Available online at www.sciencedirect.com
ScienceDirect
Fusion Engineering and Design 83 (2008) 539–544
Fusion Engineering and Design
www.elsevier.com/locate/fusengdes

A vision for a collaborative control room for ITER

D.P. Schissel *
General Atomics, P.O. Box 85608, San Diego, CA 92186-5608, USA
Available online 20 February 2008

Abstract

A vision for an onsite control room for ITER that will support worldwide experimental collaboration and operation is presented. Fusion experiments place a particular premium on near real time interactions with data and among members of the team. Enabling effective international collaboration on this scale is technically demanding, requiring powerful interactive tools and provision of a working environment for offsite personnel engaged in experimental operation that is every bit as productive as what is onsite. Expanding the view of the control room to include worldwide real time resources, both computational, data, and human, allows for a collaborative design that will significantly benefit ITER's scientific productivity. While the worldwide fusion program has a significant track record for developing and exploiting remote collaborations, the community should recognize that the collaborative needs of other communities share some similarity and therefore joint or shared research into collaborative technologies will increase the benefit to all concerned.

Keywords: Collaboration; Remote participation; Tokamak control room; Shared large display; Grid computing; Fusion experiment

1. Introduction and challenges

The future visitor to the ITER control room will most likely observe a large room with individual and large wall computer displays being used by approximately 100 individuals to operate the world's largest, most advanced, and expensive scientific instrument ever built for fusion research. The naive visitor might assume that the scientific team leading the day's operations forms in the morning, works for the day, and then disperses at night. In reality, on ITER just as on today's fusion experiments, scientific teams coalesce well before they walk into the control room. To create a day's experimental plan, scientists review previous data, create a scientific hypothesis, debate ideas on how to best test the hypothesis, run theoretical simulations to test these ideas, write a detailed plan for the day's experiments, work with the plasma control team to insure that the plan can be realized, have the plan reviewed by peers, and then modify the plan as required. It is only at this point, with the plan in hand, that the team is ready to walk into the control room and perform the day's experiment.

Therefore, when the ITER control room is discussed, and the discussion turns to making it a collaborative enterprise allowing remote participation, the discussion needs to include the entire process, not just the day's events. A remotely distributed group of scientists must jell early on in the process that leads to the creation of the team that develops the experimental plan.

Returning to the control room, ITER will require near real time interactions with data and among members of the team allowing for efficient between-pulse analysis, visualization, and decision-making. This mode of operation places a large premium on rapid data analysis that can be assimilated in near real time. However these requirements align very well with the requirements of the team before they enter the control room. Activities such as debating scientific ideas and comparing results of data analysis are performed while creating a scientific plan; they are just not done in such a time critical environment. The one activity that is unique to the precontrol room work is the usage of very large-scale plasma simulations. It is clear that massive amounts of data analysis including simulations will be performed between ITER pulses. However, it seems unlikely that the largest whole device modeling simulations that will be running on petascale computers for experimental planning will also be run between pulses. The approximate 1 h between ITER pulses will not provide sufficient time for the largest simulations. Therefore, if the collaborative requirements of the control room are satisfied, including access to datasets from the largest scale simulations, the requirements of the entire process, from origin of the experimental idea to execution, will be satisfied.

Carrying out these activities in an international collaboration and on the scale of ITER is a technically demanding problem,

* Tel.: +1 858 455 3387; fax: +1 858 455 3586.
E-mail address: schissel@fusion.gat.com.

0920-3796/$ – see front matter
doi:10.1016/j.fusengdes.2007.12.011

requiring a working environment for off-site personnel that is every bit as productive as what is onsite. One of the greatest challenges will be the provisioning of systems for analyzing, visualizing and assimilating the data to support decision making in support of ITER experiments. ITER's scientific productivity will be inextricably linked to the capability and usability of its collaborative infrastructure. Such an effective infrastructure is required for the success of the entire ITER project and will maximize the value of ITER to the home fusion programs as well.

2. Policy and infrastructure

Implicit in the design of the collaborative control room are aspects associated with policy and infrastructure more than computational or collaborative capability.

