

Towards a Standards-based, Message-oriented Advanced Collaboration System

Steve Smith

ViSLAB, The University of Sydney
J12 School of IT Building
Sydney 2006 NSW
ssmith@vislab.usyd.edu.au

Masahiro Takatsuka

ViSLAB, School of IT, The University of Sydney
J12 School of IT Building
Sydney 2006 NSW
masa@vislab.usyd.edu.au

The active use of Access Grid systems by many scientific and engineering communities over several years has uncovered a number of architectural and usability problems of this technology. After examining the system architecture of the current Access Grid and the recent trend in the instant messaging systems, we devised an alternative advanced collaboration system using an established standard protocol.

According to our approach, the Access Grid is seen as a set of functions and services rather than a defined system architecture. We used the Extensible Messaging and Presence Protocol to implement Access Grid functionality with a much simpler architecture. This system allows us to easily extend its capability using the plug-in mechanism. Furthermore, it allows developers to design various different client applications with appropriate user interfaces. These features are significantly important for supporting a wide range of communities with different requirements.

Key Words and Phrases: Design, Theory, XMPP, Jabber, IETF, Access Grid

ACM Classification: H.5.2 (Information Interfaces and Presentation); H.1.2 (User/Machine Systems); D.2.11 (Software Architectures)

1. INTRODUCTION

Advances in network computing infrastructure has spawned the concept of computing grids (Foster *et al*, 2001). This in turn has become the basis of Version 2 of the Access Grid Tool-Kit (www.accessgrid.org) which provides coordination and authentication for Access Grid sessions, providing the framework for configuring and controlling client applications. It has been used to support many national and international scientific projects. Moreover, communities from non-scientific fields are now investigating the use of this technology to support their research and business activities.

However, several years of usage of the Access Grid Toolkit has brought to light a number of problems with the existing architecture. In addition to the architectural issues, problems in security

Copyright© 2007, Australian Computer Society Inc. General permission to republish, but not for profit, all or part of this material is granted, provided that the JRPIT copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Australian Computer Society Inc.

Manuscript received: 13 December 2006

Communicating Editor: Professor Bill Appelbe

and usability have also been identified. A working-group has recently been formed to evaluate the existing infrastructure and recommend a new design (Party, 2005). However, it is anticipated that it will be some time before there is a substantial outcome from this working-group.

In this paper, we present an alternative view of what the Access Grid is by focusing on the functionality that connects physically/conceptually separate communities. This alternative view allows us to consider different architectures rather than the Grid and web-services frameworks used in the current Access Grid Toolkit. We examine a Client/Server architecture, which is employed in most Instant Messaging systems. The proposed implementation presented in this work is based on a well established internet standard.

2. REVIEW OF ACCESS GRID ARCHITECTURE

The Access Grid-2 architecture is based on Version 3 of the Globus grid toolkit. Globus is based on web-architectures such as HTTP and SOAP (Foster *et al*, 2001). Although layered on-top of the reliable, connection-oriented TCP/IP transport these protocols are stateless (Vogels, 2003), thus additional state mechanisms must be supplied by the grid architecture (Foster *et al*, 2004). Globus provides authentication and authorisation via an X.509 certificates scheme, the current Certificate Authority for the Access Grid is the Argonne National Laboratory. The recent efforts put into the new release of Access Grid Toolkit removed the Globus layer. It is, however, still based on the same architecture.

Conceptually the Access Grid consists of venue-servers which contain venues, roughly analogous to a building with meeting-rooms. Groups of meeting rooms are aggregated into 'lobbies', which are themselves rooms. Venues store meta-data (e.g. video and audio multicast addresses, details of other users present) and provide the conduit for shared applications, and a rudimentary text-chat system. This data is acted on by the venue-client which is responsible for starting the applications. The client system may span multiple machines (so-called multi-machine nodes).

Except the efforts in separating the Globus layer from the core Access Grid architecture, much of the recent work on the Access Grid toolkit has focused on patching up the deficiency of collaborative functionalities with existing widely used collaboration tools such as Virtual Network Computing (VNC), Jabber-based Instant Messaging systems. The share application capabilities built into the current Access Grid is based on the event-based synchronization of multiple instances of applications. This collaborative architecture supports only simple and primitive collaboration applications, for instance, a web-browser, an image viewer and a powerpoint. This approach fell short in addressing more complex and sophisticated collaboration sessions involving massive datasets and much more complex media rich applications.

