Grid Computing for the Masses: An Overview

Kaizar Amin,[1,2] Gregor von Laszewski,[1,*] and Armin R. Mikler[2]
[1]Argonne National Laboratory, Argonne, IL, U.S.A.
[2]University of North Texas, Denton, TX, U.S.A.
[*] Corresponding Author: gregor@mcs.anl.gov

Contents

Grid Computing for the Masses: An Overview

Kaizar Amin,[1,2] Gregor von Laszewski,[1,*] and Armin R. Mikler[2]
[1]Argonne National Laboratory, Argonne, IL, U.S.A.
[2]University of North Texas, Denton, TX, U.S.A.
* Corresponding Author: gregor@mcs.anl.gov

Abstract

The common goals of the Grid and peer-to-peer communities have brought them in close proximity. Both the technologies overlay a collaborative resource-sharing infrastructure on existing (public) networks. In realizing this shared goal, however, they concentrate on significantly contrasting issues. The Grid paradigm focuses on performance, control, security, specialization, and standardization. On the other hand, the peer-to-peer paradigm concentrates on fault tolerance, resilience, decentralization, and peer cooperation. In this paper, we discuss Grid usage models including traditional Grids, ad hoc Grids, and federated Grids. We compare these approaches to peer-to-peer computing and discuss the issues involved in the convergence of the two paradigms.
Keywords: *Peer-to-Peer Grids, Ad-hoc Grids, Federated Grids, Public Grids, Pervasive Grids*

1 Introduction

The term "Grid computing" [1] is commonly used to refer to a distributed infrastructure that promotes large-scale resource sharing in a dynamic multi-institutional "virtual organization" (VO). A computational Grid is conceptually based on the principles of an electric power Grid. A large number of electric power generating plants interconnect with one another, providing standardized, reliable, cheap, and ubiquitous access to electric power. Similarly, a computational Grid forms a closed network of a large number of pooled resources providing standardized, reliable, specialized, and pervasive access to high-end computational resources.

Typically, in order to establish a computational Grid, several institutions pool their resources such as computational cycles, specialized software, database servers, network bandwidth, and people. Thereafter, global policies for the VO are established that identify the role and responsibilities of participating entities. Well-trained professional administrators associated with the participating institutions enforce the global VO and local domain policies. Based on these policies, the Grid administrators provide security credentials to the Grid users, who can access the distributed Grid resources within the scope of their credentials irrespective of their geographical positions.

Several applications and infrastructures have been proposed in the literature that can significantly benefit from the Grid concept [1]. Applications that can leverage from more than one supercomputer can benefit from a distributed supercomputer created as a Grid by pooling several supercomputers. Applications containing "pleasantly parallel" subtasks can take advantage of the Grid to co-allocate a large number of distributed compute resources in parallel [2, 3]. Data-intensive applications can use specialized data stores and replica systems available in the Grid to store and retrieve large number of datasets. Advanced collaborative applications [4] can use the interactive feature of the Grid to provide an enhanced human-to-human interaction.

Irrespective of the application, these Grid infrastructures have some common characteristics [5].

A computational Grid is "collaborative". It comprises heterogeneous resources that are managed by more than one entity. Such a distributed pooling of resources within the same institution or across multiple institutions requires significant collaboration among the participating entities. The Grid architecture respects the institutional policies of its collaborators by giving such policies preference over the global Grid policy. Such an enforcement not only allows collaborating institutions to protect intellectual property but also provides enough flexibility to allow other institutions to participate within the Grid.

A computational Grid provides non-trivial "quality of service" (QoS) assurances. High connectivity is maintained between resources via dedicated high-speed networks. Further, the Grid services themselves offer advanced high-level functionality that enables sophisticated science and commerce. A well-established Grid administration by dedicated human-resources facilitates constant connectivity, monitoring, and fault tolerance in a Grid.

The Grid architecture is "standardized". In the early days of the Grid, the lack of standards for Grid service development resulted in several non inter-operable Grid middleware frameworks [6, 7, 8, 2]. Recently, however, the Global Grid Forum (GGF) [9] has assumed the responsibilities for coordinating the standardization of Grid developments. The Open Grid Services Architecture (OGSA) [10] initiative of the GGF defines the artifacts for a standard service-oriented Grid framework based on the emerging W3C-defined Web services technology [11]. A service-oriented Grid framework provides a loosely coupled technology- and platform-independent integration environment allowing different vendors to offer Grid-enabled services in a variety of technologies, yet conforming to the GGF-defined OGSA standards and thus making them inter-operable.

