On Using Cloud Platforms in a Software Architecture for Smart Energy Grids

(Poster)

Yogesh Simmhan[*‡], Mike Giakkoupis[*†], Baohua Cao[*†], and Viktor Prasanna[*†‡]
[*]*Center for Energy Informatics* [†]*Computer Science Department*
[‡]*Ming Hsieh Department of Electrical Engineering*
University of Southern California, Los Angeles CA 90089
Email: {simmhan, giakkoup, baohuaca, prasanna}@usc.edu

Abstract

Increasing concern about energy consumption is leading to infrastructure that continuously monitors consumer energy usage and allow power utilities to provide dynamic feedback to curtail peak power load. Smart Grid infrastructure being deployed globally needs scalable software platforms to rapidly integrate and analyze information streaming from millions of smart meters, forecast power usage and respond to operational events. Cloud platforms are well suited to support such data and compute intensive, always-on applications. We examine opportunities and challenges of using cloud platforms for such applications in the emerging domain of energy informatics.

1. Introduction

The growing importance of energy security and sustainability is resulting in a heightened need to monitor and control energy assets for their optimal use. Smart power grids, which leverage large scale deployment of Smart Meters and sensors connected to pervasive communication infrastructure, are being deployed globally. These support real-time, two-way communication between utilities and consumers, and allow software systems at both ends to control and manage power use.

Availability of live, fine-grained information on electricity demand helps utilities to forecast energy usage and limit their peak energy demand to better use their available energy sources and avoid rolling blackouts to consumers. The usage forecast is done through computational models that combine historic and realtime data, while load curtailment is performed by a combination of dynamic pricing of power and targeted incentives for consumers. A software system that performs *demand-response optimization* at scale requires platforms and programming models that support data and compute intensive applications.

For example, the demand-response optimization for the Department of Water and Power (LADWP) in the City of Los Angeles, that we consider in this article, will need to continuously analyze information streaming in from millions consumers. This is integrated with continuous weather and traffic information to accurately forecast energy usage. Consumer behavior models developed through mining historic usage data and realtime clustering techniques will inform active and passive responses, such as pricing signals and energy conservation suggestions, with the goal of reducing peak consumption by 200MW within 5mins of a potential overload "event".

Clouds offer advantages of scalable and elastic resources to build a software infrastructure to support such dynamic, always-on applications. But the unique needs of energy informatics applications also highlight challenges in using cloud platforms, such as the need to support efficient and reliable streaming, low-latency scheduling and scaleout, and effective data sharing.

In this article, we discuss opportunities for and challenges of using cloud platforms for the emerging domain of energy informatics. The issues we raise pose interesting research problems for cloud practitioners to address.

2. Scope for Using Cloud Platforms

A software architecture to support smart grid applications will need Cloud platforms as an intrinsic component due to the many benefits they offer. Information sources from smart energy grids will approach internet scales, with millions of consumers acting as data sources. Clouds data centers were built for such scales of data interactions, and can scale better than centralized systems at utilities (See Figure 1).

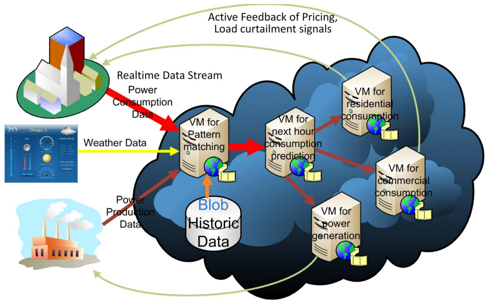

Figure 1. Smart Meters at consumers interacting with demand-response applications in the Cloud

Also, the computational demand for demand-response applications will be a function of the energy deficit between supply and demand. This typically fluctuates based on the time of the day and possible weather conditions. This translates to a need for compute intensive, low latency response at midday, and limited analysis at nights. The elastic nature of Cloud resources makes it possible to avoid costly capital investment by the utility for their peak computation needs.

Information on realtime energy usage and power pricing will need to be shared with consumers through online portals. The web presence of Cloud platforms again is well suited for this. In addition, data collected and integrated from various sources will need to be accessible to third-party applications, after meeting data privacy concerns. This will benefit developers and businesses who develop intelligent applications[1,2] for customized consumer needs, and need access to data available to the utilities in a deidentified form. In particular, in a market based power tariff model that is soon possible, making such data available to the public is necessary for transparency. Clouds provide a ready platform for data sharing and also allow third-party applications to be collocated with the data source.

3. Research Challenges for Clouds

Many of the above benefits of Clouds come with their own set of challenges, several of which are unique to the smart grid domains. We highlight some of these research problems here.

