Leveraging Computational Grid Technologies for Building a Secure and Manageable Power Grid

Himanshu Khurana, Mohammad Maifi Hasan Khan and Von Welch
National Center for Supercomputing Applications, University of Illinois
{hkhurana, mmkhan, vwelch}@ncsa.uiuc.edu

Abstract

The US Power Industry is in the process of overhauling the Power Grid to make it more secure, reliable, and available. The IT systems that comprise a major part of this change need to address a number of security and management challenges for the data generated by power systems and for resources consumed by distributed applications. In this work we leverage the technologies developed by the Computational Grid to address these challenges. We identify striking similarities between the systems and use these similarities to suggest ways of addressing challenging problems; e.g., securing inter-control center communications and power grid infrastructures, and managing data and resources at control centers.

1 Introduction

The US Power Grid (P-Grid) is in the early stages of a revolutionary makeover, which, when completed, will result in a secure service that provides everyone with abundant, affordable, clean, efficient, and reliable power [10]. This vision is being driven in part by the development and integration of advanced IT systems that provide effective communication, information management, and decision making at all levels of the power grid; e.g., data acquisition in SCADA (Supervisory Control and Data Acquisition) networks, substation automation, communication links to control centers, EMS (Energy Management System) and BMS (Business Management System) data processing, inter control-center communications, and market operations and reliability assurances at ISOs/RTOs (Independent System Operator/ Regional Transmission Organization).

A crucial component of the advanced IT systems of the future P-Grid will be technologies that provide security for and manage both the resources of the P-Grid and the data that they generate, process and store. These needs are driven by various factors including increasing (1) threats from adversaries, (2) data generation capabilities of SCADA networks and data processing requirements of EMS/BMS systems, (3) bandwidth requirements that are prompting SCADA and EMS connectivity to the public Internet, (4) standardization efforts that will lead to a common suite of protocols and hardware being deployed across multiple control areas, and (5) business requirements that prompt control centers to share information with each other and with ISOs/RTOs.

In the last decade another grid that has undergone extensive design and development with advanced IT systems in place is the Computational Grid (C-Grid). The C-Grid was envisioned in the late 1990s to provide scientists and engineers across the world access to interconnected computational resources. This vision has been achieved and allows the seamless flow of computation much like power flows in the P-Grid.

In this work we look at the security and management requirements for both the data and the IT resources that process/store the data in the future P-Grid. In these broad categories of requirements we focus on ones for which tools and technologies are maturing in the C-Grid. In general, an approach of seeking mature data and resource management technologies would be natural for the P-Grid in an effort to hasten the process of deployment of these technologies on the P-Grid. In particular, seeking technologies from the C-Grid allows us to explore solutions to challenging problems that have yet to be addressed in the P-Grid.

In exploring C-Grid technologies in the broad categories of security and management for data and resources, we take a comprehensive approach in studying how these technologies can help solve problems in the P-Grid. First, we develop a qualitative framework to identify the extent to which a particular C-Grid technology/solution can be leveraged in the P-Grid. This framework defines four abstract levels of leverage, namely, conceptual, policy, protocol, and implemented tools. Second, we provide a summary of potentially useful technologies and solutions in the C-Grid. Third, in each of the broad categories we undertake a case study of a challenging P-Grid problem for which C-Grid solutions

may be applicable. We discuss specifics of the applicability of the solution and identify the extent of the applicability using our framework. In contrast, Wu *et al.* [27] suggest an exploration of C-Grid technologies for a narrower scope of increased computational performance and also do not study the applicability of C-Grid solutions in depth. To the best of our knowledge this is the first work that studies both C-Grid technologies and specific P-Grid problems in depth to identify how the technologies can be used to solve these problems at multiple levels. Earlier works simply identified the general potential usefulness of the C-Grid at the conceptual level.

To identify challenging P-Grid problems that can be studied under this framework we have looked at both academic and industrial initiatives in re-designing the future P-Grid. In particular, the EPRI (Electric Power Research Institute) Intelligrid architecture[1] provides useful insights into how the future P-Grid may be used. Based on the analysis of these initiatives we have identified the following four challenging problems for this work. First, how can we secure inter-control center communications? Second, how can we protect the IT infrastructure from cyber attacks? Third, how can we manage and access large amounts of power and business related data at control centers? Fourth, how can we enable effective load management at a control center?

The rest of this paper is organized as follows. In Section 2 we discuss our qualitative framework for gauging the extent of a technology's applicability to the power grid. In Section 3 we look at data and resource security requirements where we give an overview of C-Grid technologies and then study their applicability to two P-Grid problems in depth. In Section 4 we do the same for data and resource management requirements. In Section 5 we conclude the paper and discuss future work.

