The National Fusion Collaboratory

A MICS/SciDAC National Collaboratory Pilot Project

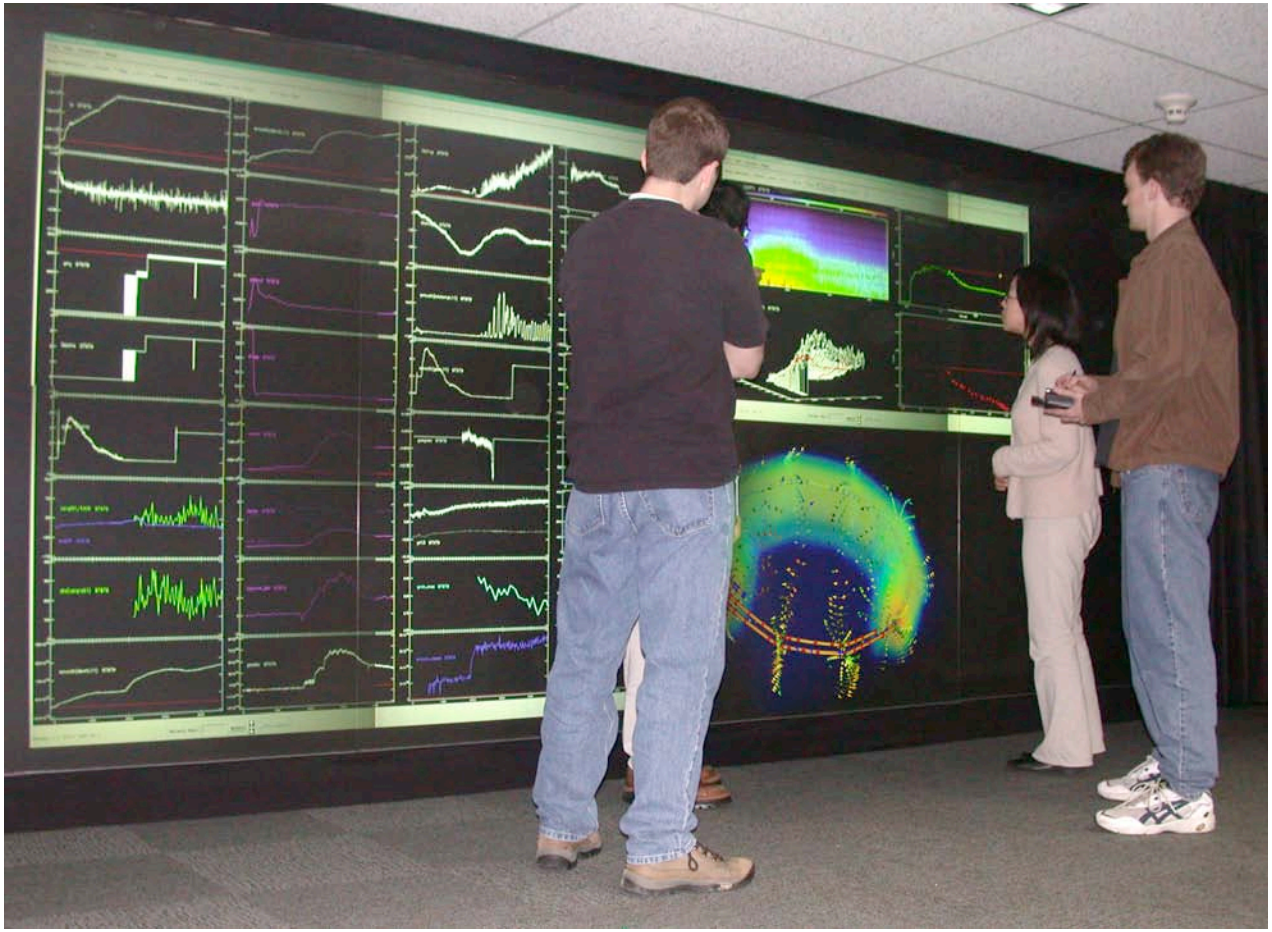

Presented by
David P. Schissel
Lead–PI

at
The PSACI PAC Meeting
June 3, 2004
Princeton, NJ

www.fusiongrid.org

PRESENTATION'S KEY POINTS

- SciDAC has changed the culture of collaborative FES research
  - FusionGrid services being used to benefit daily FES research
- The collaborative control room is being realized
  - Secure computational resources scheduled as required
  - Rapidly compare experimental data to simulation results
  - Share individual results with the group via shared displays
  - Fully engaged remote scientists with audio, video, shared displays
- Collaborative technology critical to the success of the FES program
  - Experimental: Fewer, larger machines in future (ITER)
  - Computation: Moving toward integrated simulation (FSP)

NATURE OF FUSION RESEARCH DRIVES REQUIREMENTS FOR COMPUTING AND NETWORKING

- Experiments
  - Real time interactions of large, geographically extended teams
  - Real time interactions between specialized small groups
  - Faster between-pulse analysis translates directly to productivity
  - Building an extended team of experts from small groups
  - Barriers to use of powerful analysis tools can be significant

Theory and Computation
  - Simulations producing very large data sets (GB=>TB=>PB)
  - Interactive visualization and analysis presents a severe challenge for computing and networking
  - Increased code sharing and collaborative development

The National Fusion Collaboratory Project (NFC) is addressing these needs

FusionGRID
www.fusiongrid.org

NFC PROJECT HAS PUSHED FES AS A HIGH-IMPACT SCIENCE FOR ASCR CSET ACTIVITIES

- Networking
  - High-Performance networks for high-impact science (2002)
  - Networking Challenge: Roadmap to 2008 (2003)
- Data Management
  - Ongoing led by Richard Mount of HENP
- Visualization
  - Workshop organized by John van Rosendale
- Collaborating with numerous SciDAC projects
  - Cover many different science applications including HENP

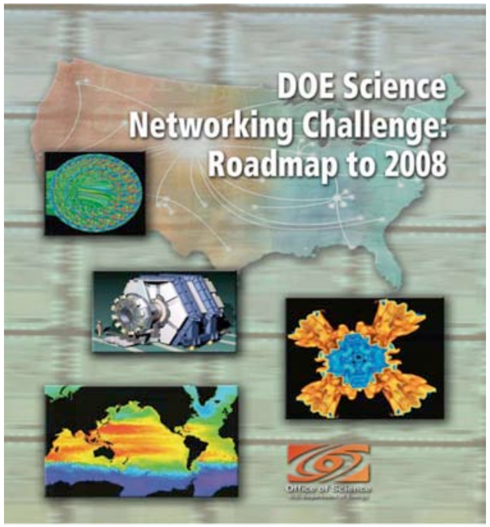

FusionGRID
www.fusiongrid.org

THE NFC IS HELPING TO ACHIEVE SCIDAC GOALS

- Create a collaboratory software environment to enable geographically distributed scientists to work effectively together as a TEAM and to facilitate remote access, through appropriate hardware and middleware infrastructure, to both facilities and data

Lead with the science:
Ultimate goal of advancing research in science central to the DOE mission

FusionGRID
www.fusiongrid.org

THE GOAL OF THE NFC IS TO ADVANCE SCIENTIFIC UNDERSTANDING & INNOVATION IN FUSION RESEARCH

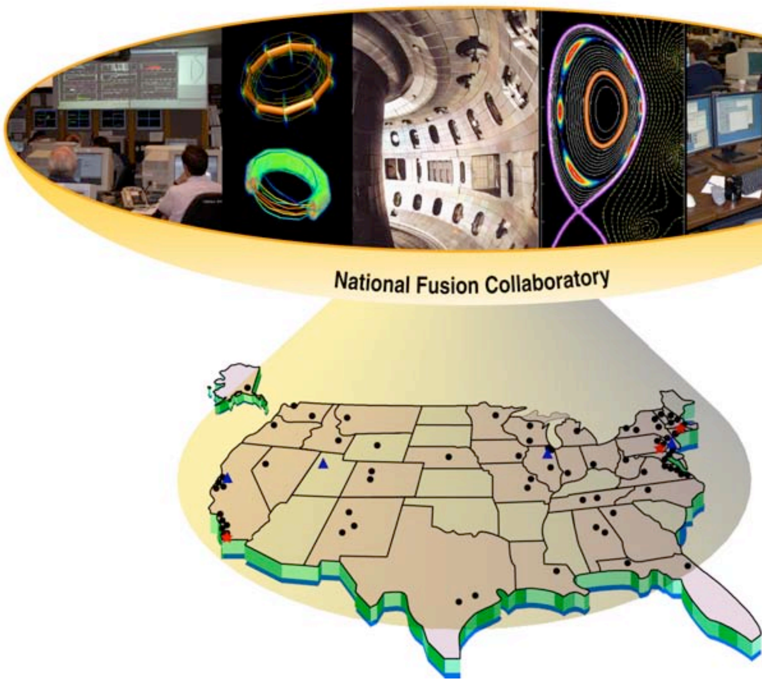

- Experimental Facilities
  - More efficient use resulting in greater progress with less cost
- Theory & Modeling
  - Integrate theory & experiment
- Facilitate multi-institution collaboration
  - Integrate geographically diverse groups
- Create standard tool set
  - To build in these services in the future

DOE ASCR:
Half-way through a total of
5 years at $1.8M/year

FusionGRID
www.fusiongrid.org

THREE LARGE U.S. EXPERIMENTAL FACILITIES AND A VIBRANT THEORETICAL COMMUNITY

Collaboratory is required to advance fusion science: geographically diverse community (37 states, 3 large experiments), leading to 1 worldwide experiment

