

PSFC/RR-06-02

**Recommendations for Remote Conferencing
for the ITER Project**

M. Greenwald¹, D.P. Schissel²,
G. Abla², and J. Stillerman¹

¹Massachusetts Institute of Technology,
Plasma Science and Fusion Center

²General Atomics, Fusion Energy Research Group

**Plasma Science and Fusion Center
Massachusetts Institute of Technology
Cambridge MA 02139 USA**

General Atomics reference number **GA/WP 261-001**

This work was supported by the U.S. Department of Energy, Grant No. DE-FC02-99ER54512. Reproduction, translation, publication, use and disposal, in whole or in part, by or for the United States government is permitted.

**RECOMMENDATIONS FOR
REMOTE CONFERENCING FOR
THE ITER PROJECT**

M. Greenwald,¹ D.P. Schissel,² G. Abla,² and J. Stillerman¹

¹Massachusetts Institute of Technology, Plasma Science and Fusion Center

²General Atomics, Fusion Energy Research Group

Technical Contact: Martin J. Greenwald (617) 253-6053 g@psfc.mit.edu	Technical Contact: David P. Schissel (858) 455-3387 schissel@fusion.gat.com
--	---

January 5, 2006

1. Introduction

This white paper proposes an evolutionary approach to remote conferencing for the ITER project. The emphasis is on reliable, low-cost commercial, internet-based solutions that can be deployed today within an architecture that can leverage future developments. The aim is to assist the ITER project in the timely completion of its goals through expedited information sharing. We have kept in mind that the international partners and institutions represent a heterogeneous user base, with diverse infrastructure and differing tolerance for capital costs. Policy decisions will also affect technical choices (e.g., the requirement to support a range of computer platforms or the need to share non-electronic presentations) and will need to be made by the ITER project organization. Our recommendations are based on what we believe are “reasonable” assumptions about project policies.

The remote conferencing recommendations include audio/video, document sharing, instant messaging, and an internal project web site. Individuals and organizations within the project will face expenses ranging from very low, to equip individual offices, to moderate, to fit-out large meeting rooms or lecture halls. Our solution scales in cost to meet the capability and needs of the organization without any degradation in remote meeting quality. The largest capital expenditures, while still modest, can be done on a one per ITER partner, or per continent basis. In fact, we have found solutions where these expenses are mitigated or eliminated by sharing existing resources with other science communities. The appendices provide details on various product options allowing the project to understand the trade-offs. It is clear that as these technologies continue to evolve, the “best” solutions today will not be optimum in the future. Fortunately, much of the hardware investments made today would continue to function with future software and services. Still, costs, especially human costs, for transitioning between technologies are not negligible and we attempt to recommend solutions with significant longevity.

Finally, it is important to note that the solutions described here support normal “business” activities. The methodologies for remote participation during ITER tokamak research operations have not been considered in this document. We expect that these will be considerably more challenging and urge the ITER organization to support a continuing activity intended to define a truly collaborative control room.

2. Findings and Recommendations

Reliable, low-cost videoconferencing technologies are now available and widely used in scientific and commercial endeavors. There are no fundamental technological or economic barriers to deploying these in support of the ITER project.

There are four basic components or functions for remote conferencing systems.

1. Venue hardware – cameras, microphones, speakers and computers installed at each location, where remote participation is needed.
2. Venue software – client software and protocols that code the video and audio signals into Internet data packets.
3. Multi-point capability – hardware or software that allows more than two participants or venues to participate in the same meeting.
4. Slide or document sharing capability.

For deployment of a basic videoconferencing system today, we recommend low-cost commercial products based on the Internet H.323 standard. By attaching a camera, microphone and speakers to a standard desktop computer and installing inexpensive (or free) client software, a full-functioned video conference system can be assembled easily. For the PC platform, PVX software (from PolyCom at \$150 US) is widely used and has been extensively tested. For MacIntosh and linux, OhPhoneX and GnomeMeeting are free, open-source client software for their respective platforms. Other client software packages exist, but need to be thoroughly tested for compatibility before they can be recommended to users. Further details can be found in Appendix B.1 and Table B-1. For large meeting rooms, it can be advantageous to install dedicated hardware CODECs (Coder/Decoder), costing \$5,000-10,000, in place of the computer/software combination, however this need not be part of an ITER deployment. Sites that have these systems can expect full interoperability with desktop systems.

For H.323 based conferencing, multi-point capability is best accomplished by dedicated hardware multi-point control units (MCUs). These bind multiple callers into a single meeting and can also connect with participants who are using an ordinary telephone. In our experience, the best MCUs are currently made by Codian. While these are expensive devices (>\$30,000), we should be able to make use of hardware already deployed in the research and education community. We have identified and gotten permission to use systems currently deployed at LBNL in California and at the IN2P3 in France. (As discussed below, a third facility in Asia needs to be located.)

