ROOTlets and Pythia: Grid enabling HEP applications using the Clarens Toolkit

Conrad Steenberg[1], Jaliya Ekanayake[2], Geoffrey Fox[2], Faisal Kahn[1], Abe Lederman [3], Vladimir Litvin [1], Iosif Legrand[1], Harvey Newman[1], Shrideep Pallickara[2], Michael Thomas[1], Frank van Lingen[1]

[1]1200 E. California Blvd., MC 356-48, Pasadena, CA, 91125

[2]Department of Computer Science, Lindley Hall 215, Indiana University, Bloomington, IU, 47405

[3] Deep Web Technology, 301 North Guadalupe Street, Suite 201 Santa Fe, NM 87501

E-mail: `conrad@hep.caltech.edu`

Abstract.

We describe how we have used the Clarens Grid Services Toolkit to develop powerful application and browser-level interfaces to the ROOT and Pythia codes used widely in the HEP Community. The Clarens toolkit is a code-base that was initially developed under the auspices of the Grid Analysis Environment project at Caltech, with the goal of enabling users to bring the full power of the Grid to their desktops, while at the same time not altering the look, feel and interface of the users chosen analysis tool. By wrapping existing applications, and providing a well documented API to the wrapper, clients are able to exchange commands, data and results using standard protocols including XML-RPC, HTTP and HTTPS. In particular, we have implemented a wrapper to the Pythia particle collision simulation code, and developed an encapsulated form of the ROOT environment, called a ROOTlet, that allows the power of the Grid to be easily applied to HEP simulation and analysis tasks. These developments, together with some usage and monitoring summaries, are covered in this paper.

1. Introduction

CERN's Large Hadron Collider is due to begin operations in 2008, and at that time several large experiments will begin to accumulate particle collision data at continuous rates from several hundred MBytes per second, up to over a GByte per second. The data will be distributed over international networks[7] to collaborating institutes and universities, where it will be processed and analyzed by physicists and engineers engaged in the experiments. The experiment collaborations each number several thousand scientists, who are spread among hundreds of institutes situated in all World regions.

Up to now, the TeraGrid has played a significant role in providing compute resources for in-depth detector simulations that have been used to study the performance of CERNs CMS experiments apparatus in particular, and its capabilities for precise measurements of LHC collision data. In this paper we describe work that is enabling the TeraGrid to offer attractive physics analysis services to the LHC community.

2. The Grid Analysis Environment

The GAE consists of a collection of components (services, portlets and portals) that provide a transparent environment for a physicist to perform batch or interactive analysis in a distributed dynamic environment. These components expose functionality to physicists wanting to perform tasks such as identifying data catalogs (and datasets), submit (complex) workflows, providing fair access to resources through priority and accounting (administration), monitoring job progress and storage and retrieval of the results.

Currently, support is available for:

Remote data access, through a file service and the ability to visualize ROOT files (format used by physicists for data storage and analysis) Browser based access to various CMS-specific production services Remote shell access for command line processing Dynamic discovery of services. Search for new or existing services and interact with these services through an associated portlet. Support for Virtual Organization Management (VO) Scheduling/Job submission. Portlet to submit jobs to multiple submission systems (e.g. Condor, PBS,) . Global Monitoring. Create global views of the system health or zoom in to a particular site for analysis and diagnostics. Job Monitoring. Interaction with running jobs (e.g view log files, suspend, kill).

The monitoring and control aspects of the GAE are managed by Caltechs MonALISA framework. This framework is designed as a set of autonomous agent-based dynamic services that collect and analyze real-time information from a wide variety of sources (Grid nodes, network routers and switches, running jobs, video devices as part of Evo[5]). MonALISA's[8] multi-threaded, self-describing agents collaborate on processing and analyzing the information they gather, so enabling them to perform a wide range of monitoring and/or control tasks. Higher level functions encoded in the agents support automated control decisions and global optimization of workflows in complex Grid systems, and other large-scale distributed applications. Agents executing in the MonALISA servers organize and collaborate with each other so as to manage access to distributed resources, to make effective workflow planning decisions, to respond to problems that affect multiple sites, or to carry out other globally-distributed tasks.

