

GFD-I.043
Category: Informational
Advanced Collaborative Environments-RG

D. Agarwal – Lawrence Berkeley National Laboratory
B. Corrie – Simon Fraser University
J. Leigh – University Illinois at Chicago
M. Lorch – Virginia Tech
J. Myers – Pacific Northwest National Laboratory
R. Olson – Argonne National Laboratory
M. E. Papka – Argonne National Laboratory
M. Thompson – Lawrence Berkeley National Laboratory

December 22, 2004

Security Requirements of Advanced Collaborative Environments (ACEs)

Status of This Memo

This memo provides information to the Grid community on the security requirements of advanced collaborative environments. This document is specifically intended for the Grid Security community, as data to guide their design work.

GGF Editor Note: This is an INFORMATIONAL document. It does not define any standards or technical recommendations. Distribution is unlimited.

Copyright Notice

Copyright © Global Grid Forum (2004). All Rights Reserved.

Abstract

This document describes the security requirements of a broad range of advanced collaborative and/or visualization environments. These include: the Access Grid (**AG**); Tele-Immersion (**TI**); Remote Visualization (**RV**); Dynamic and Asynchronous Environments (**DAE**); and Collaborative Experiments (**CE**).

The contributors of this document are experts in their respective fields of ACE. The decision on the content of the document was guided by experts in Grid security. The goal of this paper is to highlight security issues that are important to the community of grid users represented by the various authors of the document.

Contents

Abstract	1
1. Introduction	3
2. Usage Scenarios.....	3
2.1 Access Grid (AG)	3
2.2 Tele-Immersion (TI).....	3
2.3 Dynamic and Asynchronous Collaborative Environments (DAE).....	3
2.4 Collaborative Experiments (CE).....	3
2.5 Remote Visualization (RV)	4
2.6 Intended community for each of the scenarios	4
3. Implementation of the Scenarios	4
3.1 Hardware	4
3.2 Specialty devices and displays.....	4
3.3 Network resources	4
3.4 Data resources	4
3.5 Access permissions.....	5
3.6 Software.....	5
3.7 Connectivity model	5
4. How Security Properties Apply to the Implementation.....	7
4.1 Assets that require protection	7
4.2 Need for extensible and dynamically definable access rights.....	7
4.3 Estimated value/cost to recover compromised systems.....	8
4.4 Anticipated Threats.....	9
4.5 Security features needed and protected items that need them	9
5. Performance requirements.....	11
Author Information and Role	12
Glossary.....	12
Intellectual Property Statement	12
Full Copyright Notice.....	12
References	13

1. Introduction

The continuing drop in cost of computing, video, graphics and high-speed networking is fast making distance collaboration not only possible but routine. It is therefore crucial now, more than ever, for developers of advanced collaboration systems to also be able to provide *secure* collaboration services. Much research is still needed to develop the correct security model for these environments.

This document is intended to assist the community of Grid security experts in their design work, by providing them with realistic requirements of ACEs.

The first half of the document provides usage scenarios of a number of ACEs including AccessGrid (AG), Tele-Immersion (TI) [Leigh97], Dynamic and Asynchronous Environments (DAE), Collaborative Experiments (CE) and Remote Visualization (RV). The second half of the document summarizes the available resources and security requirements of each of these application areas. In each section, there has been an effort to provide the requirements that apply across all of the usage scenarios and additional requirements specified by specific usage scenarios. Some usage scenarios did not provide specific input to some sections. The usage scenarios listed reflect the authors' experiences. These requirements may not generalize to all installations of a particular usage scenario and all usage scenarios.

2. Usage Scenarios

In order to lay the foundation for future sections each of the focus areas wrote a brief usage scenario. The following sub-sections lay out these scenarios.

2.1 Access Grid (AG)

A meeting of geographically distributed collaborators, each using a "designed space" that explicitly supports the high-end audio and visual technology needed to provide a high-quality compelling and productive user experience; these spaces are referred to as Access Grid nodes [Stevens03]. In a typical usage, several collaborators are presenting results to their colleagues using a distributed presentation application. This application requires the distribution of a presentation data file to each participating site. Meetings may be recorded by a multimedia server for later playback [Disz97].