2 Framework for Leveraging Computational Grid Technologies

When discussing the usefulness of a C-Grid technology/solution it is important to be able to answer the question, "how useful ?". We have developed the following simple framework that allows us to classify the usefulness at four abstract levels. These levels are broad but provide increasing level of usefulness with minimal overlaps. This framework does not include problem specification and characterization because similarity between a pair of specific problems in the C-Grid and the P-Grid is considered a prerequisite for the solution technology to be useful.

1. **Conceptual Approach.** We identify a solution/technology as being useful at the conceptual approach level if it inspires an approach for solving an identified problem in the P-Grid. The inspiration may include elements of the solution architecture, desirable properties of the solution, processes for system design, etc. Intuitively, this is the most basic level of usefulness and it is possible that the usefulness of particular technology stops at this level (while still being very helpful).

2. **System Policy.** A solution/technology is useful at this level if the policies associated with the solution are applicable to the P-Grid problem. These policies typically capture desirable system properties. For example, policies associated with (1) access control systems that limit access to authorized users, (2) system automation that specify processes requiring human intervention, (3) storage systems that specify replication and backup schedules, and (4) communication systems that specify network layer connectivity between systems. Such policies are usually defined as part of the architecture design process but since they capture system constraints they require careful consideration and consensus among the designers. Consequently, if the policies can be applied to the P-Grid problem, significant effort can be saved.

3. **Formats, Algorithms, and Protocols.** In distributed systems like the P-Grid, architected solutions typically involve components that process data. To do so, systems need to specify data formats as well as the algorithms and protocols used for processing and sharing the data. Furthermore, for interoperability these formats, algorithms, and protocols need to be standardized. This next level of usefulness applies to solutions/technologies where the specific formats, algorithms, or protocols can be used in architected solution of the P-Grid. Clearly, there has to be a great deal of similarity between the C-Grid and P-Grid solution for this to be the case. Furthermore, this is restricted today by existing legacy P-Grid protocols and standards.

4. **Implemented Tools.** This self-explanatory level identifies technologies for which implemented tools from the C-Grid can be used in the P-Grid with little or no modification. These tools can include both hardware and software systems. If a technology is useful at this level then the effort of the entire product development cycle can be saved.

3 Data and Resource Security

As P-Grid resources grow to process and share increasing amounts of data, the need for protecting the data and resources grows as well. Protecting data will maintain its sensitivity (e.g., due to reasons of intellectual property or

[1] http://www.epri.com/IntelliGrid/

customer privacy) and will require establishment of security policies and mechanisms that allow access to only authorized systems and users along with auditing capabilities. Protecting resources will ensure availability and trustworthiness of the P-Grid systems and will require establishment of security policies and mechanisms that prevent, detect and respond to intrusions and attacks.

These policies and mechanisms for protecting data and resources will have to address a variety of threats such as cyber attacks and insider misuse. At the same time, they will have to enable users access to a multitude of resources across organizational boundaries. For example, users in an EMS or a business unit at a particular control area need access to the reliability coordinator's system for delivery of power grid and market data. Scaling such accesses will require federated approaches where the organizations agree on a common authentication and authorization system for accessing data and resources. A federated authentication and authorization system, for example, provides users with single sign-on capabilities on multiple resource servers across organizational boundaries. Such a system simplifies management of identity and access permissions for users and for resource managers. The underlying technology enabling such single sign-on capability is a federated identity management system that provides users with a single identity that she can reliable reuse in a federation of trusted resource providers. Example of such federated identity management systems include Liberty Alliance (`http://www.projectliberty.org/`), Shibboleth (`http://shibboleth.internet2.edu/`), and WS-Federation (`http://schemas.xmlsoap.org/ws/2003/07/secext/`).

In this section we review some technologies for data and resource security developed by the C-Grid community. We then provide two case studies, one each for data and resource security, where problems in the P-Grid can benefit from solutions provided by the C-Grid.

3.1 C-Grid Data Security Technologies

The C-Grid is aimed at providing scientists access to distributed high performance computational and data systems for running their experiments. Characteristics of these experiments include high volumes of data that need to be shared between physical and virtual organizations, jobs that need to be run with a multitude of privileges on high performance compute systems, and long execution periods. To support the nature of such experiments the security architecture has adopted a federated approach that provides federated authentication and authorization services to users.