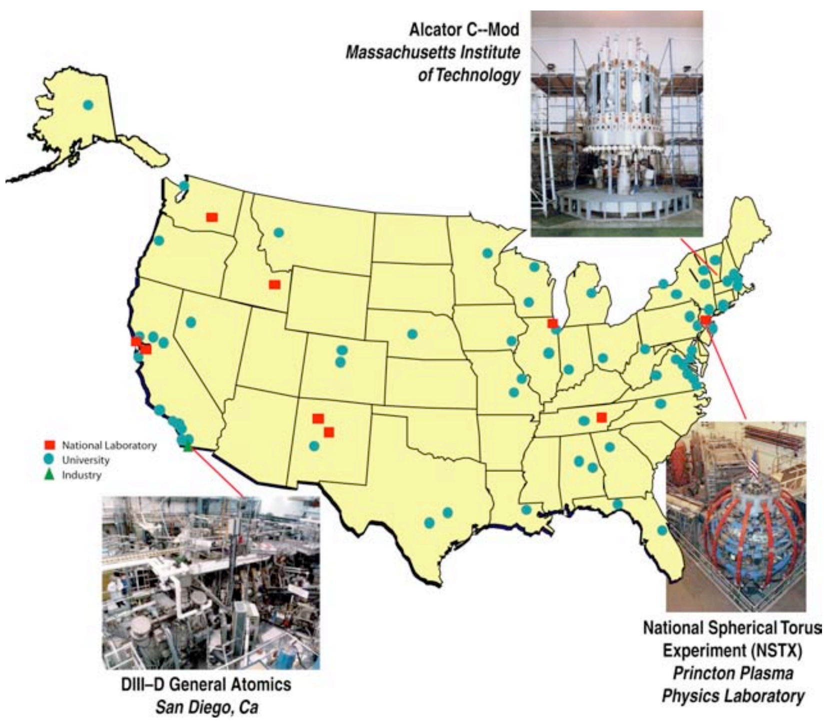

- 3 Large Experimental Facilities
  - ~$1B replacement cost
- 40 U.S. fusion research sites
  - Over 1500 scientists
- Efficient collaboration is required!
  - Integrate geographically diverse teams
- One future worldwide machine
  - Not based in US
  - US needs collaboration tools to benefit

FusionGRID
www.fusiongrid.org

THE NFC IS BEING SUCCESSFUL FOR FUSION SCIENCE

- Success is not defined by demos or papers published
  - Although these are important aspects of the project
- Success is defined by positive impact on the science
  - True metric is the actions of the fusion scientists
- A main fusion code is *only* available via NFC Grid computing
- The fusion program is spending its *own* money to purchase equipment and to support the usage of new NFC functionality
- A fusion scientist decided to *remain in San Diego* and use NFC technology to lead the fusion experiment in England

There is no going back:
SciDAC has changed FES collaborative culture

FusionGRID
www.fusiongrid.org

THE NFC PROJECT BENEFITS FROM A DIVERSE TEAM

Synergistic benefits derived from interdisciplinary interactions

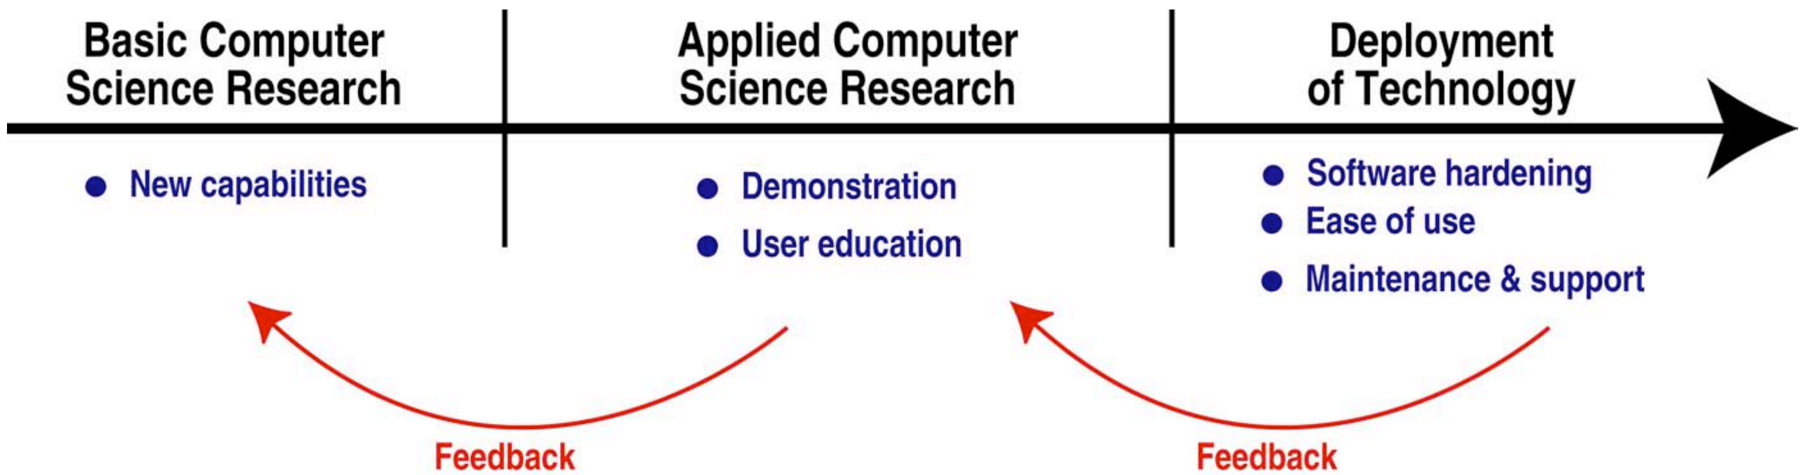

- ANL: Distributed Systems Lab
- ANL: Futures Lab
- General Atomics: DIII-D Fusion Lab
- LBNL: Distributed Systems
- MIT: C–Mod Fusion Lab
- Princeton Computer Science
- PPPL: NSTX Fusion Lab
- Utah: Scientific Computing & Imaging

NFC'S RELATION TO SCIDAC AND OTHER DOE PROJECTS

- DOE/OFES base program for FES research
  - Skilled computer scientists and computational resources
- Fusion science based SciDAC Programs
  - MHD, Mircoturbulence, and rf heating
- Plasma Science Advanced Computing Institute (PSACI)
  - Strong endorsement of NFC plans & accomplishments by PSACI PAC
- Data Grid Toolkit; Security & Policy for Group Collaboration; Distributed Security Architecture
  - Secure access, authentication, authorization, Globus GSI/Akenti
- Particle Physics Data Grid; DOE Science Grid
  - Site security,Firewalls, and Grid security; CA for FusionGrid
- Middleware to Support Group to Group Collaboration
  - AG development: user education & testing & feedback
- eServices Infrastructure for Collab Science; Portal Web Services
  - NFC & Fusion science as customer
- Global Grid Forum and Common Component Architecture Forum
  - Community wide standards

FusionGRID www.fusiongrid.org

NOT FOCUSING ON TRADITIONAL GRID APPLICATIONS – CYCLE SCAVENGING & DYNAMIC CONFIGURATION

- Traditional computational Grids, arrays of heterogeneous servers
- Machines can arrive and leave
- Adaptive discovery where problems find resources
- Workload balancing and cycle scavenging
- Bandwidth diversity where not all machines are well connected

This model is not well suited to fusion computation:
We are aiming to move high-performance distributed computing out onto the wide area network

FusionGRID www.fusiongrid.org

PLACING DISTRIBUTED APPLICATIONS OUT ON THE WAN PRESENTS SIGNIFICANT CHALLENGES

- Crosses administrative boundaries
- Increased concerns and complexity for security including authentication and authorization
- Resources not owned by a single project or program
- Distributed control of resources by owners is essential
- Needs for end-to-end application performance & problem resolution
  - Resource monitoring, management & troubleshooting not straightforward
  - Higher latency challenges network throughput & interactivity
- People are not in one place for easy communication

FusionGRID
www.fusiongrid.org

THE VISION FOR THE FUSION COLLABORATORY

- Data, Codes, Analysis Routines, Visualization Tools should be thought of as network accessible services

- Shared security infrastructure with distributed authorization and resource management

- Collaborative nature of research requires shared visualization applications and widely deployed collaboration technologies
  - Integrate geographically diverse groups

- Not focused on CPU cycle scavenging or “distributed” supercomputing (typical Grid justifications)

Optimize the most expensive resource - people’s time

FusionGRID
www.fusiongrid.org

VISION – RESOURCES AS SERVICES

- Resources are computers, codes, data analysis routines, visualization tools, experimental operations
- Access is stressed rather than portability
- Users are shielded from implementation details
- Transparency and ease–of–use are crucial elements
- Shared toolset enables collaboration between sites and across sub–disciplines
- Knowledge of relevant physics is still required of course

VISION – SECURITY INFRASTRUCTURE

- Strong authentication identifies users currently based on X.509 certificates from DOE science Grid
  - Interoperability with international Grid Certificate Authorities

- Distributed authorization allows stakeholders to control their own resources
  - Facility owners can protect computers, data, and experiments
  - Code developers can control intellectual property
  - Fair use of shared resources can be demonstrated & controlled

