Tier-3 Implementation Committee Summary and Open Questions

US ATLAS Software Planning Meeting
LBNL
August 21, 2014

Mark Neubauer

University of Illinois

The US ATLAS Tier-3s

- 90% of US universities consider their Tier-3 computing as critical to their group's productivity in ATLAS
- The current US Tier-3 computing resources were purchased using DOE ARRA, NSF MRI and university funds
  - Accompanied by substantial in-kind contributions from the institutions
- Tier-3s have been very successful in facilitating the analysis productivity by US physicists in ATLAS during Run 1

➔ **Now is the time to chart their future course for Run 2** (and beyond)

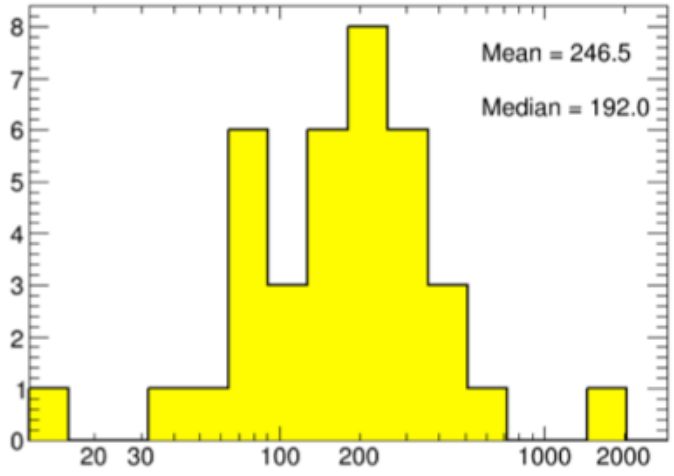

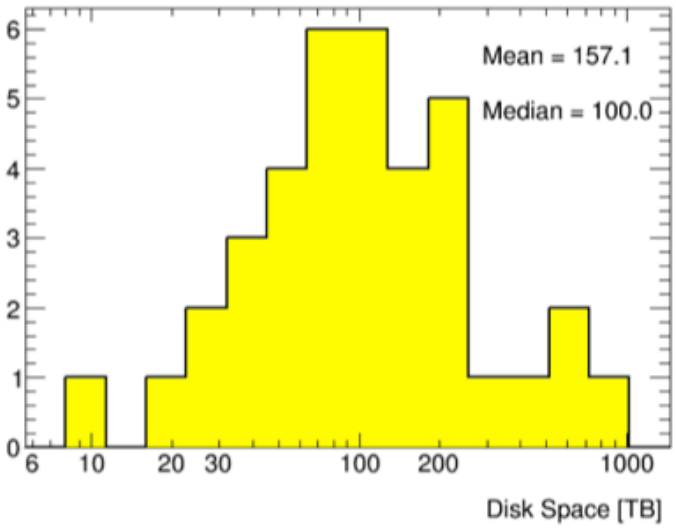

Prior Tier-3 Task Forces

- Two prior committees reviewed the T3s and produced a report
  - 2009 Tier-3 Task Force. Chair: Chip Brock
  - 2013 Tier-3 Task Force. Chairs: Jianming Qian and Gabriella Sciolla

Tier-3 Implementation Committee (T3IC)

- Formed in November 2013 by US ATLAS Operations Program
- Chairs: Mark Neubauer (Illinois) and Jason Nielsen (UCSC)
- Need for such a committee driven primarily by desires from
  - the funding agencies for a clear message from US ATLAS regarding its Tier-3 needs and how future Tier-3 funding should be provided (if at all)
  - the US ATLAS Operations Program to formulate a specific plan-of-action to implement the core T3TF report recommendations and estimates of any additional Operations Program-supported resources required
- Findings and Recommendations to JOG in late September

