The National Fusion Collaboratory

A DOE National Collaboratory Pilot Project

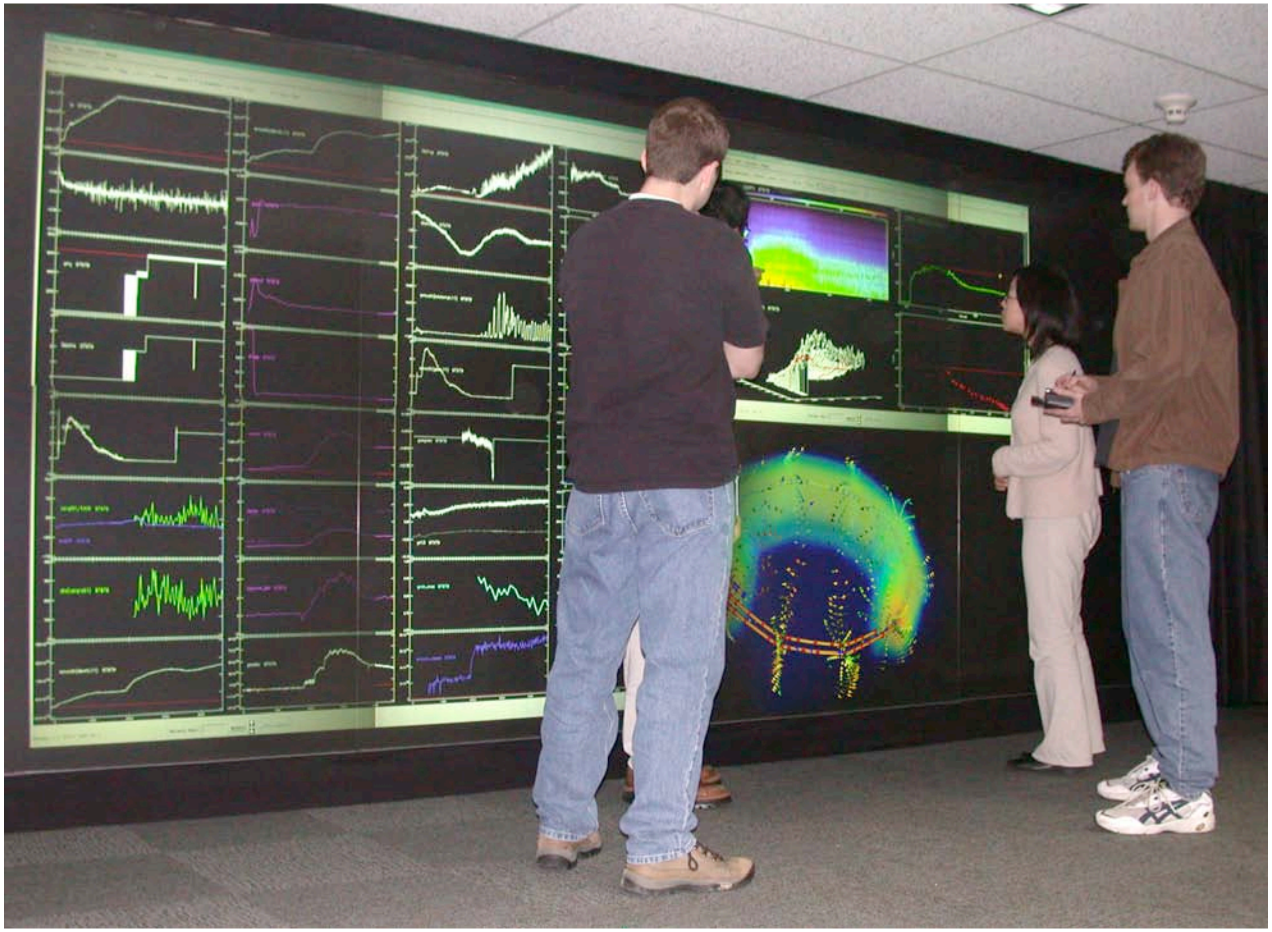

Presented by
David P. Schissel
Martin Greenwald
Doug McCune

at
The US–KSTAR
Workshop
May 20, 2004
San Diego, CA

PRESENTATION'S KEY POINTS

- Collaborative technology critical to the success of the FES program
  - Experimental: Fewer, larger machines in future
  - Computation: Moving toward integrated simulation

- The National Fusion Collaboratory Project is implementing and testing new collaborative technologies for fusion research
  - FusionGrid services being used daily to benefit FES research

- The work and experience gained by the National Fusion Collaboratory Project is directly applicable to KSTAR

NATURE OF FUSION RESEARCH DRIVES REQUIREMENTS FOR COMPUTING AND NETWORKING

- Experiments
  - Real time interactions of large, geographically extended teams
  - Faster between-pulse analysis translates directly to productivity
  - 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

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

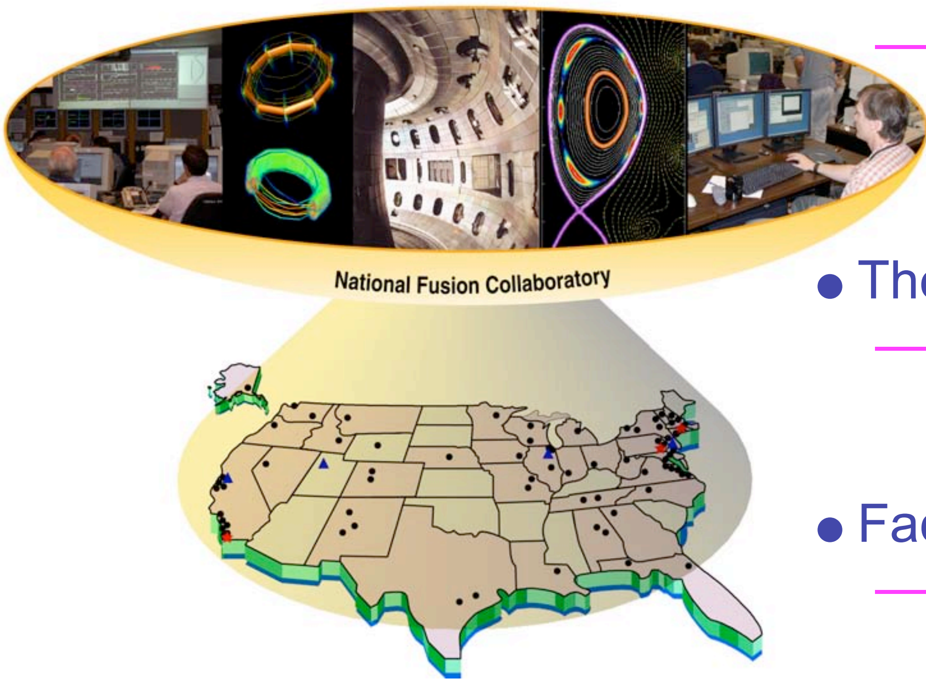

- Experimental Facilities
  - More efficient use resulting in greater progress with less cost
- Theory & Modeling
  - Integrate theory & experiment
- Facilitate multi-institution collaboration
  - Integrate geographically diverse groups
- Create standard tool set
  - To build in these services in the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

