



# A Complete History of the Grid

(abridged)

**Dr Rob Baxter**

Software Development Group Manager

EPCC and NeSC

University of Edinburgh

[R.Baxter@epcc.ed.ac.uk](mailto:R.Baxter@epcc.ed.ac.uk)

---

- ▶ We have vast numbers of distributed devices
  - high performance computers, data storage facilities, visualisation equipment, scientific instruments (telescopes, electron microscopes, ...), TVs, games consoles, mobile phones, PDAs
- ▶ Connected by high-performance networks
  - with many different access modes and protocols
- ▶ Let's put 'em to work
  - for science
  - for business
  - for global collaborations in 'virtual organisations'
- ▶ What the Web is to information, the Grid is to services

- ▶ “Toolkits”
  - Globus, Avaki, ICENI
- ▶ “Schedulers”
  - Condor, Nimrod, SUN Grid Engine
  - Entropia’s DCGrid, Platform’s LSF
- ▶ “Data management”
  - Storage Resource Broker
- ▶ “Portals”
  - Grid Portal Development Kit, Unicore
- ▶ Web Services
  - WSDL, SOAP, XML and the rest

- ▶ 1967 - World's first packet-switched network, the NPL Data Network built by Donald Davies at 768 kbps
- ▶ 1969 - ARPANET built at 50 kbps
- ▶ 1972 - International Networking Group formed
- ▶ 1973 - First international connection (UCL) to ARPANET
- ▶ 1973 - Ethernet demo'ed at Xerox PARC
- ▶ 1975 - TCP over satellite
- ▶ 1976 - First Cray-1 (at LANL)
- ▶ 1979 - First multiprocessor computer at CERT-ONERA, Toulouse
- ▶ 1979 - First ICL DAP (at Imperial College)

- ▶ 1981 - Bruce Nelson defines RPC (remote procedure call)
- ▶ 1981 - First Cyber 205
- ▶ 1982 - "Internet" defined as connected TCP/IP nets
- ▶ 1984 - DNS defined
- ▶ 1984 - JANet built
- ▶ 1984 - "Neuromancer" by William Gibson
- ▶ 1984 - 1,000 Internet hosts
- ▶ 1986 - IETF created
- ▶ 1986 - NSFNET created (56 kbps backbone)
- ▶ 1986 - First Connection Machine CM-1

- ▶ 1987 - 10,000 internet hosts
- ▶ 1987 - First CM-2
- ▶ 1988 - 60,000 internet hosts
- ▶ 1988 - NSFNET backbone 1.5 Mbps (T1)
- ▶ 1988 - IRC developed by Jarkko Oikarinen
- ▶ 1988 - First Cray Y-MP
- ▶ 1988 - Condor project begins
- ▶ 1989 - 100,000 internet hosts
- ▶ 1989 - Gordon Bell prize: 400 Mflops/\$1M (Mobil/TMC)



- ▶ Earliest “cycle scavenger” from U. Wisconsin
  - designed to run, eg, across all the machines in a department
  - users see a single Condor “pool” of many machines
  - not designed for tightly-coupled parallel jobs
- ▶ Designed to cope with
  - finding appropriate resource among many different machines
  - fault tolerance, eg someone turning off their PC
- ▶ Seems ready-made for the Grid!

- ▶ Very general resource discovery system
  - machines advertise capabilities via “Classified Adverts”
  - Condor matches to user-specified requirements
  - a two-way process giving control to users and sys admins
- ▶ Robust
  - can deal with jobs that never finish
  - or machines that appear on the Grid
- ▶ Free to academics :-)
  - see [www.cs.wisc.edu/condor](http://www.cs.wisc.edu/condor)



- ▶ Condor appears to be an ideal Grid tool
  - originally limited to a Local Area Network
  - assumed all machines are in the same administrative domain
- ▶ Needs to able to
  - find out about machines in other continents
  - determine their power, capabilities and current load
  - launch and track jobs on these resources
- ▶ Globus enables exactly these features...
- ▶ ...thus Condor-G
  - Condor that knows how to submit jobs to a Grid resource via Globus

- ▶ 1990 - First commercial dialup access (world.std.com)
- ▶ 1990 - Gordon Bell prize: 800 Mflops/\$1M (Intel iPSC)\*
- ▶ 1990 - PVM project begins
- ▶ 1991 - WAIS created by Brewster Kahle at TMC
- ▶ 1991 - Gopher created by Paul Lindner and Mark P. McCahill at U. Minnesota
- ▶ 1991 - PGP released by Philip Zimmerman
- ▶ 1991 - NSFNET backbone 44.7 Mbps (T3)
- ▶ 1991 - WWW created by Tim Berners-Lee at CERN

\*(my laptop does ~400,000 Mflops/\$1M, a factor of 500 in 12 years...)