T3IC: Run 1 Analysis Findings

- T3IC examined three analysis workflows to determine the range of resources needed to analyze the 25 fb^{-1} Run 1 data set:
  - (i) Dilepton/Multilepton, (ii) lepton+jets, and (iii) jets + MET analyses
- Most analyzers process data in a two- or three-step process using a combination of remote and local computing resources
- We find that a typical **Run 1** analysis workflow requires *roughly*:
  - 5 TB of user-defined storage to store the reduced datasets from 1st step
  - 100 cores to achieve a Step 2 turnaround of < 2 hours (wall-clock time)

(It is interesting to compare these numbers with the 2009 U.S. ATLAS T3TF report, which predicted the need for a typical "T3g" system having 80 CPU cores and 20 TB of storage, which would be used to perform several analyses. This has turned out to be very close to the actual minimum scale of resources needed for ATLAS Run 1 physics analysis.

New Analysis Model for Run 2

- The new ATLAS analysis model in **Run 2** allows analyzers to produce results directly from a common analysis format (xAOD) that can be read outside of the Athena framework
- It also provides a centralized skimming production service to be used by physics groups. The skimmed/slimmed common analysis datasets can be further reduced in steps to a final ntuple.
  - These reductions correspond to the first two reduction steps previously mentioned

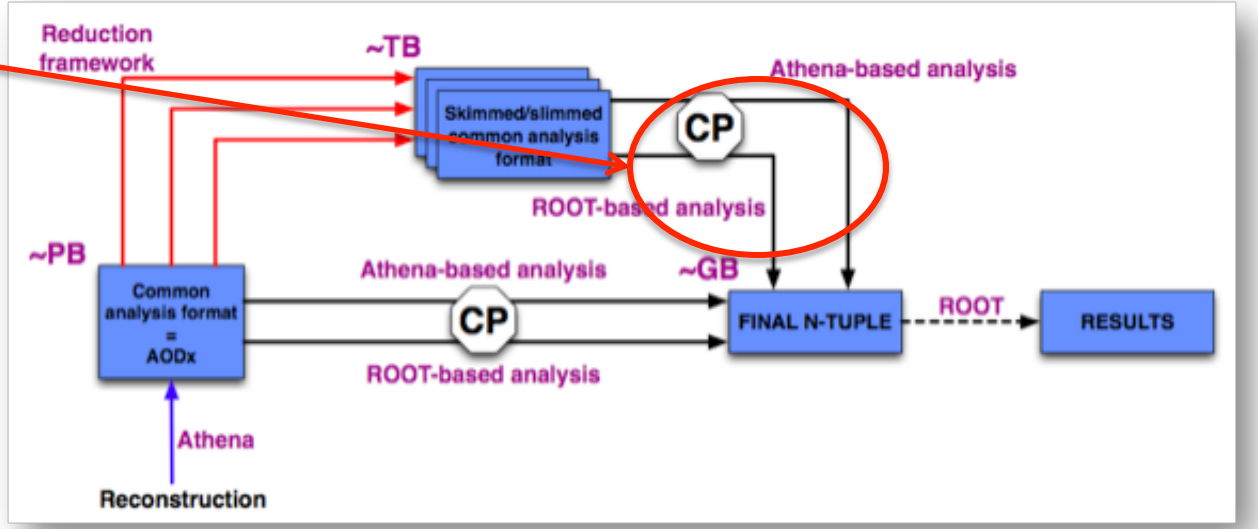

Resource Estimates for Run 2 and Plans

- The 2013 US ATLAS T3TF estimated that the requirements to analyze 100 fb^{-1} of Run 2 data will entail about 3-4 times the resources used to analyze the Run 1 data sample
  - This estimate is based on the expected increase in peak luminosity, the increased rate of trigger system output, and other such factors which are known to scale over Run 1
  - ***Our committee concurs with this estimate***
  - These estimates are consistent with results from a US ATLAS Analysis Support survey conducted in 2012. At that time, the median expectations were that Run 2 analysis would require 2-3 times the Run 1 resources
- In isolation, we expect a **Run 2** analysis to need 20 TB of user-defined storage and 400 CPU cores to satisfy these constraints
- We developed 3 candidate plans for Tier-3 evolution into Run 2
  - These plans were compared for cost-effectiveness → one plan emerged