The core security technologies that address data security requirements in C-Grid are those that provide authentication, single sign-on, delegation, authorization, credential management, confidentiality and integrity [19, 25, 13]. Authentication is needed to verify the identity of a user or process that is requesting access to data and resources. Single sign-on is needed to allow access to multiple coordinated resources after a single authentication step. Delegation is needed to grant access rights to resources from the resource requestors. Authorization is needed to control access to data and resources based on specified access policies. Credential management allows users to engage in secure communication without the need to manage sensitive credentials. Confidentiality and integrity are needed to ensure that data remains private and cannot be modified. These technologies have been implemented with versatile components that allow compatibility with systems on differed sites [24, 2, 25]; e.g., X.509 based Public Key Infrastructures (PKI) and Kerberos for authentication and delegation, Community Authorization Service (CAS) for authorization, X.509 proxy credentials for delegation, MyProxy server for credential management, and SSL/TLS for confidentiality and integrity. These technologies have been integrated to provide comprehensive data security in several deployed Grid projects. For example, the Bridges [21] project facilitates secure biomedical data sharing with the data being located in different autonomous security domains and maintained in multiple databases.

3.2 C-Grid Resource Security Technologies

In order to provide easy access to scientists, C-Grid resources are connected to the Internet. Consequently, these resources are vulnerable to intrusions and attacks faced by systems on the Internet including Distributed Denial of Service (DDoS) attacks and worms and viruses that exploit code vulnerabilities. The threat to C-Grid resources is further compounded by the homogenous nature of these systems. The homogeneity in terms of hardware, operating systems, and applications is needed to enable fast execution of distributed experiments but also enables fast propagation of attacks through the resources.

The availability of data and computational services to scientists is crucial for C-Grid resource providers. In order to ensure availability C-Grid resource providers use a multitude of technologies that prevent, detect, and respond to attacks including firewalls, network and host-based Intrusion Detection Systems (IDSs), 24x7 staffed on-site monitoring systems, and remote configuration engines that allow rapid modification of system configurations. Comprehensive security policies and mechanisms are also being developed that address unique threats to C-Grid systems. For example, the GridSec [14] architecture uses distributed, collaborative IDSs combined with alert correlation techniques and overlay networks for efficient worm containment and pushback

of DDoS attacks. The Mithril [1] architecture uses adaptable security mechanisms that vary levels of security of the C-Grid resource provider depending on perceived threats.

3.3 Case Study I: Inter Control-center Communication

Communication between control centers in the P-Grid is becoming an important requirement; e.g., to exchange fault information for contingency analysis and emergency operations[2]. The Telecontrol Application Service Element (TASE.2) protocol (also known as Inter-Control Centre Communications Protocol, ICCP) has been developed to address this communication need. Since the data exchanged using this protocol crosses organizational boundaries, security becomes an important concern. This concern has been recognized and the IEC's Technical Committee on Power Systems information exchange suggests the use of SSL/TLS to secure such communications [6].

For TLS to be used a PKI needs to be in place with certificates and private keys available to users for initiating secure communications. Management of private keys is a challenging problem because, if stolen, they can be used by adversaries to impersonate system users. Trusting users to manage private keys is risky given the potential consequences of compromise. Therefore, solutions for managing private keys and enabling their use in protocols such as ICCP is an important security requirement for the future.

A similar security requirement exists in the C-Grid in that scientific users cannot be burdened with the management of private keys and, at the same time, compromise of these keys affects the security of C-Grid resources. To address this concern, the MyProxy credential repository has been designed and developed that manages private keys for users [2]. The repository then issues short-lived proxy credentials that the users can use for initiating secure communications. One of the most commonly used communication tools in the C-Grid is OpenSSH that provides secure SSL channels. In order to enable the use of proxy credentials a patch to OpenSSH, called GSI-SSH[3], has been developed that supports authentication with proxy credentials for establishing secure SSL connections [24]. We now give an overview of MyProxy and the proxy credential based secure communication solutions followed by an analysis of the potential usefulness of this solution for the P-Grid.

MyProxy is open source software for managing X.509 Public Key Infrastructure (PKI) security credentials (certificates and private keys) that combines an online credential repository with an online certificate authority to allow users to securely obtain credentials when and where needed [2]. Storing credentials in a MyProxy repository allows users to easily obtain proxy credentials, without worrying about managing private key and certificate files. A professionally managed MyProxy server can provide a more secure storage location for private keys than typical end-user systems. MyProxy can be configured to encrypt all private keys in the repository with user-chosen passphrases, with server-enforced policies for passphrase quality. Furthermore, the server can use Hardware Security Modules for enhanced protection of private keys [15]. By using a proxy credential delegation protocol, MyProxy allows users to obtain proxy credentials when needed without ever transferring private keys over the network.