FusionGRID
www.fusiongrid.org

VISION – VISUALIZATION AND A/V TOOLS

- Maximum interactivity for visualization of very large datasets

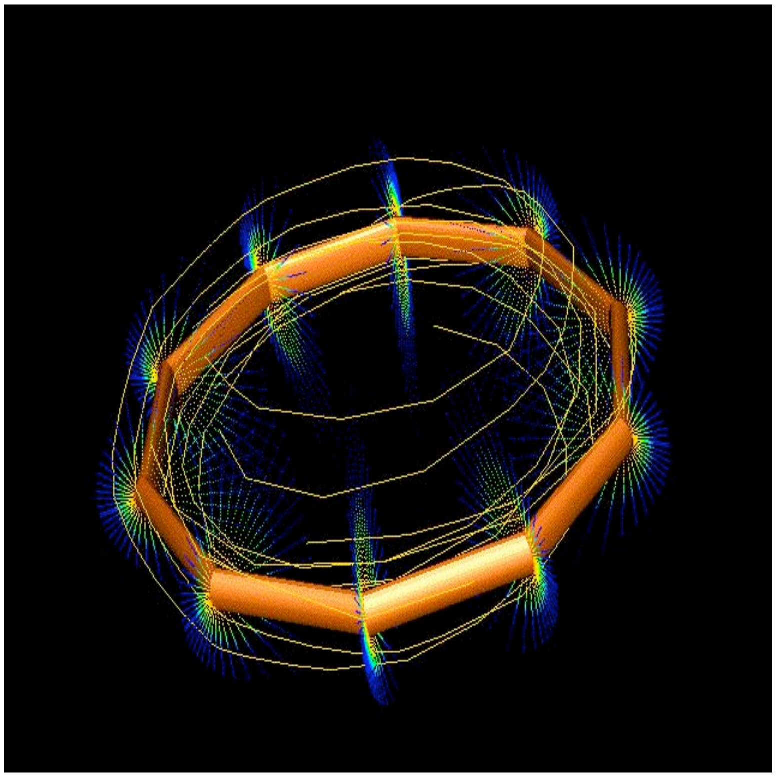

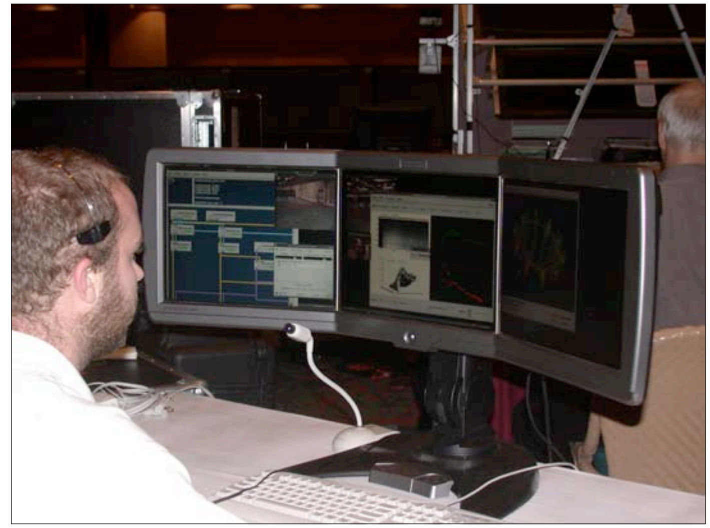

- Use of extended tool sets for remote collaboration
  - Flexible collaboration environment
  - Shared applications

EXPERIMENTAL SCIENCES PLACES A LARGE PREMIUM ON RAPID DATA ANALYSIS IN NEAR-REAL-TIME

DIII-D Control Room

- Pulsed experiments
  - 10s duration plasma every 20 minutes
- 20-40 people in control room
  - More from remote locations
- 10,000 separate measurements/plasma
  - kHz to MHZ sample rates
  - Between pulse analysis

- Not batch analysis and not a needle in a haystack problem
  - Rapid “real-time” analysis of many measurements
- More informed decisions result in better experiments
  - The collaborative control room

FusionGRID
www.fusiongrid.org

THE COLLABORATIVE CONTROL ROOM IS FUNDAMENTAL TO ADVANCING FUSION SCIENCE

- Secure computational resources that can be scheduled as required
- Rapidly compare experimental data to simulation results
- Share individual results with the group via shared displays
- Fully engaged remote scientists with audio, video, shared displays