2.2 Tele-Immersion (TI)

This scenario involves a small group of remote collaborators viewing data sets on immersive displays such as the GeoWall [GW], ImmersaDesk or CAVE. Participants may appear to each other as a puppet-like avatar, or as static photographs, or dynamic video streams. Participants are able to hear each other via streaming audio. In the case of the GeoWall it is typically placed next to an Access Grid. Participants may be able to bring remote data sets into the environment. In some cases the data may be large static data sets (like 3D models). In other cases, the data may be a continuous stream of data from a computer simulation or a collection of sensors. These data generators are mediated by a state server that is responsible for making sure these changes are propagated across all collaborating clients.

2.3 Dynamic and Asynchronous Collaborative Environments (DAE)

An on-going or ad hoc collaboration that spans time zones and involves a changing set of participants. This scenario describes collaboration where the participants and duration are not necessarily known in advance and the set of participants is changing often. It also describes situations where the collaborators are often unable to interact in real time.

2.4 Collaborative Experiments (CE)

Users of an instrument facility submit proposals for collaborative work. Once accepted, they send any required physical material to the lab and work with facility instruments and staff to complete their experiments. Interactions include remote/collaborative control of the instrument's computer-

based controls, discussions with facility staff, creation and retrieval of data and experiment notes, and viewing the laboratory activity.

2.5 Remote Visualization (RV)

Users are located at a remote location that is not collocated with the generation of the visualization, which is not to say that the user is geographically far away from the visualization. Generation of the visualization may happen on a machine that is located on the same local area network; the point is the visualization is not generated on the local machine.

2.6 Intended community for each of the scenarios

In the case of TI, the community consists of scientists who want to view 3D data sets collaboratively. The Access Grid community includes both academic and industrial users in many countries. The DAE scenario is intended to support on-going and ad hoc collaborations in the business and academic community. The CE environment supports users of shared scientific instrument facilities.

3. Implementation of the Scenarios

This section outlines the hardware and software that are common to the various focus areas to set a baseline for security discussions. Each of the application areas makes use of standard commodity PC's and Workstation class machines. The machines are used either in a single desktop setting or sometimes as part of a larger cluster. The displays are either standard desktop display solutions or projectors. In some application settings the machines have additional devices connected to them, such as audio and video input and output devices. In all cases the machines interact with other machines in the same or different administrative domain. Direct physical access to most machines used is relatively public and several people might use them on a regular basis. Some of the machines are only generally accessible to one person. In the case of TI, AG, and RV some output devices and machines might be in specialized laboratories with controlled access to the room.

3.1 Hardware

All major hardware platforms are in use. Almost all of the identified scenarios are making use of similar hardware platforms representing the three major OS platforms (Mac, Linux, and Windows) with Solaris also being used in the CE area. Hardware based security solutions such as smart cards or cryptographic accelerators are not common, large facilities or specialty hardware associated with TI, AG, and RV are often controlled via locked doors. Some CEs are making use of SecureID tokens.

3.2 Specialty devices and displays

For VR, TI, AG, and CE large scale displays and specialty devices are often but not always located in specialized laboratories and limited access locations. These devices include very large high resolution displays and immersive environments. In the case of AG, the specialty devices also include cameras, microphones, echo cancellers, and projectors. Because the display systems are large scale, all users in the room need to be considered to have access to the environment. Thus if strict authorization/authentication is required then all people in the room would need to be authorized/authenticated.

3.3 Network resources

Since in this document we are discussing the requirements of collaborative applications, it is clear that all the focus areas involve networking between participating users and machines. Networking resources required vary depending on the application. Applications such as TI, AG, and VR sometimes require special capabilities from the network such as bandwidth, latency and jitter guarantees. The AG scenario also requires multicast capabilities.

3.4 Data resources

In the case of TI, access to distributed data sets on multiple computers is required. Permissions for these data sets may be open or may be restricted to certain users and/or at certain locations.