2 Grid Usage Models

In this section we discuss the usage model attributed to contemporary Grid frameworks. We also explore other usage patterns that can enhance the benefits of the Grid to a larger community.

2.1 Traditional Grids

The popularity of the Grid architecture is evident from the large number of advanced scientific and commercial projects that are deploying the Grid framework. Some success stories include the DOE Science Grid [12], the European Union DataGrid [13], the Grid Physics Network (GriPhyN) [14], the Information Power Grid [15], the National Fusion Grid [16], the National Research Grid Initiative [17], the Network of Earthquake Engineering Simulation (NEES) Grid [18], the Particle Physics Data Grid [19], and the TeraGrid [20].

A prominent characteristic of each of these Grid applications is that it constitutes a "closed" network of resources. For example, the NASA Information Power Grid is available only to engineers and scientists employed by NASA for official business and research. Similarly, GriPhyN is open only to experimental physicists. In other words, the current Grid usage model caters to the needs of certain "classes". Unless one belongs to a research or commercial organization, it is quite difficult to get access to one of the high-performance Grid infrastructures. Further, the administrative overhead involved in the initial Grid setup makes it non-trivial for an individual or small organization to establish personal Grids at will. Proponents of the current usage model justifiably argue that the philosophy of these collaborative Grids is to support the computational and data-intensive needs of an elite few, rather than providing Grid access to the "masses". The tightly controlled administrative mechanism enables Grid service providers to offer the promised QoS guarantees.

Further, contemporary Grids have a highly segregated role-based usage [1]. An individual interacting with the Grid can be conveniently categorized as one of the following: service provider, service developer, administrator, or end-user [21]. Although there can be some degree of overlap between associated roles, an extreme mixture of these roles in explicitly avoided. Such a well-defined role-based interaction facilitates a separation of concerns, allowing the end-user to concentrate on science or commerce, enabling the service provider to concentrate on QoS assurances, the service developer to concentrate on Grid protocols, and the administrator to focus on enforcement of access control and organizational policies.

2.2 Ad hoc Grids

Even though a tightly controlled Grid framework provides the required QoS parameters, it is highly restrictive in expanding its usage beyond the traditionally proposed model. As discussed in previously, it is non-trivial for individuals not belonging to advanced scientific, academic, or commercial institutions with a Grid-vision to collaborate with fellow peers at random in a Grid environment. In other words, the current usage model does not facilitate "ad hoc" Grid establishment [22]. Advocates of the conventional Grid architecture may reasonably argue that, similar to the power Grid, the computational Grid provides persistent and reliable service to its users. However, a large number of scenarios require transient, short-lived collaboration that needs to be supported by the Grid [23].

For example, consider the following case. A group of geographically separated scientists require ad hoc short-term collaboration and resource sharing in a secure environment to evaluate different experimental simulations of a thermochemistry application. One scientist contributes the simulation service, one pools a visualization service to render the results of the simulated experiment, another scientist provides the data repository storing the input datasets for the experiment, and a few others want to interactively discuss the final results in an educational setting. Although simple, this example is representative of a large class of collaborative applications developed as a part of multi-domain sciences.

To implement such an application using the current usage model, the scientists would need to formally establish a Grid virtual organization (VO) defining appropriate use policy and describing individual contributions and responsibilities. The VO would then assign administrative privileges to a dedicated entity, who would then create Grid credentials for every other entity within the VO. All of the participating entities would need to support the appropriate Grid middleware [24] and expose their services as a part of this middleware. Once the administrative functions were performed, each user would be able to interact with the Grid within the context of his own rights. Although this setup provides the required functionality, the administrative overhead to establish such a short-term community (possibly one-time collaboration) surpasses its utility. Clearly, there is a need for a set of Grid tools to realize such an ad hoc Grid, thereby increasing the user base of the Grid community.

2.3 Federated Grids

Motivated by the success of volunteer computing architectures such as seti@home [25], distributed.net [26], and grid.org [27], we further extend the concept of ad hoc Grids to form "federated Grids". A federated Grid is a generic Grid architecture where a resource contribution is not limited to active collaborators alone. Individuals can pool their resources (idle computational cycles) to enhance the computational power of the Grid.

One of the greatest problems with the current volunteer computing model is the lack of motivation to encourage the "masses" to contribute their resources. Success of some initial applications can be attributed to the "cool" factor of the underlying science. However, without providing sufficient incentive to the resource contributor, this does not constitute a viable economic model. Further, these architectures represent a master/slave paradigm. A single privileged master has the authority to attribute independent tasks to a large number of slaves. The slaves do not have the appropriate authorization to submit their compute-intensive tasks to the pool of available resources.

