


# GRID PERFORMANCE AND FUSION SCIENCE



**PRESENTED BY  
JUSTIN BURRUSS**

**GRID PERFORMANCE  
WORKSHOP 2005**

**EDINBURGH, SCOTLAND  
JUNE 22-23, 2005**

burruss@fusion.gat.com  
<http://web.gat.com/~burruss/>

# ACKNOWLEDGEMENTS



- **U. S. Department of Energy**
  - OFES & OASCR (SciDAC)



- **DIII-D National Fusion Facility**
  - Operated by General Atomics



- **FusionGrid collaborators**
  - MIT, PPPL, LBL, ANL, Utah CS, Princeton CS



- **National eScience Centre**

# OUTLINE

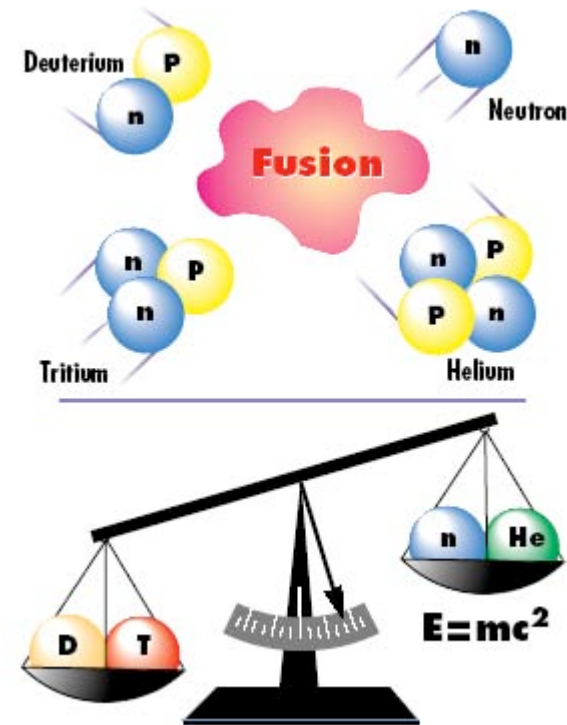
- **Background**
  - What is Fusion?
  - How is Fusion Science advanced?
- **Application**
  - Between-shot analysis
  - Monitoring & failure recovery
  - Moving from LAN to WAN
- **Performance Metrics**
  - Data analysis throughput
  - Reliability
- **The Perfect Performance Monitoring & Management System?**

