



HEPiX From the Beginning

Alan Silverman

28th October 2011

HEPiX Vancouver 2011



Agenda

- The birth
- Meetings, some highlights
- Working Groups
- Some successes, or at least some successes in which HEPiX played an important role
- Are we value for money



The Birth



- CHEP Tsukuba, Japan, March 1991
- Four lost “mainframe” people in the land of Unix
Les Cottrell/SLAC, Judy Nicholls/FNAL and
Harry Renshall/CERN from VM-land
and Alan Silverman/CERN from VMS-land.



Anyone seen my Apollo?

- Experience of previous user groups such as
HEPVM and DECUS
- Leverage this experience to learn more
about how to deal with Unix in all its flavours



man, grep, chmod, ... ???

Les's presentation to CHEP - 1

- Independent activities are in progress at many HEP sites to understand and define how to support Unix
- At this meeting, people are ... recognising common challenges & solutions in supporting Unix in the HEP community
- We are learning much from one another & feel there is much to gain from further sharing of experiences and from coordinating or aligning our Unix activities
- Most of those of us who have discussed this at this meeting feel that HEP is at an appropriate place in the Unix learning curve that it would be extremely valuable to:-
 - **Share Unix experience** between HEP sites active in Unix
 - **Encourage alignment** in the Unix user environment in the HEP community
- There is a critical need for mechanisms to be set up to provide information exchange:-
 - Within HEP via **real face-to-face meetings**, e-mail, netnews, info-servers, remote conferencing, documentation, newsletters
 - With vendors – make sure they know about our needs
 - With standards bodies – encourage HEP representation on critical bodies

Les's presentation to CHEP - 2

Proposal – set up a HEP Unix coordination group

Brief

- Provide information exchange
- Promote alignment of the Unix environment in the HEP community

Recommended regular meetings

- Half yearly
- Aligned with CHEP & other meetings of relevance
- 1-2 people from HEP sites working on providing a Unix environment to users

Name of group – HEPiX

First meeting – FNAL, 1st week of October 1991

Meetings summary

- This is the 44th meeting
- Hosted by
 - 8 sites in North America
 - 18 in Europe
 - 1 in South America (Rio, after CHEP)
 - 1 in Asia (Taiwan)
- 14 countries* visited

* Where Scotland and England are different countries

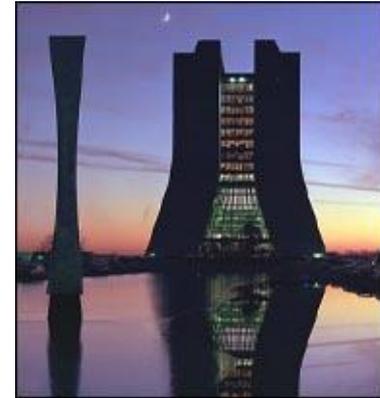


Have HEPiX, will travel

Some Meeting Highlights - 1

- First meeting: FNAL, Sep 1991:
 - 37 attendees
 - 17 sites
 - Spread over 2½ days
 - Attendees included Sandy, Les, Corrie, Rochelle
 - Lots of ideas (e.g. half-yearly meetings) but not so much follow-through.