Recommended Plan for Tier-3 Evolution

- Tier-3 sites repurpose their existing Tier-3 hardware to provide a more optimal mix of interactive and batch
- Each institution deploy a *modest* Tier-3 resource locally using the *latest hardware*, used for interactive and small-scale batch jobs characteristic of the last few steps of analysis workflows
  - We define this resource in terms of a Local Resource Unit (LRU)
    - LRU: minimum new Tier-3 resource at an institution w/ single active analyzer
    - The # of LRUs at an institution should be scaled by # of active analyzers
- Deploy new tools and technologies that transparently expand the Tier-3 resources into the beyond-pledge resources on the facilities when the workload is such that it would not complete in a reasonable amount of time on the local Tier-3 resource
  - Local resources are used first for workflows in support of analysis and are augmented (as needed) by resources available on the US facilities

University of Illinois

The Recommended Plan calls for new funds to deploy a modest, but modern, resource locally at each institution. We refer to the building block of this resource as a ***local resource unit*** **(LRU)**

- From the Run 2 estimates presented, one might naively think that the LRU is 4x the resources needs for an isolated analysis. This would neglect that:
  - people at an institution collaborate on a given analysis (share resources
  - not all Tier-3 computing is local in this plan
- We define the LRU as **10 TB of user-managed storage** and **100 CPU cores**. The LRU the minimum new Tier-3 resource for a single active analyzer at an institution
- We recommend that the number of LRUs at a given institution be scaled by the number of active analyzers at the institution
  - The exact prescription for this scaling being finalized by the T3IC

Possible LRU Configuration

- The LRU is the minimum new, local Tier-3 resource for a single active analyzer at an institution
  - **10 TB of user-managed storage** and **100 CPU cores**
  - In reality, no institution has a single activle analyzer
- Once one considers sensible, cost-effective architectures that implement the LRU, the actual resource will be different
- One possibility under consideration by CMS (aka “Frankiac”):
  - Colfax CX22850i-X5 2U Rackmount Server
  - Four compute nodes in one 2U chassis
    - Dual Xeon E5-2650V2 processors (8C/16T) / node
      - 16C/32T / node ➔ 48C/96T / chassis
    - 12 x 4 TB drives total gives total of 48 TB
    - Costs ~$20k total

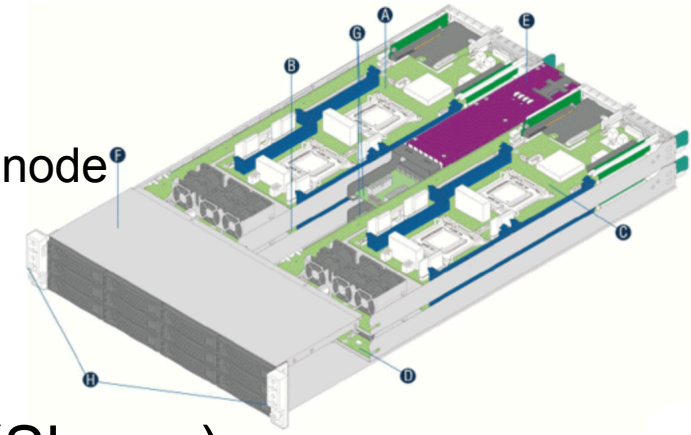

- Another possibility based on Dell R720 (Shawn)

T3 *Local* Resource Extension: ATLAS Connect

- Allows a client machine running Condor to be logically extended to other resource targets (Tier-2s, Cloud, Campus clusters, …) through HTCondor flocking mechanisms
- Numerous resource targets already enabled
- The system is being used by several groups in US ATLAS to get more cycles
  - Also, a CMS Connect now
- Planned role in US ATLAS Tier-3 evolution to extend the local Tier-3 resources into the facilities