Streaming Applications in the Cloud. Like many eScience domains, a major source of data for demand-response analysis in Smart Grids comes from sensors and Smart Meters. These sample power usage and quality data at different granularities – from the building to the appliance level, and from few seconds to several hours intervals – and transmit the data to the utility or their authorized agents using internet protocols [1], [2]. While stream computing is well studied for applications like video and audio streaming [3] and intelligent transportation services [4], these are unlike Smart Grid applications that run across distributed consumer meters, Cloud platforms for data integration and clusters at utilities which introduce heterogeneity in the platforms across which efficient streaming has to be supported. There has been limited work on stream processing in public and private Clouds, and across diverse platforms [5]–[7]. At present, Cloud providers do not provide specialized data abstractions for streams, other than TCP sockets, that allows reliable data access and efficient sharing of streams across virtual machines [8], [9].

Scheduling Latency Sensitive Applications. Demand-response applications need to be responsive to the current load on the power system. The latency for responding to a demand surge will be inversely proportional to the available power capacity. The computational cost for such algorithms to reduce power demand will depend on the latency requirements and the amount of load to curtail. The execution and scheduling model for such applications will need to plan the availability of Cloud resources accordingly since the elasticity of the Cloud comes with an overhead. This planning may include starting additional virtual machines as the power begins to peak, or increasing cumulative bandwidth capacity to support higher sampling rate of streaming data. These execution policies will also have to intelligently use available compute capacity at the private Cloud and clusters available with the utilities, and make dollar cost tradeoff of enlisting additional compute power in commercial Clouds against the KW or power curtailed.

Scalable Data Sharing and Privacy Preservation. Information on energy assets that are aggregated and integrated within Cloud environments need to be communicated to external entities. These include push mechanisms on pricing and usage to Smart Meters and mobile devices, and pull mechanisms from web portals and third-party applications. Scalable access to this data will have to balance the need for open access to deidentified data while not affecting the performance of mission-critical, demand-response applications. Naively replicating all data is insufficient, given the sizes of data involved (O(1TB/day)) and the cost for incremental updates. Selective movement of data hosted on private Clouds at the utility to

1. http://www.microsoft-hohm.com
2. http://www.google.org/powermeter

the public storage space on-demand will have to be supported. Also, typical public Cloud storage platforms do not provide fine grained authorization control for data. Models for using the shared Cloud repository by multiple users and their software agents, with different levels of access, need to be examined. Some of this can be informed by similar concerns that exist in health informatics.

Acknowledgment

The authors would like to thank members of the Department of Energy sponsored Los Angeles Smart Grid Demonstration Project, including from the Los Angeles Department of Water and Power (LADWP) and USC Facilities Management Services, for discussions on the domain challenges for demand response in Smart Grids.

References

[1] "Why ip is the right foundation for the smart grid," Cisco Systems Inc., Tech. Rep. C11-581079-00, January 2010.

[2] "The role of the internet protocol (ip) in ami networks for smart grid," National Institute of Standards and Technology (NIST), Tech. Rep. NIST-PAP-01, October 2009.

[3] T. P. Nguyen and A. Zakhor, "Distributed video streaming over internet," M. G. Kienzle and P. J. Shenoy, Eds., vol. 4673, no. 1. SPIE, 2001, pp. 186–195. [Online]. Available: http://link.aip.org/link/?PSI/4673/186/1

[4] A. Biem, E. Bouillet, H. Feng, A. Ranganathan, A. Riabov, O. Verscheure, H. N. Koutsopoulos, M. Rahmani, and B. Güç, "Real-time traffic information management using stream computing," *IEEE Data Engineering Bulletin*, vol. 33, no. 2, pp. 64–68, 2010.

[5] Y. Gu and R. L. Grossman, "Sector and sphere: the design and implementation of a high-performance data cloud," *Philosophical Transctions of the Royal Society A*, vol. 267, no. 1897, pp. 2429–2445, 2009.

[6] C. S. Liew, M. Atkinson, J. van Hemert, and L. Han, "Towards optimizing distributed data streaming graphs using parallel streams," in *International Workshop on Data Intensive Distributed Computing (DIDC)*, 2010.

[7] D. Zinn, Q. Hart, B. Ludaescher, and Y. Simmhann, "Streaming satellite data to cloud workflows for on-demand computing of environmental products," in *Workshop on Workflows in Support of Large-Scale Science (WORKS)*, 2010.

[8] Microsoft Patterns and Practices, "Introduction to the windows azure platform," 2010. [Online]. Available: http://msdn.microsoft.com/en-us/library/ff803364.aspx

[9] Amazon Web Services LLC, "Amazon web services: Elastic load balancing," 2010. [Online]. Available: http://aws.amazon.com/elasticloadbalancing/