

CERN openlab for DataGrid applications

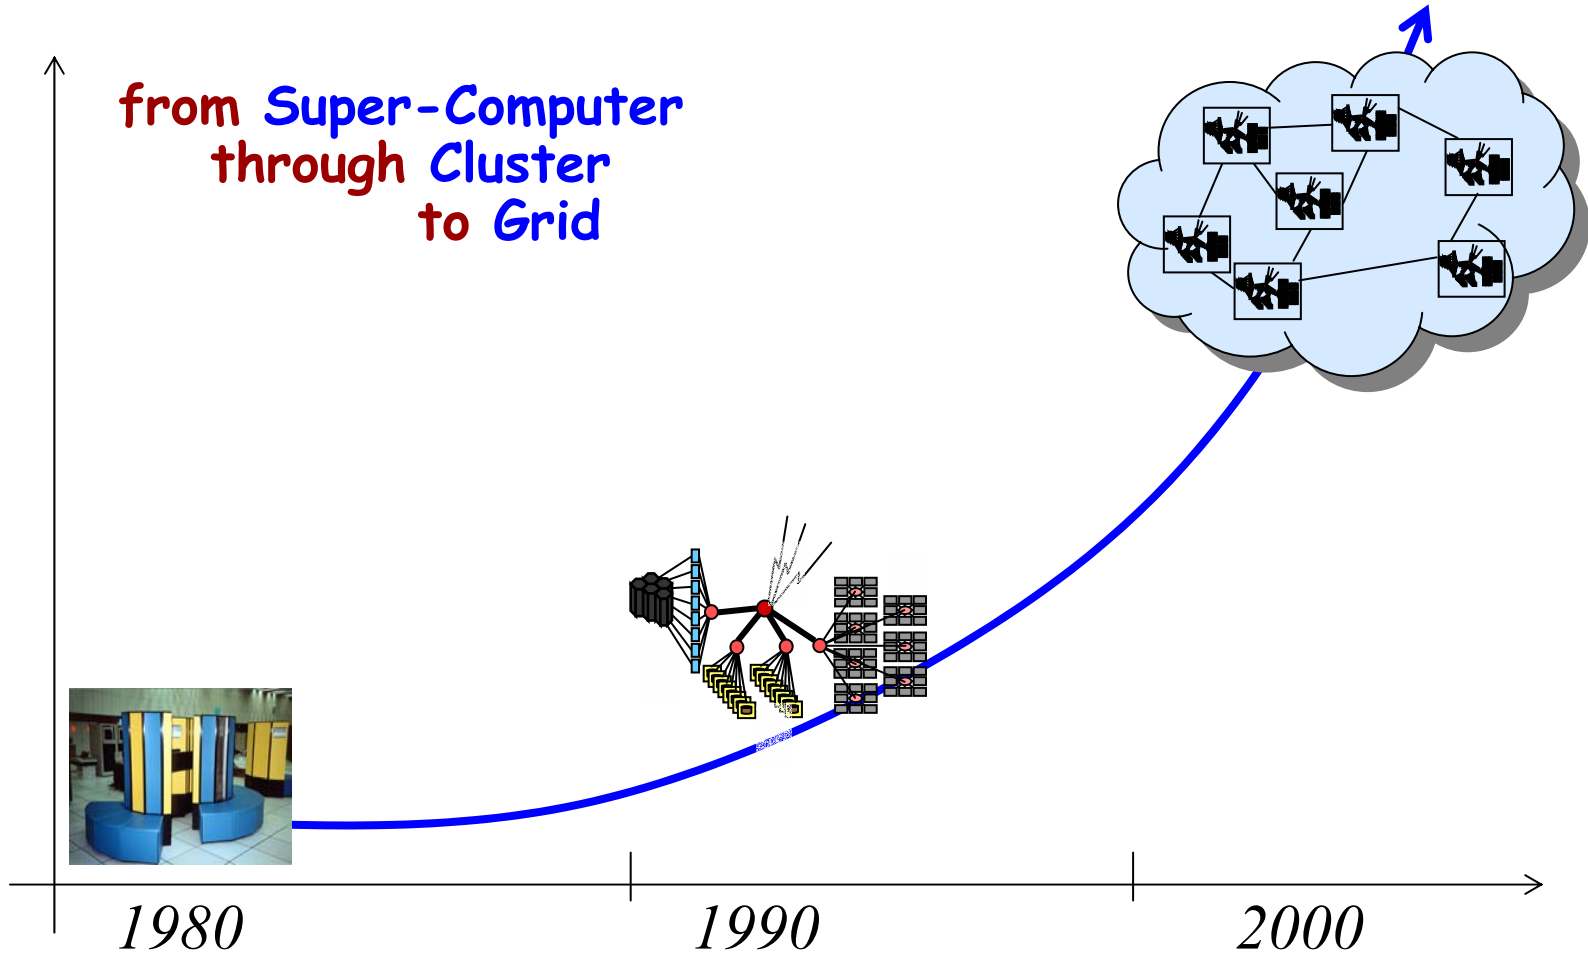
2nd Workshop on Fabric Management
July 7-8, 2003