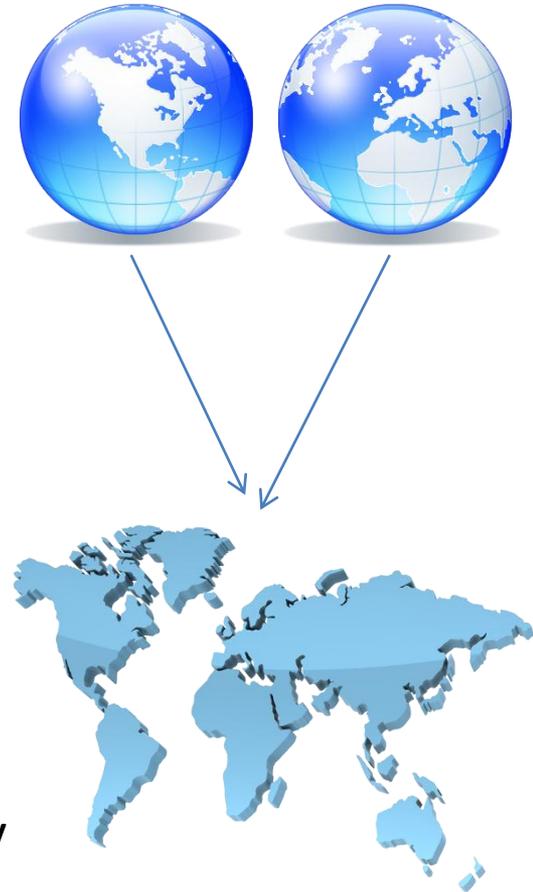
- CERN, Sep 1992:
 - 65 attendees
 - 30 sites
 - Agreed to go with regional groupings (Europe and North America) and each regions should meet every 6 months
 - Judy Nicholls takes the American chair
 - Alan takes on European chair



Spot the deliberate mistake

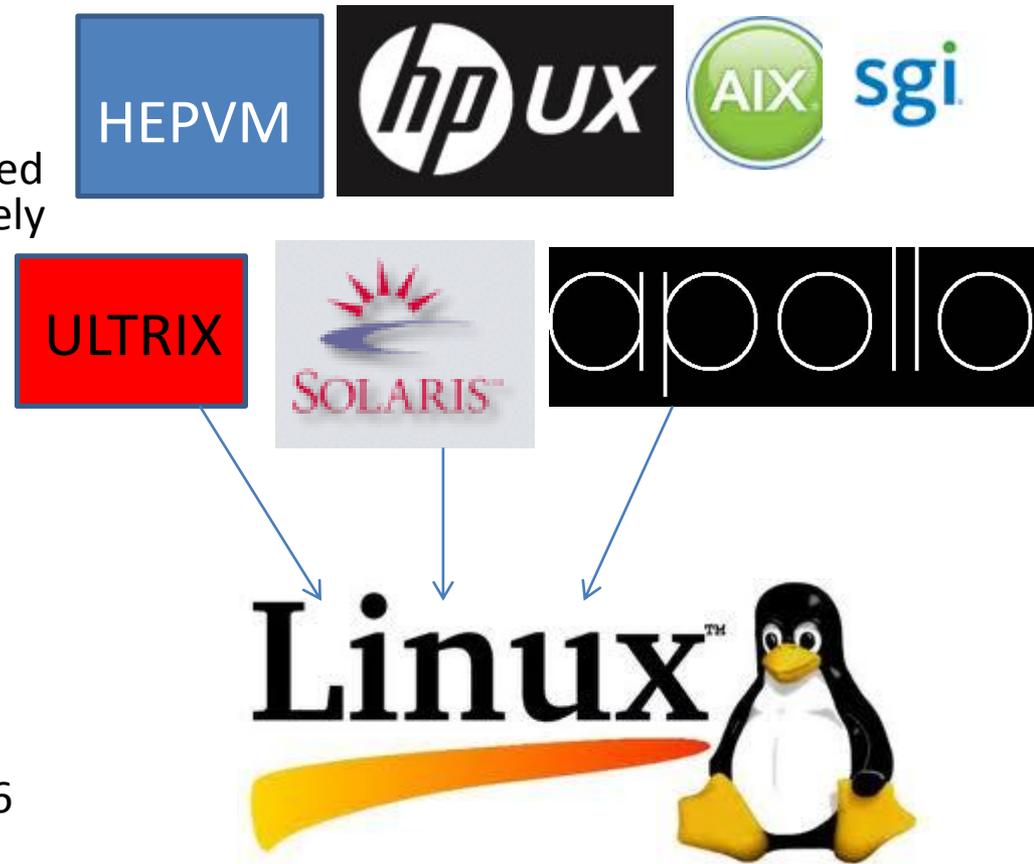
Some Meeting Highlights - 2

- Parallel meetings for the next 2 years until one universal meeting in Prague in May 1995. Agreed to merge the regions and Alan appointed chairman.
- Next meeting was a 3 day affair the week following CHEP in Rio in September 1995.
- From then on, every 6 months on opposite sides of the Atlantic apart from one excursion to Asia (Taiwan, Oct 2008).
- Along the way, it was decided to stop having meetings next to CHEP because people were absent too long from their work
- Also it was agreed to replace a single chair with two coordinators, one in North America, currently Sandy, and one in Europe, currently Michel.



