BABAR DATA PRESERVATION STATUS UPDATE

Concetta Cartaro
SLAC

DPHEP @ CHEP2013
Amsterdam, October 16th, 2013

Outline

- LTDA cluster status
- New developments of *BABAR* LTDA
  - *BABAR*-To-Go
- Discussing the DPHEP collaboration agreement

LTDA STATUS

- The LTDA is very stable
  - Simulation and physics streams production are running smoothly on LTDA.
  - We are preparing to move CVS repository to LTDA.
- On the long term all major services will be made provided through the LTDA cluster
  - It is under constant usage by analysts.
  - The importance of the LTDA is increasing as *BABAR*'s share of other computing resources at SLAC and remote sites is decreasing.
  - Now the LTDA is the only system which data reconstruction can be done on.

DATA PRESERVATION

- Almost all Legacy Data now backed up at CC-IN2P3.
  - 1.8PB transferred, ~0.5PB still to be copied
    - 2-4TB/day

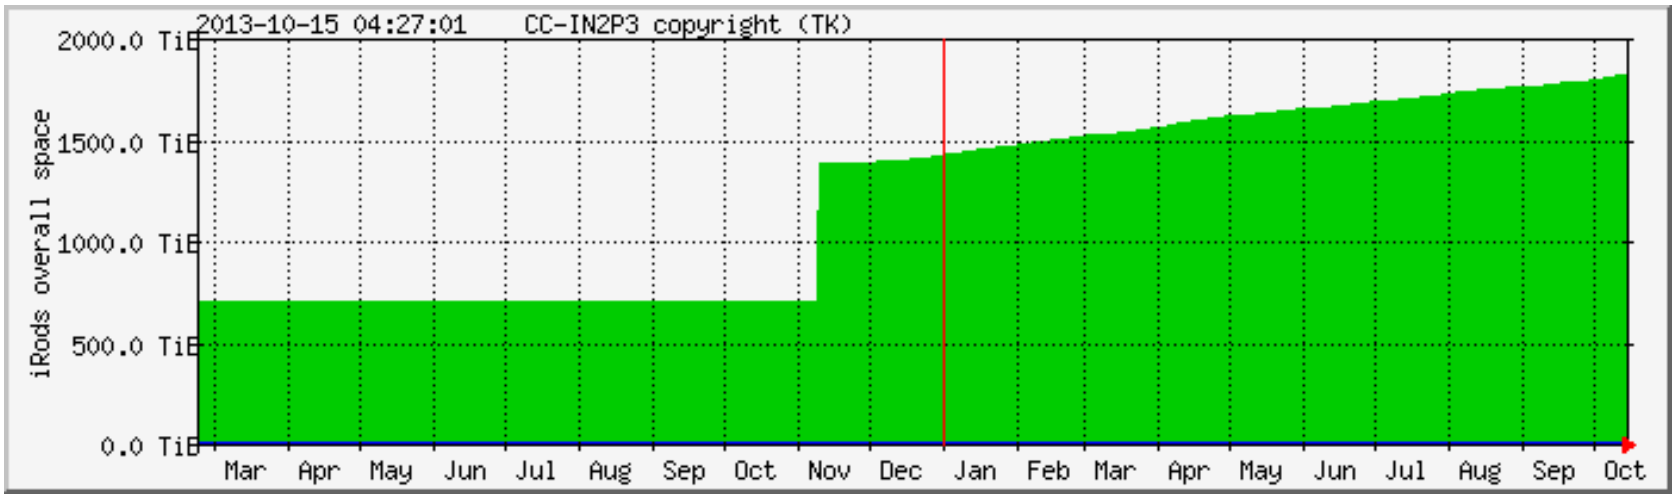

BABAR-TO-GO

- Request made by a collaborator Meeting in October 2010.
  - Never forgotten, but many things were ahead of us before we could think about it.
- Basic idea:
  - A virtual machine, in raw format, that can run with multiple software applications on multiple platforms (KVM, VMware, VirtualBox, …) , with one or more BaBar analysis releases fully installed and ready to use.
  - Take it home with you and run it on your laptop.
  - Then you only need a (large) box with the data.
- It seems easy enough but need to define the requirements.

BABAR-TO-GO BASICS

- How should it work?
  - Many possibilities and combinations of features.
  - Single interactive machine with user login.
    - Useful for development work and to run interactive jobs.
  - Run batch jobs.
    - If you need to run batch jobs we will have to provide an easy way to configure it.
    - The home institution may have a local batch service to use directly.
- Even if the VM is based on SL5 or SL6 which will receive support until 2017 or 2020, without updates there is a security risk associated to the VM if it can access the network.

How to Access the Code

- The portable VM will have one or more analysis releases installed and there are several possibilities.
  - Local BFROOT/ (*) or
  - External BFROOT/
    - Local disk, NFS mount point, remote access (afs, cvmfs?).
- Tools.
  - Example: BbkDatasetTcl needs the bookkeeping db to work.
    - Need to setup a site connection parameter as you do from remote sites.
  - Local copy of the bookkeeping db.
  - Prepare parametric TCL files for the most common datasets.
- Remote access: security concern if the VM is allowed to connect to the network.

(*) BFROOT/ is the afs top level of BaBar filesystem. Underneath there are the code repository, the releases, the tools, and even the web pages.