The Grid Portal / Web Service middleware used is Clarens, (clarens.sourceforge.net), which is described in the following section.

3. The Clarens Grid Portal Toolkit

The Clarens toolkit[14], available in C/Python[10] and Java[13], enables X509 based, secure, authenticated, and authorized access to resources through web services supporting multiple protocols such as XML-RPC, SOAP, and JSON. Clarens offers hierarchical group based access control in addition to perform authorization using VOMS[4], and has several ready to use clients (ROOT, python, java, javascript in a browser) for various portal developments. Clarens is used by several other projects within High Energy Physics including CMS data production[15], HotGrid[12], LambdaStation[2][3], and JobMon[11].

In addition to being made available in the Virtual Data Toolkit used by the Open Science Grid and other HEP projects, the Clarens source code is also integrated in the main source code repository used by the CMS experiment, which eases the deployment of the Clarens toolkit by HEP institutes worldwide, and helps the development of useful services both inside and outside the HEP community.

A need was identified for a lightweight version of the Clarens server that could be deployed as a "personal server" or easily embedded into other projects. Thus a ClarensLight server was developed with minimal hardware and software requirements - any system with a Python language interpreter. This server was subsequently embedded in the CMS experiments production software where it is used for inter-component communication.

4. Portal Applications

4.1. ROOTlets

ROOT (root.cern.ch) is the de facto analysis tool of choice in HEP, and is widely used in that community. The vast majority of physics analyses undertaken by LHC physicists use ROOT as a tool for I/O, processing, filtering, fitting, histogramming, plotting and production of publication-ready graphics.

In collaboration with Indiana University we have been working on facilitating ROOT-based analysis in the Grid, by the use of Clarens and the NaradaBrokering[9] publish/subscribe framework. The "ROOTlet" service makes one or more instances of the complete, unadulterated, ROOT analysis application available to lightweight client environments, most notably ROOT itself.

It should be noted that this is different from the currently available PROOF parallel processing scheme in ROOT, which is designed for operation on a cluster, and relies on low latency inter-node communication to form a virtual ROOT supercomputer. Instead, the ROOTlet service is latency tolerant and scalable to the worldwide Grid, allowing for independent execution of ROOT code by authenticated and authorized users in a sandbox on a Clarens server, with processing by either the server itself or compute nodes in a cluster controlled by a batch scheduler.

Monitoring and status notification is handled by the NaradaBrokering publish/subscribe infrastructure to provide a scalable Grid Analysis environment as shown in Figure 1. ROOT clients and the individual ROOTlet instances connect to the broker network, which is also available to outside entities like MonALISA to provide monitoring and notification functionality.

The ROOTlet service has been deployed on the Caltech Tier 2, as well as the 64 bit environment of the TeraGrid at the San Diego Supercomputing Center. On the TeraGrid the ROOTlet service is currently deployed at the San Diego Supercomputing Center using ROOT version 5.12 and the open source Torque batch scheduling system derived from PBS. A limited number of users are allowed to execute jobs using their own TeraGrid acccounts.

At Caltech a prototype deployment uses the HotGrid graduated security implementation to also provide limited access to weakly authenticated users through a group account. The Torque/PBS scheduler is used to enforce CPU and storage quotas for these users.

Figure 2 shows in a dramatic fashion the effect of an increase in computing power available to an individual physicist: with only a desktop computer, the analysis produces a signal that is hardly discernible (left panel), whereas the ROOTlet analysis shows a clear signal for the detection of a Higgs particle (right panel).

