HDF5 at ITER

Lana Abadie [1], Rodrigo Castro [2], Yury Makushok [3], David Muir [4], Jonathan Hollocombe [4], Simon Pinches [1], Mikyung Park [1], Anders Wallanders [1]

[1] ITER Organization
[2] CIEMAT
[3] INDRA
[4] Culham

Disclaimer: The views and opinions expressed herein do not necessarily reflect those of the ITER Organization

ITER

- 7 ITER members (China, E.U, India, Japan, Russia, South Korea, U.S.A) make cash in kind-contributions to the ITER project
- Demonstrate feasibility of fusion energy for peaceful purposes
- A plasma of deuterium and tritium (hydrogen isotopes) is heated to more than 150 millions ◦C
- The hot plasma is shaped and confined by strong magnetic fields
- Helium nuclei sustain burning plasma
- Neutron transfers energy to blanket
- Conventional steam generator, turbine and motor will transform the heat into electricity

1gr of fusion fuels=8t of oil

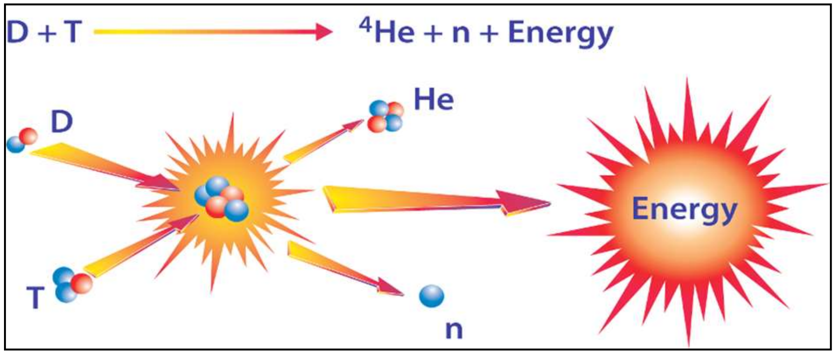

ITER TOKAMAK

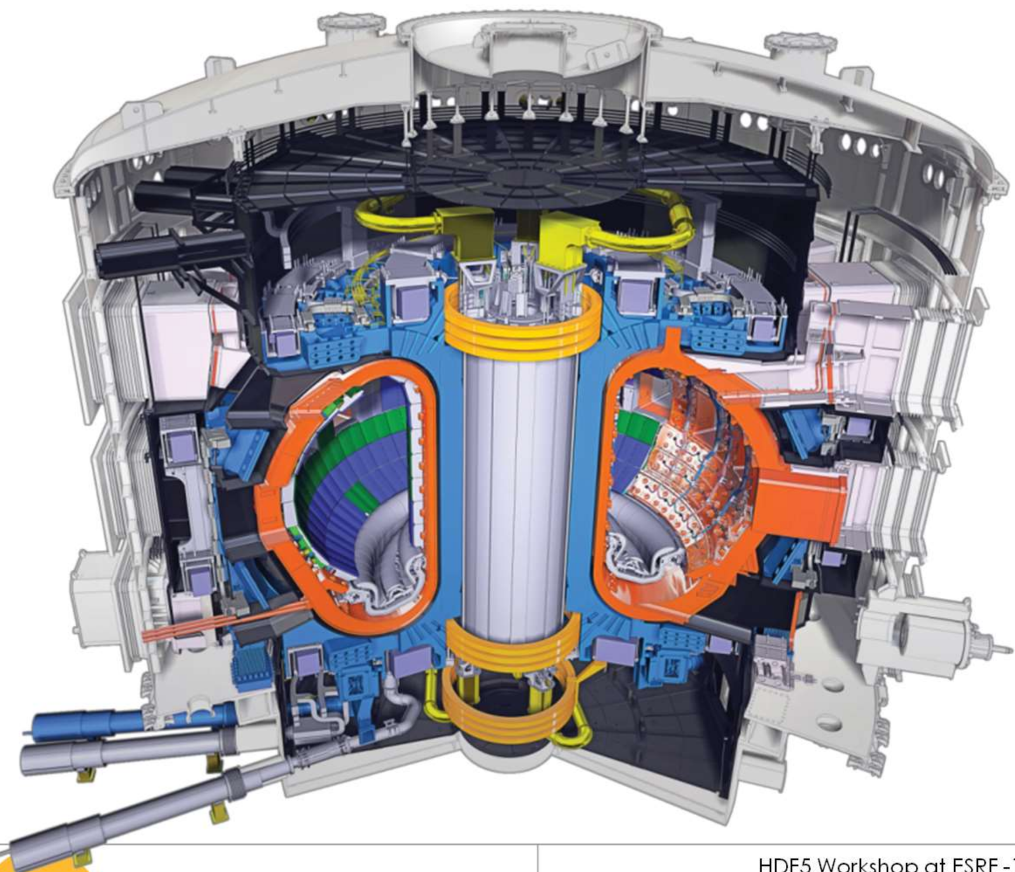

- Vacuum Vessel: ~8 000 t.
- TF Coils: 18 x ~360 t.
- Central solenoid: ~1 000 t.
- Total: ~23 000 t.
- ~3.5 the Eiffel tower…
- R=6.2m,a=2.0m, Ip=15MA, B_T=5.3 T