We note that sometime in the near future the developers of VRVS will release a production version of EVO, an integrated conferencing utility. It could augment or even replace the H.323 systems for some small meetings. We will continue testing this software and evaluate the production version for reliability, quality, ease of use and interoperability. One of the potential advantages of EVO is that the multi-point function is implemented in software, eliminating the need for expensive hardware MCUs. EVO is not based on industry standards, is somewhat more difficult to use, does not scale well to large venues and may require more support. For these reasons, and because the recommended investment in H.323-specific components is small, it should not be considered a substitute to a general deployment of H.323 systems.

For sharing slides or other presentation material, such a clear-cut recommendation is not possible. The most functional system that we have employed comes from WebEx. Slide sharing is a component of its full-featured, integrated conferencing utility. As the name implies, WebEx uses web protocols and operates out of a browser with no software installation required. Impact on site security should be minimal. However, the cost for WebEx may be an issue. The ITER project would need to spend at least \$2,000 a month for the service. (There is considerable interest in this product from the High Energy Physics community, so some sort of cost-sharing arrangement might be negotiated. We have not pursued this avenue.) A second option is VNC, which is open source and free. However VNC has relatively low performance and may not be acceptable at some sites due to security concerns. A third option is to use tools based on the Internet H.239 standard. These look promising but are currently available only from a single vendor and cannot be recommended at this time. Further options exist, but are not recommended either for a variety of reasons. See Appendix B.2 and Table B-2 for further details.

A useful adjunct to these conferencing tools would be an “open” web site for project workers. We recommend using the open source Wiki software, which runs on a central server and allows authorized users to freely create and edit web content through their browsers. Wiki supports shared text, images and media files in a persistent but flexible format. Text-based Instant Messaging (IM) technologies could provide another lightweight collaboration channel for the ITER project. It is more immediate than e-mail, less intrusive than a telephone call, and can be used as a collaboration tool in non-formal meetings. Jabber-based IM is the preferred technique for the ITER organization over proprietary IM systems, since ITER can deploy a Jabber-based open-source server and have full control over user authorization, management and security.

3. Next Steps and Concluding Comments

This white paper describes an approach to remote conferencing for the ITER project; it is not an implementation plan. If the project accepts this approach, or something like it, development of an initial plan is the logical next step. We would be pleased to help in the development of the plan and to assist the project in its implementation including leveraging some resources from the USDoE National Fusion Collaboratory Project. Noting that in some areas research and testing will be needed on an ongoing basis, we would also be available to help qualify additional hardware and software as the deployment proceeded.

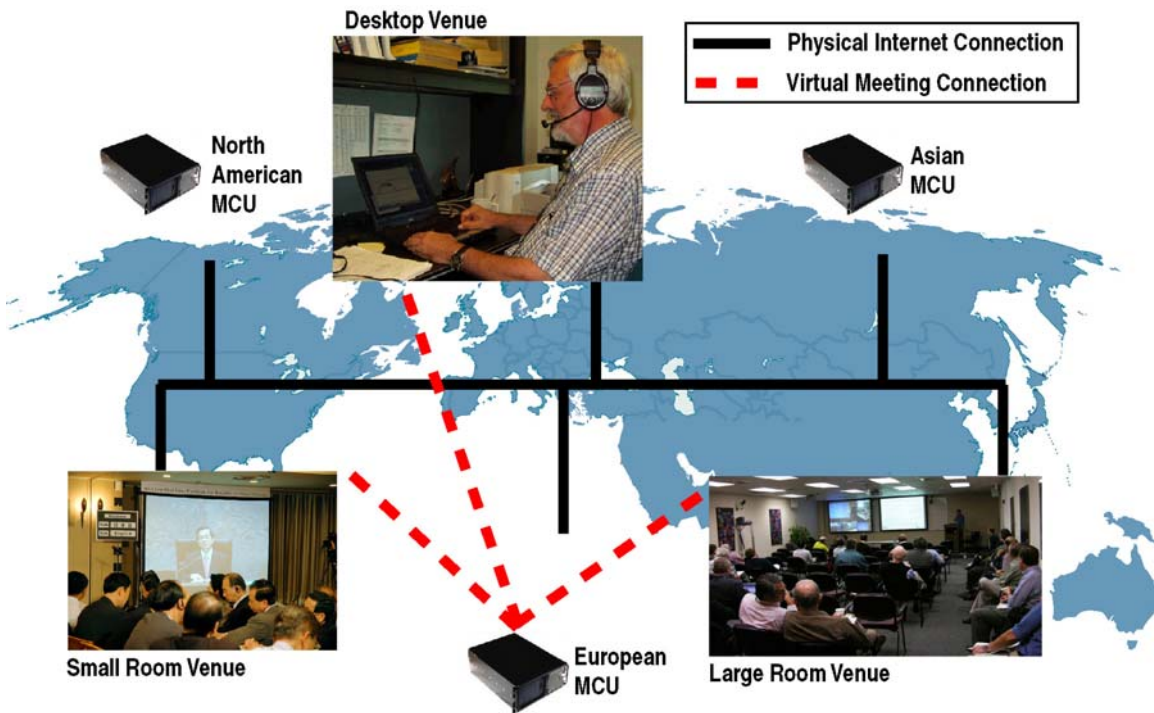


Fig. 1. Utilizing the H.323 protocol with dedicated MCUs will enable remote conferencing for the ITER community from desktops and meeting rooms. While all devices connect physically to the Internet, the MCUs create a virtual “hub and spoke” conferencing network for each meeting.