- ▶ 1991 - UK JANet goes IP
- ▶ 1991 - nCUBE running Oracle PS achieves 1,037 Tps (2x mainframe speed, 0.05x cost)
- ▶ 1991 - CORBA v1.1 introduced
- ▶ 1992 - 1,000,000 internet hosts
- ▶ 1992 - CODINE project underway
- ▶ 1992 - MPI Forum begins standardisation of IPC
- ▶ 1992 - Platform Computing Inc launch
- ▶ 1993 - Mosaic revolutionises the WWW
- ▶ 1993 - First Cray T3D
- ▶ 1993 - Legion project launch

- ▶ The Grid Object Model view
  - takes advantage of abstraction of Object Oriented model
  - everything is an object with defined access methods
  - represent files, computing resources, data storage facilities, ...
- ▶ Entire Grid looks like a single virtual machine
  - all the complexities are hidden from the user
  - user sees, for example, as single file space
- ▶ Attractive in principle
  - very difficult to implement in an efficient manner

- ▶ Originally a research project at U. Virginia
  - commercialised via Applied Meta (now Avaki Corp)
- ▶ Some successful applications
  - used to perform a metacomputing run of an ocean model
  - ran a task-farm of computational chemistry calculations
- ▶ Grid Object Model not really found favour
  - in development phase, Grid must run existing applications
  - hard to do this within an integrated environment
  - major applications are not themselves object oriented
  - however cf. OGSA...

