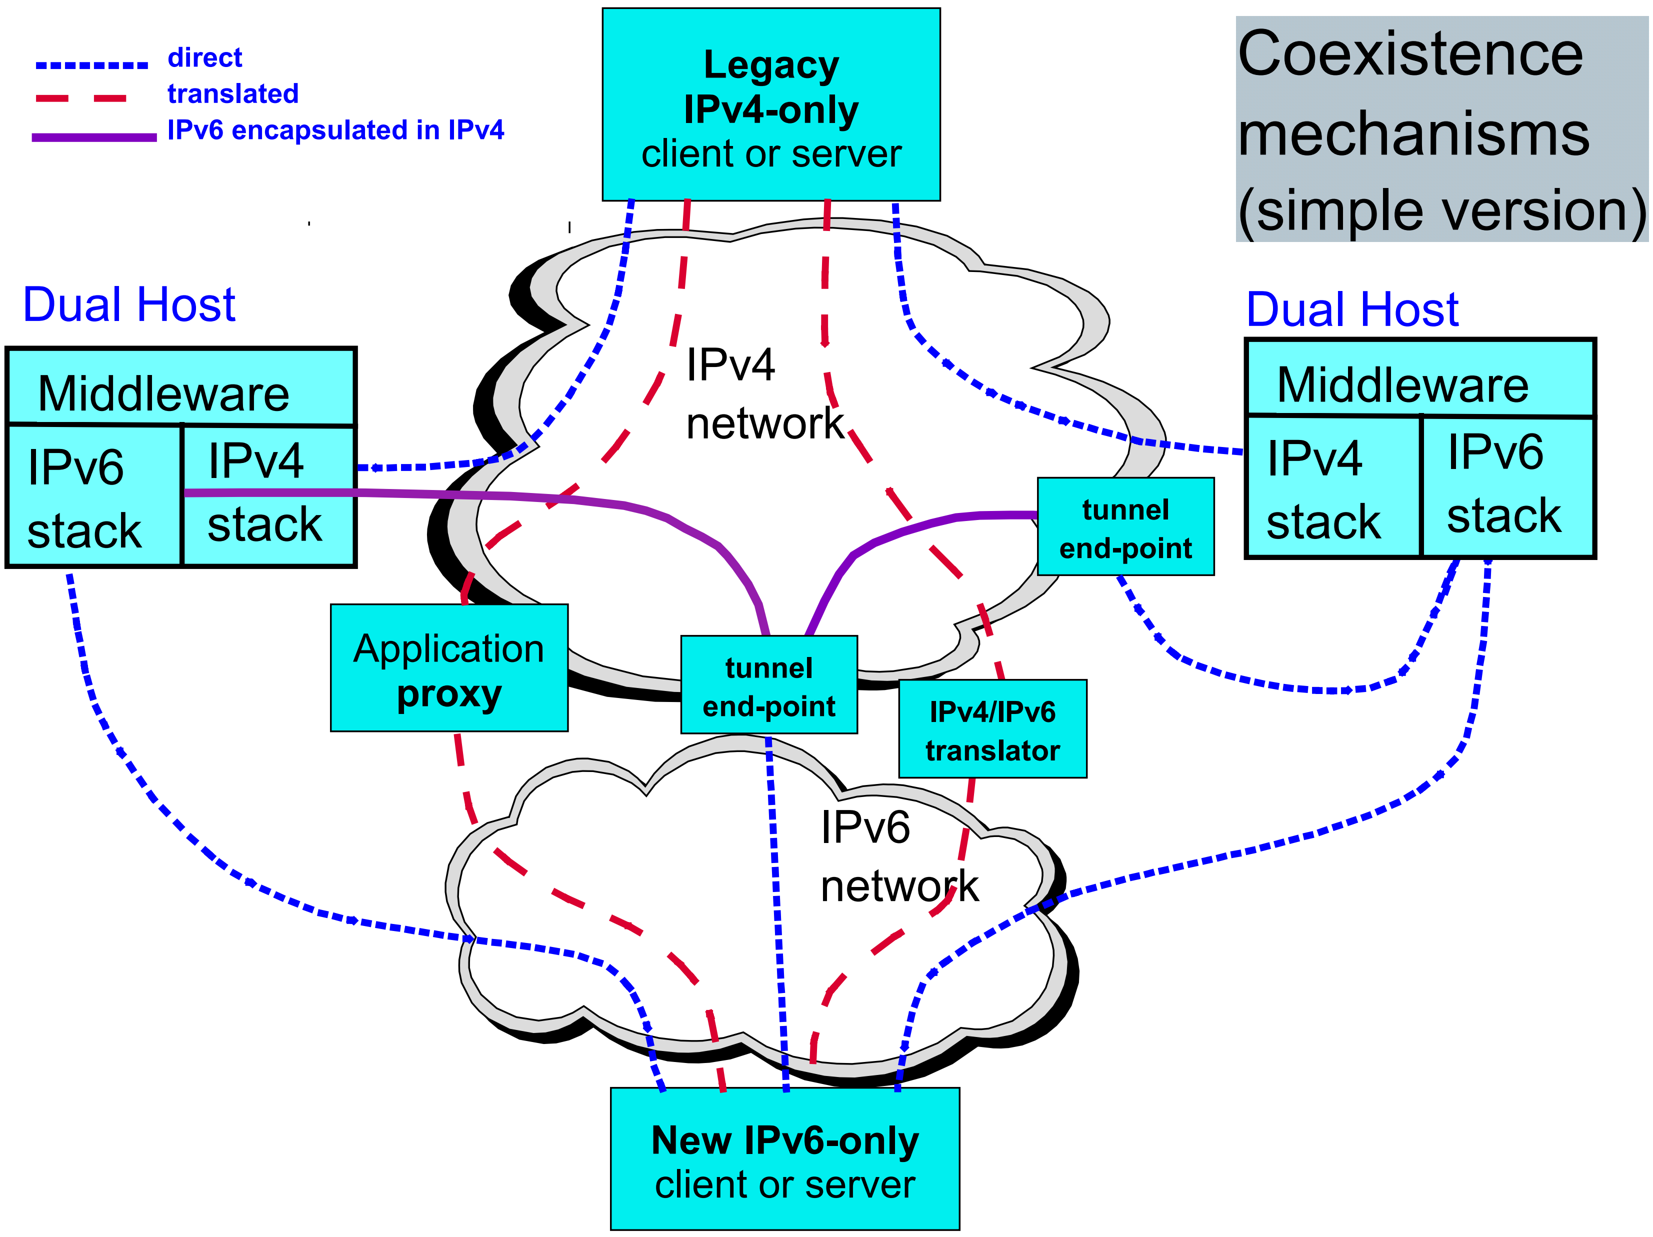

direct
translated
IPv6 encapsulated in IPv4
Coexistence mechanisms (simple version)
Legacy IPv4-only client or server
IPv4 network
Dual Host
Middleware
IPv6 stack
IPv4 stack
Dual Host
Middleware
IPv4 stack
IPv6 stack
tunnel end-point
Application proxy
tunnel end-point
IPv4/IPv6 translator
IPv6 network
New IPv6-only client or server