The Grid Security Infrastructure (GSI) implements a suite of security tools for C-Grid users[4] and implements Proxy Certificates to provide authentication and delegation capabilities. These proxy credentials are X.509 based and have been standardized [23]. A proxy credential delegation protocol has been implemented that allows proxy credentials generated by MyProxy to be used for establishing authenticated connections between users and resources with OpenSSH. This delegation protocol allows single sign-on as well as execution of long running jobs. An important policy issue is the decision of the set of privileges granted to the proxy certificate. Three different authorization models are available, full rights, restricted rights, or no rights but with additional attribute assertions. All of these models have consequences with integration in communication protocols and, in practice, the full rights option is typically chosen. Collectively, the MyProxy and GSI-SSH toolkits are used at hundreds of C-Grid sites all over the world for protecting private keys and enabling secure access to resources.

Analysis. We argue that the technologies of credential management combined with authenticated secure communication tools developed by the C-Grid may be helpful to the P-Grid community at at least the first three levels, namely, conceptual approach, system policy, and format, algorithms, and protocols. At the conceptual level, these solutions will drive the design of more secure inter control center communications with effective protection of users' private keys. At the system policy level, the solutions will drive establishment of policies for protecting user credentials at the server as well as design extensions to communication protocols that allow for integration with credential repositories. At the format level, it is possible that the Proxy Certificate profile and standard may be adopted by the P-Grid. Since the compute systems of the P-Grid might be sufficiently different from those of the C-Grid it is not clear whether the implemented tools will be directly useful.

[2] http://www.intelligrid.info/IntelliGrid_Architecture/Environments/Env8_Inter-Control_Center.htm

[3] http://grid.ncsa.uiuc.edu/ssh/

[4] http://www.globus.org/toolkit/docs/4.0/security/

3.4 Case Study II: Security Incident Response Operator

The reliability of power operations is critical to the P-Grid. Ensuring reliability in a given geographic region is the responsibility of reliability coordinators. These reliability coordinators handle the challenging task of coordinating power resources managed at multiple control centers that often have a competing marketing focus. However, the same control centers must cooperate with the reliability coordinators to ensure reliability because of the connected nature of the P-Grid.

As advanced IT systems are deployed in the P-Grid for supporting both energy and business management functions, a similar inter-connected and inter-dependent IT infrastructure will emerge. The infrastructure will comprise, for example, high-bandwidth networks, high performance compute and storage systems, and large numbers of field devices. Ensuring the security of this infrastructure will become paramount to the security of the P-Grid. Since the infrastructure spans multiple autonomous organizations the security process needs to be coordinated by an entity such as an reliability coordinator but implemented using a distributed approach that encompasses all the involved organizations. The security coordinator, or operator, would need to establish policies for intrusion prevention and monitoring at each organization in its jurisdiction and define a procedure for sharing incident data as well as for responding to incidents. Though first steps towards cyber security have been undertaken by the North America Electric Reliability Council (NERC) via establishment of basic cyber security guidelines[5], the establishment of an effective security and incident response process will become an important challenge for the future power grid.

A similar need exists in the C-Grid where a process is needed to ensure protection of resources at various interconnected sites; e.g., at the TeraGrid[6]. The security process as briefly outlined by the TeraGrid Security Working Group[7] involves several steps. First, the organizations need to define an organizational security policy for the use and protection of their resources. This is an extensive undertaking that involves the participation of and conformance by almost all employees of the organization. An example of such a security policy is that of the National Center for Supercomputing Applications[8] (NCSA).

Second, the organizations need to deploy a system for monitoring the network for intrusions. Effective monitoring requires a 24x7 staffed operation center that is supported by advanced Intrusion Detection Systems (IDSs). Several organizations including NCSA have such a security operation center. IDSs are a crucial technology for detecting incidents but face the challenge of reducing the number of generated alerts and coping with high bandwidth networks. Addressing these challenges requires a trained staff that can identify interesting alerts and experimentation with advanced hardware based IDSs that can cope with gigabit networks; e.g., the Oak Ridge National Laboratory is deploying the Force10 P-Series IDS for 10 GBPS network monitoring[9].

Third, organizations need to set up mechanisms for forensic analysis and for responding to successful intrusions. Forensic analysis at C-Grid sites typically begins with an IDS alert followed by a detailed analysis of various logs that record the intruders path. To enable this analysis extensive logging is undertaken along with mechanisms such as distributed replication of logs to prevent modification of logs by the intruder. Upon discovery of successful intrusions a response is undertaken with mechanisms that include capabilities for remote modification of firewall rules and system configurations, killing processes and disconnecting machines, etc. This is followed by a "clean up" of the attacked system and restoration of accesses. Currently, this process is only partly automated because fully automated responses can affect legitimate users and processes based on false IDS alerts. However, efforts are underway to automate response to certain types of intrusions such as worm attacks as well as raising of site security levels that allow availability of services even while the attack is in progress [5, 1].

