Viewpoints on Grid Standards

Andrew A. Chien[1], Xian-He Sun[2], and Zhi-Wei Xu[3]

[1] *Computer Science and Engineering, University of California, San Diego, U.S.A*

[2] *Department of Computer Science, Illinois Institute of Technology, U.S.A*

[3] *Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100080, P.R. China*

E-mail:

Received December 21, 2004.

Abstract At GCC 2003 in Shanghai in December 2003, a panel discussion was held on the future of Grid computing and on the role of the Globus Toolkit in future Grid standards. Panelists include Andrew Chien (UCSD, USA), Wolfgang Gentzsch (Sun), Francis Lau (HKU, China), Carl Kesselman (USC, USA), Satoshi Matsuoka (TIT, Japan), Xian-He Sun (IIT, USA), Richard Wirt (Intel), Liang-Jie Zhang (IBM Research) and Song-Nian Zhou (Platform Computing). The panel talks were stimulating and well received. Three of the panel talk notes are selected and included in this viewpoint.

1 Reflections on Grid Middleware — Lessons from the History of Unix ① and Linux ②③

It has recently been suggested that the Globus Toolkit ④ has or is likely to become the "Unix/Linux of the Grid". What is usually meant by this is that Globus would become a nearly pervasive and open and interoperable technology standard. When the term Linux is used, the idea is taken further to mean an open implementation standard. I present some observations on the bumpy road which we have traveled to convergence for Unix, and conclude with a few observations on the challenges that Globus must surmount to become an open technology standard of similar scale and impact.

1.1 Unix Fragmentation and Incompatibility (1980's)

Unix began life as an operating system focused low-end systems (PDP-11) being used for research and particularly software development environments. It was designed to be portable, and run on nearly any hardware with a modest porting effort. It was also built as a small, minimal system, with much of the functionality provided in libraries which were not generally expected to be standard. Unix was made widely available in source form through ATT and later UCB, resulting in broad diffusion of knowledge and expertise, and a large, passionate set of supporters and users.

As Unix moved into the commercial world, it became popular in workstations sold by companies such as Sun, Hewlett-Packard, IBM, SGI, as well as several software vendors who sold x86-based versions. As use increased, so did pressure for investment in the technology to meet an increasing spectrum of needs. Companies competing directly with each other also engaged in "customization or improvement" of their version of Unix. Both of these trends led to fragmentation with major efforts behind a BSD thread and ATT's SVR4 which later spawned efforts such as OSF which largely failed to converge commercial Unix systems. While Unix systems share many common features, the major commercial dialects that were created by this fragmentation are still distinct and incompatible.

1.2 Enter Third Party Applications, Linux, and Microsoft (1990's)

Three distinct forces have been pushing the Unix/Linux community to a new focus around shared interfaces, function, and implementations. First, in the 1990's major third party software vendors emerged (e.g., Oracle, Peoplesoft, SAP, etc.), building multi-billion dollar businesses, and exerting increasing influence over providers of operating systems and hardware. For such companies, a large number of incompatible Unix systems were a porting and support cost and an impediment to growth. Second, Microsoft began a very aggressive development of Windows NT, with the explicit public goal of supplanting the Unix dialects. By promising compatibility across many of these platforms, Microsoft was able to gain porting commitments (and even pri-

*Note

①Unix is a registered trademark of The Open Group.

②Linux is a registered trademark of LMI.

③This section is contributed by Andrew Chien. Andrew Chien is supported in part by the United States National Science Foundation under awards NSF EIA-99-75020 Grads and NSF Cooperative Agreement ANI-0225642 (OptIPuter), NSF CCR-0331645 (VGrADS), NSF ACI-0305390, and NSF Research Infrastructure Grant EIA-0303622. Support from Hewlett-Packard, BigBangwidth, Microsoft, and Intel is also gratefully acknowledged. The opinions expressed are the author's, and not those of the sponsoring organizations.

④Globus Toolkit is a registered trademark of the University of Chicago.

ority) from these third party software companies over many of the Unix dialects. Third, Linux had begun to grow like wildfire based on tight control by a strong technical team, innovative licensing (the Gnu Public License) which prevented fragmentation, and a strong free software ideological culture. These changes established a strong positive cycles amongst application vendors, hardware vendors, and end users around Linux. Remarkably, these changes led to major commitments behind Linux by all of the major Unix system vendors (IBM, Hewlett-Packard, SGI, and Sun) by the beginning of the 2000's decade.