# KEY POINTS

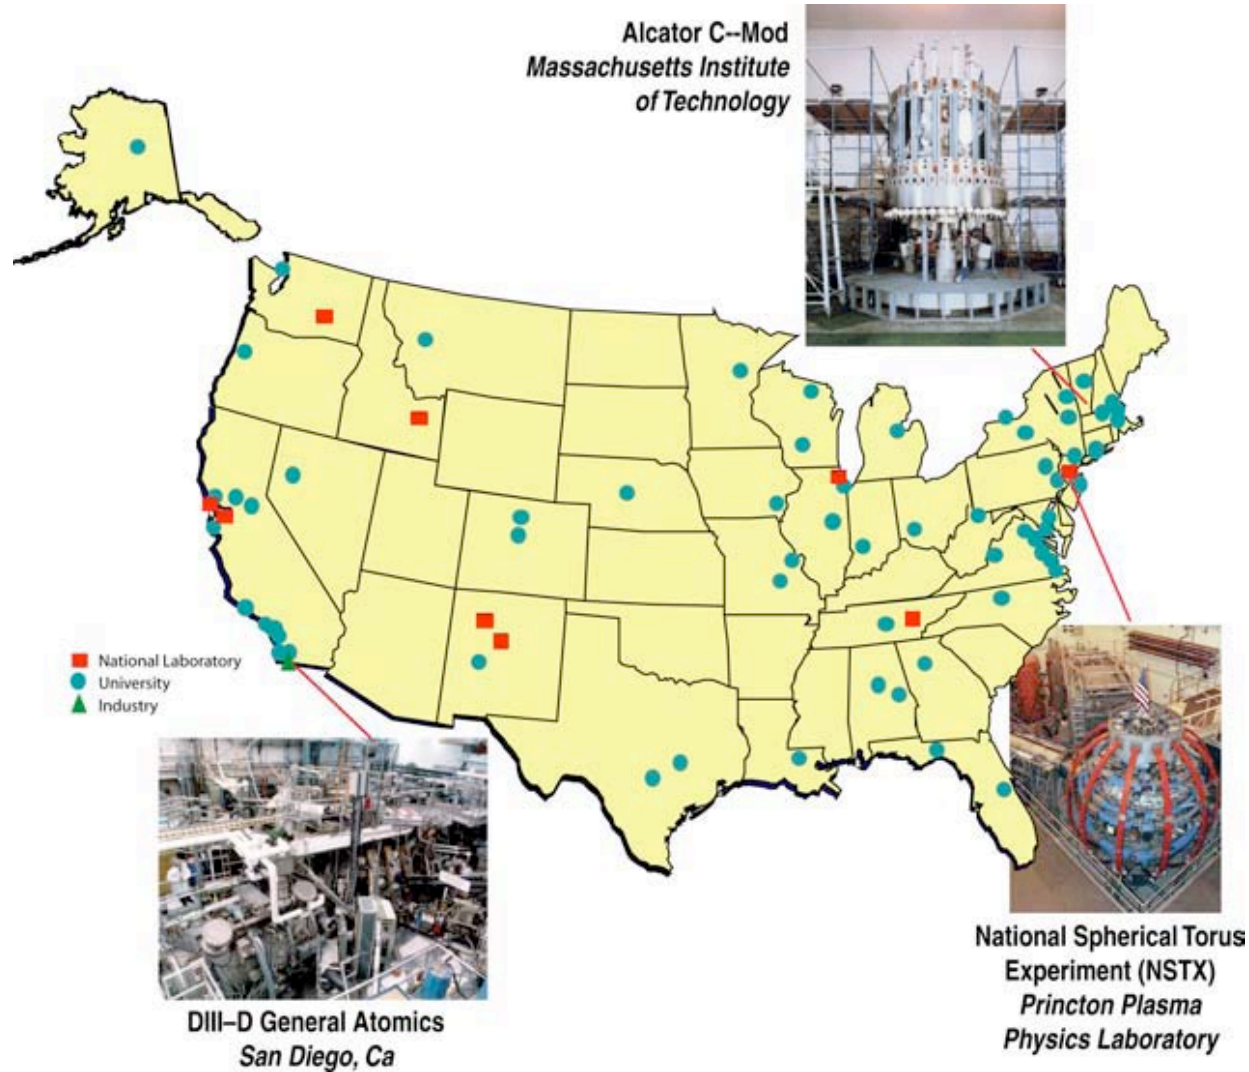
- **Fusion science seeks a new power source and is advanced by experiments on fusion devices (tokamaks) located worldwide**
- **Experiments require rapid data analysis**
  - Not batch analysis
  - Not “needle in the haystack” searching
  - Thousands of measurements that build a coherent picture
- **Data analysis throughput and reliability are the important performance metrics**
- **Moving data analysis from LAN to WAN brings many benefits, but introduces new performance issues**
- **We look forward to collaborating to make this happen**

# FUSION SCIENCE SEEKS ENVIRONMENTALLY & ECONOMICALLY ATTRACTIVE POWER PLANT


- Fusion is when you combine two atoms into one atom
- Energy is released from this fusion reaction
- **An attractive power source**
  - Abundant fuel available to all nations
  - Environmentally friendly
  - No proliferation risk
  - Can't blow up/melt down
  - Not subject to weather/seasonal issues
  - Concentrated relative to wind/solar



# FUSION RESEARCH TAKES PLACE ACROSS THE U. S.




# FUSION RESEARCH TAKES PLACE WORLDWIDE




*DIII - D*



*JET*



*The ITER project*



*JT - 60 U*

## Fusion Centres in the World



# FUSION RESEARCH TODAY IS A TEAM EFFORT



- 90 institutions participate
- 425 active users
- 317 scientific authors
- Students and faculty from
  - 65 universities
  - 28 states

## Active Collaborations 2004

### US Labs

ANL (Argonne, IL)  
 LANL (Los Alamos, NM)  
 LBNL (Berkeley, CA)  
 LLNL (Livermore, CA)  
 ORNL (Oak Ridge, TN)  
 PPPL (Princeton, NJ)  
 SNL (Sandia, NM)