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

FusionGRID
www.fusionqrid.org

THE NFC PROJECT BENEFITS FROM A DIVERSE TEAM

Synergistic benefits derived from interdisciplinary interactions

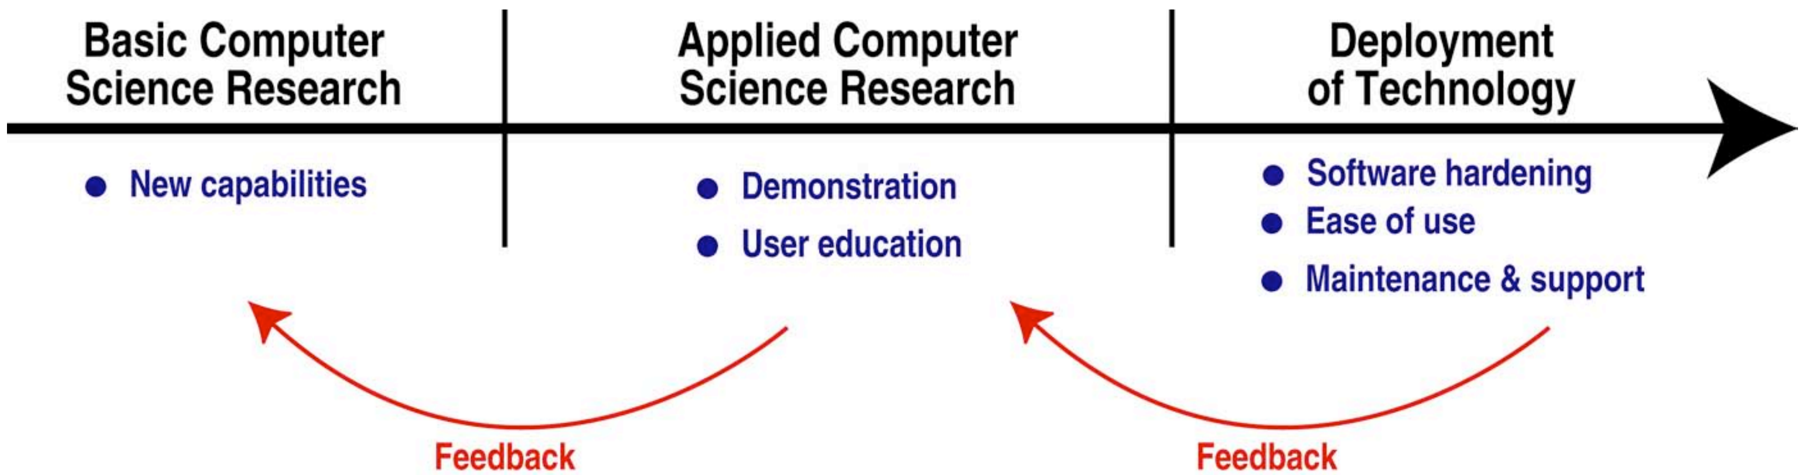

- ANL: Distributed Systems Lab
- ANL: Futures Lab
- General Atomics: DIII-D Fusion Lab
- LBNL: Distributed Systems
- MIT: C–Mod Fusion Lab
- Princeton Computer Science
- PPPL: NSTX Fusion Lab
- Utah: Scientific Computing & Imaging

FusionGRID www.fusiongrid.org

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

- 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

FusionGRID
www.fusiongrid.org

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

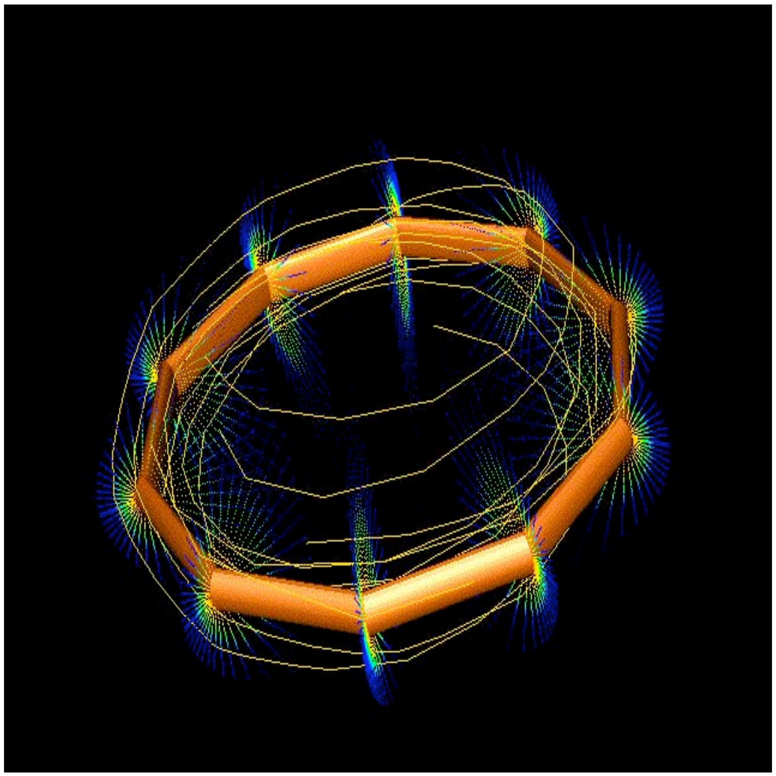

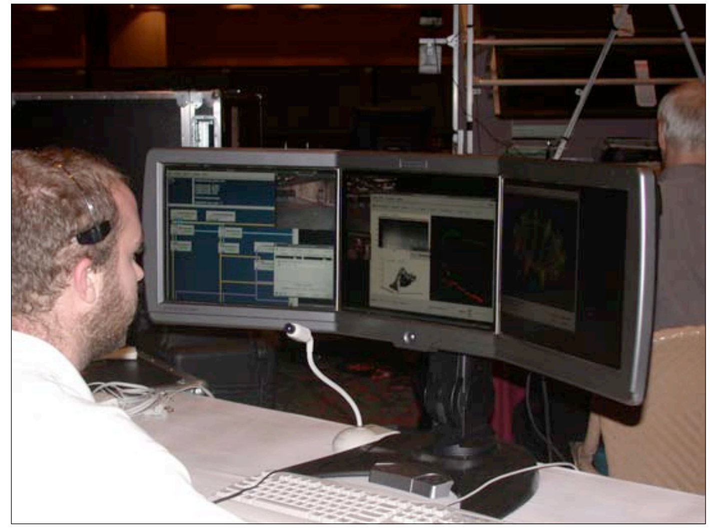

- Use of extended tool sets for remote collaboration
  - Flexible collaboration environment
  - Shared applications