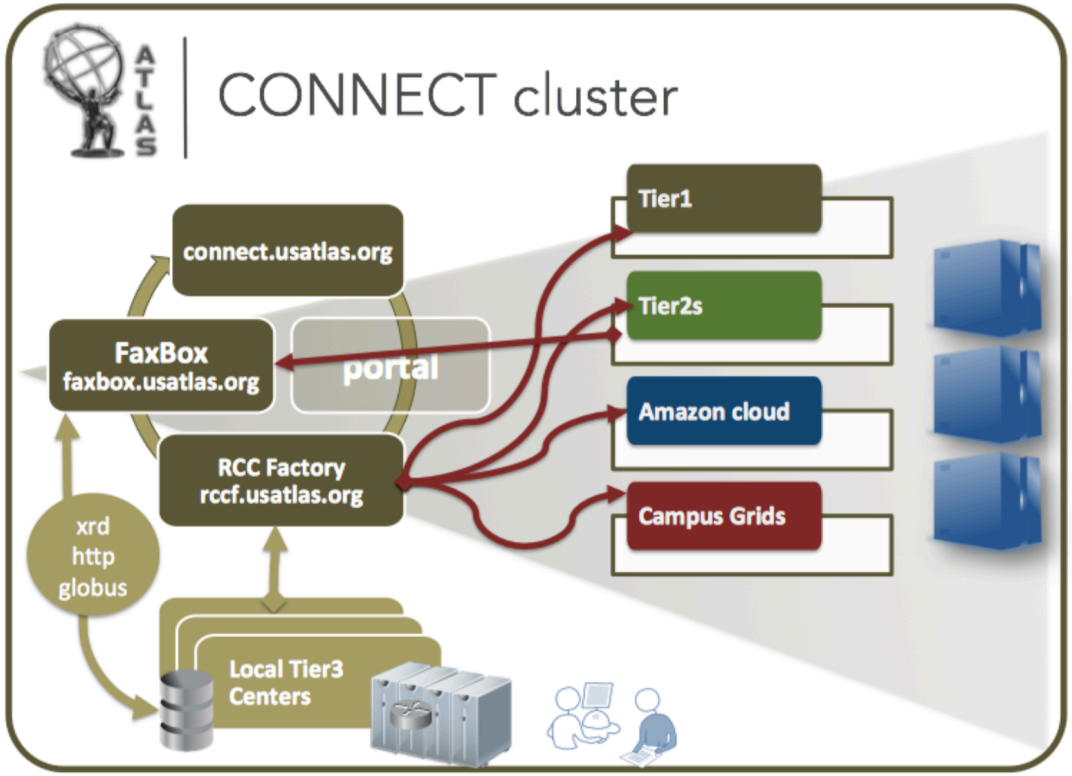

Impacts on US ATLAS Operations

- The Recommended Plan requires additional support from the US ATLAS Operation Program
- The tools, technologies and support are highly leveraged from existing software, facilities, and ongoing efforts within the US ATLAS Physics Support, Software, and Computing Program
- A *modest* amount of additional US ATLAS Operations Program funds will be needed to support US-wide analysis activities with these new elements. We have identified the areas where additional support is needed and are finalizing the estimates
- Training and documentation will be ***very important*** to support the Run 2 Tier-3 activities in the US since they will have a different feel in Run 2 given the new tools in the Recommended Plan and the new analysis model

Large Shared Tier-3s and Tier-3 Consortia

- *Large Shared Tier-3s*: Some institutions have large computing capabilities along with potential to bring in substantial new resources & expertise for Tier-3 computing at *no cost to ATLAS*
  - We recommend assessing the level interest among US ATLAS institutions and the funding agencies to incentivizing a small number of large Tier-3 that are shared among US ATLAS analyzers
- *Tier-3 Consortia*: Several groups have jointly provisioned large shared Tier-3 systems, pooling resources to locate equipment at a single host institution. These "Tier-3 consortia," which enjoy economies of scale and foster close collaboration on physics analysis, were highlighted in the 2013 T3FT report
  - While the formation and organization of the consortia themselves is outside the scope of our report, we endorse this approach and suggest investigating ways in which such consortia can be further incentivized