### Industries

Calabasas Creek (CA)  
 CompX (Del Mar, CA)  
 CPI (Palo Alto, CA)  
 Digital Finetec (Ventura, CA)  
 DRS (Dallas, TX)  
 DTI (Bedford, MA)  
 FAR Tech (San Diego, CA)  
 IOS (Torrance, CA)  
 Lodestar (Boulder, CO)  
 SAIC (La Jolla, CA)  
 Spinner (Germany)  
 Tech-X (Boulder, CO)  
 Thermacore (Lancaster, PA)  
 Tomlab (Willow Creek, CA)  
 TSI Research (Solana Beach, CA)

### US Universities

Auburn (Auburn, Alabama)  
 Colorado School of Mines (Golden, CO)  
 Columbia (New York, NY)  
 Georgia Tech (Atlanta, GA)  
 Hampton (Hampton, VA)  
 Lehigh (Bethlehem, PA)  
 Maryland (College Park, MD)  
 Mesa College (San Diego, CA)  
 MIT (Boston, MA)  
 Palomar (San Marcos, CA)  
 New York U. (New York, NY)  
 SDSU (San Diego, CA)  
 Texas (Austin, TX)  
 UCB (Berkeley, CA)  
 UCI (Irvine, CA)  
 UCLA (Los Angeles, CA)  
 UCSD (San Diego, CA)  
 U. New Mexico (Albuquerque, NM)  
 U. Rochester (NY)  
 U. Utah (Salt Lake City, UT)  
 Washington (Seattle, WA)  
 Wisconsin (Madison, WI)

### Russia

Ioffe (St. Petersburg)  
 Keldysh (Udмурtia, Moscow)  
 Kurchatov (Moscow)  
 Moscow State (Moscow)  
 St. Petersburg State Poly (St. Petersburg)  
 Trinit (Troitsk)  
 Inst. of Applied Physics (Nizhny Novgorod)

### European Community

Cadarache (St. Paul-lez, Durance, France)  
 Chalmers U. (Göteborg, Sweden)  
 CFN-IST (Lisbon, Portugal)  
 CIEMAT (Madrid, Spain)  
 Consorzio RFX (Padua, Italy)  
 Culham (Culham, Oxfordshire, England)  
 EFDA-NET (Garching, Germany)  
 Frascati (Frascati, Lazio, Italy)  
 FOM (Utrecht, The Netherlands)  
 Helsinki U. (Helsinki, Finland)  
 IFP-CNDR (Italy)  
 IPP (Garching, Greifswald, Germany)  
 ITER (Garching, Germany)  
 JET-EFDA (Oxfordshire, England)  
 KFA (Jülich, Germany)  
 Kharkov IPT, (Ukraine)  
 Lausanne (Lausanne, Switzerland)  
 IPP (Greifswald, Germany)  
 RFX (Padua, Italy)  
 U. Dusseldorf (Germany)  
 U. Naples (Italy)  
 U. Padova (Italy)  
 U. Strathclyde (Glasgow, Scotland)