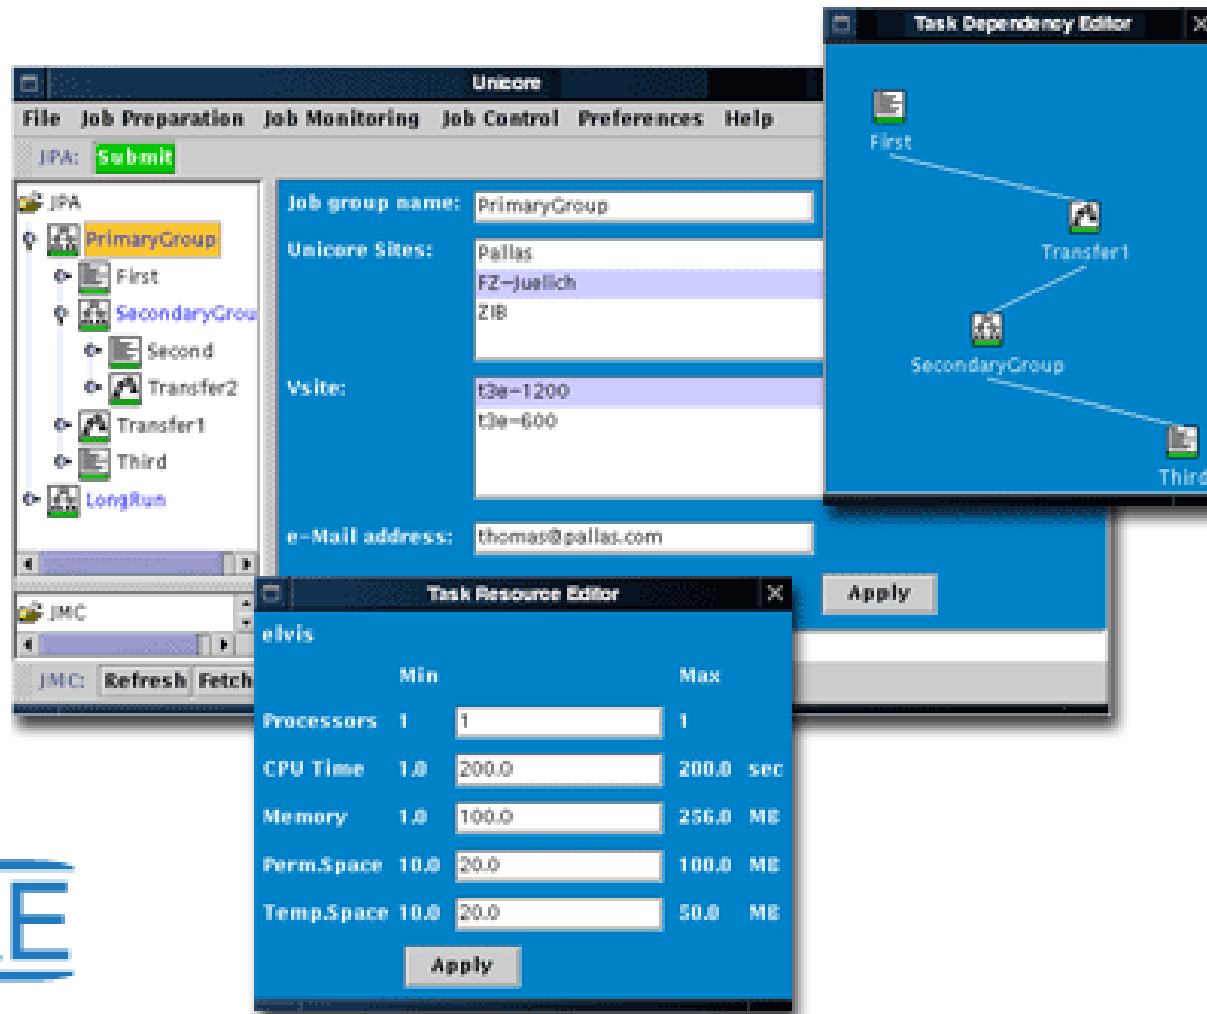
- ▶ 1994 - First commercial spam (an Arizona law firm)
- ▶ 1994 - Nimrod project launched
- ▶ 1995 - Netscape IPO 3rd largest in NASDAQ history
- ▶ 1995 - Richard White declared a munition under US arms control laws (RSA encryption program tattoo...)
- ▶ May 1995 - Java launched by SUN
- ▶ Jul 1996 - SETI@home launched
- ▶ 1997 - Globus under development
- ▶ 1997 - UNICORE project launch
- ▶ Mar 1997 - Condor deployed at NCSA
- ▶ 1997 - Entropia Inc founded to commercialise PC cycle scavenging

- ▶ David Abramson at Monash U. (Melbourne)
  - <http://www.csse.monash.edu.au/~david/nimrod/>
- ▶ Imagine the most trivial task farm
  - exactly the same code but many different input parameters