How to Access Data and Conditions

- A natural option: XROOTD.
  - Read directly from SLAC:
    - Clients read over the WAN from Xrootd at SLAC (or CC-IN2P3).
      - CC-IN2P3 does not by rule allow incoming requests.
- Local XROOTD cluster.
  - XROOTD is setup at the local site and clients read from it.
  - Users can copy files to the local XROOTD cluster.
  - There are two option to handle files that are missing:
    - XROOTD stages a missing file from SLAC (xrdcp or bbcp)
    - A client gets redirected to SLAC to read a missing file.
- The import/export tools can be used to transfer files/collections/datasets from SLAC to a local site (applies to using XROOTD or local file system).
- Allowing read-only access to XROOTD from outside of SLAC would require some authentication (maybe not?).
  - XROOTD supports krb5 and gsi
- Other options: direct download/copy.

BABAR-TO-GO FIRST CUT

- Create a base image with the OS (SL5, SL6) and an auxiliary image containing only BFROOT (independent from OS) with a minimal structure
  - package/, repo/, dist/releases, dist/packages, … with only the bin directories of the supported OS's
- Conditions not included in the image
  - Cond24boot11 ROOT snapshot is 24GB in 220 files
- Basic approach for data and conditions
  - Fixed NFS mount point in the VM for conditions and data
  - The user provides a NFS disk with a specific path to the data
    - /store/…

MINIMAL DATASETS

- What is the minimal amount of disk needed for the data by a user who wants to use *BABAR*-To-Go ?
  - Need micro (reconstructed events info) or mini (detector level info) ?...

Components	R24g data + MC + skims	R24c data + MC + skims	R24d Y(2S,3S) + MC +Skims	R26 Gen MC (1237,1235,1005)
Micro (TB)	113	216	122	81
Micro+Mini (TB)	280	384	290	160

- Consider only one deepCopyMicro (or Mini) skim
  - Very light weight
  - Distribute copies of deepCopy skims and reach them through a federated XROOTD cluster…
    - Impractical on the longer term?
- Examples →

Skim	R24c skim size (TB)
BToDlnu	7.7
BSemiExcl	4.0

An Occasion for Common Project?

- Aleph (CHEP13 poster P3.74 (?)) has deployed a SL4 VM working in VirtualBox that is used by INFN BARI cloud service.
  Isn't this a perfect example for a common project where the experience of many can be really put into profit?
  - A common way of building VMs.
    - Raw format so that can work anywhere (KVM/QUEMU, VB, VMWare)?
  - A common way of accessing data (and/or code).
    - Federated datasets? Code repositories? CERN contextualization?
  - A common way for users to access services
    - User login? Collaboration login? GRID certificate?

DPHEP Collaboration Agreement

- *BABAR* management group is examining the document.
  - *BABAR*'s involvement considered at next Executive Board in 2 weeks.
- Some perplexities about the document and how to present it:
  - It appears too general and vague, making people miss the point.
  - It addresses the labs.
    - Do Collaborations fit in? People assume NO.
    - If Collaborations are candidates for the agreement then DPHEP management should contact the collaborations spokespersons and invite them to participate and sign.
  - DPHEP role and interaction with other projects (RDA, H2020, DASPOS, …) is not clear.
  - Data preservation and open access seem to go hand in hand. *BABAR* like other collaborations has not yet decided in favor of making data public.
    - With the added complication of providing support/code/tools for the data.

ACCESS POLICY

- You need to be a BaBarian to access *BABAR* data.
- Point is… anyone can be a BaBarian.
- If you have an idea to test, a theory to verify, or a new measurement that can be done with our data come and join the Collaboration.
  - An institutional PI or the management group will endorse your association.
  - You sign only your paper(s) together with the Collaboration as an Associate and then you may decide to become a Member (Collaboration approval and community service requested) and sign all the papers.
  - Very fruitful interaction with theorists.
    - Example: T-Violation measurement (see later) done by a theorist (now BaBar Associate) and a BaBar PhD student.
- Public access policy still not faced by the Collaboration.
  - Even if we made the data public without the Framework is quite useless.
    - Need manpower to support.
    - Small data samples used by affiliated groups for education (undergrad students).

More on BaBar and LTDA

- LTDA technical details, performance tests, and problems & solutions (DPHEP 6)
  - https://indico.cern.ch/getFile.py/access?contribId=30&sessionId=9&resId=0&materialId=slides&confId=209688
- *BaBar* "*ingest*" talks (DPHEP 7)
  - https://indico.cern.ch/getFile.py/access?contribId=28&sessionId=5&resId=1&materialId=slides&confId=233119
  - https://indico.cern.ch/getFile.py/access?contribId=2&sessionId=1&resId=0&materialId=slides&confId=233119

LTDA Developers

- Coordinator: Tina Cartaro
- BaBar software expert: Homer Neal
- Development, virtualization and LTDA system administration: Marcus Ebert
- Network design: Steffen Luitz
- System performance and CDB: Igor Gaponenko
- Databases, tools and production: Douglas Smith and Tim Adye
- Computing Division experts
  - System setup and adminstration: Booker Bense, Lance Nakata, Randall Radmer and all the Unix-Admin team
  - Xrootd expert: Wilko Kroeger
  - Network setup: Antonio Ceseracciu
  - BaBar-SLAC Computing Division liaison: Len Moss → Andrew May