

# HEPIX Backup Survey

David Asbury CERN/IT/FIO

HEPIX, Rome, 6 April 2006




# Survey by email

- ◆ Harry Renshall sent questionnaire to HEPIX sites
- ◆ 15 questions
- ◆ *10.76 responses*
- ◆ **Chocolate** bar for each site which responded
- ◆ **Collect from Andrei tomorrow!**



# Who responded?

- ◆ CERN, Geneva, Switzerland
- ◆ DESY-H, Hamburg, Germany
- ◆ DESY-Z, Zeuthen, Germany
- ◆ FNAL, Fermilab, Chicago, USA
- ◆ FZK, Forschungszentrum, Karlsruhe, Germany
- ◆ INFN, (several sites), Italy
- ◆ IN2P3, Lyon, France
- ◆ JLAB, Jefferson Nat. Accelerator Lab, VA, USA
- ◆ RAL-HPC, Rutherford Appleton Lab, UK
- ◆ TRIUMF, Vancouver, Canada




# Which backup Product?

- ◆ 9 Home made (INFN, RAL-HPC)
- ◆ 7 TSM (CERN, DESY-H, FZK, IN2P3, 3xINFN)
- ◆ 3 Legato (DESY-Z, 2xINFN)
- ◆ 2 Time Navigator (2xINFN)
- ◆ 3 HP Data Protector (3xINFN)
- ◆ 1 TiBS (FNAL)
- ◆ 1 AMANDA (TRIUMF)
- ◆ 1 Oracle Reliaty (JLAB)




# Types of Data handled

- ◆ Home directories – all
- ◆ File systems - all
- ◆ Databases – 5 (CERN, DESY-H, IN2P3, JLAB)
- ◆ Mail – 2 (1 does the spam too!)
- ◆ Experiment data – 1 (INFN)



# Incremental Frequency

- ◆ Most do **daily** incremental backups
- ◆ Volumes of data ranges...
  - ◆ ~30GB/d: DESY-Z
  - ◆ 200GB/d-600GB/d: FNAL, IN2P3, JLAB, RAL-HPC, TRIUMF
  - ◆ > 1TB/d: CERN, FZK



# Full backup frequency


- ◆ Wide range...
- ◆ Weekly, each 10 days, monthly, 6-monthly
- ◆ Ranges ...
  - ◆ TRIUMF: 2TB/10days
  - ◆ JLAB: 4TB/week
  - ◆ CERN, RAL-HPC: 6TB/weekend



# Restore period (backup)


- ◆ **Varies with type of data and site**
- ◆ **RAL-HPC – one month**
- ◆ **Often 1 year for normal files**
- ◆ **FZK & JLAB keep backup files “forever”**
- ◆ **DESY-H keeps DBs for a week**
- ◆ **2 copies of DBs kept in CERN**





# How Much Restored?


- ◆ **Not in questionnaire!**
- ◆ **CERN:**
  - ◆ **~100 files (~ 1GB) per day**
  - ◆ **Big restores (> 250,000 files, 1TB) every month**
  - ◆ **Much more starting in 2005**
  - ◆ **Increasing**



# Total Volume stored

## ◆ Wide range according to site

- ◆ **< 100TB:** TRIUMF, DESY-H, DESY-Z, IN2P3, RAL-HPC
- ◆ **100TB-200TB:** FNAL, JLAB
- ◆ **> 400TB:** CERN, FZK



# Archived total

## ◆ According to site:

- ◆ **RAL-HPC:** none
- ◆ **< 2TB:** DESY-Z, IN2P3
- ◆ **2TB – 5TB:** CERN, DESY-H, TRIUMF
- ◆ **280TB:** FZK




# Restore period (archive)

- ◆ Several sites keep files forever
- ◆ DESY-H: for 3y (user), or 12y (expt)
- ◆ IN2P3: 25y




# Clients by OS

- ◆ Most have PCs with Linux or Windows
- ◆ Solaris: almost all sites
- ◆ AIX: FZK, INFN, IN2P3
- ◆ HPUX, OSF: CERN



# Backup Servers by OS

- ◆ All TSM sites use AIX
  - ◆ TSM on Linux coming in CERN
- ◆ Solaris used for other products - Legato
- ◆ Linux used for AMANDA (TRIUMF)
- ◆ RAL-HPC uses their Atlas Datastore



# **Tape Drives – wide variety**


- ◆ **LTO (FZK, DESY-Z, JLAB, INFN)**
- ◆ **SDLT (TRIUMF)**
- ◆ **IBM 3590 & 3592 in IBM 3594 robot (FZK, INFN)**
- ◆ **STK 9940, 9940B in Powderhorns (DESY-H, CERN, IN2P3, RAL-HPC)**
- ◆ **SAIT in T950 robot (FNAL)**
  
- ◆ **Total Cartridges 1000 (IN2P3) – 5600 (DESY-H)**



# Why This backup system?

- ◆ **TSM** <http://www-306.ibm.com/software/tivoli/products/storage-mgr>
  - ◆ **Functionality, wide range of clients, good support**
- ◆ **AMANDA** <http://www.amanda.org>
  - ◆ **works and is cheap**
- ◆ **Legato** <http://www.legato.com>
  - ◆ **Good offer for our Solaris requirements when started**
- ◆ **TiBS** <http://www.teradactyl.com>
  - ◆ **Scalable, easily partitioned, met requirements**
- ◆ **Oracle Reliaty**
  - ◆ [http://www.oracle.com/technology/deploy/availability/pdf/1168\\_Bednar\\_PPT](http://www.oracle.com/technology/deploy/availability/pdf/1168_Bednar_PPT)
  - ◆ **Best fit to requirements in 2001**





# **FTEs to run it**

- ◆ **Wide range according to site**
- ◆ **RAL-HPC uses shared site-wide store**
- ◆ **< 1 FTE: DESY-Z, JLAB, INFN, TRIUMF**
- ◆ **> 1 FTE: CERN, DESY-H, FZK, FNAL**



# Problems and Plans

- ◆ No obvious trend
- ◆ Backup is not very exciting
- ◆ Plan more Linux servers at CERN
- ◆ RAL-HPC data rate too slow into store



# Conclusions

- ◆ **Several large sites use TSM on AIX**
- ◆ **Legato also used on Solaris**
- ◆ **Smaller sites have more freedom of choice**