Grid Tutorial - Introduction

Adrian Cockcroft
Availability and Performance Engineering
eBay Inc.
acockcroft@ebay.com
October 1, 2004

Abstract:

Evolution of Grids in Scientific, Technical and now Commercial arenas

Analogies to clarify the concepts

Definitions of Cluster Grids, Enterprise Grids and Global Grids

Grid standards from the Global Grid Forum, Enterprise Grid Alliance and OASIS

Cluster Interconnects Overview and Performance

Dynamic Capacity Planning in a Virtualized Datacenter

Disclaimer – this information is based largely on work done while working for Sun Microsystems Inc. My current employer eBay Inc. does not endorse any of this content. eBay has its own “Grid-like” in-house datacenter management and virtualization technology which is not discussed here.

Evolution - Scientific Grids

- Core High Performance Computing market
- Universities and national research labs
- Grid platform trend starts in 1994
  - Clusters of commodity Intel PC hardware nodes
  - Linux based clusters e.g. “Beowulf”, thousands of nodes
  - Growth of capacity per node engulfs more and more applications over time
  - API based on Message Passing Interface – MPI for applications that span multiple nodes
  - High speed interconnects improve performance
- “Grid” as a software trend surfaces in 1998
  - Resource scheduling and sharing begins to move from the department to Global Grid

The Early Definition of Grid - 1998

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasiv and inexpensive access to high-end computational capabilities"

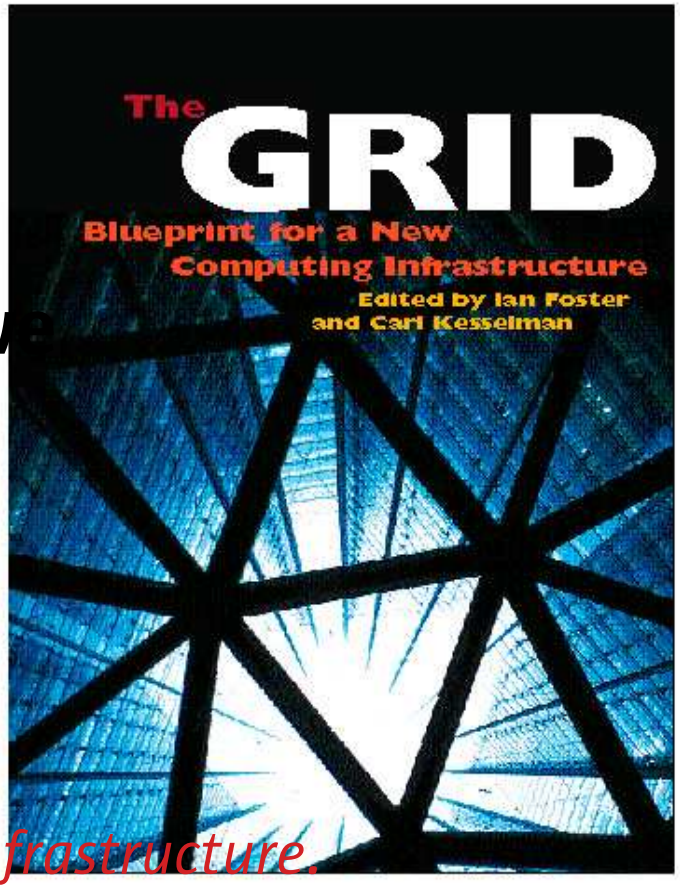

The Grid: Blueprint for a New Computing Infrastructure. Foster and Kesselman 1998