Example coexistence cases from 3G phone world

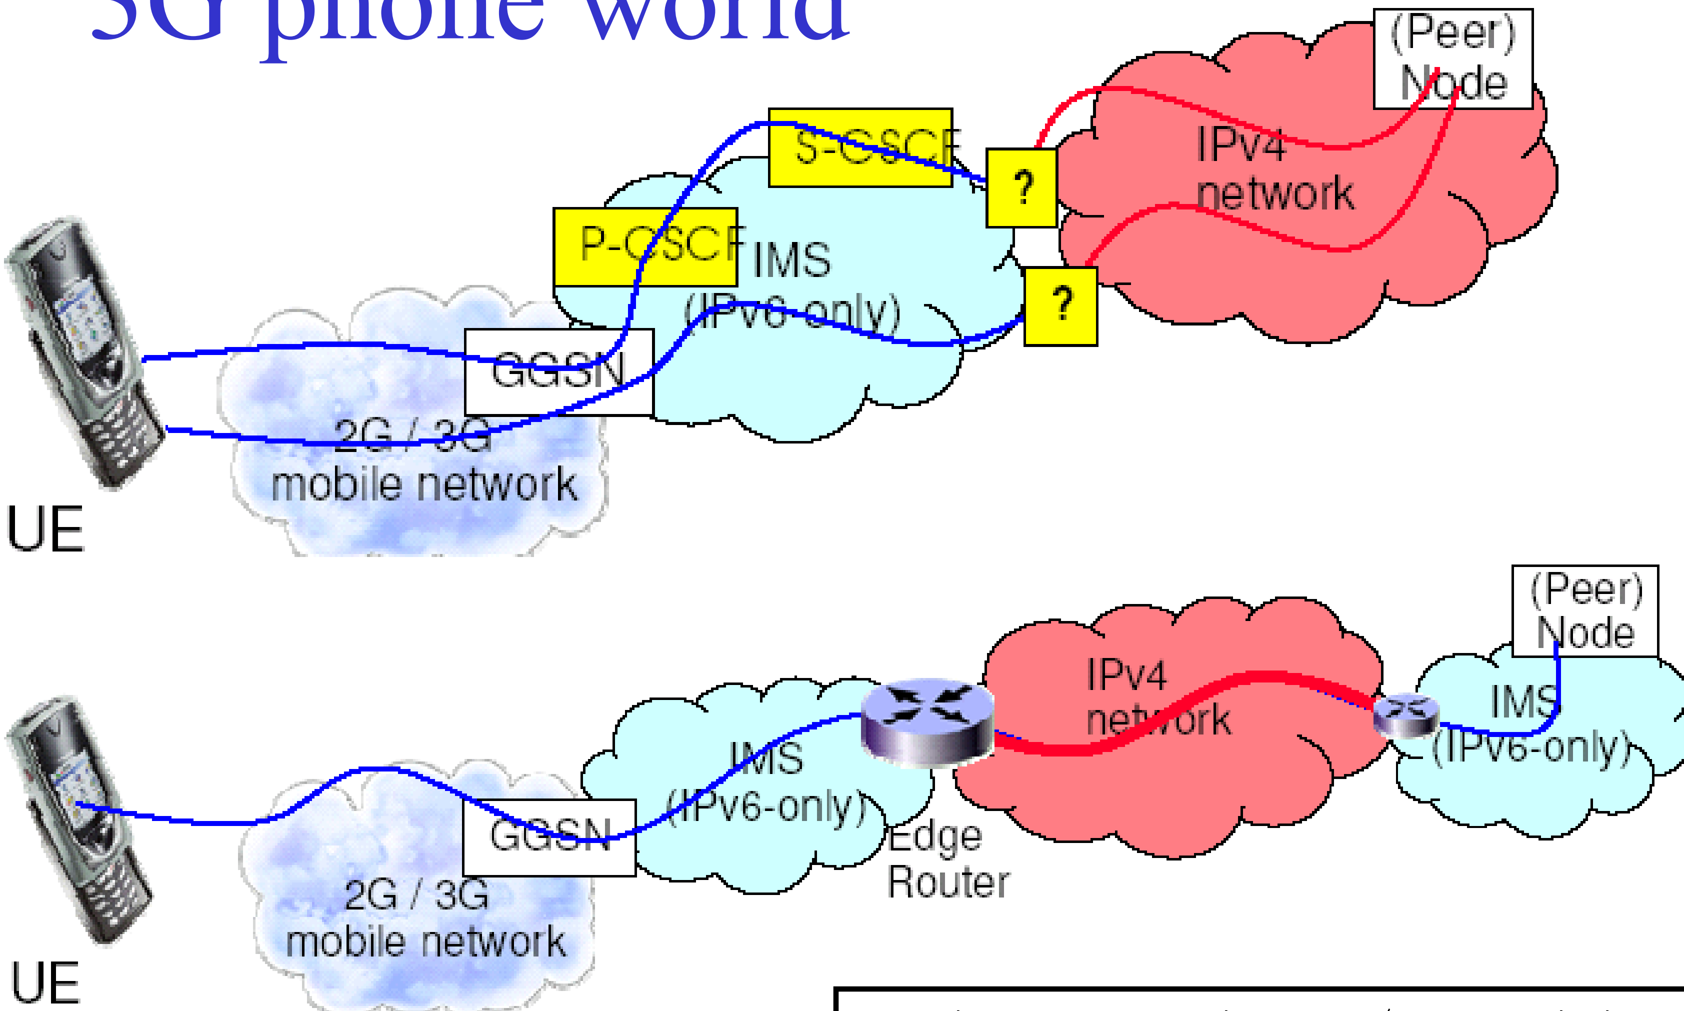

credit: Jonne Soinenen / Juha Wiljakka

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 and by US DoD and several NATO MoDs

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

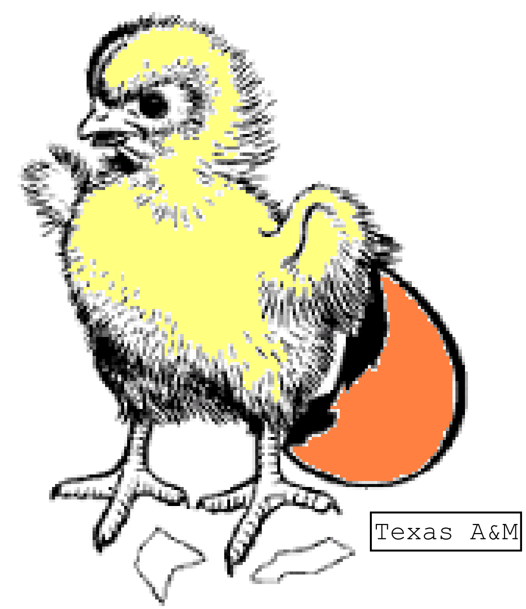

Deployment status (2)