Facilities Impact and Open Questions

- Facilities impact of more extensive use of beyond-pledge resources for analysis is minimal from a technical perspective
  - T2 ATLAS Connect targets, FAX/FAXBox, PanDA for T3 submissions,…
  - Recommend lab-hosted interactive platform (BNL, others?)
- Funding open questions
  - Exact prescription for scaling off of LRU based on active analyzers
  - Funding for large, shared T3 and consortia. Interest? How do we frame?
  - Extra support level from Ops Program to support new tools/technologies
  - Need to crisp arguments for value-added by local T3 computing for JOG
- Implementation open questions
  - How do T3 jobs access data? Accounting of resources? Support level?
    - Should T3 sites in general stand up DDM-aware storage (e.g. Rucio)?
      - Given level of admin T3 support, seems “no”, unless there is a way to get benefits of Rucio caching w/o the support burden

Bonus Material

T3IC Charge (Needs)

- To carry out a comprehensive study that proposes a cost-effective implementation plan to address the T3 challenges. […] Among the questions and issues you should address are:

  - Provide a best estimate of the computing capacities required to satisfy the physics analysis activities in the U.S. over the next five years
  - Address how far the existing T3 infrastructure goes to accomplish these goals and how to make better use of all existing resources to support the U.S. physics analysis requirements
  - Address the incremental capacities that are needed to provide adequate support for U.S. physics analysis

T3IC Charge (Plans)

- […] Among the questions and issues you should address are:
  - Identify potential implementation plans that address the T3 needs of U.S. ATLAS physicists and evaluate their cost-effectiveness. Your evaluation of the cost-effectiveness should take into account and identify any synergies, efficiencies, institutional or laboratory leveraging, potential for existing or additional funding sources, and possibly other intangibles. A final ranked comparison table should summarize the cost effectiveness of these plans.
  - Identify how the T3 computing resources and personnel would be managed for (centrally, institutionally, etc.) for each plan. This is important for any solutions that require Operations Program funds.
  - Identify what can be accomplished within the current Operations Program budget guidance and prioritize additional requests if supplemental funding materializes (from Operations or other sources)

Candidate Plans for Tier-3 Evolution

- We developed three candidate plans for Tier-3 evolution into Run 2
- These plans were chosen to have elements of realism from *both* funding and implementation perspectives and be sufficiently broad to include the range of possibilities discussed both inside and outside of the T3IC
  - More granularity in the plans and additional plans-of-action are possible, of course, but we aim keep things are simple as possible, but no simpler
- The plans are described in detail and compared for cost-effectiveness in our report

Aside: US Beyond-Pledge Resources

Besides Simulation and Analysis a variety of Workflows are executed at T2s: MC reco (pile), Evgen, validation