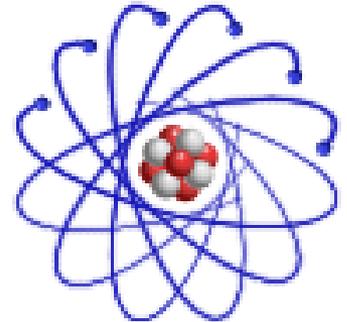
Some Meeting Highlights - 3

- Some of the early meetings (e.g. at NIKHEF in April 1993) were shared with a HEPVM audience, the so-called HEPMIX meetings. But VM had largely disappeared as a topic by 1995-6, except as some VM tools and processes were being migrated to UNIX.
- During the 90s, HEPiX covered all flavours of UNIX although pressure grew steadily to adopt Linux as the preferred HEP standard O/S. Many heated discussions took place, for example at CASPUR in Rome in 1996 (Andrei versus Alan in particular).



Some Meeting Highlights - 4

- In 1997-8 we merged with a HEPNT working group to discuss Windows-NT topics and at the SLAC meeting in Oct 99 we renamed ourselves HEPiX-HEPNT. But Windows was always a “junior” player in HEP and HEPNT disappeared from the organisation’s name in May 2004 at Edinburgh (although we still cover the subject in a few talks each meeting)
- The other notable event which took place in Edinburgh was the emergence of Scientific Linux as a possible standard for the HEP preferred Linux flavour.
- Furthest south meeting – Rio
- Furthest north – Umea, Sweden (nearly inside the Arctic circle)
- Furthest west – TRIUMF (just, over SLAC)
- Furthest east - Taiwan



Some Personal Highlights

- My first meeting (apart from the birth in Japan) was at the CERN meeting in Sep 1992.
- My first HEPiX trip report was from the NIKHEF meeting in April 1993 – only 12 pages! But the HEPiX part was only 2 days!!

HEPiX Nikhef Meeting Minutes

Alan Silverman

19-20 April , 1993

These are the minutes of the HEPiX Europe meeting held on April 19th and 20th in NIKHEF, Amsterdam. Overheads of the speakers are available for most sessions, either from the author of the session; or via WWW

Some Personal Highlights - 2

- The one I really missed – Braunschweig, April 2000, where Lisa pronounced my epitaph (and made at least two people almost believe enough to call me the next day to check I was still alive)
- The one I was present at – Lisbon, April 2010, where the Icelandic volcanic ash cloud closed European airspace and the meeting had to be overnight converted into a video-conference event – many congratulations to Michel and the LIP team for a successful event, and commiserations to the meeting organisers for all their hard work to prepare it



Some Personal Highlights - 3

- The most dangerous – Rio – went there, spent the weekend between CHEP and HEPiX sightseeing and did not get mugged or robbed!
- The most secure – BNL, armed guards at the gate!
- The most exotic – Taipei – I only just survived all Simon's banquets!!

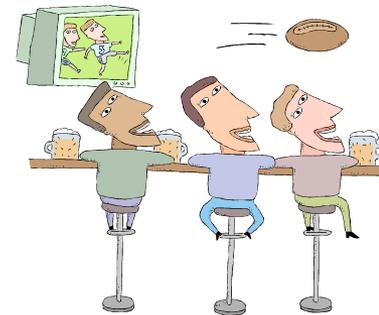


Maracanã
football stadium



Some Personal Highlights - 4

- My favourite HEPiX destination – Vancouver; this is my third HEPiX here (plus a CHEP in Victoria)
- My least favourites – Fermilab or RAL in the winter – the wind!!!
- HEPiX “traditions”
 - HQ hotel for the “hard core”
 - “Meet in the bar at 7”
 - “Table for 20” (it’s always 20)
 - Fine dining (pushed by Tony, not me)
 - “don’t give Alan the wine list” (Lisa and Sandy)