Welcome !

GRIDs - The next step in the evolution of large scale computing

performance, capacity

from Super-Computer
through Cluster
to Grid



European Commission for Nuclear Research

CERN "Fabric Management" history

- 1980's: Mainframe era
 - ~5-10 analysts or system programmers + help of the company on site
 - Dedicated room for ... system manuals
 - The network is the proprietary backplane/switch/interconnect
- 1990's: Unix RISC -> PC Linux
 - 1 system administrator/ 100's machines
 - Started to invest in automation/recovery
 - Outsourcing model (MCHF/Year) – now being insourced
 - High speed multivendor local area network (Gigabit networking as of 1990!)
- 2000's: Grid's
 - Must go one step further in automation/configuration...
 - Complex infrastructure problems
 - Operation centers
 - High Speed/high latency multivendor wide area networks

Integrated Management

- 1994: Proposal to install a Convex Exemplar
 - Was ~100% HP/UX compatible
 - Single system image
 - Single Management
 - In theory 1 system administrator for 1000 nodes
 - But 2-3 times more expensive than HP Kittyhawks
 - Project finally abandoned but triggered the adoption of outsourcing

Integrated Management (2)

- 1994: IBM SP2
 - Not single system image
 - Very similar to a single-vendor Unix/RISC distributed system
 - Integrated Management
 - Set of (customizable) tools to give a single system management *view*
 - A few people to manage the 64-node system (but probably could expand with many more nodes without noticeable manpower increase)
 - Things started to become difficult as we were trying to integrate non-IBM equipment (e.g. standard SCSI disks, difficult to enable TAG queing)
 - More expensive than commodity RISC (~ x2)

Integrated Management (3)

- 1998: SGI Origin series
 - Single system image
 - 1 system administrator/1024 CPU's
 - But actually ~ 20 CPU's installed at CERN
 - Complex performance issues
 - Poor support from the company
 - Too expensive

Today

- 1000's of PC/Linux
- Complex Interdependencies
 - Still not able to shutdown the computer centre in 5 minutes
 - Took > 1 day to recover from last Computer Centre power failure
- But ...
 - Able to reinstall all systems overnight
 - EDG WP4 tools in real use

Monitoring

- Cray: 6 monitoring consoles/machine
- IBM: could call IBM in Montpellier
- Distributed Unix/Linux
 - Sure 1/2/3, PEM, PVSS, EDG WP4
- Management Suites
 - Looked at CA Unicenter, IBM Tivoli, BMC
 - Too complex, too expensive
- New initiatives
 - Intel WfM, PXE
 - DMI (DMTF)/WBEM
 - HP UDC, IBM Eliza, SUN N1

Reality check

Logged Events for FULL - Message (HTML)

File Edit View Insert Format Tools Actions Help Type a question for help

Reply Reply to All Forward         

From: owner-service-logger-full@listbox.cern.ch on behalf of Logger.Support@cern.ch

Sent: Mon 7/7/2003 8:00 AM

To: service-logger-full (All events logged by operators)

Cc:

Subject: Logged Events for FULL

edanis	04 JUL 2003 22:00	lxbatch273	41621 S	Assigned	NO_CONTACT alarm received. Able to ping. Unable to connect. Magic key ok. Back in service.
ggrandon	05 JUL 2003 0:03	lxcvs02	41624 S	Assigned	NO CONTACT alarm, unable to ping. Black screen, machine reseted. Machine back in prod.
ggrandon	05 JUL 2003 2:09	web13	41627 L	New	BOOTED alarm. Taken for A. Wagner, in site. Rebooted because Web-server was not responding
ggrandon	05 JUL 2003 2:11	lxb0440	41628 S	Assigned	NO CONTACT alarm, unable to ping. On console: unexpected IRQ trap at vector d9. No keyboard response. Machine reseted. Back in prod.
ggrandon	05 JUL 2003 4:17	lxcvs02	41629 S	Assigned	cvsweb_httpd_fail and cvsudp_fail and local filesystem and root filesystem.....alarms. Machine rebooted....back in prod. All alarms disappears.
edanis	05 JUL 2003 8:24	lxbatch309	41630 S	Assigned	NO_CONTACT. Pingable. Unconnectable. Msg 'Uhhuh. NMI received. Dazed and confused, but trying to continue', 'You probably have a hardware problem with your RAM chips'. Magic key failed on DUMP. Reset. Need to manually fsck for /usr/vice/cache. Left down.
edanis	05 JUL 2003 11:13	lxbatch469	41632 S	Assigned	NO_CONTACT. Able to ping. Unable to connect. Msg a lot off call trace. /usr/vice/cache need manual fsck. Left down.
edanis	05 JUL 2003 12:02	lxbatch277	41634 L	New	NO_CONTACT. Able to ping Unable to connect.
ggrandon	05 JUL 2003 22:30	lxbatch476	41658 L	New	POOL_FULL alarm, 91%/pool. Procedure done. After pool occupancy 32%. See detail.
ggrandon	05 JUL 2003 22:42	lxbatch575	41659 S	Assigned	NO CONTACT alarm, unable to ping. On console: unexpected IRQ trap at vector d9. No keyboard response. Machine reseted. Back in prod.
ggrandon	06 JUL 2003 1:55	lxbatch032	41660 S	Assigned	SBATCHD_DEAD alarm