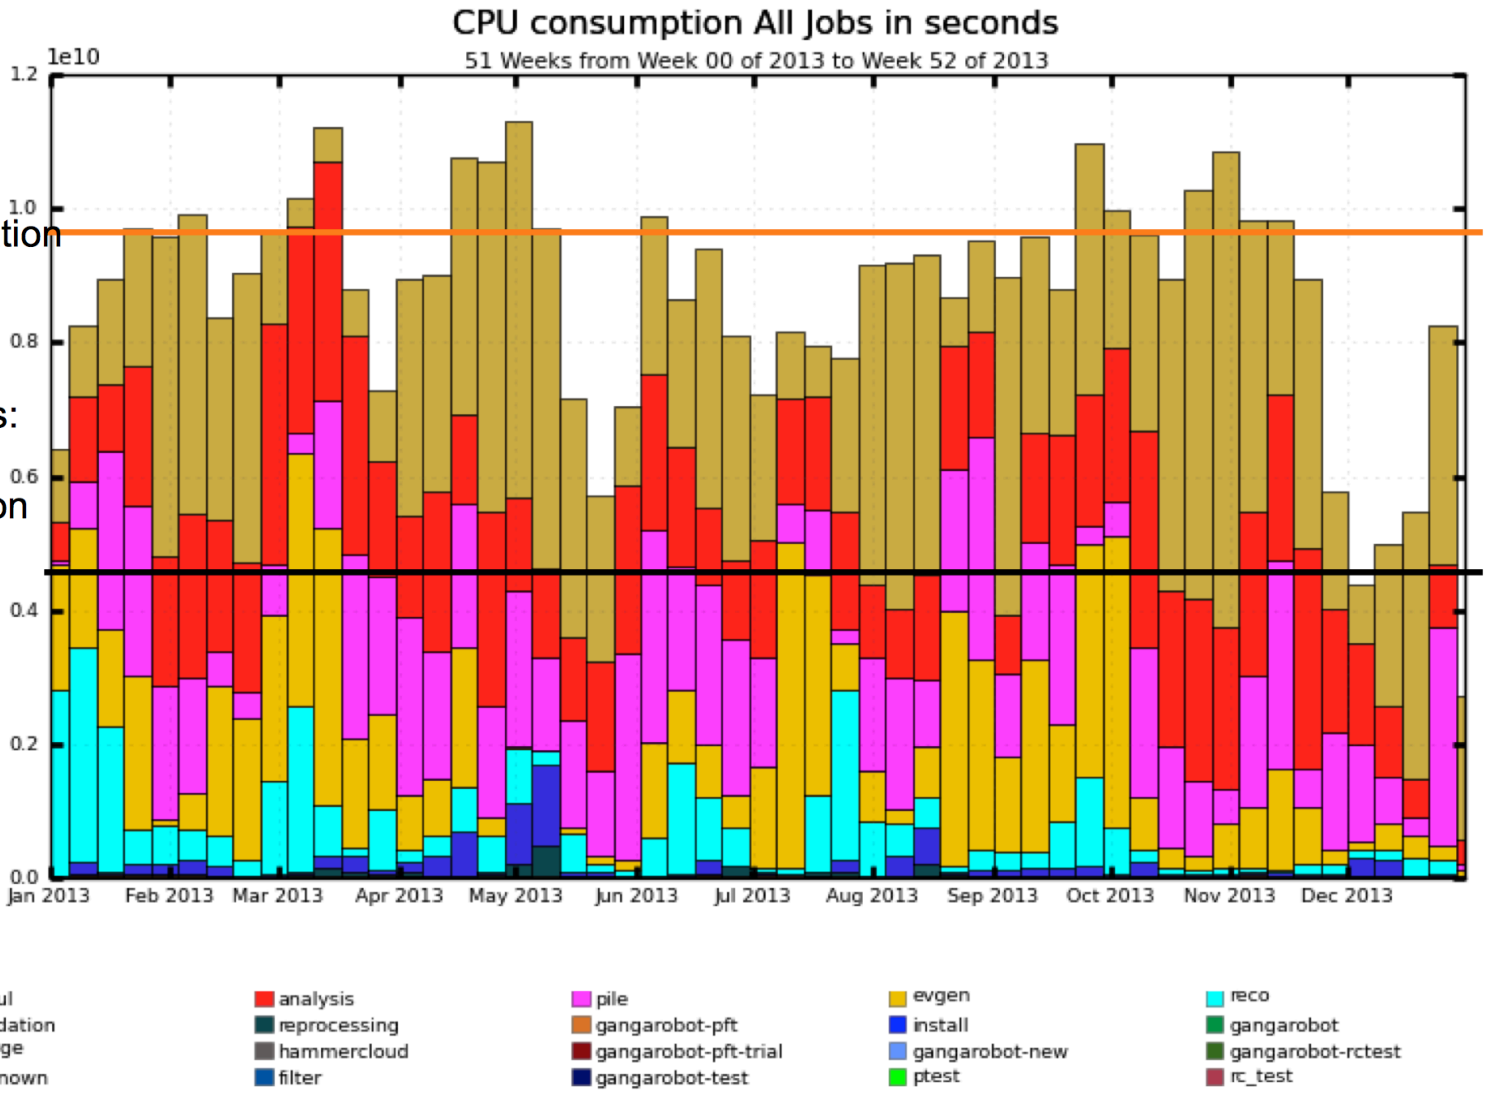

Candidate Plan 1

- Execute an approximate refresh of the *current* Tier-3 system and utilize *existing* tools for Tier-3 analysis
  - At most sites, the Tier-3 computing hardware is either out of warranty or very nearly in that state. New funds would primarily be spent to replace the aging Tier-3 hardware at the existing sites
  - This plan would require new funds at a similar level as the funding used to procure the present Tier-3 system
    - Primarily DOE ARRA + NSF Collaborative MRI awards
    - A similar or greater amount was contributed by universities through on-campus funds or in-kind contributions
  - Involves no new US ATLAS operations program investment in tools and technologies to more efficiently utilize local Tier-3 resources or Tier-1/2 facility resources for US physics analysis
  - Hardware bought in 2015+ will have more CPU and storage capability / $
    → covers sizeable fraction of x4 in extra resources needed for Run 2

Candidate Plan 2

- Repurpose existing Tier-3 hardware at each site to provide a more optimal mix of interactive and batch computing within the Tier-3 complex. Sites with Tier-3 resources would use them for interactive and small-scale batch jobs characteristic of the last few steps of analysis workflows
- Heavier analysis workloads that cannot be run on the local resource in a reasonable amount of time would be submitted to the Tier-2 facilities *using existing tools* and use the beyond-pledge CPU and storage resources there for U.S. physicists
- *Does not* require new funding for Tier-3 hardware or US ATLAS operations program investment in tools and technologies
- *Does* require that the US ATLAS facilities continue to receive a similar level of funding to provide beyond-pledge resources

Candidate Plan 3

- As in Plan 2, Tier-3 sites repurpose their existing Tier-3 hardware to provide a more optimal mix of interactive and batch
- Each willing institution deploy a *modest* Tier-3 resource locally using the *latest hardware*. The local resource is to be used primarily for interactive and small-scale batch jobs characteristic of the last few steps of analysis workflows
  - We define this resource in terms of a Local Resource Unit (LRU)