Data is imported into a persistent collaboration session for the purposes of visualization. Data is not archived in the environment (the data in the persistent store is a copy of the original data) and is only a copy to be used for visualization purposes. Ownership of the data is maintained by the creator/importer of the data set. The owner and the administrator are the only users who can delete the data set. Data is readable by all users in the group. Permissions to modify/write the data can be granted to other users in the group. Only those users with modification privileges can change the data set.

3.5 Access permissions

Resources typically involve one or more computers, with permission to access those computers typically provided by standard OS-level authorization techniques. The local resources of the AG node computing gear are dedicated to the node; hence permission to use them is implicit.

An AG session is in the context of a Venue. Each participant must have been granted access to the Venue in order to join the session. An open session would allow anyone with a valid AG credential to enter the venue; a closed session requires prior configuration of authorization of the venue to enable entry only of the desired participants. In the case of DAE, the resources in the environment are added and removed dynamically so the access permissions need to be managed dynamically. In the case of RV, the output is often an image or images displayed of some sort of display surface in which securing the output requires securing who can look at the display.

3.6 Software

The software for these focus areas is typically custom built but often leverages existing frameworks when they are available.

3.6.1 Applications

In the case of TI, Some packages (such as AVS) have modules that support immersive displays and can be used in a collaborative environment but to support the richness of immersive collaboration, custom applications that provide collaborative interaction are typically required.

The core AG software is composed of a combination of custom applications and third-party applications (for media tools, etc). The AG system allows users to install additional third-party extensions to the Venue Server and client-side applications. In the case of DAE, core applications include messaging, shared file, and shared application tools. In the case of CE, domain specific applications for data acquisition, analysis and visualization (e.g. for NMR - Varian vNMR, Insight) are in use. General visualization tools and presentation graphics software are also usually available locally. In the case of RV, besides custom built applications, sometimes pixel scraping technologies like VNC are also used to share images.

3.6.2 Frameworks (servlets,web services)

In the case of TI, the CAVE library, VR Juggler, Aura [Germans01], and CAVERNsoft [Park00] are examples of open source frameworks in use. There are many others also in use. The AG software environment defines its own framework for extensibility, this framework is based on the Globus Toolkit, the pyGlobus Python COG Kit, and SOAP-based web services. In the case of DAE, Java, web services, peer-to-peer and peer-to-peer based on group communication are in use. In the case of CE, Java and CGI are in use.

3.7 Connectivity model

The focus areas make use of client/server connections and some of the application areas make use of peer-to-peer and group communication mechanisms. The sequence of how the connections are made is specific to each application.

3.7.1 Topology (server/client, peer-to-peer, how protocols are used)

Each of the focus areas make use of client/server architectures. The primary difference between the areas is in the number of clients expected and whether they are dependent on the servers to be continuously available.

In the case of TI, typically a client/server architecture is used with a small number of collaborators. UDP protocols transport time critical data while TCP transports provide for reliable control and data channels. For large bulk data transfers (especially over long fat networks), either parallel TCP [Leigh01] or reliable UDP is used [Leigh01, He02].

In the case of AG, the connectivity between the venue server and the AG nodes is client/server. This communication uses SSL/TLS (via the Globus Grid Security Infrastructure). Media data flows are conceptually peer to peer. That is, media traffic is generated and injected into a communication group, and clients extract data from the communication group for local consumption. This traffic uses multicast (where possible) or unicast (through multicast/unicast bridging technology). Control of distributed presentations is client/server between the presenter, the presentation service, and the presentation clients. Data transfer between clients and the venue server uses HTTP over Globus-secured TCP.