In a federated Grid framework, we can overlay an economic model on the Grid, providing sufficient incentive to the resource providers to contribute their resources. Further, we shift from the master/slave model of volunteer computing to a more flexible model allowing resource providers to also become resource consumers. For example, commercial institutions can contribute their idle resources to a federated Grid in return for computational power on demand at a later time. This paradigm exposes the Grid to an important research domain of usage economics and brokering [28].

A similar ideology has been adopted in electrical power Grids with great success [29, 30]. Sophisticated solar and wind electricity generation devices are installed in homes and other establishments. Advanced control mechanisms feed all excess electricity generated by these devices into the power Grid. When the system does not produce enough power, electricity is extracted from the Grid. The users pay only for the net electricity used by them. Such an economic use-case serves as a model for computational Grids too.

On a larger scale, one can envision a universal version of federated Grids as ubiquitous "public Grids" or "pervasive Grids". The concept of public Grids in computing is analogous in principle to the vision of the Internet in information science. However, the degree of complexity in public Grids is extremely high when compared with that of the contemporary Internet. Users could access the public Grids ubiquitously as service providers or consumers or both. Further, these users could dynamically establish virtual organizations in an ad hoc fashion, thereby overlapping the functionalities of traditional as well as ad hoc Grids. Based on their personal credentials and preferences, peers could participate in these ad hoc transient groups enabling collaborative resource sharing.

3 Grid and Peer-to-Peer Computing

The term peer-to-peer (p2p) computing is used to refer to an ad hoc, dynamic, unstable, and self-organizing distributed model that assists in collaborative community formation and resource sharing at the edge of the network. In a p2p environment, there is no distinction between a resource user (client) and a resource provider (server); each is commonly referred to as a "peer".

Several aspects of ad hoc Grids and federated Grids have already been showcased within the p2p community. Ad hoc collaborative file-sharing applications such as Napster [31] and Gnutella [32] have been successfully implemented and widely used by a large p2p community. However, ad hoc Grids focus on issues that go beyond file-sharing mechanisms and addresses aspects such as advanced security, trust and reputation management, and quality-of-service assurances. As discussed in Section 2.3, several volunteer computing frameworks such as seti@home and distributed.net have demonstrated the usefulness of distributed resource pooling using p2p technologies. However, federated Grids address a much larger problem of converting a master/slave model to a pure p2p model.

It is evident that the Grid paradigm and the p2p paradigm both have the same end goals: collaborative community formation and shared access of distributed resources. Despite of the similarity in philosophy, there some fundamental differences between the two technologies [33, 34].

3.1 Organization

The Grid computing model is a client-server model where the Grid servers offer specialized, reliable, highly advanced, and sophisticated scientific and commercial applications. Grids require a pre-established administrative infrastructure enforcing the VO policy. In other words, the roles, responsibilities, and privileges of the collaborating institutions and users are pre-defined in a Grid environment. These responsibilities and privileges do not change frequently and are maintained by well-trained administrators.

On the other hand, p2p paradigm provides direct communication between peers without warranting any pre-established management infrastructure. The responsibility and privileges of the participating entities are not defined *a priori* and are often in flux. Every peer is responsible for defining and maintaining the access policies for his resource within the community.

3.2 Security

Grids are centrally controlled by dedicated administrators enforcing centralized security policies. The trust level between participating entities is high, hence alleviating the requirement of complex reputation and trust enforcement models.

Popular peer-to-peer applications such as Napster, seti@home, and Gnutella lack the concept of a predefined trust relationship between participating entities [35]. The absence of a centralized policy enforcement architecture warrants p2p applications to deploy advanced distributed trust and reputation management services. Most p2p applications used by the masses assume an unsecured environment at all times. Peers cannot trust fellow peers due to the lack of accountability for wrongdoings.

3.3 Scalability

The Grid paradigm concentrates on providing a suite of advanced services to a moderate number of privileged users. Its goal is to provide high quality of service to a small community, rather than providing scalability assurances to a large group. Moreover, the VO membership in a Grid requires significant administrative overhead, which makes it difficult to manage a large user-base and thus further restricts the scalability of the Grid.

The p2p architecture focuses on integrating simple resources for the masses. Popular p2p applications such as Napster and seti@home have reached a user-base of millions [35]. The absence of any centrally controlled administrative infrastructure and the distributed nature of resource utilization make p2p applications extremely scalable.