- Deploy new tools and technologies that transparently expand the Tier-3 resources into the beyond-pledge resources on the facilities when the workload is such that it would not complete in a reasonable amount of time on the local Tier-3 resource
  - Local resources are used first for workflows in support of analysis and are augmented (as needed) by resources available on the US facilities

Comparisons of the Candidate Plans

- We compared the three candidate plans for cost-effectiveness. The details of that assessment are in our report
- Plan 1 is the simplest but does not optimize the distribution of Tier-3 resources across the US nor make investments to better use the US beyond-pledge resources (BPRs) for analysis
- Plan 2 optimizes existing Tier-3 hardware and recognizes the potential to use BPRs for analysis. However, it does not make investment to better use those resources nor does it address the issue of ageing Tier-3 hardware
- Plan 3 optimizes existing Tier-3 hardware, deploys new hardware at each institution at the appropriate level, and invests in new tool and technologies to better use the BPRs
  - Hardware/Software/Support investments are appropriate & modest

Enabling Tools in the Recommended Plan

- Central to the Recommended Plan is deploying tools that follow the commercial trend of leveraging high-speed networking to enable transparent access to *remote* resources for US physics
  - **Federated Data Access (FAX):** Provides direct read access to storage for authenticated users from any participating location
  - **ATLAS Connect:** Allows a Tier-3 cluster to be logically extended to use the facilities or cloud resources through HTCondor flocking mechanisms
  - **Rucio Cache:** Cache that may be used as part of a local Tier-3 storage system to receive output datasets from remote computing jobs
  - **PandaOnTier3:** For Tier-3s set up as PanDA sites, spare CPU cycles not utilized 100% by local users can be used transparently by ATLAS. Also provides uniform support for datasets known to PanDA
  - **Agile Infrastructure:** An approach that allows for increased flexibility in provisioning computing resources in public and private clouds
- First four initiated by US ATLAS. Agile actively used in the US