First is the realization that the remote participant is as valuable as the local scientist. Therefore, ITER's policy must be to allow an individual scientist to retain the same rights and privileges as they move from on site to working remotely. A scientist should not be *de facto* penalized for being at a remote site.

ITER's data archives need to contain all raw, processed, and simulation data and should be available in a timely manner to all members of the scientific team. Thus, data analysis performed by off-site researchers should be written back into the ITER repository. For complex and highly coupled systems like ITER, the scientific team and program must not be fragmented or impeded by data access limitations.

Clearly, the opening of the ITER control room and data repositories to off-site collaborators will require sufficient security to ensure protection of these valuable resources. However, at the same time, security must not be so onerous as to restrict the ability of remote scientists to contribute. ITER's cyberspace security needs to allow users from administratively and geographically distributed organizations to access resources (data, codes, visualizations, people). This imposes the need for users to have a unique cyber-identity and a means by which they can authenticate themselves as that entity at each site. Additionally, the user needs to be able to use a unique passphrase once, referred to as single sign-on and have authentications at other sites derived automatically from a limited lifetime token granted by sign-on. Additionally, the security system must enable multiple stakeholders for a single resource to set access policy for that resource.

Network connectivity to the ITER site will be an enabling technology and must be sufficient to allow for fully active remote participation. Estimates show that ITER will require network connectivity from the site in multiples of 10 Gb/s to fully support remote scientific collaboration [1]. Even with this network connectivity, ITER's local and wide area network should be able to support quality of service (QoS) guarantees so that outgoing and incoming time critical data is available to support experimental operations. Additionally, ITER will most likely need to adopt IPv6 due to the limitations of address space in IPv4 [2]. A number of methods have been invented to circumvent this limitation, with network address translation (NAT) being the most common. But NAT breaks the IP architectural model that assumes globally unique address for each host and breaks peer-to-peer connectivity symmetry required for many applications that support collaborative computing. The presence of NATs in the communication path hinders the ability to encrypt traffic at the network layer, which would otherwise provide security that is transparent to the application.

3. Collaborative control room

The vision for ITER's onsite control room is one that supports worldwide experimental collaboration and operation. The concept is modular, in that any function performed onsite can also be performed remotely. Therefore, our vision supports one onsite control room and an arbitrary number of remote control rooms. The vision pushes in two directions. First, to view applications and data as services to be run on geographically dispersed resources thereby increasing the amount and detail of both analysis and simulation results that are made available to the scientific team. Second, by bringing in expertise from geographically remote teams of experts, the depth of interpretation can be increased leading to improved assimilation of those results. The following five subsections define the collaborative control room.

3.1. Secure computational services

The model of computation for ITER is that of an application service provider (ASP) where computer-based data analysis and simulation services are provided to end-users over the wide area network [3]. This model does not preclude local scientists from writing and running their own applications locally. In fact, such codes can be offered in an ASP model to fellow local scientists as well as those stationed remotely. What the model does eliminate is the requirement for remote scientists to bring their application to the ITER site for it to be used. Instead, the scientist locally maintains the code as well as the resources that the code would execute on. This mode of operation has advantages for both the client and the provider: it frees the clients from maintaining and updating software and the providers from porting and supporting it on a wide-range of platforms. In this environment, access is stressed rather than data or software portability. In some cases, the ability to bring the code to ITER is not feasible. Imagine an integrated burning plasma simulation code that operates on one of the world's fastest supercomputers. Portability is not an option, yet by using the ASP model, the code and computer systems can be made available to the distributed ITER team.

This type of computing paradigm, the decoupling of production and consumption, is often associated with grid computing [4]. However, our usage of the term does not refer to computer cycle scavenging or distributed supercomputing. The proposed computational grid has been used in a variety of scientific disciplines and in the United States to support fusion research on FusionGrid [5]. The codes TRANSP, GATO, and ONETWO operate as FusionGrid computational services under the ASP model. The code TRANSP, used for time-dependent analysis and simulation of tokamak plasmas, was