



Virtualized Tier3 project

Sergey Panitkin

BNL

ATLAS



BROOKHAVEN
NATIONAL LABORATORY



Outline

- Introduction
- Current status
 - People involved
 - Test farm at BNL
 - Progress so far
- Summary



Virtualization and Atlas Tier 3 centers

- The goal of the current project is to try to understand modern virtualization technology and utilized it in the US Atlas Tier 3 context.
- We expect that most of the T3 will be local, relatively small clusters, with some modest level of local technical support.
- US Atlas T3 task force identified several possible T3 models
- Minimization of a local group effort that is needed for hardware and software maintenance will be an important, if not a decisive, factor for a successful deployment of the Tier 3 centers.
- The main area of concern right now is installation, support and maintenance of various software components needed for productive T3 operations. Think of DDM related software (both Grid related and Atlas specific), Atlas software proper, batch software, PROOF, Xrootd or some other distributed storage solution, monitoring tools, administrative tools, etc.



Virtualization and Atlas Tier 3 centers

- Tier 3 machines will have multi-core CPUs
- Virtualization may help to utilize multi-core hardware.
 - In many cases it makes no sense to have a separate machine allocated for a given service
 - Several services may be co-located as independent VMs on the same physical machine
- One can hope to increase robustness of the services by isolating them in a VM.
 - Failed service will not bring down physical machine and will not affect other services
- Virtualization may simplify software configuration and packaging
 - Sites need to install a hypervisor – software is distributed as ready to run VMs With (hopefully simple) customization instructions or perhaps auto configuration Mechanism
- It may improve security. Central VM repository maintains up to date, patched VM OS.
- Central VM repository does application software validation, etc
- Several groups in Atlas were already studying and using virtualization technology



People involved so far

People and areas of interest/expertise

- Doug Benjamin (Duke) – Condor, Grid services, documentation, coordination
- Sergey Panitkin (BNL) – sysadmin, infrastructure, XRootd/PROOF, clouds
- Alden Stradling (UTA) – CernVM, Atlas software
- Yushu Yao (LBNL) – CernVM, Atlas software, Condor

• Many thanks to ACF@ BNL people for their support of the project and for hardware allocation

• Note that Duke, LBNL and UTA have their own virtualization testbeds



Test farm at BNL

- Currently the farm has 8 machines allocated to us by RACF.
- They are outside of RACF firewall
 - Many thanks to Michel Ernst and Chris Hollowell for making it happen!
- Each machine has:
 - 2 dual core AMD Opteron 265 CPU's running at 1.8 Ghz
 - 16 GB of RAM
 - 2TB (4x500GB) HDD
 - 1GB network interface
 - They run SL53 x86_64
- Right now 5 machines run Xen hypervisor
- 3 machines are left “un-hypervised” for tests of:
 - Other hypervisors like KVM or VMware
 - Performance comparisons of virtualized solutions to “bare” ones, etc.
 - Test of the “cloud” tools and/or management tools for virtualized clusters.
- Farm partitioning will be changed when necessary



Progress so far

- Kick off meeting at BNL - June 17, 2009
- Dedicated hardware was allocated at BNL - early July 2009
- Hypervisors installed and first VM's built at BNL - mid July 2009
- User accounts set up at BNL - early August 2009
- Project Wiki page set up - early August 2009
<https://www.usatlas.bnl.gov/twiki/bin/view/AtlasSoftware/VirtualTier3>

- Several SL based virtual Atlas worker nodes at UTA - mid August 2009
- First tests of virtualized Condor batch queue at LBNL - mid August 2009
 - Port to BNL test bed soon
- First presentation to a wider Atlas T3 community - Today



Summary

- We recently started a project aimed at application of virtualization technologies to Atlas Tier 3 model
- Four people are involved so far. More developers are welcomed!
- Right now it's more R & D project rather than production type of activity
- We need to gain experience and expertise with virtualization, learn about what works and what doesn't.
- Common development farm is available now at BNL .
- First VM prototypes exist – Condor queue, Atlas worker nodes
- The project (hopefully) brings together existing VM related activities in the US Atlas organization.
- We hope to leverage expertise in VM existing in Atlas in order to make rapid progress and provide useful VM based solutions for the Atlas Tier 3 community
- Stay tuned....