3.4 Quality of Service

Like most client-server models, Grid services are hosted on specialized “high-end” resources including expensive scientific instruments, clusters, and data storage systems. High connectivity is maintained between resources via dedicated high-speed networks. A well-established resource administration facilitates constant resource connectivity, resource monitoring, and fault tolerance. The required quality of service (QoS) is provided by the committed members of the VO based on their pre-agreed upon Grid policy and their dedication in the overall collaboration.

Peer-to-peer systems cannot guarantee any QoS. They provide their services on a best-effort basis. Ongoing efforts have been made in the p2p community to improve the QoS by deploying dedicated peers (rendezvous peers) to improve resource connectivity and service discovery. However, such peers cannot be guaranteed and any QoS assertion is entirely dependent on the organization of connected peers.

4 Challenges

The concept of ad hoc Grids and federated Grids is intriguing and the possibilities of its applications are many. Nonetheless, several technical challenges need to be addressed within the Grid community before such a paradigm can be adopted.

The biggest challenge in p2p Grids is the enforcement of a dynamic and adaptive security model that overlays a secure framework over an insecure network. In a Grid approach the notion of trust and reputation in implicit within a closed VO formation. As discussed earlier, every collaborating entity trusts other entities. Grid resources are not anonymous; they are accountable for any misconduct. Further, Grid resources adhere to well-established mechanisms for authorization and authentication, restricting the use of resources by non collaborating or rogue users. Several security enforcement and policy maintenance models have been proposed within the Grid community [36, 37, 38, 39]. However, all these architectures are based on the assumption of a stable, persistent, and long-term Grid establishment with a small set of seldom-changing users. Hence, these frameworks cannot be applied "out-of-the-box" to the proposed p2p Grid framework. Ad hoc and federated Grids require an adaptive security model that incrementally builds a secure Grid community based on the notion of trust and reputation. Similar security models are being investigated by the p2p research community [40, 35]; however, their application in the Grid domain needs to be studied.

Providing the adaptive authentication and authorization model described above does not guarantee the safety of Grid resources against spying, sabotage, and destruction. This issue has been addressed in detail by the mobile agent community in the context of protecting a mobile agent from a malicious host and vice versa. However, it needs to be verified whether these principles can be easily extrapolated from the agent context to the Grid context.

Grid computing specializes in providing different levels of QoS guarantees. Without such assurances a p2p Grid infrastructure will lose its utility as a Grid environment. Hence, it is imperative to overlay an advanced QoS framework on p2p Grids. A sophisticated QoS environment that delivers the promised assurances despite unreliable, dynamic, and insecure Grid resources is yet to be researched. Peer-to-peer technologies assume the occurrences of system failure and service unavailability and provide mechanisms to adapt to such occurrences. However, the impact of these events on the performance of Grid applications and the effects of their fault-tolerance measures in a Grid context need to be thoroughly investigated.

Many computational economic models have been proposed in literature by the Grid and p2p communities [28, 35]. Several of these models could be used in the p2p Grid framework for fair sharing of resources. However, their feasibility and utility in the context of a strict yet dynamic security and QoS requirements need to be analyzed.

Further, in the absence of a centrally controlled administrative framework, policing and monitoring these self-adaptive and self-configuring Grids is a formidable task. Mechanisms must be established that allows dynamic monitoring of such p2p Grid frameworks.

Some researchers believe that social acceptance of such federated Grids is far from reality [41]. They argue that irrespective of the security, economic, and QoS assurances given by the Grid community, it would be extremely unlikely for research and commercial agencies to execute critical applications in such an unpredictable environment. Although the initial success of volunteer computing infrastructures in low-risk research applications seems encouraging, its wider sociological acceptance by a much larger community needs to be seen.

5 Conclusion

Irrespective of the inherent differences, both the Grid and p2p paradigm have some distinct characteristics that complement each other. Grids enable sharing of specialized and advanced services in a secure environment providing high quality of services. The p2p environment offers dynamic, self-organizing, and self-configuring ad hoc collaboratories sharing a large number of resources at the edge of the network. A large number of applications can inherently benefit by combining the characteristics of both these paradigms. Indeed, the amalgamation of the two technologies, resulting a new computing paradigm, the "peer-to-peer Grid", will provide the necessary standards, security, and QoS assertions of the Grid paradigm and the ad hoc, self-organizing, and self-configuring attributes of the p2p architecture. This new paradigm can be further enhanced to form ubiquitously available universal public Grids that can provide a computing infrastructure similar to the Internet.

