

Offloading peak processing to virtual farm by STAR experiment at RHIC.

**Jan Balewski¹, Jerome Lauret², Doug Olson³, Iwona Sakrejda³,
Dmitry Arkhipkin², John Bresnahan⁴, Kate Keahey⁴, Jeff Porter³,
Justin Stevens⁵, and Matt Walker¹**

¹ Massachusetts Institute of Technology, Cambridge, MA 02139

² Brookhaven National Laboratory, Upton, New York 11973

³ National Energy Research Scientific Computing Center, 1 Cyclotron Road Berkeley, CA 94720

⁴ Argonne National Laboratory, Argonne, Illinois 60439

⁵ Indiana University, Bloomington, Indiana 47408

E-mail: balewski@mit.edu

Abstract. The Virtual Machine framework was used to assemble the STAR-computing environment, validated once, deployed on over 100 8-core VMs at NERSC and Argonne National Lab, and used as a homogeneous Virtual Farm processing events acquired in real time by STAR detector located at Brookhaven National Lab. To provide time dependent calibration, a database snapshot scheme was devised. The two high capacity filesystems, localized at the opposite coasts of US and interconnected via Globus-Online protocol, were used in this setup, which resulted with a highly scalable Cloud-based extension of STAR computing resources. The system was in continuous operation for over 3 months.

1. Introduction

The Relativistic Heavy Ion Collider (RHIC), situated at Brookhaven National Laboratory, is the highest energy collider in the world that accelerates polarized proton beams in which the spin state of the majority of the protons will be aligned with the direction of the beam. This allows nuclear physicists to study spin properties of quarks and gluons - the constituents of the proton, schematically shown in inset a) of figure 1. The STAR experiment [1] records particles produced in the collision of two protons, shown in figure 1, which yields valuable information on the spin structure of the proton⁶. In particular, our understanding of the polarization (alignment of rotation axis) of the anti-quarks inside the proton is still very limited. By studying the relatively rare process where a quark and anti-quark annihilate during the collision of the two protons and produce a W boson (see figure 2), we can infer information about spin properties of the incident anti-quark. We measure the probability of such W production processes, by reconstructing W boson events shown in figure 2.a.

⁶ Spin is central to a variety of scientific concepts and technologies, including Magnetic Resonance Imaging machines used for medical imaging in hospitals around the world.

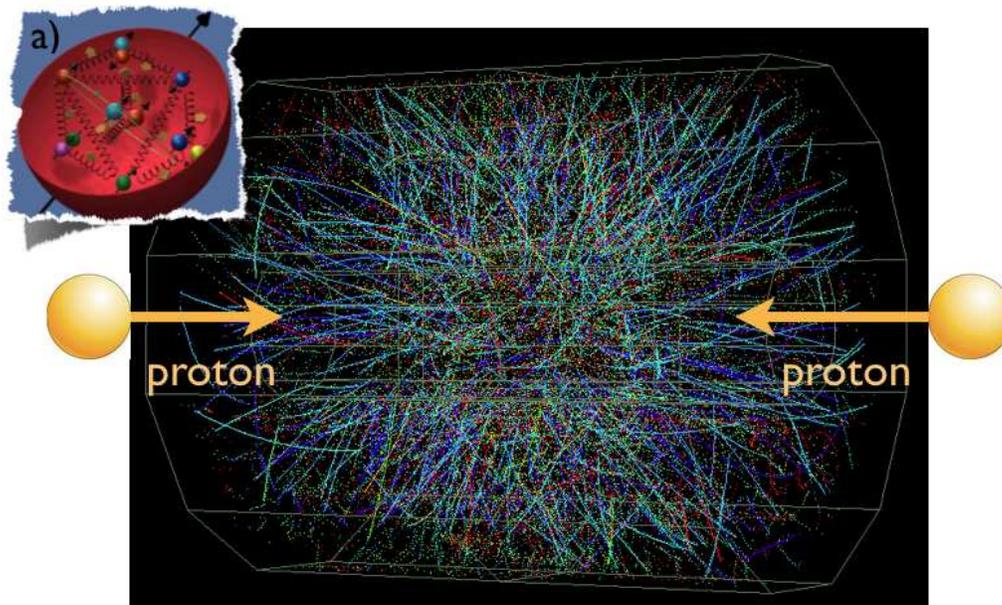


Figure 1. Example of ionization tracks recorded by STAR detector after the collision of two high energy protons. The track colors denote particles of different energy produced in the collision. The inset a) illustrates our current understanding of the structure of proton before collision: quarks, anti-quarks, and gluons interacting with each other.