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

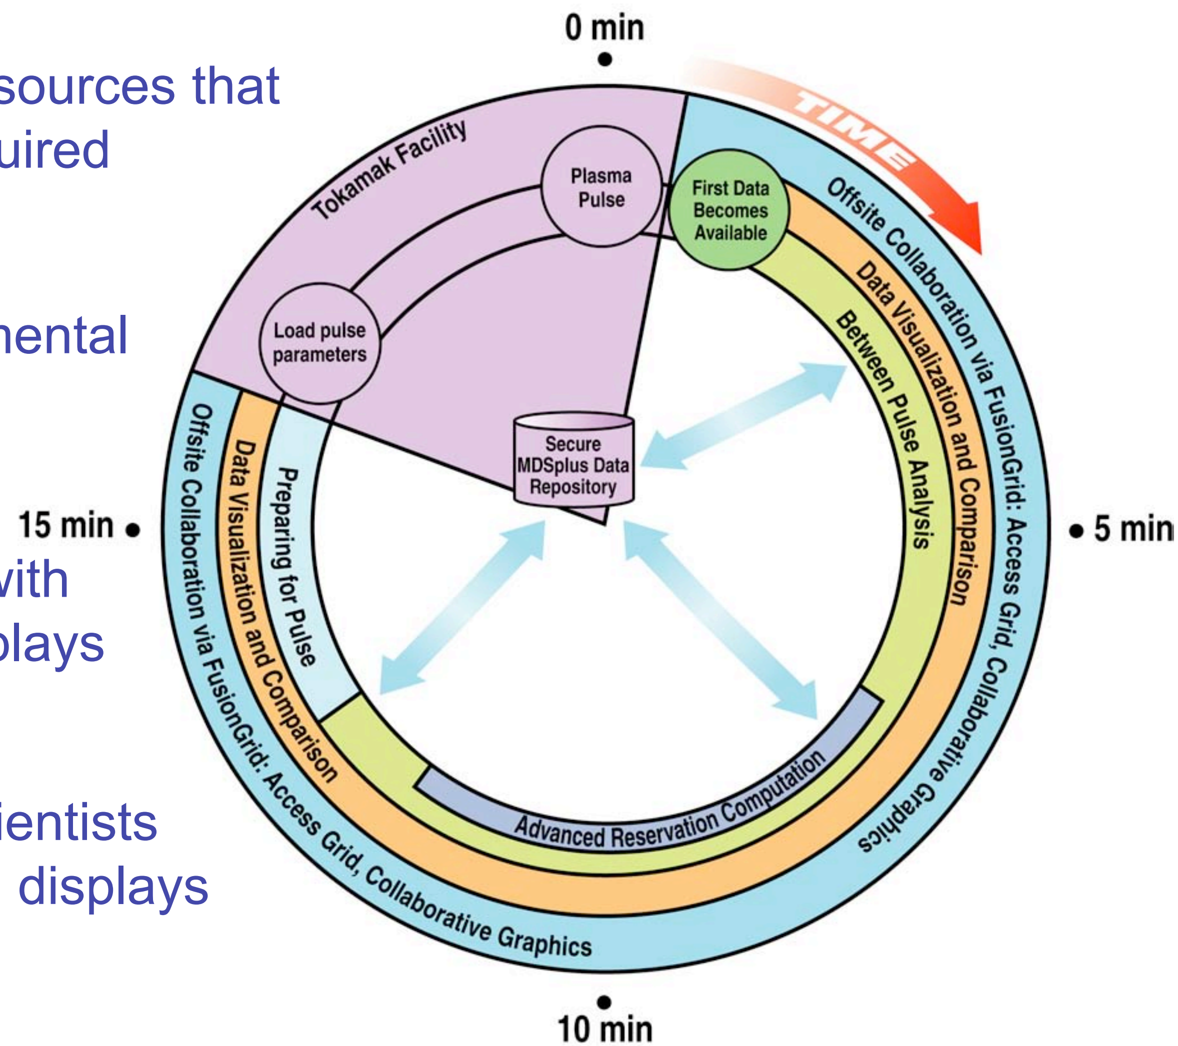

FusionGRID www.fusiongrid.org

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM

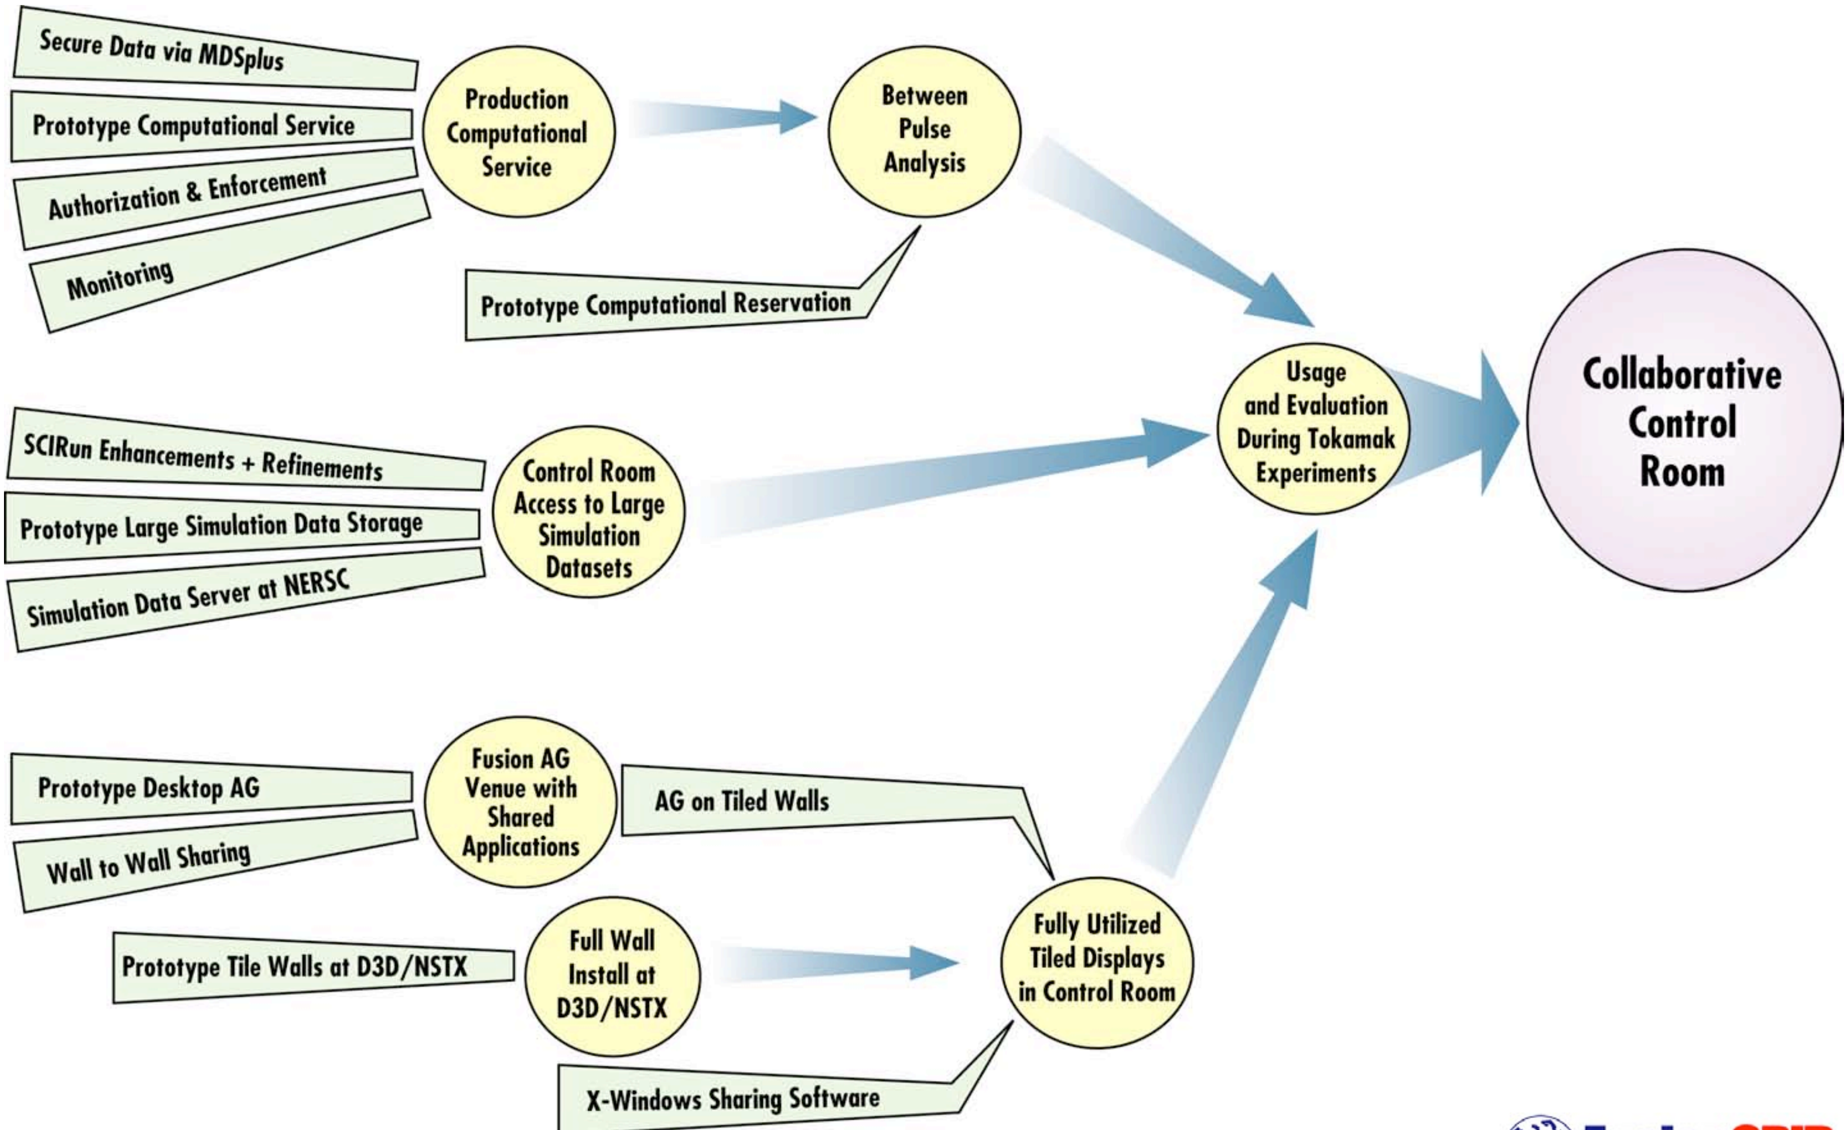

FusionGRID www.fusiongrid.org

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