Working groups - 1

- HEPiX Scripts – DESY + CERN
 - Original idea began at DESY in 1992-3
 - a DESY/CERN collaboration working at CERN produced a production version which then spread across European HEP
 - used at DESY until 2007-8 and still used at CERN
- AFS (CERN, CASPUR and many others) – how best to use it; negotiations with Transarc

Many CERN/IT managers have tried to replace it, without success.

Size of the current CERN installation is –

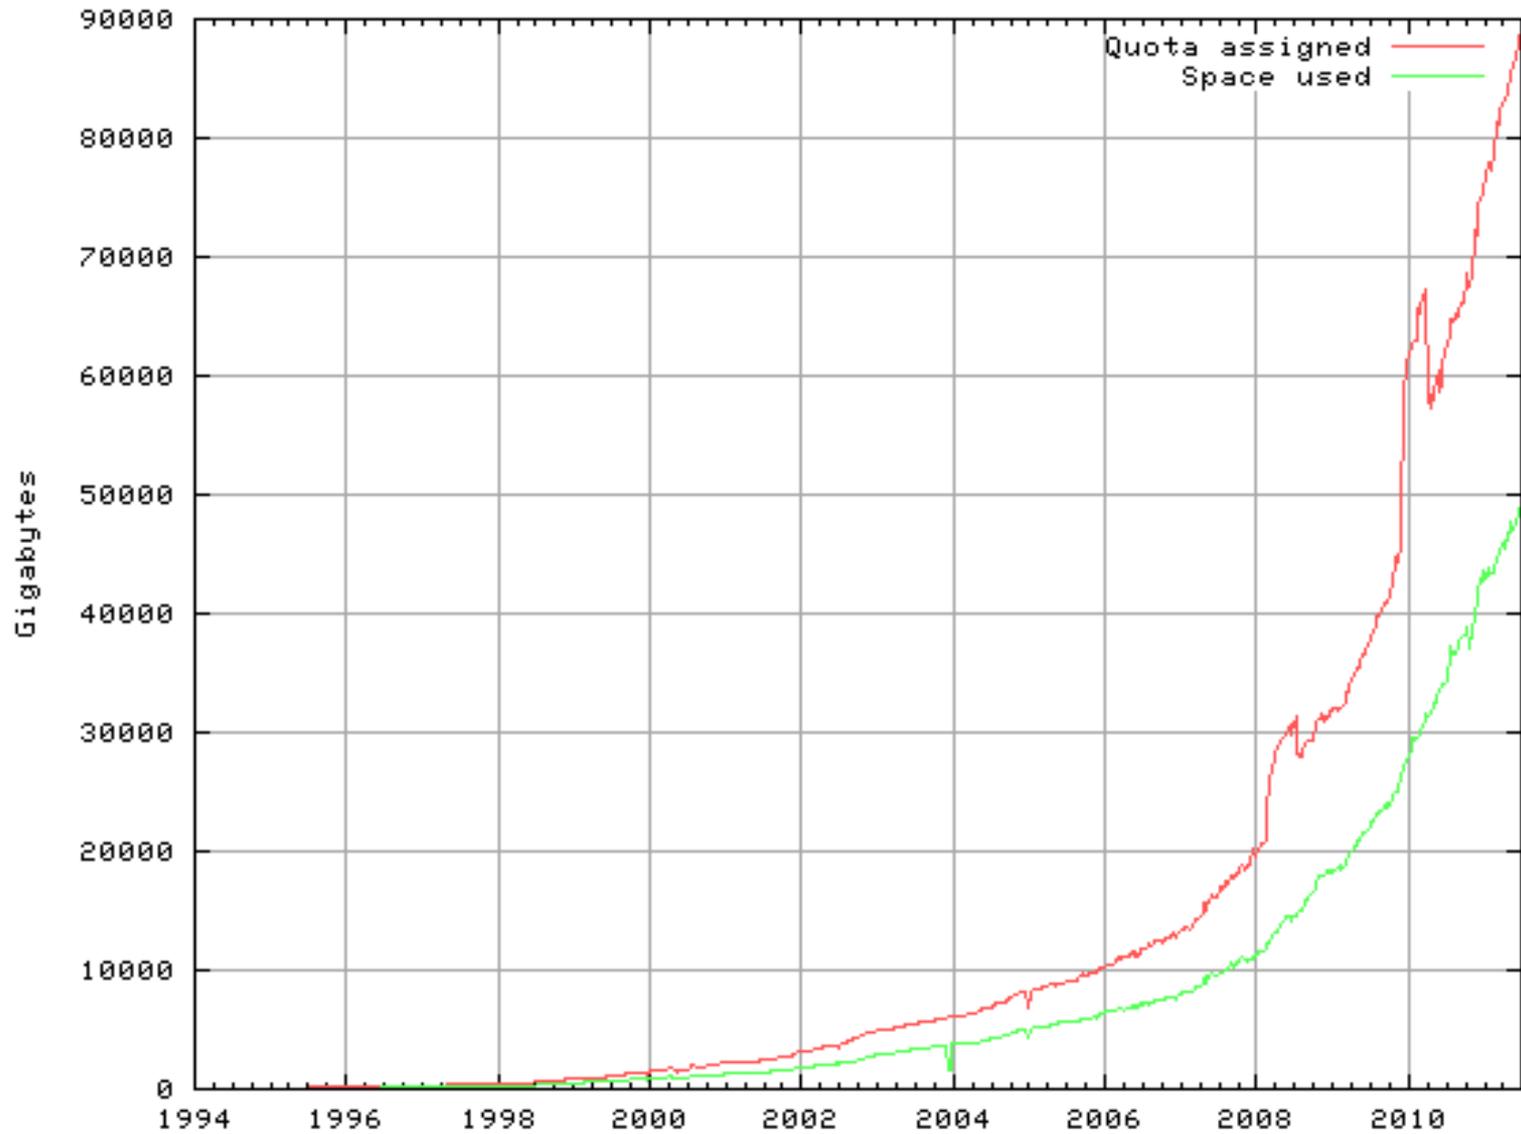
 - 50TB on 50 file servers, 750 million files
 - 31,000 home directories
 - 15,000 client nodes connecting weekly,
 - one-third from off-site
 - 5-6000 active users per week

```

$HOME/.hepixonkbd.NCD
#####
! File: $HOME/.hepixonkbd.NCD
! Description: This file contains
!   xmodmap modifiers.
! Author:  Arnaud Taddei
! Date:    03 october 1995
  