What you would like to see

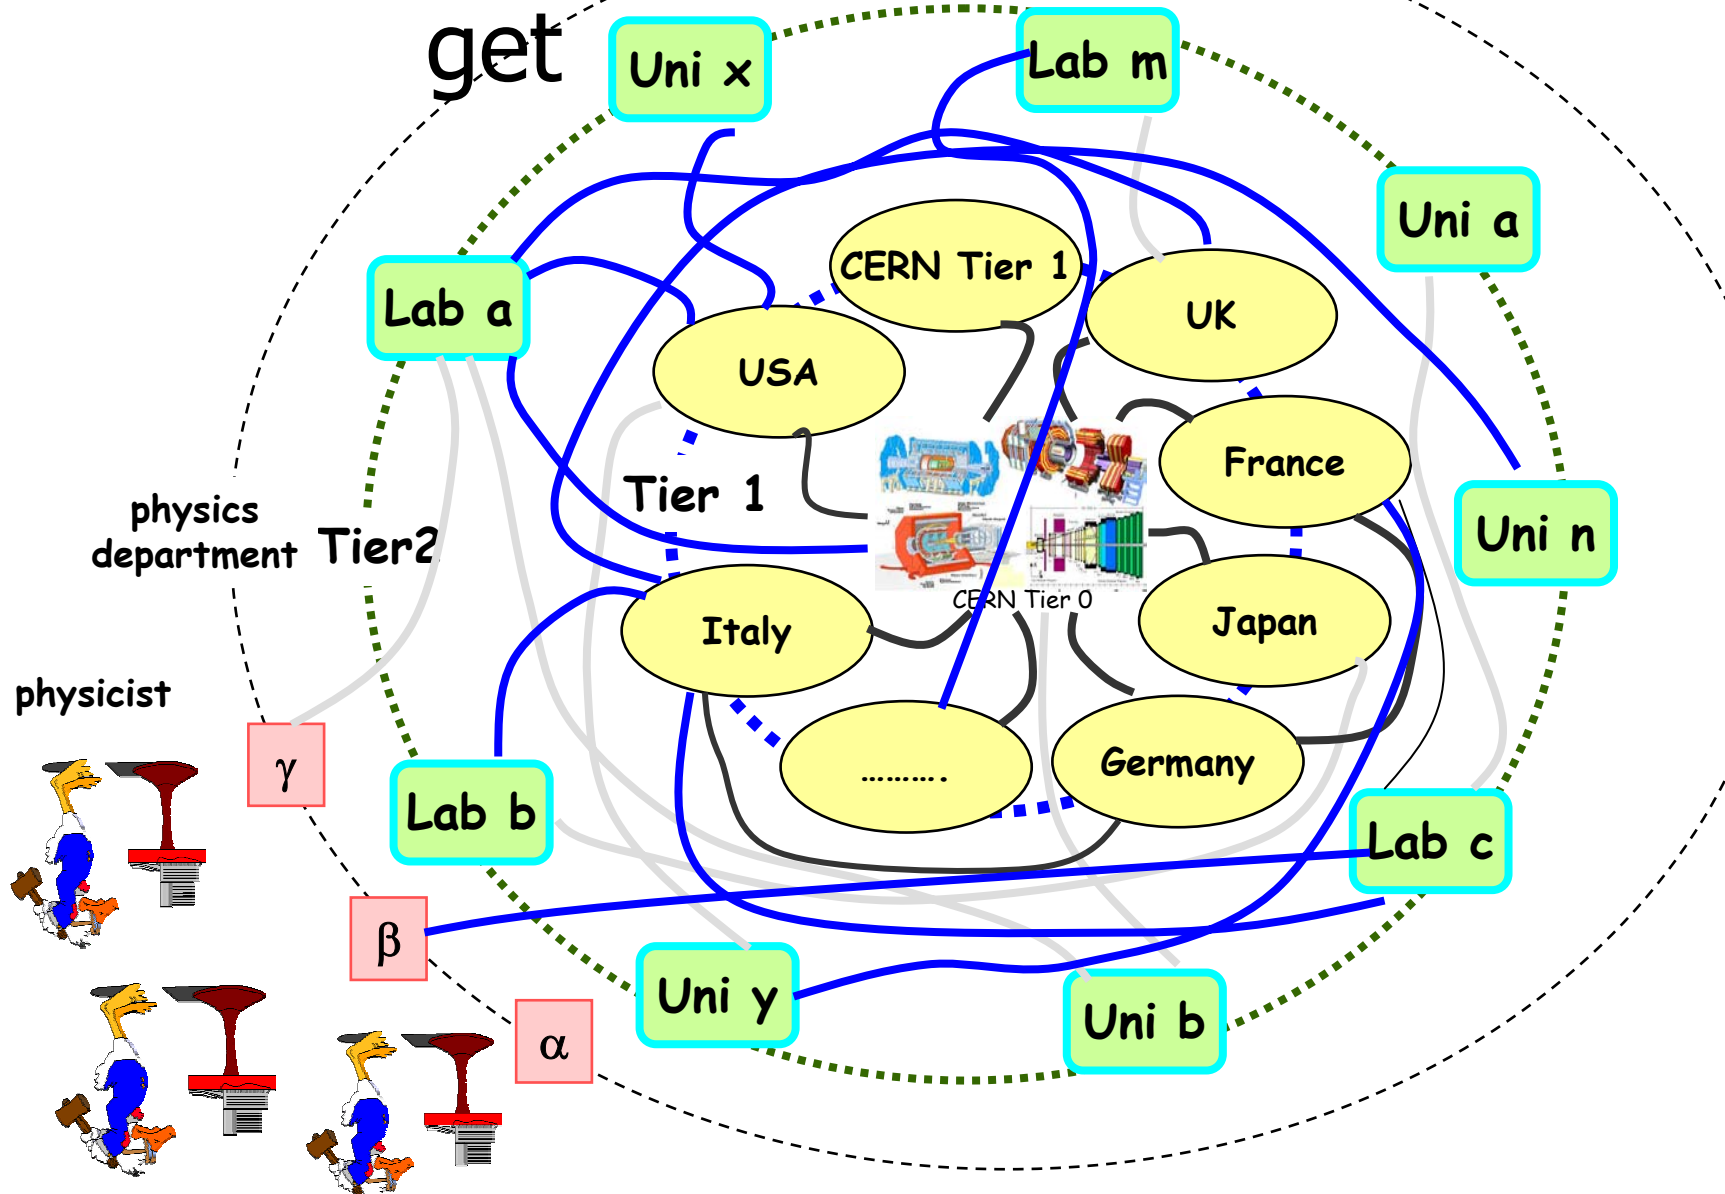


reliable
available
powerful
calm
cool
easy to use

.... and nice to look at

European Organization for Nuclear Research

What you get



Summary

- Fabric Management is an issue we need to look at
 - The problem has been around for > 10 years
 - Will be of importance for most Tier 0/1 centers
 - New ideas/initiatives are popping up
 - Costs are not negligible
 - E.g. infrastructure costs
 - Manpower costs even more important
 - Will Launch a TCO study later in the year

Let's try to address some of the issues in the next two days..