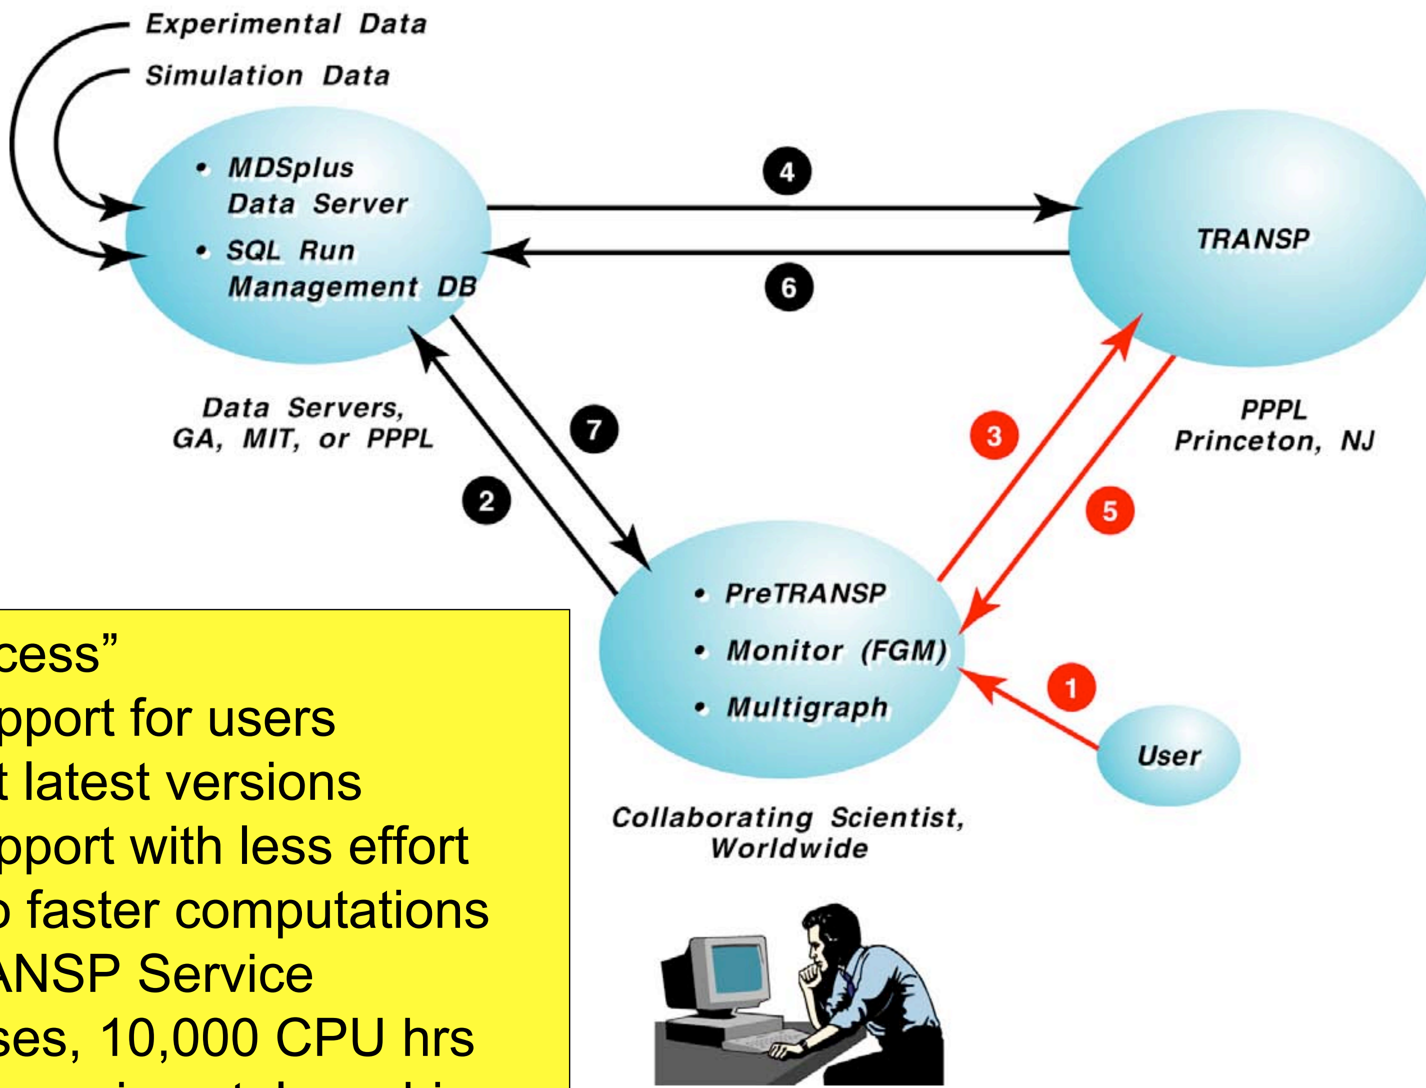

"This is a success"
- Better support for users
- Users get latest versions
- Better support with less effort
- Access to faster computations

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

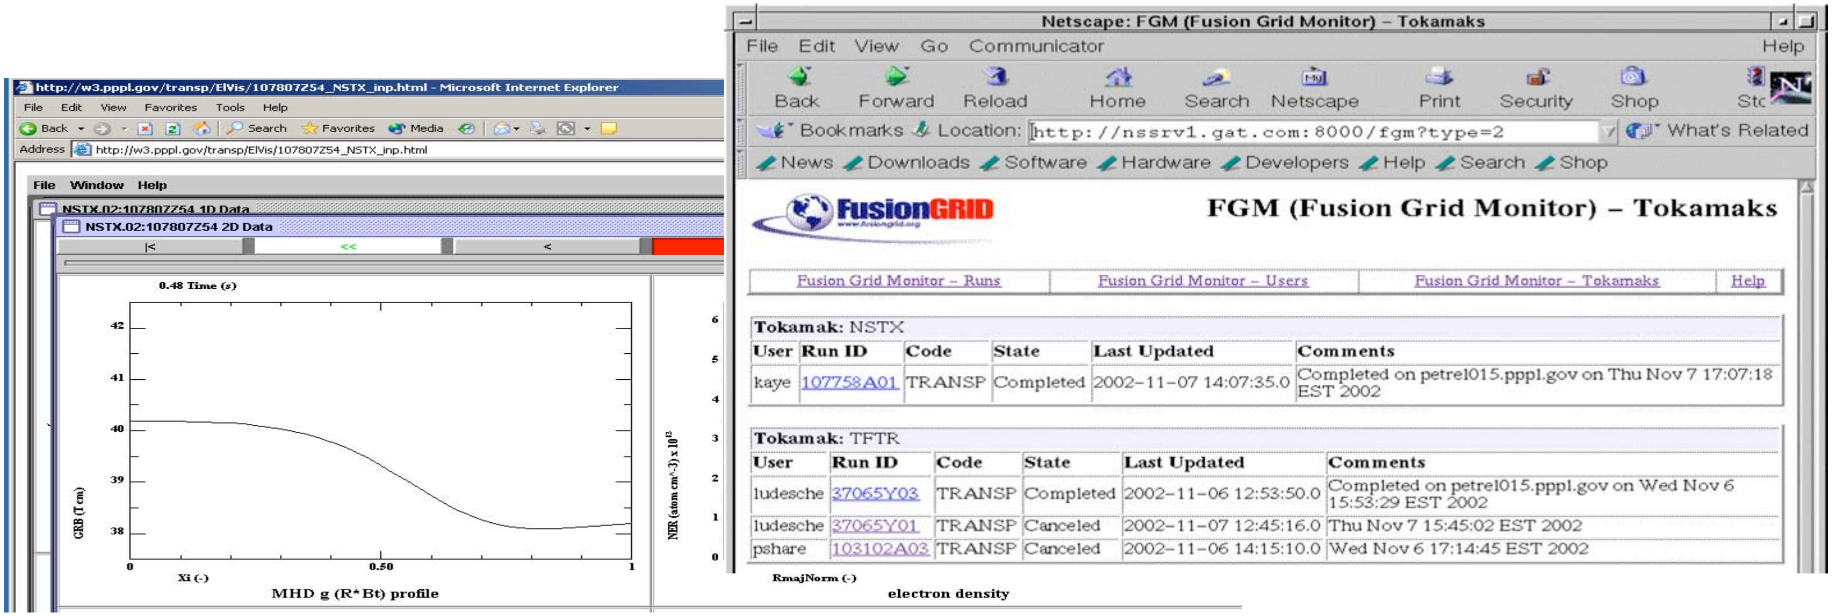

- Users track and monitor the state of applications on FusionGrid
  - Output dynamically via HTML, Built as Java Servlet (JDK2.1)
- Code maintenance notification
  - Users notified, queuing turned off, code rebuilt, queue restarted
- Results of simulation visualized during run
  - Both input and output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