1.3 Challenges for Globus

In drawing the parallels to Unix's development and adoption, I would encourage caution in drawing conclusions. Historical retrospectives often make the actual outcomes seem inevitable, and rarely capture the uncertainty which existed as events actually occurred. In rough terms, Unix/Linux passed through three major phases: Shared Technology, Vendor Differentiation, and then Convergence. Globus has been introduced as an open source, shared technology, already used by a large community of scientific grid applications developers and users. It has gained significant support from the research community, federal research programs, and endorsement by major computing vendors — IBM, Hewlett-Packard, Sun, SGI, and even early grid computing companies such as Platform. It would appear that all of these constituencies would like to jump ahead to Convergence.

Globus faces three major challenges to become a converged technology forming the foundation of the grid. First, Microsoft remains as a major holdout in the support and adoption of Globus. Second, web services represent another major thread of interoperation and federation technology which competes in role. While substantial efforts continue to knit these streams together, the possibility of "web service solutions" to grid problems already addressed by Globus remain a possibility. Third, commercial products based on Globus source and technologies are just becoming available. As such, the competitive pressures for differentiation and fracture are just coming into play. The use of BSD licensing means that the benefits of interoperability must win out over the marketplace pressures for differentiation — the licensing cannot force coherence. A hopeful note is that more than ever, computing vendors and their customers understand the value of robust interoperability. Perhaps this time the community can exploit its experience to jump ahead to Convergence.

2 Grid Verses Virtual Organization ⑤

Grid computing is revolutionary by enabling access to unprecedented computing power and shared information. The Globus Toolkit is the leading software system for Grid computing. While industry communities are increasingly accepting the concept of Grid, can Globus continue to serve as a community leading open source system? To answer these questions, we need to have a better understanding of what Grid computing is.

Foster and Kesselman have given a three-point criteria to define a Grid: Coordinates distributed resources, using standard, open, general-purpose protocols and interfaces, and the ability to deliver nontrivial qualities of service[1]. On one hand, this definition is general and is well accepted, but on the other hand, it is general and many questions remain open when digging into detail. For instance, to coordinate distributed resources, how many resources need to be coordinated and how are they coordinated? By the three-point definition of Grid, the quantity of resources and "how to coordinate" are not a concern to be a Grid. We have always been told that Information Grid is mimicking the electrical power grid. Recently, some leading researchers also claim that Grid technology and Peer-to-Peer technology are equivalent and merging[2]. Peer-to-Peer technology uses *existing resources* as peers to *serve each other* and requires a *massive number* of peers to make it work. To mimic the power grid, a Grid should *build few* powerhouses to *deliver services* to the users, similar to the concepts behind utility computing. To create a view of a single system to cover everything, from mimicry of the power grid to superset of peer-to-peer technology, is a hard task.

Using standard, open, general-purpose protocols and interfaces does not guarantee resource sharing and problem solving in dynamic, multi-institutional virtual organizations, rather it may just lead to another virtual organization. Assuming GPS is developed based on standard, open, general-purpose protocols and interfaces, can GPS considered as a Grid? ⑥ GPS must coordinate three to four distributed resources to provide the location service. In fact, following the three-point definition, GPS is not only a Grid, but also a very successful one, in the sense that it is pervasive and has already melted into our day-to-day life. However, it is clear that GPS is not the kind of Grid Globus intends to support. Applications such as the DoD BioWar and virtual reality emulation benefit from high computing power, but requires a "single private system" view for individualized quality of services. Current industrial developments, such as utility computing and business on demand, use technologies more along the lines of virtual organizations

⑤This section is contributed by Xian-He Sun. Xian-He Sun is supported in part by the United States National Science Foundation under awards NSF. The opinions expressed are the author's, and not those of the sponsoring organizations.

⑥We use GPS since GPS is well known and serves the purpose. If government control is a concern we can equally use WiFi locator services as an example here.

rather than Grid to enhance their policies and services. There is a disparity in perceiving Grid. Academia tries to provide a general infrastructure, which can coordinate all the virtual organizations. Industries are profit driven. They may or may not follow OGSA or Web Service standards to meet their set milestones. They could build a private Grid to provide a single service, which in general is more effective and reliable. Their private Grids may be a success and stay. These private Grids, with their own open protocols and interfaces, may be seen as new virtual organizations. Unless these private Grids are developed on top of the same Grid infrastructure, the continued success of these private Grids would mean an endless effort to coordinate existing virtual organizations.

Panelist Song-Nian Zhou used the Chinese proverb "United long would lead division. Divided long would lead union" to start his talk. Panelist Andrew Chien discussed the bumpy road of Unix/Linux development to illustrate the mixed driving force of uniformity and distinctness. All lead us to the same viewpoint, the success of (local) distributed computing leads to the demand of uniformity; the initial success of uniformity (Grid) motivates diverse, distinguished services. The question is how to support uniqueness under uniformity. Coordinating all existing and future virtual organizations does not look very promising, developing new virtual organizations under the same frame might be. Does current Grid infrastructure support the development of private Grids? Lifting the Grid technologies to meet user-specific requirements is extremely challenging. More importantly, it requires a fundamental change toward what Grid is. Our original understanding of Grid is as that of middleware designed to coordinate virtual organizations[3]. To support private Grids, we claim that the Grid is an infrastructure on which virtual organizations can be created based on users' need. Are we ready for the change?