- Achieve fusion power of 500 MW with P_{fus}/P_{in} (Q) ≥ 10 for 300-500 s

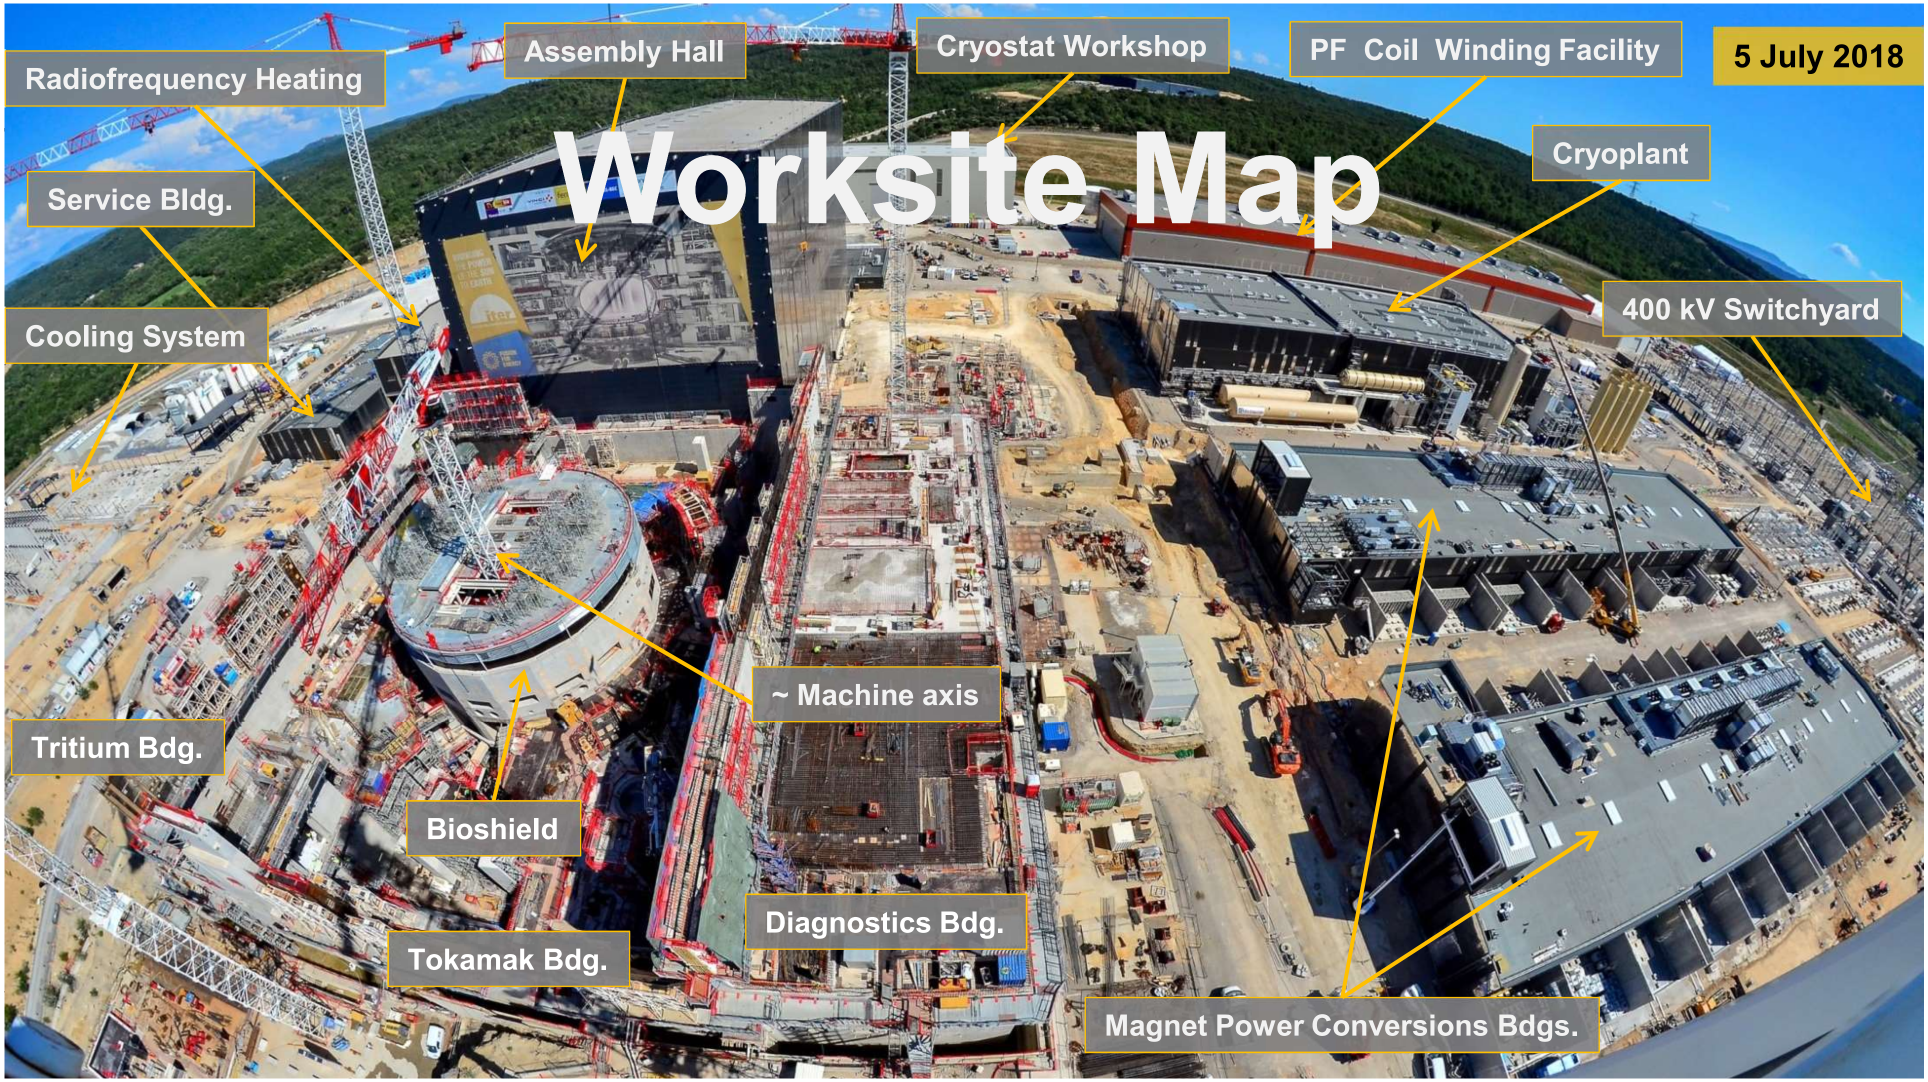
Worksite Map
5 July 2018
Radiofrequency Heating
Assembly Hall
Cryostat Workshop
PF Coil Winding Facility
Cryoplant
Service Bldg.
Cooling System
400 kV Switchyard
~ Machine axis
Tritium Bdg.
Bioshield
Diagnostics Bdg.
Tokamak Bdg.
Magnet Power Conversions Bdgs.