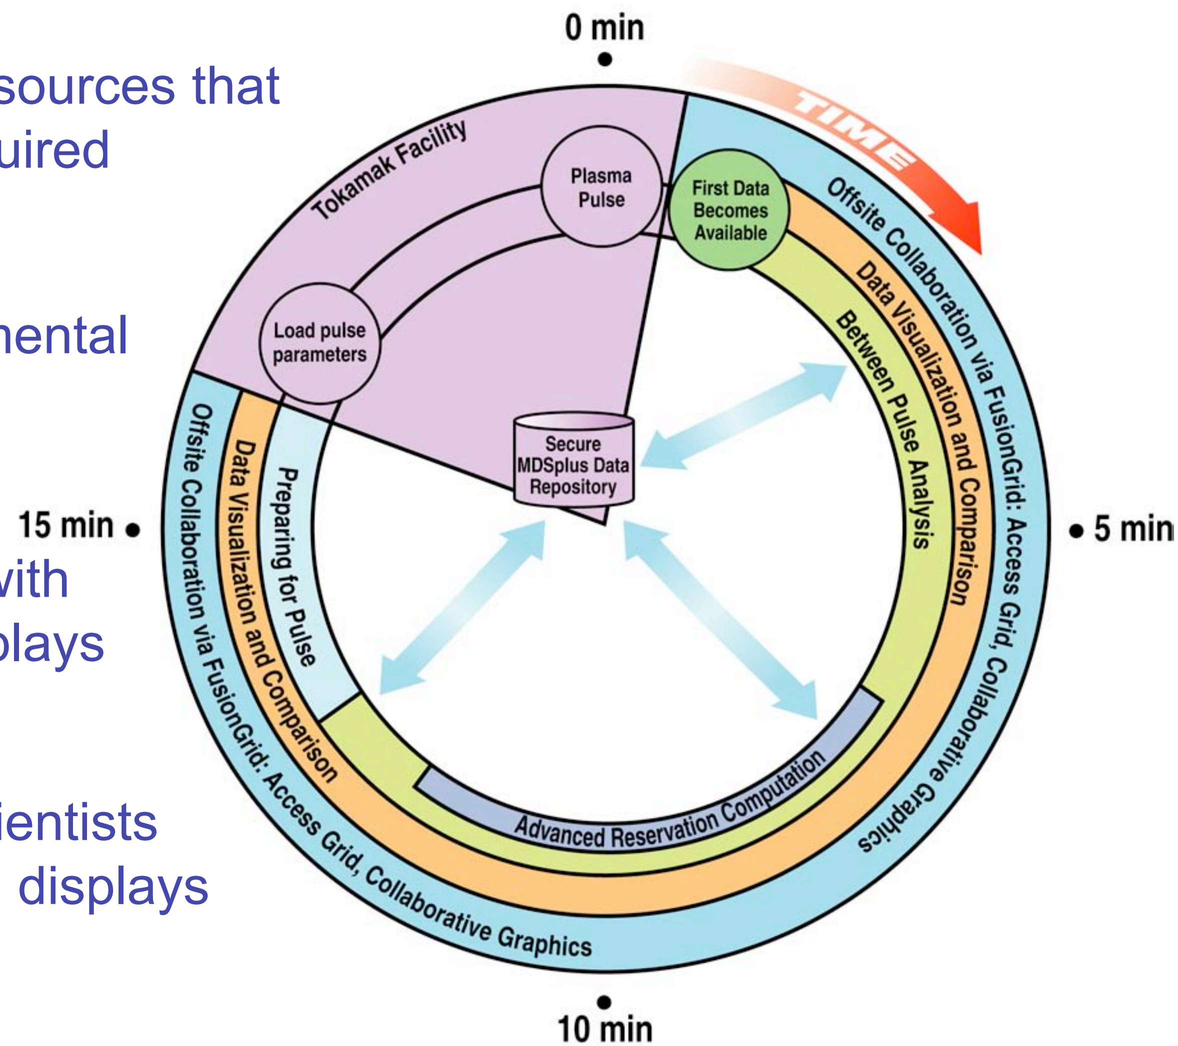

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM

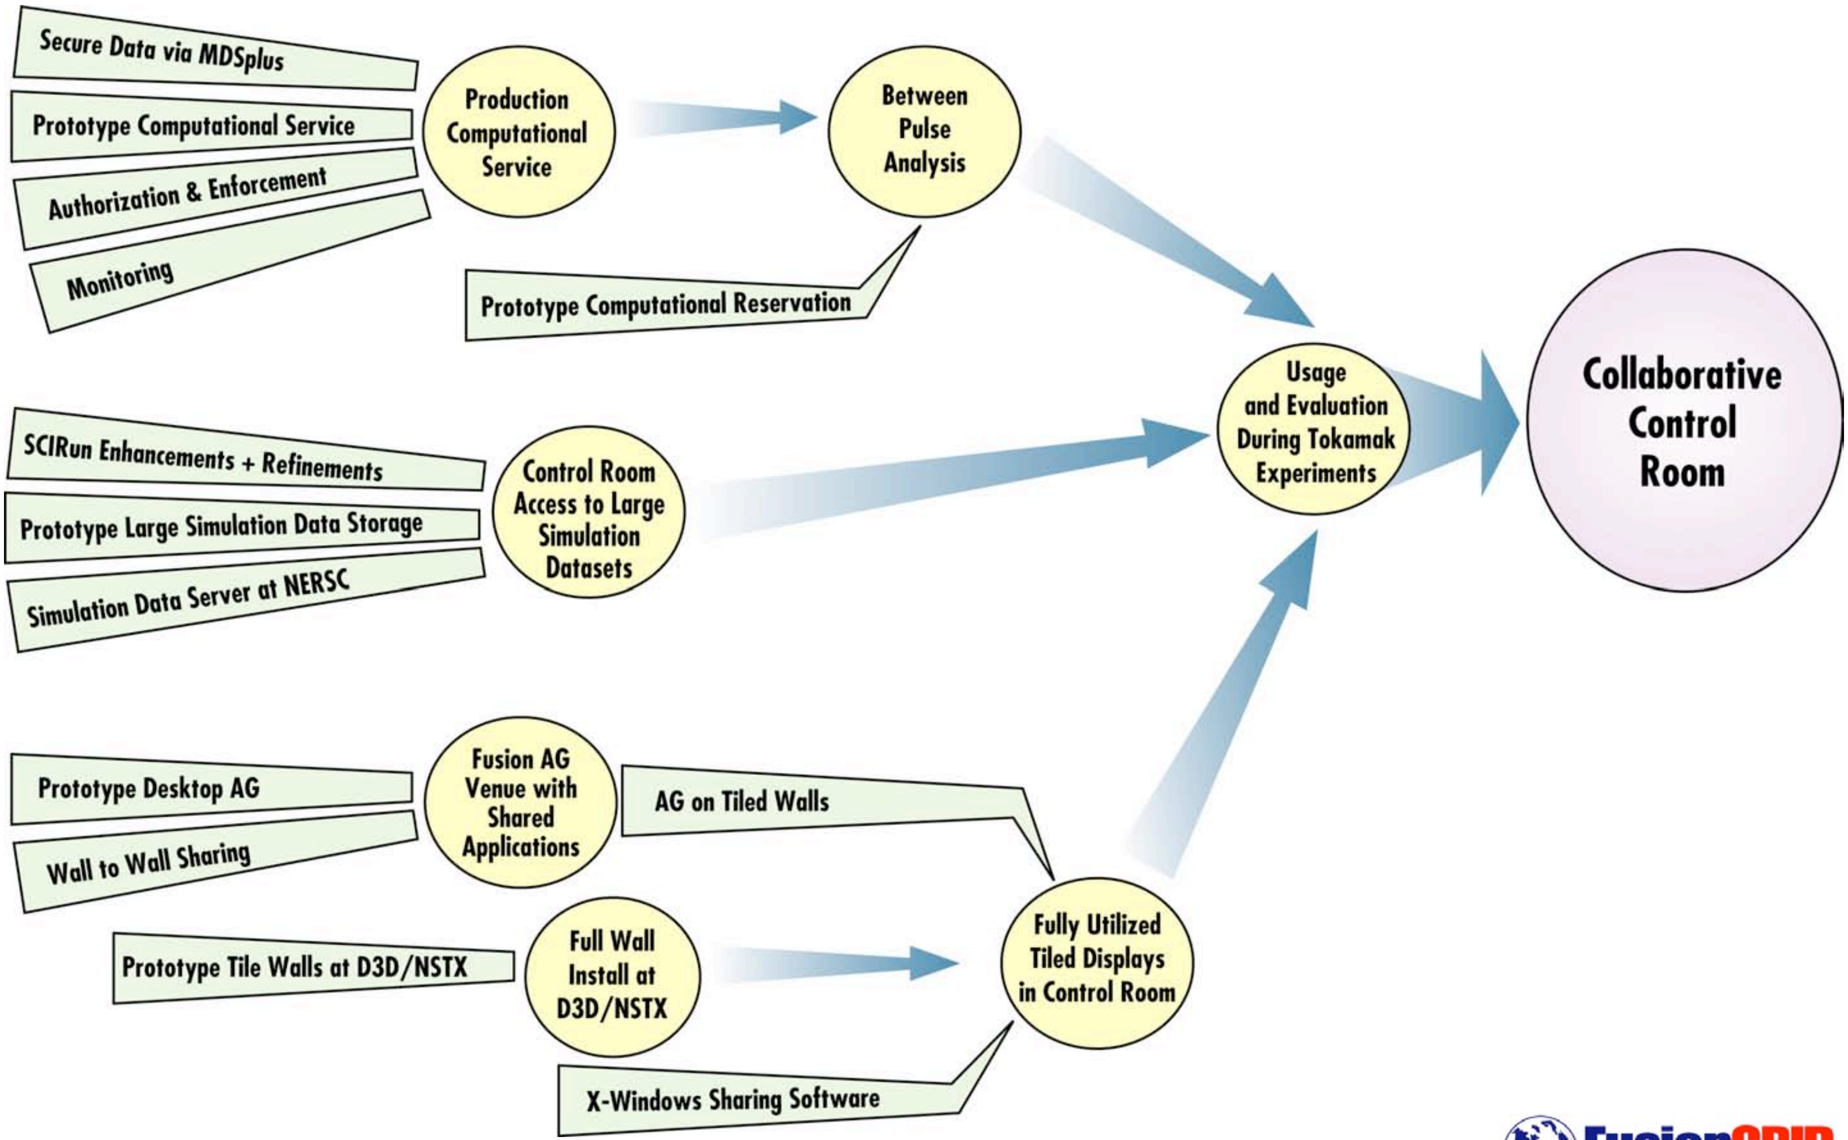

FusionGRID
www.fusiongrid.org

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM

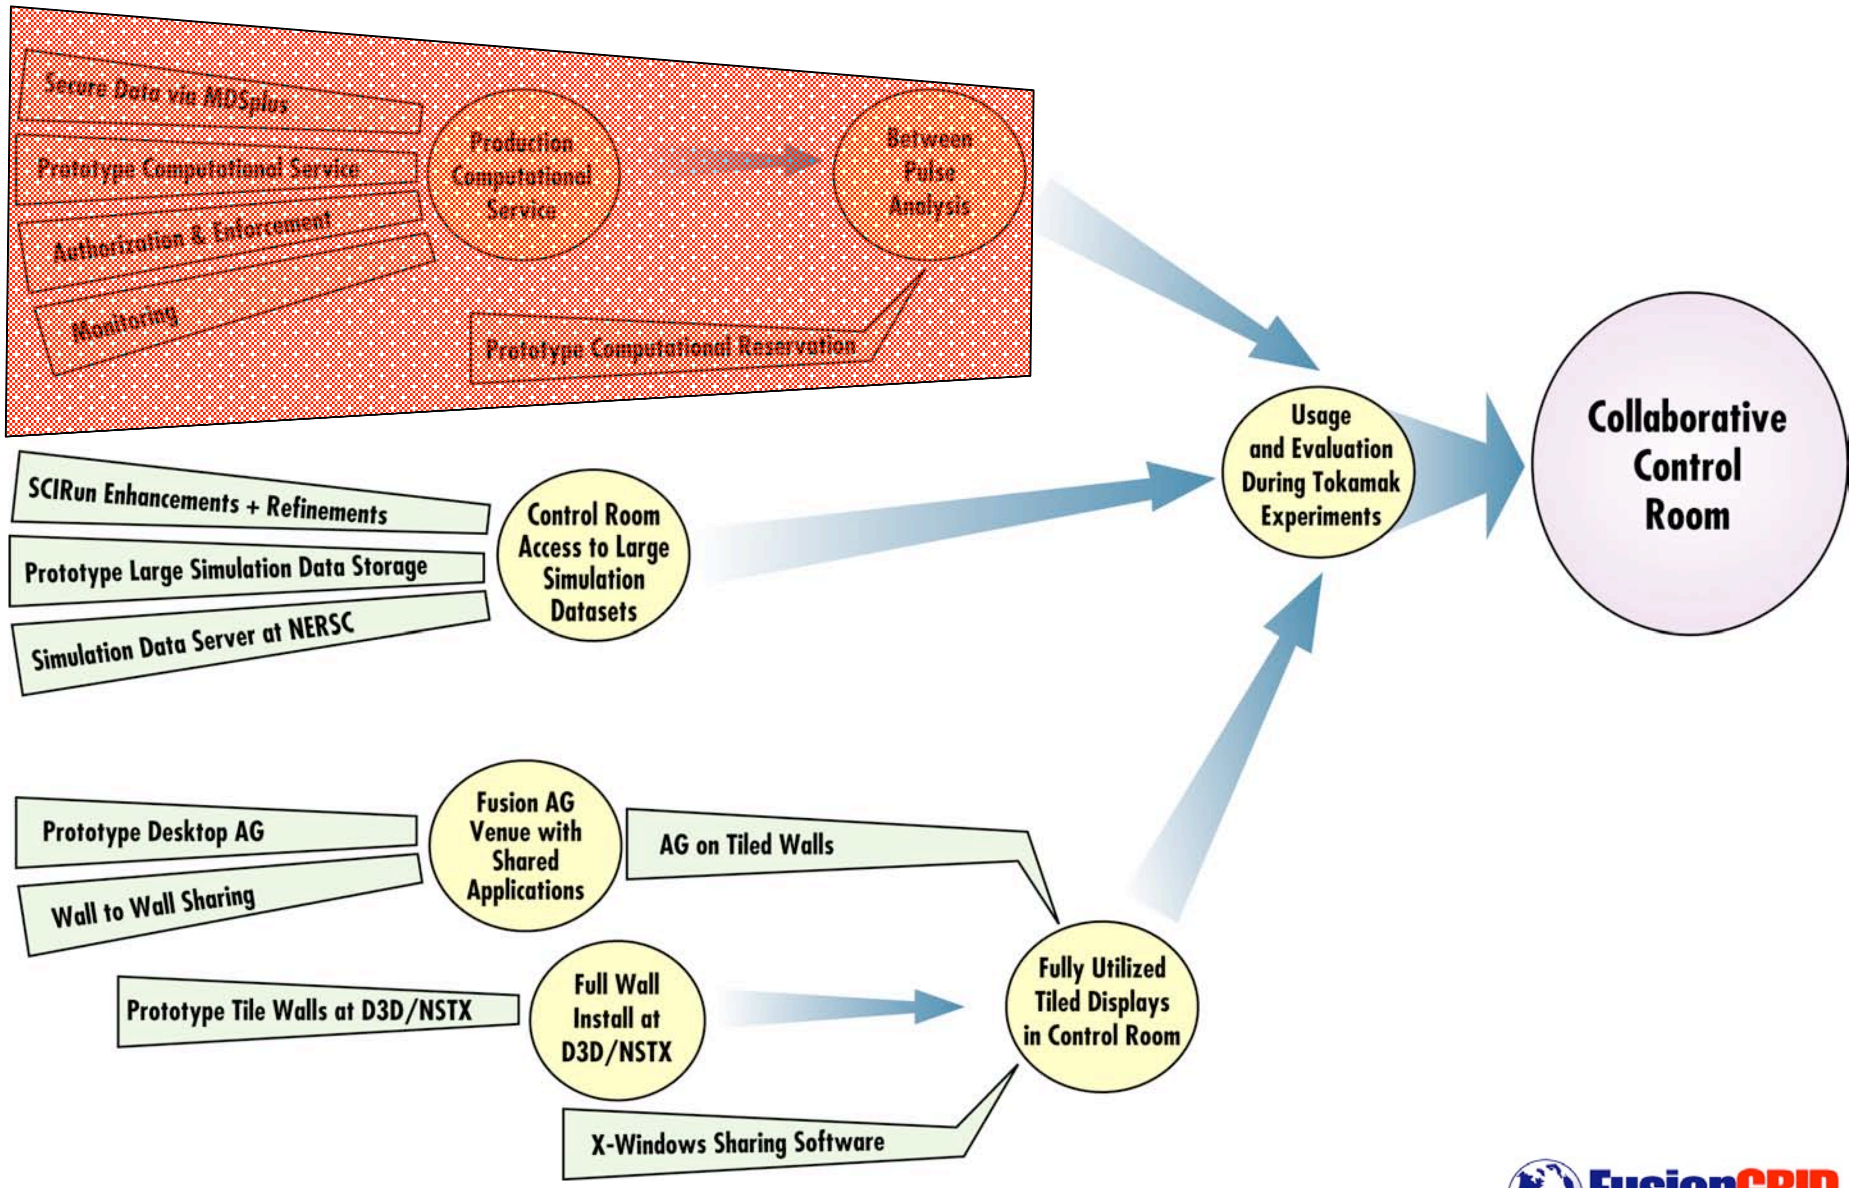

FusionGRID
www.fusiongrid.org

SECURE ACCESS TO FUSION DATA VIA MDSplus

- MDSplus: remote access based on client-server model
  - Used at more than 30 sites (robust)
- Service rather than file oriented
- Hierarchical, self descriptive, extensible, scalable, simple but powerful API
- MDSplus secured on FusionGrid via Globus GSI
  - Underlying technologies are X.509 certificates and OpenSSL
- Parallel network transfer via XIO - useful for high latency networks