  - eg performance of an aircraft wing at many angles of attack
  - quite a common situation in many scientific applications
- ▶ Nimrod system designed to automate process
  - master process coordinates the individual tasks
  - locates free machines and sends jobs to them
  - originally developed for local area network
- ▶ Obviously applicable to the Grid

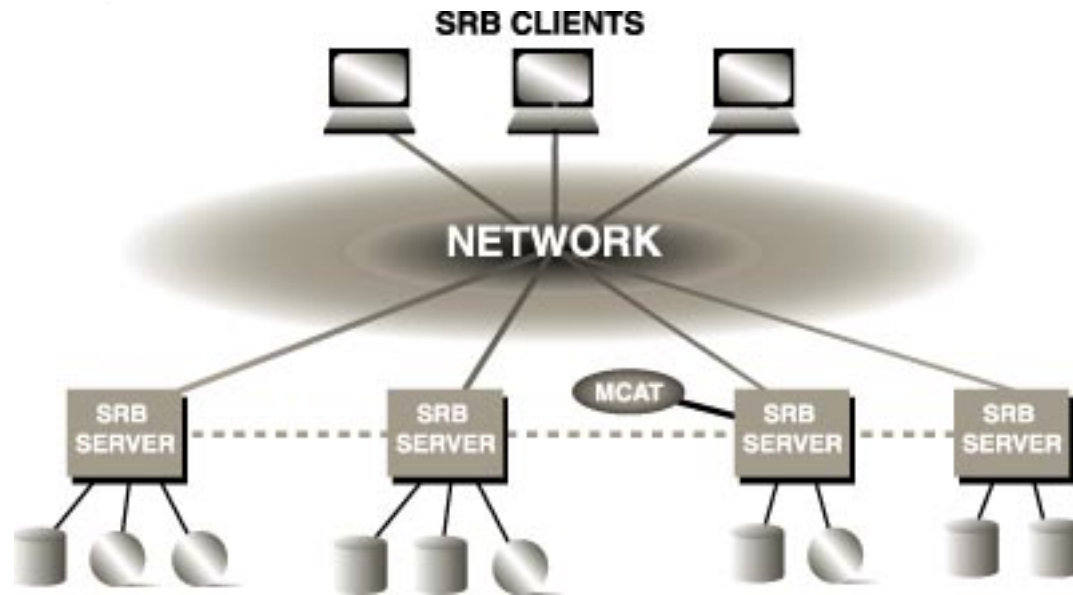
- ▶ Ported to Grid using Globus
  - major effort in developing flexible resource broker
  - discovers free resources and continually monitors job progress
- ▶ Has a concept of charge per CPU cycle
  - time on an HPC machine more expensive than a Beowulf
  - broker can consider cost as well as time
- ▶ Generally, Grid economics has not yet been addressed
  - on many people's To Do list...