2.1 Limited Communication Scope

The role of the server in Access Grid world-view is minimal; users connected to a server can only communicate within the same venue (room), and the design of the client is such that a user can only be present in one room at a time. Thus the Access Grid can be seen as isolated islands of communication scope as shown in Figure 1.

In practice, this means that coordination/organization of Access Grid sessions must take place over other communication mediums, such as phone and email. This can be problematic in the case where there is a misunderstanding or misconfiguration. It is not uncommon for users be logged-in to a server but be in the wrong venue; despite being on the same server there is no method of knowing where that other user is on the server, or if they are on at all. When this happens the users usually fall-back to phone or other communication medium to correct the situation. This, in fact,




Figure 1: Islands of communications spaces in the current Access Grid

motivated the inclusion of a Jabber-based instant messaging client in the latest release of the Access Grid. However, it was provided without addressing the fundamental architectural problem of the Access Grid.

2.2 Limited Presence

Related to the above issue is the problem of user presence, which is poorly modelled in the disconnected web-services architecture of the Grid. The Access Grid metaphor is connection-oriented, in that users are connected to a room where they can communicate. However the underlying implementation based on Globus is based on web-services, which utilise transient HTTP connections. This can lead to errors in a person's state on the server, as can be demonstrated with the following exercise:

- Alice and Bob are connected to the server. Alice can see Bob in the user list.
- Bob's client crashes and he is disconnected. However, as his session has not timed out yet he remains in Alice's user list.
- Bob restarts his client and reconnects. His previous session is still active, so there are now two Bobs in the list.

The authors have seen cases where a buggy client has crashed and reconnected so frequently that the server has locked up and required a restart. Although this issue may be exposing implementation problems as much as architecture problems, it wouldn't occur in architecture that utilises the connectivity of the underlying TCP/IP protocol. For instance, when Bob's session crashed, the underlying operating system has responsibility for terminating the TCP/IP session.

Another side effect of the use of web-services in the Access Grid can be seen in the performance problems of the current implementation. The necessity of reauthenticating with each transient connection creates a large overhead in the protocol. Utilisation of a single persistent connection allows a single authentication phase for a session.

3. AN ALTERNATIVE VIEW: ACCESS GRID AS INSTANT-MESSAGE SERVICE

One view of the development of Access Grid Version 2 can be seen as introducing functionality from the world of instant messaging. User presence and file transfer were added, and chat functionality, previously implemented with a separate ‘Mud’ application, was integrated into the client (although it is still limited, and large sessions will frequently utilise Jabber clients and servers for chat and coordination).

Meanwhile, on the other side of the fence the world of instant-messaging is increasingly moving into the field of video-conferencing and online collaboration. Consumer-level video-chat has existed since the introduction of CU-SeeMe in 1992 (Park, 1992), and Apple’s iChat, AOL’s AIM, Microsoft’s Messenger and Skype have all added video and voice communication capabilities. More recently iChat has added the ability to have up to four participants connected with audio and video. The similar features are scheduled for Microsoft’s Messenger related products.

3.1 An Architecture Based on the XMPP Protocol

All of the IM protocols mentioned above are proprietary, but an open standard equivalent exists in the form of the IETF Extensible Messaging and Presence Protocol (XMPP), better known as ‘Jabber’. From the Jabber Foundation homepage:

Jabber is best known as “the Linux of instant messaging” – an open, secure, ad-free alternative to consumer IM services like AIM, ICQ, MSN, and Yahoo ... Under the hood, Jabber is a set of streaming XML protocols and technologies that enable any two entities on the Internet to exchange messages, presence, and other structured information in close to real time.

The base XMPP RFCs (3920 and 3921) was approved by the IETF in October 2004 (Saint-Andre, 2004a; 2004b). It was designed from the ground up to be extensible, and as such a proposal system has been developed around it, borrowing heavily from other development organisations such as IETF, Python, and the W3C (Saint-Andre, 1999). From this system a large number of Jabber Enhancement Proposals (JEPs) have been authored and are in various stages of standardisation.