Fourth, and perhaps the most challenging, TeraGrid needs a communication process that enables sites to share incident data and to coordinate a response. Organizations are typically reluctant to share such data for business reasons. TeraGrid has initiated and successfully established such a communication process that allows system administrators at various sites to trust each other and share incident data to protect TeraGrid as a whole. A consensus based memorandum of understanding forms the basis of this communication, which has averted several intrusions in recent times.

Analysis. We argue that this security process and associated technologies may be helpful to the P-Grid at the conceptual and policy levels. At the conceptual level, it drives the need for establishing such a security process as well as an approach for doing so. At the policy level, it may help the establishment of policies for sharing incident data. Note that sharing of incident data cannot be mandated as the failure to do so by a member organization is difficult to identify. Therefore, a community based approach is more appropriate.

[5] http://www.nerc.com/~filez/standards/Cyber_Sec_Renewal.html

[6] http://www.teragrid.org/

[7] http://security.teragrid.org/

[8] http://www.ncsa.uiuc.edu/UserInfo/Security/policy/

[9] http://www.gridtoday.com/grid/629727.html

4 Data and Resource Management

The P-Grid of the future will have significantly greater data and resource management requirements driven by a large number of data sources and resource-consuming applications that process, store, and share data. The data will be generated by an increasing number of field devices such as IEDs (Intelligent Electronic Devices), automated stations that pre-process the data, integrated EMS and BMS systems at control centers, and inter control center communications. Even today IEDs generate a lot of information that is stored in substations but not sent to control centers due to bandwidth limitations and proprietary non-compatible data formats, and control centers archive a lot of historical data that is available but rarely used.

Many applications are being developed and envisioned that will make use of this data and will require resources for execution. For example, (1) using historical archives to fulfill the business needs such as real time pricing instead of flat rate and operational needs such as proactive maintenance or identifying the weak points in the power grid [16], (2) advanced accounting and billing that allow consumers to charge their electric cars at a friend's house but pay for that charge themselves, (3) business functions integrated inside SCADA and IEDs, (4) temporal analysis of time-stamped SCADA data [27], (5) remote load management[10], and (6) smart meters [10]. Efficient execution of these applications will require efficient management of resources that run the applications.

Currently P-Grid lacks tools and techniques to support the management and efficient use of these data and resources. In this section we review some technologies for data and resource management developed by the C-Grid community. We then provide two case studies, one each for data and resource management, where problems in the P-Grid can benefit from solutions provided by the C-Grid.

4.1 C-Grid Data Management Technologies

In C-Grid data requirements are of the order of petabytes spread over millions of files and database records in multiple physical locations. Furthermore, data is often stored in different formats raising the requirement of presenting the data from heterogeneous sources as a single virtual data source. To access and manage this huge volume of data a number of techniques have been developed. Some of the core technologies that address data management in C-Grid are data virtualization, data federation and integration, and data location services [8, 12, 22, 20, 3, 11]. Here, metadata is often used to describe data and metadata services are used to locate data based on descriptive attributes stored in the metadata.

Many of these data management techniques have been implemented in tools and the tools have been integrated into several large Grid applications. The Grid Data Mediation Service (GDMS) [26] is a prototype implementation of a Grid service that can present distributed, heterogeneous data sources as a one logical virtual data source. The Metadata Catalog Service (MCS) [20] provides a scalable solution for storing and accessing descriptive metadata. The GDMS toolkit is used by the GridMiner application [4] and MCS is used by the Earth System Grid[11] application. In many cases metadata catalogs themselves are often distributed and need to be managed. To address this issue, the Artemis system [22] has been developed to integrate distributed metadata catalogs on the C-Grid.

4.2 C-Grid Resource Management Technologies

In CGrid resources are distributed over multiple administrative domains, and often consist of hundreds of machines at a single site. Some examples of Grid resources are computing elements (e.g., processors), storage spaces (e.g., hard disks), application data, network connections, and software components [18]. As Grid resources are shared by multiple users and multiple applications at the same time, running these applications efficiently requires efficient Grid Resource Management Systems (RMSs). Resource Discovery, Access to Resource Information, Status Monitoring, Brokering/Scheduling, Reservation management, Execution Management/Provisioning are some of the important services required for managing Grid resources as identified by the Grid Scheduling Architecture Research Group (GSA-RG) of the Global Grid Forum[12].