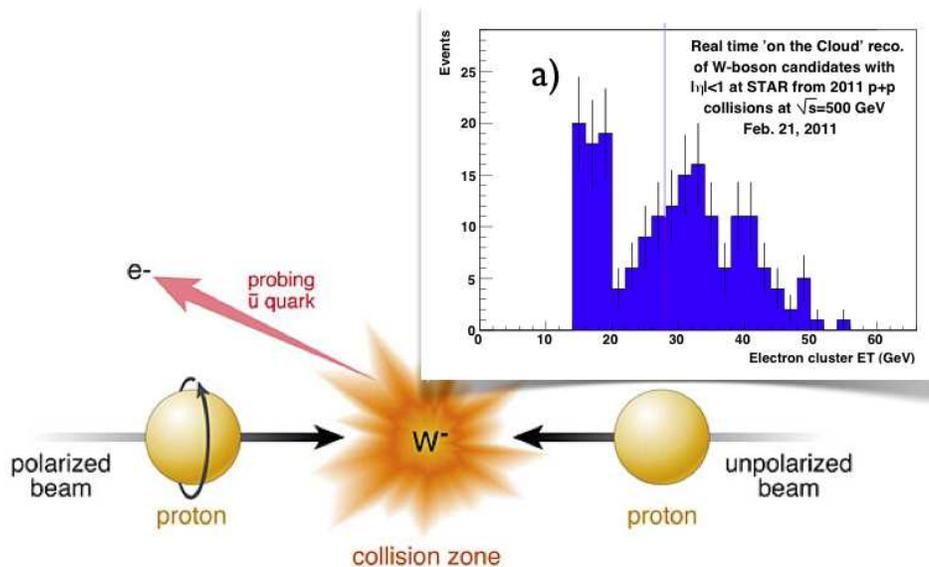


Figure 2. Illustration of the rare process of W boson production in the collision of two high energy protons. The W^- boson, shown here, decays into an electron, which if emitted in the opposite direction of the incident polarized proton, allows us to study polarization of anti- u quark inside the proton. The inset a) shows the energy distribution of electrons measured by STAR and reconstructed in almost real time using cloud resources - which are discussed later in this paper.

2. Computational challenge of STAR data mining

To study a very rare process, such as W boson production, one needs to give Nature many chances to produce it, hence RHIC is set up to provide a few million proton-proton collisions per second at the center of the STAR detector. We select and record on tape for several weeks only 1/1000-th of the most promising candidate events, still acquiring millions of events. The process of reconstructing recorded events involves sorting a hundred thousand space points, shown in figure 1 in to thousands of distinguishable particle tracks, and requires significant computing resources on the order of tens of CPU years. Furthermore, the measurement is completed in a broader sense, only after results are published. Therefore, one needs to fold in also the time needed to complete all the necessary aspects of data analysis, namely calibration of the STAR detector response, development and vetting of the W event reconstruction algorithm, discussion of the results within the STAR collaboration (consisting of ~500 members), and preparing a manuscript for publication.