### Japan

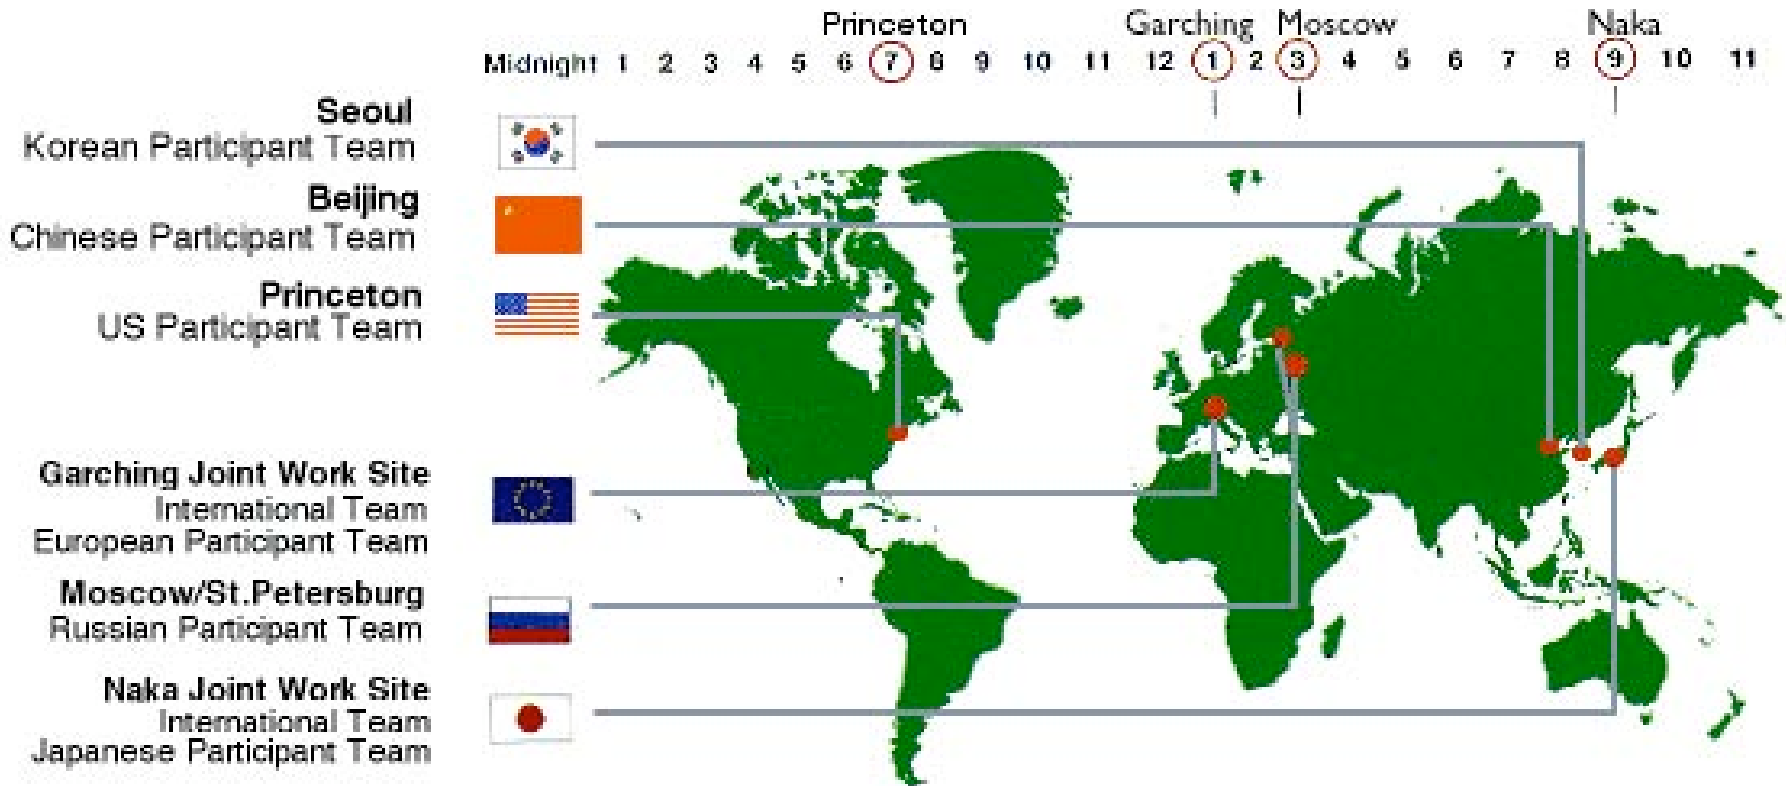
JAERI (Naka, Ibaraki-ken, Japan)  
 JT-60U  
 JFT-2M  
 Tsukuba University (Tsukuba, Japan)  
 NIFS (Toki, Gifu-ken, Japan)  
 LHD

### Other International

Australia National U. (Canberra, AU)  
 ASIPP (Hefei, China)  
 Dong Hau U. (Taiwan)  
 KBSI (Daegon, S. Korea)  
 KAERI (Daegon, S. Korea)  
 Nat. Nucl. Ctr. (Kurchatov City, Kazakhstan)  
 Pohang U. (S. Korea)  
 Seoul Nat. U. (S. Korea)  
 SWIP (Chengdu, China)  
 U. Alberta (Alberta, Canada)  
 U. of Kiel (Kiel, Germany)  
 U. Toronto (Toronto, Canada)

**DIII-D has active collaborators on four continents**


# FUSION RESEARCH WILL CONTINUE TO BE A TEAM EFFORT



The Six ITER Partners

# FUSION EXPERIMENTS CONSIST OF PULSES OR “SHOTS”


- **Fusion devices called tokamaks operate in a pulsed mode**
  - Turn on, then back off again
  - Called a “pulse” or “shot”
- **Raw data from sensors is collected during each shot**
- **Raw data is analyzed by codes, producing analyzed data**
- **Summary data is collected from analyzed and raw data**
- **The scientists can “see” what happened during a shot by looking at this raw, analyzed, and summary data**