In the case of DAE, the connection topology may consist of a mix of client/server connectivity for centralized components and peer-to-peer connections including group multicast among the collaborating members. In most applications of this scenario client/server connections are avoided and a fully distributed peer-to-peer model is employed, however, for the synchronization, authorization and discovery services a client/server architecture may be used. The sequence of connection establishment may vary and there are two basic models. The first is a model that establishes secure channels between communicating parties, which follows the general schema of a signaling phase followed by authentication and secure channel creation before the transmission of data. It is likely that such connections, due to their setup overhead are persistent for the duration of a collaboration session. The second model neither establishes nor relies on a secure channel; messages are sent independent of each other and contain all necessary information to ensure their authenticity, integrity, and confidentiality.

In the case of CE, small groups of clients share a VNC server instance to collaborate. The communication is via TCP. Electronic notebooks are also in use and these are based on a client-server architecture.

In the case of RV, custom solutions are generally in use and these often use a client/server model with multiple clients. TCP is most commonly used in today's systems, although this may change over time for certain scenarios where something like reliable UDP performs better. Control information needs to be reliable.

3.7.2 Sequence in which connections are made

In focus areas where client/server architectures are in place, the first connection is usually a TCP connection to the server. This connect is used to establish information about the session at the client and information about the client at the server. This connection also allows the client and server to authenticate with each other and is usually the long-lived primary information channel between the client and server. In most of these applications, when the initial connection with a new client is established, messages are sent out to the other clients informing them of the presence of the new client and they or the server then send their information to the new client in return. The duration of each connections can range from minutes to days.

In the case of TI and RV, a UDP or multicast connection is established next to send the data. Additional TCP connections are also established to send updates and connect to other resources such as data servers.

In the case of AG, media connections are then established via IP Multicast or bridged UDP and are peer-to-peer. Distributed presentations require a separate server which all other sites connect to as clients using long-lived connections.

In the case of DAE, when client/server is not in use the first connection is the joining of a group communication channel. Subsequent connections are made via either reliable group

communication or TCP/TLS. Connections to servers are normally via TLS and connections to peers are normally through reliable group communication.

3.7.3 Protocols (e.g. TCP, UDP, Multicast)

All of the focus areas make use of TCP and UDP unicast communication. Many of the applications also make use of a security layer for their communication such as TLS or GSI. Most of the application areas also make use of IP Multicast (e.g. for streaming media and wargame simulations). When these multicast streams are secured, it is generally through encryption using a shared key. In addition, a few of the focus areas are making use of reliable multicast mechanisms and are interested in using secure multicast communication channels as they become available.

4. How Security Properties Apply to the Implementation

4.1 Assets that require protection

In the case of the AG, access to the Venue and the data, media streams, and membership associated with the Venue need to be protected from unauthorized users. In the case of CE, expensive experiment equipment must be protected. Also, data, computers, etc. on internal facility network (i.e. resources not part of the experiment) need protection. In the case of RV, the data needs to be protected.

In the case of DAE, we have both temporary and long-term data. This data is divided into the data about the collaboration itself (e.g. user lists, locations, tools in use, etc), the data produced by the collaboration (e.g. conversation archives, activity logs, etc.), and data that is the subject of the collaboration (e.g. shared files, drawings, etc.). All of this data will need to be protected. Some of this data is created dynamically during a session and protections and configuration mechanisms need to be available to regulate access "on the fly." The minimum level of protection expected is that only users authorized to access the collaboration can read, write, and create this data. Permission to delete the data is likely restricted to a smaller set of users particularly in the case of data about the collaboration and produced by the collaboration. At an enhanced level of security, a user creating any particular data could dynamically specify more restrictive access rights to that data.

The DAE focus area will use a combination of secure and insecure data and control channels. Ideally this can be decided for each transmission with a configurable default for the collaboration. Distributed Data may be stored in distributed data storage facilities. In an enhanced security environment, an audit trail for each data file would need to be kept in a secure way so that different versions of the same file are recognizable and their history, authenticity, and modifications can be identified and verified. The names and addresses (e.g. eMail, IM) of collaborators require protection so that they are not exposed to external entities (i.e. spammers) even if the data itself does not need protection. The communication media (lists, channels, etc) also require protection from unauthorized access.