FusionGRID www.fusiongrid.org

TRANSP WAS FIRST GRID SERVICE DEPLOYED: ANY CODE CAN BE A GRID SERVICE

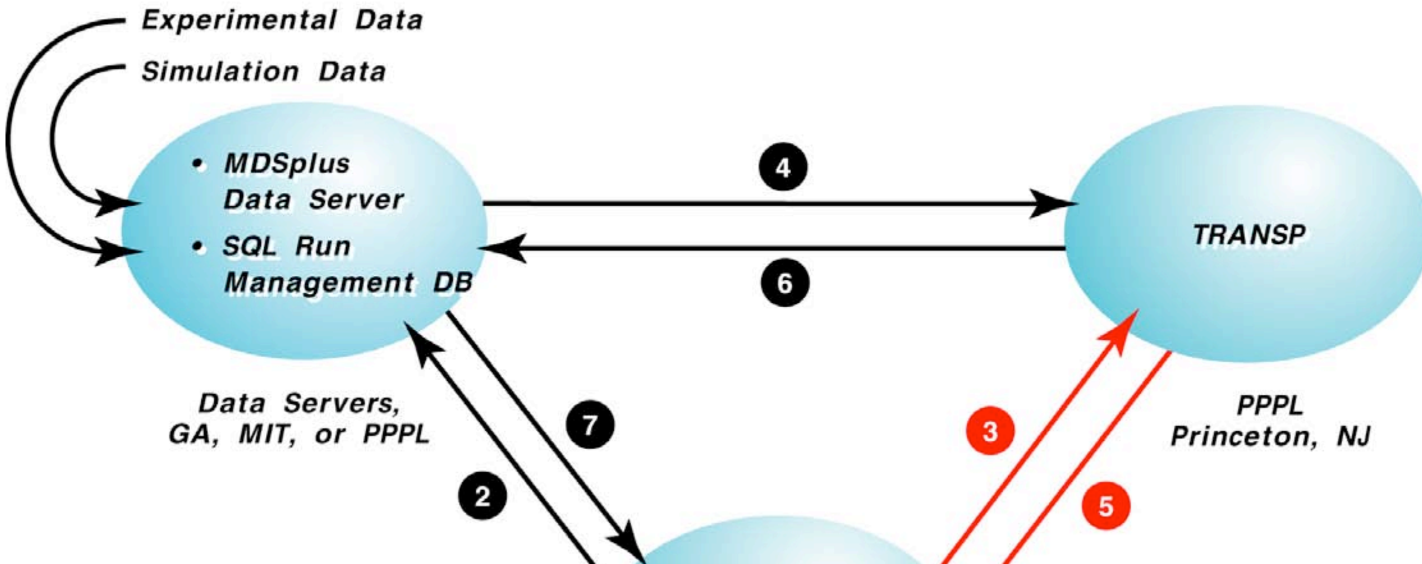

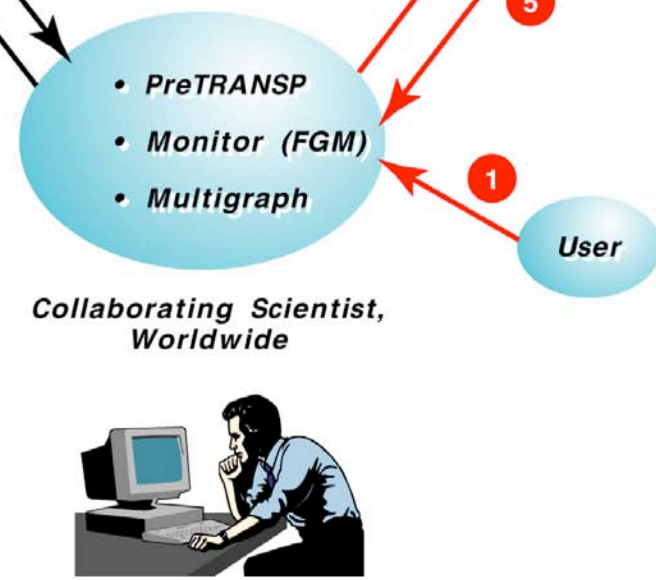

"This is a success"
- Better support for users
- Users get latest versions
- Better support with less effort
- Access to faster computations

The U.S. TRANSP Service
- 1,800 cases, 10,000 CPU hrs
- 9 fusion experimental machines

FusionGRID
www.fusiongrid.org

FUSION GRID MONITOR: AN EFFICIENT APPLICATION MONITORING SYSTEM FOR THE GRID ENVIRONMENT

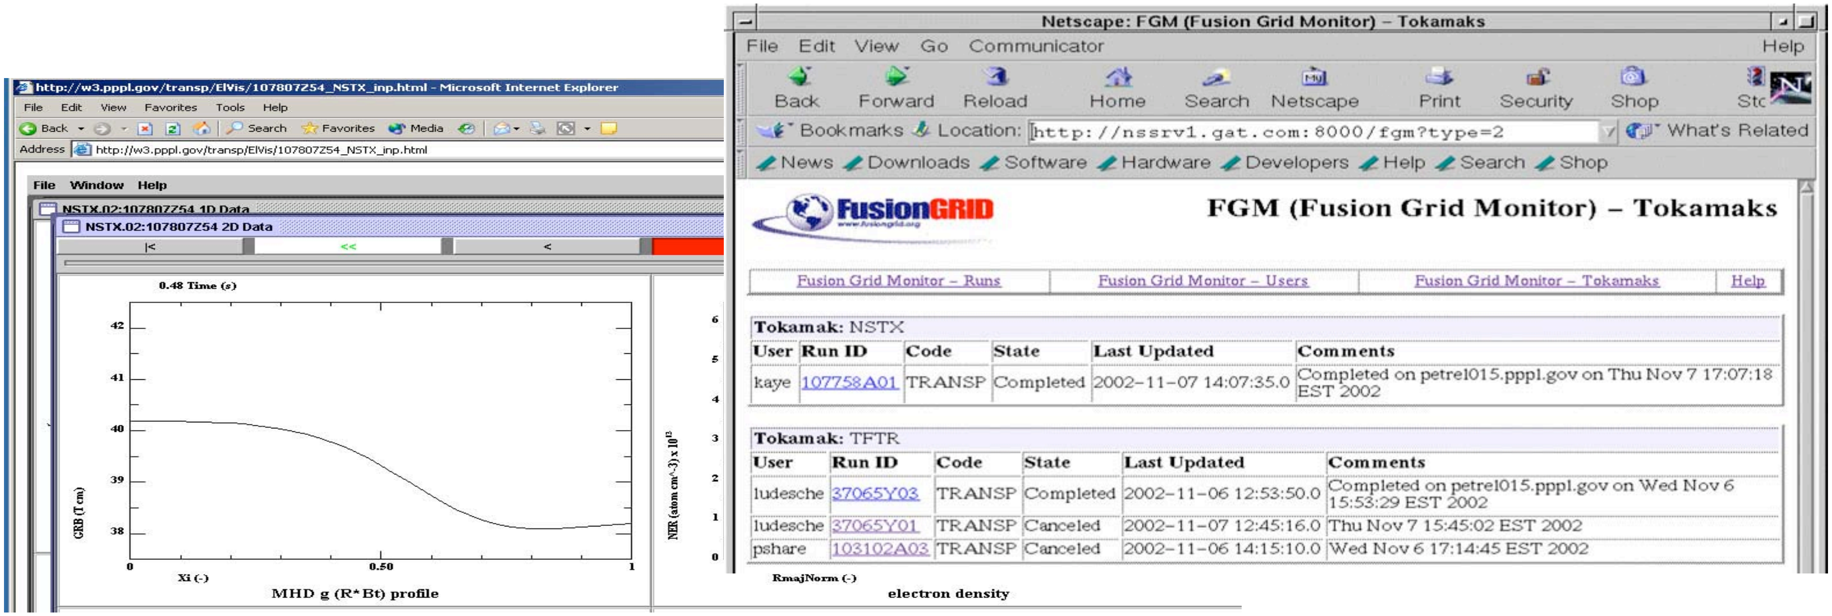

- Users track and monitor the state of applications on FusionGrid
  - Output dynamically via HTML, Built as Java Servlet (JDK2.1)
- Code maintenance notification
  - Users notified, queuing turned off, code rebuilt, queue restarted
- Results of simulation visualized during run
  - Both input and output quantities

ADVANCED RESERVATION COMPUTATION FOR DATA ANALYSIS TO SUPPORT EXPERIMENTAL SCIENCE

- Long-term vision: SciDAC code on supercomputer between pulses
  - Data management
  - Network QoS
  - Visualization
  - CPU scheduling
  - Faster CPUs and algorithms

- End-to-end agreement being prototyped in the NFC project
  - CPU reservation
  - Network transfer agreements based on simple prediction

- FusionGrid service TRANSP will be tested between pulses
  - First such capability for FES research