Globus Toolkit® History

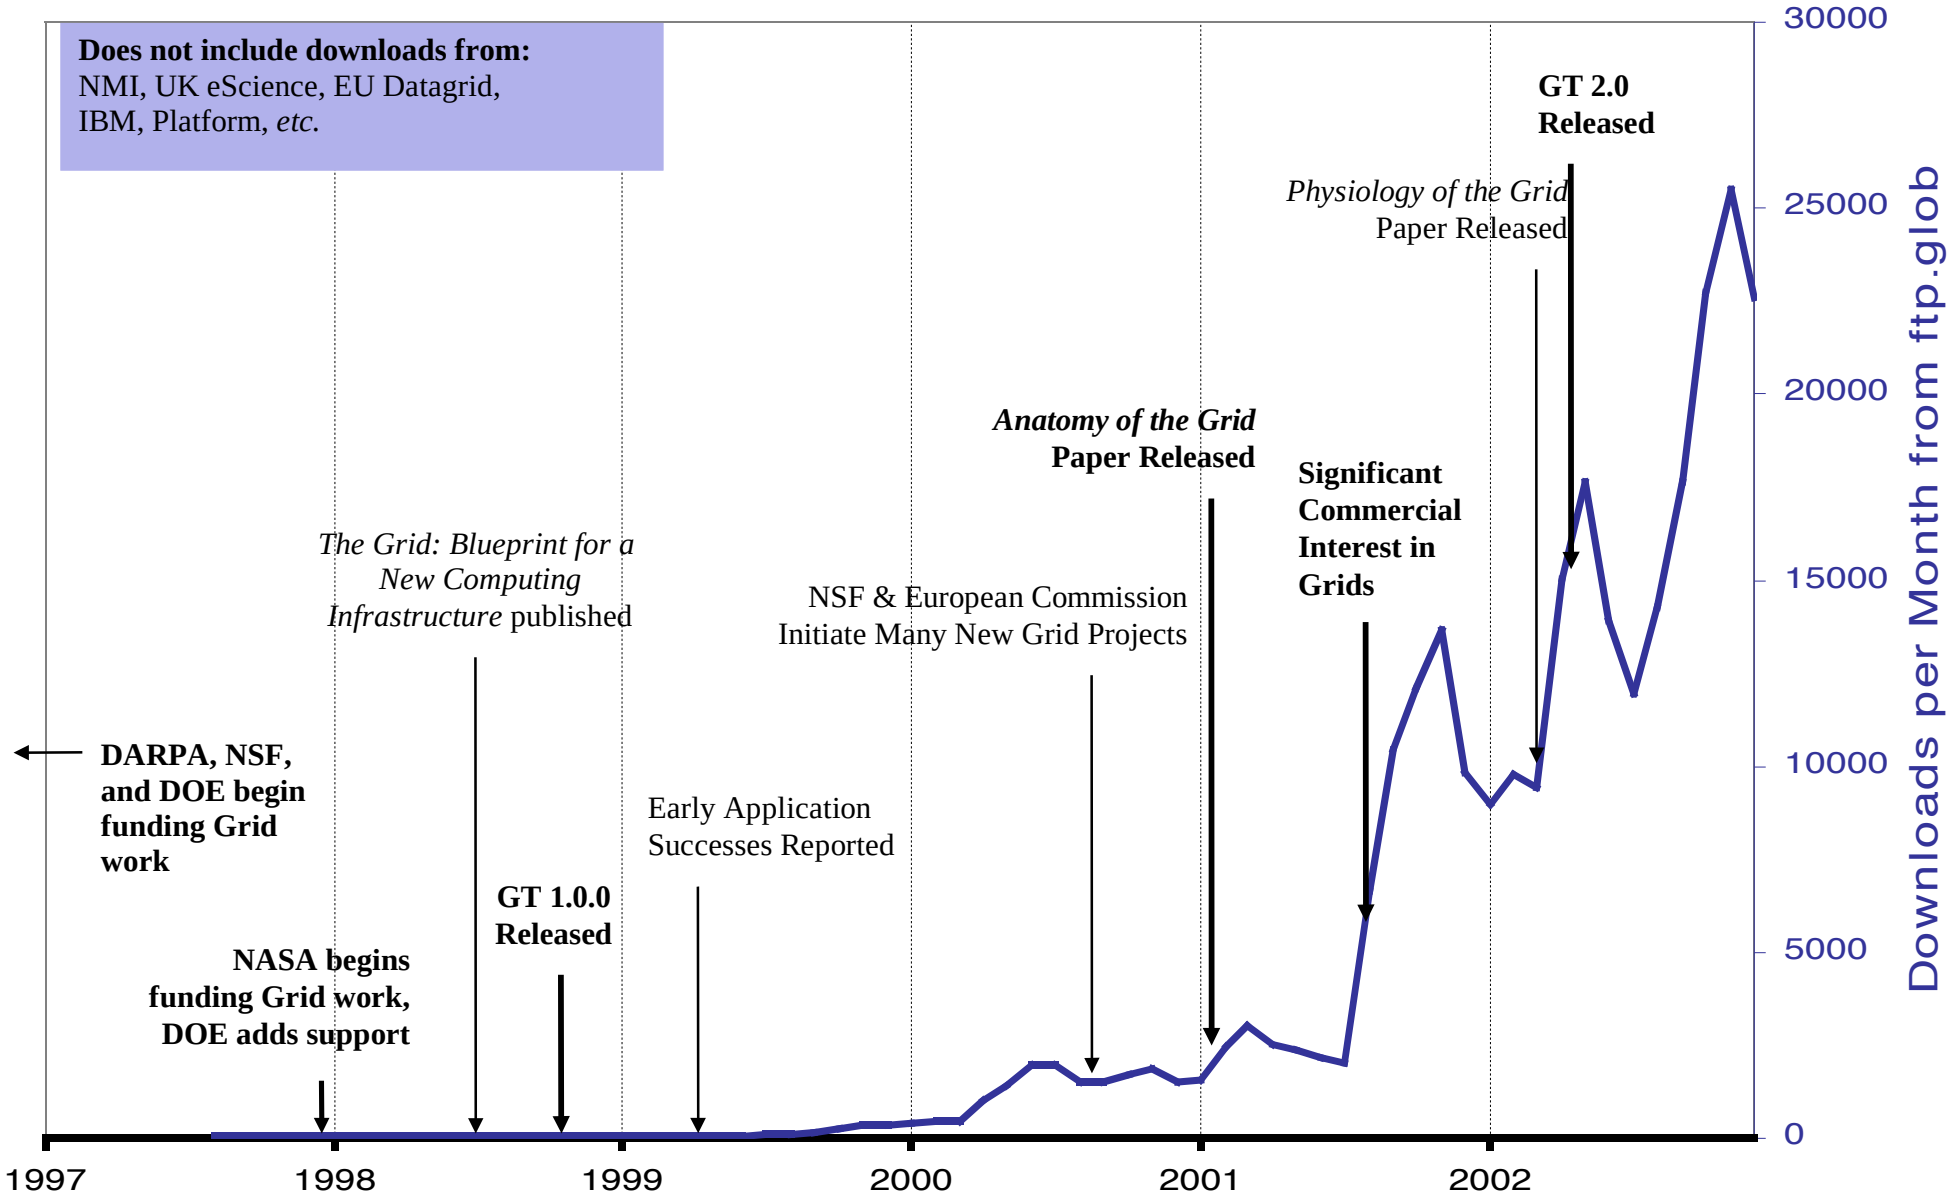

Scientific Grids Today

- Creating virtual communities that span research institutions and countries – e.g. CERN
- Solving problems that are bigger than any one datacenter can hold – e.g. Teragrid
- Cutting edge is data grid issues
- Globus Toolkit is built to solve this problem
  - Globus 2 – 2002 – quite widely deployed
  - Globus 3 – 2003/4 - based on core web services
  - Globus 4 – 2004/5 - refactored using WS standards

What is in GT4?