- About 320 "production" IPv6 prefixes allocated, which mainly belong to ISPs.
  - plus ~100 legacy 6BONE prefixes
  - 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-complete.html

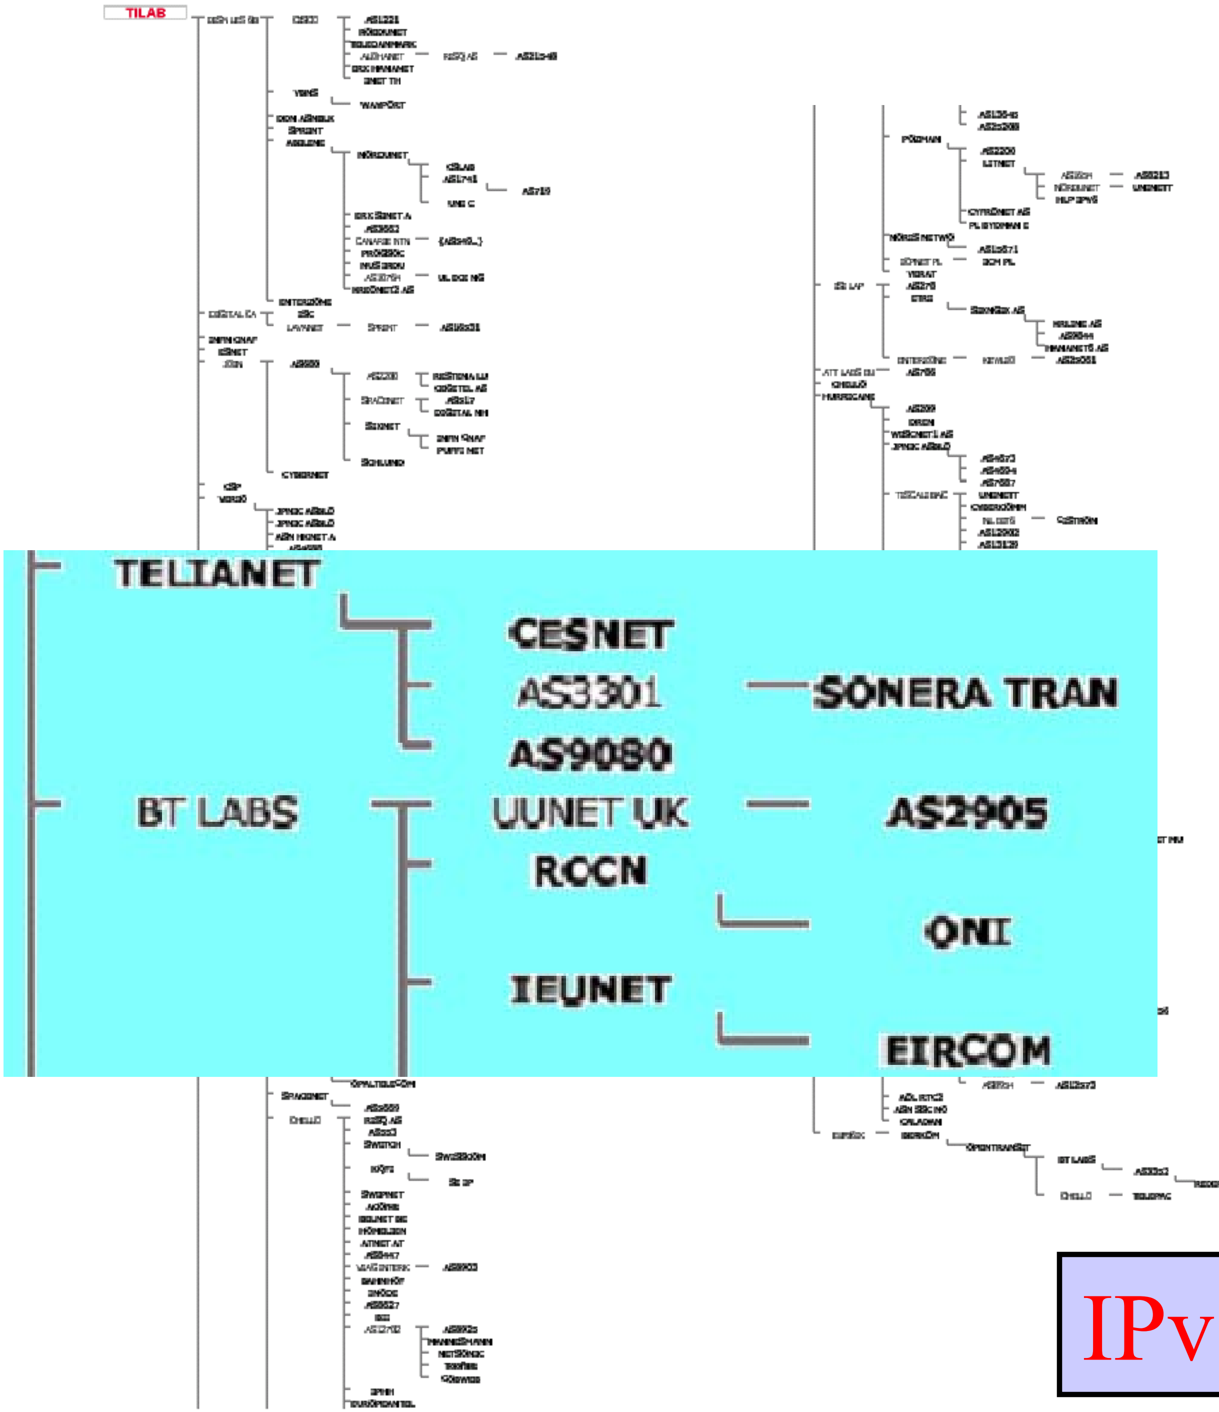

Active IPv6 topology, 2003-04-09.

Much better route aggregation than IPv4.

IPv6 is real!

Testing, testing

- 6NET is a three-year European Union funded project to demonstrate that continued growth of the Internet can be met using IPv6.
  - SURFnet participates
- It includes a work package for ***IPv6 Middleware and User Application Trials*** (led by IBM).
  - Telematica Instituut participates
- Globus is 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*

Further plans (evolving daily)

- Plan is to make more extensive tests with successive GT3 alphas, with about 10 nodes
  - Issues with IPv6 are 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

Remove the barriers to… (1)

- VoIP, p2p applications, etc.:
  - stop wasting resource on NAT beating
- 3G:
  - start with a clean addressing & routing scenario for "Internet on the run"
- Web Services, Grids & e-business in general:
  - stop using HTTP as a Trojan Horse
  - enable all nodes to be providers
  - **let e-business pervade every SME**

Remove the barriers to… (2)

- Distributed and virtual enterprises:
  - enable true end-to-end network security
  - simplify mergers & acquisitions (merging two Net 10s is a major cost; merging IT systems is an enormous cost)
  - enable massive scale Grids and generalised on demand computing: **everybody wins economies of scale as the IT market grows**