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM

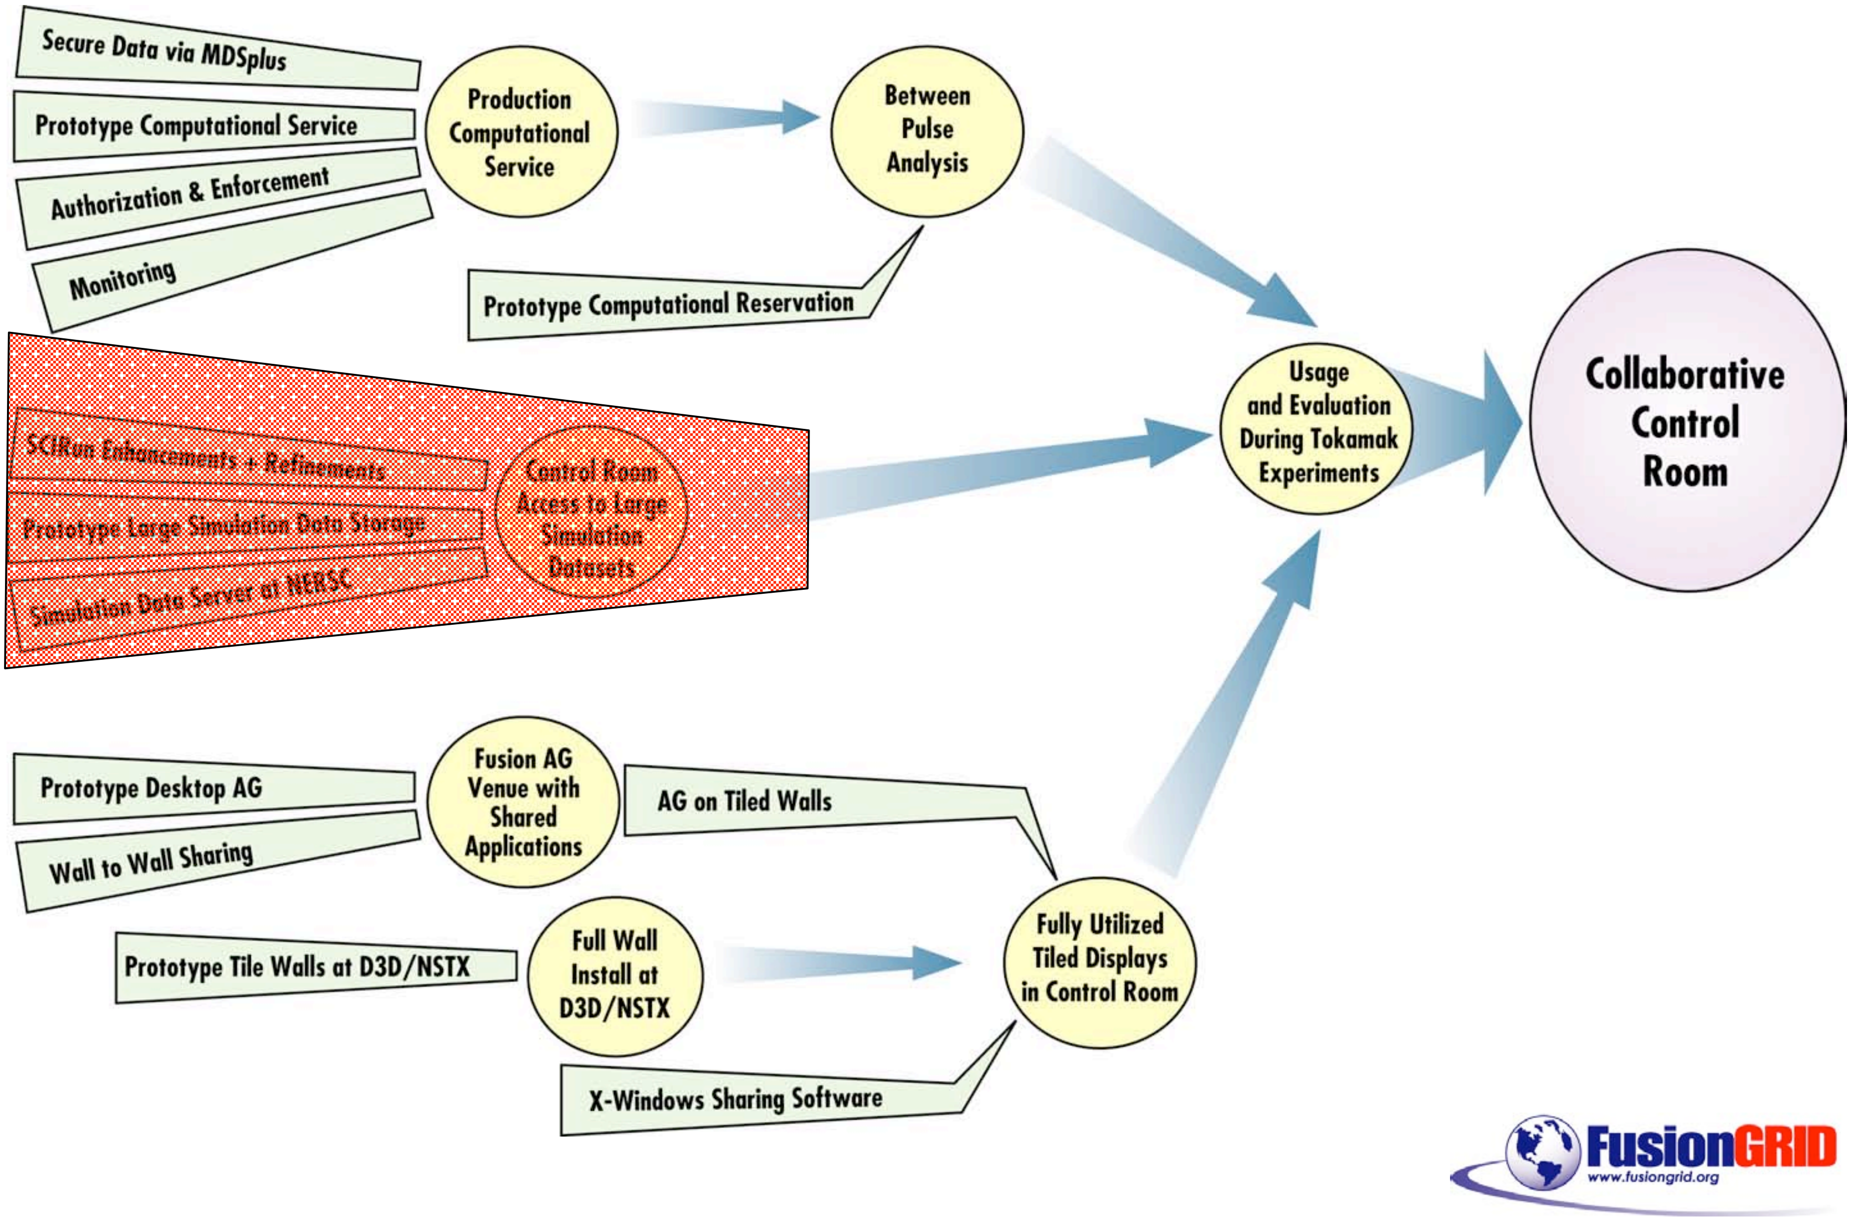

SCIRUN TO VISUALIZE COMPLEX SIMULATIONS FOR BETTER UNDERSTANDING

- Open source, multi-platform capable for a wide user base
- To facilitate quantitative comparison of simulations & experimental results

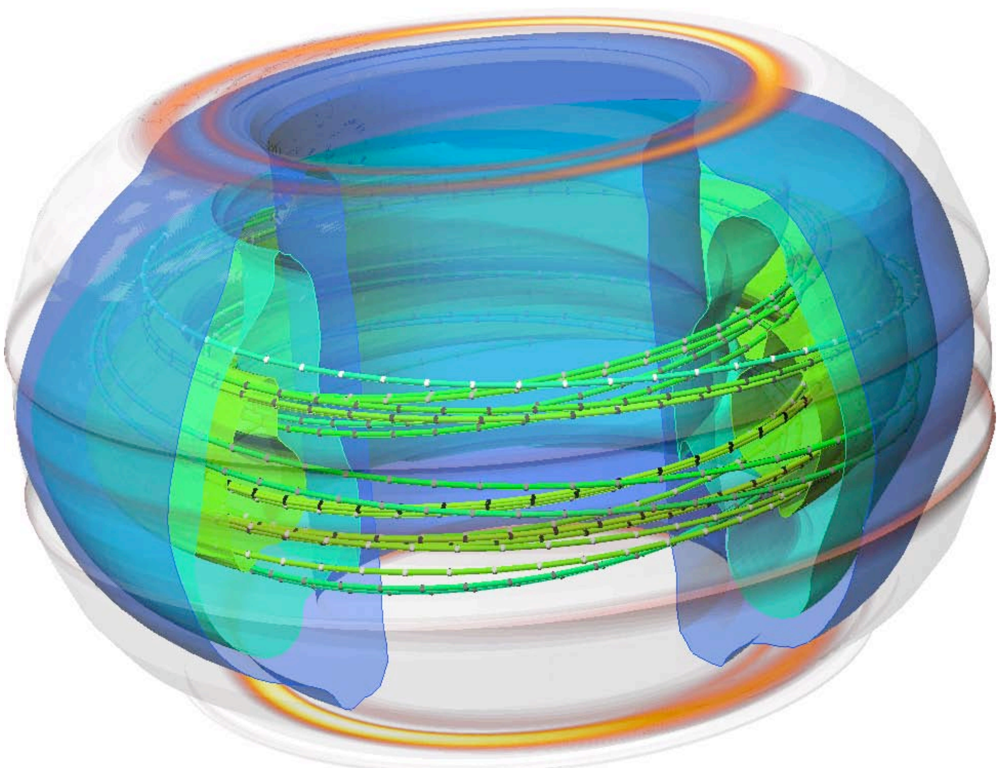

SciDAC CEMM NIMROD Simulation of a DIII-D Plasma

Raising the challenge of very large datasets
- MDSplus
- Storage method
- Data location
- Parallel I/O