4.2 Need for extensible and dynamically definable access rights

All of the focus areas except CE require some ability to extend rights and define them dynamically. In the case of TI, access to data and collaborative spaces needs to be extensible. In the case of AG, the identities of users of the AG system can be widely changing. The dynamic nature of the DAE area (in terms of participants and resources) requires that rights need to be extensible, inheritable, and dynamically definable. In the case of CE, human interactions are sufficient to manage scheduling, etc., instrument software enforces signed user access (a second user cannot log in and change a running experiment)

4.2.1 Access (read/write) to stored data

All of the focus areas require some level of protection from read/write access to files. In the case of TI, data that is posted to the central server is readable by the group. Write access should be limited to the author of the data and the administrator unless access is granted to other users. In the case of the AG, the data stored in a venue is shared by the users in the venue. AG requires

access controls to both limit access to venue users and to limit the destruction or modification of data files by non-owners of the data. In the case of DAE, the files are shared from local systems with the local user controlling access to the local files for other collaborators. Access to archives of sessions also needs to be controlled. In the case of CE, the file system access restrictions are leveraged.

4.2.2 Access (execute/modify) to shared applications

In most of the focus areas there is a component of shared applications. The integrity of these applications is critical to the environment. The different focus areas use different means of providing this. For example in TI, it is assumed that all participating sites already have access to the programs needed to participate in the collaboration. Source code may not, however, necessarily be available to everyone. In the case of the AG, it is unclear exactly how shared applications which may have components that are distributed via the venue shared data mechanisms fits into the picture. There is certainly a need for the verification of the provenance of this code and management of the trust in executing potentially unknown code.

4.2.3 Communication channels (send/receive)

In all of the focus areas control channels need to be protected so that traffic from an unauthorized source is ignored. In all the areas except TI authentication and encryption of control channels is also necessary. In all of the focus areas the ability to encrypt data channels is needed. This capability might not be used in situations where the security requirements do not indicate the need for encryption and should be configurable for each collaboration space. Authentication of the senders of media data is important.

4.2.4 Mailing lists/addresses (read/send)

In the TI focus area mailing lists also require protection. All members of the group have read access to the mailing list. Outside users do not have access to the list.

4.2.5 Access to transient resources (e.g. collaboration sessions)

In all of the focus areas there is a need for protection of transient resources. Collaboration sessions themselves are often transient. All members of the group should have access to appropriate collaboration sessions. Often sessions are persistent, but transient sessions to foster spontaneous collaboration are required. Non-group members do not have access to the collaboration sessions. In the case of DAE, authorized users also need to have a way of temporarily authorizing unauthorized users into restricted areas.

4.2.5.1 Special vulnerabilities for this scenario

In the case of RV and TI, interactivity and security are often incompatible goals forcing choices. Critical data streams need to be encrypted but the vulnerability of other data streams will depend on the security required by the application. In the case of DAE, denial of service is the greatest concern. Users will revert to less secure means rather than be denied access to collaborators.

4.3 Estimated value/cost to recover compromised systems

This section attempts to articulate, for each asset (data, code, communication, etc), the estimated value in terms of cost to recover the asset if it is lost, damaged, or compromised. These are formulated as a set of questions.

4.3.1 Anticipated losses in terms of intellectual property control and applicable regulatory penalties

If sensitive information is inappropriately disclosed, what regulatory penalties may apply? What losses are anticipated in terms of intellectual property control?

In general collaborative data is not subject to regulatory controls (i.e is not sensitive or classified, etc.) IP rights and laboratory reputation are at stake. However, since the use of the data is usually to create a peer-reviewed publication, it is likely that any data theft will be discovered and dealt with through standard scientific ethics procedures. In the case of RV, data could be sensitive either in a classified sense or proprietary, though classifying loss is situation dependant.

In DAE the value of each asset and the cost to recover it is difficult to estimate since these items are dynamically created. Their value will also be determined only at the time the item is created. It is likely that in most cases the value will be relatively low and the cost to replace it will not be excessive.