Remove the barriers to… (3)

- Enable the networked home & school
  - Entertainment becomes on-demand and largely interactive
  - Education… ditto
- **Expand the IT market into every corner of life**
    - Needs broadband, but needs addresses and transparency too (interactive groups for learning or playing require peer-to-peer transparency)

Remove the barriers to…(4)

- Encourage emerging markets
  - Only a tiny percentage of the world population have Internet access today
  - Over the next 50 years, let's aim to get to all of them: make our market **20 to 50 times bigger.** Good for business, but good for society too.
    - Needless to say, we can't do this without enough address space

The barriers are not permanent

- IPv6 is ready to roll
- Architected middleboxes (Web Services, OGSA, MIDCOM, OPES, etc) are coming
- Any2any trust models will emerge (for example as part of OGSA)

Summary

- We’ve managed to get as far as Web Services, just, with IPv4 and some kludges (NAT-beating, HTTP as a Trojan Horse).
- As growth continues, the ***Open Grid Services Architecture*** will transform the Internet into an on demand computing platform, but it too will get stuck on rough edges of NAT boxes, firewalls, and layer-violation boxes
- Let’s tear down these barriers

For student opportunties in NL

ibm.com/education/students

fellowships, education & software

Cool projects
ibm.com/nl/extremeblue
Djeevan Schiferli, schiferli@nl.ib

ibm.com/nl/jobs

Pointers

www.globus.org

www.ggf.org

www.ipv6forum.org

www.ietf.org *(for RFCs and Internet Drafts)*

brc@zurich.ibm.com