

# Xrootd and clouds


Doug Benjamin  
Duke University



# Introduction



- Cloud computing is here to stay – likely more than just Hype (Gartner Research Hype Cycle 2011)





# Storage and clouds



- What is our Cloud compute/storage model?
  - Mostly computational (MC generation/simulation)
    - Xrootd does not play much of a role here
  - User analysis?
    - Storage is clearly an issue
      - Does one use the Storage in the cloud (EC3)?
      - Does one use the Storage provided by each VM
    - How do we stitch it all together?
      - Should federate the storage within a cloud cluster
      - Xrootd could be a solution here (multi CPU VM's for data serving and computation)
  - Initial work focusing on federating the worker node storage



# Current status



- Using Fermi-Cloud for initial testing
  - Open Nebulus configuration
  - “static” VM’s
    - Instantiated by hand based on configuration files
    - OS is an approved image (1.5 GB SL 5 based)
      - Preconfigured and “static”
    - Static / dynamically assigned IP
    - External address (at least tested within FNAL campus)
  - Given 6 VM’s
    - Redirector
    - 3 data servers
    - Xrootd proxy machine
    - Interactive machine
- Have begun to reconfigure VM (add more disk space) and reinstantiate



# Lessons Learned so far



- Need to have contextualization/configuration mechanism in place
  - Several options:
    - CernVM Copilot approach
      - Jabber/XMPP messaging approach (used by Google Chat etc)
    - Puppet to configure the machines on the “fly”
      - Need to setup the puppet master first (likely on Xrootd redirector for simplicity)
        - All nodes need to know the redirector IP and IP addresses are dynamic so need to discover this right away and place information into configuration files
      - Need to fetch in small script to run setup puppet clients before puppet
      - Install Xrootd on the fly through puppet
      - Teach Puppet to setup Xrootd
        - Need to understand the security and account issues
    - Some combination of Both



# Next steps



- Collaborate with CernVM team on Lxcloud to study the contextualization issues
  - Contextualize a VM creation
  - Contextualize with some information at start and connect to “central” server (puppet master)
- Prepare for the dynamic nature of the cloud
  - Could/will have more than 64 data servers in a cloud cluster
  - Need to dynamically setup Xrootd supervisors as the data server count grows
- Test effect of Xrootd redirector/Xrootd supervisors on VM's vs all on one physical machine
- Understand virtualization I/O penalty better
  - Use the suite of typical user analysis jobs being collected within US ATLAS analysis support team



# Longer term goals



- See if can use Xrootd to cluster ephemeral cloud worker storage
  - Automatic discovery and configuration the key here.
- Can we then federate this storage into the Xrootd federation?
- What are the cost model for this useage
  - Network/Storage costs
- Can we launch a cloud cluster/ load data / grow the cluster as needed, extract the precious data and shutdown the cluster?
- See if this all can be part of a Tier 3 cluster in the cloud (local analysis site, with interactive access)



# Conclusions



- Cloud computing is here to stay
- Just starting the process of evaluating federating storage on cloud services –
  - Are we at the “peak of inflated expectations”?
  - “Through of disillusionment”?
- Time is short – 2014 is closer than we think
- Is this something that there is multi experiment appeal?
- How should we proceed?