4.2. Pythia

A CMS Pythia generator service was developed for simulation of CMS data using input provided by the user. As with the ROOTlet service, user jobs can be run either on the server or on worker nodes controlled by a batch scheduler. Since the input for the Pythia service consist not of user code but of so-called datacards, or parameter files, it was possible to develop a web browser based interface, or "portal" to the service. The portal allows user-friendly interaction to Pythia, including job submission, datacard upload and result data download functionality, as shown in Figure 3.

The portal web interface is hosted by the Clarens server, and is loaded into the users browser as a standalone AJAX (Asynchronous Javascript and XML) application communicating with the Pythia and related web services on the service host using the JSON (Javascript Object Notation) message format.

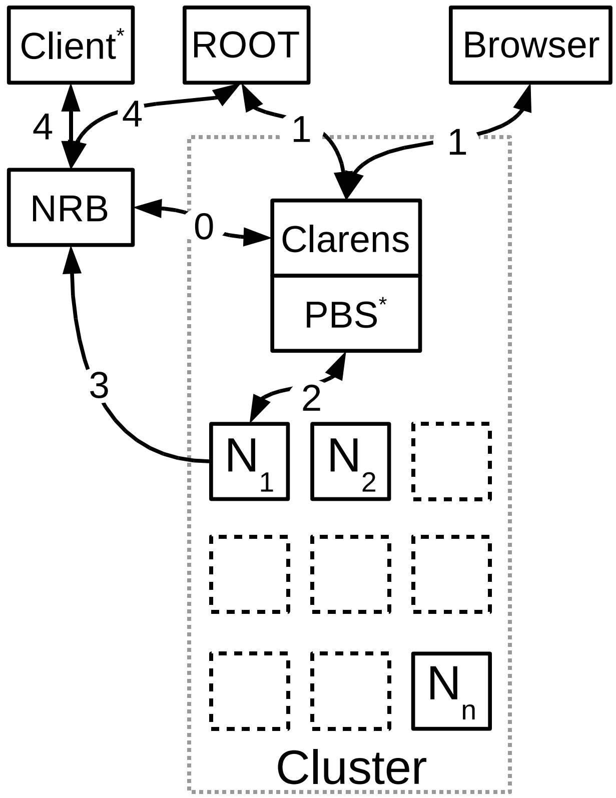

Figure 1. The ROOTlet service architecture: 0. The Clarens server publishes its service descriptions to a Narada broker, 1. Using these descriptions, clients (ROOT, AJAX Web browser, Java GUI, Python script etc.) on the Grid contact one or more Clarens ROOTlet services that 2. create ROOTlet instances (e.g. using the PBS scheduler) on compute nodes in one or more clusters. 3. Once started, ROOTlet instances publish their status via a Narada broker 4. to interested clients, e.g. a ROOT visualization.

5. Future Plans

A particularly advantageous aspect of the US-LHC, OSG and TeraGrid computing infrastructure is the co-location and/or very high bandwidth access between Tier2 and TeraGrid RP facilities at some sites, for example at SDSC. We intend to use substantial dedicated storage (10-20 TBytes initially) at the SDSC and Indiana TeraGrid RPs to mirror hot parts of the analysis datasets from CMS, so providing Tier2/3 users the availability of TeraGrid-based ROOTlet c̈rowdsf̈or remote analysis.

6. Conclusions

Substantial progress has been made in deploying High Energy Physics distributed analysis and data production on the TeraGrid. These services allow physicists from around the world to leverage the TeraGrid resources from familiar analysis applications as well as the ubiquitous web browser, user-written scripts and full blown client applications if they so choose. This flexibility meets the needs of a demanding class of power users that are used to having low level access to resources, but also need the convenience of user friendly tools to help improve overall productivity.