*DIII-D Tokamak*

# DECISIONS MADE DURING FUSION EXPERIMENTS REQUIRE RAPID DATA ANALYSIS ON TOKAMAK PULSES

- Fusion experiments are not done in a batch mode: scientists “take a shot”, analyze the results, and move on to the next shot throughout the experiment
- This requires rapid data analysis
  - Scientists view results of previous shot to make informed decisions about how to configure the next shot
  - Decisions must be “locked in” at some point before shot
- Each tokamak shot is like a space shuttle launch—but done every 20 minutes




# EXPERIMENTAL FUSION SCIENCE IS AND WILL CONTINUE TO BE A DEMANDING ACTIVITY



Thousands of measurements and a team of specialists needed to build a coherent picture of what happened during a shot

# TODAY MOST BETWEEN-SHOT ANALYSIS IS LOCAL


24-CPU DIII-D Linux Cluster





# FUSIONGRID CREATED FOR BETTER USE OF RESOURCES

- **U. S. Fusion Grid (FusionGrid) aims to make more efficient use of computing resources**
  - Access is stressed rather than portability
  - Not CPU cycle scavenging or “distributed” supercomputing
- **Share resources between sites**
  - Reduce duplication of effort
  - Exploit comparative advantage
- **Develop a common tool set for fusion**
  - Globus Toolkit (GRAM & GSI)
  - FusionGrid CA & credential manager
  - ROAM
  - Secure MDSplus
  - Access Grid and VRVS



# FGM USED FOR GRID-WIDE MONITORING

- For general grid-wide monitoring, the FusionGrid Monitor is used
- Like DAM, but more generalized


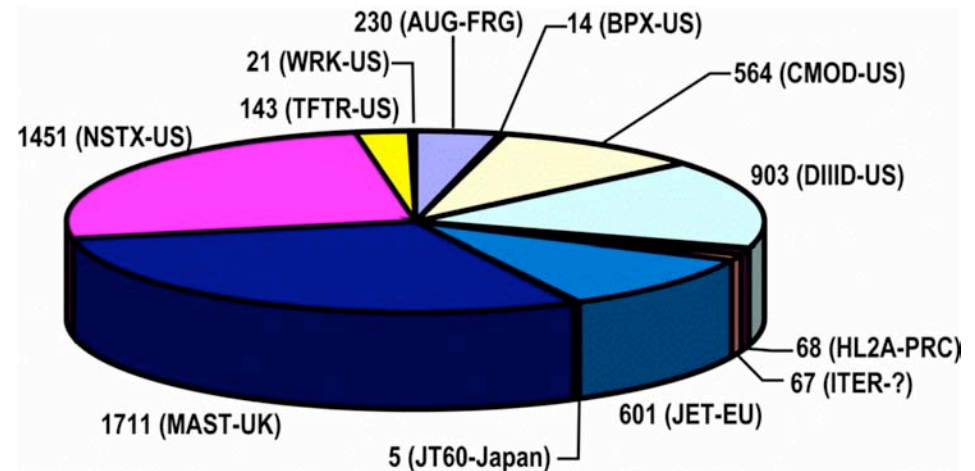


## FGM (Fusion Grid

<a href="#">Fusion Grid Monitor - Runs</a>		<a href="#">Fusion Grid Monitor - Users</a>		<a href="#">Fusion Grid Monitor - Tokamaks</a>	
<b>Tokamak: WRK</b>					
User	Run ID	Code	State	Last Updated	Comments
ludesche	<a href="#">37065Y01</a>	TRANSP	Completed	2005-06-13 10:11:05.0	Completed on falcon133.pppl.gov on
<b>Tokamak: MAST</b>					
User	Run ID	Code	State	Last Updated	Comments
psh0039	<a href="#">13035C02</a>	TRANSP	Completed	Tue Jun 14 09:21:18 PDT 2005	Completed on sunfire05.pppl.g
<b>Tokamak: D3D</b>					
User	Run ID	Code	State	Last Updated	Comments
murakami	<a href="#">117740T06</a>	TRANSP	Stopped	Tue Jun 14 01:42:03 PDT 2005	Tue Jun 14 04:41:14 EDT 2005: c
burruss	<a href="#">527</a>	GATO	Running	2005-06-13 17:05:14.0	GATO failed, status=1
burruss	<a href="#">526</a>	GATO	Completed	2005-06-13 17:06:01.0	GATO completed
burruss	<a href="#">525</a>	GATO	Running	2005-06-13 17:04:55.0	GATO failed, status=1
burruss	<a href="#">524</a>	GATO	Completed	2005-06-13 17:13:30.0	GATO completed