Appendix A. General Considerations

A.1. Multi-Platform Issues

The worldwide scope of the ITER project, combined with its in-kind and cash-based funding, likely precludes the organization's ability to dictate to all parties, institutions and contractors, that only one computer platform will be supported. So it is our assumption that any solution offered to the users (clients) in this paper be designed for multiplatform environments. In our context, multiplatform means Microsoft Windows, Apple OS X, and Linux since the vast majority of client computers that users have on their desk or carry with them are of one of these types.

Although we assume that client systems must be designed for a multiplatform environment, the server systems that create the conferencing infrastructure might be standardized. Examples of such infrastructure include fixed setups for large conference rooms and reflecting computers for audio and video. These are systems that the users rely on, but never interact with, and are also orders of magnitude fewer in number compared to client computers. Such standardization could be a decision made on technical grounds by IT professionals. As solutions for ITER are designed and deployed, such server standardization should be attempted where possible, as it would benefit system reliability and simplify support.

A.2. Support Issues

Some level of support will be required for any deployed remote conferencing system. It is clearly desirable to keep support costs low as organizations may choose to abandon a technology if the manpower lost in support outweighs the manpower saved by using it. The most critical area of support is for the basic infrastructure, that is, servers, MCUs or reflectors. It does not seem practical for a worldwide effort like the ITER project to derive that support solely from the central site. Time zone issues alone would make this problematic, even if the workload was not overwhelming. Instead, experts for each party or from each continent should be identified for infrastructure support. This distributed set of experts would work together to insure a smoothly operating system.

The next area of support is at the user level. A distributed set of experts (who might overlap with those above) should be tasked to provide adequate user support and consultation in the form of telephone, email or instant message. In time, with more users participating in remote meetings, some will grow into power users, who can become part

of an informal help team. The central team should create a set of web pages with up-to-date information on ITER standard practices, instructions for participating in meetings, troubleshooting advice, frequently asked questions and links to support personnel. Help for setting up a venue could be provided, with suggestions for hardware and links to vendors and download sites. If desired, the site could maintain a time-zone aware schedule of meetings, dialing information and be linked to project-wide calendars and directories.

A.3. Security Issues

Security for a distributed organization like ITER falls into two broad categories, site security and resource security. The first involves protecting a site's LAN from attack or malicious intrusion from the outside and is usually implemented by software or hardware firewalls at the site perimeter. The second entails management of computational resources in a secure fashion through authentication and authorization of remote (and local) users. Authentication is the process, in a computerized transaction, that gives assurance that a person or computer acting on a person's behalf is not an imposter. Authorization is the process of determining, by evaluating applicable access control information, whether an identified person is allowed to have the specified types of access to a particular resource. If required, an additional level of security can be provided by data encryption, which ensures the privacy of network communications. These security measures should be developed and implemented in concert, taking cognizance of the risks and benefits inherent in network based applications.

For video conferencing, we restrict our comments to H.323 based tools. The H.323 protocol is designed in such a way that the communication ports used for audio, video, and control are dynamically negotiated between two clients. Further, this port negotiating covers the entire upper port range (1024-65535). Since ports are negotiated on the fly, a firewall cannot be configured ahead of time to open specific ports for audio and video streams. The requirement to open a large range of ports to the outside world is often resisted by security managers and may violate security policy at some sites. There are three possible methods to solve this problem. The first is to utilize H.323 aware intelligent firewall products. These firewalls can detect the H.323 signaling requests and take appropriate action to allow the traffic to traverse the router or firewall. This approach may require a significant capital investment if new hardware is required. The second solution involves the use of an H.323 proxy server. In this case, all the H.323 traffic in the network goes through the proxy, and firewall ports are opened only for the proxy. This method increases the overall system complexity and takes time to configure and test. However, the decrease in risk may warrant the effort. The third approach is to restrict H.323 access to a part of the site's network that has relatively low perimeter protection. This typically would mean that H.323 conferencing would only be available

from designated conference rooms, not from staff offices. Resource protection under H.323 is provided by a gatekeeper, an essential component of the protocol. Approval of individual venues is possible, as each is required to register (once) with the gatekeeper. In addition, each meeting can be password protected. Capability for encryption is becoming available in some of the newer clients.

Recommended web based tools like WebEx and Wiki run from a standard browser and do not present security problems any greater than those posed by the web itself. The Jabber IM system presents minimal security risks and is a firewall friendly protocol utilizing secure socket layer (SSL) for encrypted communications. Both incoming and outgoing messages as well as presence information uses this single port. Jabber also supports HTTP binding that can be utilized by users behind strict firewalls. Overall, messaging systems like Jabber present risks similar to email. IM worms can spread rapidly, so properly patched server software and anti-virus protection on client systems is important.