- ▶ A unified environment for HPC users
  - currently in operation across German HPC centres
  - offers a single point of contact
  - written in Java for portability
- ▶ Combination of toolkit and portal model
  - like a toolkit
    - users run their own jobs - not precompiled applications
    - middleware exists to schedule jobs to HPC resources
  - like a portal
    - access is via a single GUI
    - complexities are hidden from the user
- ▶ Currently being integrated with Globus
  - the EU GRIP project



- ▶ Sep 1997 - “Building a computational Grid” workshop, Argonne National Lab
- ▶ Oct 1997 - SRB v1.0 released
- ▶ Jul 1998 - Foster/Kesselman: “The GRID book”
- ▶ Aug 1998 - Applied Meta Inc commercialises Legion
- ▶ Oct 1998 - XML 1.0 recommendation
- ▶ Oct 1998 - Globus v1.0.0 released
- ▶ Jun 1999 - Grid Forum 1
- ▶ Jan 2000 - UNICORE stage 2 launch
- ▶ Jul 2000 - SUN buys Gridware Inc → Grid Engine
- ▶ Oct 2000 - NASA IPG prototype completed

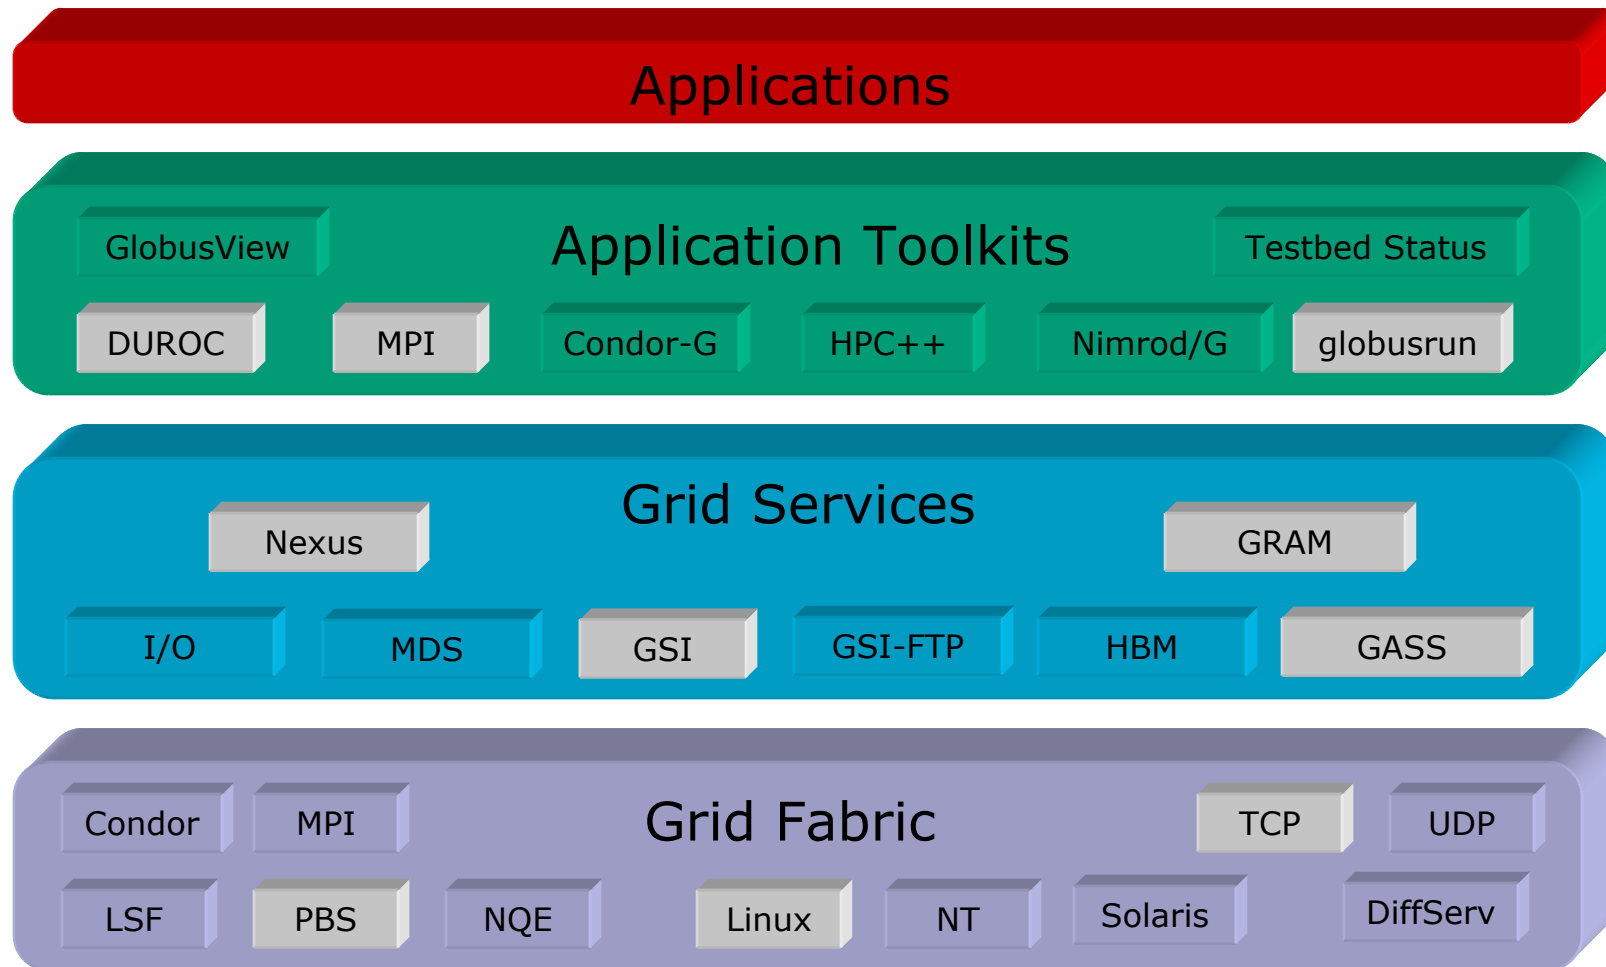


– The SDSC Storage Resource Broker (SRB)

- client-server middleware
- uniform interface for connecting to heterogeneous data resources over a network and accessing replicated data sets
- provides a way to access data sets and resources based on their attributes rather than their names or physical locations

- ▶ Science is becoming data-limited
  - spend time doing disk IO for multi-Gigabyte/Terabyte files
  - bottleneck may be retrieving data from tape or other storage
  - for HPC it is usually a bandwidth problem (always use large files)
- ▶ All major projects have to cope with
  - data format: byte order, single/double precision, binary/ASCII, ...
  - file format: order of data in file, addition of extra fields, ...
  - file transfer: must copy file to local disk for efficient IO
  - discovery: where is that file I created last year?
- ▶ SRB is one solution, OGSA-DAI (q.v.) is another
- ▶ Data will drive the development of the Grid