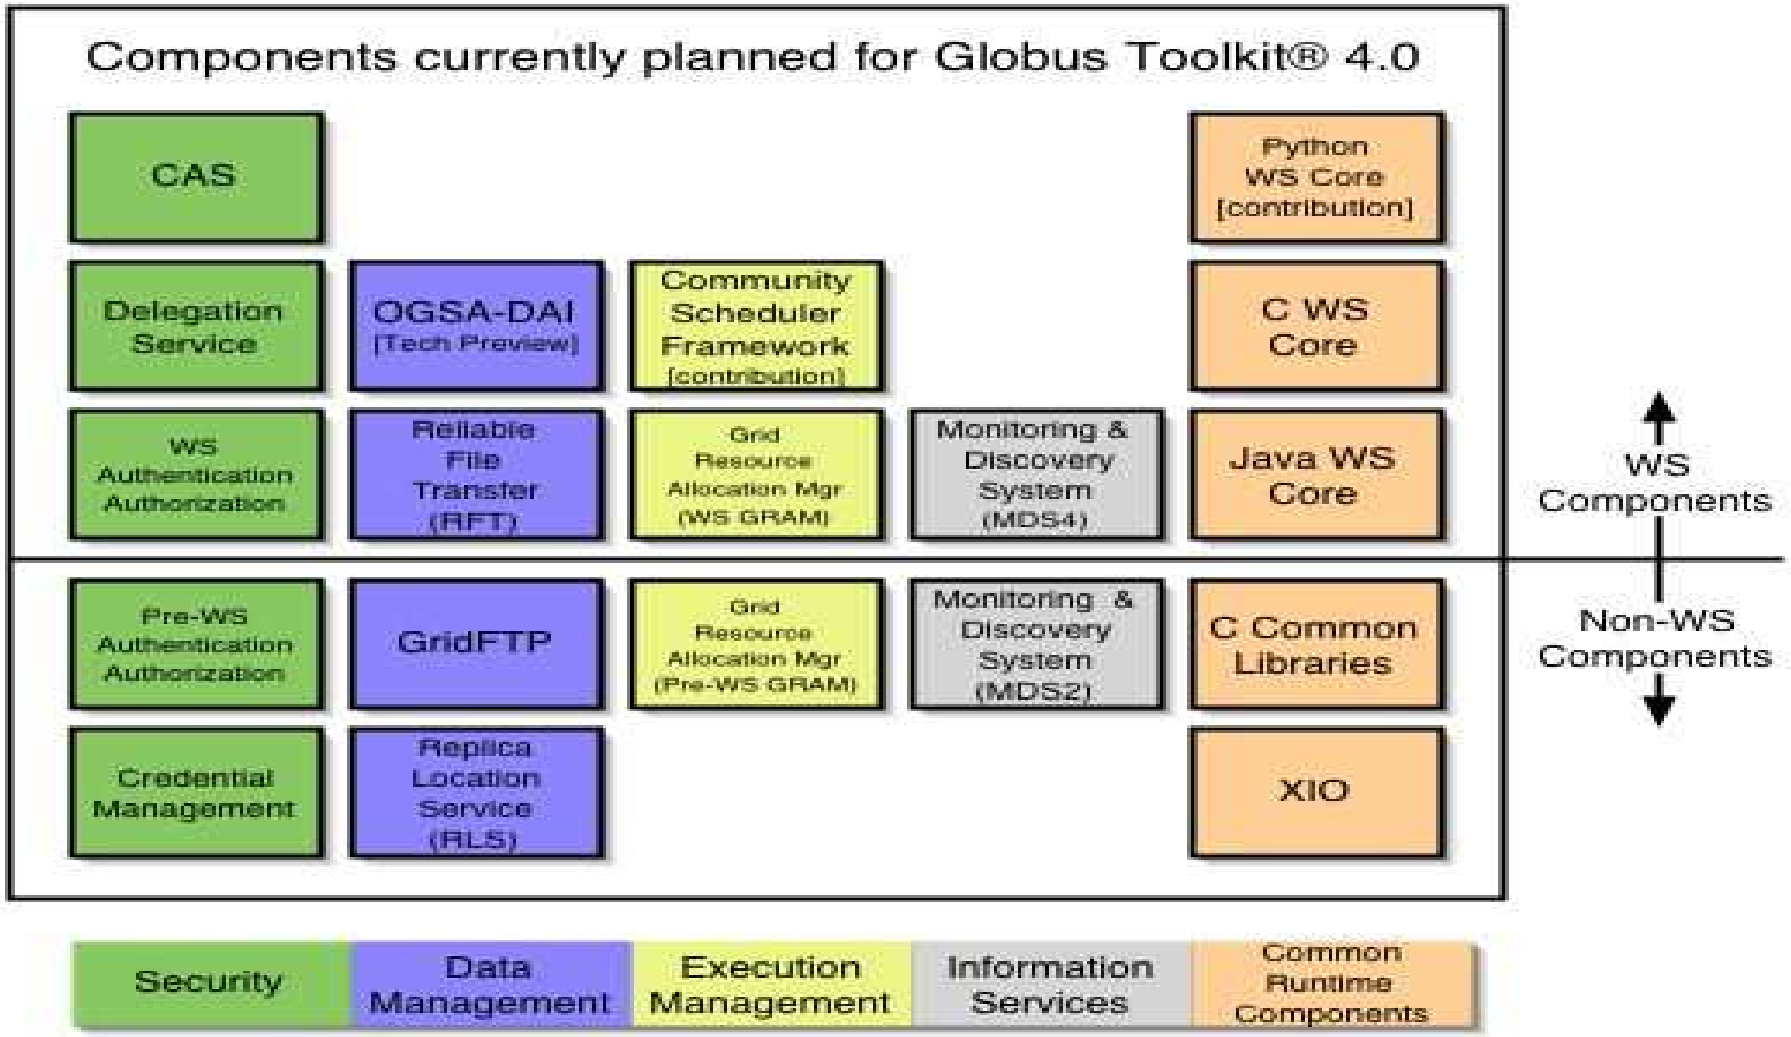

Opportunities and Problems

- Opportunities
  - Powerful solution set attractive to technical and commercial markets as well as scientific
- Problems
  - Very complex, hard to install and configure
  - Each release is a major rewrite needing new applications and introducing new concepts
  - “Religious debate” about whether to manage state explicitly divides the community
  - Fundamental trust issues, especially relating to shipping commercial/confidential data around