FusionGRID www.fusiongrid.org

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM

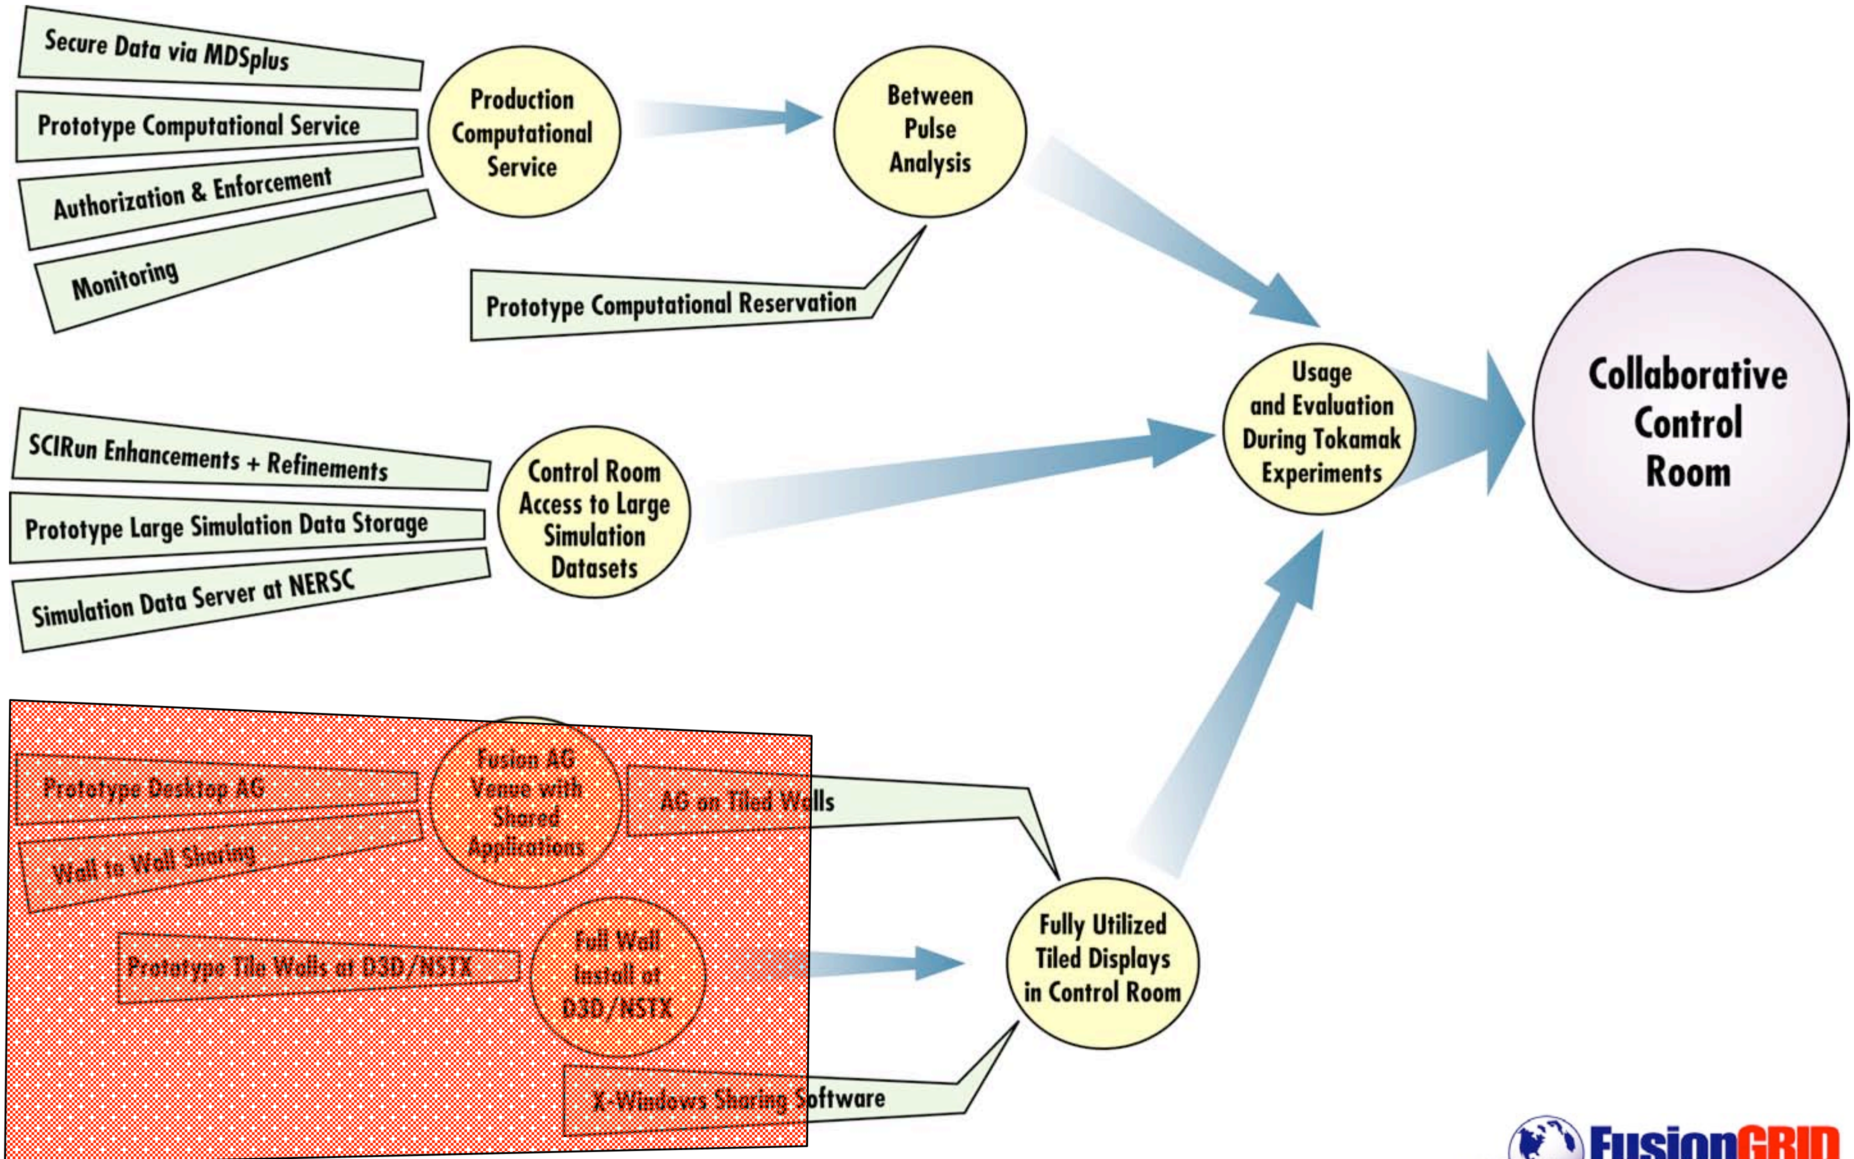

TILED DISPLAYS INSTALLED IN FUSION CONTROL ROOMS

DIII-D Tokamak Control Room

NSTX Tokamak Control Room

- Enhanced collaboration within the control room
  - Software for application sharing to tiled walls

- Very well received by fusion scientists
  - Fusion research funds used to purchase tiled walls for control rooms

FusionGRID
www.fusiongrid.org

ACCESS GRID: REAL TIME COMPLEX COMMUNICATION

AG in DIII-D Tokamak Control Room - July 2003

- Multi-site participants
  - Rich collaborative environment
  - Includes application & data sharing
- Modest cost of entry
  - Open source software
  - Commodity hardware

- Being used for seminars, working meetings, tokamak operations

Personal Interface to the Grid (PIG) motivated by Fusion research

FusionGRID
www.fusiongrid.org

SC03 DEMO: COLLABORATIVE CONTROL ROOM

- Fully interactive discussions utilizing AG
  - Includes shared applications
- Presence beyond AG communication
  - What one "sees and hears" in the control room
- Enhanced collaboration within the control room
  - Tiled displays and a shared X environment
- Advance reservation computation
  - Between pulse data analysis