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

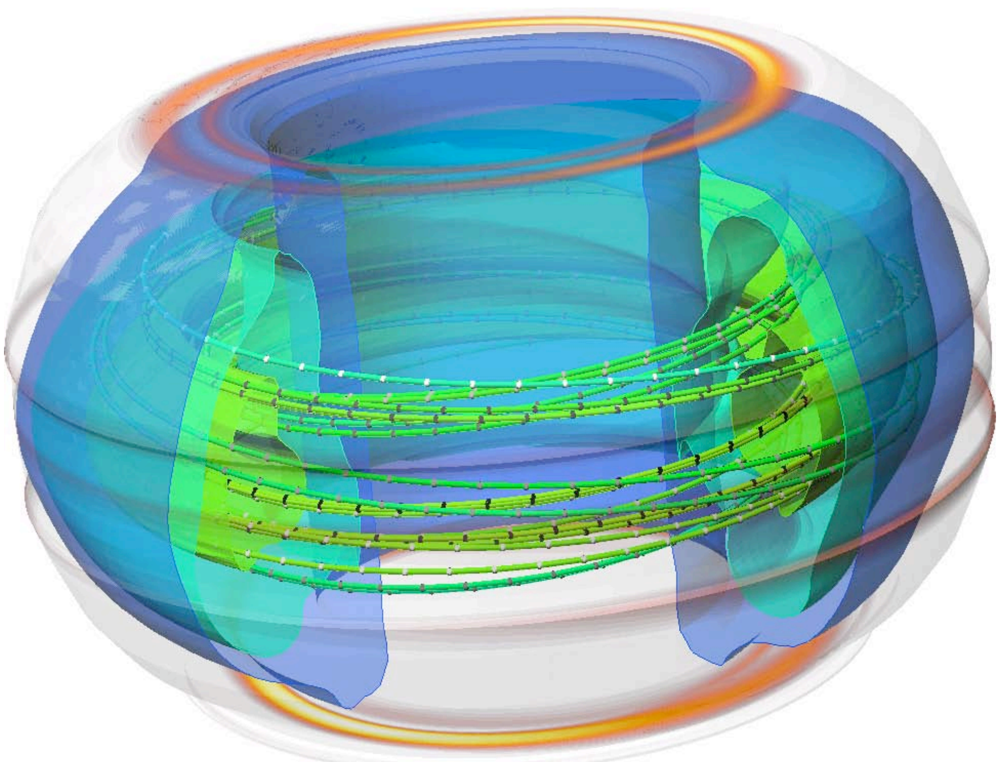

SciDAC CEMM NIMROD Simulation of a DIII-D Plasma

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

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

ACCESS GRID: REAL TIME COMPLEX COMMUNICATION

DIII-D Tokamak Control Room - July 2003

- Tested with off-site scientist to control room
  - Includes application sharing
  - Detailed data analysis discussion
- Feedback indicated the need for a greater control room presence for the off-site scientist
- 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

Shared Application

Video & Audio

Real Time Data Display

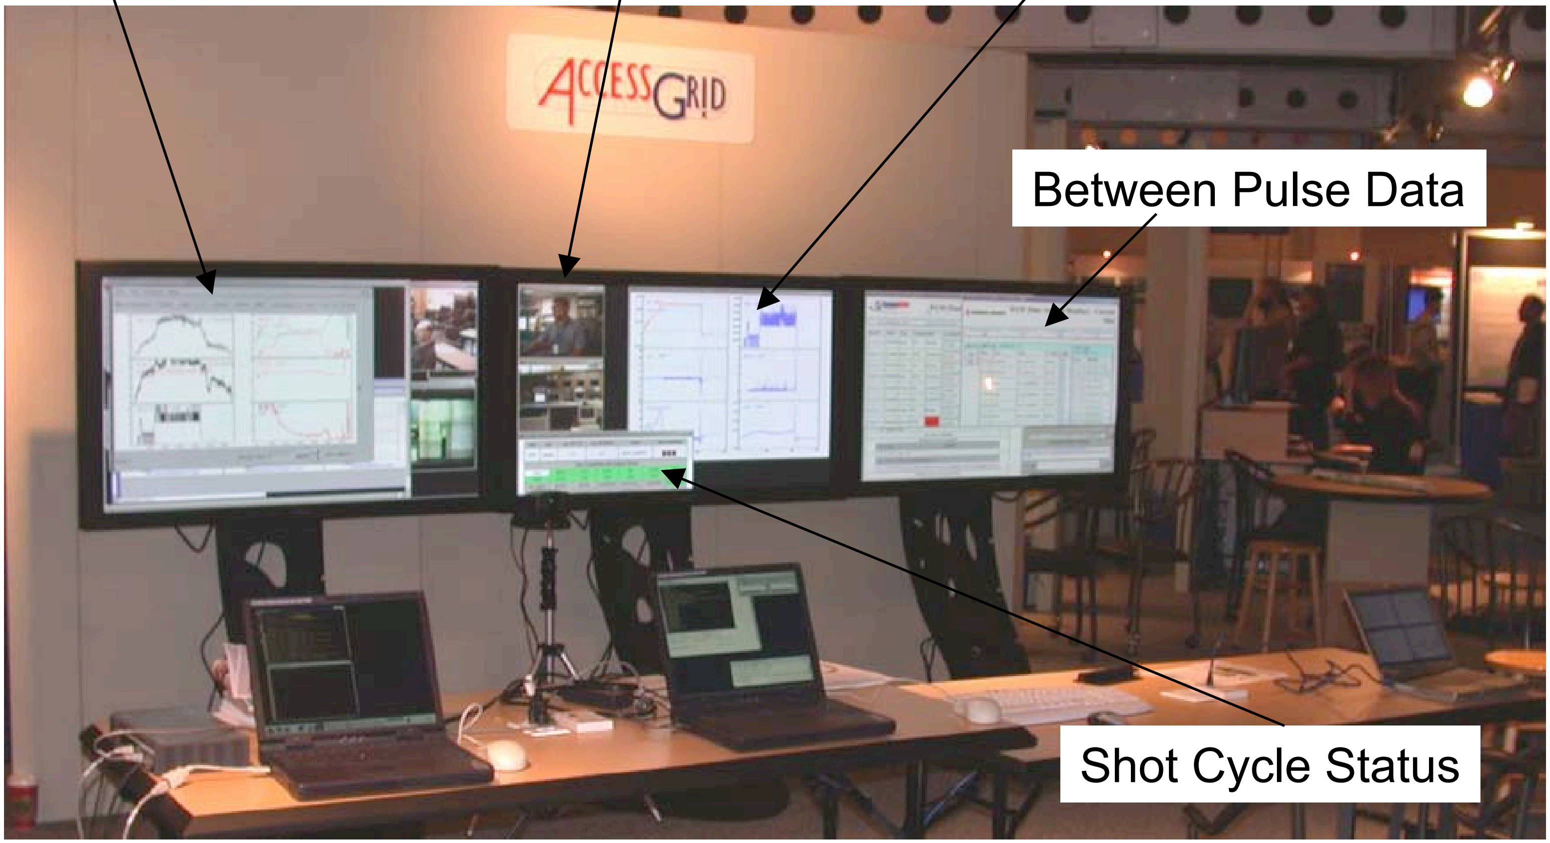

Between Pulse Data

Shot Cycle Status

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
- Similar collaborations for Japan - US and US-Germany