Evolution - Technical Grids

- Engineering driven - Energy, Biotechnology, Mechanical and Electrical design automation, and Financial optimization markets
- Cluster grids are in normal production use
- Supply chain design integration is moving from private networks to global grid
- Single node batch compute jobs, some MPI
- Batch managers LSF, GridEngine, PBM
- Relatively low usage of Globus toolkit

Evolution - Commercial Grids

- Vendor and IT driven - Datacenter automation, virtualization, utility computing and outsourcing
- Try to learn from science and engineering communities to leverage low cost hardware and sophisticated web services
- Enterprise Grid Alliance formed in 2004 to focus on commercial grid issues
- Lack of stable standards is main technology issue slowing adoption – waiting for GT4 to ship
- Trust issues are fundamental business inhibitor to inter-company resource accesses
- Many companies have their own in-house automation and virtualization technology

Analogies

- A good analogy stretches your vision
- A bad analogy distorts your vision
- Electrical Power - bad analogy!
- Telecommunications – some good analogies
  - Virtual Private Networks
  - Cellphone services

VPN Analogy

- History – Corporations owned dedicated wide area networks – private, safe, secure, expensive
- Optical fiber capacity increased to the point where one fiber could support many corporations at lower cost
- Telco's created Virtual Private Networks which multiplex traffic over a fiber – Corporations learned to trust that their own data would be safe
- The Internet provides a very low cost alternative with quality that is good enough for low-end WAN tunnel VPN usage

- Analogy: Dedicated computers grow in capacity until sharing them is more cost effective. Utility computing is the VPN of Grid, but its not yet fully trusted
- Commodity hardware based Grids are good enough for low-mid range applications

- Trend is that both Internet VPNs and Commodity hardware are moving up market fast, replacing more expensive alternatives

Cellphone Analogy

- Cellphone hardware is a commodity
- Cellphone service bundles are commodities
- Cellphone business model is where differentiation occurs
- Imagine being able to compose your own Cellphone service bundle!
  - Sprint Internet Access and Microsoft Exchange integration
  - T-mobile customer support and billing (one centralized bill for this bundle)
  - AT&T international calling plan and Cingular rollover minutes
  - Verizon reception/coverage
  - Kodak photo processing and storage
  - Accessline voicemail and addressbook
- Why can't this be done?
  - Because these services are not (yet) composable web services
  - They have no way to talk amongst themselves to let you look up an address and send a photo in an email then charge you for it
  - There are no tools to let you securely build and edit your bundle
- Analogy: the move to web services lets you compose your own service from commodity components located anywhere on the Global Grid

Definitions

- Cluster Grids
  - Racks of small commodity nodes
  - Best price/perf/density is 2-CPU Intel/AMD interconnected using Ethernet, Myrinet, Infiniband etc.
  - Many management automation tools
    - e.g. Cassat, ROCKS for scientific/engineering use
    - e.g. Sun N1-Grid, HP Adaptive Enterprise, IBM On Demand for commercial
- Enterprise Grids
  - Policy based batch scheduling that shares resources between departments and datacenters inside an organization
  - Scientific/Engineering – Globus, Platform LSF, Sun Grid Engine, PBS
  - Commercial – tends to be in-house custom framework
- Global Grids
  - Secure resource brokering and sharing over the Internet between organizations
  - e.g. Globus Toolkit, Data Synapse, United Devices

Grid standards

- Global Grid Forum
  - Based on the scientific grids community with some technical and commercial grid activity
  - Lots of research groups and working groups
- OASIS – Web Services standards
  - WS-Resource Framework – GT4 is an implementation
- DMTF – Utility Computing Standards
  - Solve for permutational complexity by making standard component management models
- Enterprise Grid Alliance
  - Focused on the needs of commercial Grid adoption

GGF and EGA

- Joint presentation made at GGF12 Sept 2004 available at http://www.gridalliance.org

[since I was Sun's Technical Steering Committee representative during EGA's founding and attended several GGF meetings I will present these slides at this point with my own commentary and insight]

Summary

- Recap
  - Global Grid Technology is maturing in large scale use for the scientific community
  - Cluster Grids are normal practice in engineering
  - Commercial users are adopting cluster grid techniques and web services but are using a wide variety of toolkits and home grown methods
  - The nirvana of commercial global grids has not emerged yet, the technology should be standardized (WS-RF) and stable (GT4) later in 2005
- Following discussions:
  - Interconnect technologies are the key “glue” to build large scale systems out of commodity nodes
  - Capacity management in a very dynamic environment is a major challenge for analysts and tools vendors