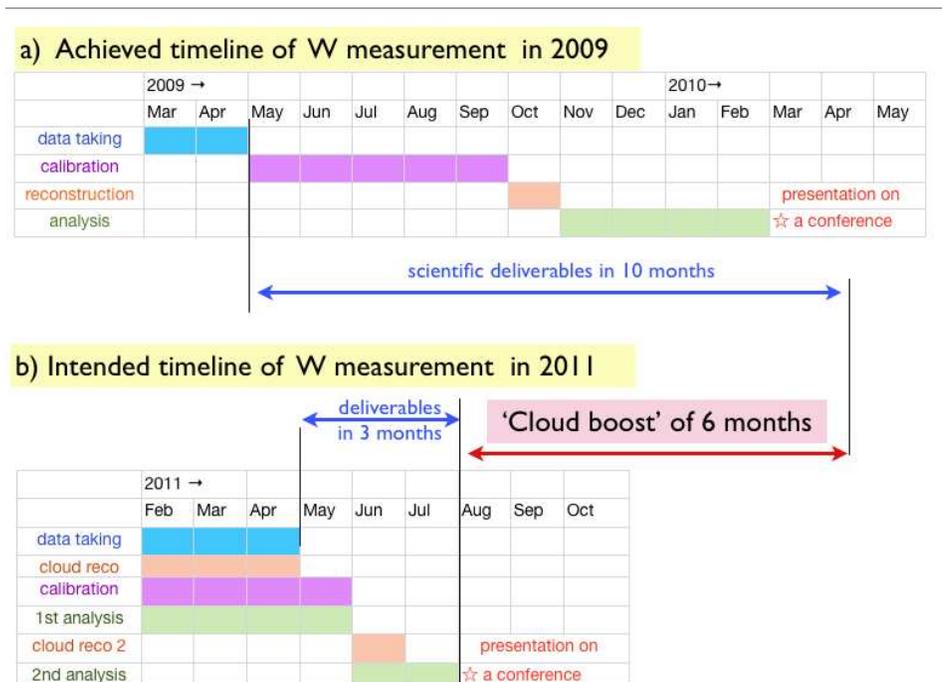


Figure 3. Comparison of the typical timeline from data collection to publication for traditional sequential data mining a) vs. parallel iterative approach benefiting from the cloud computing b).

2.1. Traditional sequential approach

The traditional approach at STAR calls for sequential accomplishing of required tasks, as illustrated by the diagram a) in figure 3. Prior to data taking the arrangement is made to reserve sufficient CPU power for the offline processing. After the data is collected on tapes, a subset of events is processed to establish detector calibrations. This calibration process takes several months. Next, all recorded events are processed on the RCF (RHIC Computing Facility) computing farm for about a month which reduces the footprint of the event files and allows for the development and tuning of the W reconstruction algorithm. Typically, it takes several months to identify and remediate problems with the data, reduce background, maximize efficiency, and vet

intermediate results within STAR collaboration. In total, it is more than a year after experiment stops collecting data, when publication quality results are released by the STAR Collaboration to the public domain. This is a long time for a student or postdoc who often needs published results from their work to advance in their career.

2.2. 'Cloud boost' of data publication cycle

The advent of Virtual Machines (VM) [2, 3] combined with abundance of cloud computing resources, such as commercial Amazon EC2 or university based farms discussed later (see table 1), allows one to restructure the strategy of extracting scientific output from the raw measurements. The diagram in figure 3.b summarizes the changes leading to perhaps a half year shorter publication cycle. The key factor is continuous access to cloud computing resources of moderate magnitude, but powerful enough to process all events of interest in real time using a temporary calibration extrapolated from data processed the previous day. The availability of reconstructed events stimulates parallel work on the refinement of the W reconstruction algorithm and enables at the same time work on the quality assurance of data recorded on the previous day. Another positive side effect of these efforts is the continuous feedback to the ongoing experiment about the performance of various detector subsystems leading to a better utilization of the limited beam time. Finally, one should not forget about psychological impact of weekly updates on the number of recorded and reconstructed W events. It changes the attitude of the STAR shift crew if they can see that with the passing weeks of monotonous data taking there is an unambiguous measure of progress, e.g. a plot of reconstructed W candidates as function of measured energy, shown in figure 2.a.

Table 1. History of VM-based computation by STAR.

Date	Facility	Tools	Type of Task	# of VMs	# of Jobs per VM	Total CPU (days)	Calend. Days	Total Input (TB)	Total Output (TB)	Remarks
2009, March	Amazon EC2	Nimbus, Globus	simu	100	1	500	5	0	0.3	globus GK grid site
2009, Nov.	Amazon EC2	EC2	simu	10	1 or 2	1	1	0	0.1	use commercial interface
2010, Feb.	GLOW †)	Condor VM	simu	430	1	130	0.6	0	0.1	'call home' model
2010, July	Clemson Uni	Ketrel KVM	simu	1000	1	17,000	20	0	7	short VM lifetime, no ssh to VM
2011, Spring	NERSC, ANL	Eucalyptus, Open-Stack	data reco	20-120	8	25,000+	120+	60	40	real time processing

† University of Wisconsin - Madison

3. STAR encounters with VMs

Over last 3 years STAR has ramped up its investment in the utilization of cloud computing resources [5, 6], as summarized in table 1.