Manufacturing Progress



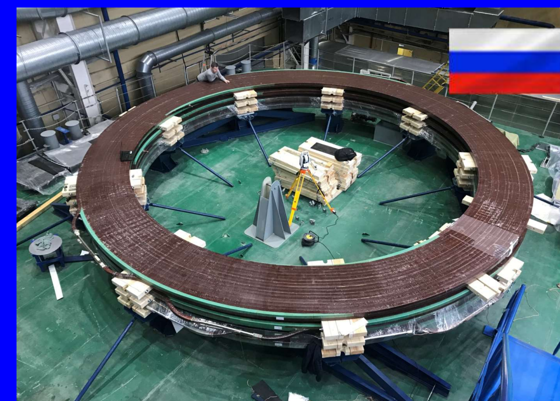

Vacuum Vessel, Magnets, Cryostat, Thermal Shields, Cryopumps, etc.


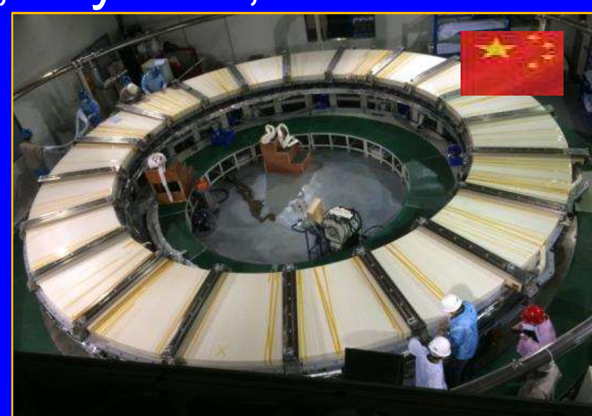


A “Staged Approach” to Full Operating Capacity

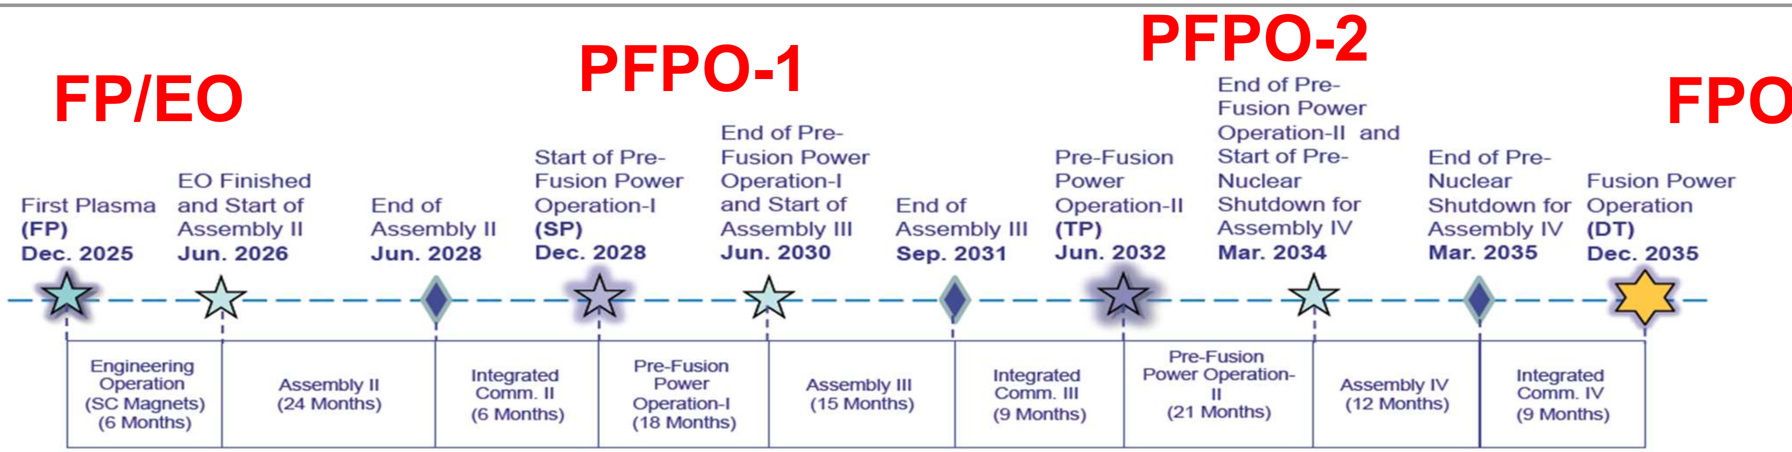

- Extensive interactions among the ITER Organization and Domestic Agencies to finalize revised baseline schedule (2015-6)
  - Schedule estimates through First Plasma (2025) up to DT operation (2035) consistent with Members’ budget and technical constraints

