Autumn HEPiX-HEPNT 2003

Report from

HEPiX Fall '03
&
LISA '03

Wolfgang Friebel
Stephan Wiesand

DESY Zeuthen,
Nov. 18, 2003

LISA '03

OCTOBER 26-31, SAN DIEGO, CA

BY & FOR SYSADMINS!

THE 17TH LARGE INSTALLATION SYSTEMS ADMINISTRATION CONFERENCE

LISA

- Large Installation Systems Administration conference
- organized by Usenix every late autumn someplace in US
- Format:
  - 3 days of tutorials
  - 3 days of parallel sessions
    - and BOFs in the evenings
- DESY participants (~ 1500 total)
  - W. Friebel, P. v.d. Reest, Stephan Wiesand
  - have printed proceedings & CD with most tutorial materials
- online proceedings available to Usenix members

HEPiX

- politically correct name: HEPiX/HEPNT
  - meeting held twice a year, in spring and fall
  - on either side of atlantic ocean every other time
- Fall '03 held at TRIUMF in Vancouver
  - DESY participants (73 total)
    - R. Baltrusch, P. v. d. Reest
    - W. Friebel, H. Schwendicke, Stephan Wiesand
  - **http://www.triumf.ca/hepix2003/**
    - all presentations (powerpoint or staroffice or PDF format)
    - audio/video capture of almost all sessions (actually very usable)
    - good summary (23 pages of text in PDF format)

HEPiX Fall '03: Format

- 3 days of site reports and general talks (many very good)
- 1 day dedicated to presentations on security (ditto)
- ½ day of parallel sessions
  - security round table
  - windows round table
    - about security, first part joint security session
  - mass storage forum (not covered in this talk, P.v.d.R. only)
- first HEPiX with commercial vendor demos
  - demos and talks by 2 vendors of advanced, global file systems
- invited talks by Red Hat & Microsoft

Talk Format

- Part I
  - by S.W.
    - Windows input from H. Schwendicke, R. Baltrusch
  - ~ 30 minutes
  - mostly along HEPiX lines
  - focus:
    - selected site report topics
    - security
    - linux distribution discussion
  - additional input from LISA

- Part II
  - by W.F.
  - ~ 15 minutes
  - focus:
    - spam fighting
    - monitoring
    - other omissions by S.W.

Topics from HEPiX Site Reports

- Operating Systems
  - Linux, Windows, Solaris/SPARC everywhere
    - some HP-UX, AIX, IRIX left, typically being phased out
    - little MacOS (X) support, typically not on agenda
    - Windows rules the desktop domain
    - Linux rules the compute server domain
    - Linux is conquering the "real services" domain at many sites
      - AFS, NFS, Oracle, TSM, ...
      - mail, DHCP, Web, DNS, ...
  - all sites concerned about linux distributions
    - some expressed interest in Solaris/x86
      - SUN was marketing it very actively at LISA

Site Report Topics: Hardware

- Complaints about P4/Xeon
  - performance/GHz much worse than PIII
    - HyperThreading helps, but issues with linux scheduler, and CPU accounting / job scheduling complicated
  - power consumption
    - "Westgrid" at UBC (1008 dual Xeon 3GHz) can not run all blades (IBM bladecenters) in a crate until power supplies replaced
- positive reports about AMD Opteron performance
  - being considered for most farm purchases next year
- one site reports SCSI-attached IDE-RAID was a desaster
- CERN seems last site settling for "white boxes"

Site Report Topics: Windows

- all sites have or are deploying AD domains
  - 2000, 2003, XP
  - NT/9x still exist at some sites
- most sites have deployed or are evaluating at least one of
    - MS SMS
      - systems management server
    - MS SUS
      - software update service
  - necessity for efficient patch deployment
  - typically, only for new domains
    - NT/9x often managed manually only

Site report Topics continued

- Windows Terminal Services
  - either already deployed
    - sites report use increasing
    - often citrix
  - or being evaluated (most other sites)
    - typically RDP
- SLAC project on AD/Heimdal password synchronization
  - working with MS on tools to allow this smoothly
  - interest expressed by DESY Windows group
- Kerberos 5 is present or most likely future at all (?) sites
  - desire for single sign on expressed by some

Security

- most major labs had a high ranking security officer present
- security officers at all sites had an "interesting" year
    - Windows worms & viruses
        - Slammer, Sobig, Lovsan, Welchi,...
        - temporarily caused up to 30% packet loss on internet
        - effectively shut down some labs (and enterprises)
        - infected systems within minutes
            - during (re-)installation
            - before systems could be patched when turned on
        - CERN hit by virus before antivirus signature available
            - exploits IE weakness, installs spam relay on random high port
            - lab faced threat of being brought to court due to nature of spam