These Pythia and ROOTlet services are able to leverage the graduated security services

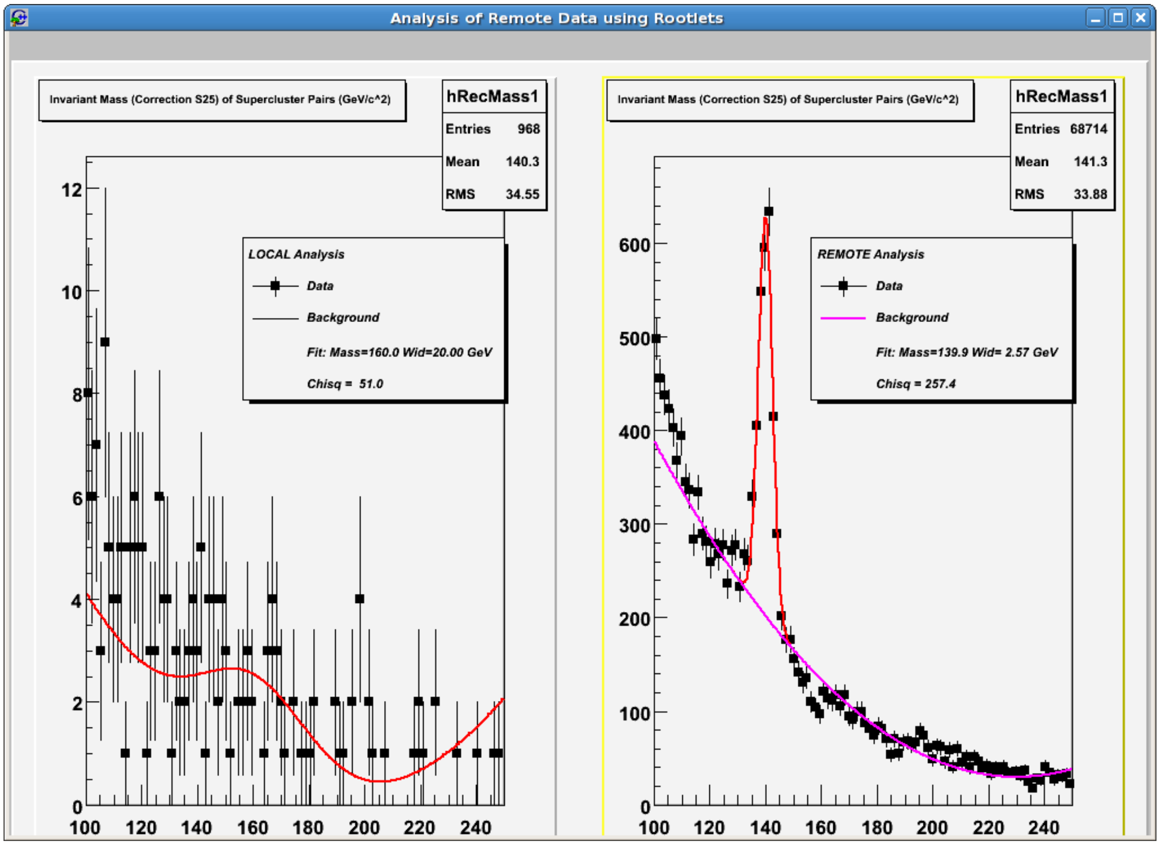

Figure 2. A ROOT-based analysis on low statistics local data (left hand plot) and the results from the same analysis when run in many remote ROOTlet services, using higher statistics event data.

developed as part of the related NESSSI portal project.

References

[1] IOP Publishing is to grateful Mark A Caprio, Center for Theoretical Physics, Yale University, for permission to include the `iopart-num` BIBTEXpackage (version 2.0, December 21, 2006) with this documentation. Updates and new releases of `iopart-num` can be found on `www.ctan.org` (CTAN).