I&C architecture

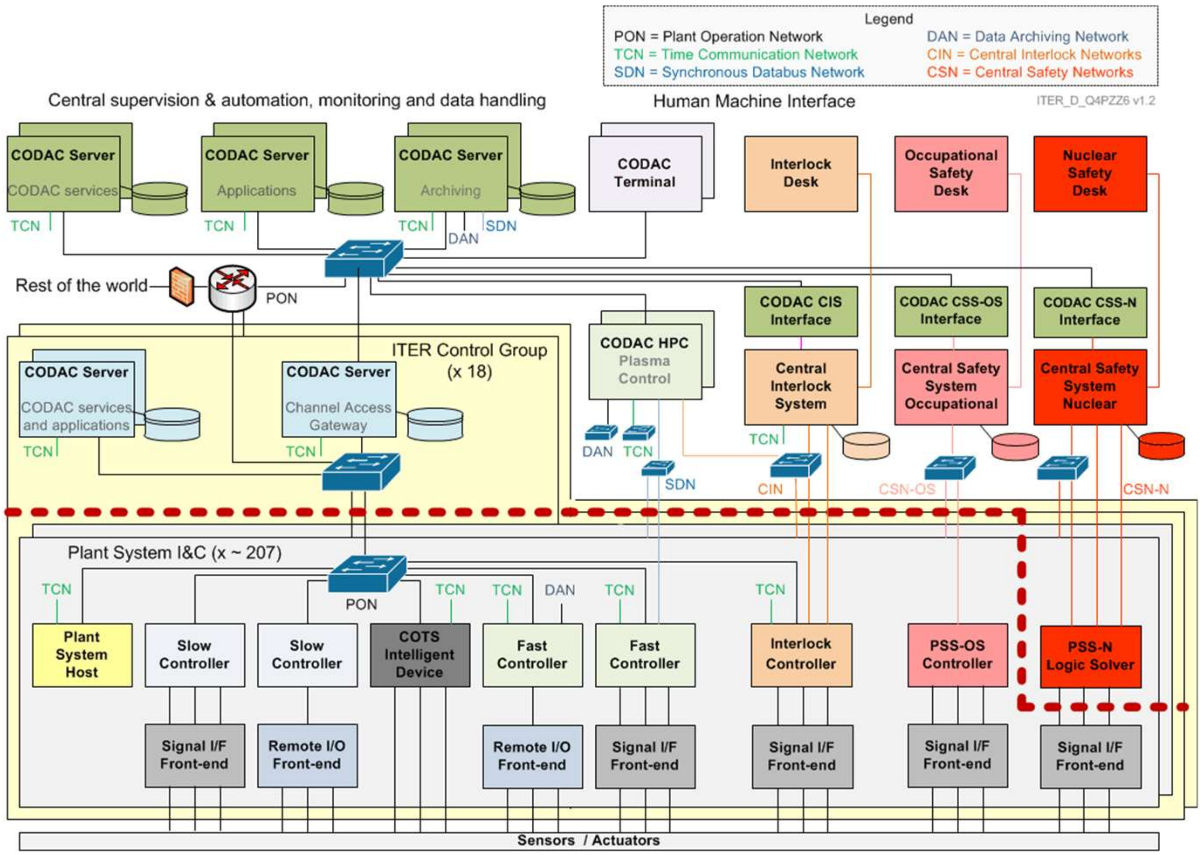

Main control parameters

Parameter	Value
3.1 Total number of computers (PON)	1.000
3.1.1 Total number of PSH and controllers	700
3.1.2 Total number of CODAC servers	80
3.1.3 Total number of CODAC terminals	150
3.2 Total number of channel access servers (IOC)	2.000
3.3 Total number of channel access clients	500
3.4 Total number of signals (wires)	100.000
3.5 Total number of process variables	1.000.000
3.6 Update rate on HMI screen (PON)	200 PV at 5 Hz
3.7 Total engineering archive rate (PON)	25 MB/s
3.8 Maximum sustained data flow on PON	50 MB/s
3.9 Total engineering archive capacity	2.2 TB/day
3.10 Maximum number of alarms per ctrl group	600
3.11 Total scientific data rate (DAN)	2-50 GB/sec
3.12 Maximum sustained data flow on DAN	2-50 GB/sec
3.13 Total scientific archive capacity	90-2200 TB/day
3.14 Accuracy of time synchronisation (TCN)	50 ns RMS
3.15 Number of nodes connected to TCN	500
3.16 Maximum latency asynchronous event (SDN)	1 ms
3.17 Number of nodes connected to SDN	100
3.18 Maximum latency sensor to actuator on SDN	500 µs
3.19 Maximum jitter sensor to actuator on SDN	50 µs RMS
3.20 Maximum sustained data flow on SDN	25 MB/s
3.21 Maximum cable length between two nodes	2 km

Different Sources for Raw Data

- PON-Data
  - ✓ EPICS traffic – conventional control & monitoring
  - ✓ 25-50 MB/sec
  - ✓ Two back-end RDB and HDF5 files
- SDN-Data
  - ✓ Real-time data for feedback control
  - ✓ Multicast topics
  - ✓ Max. 10Khz
  - ✓ In-house implementation
  - ✓ Stored in HDF5 files
- DAN-Data
  - ✓ Archiving of data with High throughput
  - ✓ Up to 50GB/sec for cameras
  - ✓ In-house implementation
  - ✓ Stored in HDF5 files
- Data access shall be transparent and independent of data source and data type
  - ✓ UDA (Unified Data Access)

	FP	PFPO-1	PFPO-2	FPO-1	FPO-2, …, -8	Total
Data (PB)	0.1-1	10-100	100-300	200-500	+500 each	5000