Before p2p Grids can become a reality, a large number of issues need to be addressed by the Grid and p2p communities. Issues dealing with decentralized resource management, dynamic security policies, adaptive trust and reputation management, robust QoS delivery, reliable yet distributed monitoring mechanisms, and resilient computational economics need to be solved in way that is acceptable to both communities. The peer-to-peer research group [42] of the GGF is one example of an initial attempt to address some of these issues. At present the group is undertaking a comparative study of conventional client-server Grids and p2p Grids suggesting the OGSA working groups to incorporate p2p requirements. We hope that our contribution will help this effort.

Acknowledgments

This work was supported by the Mathematical, Information, and Computational Science Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-Eng-38. DARPA, DOE, and NSF support Globus Project research and development. The Java CoG Kit Project is supported by DOE SciDAC and NSF Alliance.

References

[1] I. Foster and C. Kesselman, Eds., *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, July 1998.

[2] D. Thain, T. Tannenbaum, and M. Linvy, *Grid Computing: Making the Global Infrastructure a Reality*. John Wiley, 2003, no. ISBN:0-470-85319-0, ch. Condor and the Grid.

[3] G. Bell and J. Gray, "What's next in high-performance computing," *Communications of the ACM*, vol. 45, no. 2, pp. 91–95, Feb. 2002. [Online]. Available: http://doi.acm.org/10.1145/503124.503129

[4] "The Access Grid Web Page," Web Page. [Online]. Available: http://www-fp.mcs.anl.gov/fl/accessgrid/

[5] I. Foster, "What is the Grid? A Three Point Checklist ," 22 July 2002. [Online]. Available: http://www.gridtoday.com/02/0722/100136.html

[6] I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *International Journal of Supercomputing Applications*, vol. 15, no. 3, 2002. [Online]. Available: http://www.globus.org/research/papers/anatomy.pdf

[7] "Unicore," Web Page. [Online]. Available: http://www.unicore.de/

[8] A. S. Grimshaw and W. A. Wulf, "The Legion Vision of a Worldwide Virtual Computer," *Communications of the ACM*, vol. 40, no. 1, pp. 39–45, January 1997. [Online]. Available: http://legion.virginia.edu/copy-cacm.html

[9] "The Global Grid Forum Web Page," Web Page. [Online]. Available: http://www.gridforum.org

[10] I. Foster, C. Kesselman, J. Nick, and S. Tuecke, "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration," Web page, Jan. 2002. [Online]. Available: http://www.globus.org/research/papers/ogsa.pdf

[11] "World Wide Web Consortium," Web Page. [Online]. Available: http://www.w3.org/

[12] "Doe Science Grid," Web Page. [Online]. Available: http://www.doesciencegrid.org/

[13] "The DataGrid Project," 2000. [Online]. Available: http://www.eu-datagrid.org/

[14] "GriPhyN - Grid Physics Network," Web page. [Online]. Available: http://www.griphyn.org/index.php

[15] W. E. Johnston, D. Gannon, and B. Nitzberg, "Grids as Production Computing Environments: The Engineering Aspects of NASA's Information Power Grid," 1999.

[16] K. Keahey, T. Fredian, Q. Peng, D. P. Schissel, M. Thompson, I. Foster, M. Greenwald, and D. McCune, "Computational Grids in Action: The National Fusion Collaboratory," Argonne National Laboratory, Tech. Rep., 2002. [Online]. Available: http://www-unix.mcs.anl.gov/~keahey/papers/FusionPaperSubmitted.pdf

[17] "National Research Grid Initiative," Web Page. [Online]. Available: http://www.naregi.org/

[18] T. Prudhomme, C. Kesselman, T. Finholt, I. Foster, D. Parsons, D. Abrams, J.-P. Bardet, R. Pennington, J. Towns, R. Butler, J. Futrelle, N. Zaluzec, and J. Hardin, "NEESgrid: A Distributed Virtual Laboratory for Advanced Earthquake Experimentation and Simulation: Scoping Study," NEES, Tech. Rep. 2001-02, February 2001. [Online]. Available: http://www.neesgrid.org/html/TR_2001/NEESgrid_TR.2001-04.pdf

[19] "Particle Physics Data Grid," Web Page, 2001. [Online]. Available: http://www.ppdg.net/

