



Cloud Data Sharing Platform

Final Project Report

S-67G

Power Systems Engineering Research Center

*Empowering Minds to Engineer
the Future Electric Energy System*



Cloud Data Sharing Platform (S-67G)

Final Project Report

Project Team

Carl Hauser, Project Leader
Anjan Bose
Ming Meng, Graduate Student
Dave Anderson, Graduate Student
Washington State University

Ken Birman
Theodoros Gkountouvas, Graduate Student
Weijia Song, Graduate Student
Cornell University

PSERC Publication 16-05

April 2016

For information about this project, contact

Carl Hauser, Project Leader
Associate Professor
School of Electrical Engineering and Computer Science
Washington State University
P.O. Box 642752
Pullman, WA 99164-2752
Phone: (509) 335-6470
Email: hauser@eecs.wsu.edu

Power Systems Engineering Research Center

The Power Systems Engineering Research Center (PSERC) is a multi-university Center conducting research on challenges facing the electric power industry and educating the next generation of power engineers. More information about PSERC can be found at the Center's website: <http://www.pserc.org>.

For additional information, contact:

Power Systems Engineering Research Center
Arizona State University
551 E. Tyler Mall
Engineering Research Center #527
Tempe, Arizona 85287-5706
Phone: 480-965-1643
Fax: 480-965-0745

Notice Concerning Copyright Material

PSERC members are given permission to copy without fee all or part of this publication for internal use if appropriate attribution is given to this document as the source material. This report is available for downloading from the PSERC website.

Acknowledgements

This is the final report for the Power Systems Engineering Research Center (PSERC) research project titled “Cloud Data Sharing Platform” (project S-67G). We express our appreciation for the support provided by PSERC’s industry members and by the National Science Foundation under the Industry / University Cooperative Research Center program.

We wish to specifically thank Eugene Litvinov and ISO New England for their sponsorship of this project, and to thank Eugene, Xiaochuan Luo, Qiang Zhang, and Feng Ma at ISO NE for their technical support and guidance throughout the project. Thanks also to industry advisors Jay Giri of Alstom Grid and George Stefopoulos and Atena Darvishi of NYPA.

Executive Summary

This final report details the findings of a study conducted by WSU, Cornell and ISO NE aimed at understanding the practical challenges associated with using cloud computing infrastructures to monitor PMU devices in a deployment intended to be similar in scope and scale to the Northeastern regional power pool. The core idea was to set up a cloud computing system secured in ways fully responsive to the relevant ISO security policies, and then send simulated PMU data to the cloud system over secured communication links. State estimation was then carried out in real-time, with the results relayed back to the ISO for visualization.

This study built on the GridCloud project sponsored by the U.S. Dept. of Energy, ARPA-E GENI program, in which Cornell University (Cornell) and Washington State University (WSU) developed technologies supporting real-time sharing of utility operational data (esp. PMU data streams) using cloud computing. GridCloud technologies overcome some apparent limitations of commercial cloud offerings, so that utilities can take advantage of the cost savings that cloud computing offers for dynamic computing load conditions.

In addition to the careful study of security requirements, options, and their degree of match, this project quantified latency, from when PMU data was captured to when the corresponding state estimate became available, and compared latency and data quality when the identical system was deployed in duplicate with one instance running on an Amazon data center on the US Northeast and the second running remotely in the Pacific Northwest.

Our main findings were as follows:

1. The ISO security team requested changes to some aspects of the GridCloud deployment as originally undertaken in the ARP Ae research prior to our PSERC study. At the ISO's request, we modified GridCloud to support a new system configuration that the ISO felt more comfortable with. The change involved a “relay” computer owned and operated purely by the ISO, playing a firewall role: PMU data from the ISO was relayed to the GridCloud system through the relay machine, and results of the LSE computation were made available to the ISO control center through connections from visualization software running in the ISO to data relay applications on the relay computer. The ISO specified that all data sent to or from the cloud should be encrypted and suggested either transport-level security (TLS) or HTTPS; we ultimately employed tunnels over the Secure Shell (SSH) protocol because they proved to be the most convenient. Directionality of these relay connections was important to the ISO: they wanted only connections initiated from the ISO system to the cloud, with none initiated from the cloud to components owned by the ISO.
2. In our prior ARP Ae study, each PMU was separately connected to three data collectors in each data center running GridCloud. ISO NE preferred a simpler scheme in which a single connection was made from the relay machine to a single data collector within GridCloud in each datacenter. We implemented this method and carried out experiments using it, but we concluded that this configuration is not ideal and that a

configuration more like the earlier ARP Ae approach should be favored in future experiments.