HDF5

- Use of HDF5 as it is a powerful self-described data format
- Data preservation : need a format for the lifetime of ITER project (Similar to NASA)
- Use of SWMR, 1.10.3
- Use low-level C API + h5py (for test script + some plotting tools)
- Env. : linux R.H 6.5, 7.4, 7.6 + CentOS
- Different HDF5 layouts per Data sources

PON layout

- Data is essentially time-series, always on
- PON layout -
  - Many datasets in one file (one dataset per variable). Limit to 100K dataset in one file.
  - Use of compound data types to store values, timestamp and alarm information
  - Implementation of file rotation
  - Main issues : cannot create the dataset on the fly, requires to disable the memory cache of the dataset
  - Usually low rate (max. 10Hz) so writing performance is OK
  - Use of attributes for static metadata such as data type, units

SDN layout

- Data is essentially time-series, on during plasma experiments
- SDN layout –
  - A few datasets in one file (one dataset per SDN topic).
  - Use of compound data types : structure different from one topic to another topic. Size of structure can be big
  - Implementation of file rotation
  - Usually higher sampling rate(max. 10KHz)
  - Use of attributes for static metadata such as units, descriptions

DAN layout

- 2 datasets
  - Block header (fixed part + user-defined part)
  - Payload (buffer is multiplexed)
- Support for atomic types + compound types
- Same layout used for images
- Optimized for writing
- Use of attributes for metadata (units, calibration factors, user-defined attributes)
- Performance limitation is mainly due to the disk speed
  - Test 1 stream at 800MB/sec for 1 hour without loss
- Error handling was not easy, especially when the storage is full, we could see from times to times some zeros being written at the end of the files
- Inform the clients that the file is not being written? No flags?