Appendix B. Specific Technical Issues

B.1. Video Conferencing

B.1.1. User agents (client applications) and protocols

Conferencing tools can be based on standard Internet protocols (H.323, SIP) or, like EVO/VRVS and the Access Grid, on non-standard protocols. Table B-1 summarizes some of the features, advantages and disadvantages of commonly used user agents. It includes both commercial and open source offerings. For our community, solutions are needed for a range of venue types from desktop to auditorium, and for a range in the technical sophistication of the users. In general, the larger the venue and the less sophisticated the user, the more attractive commercial tools are. Where possible, we favor solutions that accommodate heterogeneous tool sets, though there can be a cost to doing so in terms of robustness and support.

The ESnet community, which covers a majority of physical science research in the U.S. and its international collaborators, has been successfully using H.323 video conferencing for many years. A replacement for earlier ISDN based systems (which were heavily used for more than ten years), this technology is mature and easy to use. Almost no meeting time is spent in setup and debugging. Video quality is good, with lip-synced audio. Most venues employ commercial user agents connected through centrally located and administered hardware MCUs. Users can choose whether to see all participants in a meeting or just the current speaker. Large sites typically have several rooms with dedicated hardware to support video-conferences, but participation by individuals with desktop systems is widespread. Integration between different sized venues is straightforward and reliable. Entry costs for participants are very low and software for the PC, Mac and linux platforms is available. ITER should be able to use existing MCUs, which are the only significant capital cost. We have already gotten approval for ITER to use the Codian MCUs at the ESnet site in Berkeley, California and at IN2P3 in France. It should be possible to locate a facility in Asia that would make a similar offer. Because the MCU hardware requires local support, the global nature of the ITER project suggests that one MCU per continent is required. Meetings would then use the MCU in the time zone where local support was most available. Aside from the potential cost for the MCUs, a disadvantage of H.323 is that it requires many ports to be open in site firewalls and deployment at some highly secure sites may pose some difficulties (Appendix A.3).

Table B-1. Comparison of Selected Remote Conferencing Technologies

Type	Commercial			Open Source		
	H.323 Based Hardware CODEC	H.323 Based Software CODEC	Web Conferencing	Internet Based Video Conferencing	Large Format Audio/Video Conferencing	H.323 Based Software CODEC
Vendor/ example	Polycom, Tandberg	Polycom/PVX	WebEx LiveMeeting	VRVS/EVO vrvs.org	AccessGrid accessgrid.org	OhPhoneX GnomeMeeting
Platforms	Self- contained	Windows	Windows, Mac, linux	Windows, Mac, linux	Windows, Mac, linux	Windows, Mac, linux
Price	\$5K-\$15K	\$150	Monthly fee \$2K-\$5K	Free	Free	Free
Infrastructure price	HW MCUs \$2KK/port (can use existing)	HW MCUs \$2KK/port (can use existing)	Fee-based hosting	Need reflec- tors on each continent ~\$2K/per	Add venue server, unicast bridges ~\$2K/site	HW MCUs \$2K/port (can use existing)
Support cost ¹	Low	Low	Low	Medium	High	Medium
Recommended venues ¹	Meeting rooms	Offices, small meetings	All	Offices, small meetings	Meeting rooms	Offices
Reliability/ease of use ¹	High	High	High	Medium	Low	Low-medium
Current availability	Yes	Yes	Yes	VRVS-Yes EVO-No	Yes	Yes
MCU	Separate commercial	Separate commercial	Hosted	Included free	Included free	Separate commercial
Multiple video streams ²	No	No	Hosted	Yes	Yes	No
Slide sharing	Limited (H.239)	Limited (H.239)	Yes	Limited	Limited	No
Instant messaging	No	No	Yes	Yes	Yes	No
Presence ³	No	No	No	Yes	No	No
Authentication ²	Optional ⁴	Optional ⁴	Yes	Yes	Yes	No
Site security impact ⁵	Yes	Yes	No	No	Yes	Yes
White board	No	No	Yes	No	Yes	No
Protocols:						
H.323	Yes	Yes	No	Yes	No	Yes
Mbone tools (vic & rat)	No	No	No	Yes	Yes	No
SIP	Some		No	Yes	No	Some
H.239	Some	Some	No	No	No	No

¹Based on experience of authors.

²Use restricted by password or certificate base log-on.

³Presence = current status/location of users (available, busy, out to lunch, ...).

⁴Some clients support encryption. Gatekeepers can restrict access to authorized clients.

⁵Based on need to open site/computers to additional internet ports or protocols.