[2] A. Bobyshev, et.al., "Lambda Station: Production Applications Exploiting Advanced Networks in Data Intensive High Energy Physics", in Proceedings of Computing for High Energy Physics Id: 235, Mumbai, India, February 13-17 2006

[3] A. Bobyshev, X. Su, H. Newman, M. Crawford, P. DeMar, V. Grigaliunas, M. Grigoriev, A. Moibenko, D. Petravick, R. Rechenmacher, M. Thomas, Y. Xia, J. Bunn, C. Steenberg, F. van Lingen, S. Ravot, "Lambda Station: On-demand Flow Based Routing for Data Intensive Grid Applications over Multitopology Networks", In proceedings of GRIDNETS 2006, San Jose, California, October 1-2, 2006

[4] G. Carcassi, T. Carter, Z. Liu, G. Smith, J. Smith, J. Spiletic, T. Wlodek, D. Yu, X. Zhao, "A Scalable Grid User Management System for Large Virtual Organization", In proceedings of Computing for High Energy Physics, Id:122, Interlaken, Switzerland, September 2004

[5] R. Cavanaugh, H. Newman, F. van Lingen, I. Legrand, Y. Xia, D. Nae, C. Steenberg, S. Ravot, M. Thomas, J. Bunn, et.al, "UltraLight: An Ultrascale Information System for Data Intensive Research", In proceedings of Computing for High Energy Physics, Mumbai, India, February 13-17, 2006

[6] P. Galvez, H. Newman, "From VRVS to EVO, the Next Generation Grid-enable Collaborative System", In proceedings of Computing for High Energy Physics, Id:154, Mumbai India, February 13-17, 2006

[7] H. Newman, J. Bunn, D. Bourilkov, R. Cavanaugh, I. Legrand, S. Low, S. McKee, D. Nae, S. Ravot, C.