At first STAR ran small batches of simulation jobs, which are easier to execute on a remote site since the input is identical for all jobs, negligible in size, and all compute nodes could be made identical. Also, produced events were transferred from the remote location to RCF which did not require ssh capability from RCF to compute nodes - the configuration preferred by IT at the remote site because of security concerns. The largest STAR simulation using cloud computing to date was completed in July of 2010. At it's peak, STAR was running on 1300 nodes, as shown in figure 4. We acquired a total of 47 CPU years over 3 calendar weeks and transferred back to RCF 7 TB of data. The practical benefit of this simulation was accelerating the work towards a Ph.D. for one of us (M.W.) by about a year.

Encouraged by this success we decided to increase the complexity and embark on the more ambitious task of almost real time processing of experimental data acquired at STAR in 2011. The challenges of two way data transfer, organization of asynchronous data processing, and updating calibrations will be discussed in the next section. Here, we only note this task lasted over 3 months, peaked at about 900 parallel running jobs, resulted in an integrated CPU time of above 70 CPU years over 3 calendar months, and transferred a total of 100 TB of data.

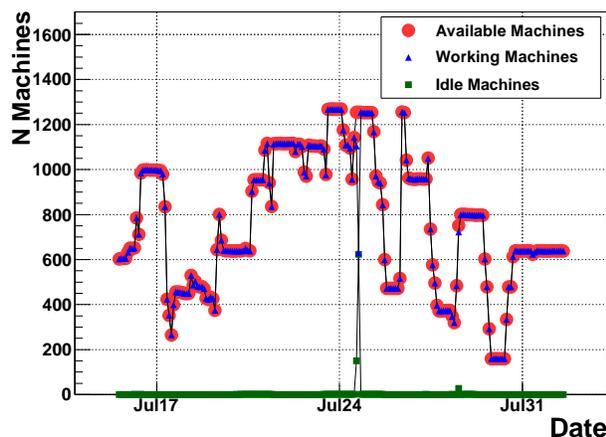


Figure 4. Load on GLOW cluster at University of Wisconsin - Madison, during 2 weeks of STAR simulations.

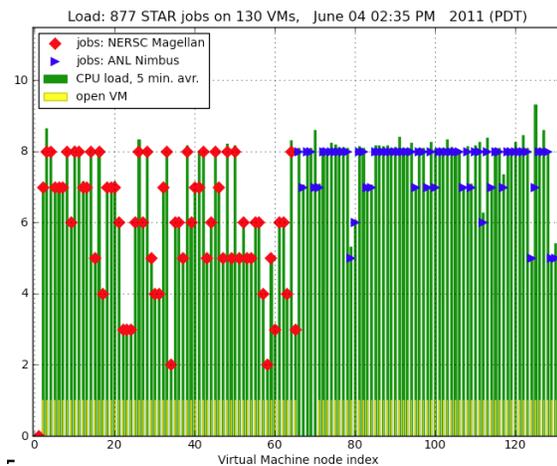


Figure 5. Load on Magellan cluster during real time STAR data processing, see text for details.

3.1. Advantages of VM

The use of Virtual Machines was an integral part of the STAR cloud computing exercises. The VM technology allowed STAR to separate details of the transient hardware/compiler features from persistent STAR code libraries. Some effort was necessary to understand the nuances between the installation of the STAR core code directly on a physical machine, and the assembling and remote deployment of a STAR VM at the computing facilities. However, the long term benefits of the use of VMs was of much larger value.

First of all it is sufficient to validate STAR VM only once and then re-use the same VM image multiple times on different machines located at different facilities. Additionally, the use of VMs has homogenized the STAR computing environment. The burden of accounting for the hardware differences at different computing facilities was shifted to the remote facility administrator assuring the VM hypervisor interacts properly with the local hardware and conforms with local

security requirements. Finally, for a multi-decadal experiment like STAR it is essential to retain the ability to execute past analysis code at any time in the future. The typical obstacles are irreversible upgrades of the operating system and hardware (architecture) retirements. Archiving of the present STAR code in the form of a VM assures we will be able to execute it in the distant future on not yet defined architecture as long as it will support virtualization as we know it today.

4. Real time STAR data processing on the cloud

The resources of Magellan - the government-funded cloud computing testbed, have been made available to STAR. Magellan resources was split between two facilities. One, based at the National Energy Research Scientific Computing Center (NERSC) in Berkeley, California uses Eucalyptus, a widely used open source cloud platform. The second, based at Argonne National Laboratory near Chicago, Illinois, hosts two clouds; one runs the OpenStack software, while the other uses the Nimbus toolkit [7]. The availability of computing power was essential but not sufficient for the real time STAR data processing. We needed to also assure the input data and calibration were sent to the compute nodes, and then transported back the results. Finally, the STAR VMs themselves have been customized to continuously occupy all available cores.