FusionGRID
www.fusiongrid.org

NFC TECHNOLOGIES SCALE TO THE NEXT DEVICE

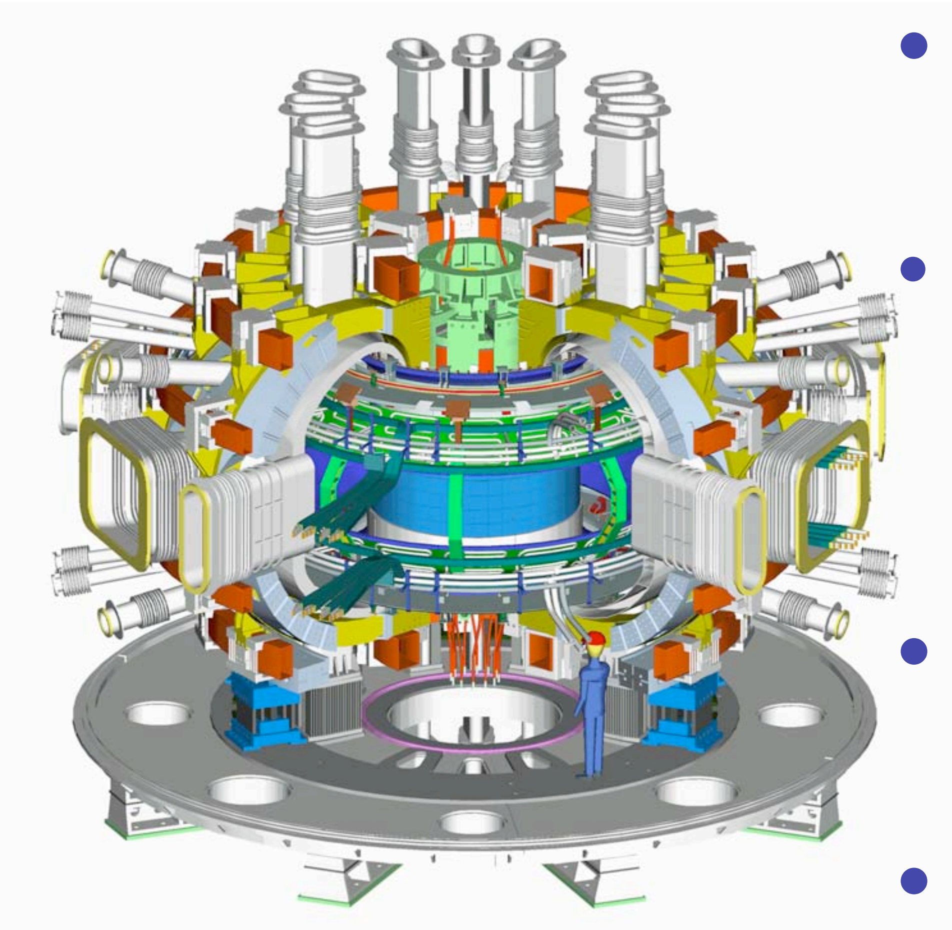

- Pulsed experiment with simulations
  - ~TBs of data in 30 minutes
- Successful operation requires
  - Large simulations, shared vis, decisions back to the control room
  - Remote Collaboration via FusionGrid
- NFC technology being discussed as the model for ITER: ITER-Grid
- KSTAR can be a proving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

FusionGRID
www.fusiongrid.org

CONCLUDING COMMENTS

- Collaborative technology critical to the success of the FES program
  - Experiment: Fewer, larger machines in future
  - Computation: Moving toward integrated simulation
- The National Fusion Collaboratory Project is implementing and testing new collaborative technologies for fusion research
  - FusionGrid services being used daily to benefit FES research
- Clear vision towards the Collaborative Control Room
  - Concept encompasses most if not all FES needs
- The work and experience gained by the National Fusion Collaboratory Project is directly applicable to KSTAR