- ▶ The leading Grid toolkit
  - Argonne National Lab, ISI at U.C. Berkeley
  - see [www.globus.org](http://www.globus.org)
  - most Grid projects are based on Globus
- ▶ Defines and implements standards for
  - security (GSI)
  - data access and transfer (GASS and GridFTP)
  - resource discovery and allocation (GIS and GRAM)
  - remote job execution (globusrun)
- ▶ Globus operates at a very low level
  - imagine DOS versus Windows 2000



- ▶ Machine details published via GIS
  - the Grid Information Service is based on LDAP
    - Lightweight Directory Access Protocol is the X.500 database standard
  - each resource runs a Grid Resource Information Server (GRIS)
  - information collated and monitored by a network of Grid Index Information Services (GIIS)
  - data published in Resource Specification Language RSL
- ▶ Applications can query a GIIS
  - queries phrased in RSL
    - not particularly easy to use and not a standard format
    - plan to re-implement using XML



- ▶ Provided by the GSI package
  - Globus Security Infrastructure becoming very widely adopted
    - Microsoft to incorporate GSI into .Net Passport
  - based on the freely available SSLeay security package
  - uses the X.509 authentication certificates (based on PKI)
  - certificates can be signed by multiple CAs
  - adds an expiry date for additional security (*proxy certificates*)
- ▶ Enables single sign-on to the Grid
  - no need to log on to every resource with a separate password
  - user identity can be proved via a single certificate
- ▶ Pretty much adopted as the standard

- ▶ Based on CODINE
  - COmputing in DIstributed Network Environments
  - now an open source project steered by SUN
- ▶ Distributed resource management system
  - single system-image of resources via GUI
  - batch queuing across distributed resources
  - provides “cycle scavenging” a la Condor, DCGrid
- ▶ Integration with Globus?
  - Sun Data/Compute Grid project at EPCC/NeSC
  - enabling SGE to talk to Globus and to data resources
  - prototype under development as we speak

- ▶ Jan 2001 - EU DataGrid project launch
- ▶ Mar 2001 - Global Grid Forum 1
- ▶ Mar 2001 - WSDL submission to W3C
- ▶ May 2001 - XML Schema 1.0 recommendation
- ▶ Jun 2001 - Applied Meta renamed Avaki Corporation
- ▶ Jul 2001 - UK e-Science Programme launch
- ▶ Aug 2001 - US TeraGrid project launch
- ▶ Nov 2001 - GEANT, the pan-EU gigabit network, activated

## ▶ Web Service?

- “a software application accessible via Internet protocols using XML for messaging, description, and discovery” (IBM)
- what people are to HTML, web services are to XML

## ▶ Built on open standards

- Web Service Description Language (WSDL)
  - defines the interfaces a WS presents to the Internet
- Simple Object Access Protocol (SOAP)
  - defines the format of messages exchanged by WS
- XML Schema, the foundation for these and more

## ▶ Perceived as the next iteration of the Web

- backing from Microsoft (.NET), IBM, Sun, Oracle

- ▶ WSDL not yet standard
  - version 1.2 is working through W3C
- ▶ Whole range of adjunct standards too
  - WS-Inspection
  - WS-Coordination
  - WS-Transaction
  - WS-Security
  - BPEL4WS
  - ....
- ▶ Has provided a sea-change in Grid thinking...

- ▶ Jan 2002 - OGSA announced
- ▶ Feb 2002 - OGSA-DAI project launch
- ▶ Jun 2002 - NEC Earth Simulator achieves 35 Tflops
- ▶ Jul 2002 - WSDL 1.2 working draft

- ▶ **Open Grid Services Architecture**
  - a marriage of Globus and Web Services
  - announced Feb 2002, GGF 4, Toronto
  - joint work between Globus and IBM
  - major GGF standards activities
- ▶ **Data Access and Integration services for OGSA**
  - Protocols for accessing databases over OGSA
    - XML; Relational/SQL; OO; other semi-structured data sources
  - Transparent integration of multiple databases
- ▶ **Best bet for a common Grid architecture...?**
  - Malcolm will tell us shortly :-)

- ▶ The Grid is growing up
  - many possible applications
  - many problems still to be solved
- ▶ The Grid has enormous potential benefits
  - in academia and industry
- ▶ Globus is the leading Grid architecture and toolkit
  - currently Globus-2.0, but Globus-3.0 is coming...
  - do we have a real chance of a single, uniform architecture for distributed systems integration?
- ▶ 2003 will be an exciting year...