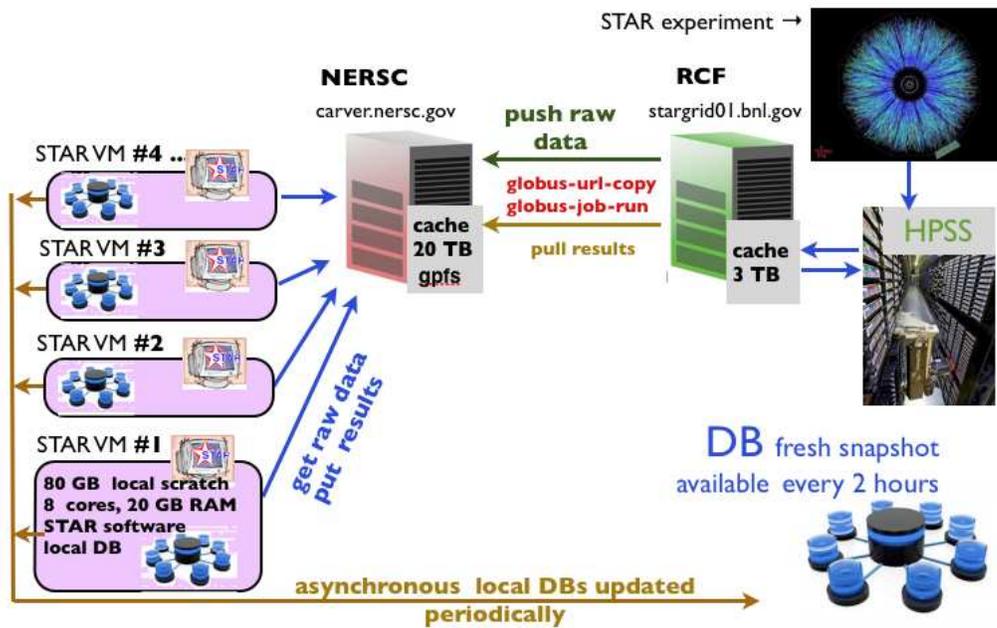


Figure 6. Example of the data flow for almost real time STAR processing on a single remote facility (NERSC) . The copies of input files produced by STAR experiment were retrieved from HPSS at RCF to a local scratch disc, sent via globus command to scratch disc at NERSC, and further using scp to individual VMs. The output files followed the same path in the reverse order.

4.1. Configuration of STAR VM

The general layout of the STAR data processing flow is illustrated in figure 6. Two independent processes check every half an hour for new files: one at RCF for input and another at NERSC for output. Both used Globus Online [8] to transfer the files they find to the scratch disc reserved at NERSC and RCF, respectively. The third process at RCF would archive output files at the

local HPSS [9, 10] and free up disc space on both scratch discs to make room for new files. Every two hours, a fourth independent process takes a snapshot of the calibration data stored at Brookhaven, which changes much less rapidly than event data taking, and makes it available for download for all VMs.

At the remote computing facility, tens of eight-core VMs were running the STAR analysis software. Each VM could run up to eight jobs at a time to occupy all its cores. Each VM acted as an independent, greedy consumer of data by requesting a new input file from the scratch disc at NERSC whenever it determined it had available CPU to spend. The main challenge was to preserve independent, unsupervised data reconstruction on different VMs without processing the same input file multiple times. It was resolved by using the atomic rename operation of Linux, which uniquely renames a selected input event file and passes the new name to the VM that requested a new file. If multiple VMs try to 'lock' the same file at the same time, only one of the atomic rename processes will succeed. The remaining VMs would continue to request files periodically until their request succeeded; the result was that eventually, either all of the VMs will be analyzing data on all cores or the pool of input files on the scratch disc will be empty.

STAR data can only be processed on a machine if the time dependent calibration constants are valid up past the timestamp of the data. Once every 24 hours, at a fixed time chosen at random per VM, a cron job running on each VM pulled the most recent calibration snapshot from the cluster back in Brookhaven. Then the local copy of the calibration data on the VM was replaced. Since each VM initiated this process at a different time of a day, it ensured that some VMs always have a fairly recent copy of calibration and were ready to process fresh STAR data. A side benefit of this approach was the reduced load on the central machine distributing calibration constants.