A variety of H.323 client software is available for different platforms. Polycom PVX is a low-cost commercial product for the PC. It features good quality audio and video and excellent compatibility with H.323 hardware devices. It is easy to install and use and is widely deployed in the ESnet community. GnomeMeeting is free, open-source software based on the H.323 protocol. It is available for PC and linux platforms and interoperates with commercial H.323 systems and MCUs. It supports all modern videoconferencing

and VoIP feature. Overall audio and video quality are somewhat lower than commercial systems, and the software is more difficult to use and less robust. OhPhoneX is a free H.323-compatible video conferencing client for Mac OS X. OhPhoneX is very stable and has good interoperability with other H.323. This client integrates with native Macintosh applications, such as AddressBook or Apple Bonjour.

Some members of the physics community have been using VRVS, a community written and supported collaboration suite for many years. VRVS will be replaced by a new tool, EVO, developed by the same group. One goal of EVO development is to improve the reliability and ease of use compared with VRVS. EVO is scheduled to be in beta testing in January 2006 and in production by July 2006. Initial testing with EVO looks promising. These tools are not commercial and do not use standard Internet protocols, but it is likely they will be supported and viable for many years. EVO does not require an expensive hardware MCU, instead, multipoint conferencing is provided by a network of “reflectors” that are implemented on standard computers distributed throughout the community. All participants in a multi-point conference can be viewed simultaneously in separate windows. Advantages of EVO/VRVS include flexibility, expandability and extensibility. New protocols can be added by “transcoding” in the reflector software. EVO and H.323 can interoperate, but we have not assessed the robustness and reliability of this capability in a production environment. Audio and video are not quite as good as the commercial H.323 products and reliability is not quite as high. The end-user software is somewhat more complicated to use, since it puts up several windows with control functions distributed between them. EVO security is good; all traffic goes through a single UDP port. (If UDP is not available, a single TCP port is used.) A simple password authentication scheme is included, and there are provisions for hooking into community provided authentication schemes. NAT and firewall traversal should not be an issue. The tool provides presence, instant messaging, as well as audio and video conferencing. Support costs are likely to be higher than for commercial H.323.

The AccessGrid (AG) is a continuous presence, open source system, developed by Argonne National Laboratory. Its applications scale from desktop/laptop based systems up to large rooms and auditoria. AG is based around the concept of shared venues containing audio, video and data pertinent to an ongoing collaboration. The concept is that users “visit” a venue to interact with their collaborators. The AG allows users to view all meeting participants in separate windows. AccessGrid software runs on all major platforms and support encrypted traffic and X.509 based authentication. Setup with the AG is complicated and technically demanding. It depends on multicast networking, which is not reliable or pervasively supported. The multicast requirement can be avoided by utilizing a unicast-multicast bridge but this increases the overall system complexity. While these characteristics are not problematic within the continuous presence paradigm, it makes the AG a difficult fit for ad hoc videoconferencing. From a firewall standpoint,

the AG requires two UDP ports per audio/video stream and some firewalls do not pass multicast traffic. The AG does not interoperate with other protocols like H.323 or SIP.

Integrated suites like LiveMeeting from Microsoft, MeetingPlace from Cisco and WebEx have the advantage of incorporating slide sharing, messaging and white board tools into audio and video conferencing. They feature ease of use and a high degree of reliability and are suitable for desktop conferencing in a homogeneous environment. They are not suitable for larger venues. They are relatively expensive and do not interoperate with open-source products or those from other commercial vendors.

B.1.2. Venue configurations

In this context, “venue” refers to any specific room from which a remote conference originates. We use this term in distinction to “site” which can refer to an entire lab or institution. The hardware required for remote conferencing at each location will depend on the size of the venue. This, in turn, is generally driven by the number of people that will want to participate from the same place at the same time. While several participants from the same site may be able to join a meeting from their own offices, it is often desirable to collect into a single room. This is particularly true for meetings that are predominately local, but for which some remote participation is needed. It is important to note that much of the investment in venues is long-lived and independent of the protocols or specific software chosen to transport the data across the network. Thus migration to newer standards should not be prohibitively expensive. It is also common to support more than one software protocol from a given venue.

Single User. The smallest and simplest installations are in individual offices supporting single users. These use inexpensive (<\$200 US) or free software running on an existing office computer. An inexpensive webcam (~\$100) must be purchased and installed. Some cameras have built-in microphones, but, for a single participant, it usually better to use a headset connected to the computer’s audio card. Shared presentations are displayed on the computer monitor. Total cost would be about \$300 per setup.

Small Groups. While it is possible for a second (or even third) participant to crowd around the computer on a fellow worker’s desk, for small groups, a more capable venue will provide a comfortable and productive environment for regular or extended meetings. The venue could be set up on a table wherever space and privacy permits. A minimal configuration would include a computer equipped with dual monitors, a camera located to cover all participants and an echo-canceling microphone. Control panels and images of remote participants would be displayed on one monitor, and shared presentations shown full-screen on the other. It is useful if a standard telephone is close by to provide “out-of-band” communication for trouble-shooting or other purposes. Such a configuration might

easily support a group as large as 5-7. Overall cost would be \$2,000-4,000 for equipment. The number of participants accommodated could be increased by adding a video projector and screen to display the presentation, additional microphones and an inexpensive sound mixer. (While it is possible to build portable conferencing systems that include the equipment described above and that can be moved from room to room, in our experience these systems are noticeably less robust and reliable.)