4.3.2 Cost associated with lost or damaged data

If data is lost or damaged, what are the costs associated with recovering it? If preparations need to be made to detect such loss or damage, and allow recovery (i.e. backups, checkpoints, etc), what are the costs to deploy and operate such recoverability means?

Most collaborative data comes from other sources (scientific data archives, computational simulation). The collaboration is mostly targeted at understanding and insight, not necessarily the creation of data. This will not always be the case but it will often be the case. Thus checkpoints and data archival of the data sets are assumed to be outside of the scope of these environments except in the case of the data that defines the persistent collaboration spaces. Persistent collaboration spaces should be checkpointed to disk regularly or otherwise archived at the server so state can be restored if the server goes down. Checkpointed data should be backed up. If a persistent collaboration space state is lost then the data from the collaboration will need to be reacquired and the collaboration will need to be repeated. There is a chance that important insight will be lost through the loss of state. Reacquiring the data may be very expensive as the data may have been created through computational simulation. If the data is not archived as part of the simulation then the computation will need to be repeated.

What costs are expected in terms of time and effort to reschedule and reprocess data?

This varies by focus area. If the underlying data that relates to a large simulation or is the results of an experiment is lost, the cost can be several weeks or months to recreate the data and some data may not be impossible to recreate if, for example, it is the result of a computational steering exercise or there are system time constraints. In the case of DAE, most of the data is a record of the collaborative interactions that took place and it can't be recreated.

4.4 Anticipated Threats

4.4.1 Disruption from random hacking

In all of the focus areas, hackers are a concern. Malicious or accidental disruption of collaboration sessions is the main threat. Ability to prevent unauthorized access and limit access to known users is the main goal. The data in many of these collaborations is not ready for public release so privacy is important. Also, corruption of data and unauthorized creation and deletion of data are concerns.

4.4.2 Explicitly excluded colleagues

In all of the focus areas, this threat is of less concern than other threats. Generally, simple protection mechanisms and peer pressure can be employed to keep this threat low. In the case of TI, physical presence is required to participate in the collaboration.

4.4.3 Outsiders who steal software, data or compute cycles

In all of the focus areas, this threat does not tend to be treated as separate from hackers breaking into the system. In the case of CE, outsiders that want to access lab resources are a concern.

4.5 Security features needed and protected items that need them

4.5.1 Integrity - messages and data can't be secretly modified

In all of the focus areas integrity of messages on the data and control channels is important. In the case of DAE, this is less of an issue but still important.

4.5.2 Confidentiality - protection from disclosure to unauthorized persons

4.5.2.1 Duration of confidentiality

For all of the focus areas, if a group is established where the initial parameters indicate that the data should be kept confidential then it should remain confidential at least as long as the persistent collaboration space exists. In the case of CE, the data should be kept confidential forever or as long as the data is retained at the facility (in practice this time period is on the order of a year). If a group desires the data to be public then a separate copy can be created that is accessible from some alternative service. In the case of DAE, new users might need to be authorized in to access previous interactions they were not authorized to access at the time.

4.5.2.2 Confidentiality of persistent data and data in transit

In all of the focus areas, if the group establishes that its data should be kept confidential then it should be confidential- including data in transit. Outsiders who are not part of the group should not be able to retrieve the data. In the case of DAE, it is likely that encrypted channels and secure data storage locations will provide an adequate solution. Length of time that data should remain confidential is likely to vary by collaboration and type of data.

4.5.3 Authorization, access control - unauthorized users are kept out

In all of the focus areas, authorization of individuals is important. Users connect to a persistent collaboration space and authorization occurs at the same time as the sign-on authentication. Subsequent user accesses to datasets within the collaboration space require additional authorization decisions at sporadic times during the collaboration. Several of the focus areas also plan to adopt role and group-based access control mechanisms to generalize the process.

In most of the focus areas authorization policy for entry to the environment is static and defined in advance to last a relatively long duration (days to years). For instance, in the AG, a venue is configured to either be open, or to be closed except to a specific list of individual identities. Subsequent authorization to other resources in the environment is granted to those individuals who have obtained entry to the venue.