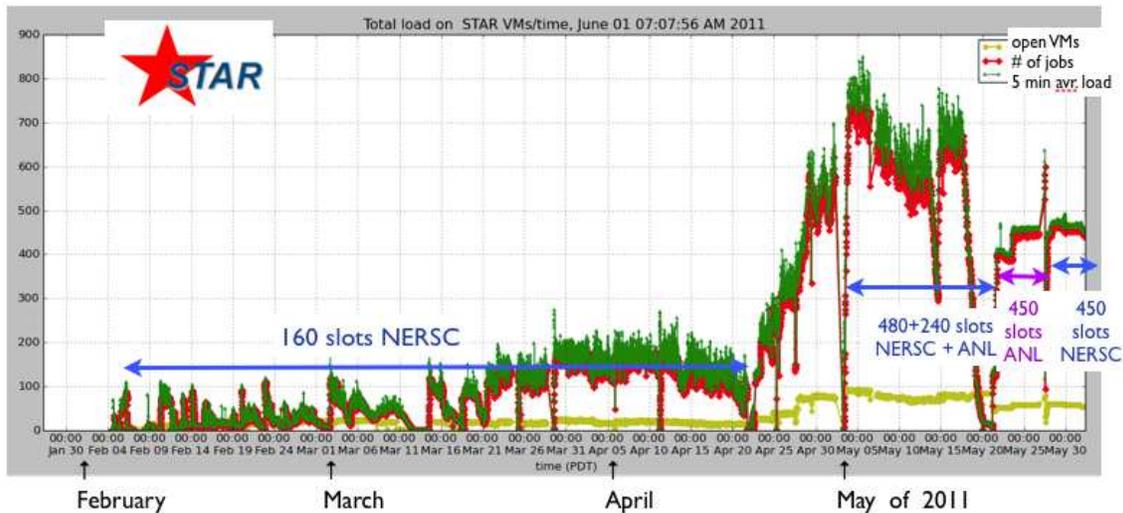


Figure 7. Total number of STAR jobs (red) and CPU load (green) for real time data processing over a period of 3 months. The total number of VMs used is shown in yellow. At different times we used resources from different facilities, as indicated by labels.

4.2. Monitoring of data processing

As with any real time process it was essential to deploy a short latency monitoring, allowing real time assessment of the performance of the data flow and pointing to possible bottle necks. A sequential survey of all responding VMs was conducted periodically and the aggregated

performance summary plots automatically posted on the web. We have used NERSC portal for that purpose. An example plot showing instantaneous load of about 900 jobs running on 130 VMs is shown in figure 5. Each VM is represented on the x-axis as a single bin. The red diamonds (blue triangles) denote the number of jobs running on a given VM located at the NERSC (ANL) computing facility. Each VM had 8 cores, so the optimal utilization of the VM is achieved if 8 jobs are running on it. The green bar marks the CPU load on each VM. At the time of data transfer or during calibration updates the load on a VM would exceed the available number of cores. It was acceptable for short periods of times, but was a sign of a trouble if it persisted over more than an hour. A video clip available in the online version of this paper shows an animated graph of the load as a function of time.

We were improving the real time STAR data processing scheme over 3 months in the spring of 2011, reaching the sustained average load of 600 jobs over several weeks, as shown figure 7.

5. Future prospects

STAR has plans to expand the existing scheme and integrate all available resources as illustrated in figure 8. The goal is to refine our real-time cloud-based data processing scheme to function as a self-adjusting assembly line and handle variable rate of throughput. No human intervention, nor supervisor process that orchestrates the entire data flow should be needed. Every stage of the process is governed by local rules designed to handle time-outs and refusals from other elements of the system by waiting and starting over. Time will show if we will be able to accomplish our plans as some of the computing facilities used so far are being re-tasked for other non-cloud related projects.