It is often the case that a particular job is executed on multiple Grid resources and needs to access data from multiple data sources. Job Scheduling involves decision-making regarding execution orders of multiple such jobs and uses information about job requirements in terms of processing cycles, memory requirements, storage space requirements, application data requirements and deadlines. To satisfy these requirements, the task scheduler often needs to negotiate with other Grid sites and locate resources (Resource discovery) and then reserve resources (Resource reservation) before starting execution. While executing a task, the status of ongoing jobs needs to be monitored (Status monitoring and Execution management) for potential problems and re-scheduling them if needed. One of the

[10] http://intelligrid.info/IntelliGrid_Architecture/Use_Cases/CS_RTP_Overview_Use_Cases.htm

[11] https://www.earthsystemgrid.org/

[12] http://forge/gridforum.org/projects/gsa-rg/

tricky issues in task scheduling is handling of irregular and time dependent performance of processors. In [28], they propose a conservative scheduling algorithm that can handle this issue efficiently.

Several toolkits have been implemented that provide Grid RMSs. The Resource Specification Language (RSL) [7] is used for specifying resource requirements and engaging in resource negotiation. The Condor-G [9] system is capable of using multi-domain resources transparently and can perform resource discovery and resource management. The Globus GRAM (Grid Resource Allocation Manager) [7] toolkit allows remote job management by providing an API for submitting, monitoring, and terminating jobs.

4.3 Case Study III: EMS Data Storage and Processing

The need for managing and sharing large amount of data from heterogeneous resources at EMSs is clearly going to increase with time. Several factors contributing to this increase have been discussed throughout Section 4. Even today, the absence of this capability is hurting the P-Grid. For example, the lack of overall system views outside one's control area and poor communication of information often contributes to system failures [16]. The ability to communicate and share information with other parties (e.g., other control centers, RTUs, market participants, large customers and suppliers, control center service providers) is required due to changes in business models and for achieving system wide reliability [27]. However, the presence of proprietary data formats as well as proprietary vendor-specific hardware and software make this communication challenging [17].

Similar needs exist in the C-Grid community for data management where it is often the case that data accessed by C-Grid applications involves multiple databases, crosses administrative boundaries and is stored in different formats. To integrate data from different sources and to provide a single, logical view GDMS [26] has been developed. GDMS uses a mapping schema that defines the mapping information between virtual data sources and participating data sources. This schema stores the actual location of a data source, the format of the data source (e.g., CSV file, MySQL), and how the join operation will be performed between different types of data sources. Using this schema the query is reformulated as required by different data sources and then executed by a query optimizer, which can choose between alternative replicas depending on availability and performance needs. The GDMS service has been successfully used in the GridMiner project.

To manage large data sets efficiently, MCS [20] has been developed. MCS contains metadata information about the data that need to be managed and this metadata service is used for data discovery and access in the Grid. As described in [20], MCS is first consulted to find data sets containing particular attribute values to which MCS responds with a list of logical name attributes for data items that match the attributes. Next, the Replica Location Service is used to get the physical locations of the data for those logical list of names. Finally, the selected replica is queried for data and the result is returned using the Grid FTP protocol. MCS has been successfully used in Earth System Grid Application by scientists to discover and query data files based on metadata attributes.

Analysis. We argue that the technologies for data management developed by the C-Grid may be helpful to the P-Grid community at at least the conceptual and policy levels. At the conceptual level, these solutions will drive the design of more efficient data management, integration, and access services in P-Grid. At the policy level, these solutions can help define the various services that would be needed. Depending on the kinds of data needed in the future, the data formats and query languages may be helpful as well.

4.4 Case Study IV: Load Management

A challenging requirement for resource management that will likely arise in the future but perhaps has not yet been fully envisioned yet is related to load management. With increasing demands for power, various load management techniques have been proposed to combat peak load demand and to avoid power failure. One such proposed scheme is load management[13] where control centers would be able to remotely control power delivered to specific appliances at customer sites. To manage load power would be cutoff during peak hours and overload conditions but this would be done at the fine granularity of individual appliances.

Remotely operating electrical appliances is not a new concept. European utilities have used "ripple control" for this purpose where they inject low frequency signals into distribution feeders to disconnect water heaters at customer homes during peak hours. In contrast, Intelligrid[13] suggests the use of special hardware that is installed alongside appliances to enable remote control and management. Extending this simple scenario to one where multiple appliances in every commercial and residential building are available for load management brings us closer to the vision of "smart meters" [10] where two-way networked communications between buildings and control centers would enable remote load management. In this extended scenario the problem becomes one of scale where the control center needs to decide the set of appliances in its area to which power needs

[13] http://intelligrid.info/IntelliGrid_Architecture/Use_Cases/CS_RTP_Overview_Use_Cases.htm

to be delivered at any given point in time based on available power, the demand, and user preferences.