3.2 Mapping Access Grid Functionality to Jabber

According to the extensive list of current JEPs, it is believed that the existing JEPs provide sufficient functionality to implement the current functionality of the Access Grid Toolkit 2.4 and 3.0. Moreover, the extensible nature of XMPP allows additional messages to be added if deemed necessary. The only prerequisite is that each server has a user with the Jabber ID (JID) ‘ag@host’, which acts as the moderator. This user can be a real-person, a software-bot, or both. Here we provide a proposed mapping of Access Grid to XMPP and JEPs:

Rooms containing users/Text chat: JEP-0045 (Saint-Andre, 2004c) defines a method of multi-user chat (MUC) within rooms. This provides all the functionality of the current chat service with many enhancements. User access to the room can be regulated by the moderator (see below).

Servers: One major divergence from the current model of the Access Grid is that the user does not connect directly to the server but instead has an account on a (possibly) different server, and

messages are passed by server-to-server communication. In this sense Jabber resembles SMTP more than traditional chat applications. Thus organisations would generally have their own Jabber server and manage users locally rather than connecting to third-party servers.

Per-room multicast video and audio: Vislab, the University of Sydney has developed a minor extension to the XMPP MUC protocol to allow the publishing of arbitrary meta-data to a chat-room. This allows us to publish the relevant audio and video meta-data. The format of this data is an XML-ised form of the Session Description Protocol format (Handley and Jacobson, 1998) based on drafts of the SDPng protocol from the SDP working-group (group). This fragment is published in the <http://vislab.usyd.edu.au/protocol/xsdp> namespace.

```
<cfg xmlns='http://vislab.usyd.edu.au/protocol/xsdp'>
  <component name='AG Video' media='video'>
    <alt format='rtp-avp-31'
name='AG-video-mcast'
addr='233.2.178.9'
ttl='127'
rtp-port='17018' />
  </component>
  <component name='AG Audio' media='audio'>
    <alt format='rtp-avp-112'
name='AG-audio-mcast'
addr='233.2.178.9'
ttl='127'
rtp-port='17008' />
  </component>
</cfg>
```

File transfer: The current Access Grid method of file-sharing/transfer is to upload the file to the venue where it remains to be downloaded by clients. This has the advantage of allowing the file to remain in the venue after the user has logged-out, but places a load in terms of network traffic and storage requirements on the server.

In contrast the Jabber protocol focuses on peer-to-peer file transfer. Users may send files directly to other users, or make files available via the Jabber publish/subscribe mechanism (JEP-0060; Millard *et al*, 2005) (or possibly the experimental JEP-0137; Miller and Muldowney, 2004). File transfer takes place via the file transfer protocol (JEP-0096; Muldowney *et al*, 2004). One problem with P2P file transfer mechanisms is the situation where the person offering the files is behind a firewall; Jabber addresses this in JEP-0096 by offering the alternative transfer mechanisms of SOCKS5 Bytestreams (JEP-0065; Saint-Andre, 2004d) and in-band bytestreams (JEP-0047; Karneges, 2003).

Shared applications: XMPP has already been used for experimental shared applications, including games (checkers, chess, etc.) and shared whiteboards by embedding XML commands (or even XML documents) in messages. New messages formats may be added for additional applications:

```
<message
  from='ssmith@vislab.usyd.edu.au/notify'
  to='apac@jabber.vislab.usyd.edu.au'
  type='groupchat' >
  <event xmlns='http://vislab.usyd.edu.au/protocol/presentation' >
    <name>APAC Presentation</name>
    <streamid>file-apac-pres.ppt</streamid>
    <page>12</page>
  </event>
</message>
```

Note that the above is a perfectly valid XMPP message due to the XML namespace qualifier. Clients not able to interpret the messages will silently drop them. These messages may be sent to individuals or chat-rooms (therefore acting as a broadcast mechanism).

More advanced applications may also wish to utilise some form of RPC mechanism. JEP-0009 provides a mechanism for using Jabber as the transport for XML-RPC (Winer, 2003).