Security continued

- Linux ptrace vulnerability
  - trivially exploited from cracked user accounts
    - success rate almost 100%, exploit widely available
- frightening root kits, like SuckIt
  - very good at concealing itself, very hard to detect
  - installs backdoor defeating all firewalling
    - listens on ALL ports for backdoor trigger packets
    - then initiates TCP connection from infected host
- users running
  - P2P filesharing software
  - IRC (and being caught by bots)
  - vulnerable sshd or httpd or... (on high ports)

Security: Common Problems

- common agreement today these are the worst problems:
  - systems not properly (professionally) managed
    - each of these measures alone almost eliminates attack potential:
      - applying patches timely
      - running antivirus software with daily updated signatures
      - running a personal firewall at least buys time
    - how could so many systems be compromised this year ?
      - fix for many attacks available weeks / months / years before !
  - firewall penetration
    - notebooks, VPN, dialup (home systems)
    - unauthorized, vulnerable services / applications
  - users downloading malware, opening unknown attachments, ...
  - notebooks that can only be updated inside their home network
    - one week can be too long these days

Security: Common Measures

- most sites now apply these or are planning to do so:
  - all devices attached to network must be registered
  - and responsible has to agree (in writing) to rules, like
    - system must be configured securely
    - patches must be applied timely, system rebooted if necessary
    - system must be running update antivirus and firewall
    - system must not be running unauthorized services
  - users of centrally managed systems must agree to rules, like
    - no P2P software or other unauthorized services / applications
  - VPN/dialup users must agree to rules, like
    - no additional software, no usage by the kids, ...

Security Measures: Exceptions

- exceptions from rules generally granted if necessary
  - if work cannot be done without violating them
- most sites require a written statement
  - why there's a need for it
  - what technical measures prevent security breaches
    - "how will you prevent unauthorized file access through your P2P filesharing application ?"
  - signed by user and responsible
- sites report almost all requests are withdrawn after pointing out this requirement

Security Measures: Scans

- major sites run scans of their network
  - detect vulnerable systems, unauthorized services
  - detect compromised systems (backdoors, ...)
  - full scans regularly
    - typically take O(1 month) to complete
  - individual scans immediately when new devices attached
  - problems:
    - scan may disrupt operation of some devices (DAQ equipment...)
      - -> first detect OS, then apply specific scan
    - feasible to quarantine new systems until scanned ?
  - vulnerable/compromised systems disabled on network level

Security Round Table Results

- HEPiX labs will agree on common set of minimal rules for systems to be attached to their networks
  - systems carried by guests from other HEP labs are expected to comply with these
- incidents and attacks should be communicated to the (closed) security mailing list
- a new security discussion list for HEPiX was created
  - not public, but open to anyone from any HEP lab
    - subscription must be approved by list owners (hosted at fermilab)
    - new members expected to introduce themselves
      - or may be removed from list

Security: Summary

- today's threats are serious
  - no major damage yet, but only matter of time
- "patch early, patch often !"
  - any system, centrally managed or not
  - including network gear, farms, desktops, notebooks, ...
  - this is a significant deviation from
    - "choose patch time wisely for optimal availibility"
    - "it's ok to patch servers only"
    - "locally only exploitable bugs aren't worth patching"
- firewalls can help, but are not a sufficient solution
  - limit exceptions as much as possible

The Linux Discussion: Background

- almost all HEP sites run some vanilla Red Hat Linux
  - many also already run a few Red Hat Enterprise Servers
    - typically for Oracle
    - significant cost per server and year
- some (DESY, GSI) run SuSE and/or debian
  - few SuSE/debian hosts at few other sites
- Red Hat early this year shortened distribution life times
  - to 12 months
- later this year they discontinued their vanilla distribution
  - superseded by Fedora, life time 6-9 months

Linux Discussion: Background

- distribution end of life:

RedHat 7.x	12/03	
RedHat 8.0	12/03	
RedHat 9	04/04	
Fedora Core 1	07/04	(at best, and limited)
SuSE 8.2	04/05	
SuSE 9.0	10/05	
debian woody	12/04 + ?	(12 months after undefined date)

- SuSE/Red Hat Enterprise distributions live 5 years
  - = unlimited in practice

HEPiX Linux Discussion

- most labs now have to find a new workhorse distro soon
  - CERN & probably others will support 7.3 until 12/04
  - but need several months for certification of new OS
- most labs have contacted distributors about volume licensing
  - we talked to SuSE and RedHat, all others to Red Hat only
  - all got similar offers around XXX $/year/node
  - no lab could negotiate acceptable conditions so far
- => try common HEP effort
  - Red Hat invited to HEPiX
  - session on this topic (w/o RedHat presence, w/o recording)