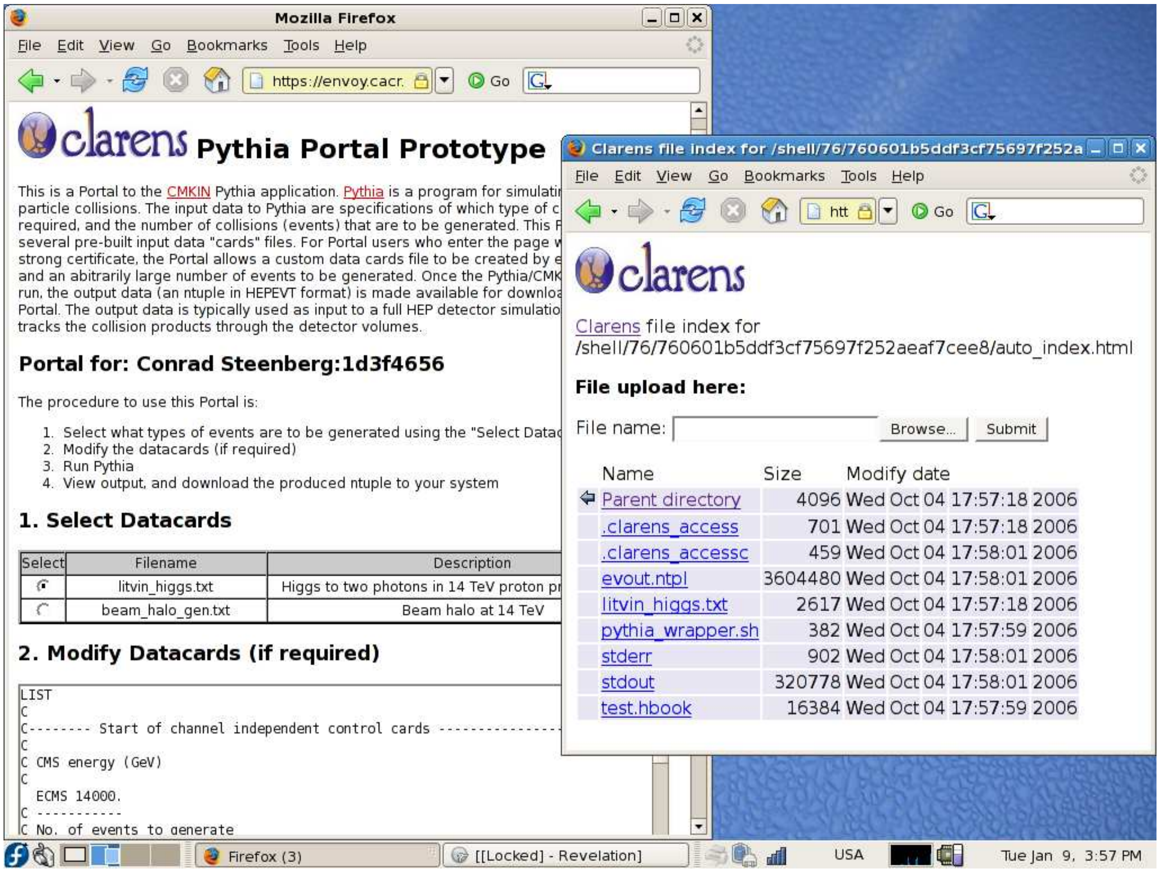

Figure 3. The browser interface to the Clarens-mediated "Pythia Portal".

Steenberg, X. Su, M. Thomas, F. van Lingen, Y. Xia, "The Motivation, Architecture and Demonstration of the UltraLight Network Testbed", In proceedings of CESNET, Prague, Czech Republic, 6-8 March, 2006

[8] I. Legrand, H. Newman, C. Cirstoiu, C. Grigoras, M. Toarta, C. Dobre, "MonALISA: an Agent Based, Dynamic Service System to Monitor, Control and Optimize Grid Based Applications", in Proceedings of Computing for High Energy Physics, Mumbai, India, February 13-17, 2006

[9] Shrideep Pallickara, Marlon Pierce, Harshawardhan Gadgil, Geoffrey Fox, Yan Yan, Yi Huang, A Framework for Secure End-to-End Delivery of Messages in Publish/Subscribe Systems. Proceedings of the 7th IEEE/ACM International Conference on Grid Computing (GRID 2006). Barcelona, Spain, 2006

[10] C. Steenberg, Eric Aslakson, J. Bunn, H. Newman, M. Thomas, F. van Lingen: "The Clarens Web Services Architecture", Proceedings of CHEP 2003, paper MONT008, 2003

[11] C. Steenberg, S.C. Hsu, E. Lipeles, F. Wuerthwein, "JobMon: A Secure, Scalable, Interactive Grid Job Monitor", In proceedings of Computing for High Energy Physics, Mumbai, India, February 13-17, 2006

[12] R. Williams, C. Steenberg, J. Bunn, "HotGrid: Graduated Access to Grid-based Science Gateways", in Proceedings of IEEE Supercomputing Conference, Pittsburgh, USA, November 2004

[13] M. Thomas, C. Steenberg, F. van Lingen, H. Newman, J. Bunn, A. Ali, R. McClatchey, A. Anjum, T. Azim,W. ur Rehman, F. Khan, J. Uk In, "JClarens: A Java Framework for Developing and Deploying Web Services for Grid Computing", In Proceedings of 2005 International Conference on Web Services (ICWS 2005), Orlando, Florida, July 11-15, 2005, IEEE Computer Society Order Number P2409 ISBN 0-7695-2409-5 pp141-148

[14] F. van Lingen, J. Bunn, I. Legrand, H. Newman, C. Steenberg, M. Thomas, A. Anjum, T. Azim, "The Clarens Web Service Framework for Distributed Scientific Analysis in Grid Projects", In Proceedings of the International Conference on Parallel Processing Workshops, Oslo, Norway, June 2005, IEEE Computer Society Order Number P2381, ISBN 0-7695-2381-1, pp45-52

[15] F. van Lingen, J. Bunn, I. Legrand, H. Newman, C. Steenberg, M. Thomas, Y. Xia, D. Bourilkov, R.

Cavanaugh, D. Evans, E. Lipeless, S. Hsu, T. Martin. A. Rana. F. Wuerthwein, ”Suporting on Demand, Policy based Monte Carlo Production, Leveraging Clarens, and RunJob”, Poster at Computing for High Energy Physics, Mumbai India, February 13-17, 2006