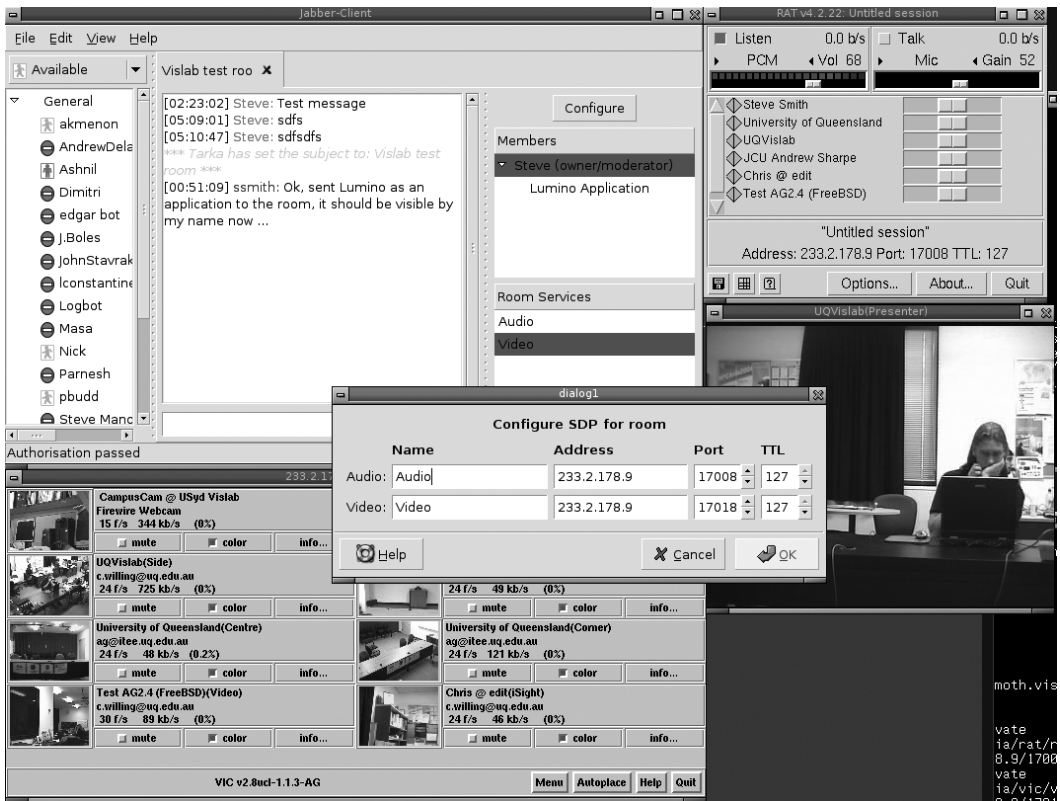


Figure 2: A screenshot of Jabber-based Access Grid illustrating the instant messaging-based main application. The audio, video and other collaboration functions are provided as services associated with the group chat session.

Figure 2 illustrates our prototype system. The left top window is the main application, which is based on XMPP. It consists of

- a members list,
- a text-based chat window,
- a list of members currently in the group chat-room,
- a list of collaboration services available in the chat-room,

In this example, the audio and video capabilities are provided at the chat-room level. It also should be noted that the moderator of this group-chat room offer a collaboration service indicated as “Lumino Application”, which provides a distributed scientific visualization capability (Stavarakakis *et al*, 2006). This example also shows the audio/video service configuration window, which sets multicast network addresses for this group chat session.

3.3 Authentication Issues

Authentication on the current Access Grid is achieved through the client-side X.509 certificate functionality of the Globus toolkit, with the certificates being issued by the Argonne National Laboratory. The XMPP protocol provides for the use of SASL (Myers, 1997), which in turn provides for certificate-based authentication through the EXTERNAL mechanism. Thus the equivalent functionality can be supported by the Jabber server.

However the security value of PKI has increasingly been called into question (Ellison and Schneier, 2000). In particular, little or no authentication of the entity requesting the certificate is done by ANL. On top of this the identity of the users connected to the server are not based on the certificate but on data input by the user separately, thus the use of the user roster in the client for identification is easily spoofed. Finally, it is also possible to retrieve an anonymous certificate, which bypasses the necessity of the entity identifying itself at all. Due to the bottleneck of having to pass through a single CA these anonymous certificates are widely accepted by servers. Thus in practice the authentication functionality is not used.

One alternative to this is to utilise a web-of-trust system as utilised in the OpenPGP system (Caronni, 2000), and some work has been to analyse their security (Jiang *et al*, 2004). One method of using this is for a user to publish their public key in a pub-sub record and for others to authenticate them against the web-of-trust. Additionally the Jabber foundation is currently working on a method of end-to-end encryption and authentication in JEP-0116 (Paterson *et al*, 2005) modelled on the SSH protocol. However, more investigation needs to be done into this before a solution is reached.