```



AFS Space Usage



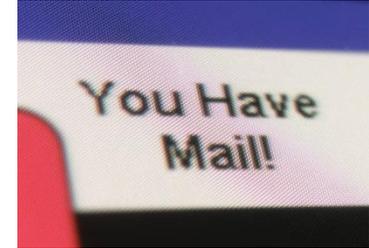
Working groups - 2

- Large Cluster SIG – created in 2000 by CERN and FNAL
 - Performed various site surveys (chip selections, operations procedures, use of video-conferencing, LSF usage, anti-virus security methods)
 - Two workshops co-convened by Dane Skow, then of FNAL, and Alan
 - Examine issues around large clusters
 - tools, procedures, useful guides and experiences.
 - First held in Fermilab May 2001 attracted 60 attendees over 4 days
 - Follow-up in Oct 2002 brought in over 100 people over 2 days
 - Results reported to CHEP 2003



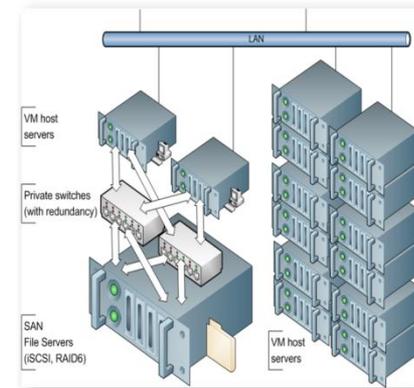
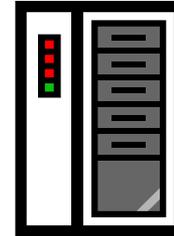
Working groups - 3