If we consider power to be a resource and delivering power to an appliance (e.g., thermostat) to be a task that is the consumer of the resource, then the problem looks very similar to task scheduling in C-Grid. That is, just like the C-Grid, there is a need to schedule a lot of tasks based on available resources. Here the resource (e.g., power) is a variable time dependent quantity and its demand (and perhaps price) is higher during peak hours. So the tasks need to be scheduled intelligently with scheduling algorithms that anticipate power availability and load requirements and manage the delivery of power keeping customer preferences in mind. There might be a tighter integration with markets for real-time pricing that makes this scheduling even more challenging.

In C-Grid, first the job description is provided in RSL, which allows for specification of resource requirements. Then, the Grid Resource Registration Protocol (GRRP) is used by a resource to notify its existence on the Grid and the Grid Resource Information Protocol (GRIP) is used to obtain information about resource status. A schedule algorithm then decides which jobs need to be run on which devices. The performance of processors (a primary resource) is irregular (due to heterogeneity) and time varying (due to unpredictability of instantaneous load). To schedule jobs more efficiently on processors, [28] proposes a conservative scheduling policy that schedules based on estimation of future variance in resource capabilities. They use a time series predictor for predicting CPU load at a future point in time and to calculate CPU load and variation for a future time interval. They propose different scheduling algorithms based on the above estimations and show the effectiveness of each technique by executing benchmark applications. Similar scenarios where EMS applications have access to resources (available power) and tasks (power requirements and user preferences for appliances), one can design appropriate scheduling algorithms for effective load management with expected loads being calculated based on historical data.

Analysis. We argue that the technologies for load management developed by the C-Grid may be helpful to the P-Grid community at conceptual approach. The concept of RSL can be used to describe electrical appliances as well as customer preferences. This information can be conveyed to control centers over communication channels established via smart meters. Load managers can then decide when to turn off a device or turn on a device based on policies and scheduling algorithms.

5 Conclusions and Future Work

In this work we have used technologies developed by the Computational Grid community to address challenging problems in the Power Grid dealing with security and management of data and resources. We suggest (1) the use of credential management services integrated with secure communication protocols for greater security of inter-control center communications, (2) the establishment of a community-based process for preventing, detecting, coordinating, and responding to intrusions in an effort to protect the P-Grid infrastructure, (3) the use of metadata catalog and data replication services for effective data management at control centers, and (4) the use of resource discovery and scheduling services for effective resource management at control centers. For each of these solutions we identify a qualitative level of the potential applicability of C-Grid technologies to the P-Grid problems.

In the future we will take an architectural approach where we place these C-Grid technologies in an appropriate P-Grid architecture so that we can study their integration and collective effectiveness. This approach will give an alternative view of the usefulness of C-Grid technologies. The end goal of the current approach and the planned one for the future is to be able find technologies than can be quickly deployed in the P-Grid because they have have already been deployed and tested in large systems like the P-Grid.

6 Acknowledgments

We would like to thank the entire Trustworthy Cyber Infrastructure for Power Grid project team for various discussions on security requirements, protocols and architectures for the Power Grid. We would like to thank the minitrack chair and anonymous reviewers for their comments and suggestions. This material is based upon work supported by National Science Foundation under Grant No. CNS-0524695. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] J. Basney, P. Flanigan, J. Heo, H. Khurana, J. Muggli, M. Pant, A. Slagell, and V. Welch. Mithril: Adaptable Security for Survivability in Collaborative Computing Sites. In *Workshop on Enterprise Network Security (WENS), part of SecureComm*, 2006.

[2] J. Basney, M. Humphrey, and V. Welch. The MyProxy Online Credential Repository. *Software: Practice and Experience*, 35(9):801–816, July 2005.

[3] S. Bourbonnais, S. Malaika, V. Gogate, I. Narang, L. Haas, V. Raman, and R. Horman. Towards an information infrastructure for the grid. *IBM Systems Journal*, 43(4):665–688, 2004.

[4] P. Brezany, J. Hofer, A. Tjoa, and A. Wohrer. GridMiner: An Infrastructure for Data Mining on Computational Grids Peter Brezany. In *APAC'03 Conference*, Gold Coast,Australia, October 2003.

[5] M. Cai, K. Hwang, Y.-K. Kwok, S. Song, and Y. Chen. Collaborative Internet Worm Containment. *IEEE Security and Privacy*, 3(5), May/June 2005.

[6] F. Cleveland. IEC TC57 Security Standards for the Power Systems Information Infrastructure Beyond Simple Encryption. Technical report, Xanthus Consulting International, October 2005.

[7] K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith, and S. Tuecke. A Resource Management Architecture for Metacomputing Systems. In *Proceedings of the IPPS/SPDP'98 Workshop on Job Scheduling Strategies for Parallel Processing*, 1998.

[8] I. Foster, J. Vockler, M. Wilde, and Y. Zhao. The Virtual Data Grid: A New Model and Architecture for Data-Intensive Collaboration. In *Proceedings of the 2003 CIDR Conference*, 2003.