One method to address encryption and authentication is to limit access to MUC rooms purely to known JIDs. In particular the room should only allow entry to JIDs, which originate from the same server as the chat-room host, as server-to-server security cannot be guaranteed. In conjunction with SSL/TLS connections to the server and encryption keys specified in the SDP meta-data a ‘reasonable’ level of security may be achieved on intranets. Given the current encryption limitations of Vic and Rat (DES-1) the Access Grid should not be relied upon for strong security in any case.

One important feature only partially supported in the Access Grid toolkit is the ability to book, secure and encrypt venues. This can be achieved in the proposed model by the moderator. The moderator has the ability to impose strict access privileges on both chat-rooms and publish/subscribe nodes. Given that the moderator may be an agent (‘bot’) rather than a human the moderator may interface to an online-booking system and restrict access to chat-rooms and associated information nodes to registered participants at certain times, and generate and publicise encryption keys for those sessions via the multicast information node.

3.4 Implementation Issues

One major advantage of adopting the XMPP protocol over building our own architecture is the immediate availability of a library of open-source code on which to base the work. In fact the currently available servers can be adopted with little or no modification. As long as an adopted server supports the Jabber Component Protocol (JEP-0114; Saint-Andre, 2005) JEPs can be implemented by the addition of an external plug-in.

On the client side there are already a large number of clients to base work on (Jabber.org currently list over 70 free and open-source clients). Additionally there are library implementations of the Jabber protocol in virtually every open-language on which to base plug-ins to share-applications.

3.5 Interoperability

It is important that we maintain a degree of interoperability with the current Access Grid implementation for coexistence and/or during any transition period. Fortunately A/V interoperability is largely a configuration issue. Venues on the multicast network are merely pre-arranged address/port combinations. By creating chat-rooms sharing the same name and multicast configuration as existing venues, we create an intuitive link to existing installations for audio and video. This is the method used in transition from Version 1 of the Access Grid.

In the case that a more advanced interoperability is needed, it is possible to create a gateway system to link to the existing Access Grid servers. The Jabber system was originally designed to be a method of interoperating with other communications systems via server plug-ins. An Access Grid-2 gateway plug-in can be created based on the current Access Grid2.3 toolkit.

4. AREAS NOT ADDRESSED

One aspect not addressed in this document is a multi-machine node. These configurations exist for two reasons; processing speed and limitations in the video transmission tool 'Vic'. Traditionally multiple machines have been necessary to maintain real-time processing speeds under heavy load. However with modern processors this is rarely a problem, and the impending move to consumer-level multi-core processors this requirement will lessen even more.

The other issue is that currently a single Vic instance can only transmit a single video stream, thus where multiple cameras are desired there must be multiple instances, which rapidly leads to a cluttered user interface. Additionally these interfaces cannot see packets transmitted by each other due to multicast implementation issues. For these reasons sites with multiple cameras often place the transmitting instances on dedicated video machines and run a display-only Vic instance on the display node, with the remote instances being controlled via Globus.

This second problem has not been addressed here for two reasons:

- This is a client-side implementation problem. Many technologies for IPC across multiple machines and these can be utilised to implement this functionality if deemed necessary.
- As shown above, this necessity is largely an implementation issue with the video processing tool, and would be better addressed by modifying or replacing this software.

5. CONCLUSION

The new alternative Access Grid implementation put forward in this paper employs the IETF Extensible Messaging and Presence Protocol (XMPP). It was inspired by the idea that the Access Grid is a set of functions, and this set can be provided by different implementations.

Adopting IETF's XMPP standard enables us to design the alternative Access Grid system with a much simpler architecture. Furthermore, the system is easy to extend using the plug-in


mechanism. This encourages the seamless integration of various services (such as shared-applications, different codec, etc.) developed by third parties. Moreover, it allows developers to design a client application with a better user interface. We believe that the use of a well established standard with the provision for easy plug-in capability is very important for supporting a wide range of scientific communities.