Red Hat at HEPiX

- Red Hat sent Don Langley
  - sales manager for california
    - including SLAC
- held a plain marketing talk for Red Hat Enterprise Linux 3
  - session not recorded
  - pdf on the web
  - no additional information
- refused to discuss HEP volume licensing
  - just stated they're "interested in creating a win-win situation"

Summary of Discussion Session

- most sites really want to use Red Hat Enterprise Linux
  - debian/SuSE/others not considered seriously
- but not with their default support model
  - HEP sites most of all want the patches
    - not per incident remedial servcies
    - after inserting an own kernel module, these are void anyway
    - on LISA, heard complaints about service from people having it
  - some sites interested in RHN satellites (->delegation)
- HEPiX believes Red Hat have not yet made up their mind
  - give them more time (how much ?)
- try negotiating on higher level

Other Linux Options discussed

- some consider rebuilding a RHEL from sanitized source
  - after all, it's GPL
  - probably legal if all trademarks and files with other licenses are removed, and the name is changed
    - situation is not really understood by anyone
    - CERN would require written permission before redistribution
- some consider using Fedora
  - and hoping for Fedora Legacy to work
    - volunteer project hosted by Red Hat to provide patches for old fedora
- hardware vendors may offer reasonable RH WS licenses
  - but what to do with existing hardware?

Linux in HEP: Next Steps

- CERN, SLAC, Fermilab will try to negotiate with Red Hat
  - objective: acceptable conditions for using RHEL
    - in all HEP (LCG?) labs, and collaborating institutes
  - no deadline set
- U.S. department of energy is negotiating for all their labs
  - what if they succeed, and HEP doesn't ?
- DESY will watch from the side line
  - we're about to roll out DL5 based on SuSE 8.2
    - buys us a year, no immediate pressure
  - but we expressed interest to buy into a reasonable solution

Email (HEPiX)

- at HEPiX two reports on Spam fighting (GSI and CERN)
- GSI:
  - did setup a new mail infrastructure based on postfix
  - input and output filters for mail with amavisd-new
    - SPAM tagging with spamassassin (2.55)
    - Virus filtering with clamav and sophie
- CERN:
  - converted the central mail servers to Exchange
    - was previously sendmail + UW-IMAP
    - spam fighting with homegrown script (.net framework based on SA)
    - proposal to use feedback mechanism for new mail senders

SPAM fighting at CERN

- Proposal to approve mails for new sender addresses
  - user receives mail from a new address
  - automatic response generated to prove identity of sender
  - only if sender replies, the sender gets whitelisted
- Much critics at HEPiX
  - similar amount of work to be done as for unfiltered mail
  - impractical for e.g. mailing lists
  - easy to forge by hackers
- Even more critics for similar concepts at LISA
  - 2-3 in favor, approx 500 against it.

SPAM mini symposium at LISA

- Very broad attendance, general trends were visible
  - most of the sites use or plan to use spamassassin
  - some other proposed methods very unpopular (see prev. slide)
- legal issues discussed
  - fairly easy to track spammers
  - spammers usually engaged by others to do the dirty work
  - would need to punish the profit making site
  - could be abused by competitors to spam in their name

SPAM and Viruses

- Active State and SOPHOS well known in this marked
  - Active State acquired by SOPHOS recently
  - come now with Spam + Virus handling
  - additionally management interface for policies etc.
- SOPHOS talked about new ideas in SPAM fighting
  - observe new tricks of spammers and have countermeasures
  - e.g white ink (print e.g. blue on blue) became almost white ink (print blue on slightly different color of blue)
  - now testing for difference in color space

HEPiX Login scripts

- Reworked by CERN
  - used also at DESY (with mods)
  - maintained compatibility with original concept
  - no longer dependency on external software
  - remarkable speedup achieved
- Reintegration at DESY?

Monitoring

- talk at HEPiX in the context of fabric management
  - work based on software written for Grid work package 4
  - also covered configuration management
  - in use at CERN already, not ready for outside labs yet
- Network telescope (invited talk at LISA)
  - great idea to observe network attacks
  - http://www.caida.org/analysis/security/telescope/

Monitoring

- Many talks and tutorials at LISA
  - alarming tools (e.g. nagios, scout (DESY))
  - intrusion detection tools (e.g. snort)
  - monitoring (system and network) (e.g. MRTG)
- Many tools for monitoring based on RRDtool
  - most major sites do have monitoring/alarming/IDS in place
  - at DESY (Zeuthen) alarming well covered, monitoring at the network level only, IDS not yet
- Work underway to do more monitoring

Famous last Words

- HEPiX/HEPNT and LISA are quite different
- both are very relevant to DESY computing
  - even if focus of this presentation was on HEPiX
- DESY staff should attend both regularly
  - next LISA: Nov. 14-19, 2004 in Atlanta
  - next HEPiX/HEPNT: May 2004 in Edinburgh