Grid technology is still dynamic and evolving. We all agree that Grid has grown from scientific computing only to general Internet computing. Does the adaptation of OGSA and Web Service mean the Globus group is trying to develop a super virtual organization which covers every one's need? Has OGSA positioned itself well for this changing concept so private Grid can be created on demand? Virtualizing Grid services alone is not enough to support private Grids. Private Grid requires the system-level service virtualization. Following the pattern of operating system development, if a private Grid stands out, it would be similar to Microsoft Windows in Grid computing, if Globus continues its leads, it would be similar to Unix in Grid computing. Can Globus maintain its leads to the end? We must exercise caution in predicting the future as we are facing a new concept changing and not sure what Grid is and what Grid infrastructure is.

3 We Need Internet-Style Grid Standards [7]

We are at an important juncture in grid standards development. With almost the entire IT industry, certainly all the major commercial companies launching large corporate-wide projects to take advantage of the coming service-oriented paradigm shift, a real danger exists that Grid could fragment into incompatible islands, forced upon us by "standards" driven mainly by short-term commercial interests.

What we need is a community of sufficient critical mass, which follows the Internet Way pioneered by visionaries such as J.C.R. Licklider. As early as 1979, Licklider described what we call Grid today as an electronic commons "featuring cooperation, sharing, meeting of minds across space and time in a context of responsive programs and readily available information". More importantly, Licklider promoted the Internet spirit: The Net should be open and public minded. It is People's Net. This Internet style has stood us in good stead, from the APARTNET, the Internet, to the Web. As Grid aims to offer connectivity and integration at application and service level, not just hardware or contents connectivity, the Internet spirit is even more important today.

The Globus Alliance has been and still is a leading international research team in our global grid community. Although it actively engages the industry, which is a good thing, the team's research and standards efforts are mainly funded by public sources. This mixture promises a public spirit and at the same time a positive influence on the mainstream IT industry.

I agree with the Globus Alliance's position: The Open Grid Service Architecture (OGSA) is a grid standard, and the Globus toolkit hopes to become a widely used reference implementation of OGSA. In our current China National Grid project (CNGird, 2002–2005), we are using the Globus Toolkit 3 to build our grid software platform. Our work complements the GT3 work by focusing on the OGSA Platform layer, on top of OGSI. Version 1 of our software has been deployed on the seven nodes of China National Grid, spanning six cities of Beijing, Shanghai, Xi'an, Changsha, Hefei, and Hong Kong. The CNGrid software supports OGSI grid services and web services.

As GT3 progresses to switch from OSGI to WS Resource Framework and WS Notification, we plan to reengineer our software to take advantage of the convergence of grid and web services. Version 2 of the CN-

[1]This section is contributed by Zhi-Wei Xu. Zhi-Wei Xu is supported in part by China Ministry of Science and Technology (Grant No.2002AA104310), National Natural Science Foundation of China (Grant No.69925205), and Chinese Academy of Sciences (Grant No.20014010). The opinions expressed are the author's, and not those of the funding agencies.

Grid software will be released in early 2005 with improved OGSA Platform level supports for virtualization, grid processes, and policies and contexts[4]. We hope the Globus Alliance will continue to follow the Internet style and spirit, encouraging cooperation and community building, by maintaining an open, non-commercial attitude and a modular Globus Toolkit architecture.

References

[1] Ian Foster, Carl Kesselman (eds.). Concepts and Architecture. In The Grid 2: Blueprint for a New Computing Infrastructure, Morgan Kaufmann Publishers, ISBN 1-55860-933-4, 2004.

[2] Ian Foster, Adriana Iamnitchi. On death, taxes, and the convergence of peer-to-peer and Grid computing. In *2nd International Workshop on Peer-to-Peer Systems,* Berkeley, CA, February 2003.

[3] X H Sun, Alan Blatecky. Guest Editors Introduction, Middleware: The key to next generation computing. *Journal of Parallel and Distributed Computing*, Special Issue on Middleware, 2004.

[4] Zhiwei Xu, Wei Li, Li Zha *et al.* Vega: A computer systems approach to grid computing. *J. Grid Computing*, to appear.

Andrew A Chien is Director of Center for Networked Systems and SAIC Chair Professor at Computer Science and Engineering, University of California, San Diego, USA.

Xian-He Sun is a professor at Department of Computer Science, Illinois Institute of Technology, Chicago, USA.

Zhi-Wei Xu is a professor and deputy director at Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China.