[9] J. Frey, T. Tannenbaum, I. Foster, M. Livny, and S. Tuecke. Condor-G: A Computation Management Agent for Multi-Institutional Grids. In *Cluster Computing*, pages 237–246, 2002.

[10] Grid 2030: A National Vision for Electricity's Second 100 Years. Technical report, US Department of Energy, Office of Electric Transmission and Distribution, July 2003.

[11] L. Guy, P. Kunszt, E. Laure, H. Stockinger, and K. Stockinger. Replica Management in Data Grids. Technical report, GGF5 Working Draft, July 2002., 2002.

[12] L. Haas, E. Lin, and M. Roth. Data Integration through database federation. *IBM Systems Journal*, 41(4):578–596, 2002.

[13] M. Humphrey, M. Thompson, and K. Jackson. Security for Grids. *Proceedings of the IEEE (Special Issue on Grid Computing)*, 93(3):644–652, March 2005.

[14] K. Hwang, Y.-K. Kwok, S. Song, M. C. Y. Chen, Y. Chen, R. Zhou, and X. Lou. GridSec: Trusted Grid Computing with Security Binding and Self-defense Against Network Worms and DDoS Attacks. *ICCS*, pages 187–195, 2005.

[15] M. Lorch, J. Basney, and D. Kafura. A Hardware-secured Credential Repository for Grid PKIs. In *Proceedings of the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid2004)*, April 2004.

[16] F. Maghsoodlou, R. Masiello, and T. Ray. Energy Management Systems. In *IEEE Power and Energy Magazine*, September/October 2004.

[17] L. Murphy and F. F. Wu. An Open Design Approach for Distributed Energy Management Systems. *IEEE Transactions on Power Systems*, 8(3), August 1993.

[18] U. Schwiegelshohn and R. Yahyapour. Resource Management for Future Generation Grids. CoreGrid Technical Report TR-0005, Computer Engineering Institute, University of Dortmund, Germany, May 2005.

[19] F. Siebenlist, N. Nagaratnam, V. Welch, and C. Neuman. Security for virtual organizations - federating trust and policy domains. In *The Grid: Blueprint for a New Computing Infrastructure*, pages 353–387, San Francisco,Elsevier, 2004.

[20] G. Singh, S. Bharathi, A. Chervenak, E. Deelman, C. Kesselman, M. Manohar, S. Patil, and L. Pearlman. A Metadata Catalog Service for Data Intensive Applications. In *SC '03: Proceedings of the 2003 ACM/IEEE conference on Supercomputing*, Washington, DC, USA, 2003.

[21] R. Sinnott, M. Atkinson, M. Bayer, D. Berry, A. Dominiczak, M. Ferrier, D. Gilbert, N. Hanlon, D. Houghton, E. Hunt, and D. White. Grid Services Supporting the Usage of Secure Federated, Distributed Biomedical Data. In *Proceedings of the UK e-Science All Hands Meeting*, 2004.

[22] R. Tuchinda, S. Thakkar, Y. Gil, and E. Deelman. Artemis: Integrating Scientific Data on the Grid. In *Proceedings of the Sixteenth Innovative Applications of Artificial Intelligence*, San Jose, CA, July 2004. American Association for Artificial Intelligence.

[23] S. Tuecke, V. Welch, D. Engert, L. Pearlman, and M. Thompson. Internet X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile. IETF Network Working Group, Requests for Comments, RFC 3820, standards track, June 2004.

[24] V. Welch, I. Foster, C. Kesselman, O. Mulmo, L. Pearlman, S. Tuecke, J. Gawor, S. Meder, , and F. Siebenlist. X.509 proxy certificates for dynamic delegation. In *In 3rd Annual PKI R&D Workshop*, April 2004.

[25] V. Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski, J. Gawor, C. Kesselman, S. Meder, L. Pearlman, and S. Tuecke. Security for Grid Services. In *Twelfth International Symposium on High Performance Distributed Computing (HPDC-12)*. IEEE Press, June 2003.

[26] A. Woehrer, P. Brezany, and I. Janciak. Virtualization of Heterogeneous Data Sources for Grid Information Systems. In *MIPRO*, 2004.

[27] F. F. Wu, K. Moslehi, and A. Bose. Power system control centers: Past, Present and Future. *Proceedings of the IEEE*, Vol. 93(11), November 2005.

[28] L. Yang, J. M. Schopf, and I. Foster. Conservative Scheduling: Using Predicted Variance to Improve Scheduling Decisions in Dynamic Environments. In *SC'03*, Phoenix, Arizona, USA, November 2003. ACM.