FusionGRID
www.fusiongrid.org

COLLABORATIVE CONTROL ROOM: A SENSE OF PRESENCE

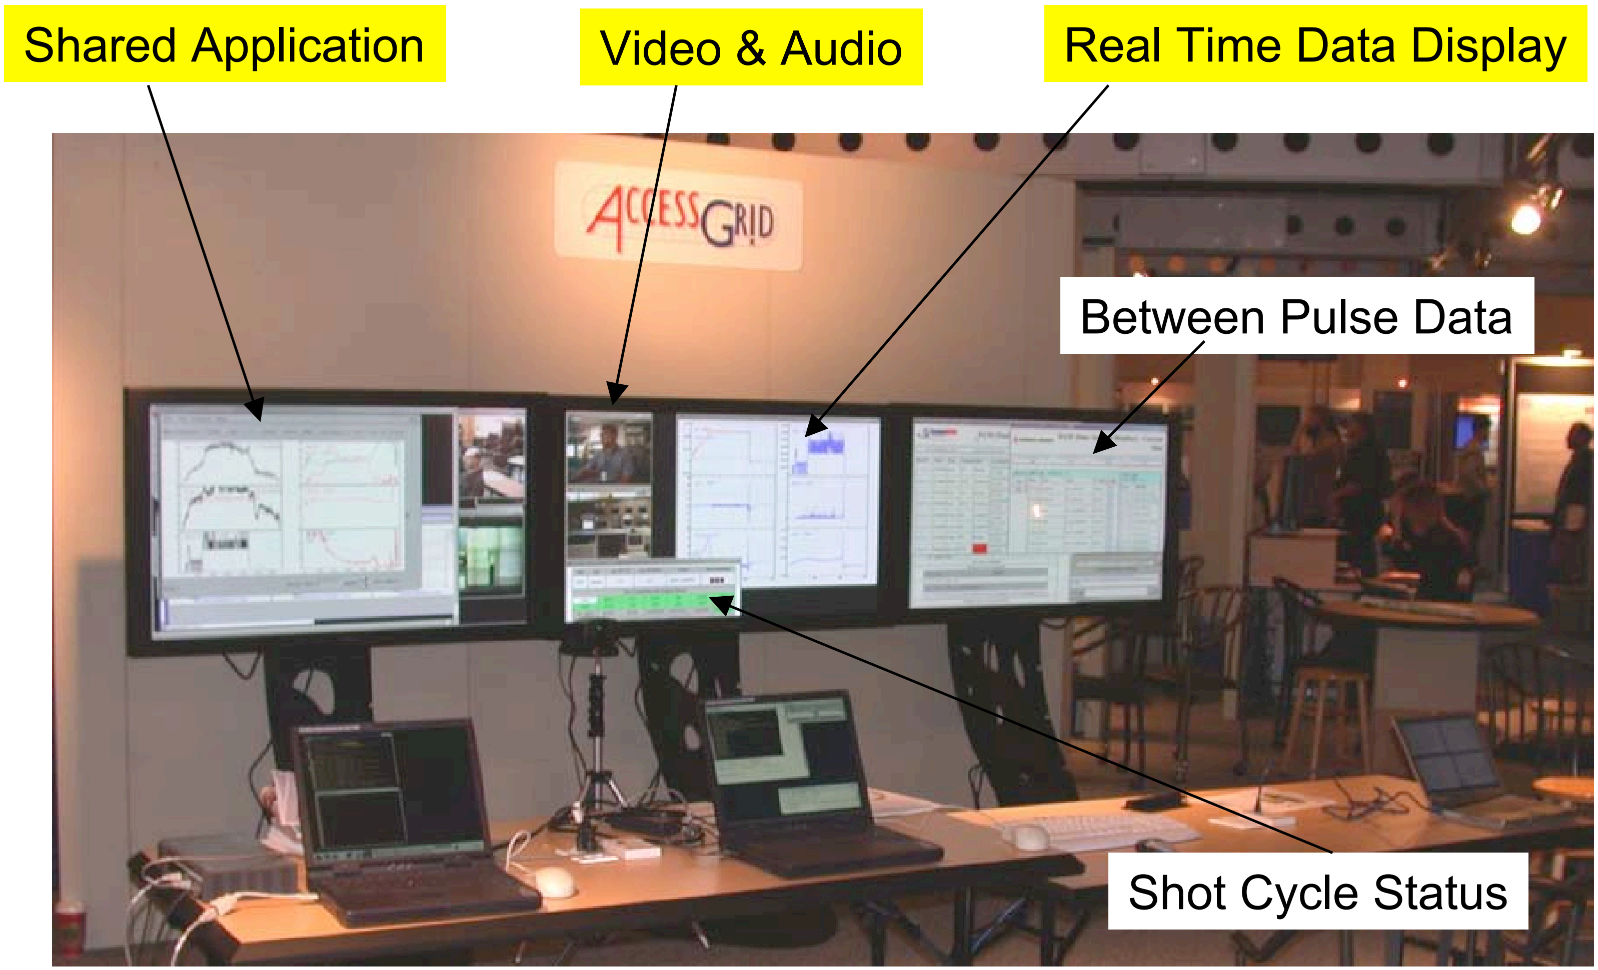

SuperComputing 2003, Phoenix AZ

REMOTE LEADERSHIP OF THE JET TOKAMAK IN ENGLAND FROM SAN DIEGO USING FUSIONGRID SERVICES

January 2004, San Diego

Working with JET and the UK e-Science Programme

- First attempt for real science and it was successful
- Successful & subsequently done: Japan - US & US-Germany

FusionGRID www.fusiongrid.org

NFC TECHNOLOGIES SCALE TO THE NEXT DEVICE

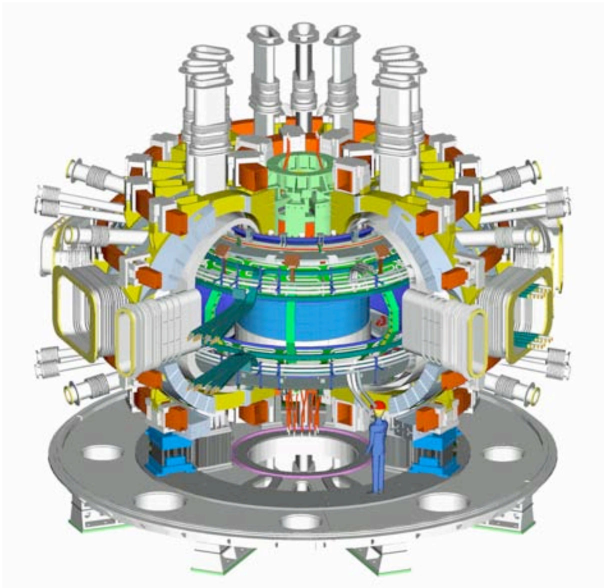

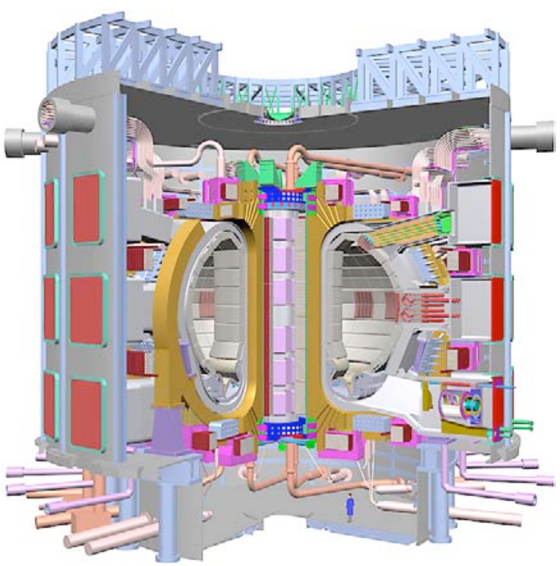

- Pulsed experiment with simulations
  - ~TBs of data in 30 minutes
- Successful operation requires
  - Large simulations, shared vis, decisions back to the control room
  - Remote Collaboration via FusionGrid
- NFC technology being discussed as the model for ITER: ITER-Grid
  - ITER workshop & OFES presentations
- KSTAR can be a proving ground
  - Software is the challenge, not networks

FusionGRID
www.fusiongrid.org

REACHING OUT TO THE ICC COMMUNITY

- Reach critical mass of people via remote participation
  - Staff multiplier if full-time travel not required

- Diagnostician, analysis code, seminars & working meetings

May 2004, ICC Madison WI

Demonstrations: Grid Computing & AG

Invited Talk with AG to MIT, ANL, GA

FusionGRID
www.fusiongrid.org

LESSONS LEARNED AND OUTSTANDING ISSUES

- Certificate management for users and developers too difficult
  - This is their first experience with FusionGrid: needs to be positive
- Software infrastructure required for a new service is too complex
  - Simple for the non-specialist (Professor & grad student)
- Difficulties combining Grid-security and Site-security (firewalls)
  - Greatly limiting the potential expansion of the FusionGrid userbase
- Manipulating large multi-dimensional datasets is still a challenge
  - Need to test new approaches
- Control room presence is more than audio/video & shared apps
  - Include things one sees & hears when physically in control room
- Users like frequent and rapid prototyping tests
  - They feel involved and it is educational to both sides

PROJECT RENEWAL SUBMITTED WITH ADJUSTED FUNDING TO REFLECT NEW PRIORITIES

- Fusion community has responded positively to the work of the NFC
  - Endorsed by the FFCC, C-Mod, DIII-D, NSTX, JET, MAST
  - PSACI PAC, User Oversight Committee
  - Velikhov as Chairman of the Board for ITER Technical Activities
  - Various SciDAC projects (both fusion and CSET)

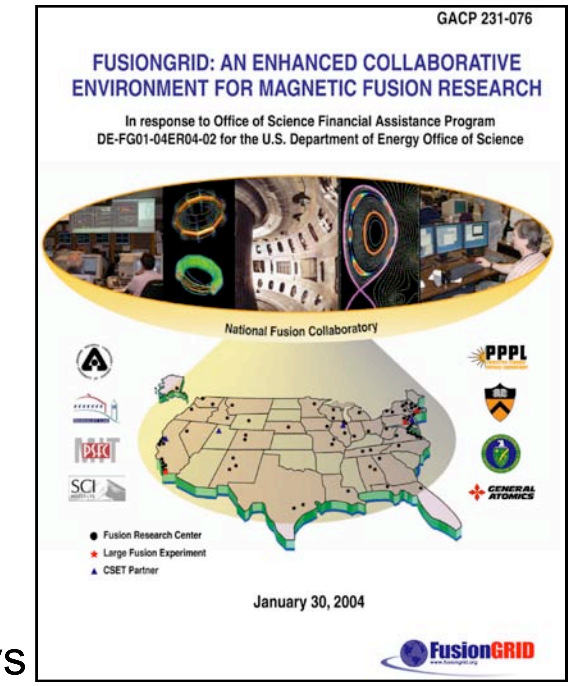

- Ease-of-use is a key area of focus
  - Lower the barriers for new users to join
  - Lower the barriers for adding new services
  - GYRO & TORIC as new services
  - Efficient data management and visualization of large simulation datasets

- Collaborative tokamak control room
  - Enhanced AG along with control room presence
  - Co-located and remote X sharing to tiled displays
  - Advanced computational reservation on FusionGrid

FusionGRID
www.fusiongrid.org

CONCLUDING COMMENTS

- The National Fusion Collaboratory Project is implementing and testing new collaborative technologies for fusion research
  - FusionGrid services being used to benefit daily FES research
- Clear vision forward to the collaborative control room
  - Concept encompasses most if not all collaborative FES needs
- Collaborative technology critical to the success of the FES program
  - Experimental: Fewer, larger machines in future (ITER)
  - Computation: Moving toward integrated simulation (FSP)

First on our list is fusion. The prospect of limitless source of clean energy for the world leads with our commitment to join the international fusion energy experiment known as ITER.
– Secretary of Energy Spencer Abraham, November 10, 2003
Introducing the Department's 20-year plan for building the scientific facilities of the future.

FusionGRID
www.fusiongrid.org