In the case of DAE, authorization policy needs to be highly dynamic to allow for the frequent changes in membership of an ongoing collaboration as well as to allow for ad-hoc creation of new collaborations. Typically collaborators have to be authorized to join a collaborative session when they attempt to enter the session as well as when they request additional services based on elevated privileges (i.e. change of policy during ongoing collaboration). The granularity of authorization is in most cases down to the individual entity requesting the service or originating the message.

4.5.3.1 Should all the collaborators know who is authorized or just the administrator

The question of whether collaborators should know who is authorized or whether only the administrator should know depends very much on the focus area and the setting. In the case of TI, RV, and CE, only the system administrators and facility staff need to know who is authorized. In the case of AG and DAE, this requirement varies with each use. Some usages will desire that only the 'owner' of the session know who is authorized and others will demand that the authorization information and decisions be open to all currently authorized users.

4.5.4 Authentication - assurance of identity of persons

A variety of mechanisms are in use today by these focus areas to provide authentication of participants. These include username/password, secureID token, and X.509 identity certificates. In many cases, the primary authentication is used to establish the initial connection and this connection provides all data transmission. In the AG, authentication is performed each time a GSI/SOAP message is sent, since connections are not cached between SOAP calls. In the CE, establishment of a SSH or VPN connection is used to authenticate remote users in combination with username/password. In the case of DAE, authentication mechanisms must also allow new users to be dynamically added to the environment very quickly. Support for multiple methods of authentication into the environment for users is essential.

In all of the focus areas users sign on as individuals by presenting their credentials but in some cases this then provides access to a group or site authentication method. Signed applets and server certificates are in use in some focus areas to authenticate applications and servers. Single-signon is important and in some cases critical to all of these focus areas. Incorporation of role-based authentication capabilities is planned for the near future in a few focus areas.

4.5.5 Auditing - permanent log kept for accountability and possible error recovery

In all of the focus areas, accountability for access to the persistent space and data sets is important. Auditing can be used to both enforce accountability of individuals as well as provide awareness to other users as to who has done what in the collaboration space. In the case of casual use of the AG, auditing is not required.

4.5.6 Non-repudiation - originator of data can't deny it later

Non-repudiation is not presently critical in any of the focus areas. Auditing will likely be the preferred method of tracking user actions and it is useful (but not critical) if the audit trail provides information that supports non-repudiation. In the case of CE, many uses of electronic notebooks will, however, require non-repudiation.

4.5.7 Survivability - recovery from attack or failure

For all of the focus areas, this is not currently a major concern. In most of the focus areas, the concentration will be on check-pointing and other mechanisms to maintain the current state on a server so that the state can be restored on recovery.

4.5.8 Availability – performance and access

Persistent collaboration spaces and servers should be available at all times. Access should be provided to any authorized user. The collaboration will only function well if collaborators seldom encounter denial of service due to failure of the system and security mechanisms. In the case of TI, a small (<10) number of users should be able to be supported at interactive speeds (<200ms latency) [Park99].

5. Performance requirements

5.1.1 Which resources are precious? (e.g. CPU, network, disk)

In the case of TI and RV, the CPU, graphics hardware, immersive spaces, networking, and disks are precious resources. These are generally data intensive applications which are often using the hardware to its current limits. In the case of the AG, the largest single user of resources in a typical AG session is the computer performing decoding and display of the multiple network streams. Sufficient network capability is also crucial, especially as the number of participants in a session grows. In the case of CE, the disks are the most precious resource. In the case of DAE, the people are the most precious resource.

5.1.2 Latency

Latency largely affects the usability of collaborative environments. Most of the focus areas include real-time interaction with another human or a computer and have difficulty with latencies which reach above 200ms (e.g. an AG audio stream or a TI interaction). Latency is not an issue for bulk data transfers and button clicks on the screen, unless the button clicks are to control remote interfaces.