Larger Groups. When the size of the meeting group at one location grows beyond the number that can comfortably sit around a table, the cost and complexity of the venue rises sharply. The most difficult problem is providing good audio for all participants. Two options exist: (1) For a general purpose room where flexibility must be maintained, an array of ceiling microphones (approximately 1 per 10 square meters) can be installed for the presenter and audience; (2) A room with fixed tables and chairs can be equipped with a separate, fixed microphone (with mute button) for each participant. The speaker would wear a clip-on wireless microphone. In both cases, sophisticated mixing, echo canceling and sound processing hardware will be required. A computer, video projector and screen needs to be provided for slides, and a second computer or commercial video-conferencing CODEC with a very large monitor or second video projector carries the images and sound of remote participants. Equipment costs would exceed \$30,000 with significant design and installation expenses.

Auditoriums. These are usually large, tiered, fixed seating venues. Requirements for venues of this size are particularly demanding. Audio probably requires professional consulting. Provisions for a control room and technician to operate equipment during meetings should be made. (For a new auditorium, if one is to be built at the ITER site, capabilities for remote participation should be part of the design specification.)

Single user systems present an environment similar to a traditional teleconference call. Most users should adapt quickly to this mode of communication. As the meeting size and complexity increase, a degree of learning is required to maintain the flow of discussions — these are issues beyond simple operation of the equipment. Small meetings can be informal and ad hoc, while larger meetings cannot. Meetings with many sites making presentations may require a moderator. On the technical side, equipment should be tested ahead of time and the meeting chairman made familiar with its operation.

B.1.3. Multipoint issues

When more than two venues are involved in a conference, a multipoint control unit (MCU) is required. MCUs can help deal with incompatibilities between user agents and help mitigate site security and firewall issues. Tradeoffs between proprietary hardware MCUs and open-source software MCUs are similar to tradeoffs with user agents. The ESnet community has experience with a variety of hardware MCUs. At this time Codian (<http://www.codian.com>) seems to be selling the most capable units. They feature automatic transcoding to match each user's connection speed with no aggregate bandwidth limits; all ports can be operated up to full bandwidth. Users can select their view of the conference (last speaker, tiled, tiled with last speaker larger, etc. ...) and streaming of meetings for broadcast to "listen only" participants is supported. MCUs often incorporate the function of an H.323 "gatekeeper". These provide the means to restrict access to registered users, allow users to create ad-hoc meetings on the MCUs, and help user agents negotiate compatible protocols. The disadvantage of hardware MCUs is that they are expensive and inflexible. A 20 port MCU costs approximately \$40K and will become obsolete as protocol standards evolve. For this reason we recommend using MCUs already deployed in the research and education communities. We have identified and been given permission to use the ESnet MCU at LBNL in California and one at IN2P3 in France. It should be possible to locate a suitable one in Asia. We recommend at least one MCU per continent since many meetings will be local and because local (time zone) support will always be required. It may be possible in the future to bridge MCUs into a videoconferencing network, but no firm date is available from the vendors.

Software MCUs are more flexible and do not require expensive proprietary hardware. They can adapt to new protocols or be replaced by new software. Further, the software can be extended to include features found in integrated conferencing suites. However, they offer less reliability and lower quality service. Support is technically demanding. For a production H.323 system, software MCUs are not stable enough today for the ITER application. It is possible that this may change in the future. The VRVS/EVO system is a good example of software based conferencing system.

B.2. Slide/Document Sharing

Most of our technical meetings require shared viewing of documents. A wide range of technologies and approaches exist for this purpose, and while many are satisfactory, none has been found to be ideal. The choice of technologies for document sharing is driven by requirements and policy. Specifically the need (or lack) to support clients on multiple platforms; the requirements for shared control or annotation; sensitivity to costs; preference for proprietary versus open solutions and the requirement to share (via real-

time imaging) hand drawn graphics. One of the critical issues for document sharing is the variation in resolution of video displays. Some products do an excellent job of resizing screen images, while others do not. Since low-cost projectors operate at the XGA (1024x768) resolution, it is a common denominator for heterogeneous environments. In general, there is a tradeoff between acceptable transmission time and resolution. Presentations should use the XGA resolution pixel space and are advised to avoid very high resolution images. Imbedded animation and animated slide transition do not always work on the remote sites as expected. Table B-2 compares available technologies and some of their features, advantages and disadvantages.