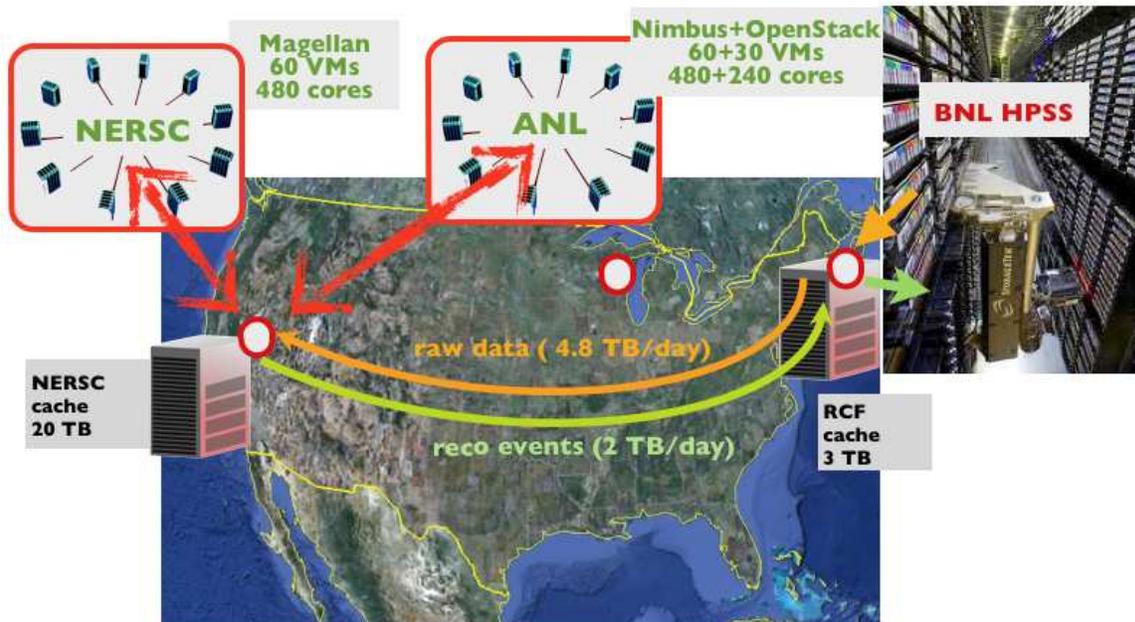


Figure 8. Target configuration of STAR data processing at multiple cloud computing facilities.

References

- [1] Adams J et al. 2005 Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR collaboration's critical assessment of the evidence from RHIC collisions *Nuclear Physics A* **757** 102

- [2] Fenn M, Goasguen S and Lauret J 2010 Contextualization in Practice: The Clemson Experience *Proceedings of 13th International Workshop on Advanced Computing and Analysis Techniques in Physics Research PoS(ACAT2010)* 027
- [3] Stour L, Murphy M and Goasgen S 2009 Kestrel: an XMPP-based framework for many task computing applications *Proceedings of the 2nd Workshop on Many-Task Computing on Grids and Supercomputers* (Portland, Oregon, November 16 - 16, 2009). MTAGS '09. ACM, New York, NY, 1- 6. DOI=<http://doi.acm.org/10.1145/1646468.1646479>
- [4] Ramakrishnan L at al. 2011 Magellan: experiences from a science cloud *Proceeding ScienceCloud '11 Proceedings of the 2nd international workshop on Scientific cloud computing* (ACM New York, NY, USA 2011 ISBN: 978-1-4503-0699-7)
- [5] Lauret J, Walker M, Goasguen S and Hajdu L 2010 From Grid to Cloud, the STAR experience *Scientific Discovery through Advanced Computing SciDAC 2010 Conference proceedings*
- [6] Lauret J, Walker M , Goasguen S Stout L, Fenn M, Balewski J, Hajdu L and Keahey K 2010 When STAR meets the Clouds Virtualization & 'Cloudified' grids experience *CHEP 2010 proceedings, J. Phys. Conf. Ser.* **331** 062016 , iopscience.iop.org/1742-6596/331/6/062016
- [7] Keahey K, Freeman T 2008 Science Clouds: Early Experiences in Cloud Computing for Scientific Applications, Cloud Computing and Its Applications (*CCA-08*), *Chicago, IL. October 2008* , see also <http://www.nimbusproject.org/papers/>
- [8] Globus Online is a fast, reliable file transfer service that makes it easy for any user to move any data anywhere, for details see <https://www.globusonline.org/about/>
- [9] LAURET J and YU D 2010 BNL Batch and DataCarousel systems at BNL: A tool and UI for efficient access to data on tape with faire-share policies capabilities *PoS(ACAT2010)* 023
- [10] Yu D and Lauret J D 2010 Tape Storage Optimization at BNL *CHEP 2010 proceedings, J. Phys. Conf. Ser.* **331** 042045