3. With two GridCloud instances running in different datacenters we found that the instances received 100% identical raw data, and gave nearly identical outputs. We observed single-point data discrepancies only in one very artificial situation, namely when we “looped” the 11s of replayed data to create a longer 30 minute run. Since the input data was in fact identical we are certain that these oddities were an artifact of the manner in which the looped reset operation was implemented and that GridCloud should be capable of perfect fidelity between two cloud computing data centers. When looped, the data contain a discontinuity that would not occur in a real power system. Since the output differed only twice and only at the exact point where this discontinuity occurred we believe that the problem is not fundamental. We did not have adequate time to track down our bug before funding was exhausted.
4. End-to-end latency was approximately 300ms with a nearby data center running the LSE, to which should be added the internet delay from the PMU to the data center, approximately 105ms each-way in our cross-country experiments. Thus, the ISO was able to receive LSE outputs at 5Hz with a delay of about 300ms from a nearby data center, and with a delay of about 510ms with LSE running on a data center on the opposite coast of the United States. As explained below, a surprisingly high portion of this delay is apparently attributable an artificial delay during data alignment in the LSE. We believe the latency could be reduced by as much as 180-200ms with more time and effort since the linear state estimation task itself is not computationally expensive for this amount of data.
5. The relay machine was found to add approximately 8-15ms of latency to the LSE computation, with an additional 2ms when using SSH for the required security. The extra delays for security are extremely modest given that the overall end-to-end latency was 300ms even with a nearby Amazon data center in Virginia.
6. During our experiments, Cornell did some network upgrades which caused high loss rates and high delays on November 16 (see Appendix B). The GridCloud experiment was not impacted by these problems, confirming that the system is highly fault tolerant. We note that in prior ARP Ae studies with redundant TCP connections, three from each PMU to each data collector, we never saw situations in which all three connections simultaneously exhibited high latency.
7. Although we did not experiment with node crashes during the ISO study (due to a lack of time) we did experiment with complete data center shutdowns. Restart required approximately 175s, during which no data was lost because the Oregon data center continued to operate while the Virginia one was recovering. The vast majority of this time was covered by booting the Windows State Estimation instance. Then, the delays of reconfiguration and initial setup of the SE instance are non-significant compared to 175s. Thus at the end of the 175s period, the GridCloud system was back to full function (state estimation included) and full redundancy. In ARP Ae experiments we also explored cases where individual compute nodes were crashed; in all such situations, the built-in CloudMake manager restarted the failed component within a similar delay.

8. Microbenchmarks revealed that the cost of AES 256 encryption is in the noise. We experiment with and without encryption both on network links from the ISO to the cloud and for network model and SE outputs stored to disk by the GridCloud data historian and in both cases, the cost of encryption was so small as to be completely negligible.
9. The ISO requested that we extend the platform to report the LSE output as a series of IEEE C17.118 data streams for subscription and visualization within the ISO data center. This extension was easily accomplished, confirming that the GridCloud system as a whole is quite easily modified and extended with new functionality.
10. The ISO requested that when using Amazon storage (S3) encryption, we take steps to ensure that the encryption key would never be stored anywhere in the Amazon infrastructure. We were not able to provide this guarantee: Amazon's normal setup does retain the given key, in a special-purpose secure key repository maintained by the company. We recommend that ISO NE negotiate with the Amazon government cloud organization to have this feature added, if ISO NE security officers feel that it is a strong requirement.
11. The work reveals a series of next steps, which include bringing a second user into the system (expected to be NYPA under funding from NYSERDA), modifying the data path from the PMU owner to the cloud so as to reduce a form of “fate sharing” that was found to be a performance and security limitation, and starting to use the data historian capabilities of the system.

We conclude that the experiment was highly successful and that the ISO NE security requirements can be satisfied at negligible overhead. The ISO's original goal was to carry out real-time LSE with a maximum delay of 2-5s; we achieved far better performance with latency as low as 300ms when the data center is reasonably near the PMU data sources. Data consistency was perfect, the system was confirmed to be highly fault-tolerant, and is easily extended.

One limitation concerned the Amazon secure key management model, and was discussed above. A second limitation arose from the ISO's request to limit the relay computer to using a single connection between the relay machine and the cloud infrastructure. This created an undesirable form of “fate sharing”: if this one TCP connection suffers a dropped data segment, the entire flow of data from the full set of ISO-managed PMUs is uniformly and negatively impacted until TCP's retransmission mechanism has recovered the dropped segment. In our prior ARP Ae experiment we employed multiple side by side connections, three per PMU, and found this to be a more performant configuration. We recommend that in future work, the ISO shift to this more stable, more robust configuration option.

Project Publications:

- [1] Anderson, Dave; Dave Bakken, Ken Birman, Anjan Bose, Carl Hauser, Theo Gkountouvas, Robbert van Renesse, and Weija Song. “Hosting the SmartGrid: The GridCloud Platform for Monitoring and Controlling the Bulk Power Grid.” Submitted to *IEEE Power and Energy*. November 2015.

- [2] Anderson, Dave; Dave Bakken, Ken Birman, David Bindel, Anjan Bose, Alan Ettlinger, Theo Gkountouvas, Carl Hauser, Eugen Litvinov, Edgardo Luzcando, Xiaochuan Luo, Weija Song, Robbert van Renesee, Frankie Zhang. “Enabling a Smarter Bulk Power Grid.” To be submitted to *Communications of the ACM*, March 2016.

Student Theses:

- [1] Meng, Ming. “Two-Level Three Phase State Estimation.” Ph.D. dissertation, Washington State University. Expected August, 2016.
- [2] Gkountouvas, Theodoros. “Dynamic Management of Distributed Systems.” Ph.D. dissertation. Cornell University. Expected May 2018.