Table B-2. Comparison of Document Sharing Software

Software	VNC	WebEx	MeetingPlace	LiveMeeting	EVO	H.239 Based Tools
Client access	Standalone app & browser access	Browser access	Browser access	Browser access	Java webstart	H.239 compatible CODECs
Recommended venue	Large meetings, desktop	Desktop	Desktop	Desktop	Desktop	Desktop-large meetings
Firewall traversal	Can be tunneled manually using SSH	Automatic tunneling	Automatic tunneling through: firewalls, HTTP, HTTPS	Automatic tunneling	Automatic	?
Platforms	Windows/Mac/Linux	Windows/Max less support on Linux	Windows/Mac	Windows	Windows, Linux, Mac	Polycom CODECs, codian MCUs, vPoint
Rights to permission levels	2 levels	Multiple levels	Multiple levels	Multiple levels	Symmetric	One at a time
Support cost	High	Low	Low	Low	Medium	Low
Ease of use	Medium	High	Medium	High	High	High
Presence info	No	Yes	Yes	Yes	Yes	No
Switching/multi-user	Difficult	Medium	Medium	Easy	Yes	Yes
Price	Free	~2K/month	Included as part of larger Cisco package	?	Free	Includes with certain video conference HW and software

At the present time, the best solution seems to be the screen sharing capabilities built into WebEx, an integrated conferencing suite. It is a web-based application, with no software installation required. Because it uses the web protocols, it runs on all platforms and has a minimal impact on site security policy. Slide sharing is easy to use and fast. The disadvantage of WebEx is its cost. A license capable of serving the entire ITER project would cost more than \$2,000 per month. Virtual network computing (VNC) is a

widely used open-source remote control client-server software suite available on all platforms of interest. However, it is slow and compute intensive. Because of its remote-control origins, it presents security and privacy issues. H.239 is a data sharing protocol standard that originated with the Polycom Corporation. It is integrated into their H.323 products, but so far has not been supported by many other vendors. It does offer the promise of high quality slide-sharing, tightly integrated into the video conferencing system. It cannot be recommended today, but may emerge as an important tool in the near future.

MeetingPlace is an integrated voice, video and document sharing suite from Cisco. It is used in production now in the ESnet community, but is somewhat difficult to use and does not do a good job of matching screen resolutions in heterogeneous environments. It is primarily a PC product, and though it is offered for Apple OS, it does not seem to offer full compatibility. The suite is fairly expensive and offers no advantages over competing technologies. LiveMeeting is an integrated conferencing suite from Microsoft. It is full featured and reliable, but is only available for PC platforms.

There are other solutions that have been employed for slide sharing, but which are less useful than those discussed above. Documents can be emailed to participants prior to the meeting and viewed independently on each user's system. While this is simple and requires no new software, the slides are not synchronized to the presentation and viewers may become lost. Further the documents are static and do not feature shared control or annotation. If multiple versions of presentations are circulated, attendees may not be able to identify the most current. Putting documents on a web site solves the last of these difficulties, but requires more preparation. In many cases, sites do not give ordinary users the ability to add documents to the web themselves (although if ITER adopts a project web site using Wiki, this problem can be eliminated — see B.5). Finally, slides can be imaged with a document camera and shared over the video link. This is the only method by which non-electronic documents can be shared ad hoc. However, the resolution is low and not suitable for many presentations. It is not at all useful for sharing non-presentation documents such as Word or Excel files. It also requires specialized hardware at each site desiring to originate presentations.

B.3. Whiteboard Tools

Virtual electronic whiteboards provide a real-time synchronous interaction environment between two or more remote parties by allowing them to write, draw and annotate the same document. Thus they provide a visual collaboration space for a group of distributed individuals to work together as one in a highly interactive fashion. The visual image (slide, document, drawing) can be captured with all of its annotations and posted on the Internet for sharing with a wider audience. Major interactive whiteboard

commercial companies include the Virtual Ink Corporation with its Mimio product line and Luidia with its eBeam products. Document annotation is included in several integrated conferencing products, such as WebEx and LiveMeeting. We do not have enough experience in this area to make a recommendation at this time.

B.4. Messaging Tools

Text-based Instant Messaging (IM) technologies provide a lightweight collaboration channel. IM requires the use of a client program that connects to an instant messaging service to provide quasi-real time communication. It is more immediate than e-mail, less intrusive than a telephone call, and can be used as a collaboration tool in non-formal meetings. Although it does not have an audio/video conferencing feature, it provides the ability to exchange information from a low-bandwidth network environment. Most services offer presence information, indicating whether people on one's list of contacts are currently online and available to chat. Besides the chatting functionality, IM can also provide for file transfer and web-link sharing.

A wide range of IM techniques and software implementations exist with similar capabilities. The major IM offerings include proprietary systems such as AOL Instant Messenger (AIM), Microsoft's Instant Messenger (MSN Messenger), Yahoo! Messenger and Google Talk, as well as open-source techniques such as Jabber from the Jabber Software Foundation. Since all IM systems use client-server architecture to send and deliver messages and content, the user management and security depend on how the IM server is deployed and where the IM server is located. Jabber-based IM is the preferred technique for the ITER organization over proprietary IM systems, since ITER can deploy a Jabber-based open-source server and have full control of user authorization, management and security.