5.1.3 Overhead due to encryption

Many of these focus areas already make use of encryption technology for their data and current levels of overhead and latency for encryption (e.g. AES) have been tolerable. If future encryption schemes were to add significant overhead, then applications that already consider the CPU or network (especially low latency) to be a precious resource may require specialized hardware such as cryptographic accelerators to offload the main processors in order to make data channel encryption possible.

5.1.4 Throughput (e.g. transfer uses 1Gb)

In the case of TI, large data transfers are required to reach a consistent state and when data sets are updated. State updates are on average 1530kbps per stream (assuming 30fps 16 floats for a 4x4 transformation matrix). In the case of the AG, the video data imposes the largest requirement for throughput. Typical use is 250 Kbps per stream, with four streams per site and 3 or more sites participating. In the case of CE, the throughput required is a few Mbit/sec for conducting experiments, but data can be gigabytes in size so greater bandwidth is useful during file transfer. The requirements of RV range from bits to gigabits per second.

Author Information and Role

Deb Agarwal Lawrence Berkeley National Laboratory DAAgarwal@lbl.gov	Brian Corrie Simon Fraser University bcorrie@sfu.ca	Jason Leigh Electronics Visualization Laboratory University of Illinois at Chicago spiff@evl.uic.edu
Markus Lorch Department of Computer Science Virginia Polytechnic Institute and State University mlorch@vt.edu	Jim Myers Pacific Northwest National Laboratory Jim.Myers@pnl.gov	Robert Olson Mathematics and Computer Science Argonne National Laboratory olson@mcs.anl.gov
Michael E. Papka Mathematics and Computer Science Argonne National Laboratory papka@mcs.anl.gov	Mary Thompson Lawrence Berkeley National Laboratory MRThompson@lbl.gov	

Glossary

None

Intellectual Property Statement

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director.

Full Copyright Notice

Copyright (C) Global Grid Forum (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE GLOBAL GRID FORUM DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

References

- [Disz97] Disz, T.L., Judson, I., Olson, R., Stevens, R., "The Argonne Voyager Multimedia Server", IEEE Computer Society 6th High Performance Distributed Computing Conference, Portland, OR, 1997.
- [Germans01] Germans, D., Spoelder, H. J. W., Renambot, L. and Bal, H. E., "VIRPI: A High-Level Toolkit for Interactive Scientific Visualization in Virtual Reality", Proc. Immersive Projection Technology/Eurographics Virtual Environments Workshop (IPT/EGVE), May 16-18, Stuttgart, Germany, 2001.
- [GW] GeoWall : www.geowall.org
- [He02] He, E., Leigh, J., Yu, O., DeFanti, T. A., Reliable Blast UDP : Predictable High Performance Bulk Data Transfer, Proc. IEEE Cluster Computing, Sept, Chicago, Illinois, 2002.
- [Leigh97] Leigh, J., DeFanti, T., Johnson, A., Brown, M., Sandin, D., "Global Tele-Immersion: Better than Being There,," proceedings of ICAT '97 Tokyo, Japan, Dec 3-5, 1997
- [Leigh99] K. Park, Kenyon, R., Effects of Network Characteristics on Human Performance in a Collaborative Virtual Environment, Proceedings of IEEE VR '99, Houston, TX, 03/13/99-03/17/99.
- [Leigh01] Leigh, J., Yu, O., Schonfeld, D., Ansari, R., et al., Adaptive Networking for Tele-Immersion Proc. Immersive Projection Technology/Eurographics Virtual Environments Workshop (IPT/EGVE), May 16-18, Stuttgart, Germany, 2001.
- [Park00] Park, K., Cho, Y., Krishnaprasad, N., Scharver, C., Lewis, M., Leigh, J., Johnson, A., CAVERNsoft G2: A Toolkit for High Performance Tele-Immersive Collaboration, Proceedings of the ACM Symposium on Virtual Reality Software and Technology 2000, Oct 22-25, 2000, Seoul, Korea, pp. 8-15.
- [Stevens03] Stevens, R., Papka, M.E., Disz, T., "Prototyping the Workspaces of the Future," IEEE Internet Computing, 7(4), July/August, 2003, pp. 51-58.