[20] "TeraGrid," Web Page, 2001. [Online]. Available: http://www.teragrid.org/

[21] G. von Laszewski, E. Blau, M. Bletzinger, J. Gawor, P. Lane, S. Martin, and M. Russell, "Software, Component, and Service Deployment in Computational Grids," in *IFIP/ACM Working Conference on Component Deployment*, ser. Lecture Notes in Computer Science, J. Bishop, Ed., vol. 2370. Berlin, Germany: Springer, 20-21 June 2002, pp. 244–256. [Online]. Available: http://www.mcs.anl.gov/~gregor/papers/vonLaszewski--deploy-32.pdf

[22] G. von Laszewski, G. Pieper, and P. Wagstrom, "Gestalt of the Grid," in *Performance Evaluation and Characterization of Parallel and Distributed Computing Tools*, ser. Series on Parallel and Distributed Computing. Wiley, 2003, (to be published). [Online]. Available: http://www.mcs.anl.gov/~gregor/papers/vonLaszewski--gestalt.pdf

[23] G. von Laszewski, M.-H. Su, J. A. Insley, I. Foster, J. Bresnahan, C. Kesselman, M. Thiebaux, M. L. Rivers, S. Wang, B. Tieman, and I. McNulty, "Real-Time Analysis, Visualization, and Steering of Microtomography Experiments at Photon Sources," in *Ninth SIAM Conference on Parallel Processing for Scientific Computing*, San Antonio, TX, 22-24 Mar. 1999. [Online]. Available: http://www.mcs.anl.gov/~gregor/papers/vonLaszewski--siamCmt99.pdf

[24] "The Globus Project," Web Page. [Online]. Available: http://www.globus.org

[25] E. Korpela, D. Werthimer, D. Anderson, J. Cobb, and M. Leboisky, "SETI@home-massively distributed computing for SETI," *Computing in Science & Engineering*, vol. 3, no. 1, pp. 78–83, January–February 2001.

[26] "Distributed.net Homepage," Web Page. [Online]. Available: http://www.distributed.net/

[27] "grid.org Homepage," Web Page. [Online]. Available: http://www.grid.org/

[28] R. Buyya, D. Abramson, and J. Giddy, "An Economy Driven Resource Management Architecture for Global Computational Power Grids," in *The 2000 International Conference on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, USA, 26-29 June 2000.

[29] "The Australian Greenhouse Office: Grid Interactive Systems," Web Page. [Online]. Available: http://www.greenhouse.gov.au/renewable/power/grid.html

[30] "Massachusetts Institute of Technology: Community Solar Power Initiative," Web Page. [Online]. Available: http://solarpower.mit.edu/

[31] "Napster Homepage," Web Page. [Online]. Available: http://www.napster.com/

[32] "Gnutella Homepage," Web Page. [Online]. Available: http://www.gnutella.wego.com/

[33] I. Foster and A. Iamnitchi, "On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing," in *2nd International Workshop on Peer-to-Peer Systems (IPTPS'03)*, Berkeley, CA, USA, February 2003.

[34] J. Ledlie, J. Shneidman, M. Seltzer, and J. Huth, "Scooped Again," in *2nd International Workshop on Peer-to-Peer Systems (IPTPS'03)*, Berkeley, CA, USA, February 2003.

[35] A. Oram, Ed., *Peer-To-Peer: Harnessing the Power of Disruptive Technologies*, 1st ed. O'Reilly, 2001.

[36] L. Pearlman, V. Welch, I. Foster, C. Kesselman, and S. Tuecke, "A Community Authorization Service for Group Collaboration," in *IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*, Monterey CA, USA, 5-7 June 2002.

[37] "VOMS Architecture v1.1," Web Page, May 2002. [Online]. Available: http://grid-auth.infn.it/docs/VOMS-v1_1.pdf

[38] M. Thompson, W. Johnston, S. Mudumbai, G. Hoo, K. Jackson, and A. Essiari, "Certificate-based Access Control for Widely Distributed Resources," 1999.

[39] D. W. Chadwick and A. Otenko, "The PERMIS X.509 Rol Based Privelege Management Infrastructure," in *7th ACM Symposium on Access Control Models and Technologies*, 2002.

[40] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The EigenTrust Algorithm for Reputation Management in P2P Networks," in *12th International World Wide Web Conference*, Budapest, Hungary, 20-24 May 2003.

[41] D. B. Skillicorn, "Motivating Computational Grids," in *2nd IEEE/ACM International Symposium on Cluster Computing and Grid (CCGrid02)*, Berlin, Germany, 21-24 May 2002.

[42] "Peer-To-Peer Working Group," Web Page. [Online]. Available: https://forge.gridforum.org/projects/p2p/