B.5. Shared Project Web Site

Shared project web sites, sometimes called enterprise web portals, unify and document communications and thinking inside an extended organization. Some features of enterprise portals include

- Single point of entry — the portal becomes the delivery mechanism for all information services.
- Collaboration — portal members can communicate synchronously through IM or asynchronously through threaded discussion and email digests (forums) and blogs. Links can also be established to H.323-based web streaming and slide/document sharing (e.g. WebEx) for more advanced collaboration.

- Content and document management — services that support the full life cycle of document creation and provides mechanisms for authoring, approval, version control, scheduled publishing, indexing, and searching.
- Personalization — the ability for portal members to subscribe to specific types of content and services. Users can customize the look and feel of their environment.
- Integration — the connection of functions and data from multiple systems in to new components/portlets.
- Security — most portals provide a single sign-on capability that requires a user to authenticate only once.

A variety of software exists for creating web portals, but the open source Wiki software has recently been gaining in popularity. In essence, Wiki simplifies the process of creating HTML pages and records each individual record change that occurs over time, so that at any time, a page can be reverted to a previous state. It supports text, images, media files, and physics formulas. Security is flexible and can be configured as required, addressing issues like authentication, read/write access, server access, communication channel specification, and administration access. It is recommended that the ITER project use Wiki to create a versatile and interactive project web site as an adjunct to the conferencing tools described in this document.

Appendix C. Outstanding Issues and Activities

The following issues will require further research, discussion and testing. None would prevent initial deployment of remote conferencing for the ITER project, but some plan for resolution of these issues should be identified.

- Identify an H.323 MCU in Asia that is available to the ITER community.
- Identify and qualify additional codec software.
- Evaluate EVO after its initial release and test H.323/EVO compatibility.
- Obtain definitive pricing information for WebEx. Define requirements for document sharing and investigate cost sharing with HEP.
- Research prospects for MCU to MCU networking, including vendor support and vendor interoperability.
- Evaluate additional MCU vendors.

Appendix D. Future Developments Affecting Remote Conferencing Solutions

Video conferencing technology is evolving. The solutions proposed here are appropriate given today's available technology. As tools and protocols change, our solutions should evolve to accommodate them and take advantage of their new features. This does not mean that conferencing systems will have to be completely replaced. First, new systems are often backwards compatible with older ones, or MCU units can transcode between protocols. The systems used by ESnet were originally H.320 ISDN based. These have been replaced by H.323 based systems, which were able to bridge to ISDN participants. A similar evolution will probably occur with the next technology switch. Secondly, for software-based units, the software can easily be replaced or modified. Some hardware CODECs can be updated to work with newer protocols via firmware. Finally, for larger venues with hardware CODECs, the total system costs are dominated by audio processing hardware, installation and configuration. These should accommodate several generations of digital technology. Multipoint Conferencing Units, however, come in more expensive, discrete, pieces, ranging in price from \$30K to \$80K. It is possible that they will be upgradeable as new protocols are adopted, but they will have to be replaced over time. (The current recommendations involve using existing units, so there is no initial investment in this equipment.) The current recommendations focus on commercial solutions. As technologies evolve, it may become appropriate to shift towards free or open-source tools. None of the decisions made today are irrevocable. The tools chosen should be able interoperate with each other. As better tools become available, and as the user community gains experience with the technology, the solutions will gradually change. As discussed above, initial experience with EVO is quite good, and it may be a good choice for ad-hoc communication when it is released. Since these tools are free and software based, this transition will be very low cost.

Most commercial video conferencing systems available today are based on the H.323 protocol. Products based on H.323 have been available for over five years, and Session Initiation Protocol (SIP), the technology behind most voice-over-IP (VoIP) implementations, is emerging as a replacement. SIP based video tools are not ready for deployment in a production environment, but are now appearing as options in commercial and open-source software. We expect that the next generation of commercial offerings will be based on SIP. (We note that when ITER makes telecom decisions for the Cadarache site, they should consider SIP based VoIP among other technologies. SIP

based phone systems would be able to interact with SIP based video conferencing, IM, presence, and data sharing.)

In the security arena, we can anticipate a shift in authentication technologies. Traditional authentication is based on username and password. An individual working with multiple sites can quickly accumulate a large number of these username/password pairs, creating obvious difficulties and security vulnerabilities. Further, the requirement for explicit password logon also reduces the ability for computers to act securely on one's behalf in distributed computing environments. Science communities have begun to explore the use of public-key cryptography and X.509 identity certificates to create a single sign-on environment that works in a highly distributed organization. Users identify themselves once per session via a multi-factor logon, creating a credential that would be accepted by other systems. This technique is showing promise, but will require further development before it would be ready for the ITER organization.

Over the life of the ITER project, we can anticipate a transition from the current version of the Internet protocols, IPv4, to the newer IPv6. Most vendors of computer operating systems and networking equipment now have working implementations of the protocol and a significant number of Internet Service Providers (ISP) offer production IPv6 services. The principal driver for IPv6 deployment is the need for more Internet address space, but it should provide improved security by obviating the need for network address translation (NAT). An assessment of the plans and readiness of various videoconferencing solutions with respect to the IPv6 transition should be made at an early date.