- Mail (CERN, Taddei) – workshop on mail (at CERN), experience sharing
- X11 scripts (DESY, CERN) – define “best” X11 environment for HEP
- Security – we created a restricted mailing list (to site security teams) to keep sites informed of intrusions; it was disbanded when more formal EGEE/OSG security initiatives took over
- Benchmarking (INFN, FZK, CERN et al) - created the HEPspec06 benchmark which is now used by CERN/WLCG and HEP-wide for specifying acquisitions, completely replacing SPECint2000.



Working groups - 4

- Storage systems – driven by CASPUR (Andrei Maslennikov) to perform scientific testing (e.g. at FZK, CASPUR and CERN) of different solutions and publicise the results. Ongoing.
- Virtualisation – ongoing, led by Tony Cass: Enable virtual machine images created at one site to be used at other HEPiX (and WLGc) sites. Some early successes reported at Spring and more at this meeting.
- IPv6 – just starting, led by Dave Kelsey: gap analysis of the challenges; create a HEP testbed; share experiences; survey tools; survey security issues; investigate a timetable for migration from IPv4



IPv6
ACT NOW

Successes

Some are pure HEPiX success stories, in others we helped and/or provided a vehicle (publicity, support, a focal point)

- HEPiX scripts
- AFS
- HEP Benchmarks
- Storage tests
- The HEP-wide spread of Scientific Linux
- Independence of the group and its flexibility in adapting to changing technology

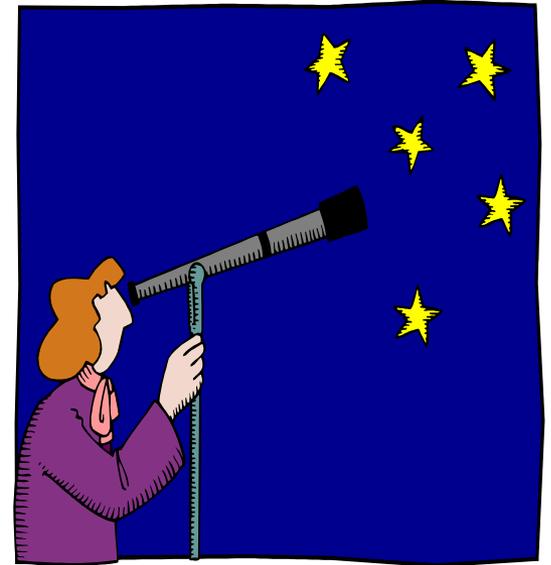
Are we value for money

- What is the value of knowledge gained, shared experience, a new idea?
- Cost of a CERN staff in Vancouver is airfare + expenses + week absent from work
- Yet there are 10 CERN staff here (plus me!), total cost ~30K CHF plus 10 FTE weeks “lost”.
- So the CERN answer is clear. Also IN2P3, RAL and DESY.
- European sites think “yes”, we are value for money.
- But FNAL often has only 2 attendees, BNL 1, SLAC 1, JLab 1. Why?
 - Competition of events
 - Funding
 - Relevance
 - Lack of managerial support

Attendance / Support

- Large European sites always present in force.
- Smaller European sites less present, at least this side of the Atlantic, much more so in Europe.
- US and Canadian sites have more competition for conference attendance, does that explain lower attendance by them, even at North American meetings?
- Do we blame US budget problems?
- Board size initially 4 people (3 labs); rose to 6 sites by 1997. Now >30 from ~25 sites). I take this as a positive sign of interest by the sites.

- Can we / should we try to increase US/Canada participation? How?
- Should we invite attendees from other sciences, for example invite again the Genome sites?
- Encourage more participation from Australasian sites (good turnout this time)





HEPIX