REFERENCES

- CARONNI, G. (2000): Walking the web of trust. In Proceedings of IEEE 9th International Workshops on Enabling Technologies: Infrastructure for Col laborative Enterprises (WET ICE2000). *IEEE Computer Society*, 153–158.
- ELLISON, C. and SCHNEIER, B. (2000): Ten risks of pki: What you're not being told about public key infrastructure. *Computer Security Journal* 16(1): 1–7.
- FOSTER, I., FREY, J., GRAHAM, S., TUECKE, S., CZAJKOWSKI, K., FERGUSON, D., LEYMAN, F., NALLY, M., SEDUKHIN, I., SNELLING, D., STOREY, T., VAMBENEPE, W. and WEERAWRANA, S. (2004): Modeling stateful resources with web services – version 1.1. Whitepaper, IBM. March.
- FOSTER, I., KESSELMAN, C. and TUECKE, S. (2001): The anatomy of the Grid: Enabling scalable virtual organizations. *International Journal of Supercomputer Applications* 15(3).
- GROUP, S. W.: Sdpng protocol. <http://www.dmn.tzi.org/ietf/mmusic/sdp-ng/>.
- HANDLEY, M. and JACOBSON, V. (1998): Sdp: Session description protocol. RFC-2327.
- JIANG, Q., REEVES, D. S. and NING, P. (2004): Improving robustness of pgp keyrings by conflict detection. In *CT-RSA*. 194–207.
- KARNEGES, J. (2003): Jep-0047: In-band bytestreams (ibb). <http://www.jabber.org/jeps/jep-0047.html>.
- MILLARD, P., SAINT-ANDRE, P. and MEIJER, R. (2005): Jep-0060: Publish-subscribe. <http://www.jabber.org/jeps/jep-0060.html>.
- MILLER, M. and MULDOWNEY, T. (2004): Jep-0137: Publishing si requests. <http://www.jabber.org/jeps/jep-0137.html>.
- MULDOWNEY, T., MILLER, M. and EATMON, R. (2004): Jep-0096: File transfer. <http://www.jabber.org/jeps/jep-0096.html>.
- MYERS, J. G. (1997): Simple authentication and security layer (sasl). RFC-2222.
- PARK, S. J. (1992): A history of video conferencing (vc) technology. <http://myhome.hanafos.com/soonjp/vchx.html>.
- PARTY, A. G. W. (2005): Future of access grid. In *Proceedings of the Access Grid 5th Annual Retreat*. Millbrae, California, USA.
- PATERSON, I., SAINT-ANDRE, P. and SMITH, D. (2005): Jep-0116: Encrypted sessions. <http://www.jabber.org/jeps/jep-0116.html>.
- SAINT-ANDRE, P. (1999): Jep-0001: Jabber enhancement proposals (jeps). <http://www.jabber.org/jeps/jep-0001.html>.
- SAINT-ANDRE, P. (2004a): Extensible messaging and presence protocol (xmpp): Core. Tech. Rep. RCF-3920, IETF, XMPP Working Group.
- SAINT-ANDRE, P. (2004b): Extensible messaging and presence protocol (xmpp): Instant messaging and presence. Tech. Rep. RCF-3921, IETF, XMPP Working Group.
- SAINT-ANDRE, P. (2004c): Jep-0045: Multi-user chat. <http://www.jabber.org/jeps/jep-0045.html>.
- SAINT-ANDRE, P. (2004d): Jep-0065: Socks5 bytestreams. <http://www.jabber.org/jeps/jep-0065.html>.
- SAINT-ANDRE, P. (2005): Jep-0114: Jabber component protocol. <http://www.jabber.org/jeps/jep-0114.html>.
- STAVRAKAKIS, J., LAU, Z.-J., LOWE, N. and TAKATSUKA, M. (2006): Exposing application graphics to a dynamic heterogeneous network. In *Proceedings of the 14-th International Conferences in Central Europe on Computer Graphics, Visualization and Computer Vision*, JORGE, J. and SKALA, V. Eds. Plzen, Czech Republic, 71–78.
- VOGELS, W. (2003): Web services are not distributed objects. *IEEE Internet Computing* 7(6): 59–66.

BIOGRAPHICAL NOTES

Masahiro Takatsuka is Senior Lecturer and Director of ViSLAB at the University of Sydney, Australia. His interests are in exploring Advanced Collaboration Technologies, in particular, the use of Service Oriented Remote Collaboration. He also works in the area of Data/Information Visualization and Distributed Computer Graphics. He is a member of the IEEE and ACM.



Masahiro Takatsuka

Steve Smith was a project manager for the Access Grid program as well as a key system administrator at ViSLAB, the University of Sydney. His interests include Network Communication, 3D Computer Simulation and Programming Languages. He currently works at Atlassian, Sydney, Australia.



Steve Smith