# CURRENTLY ADDING FUSIONGRID SERVICES


- **First FusionGrid service: TRANSP transport code**
- **The TRANSP code runs at PPPL and used throughout fusion**
  - Improved support
  - No code distribution issues
  - Result is users are more productive
- **But...what happens when you move to between-shot analysis?**



*FusionGrid TRANSP runs*

# FAILURES ARE SENT TO DAM AND FGM

- **Analyses post errors to DAM and FGM**
  - User can click link to view details in the log file
- **Can post “service is down” messages**
- **Authorization failures also sent to ROAM**
- **Errors may require closer debugging work to figure out underlying problem**
  - e.g. “failed to connect to server”
- **The ability to view logfiles through DAM/FGM has proved very useful for debugging**




# FAILURE RECOVERY IS NOT GENERALIZED

- **Failure recovery is specific to each component of FusionGrid**
- **Individual analyses may have their own mechanisms specific to that code**
  - e.g. try to connect three times before quitting
- **Secondary MyProxy and authorization servers are used**
  - Works like primary/secondary DNS
  - If clients fail to read from primary, they try secondary




# MIGRATE BETWEEN-SHOT COMPUTING TO THE WAN

- **Deploying a supercomputer to support pseudo real-time analysis**
  - Network QoS
  - CPU scheduling
  - Faster CPUs and algorithms
  - Data management
  - End-to-end performance
- **Substantially enhanced data analysis**
  - Historically this had made a huge scientific impact
- **Can have a safety impact for future devices**
  - e.g. ITER: <10% of high power discharges can disrupt



# FIRST PROOF-OF-CONCEPT BETWEEN-SHOT GRID ANALYSIS REVEALED PERFORMANCE ISSUES

- A between-shot TRANSP proof-of-concept was done in March
- TRANSP (typically 12-36 hour code runs) adapted for shot cycle of 15 to 20 minutes
- Revealed that data transfer took more than 45 minutes!
  - Negligible for 24 hour runs, but not for between-shot analysis
- **First round of optimization: 12x**
  - 48 minutes -> 4 minutes
- **Second round of optimization: 12x**
  - 4 minutes -> 20 seconds



# DATA ANALYSIS THROUGHPUT IS THE MOST IMPORTANT PERFORMANCE CONSIDERATION

- **Anything that can increase data analysis throughput is useful**
  - More volume, less time
  - Anything else that helps scientists work faster
    - e.g. easy-to-use, well documented codes



- **While it would be nice to have a supercomputing center at each tokamak site, that is too expensive**
- **We seek to increase data analysis throughput through resource sharing and more efficient use of resources**

# RELIABILITY IS ANOTHER IMPORTANT CONSIDERATION

- **To be useful during an experiment, computational infrastructure must work**
- **A lost shot is a lost opportunity**
  - Example: the next-generation ITER device will be ~ \$1 M per shot
  - Miss a shot, miss a million-dollar opportunity
  - DIII-D is an order of magnitude less expensive, but still not cheap
- **Must be reliable, dependable**
- **Would you take the bus to work if it showed up only 95% of the time?**



# THE PERFECT PERFORMANCE TOOL?



**The perfect performance monitoring & management tool would...**



**Monitor Network QoS, CPU, mass storage, status of analyses, physics data consistency**

- Automatic detection



**Easy enough for scientists to use**

- Tells them if something is wrong and who to contact for help




**Easy enough for developers to use**

- Open source, easily extended & have a simple API




# WE SEEK PARTNERS FOR DEVELOPING AND IMPLEMENTING THESE IDEAS

- Develop application requirements
- Determine relevant computer science components
- Examples of good application and computer science partnerships are the National Fusion Collaboratory and Network QoS Projects
- We would be eager to participate in exploring potential ideas



# CONCLUSION

- **Fusion science is advanced through experiments**
- **Requires rapid between-shot analysis**
- **To increase data analysis throughput, moving from local-only data analysis to the WAN**
  - Performance issues
- **Important performance metrics are data analysis throughput and reliability**
  - Increase volume of analysis between shots
  - Help scientists to work faster
  - Must be a reliable infrastructure
- **We look forward to teaming together**



*Inside a tokamak during a shot*

# QUESTIONS?