FusionGRID
www.fusiongrid.org

Potential U.S. Role in Deploying an Advanced Collaborative Environment for KSTAR

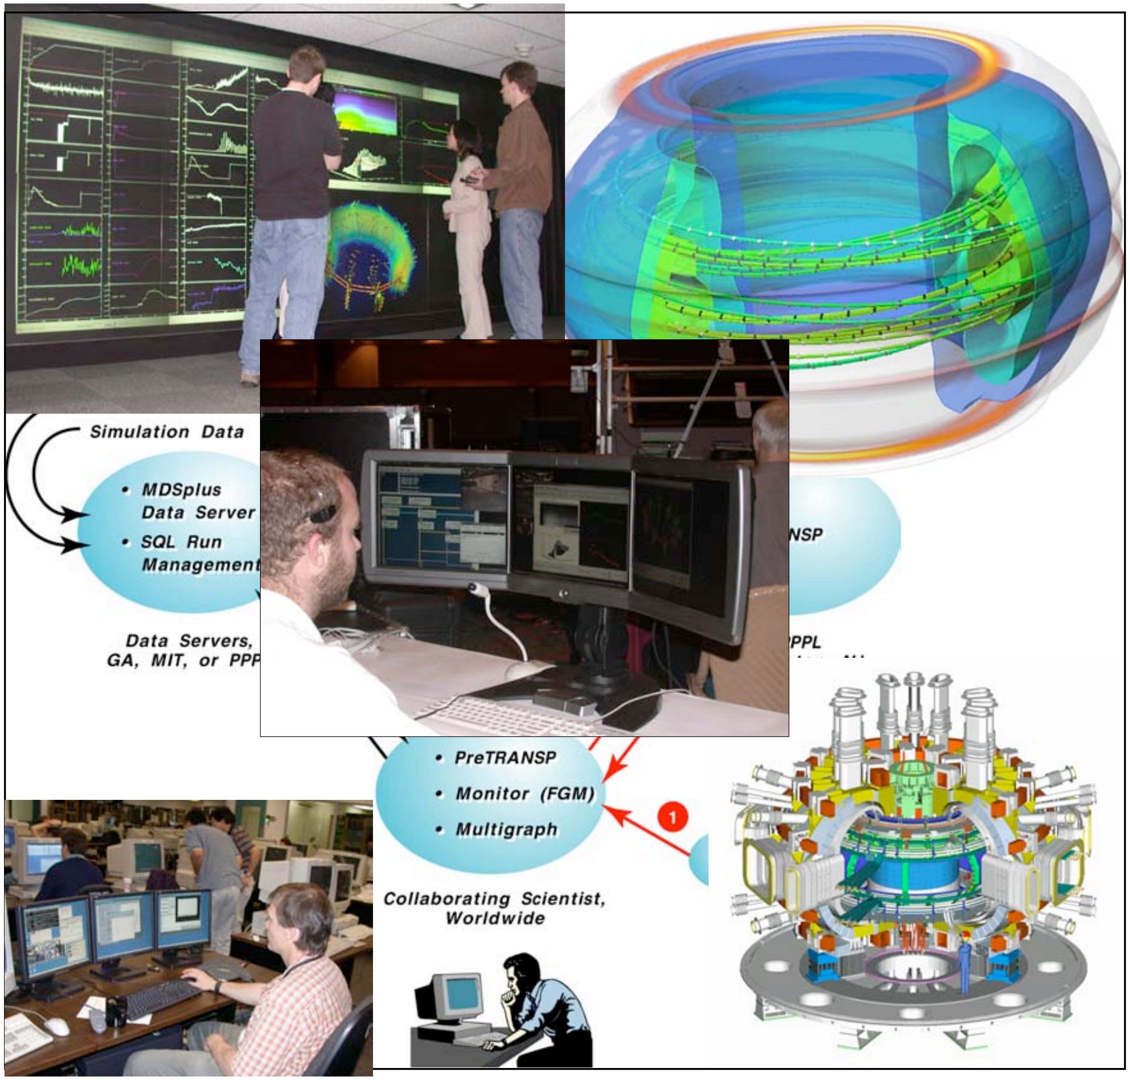

Presented by
David P. Schissel
Martin Greenwald
Doug McCune

at
The US–KSTAR Workshop
May 20, 2004
San Diego, CA

FusionGRID
www.fusiongrid.org

OVERVIEW OF POSSIBLE U.S. ROLES

- MDSplus for data acquisition, storage, management for KSTAR
  - Deployment and training in the near term
  - Long-pulse extension in concert with other projects
- Grid computing to support KSTAR design and data analysis
  - Deployment and training of Grid infrastructure
  - TRANSP or other codes critical to KSTAR success
- Remote collaboration capability
  - Deployment and training of Access Grid, VNC, related technologies
  - KSTAR control room design and implementation
- Although these integrate together than can be done separately
- KSTAR work can leverage from the Fusion Collaboratory Project and work from other FusionGrid participants - C-Mod, DIII-D, NSTX, JET, MAST, Asdex-U, …

FusionGRID www.fusiongrid.org

MDSplus IS A COMPLETE DATA MANAGEMENT SYSTEM

- MDSplus is a set of tools
  - Data acquisition
  - Data storage
  - Data management
  - Data Grid interface
  - Secure remote access
- Designed for pulsed experiments
  - Presently not able to look at data as it is streaming in
- Collaboratively developed
  - Initially MIT, IGI (RFX), and LANL
  - Unix port done by MIT, IGI (RFX), GA, LLNL, and PPPL
  - Today a worldwide open source development community
- Widely used in the fusion community
  - Over 30 sites

FusionGRID
www.fusiongrid.org

MDSplus CAN BE DEPLOYED TO KSTAR TODAY

- Short pulse required no significant modifications
  - C-Mod presently takes 1.2 GB/shot in ~4 minutes
    - Limited by CAMAC I/O speed
  - Distributed implementation allows for scalability

- MDSplus for KSTAR
  - U.S. can provide consulting, training, and documentation
  - If desired, U.S. can take a more active role integrating MDSplus into the KSTAR environment by writing code

EXTENDING MDSplus to LONG PULSE DISCHARGES

- Long pulse capability not needed in the short term by KSTAR

- Significant but tractable effort required
  - Approximately 10 man years, 3 to 4 people for 3 years
  - Could be carried out as part of the U.S. – KSTAR collaboration

- Long pulse capability desired by the international MDSplus community
  - Long pulse extension could be co-supported by all interested parties
  - W7X and ITER have expressed interest

GRID COMPUTING FOR KSTAR

- Deployment of the infrastructure for TRANSP as a Grid service
  - Include documentation and training

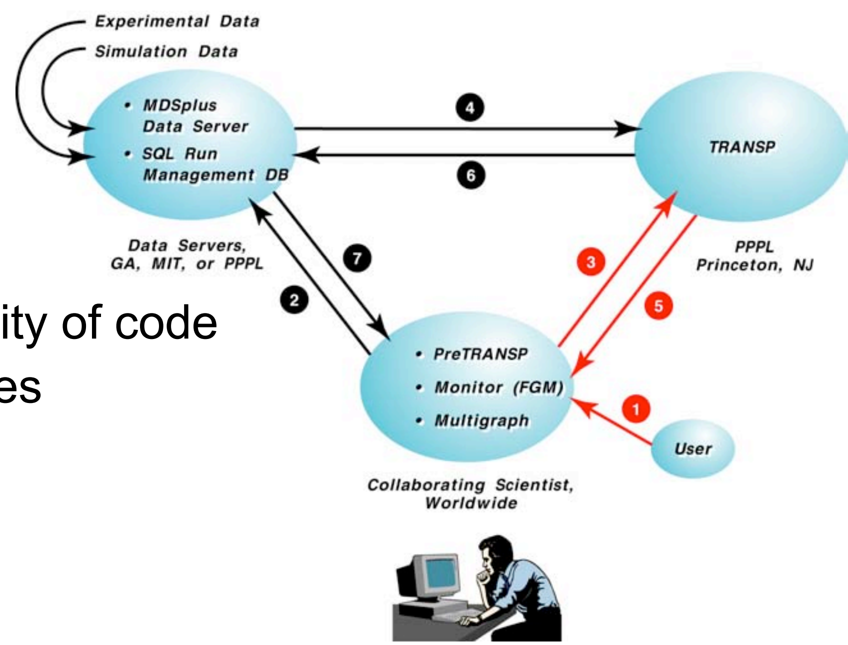

- Other codes can be deployed in addition or instead
  - Time required depends on complexity of code
  - Grid computing supports many codes
  - Support coupled services (e.g. parallel TORIC & TRANSP)

- Grid computing infrastructure includes
  - Security infrastructure: authentication and authorization
  - Data management and run monitoring
  - Code GUI

ADVANCED COLLABORATIVE ENVIRONMENT FOR KSTAR

- Deployment of the Access Grid
  - From desktop to room size as required
  - Include documentation and training
  - Network configuration, hardware

- Collaborative Control Room
  - KSTAR design
  - Software implementation
  - Testing, documentation, & training