DATA ACCESS - UDA

- Access the different source of raw data transparently
- Pre-compute min/max/avg to ease interactive zooming and speed data access
- Based on a modular solution (plugin-based), C and C++ based
- Basically two types of data access
  - Given a time range
  - Given a pulse number
- Many data types to support
  - Scalar value
  - Structure
  - Profiles (2-D)
  - Images/video
- Access the metadata (at given time)
- Performance issues found with hyperslab method reported to HDF group

DATA ACCESS - UDA

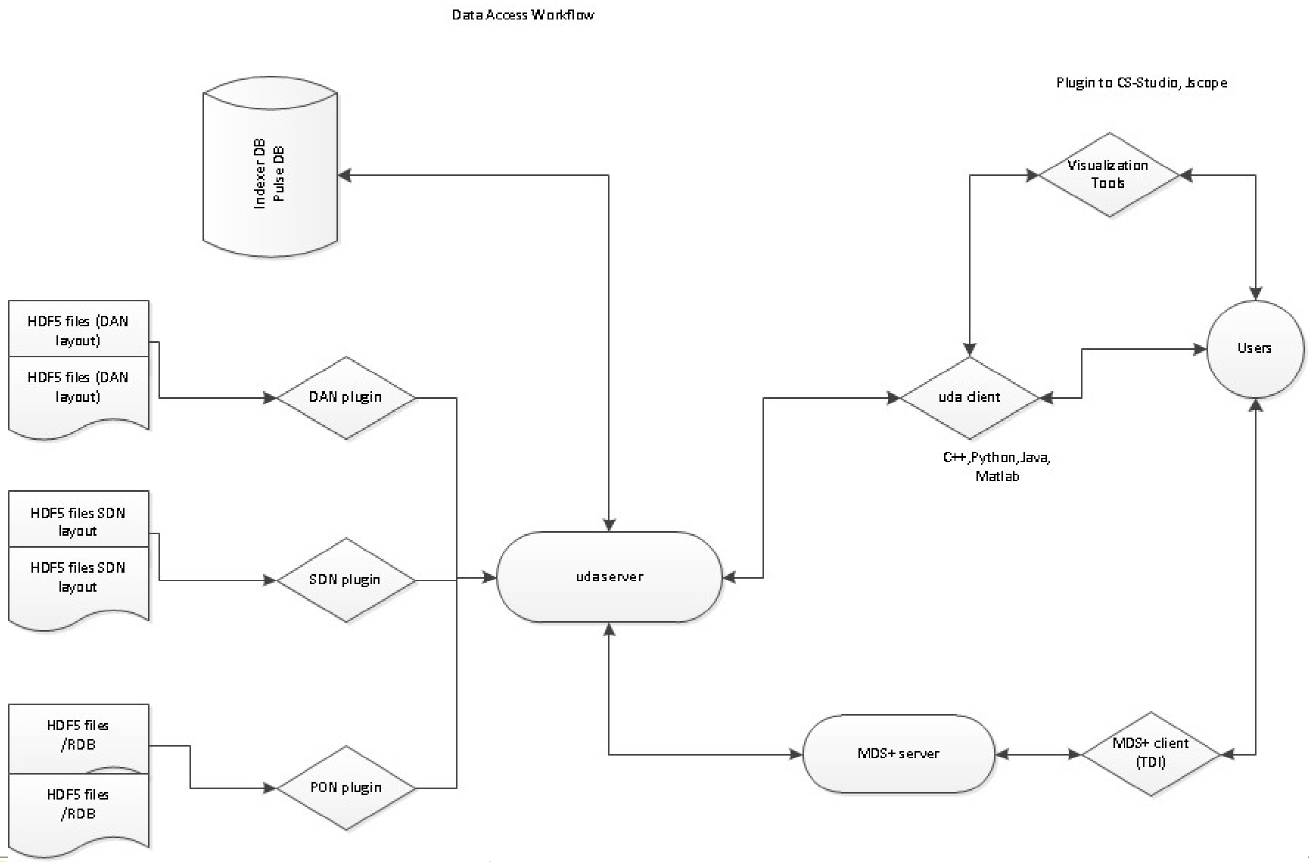

Matlab Windows UDA interface

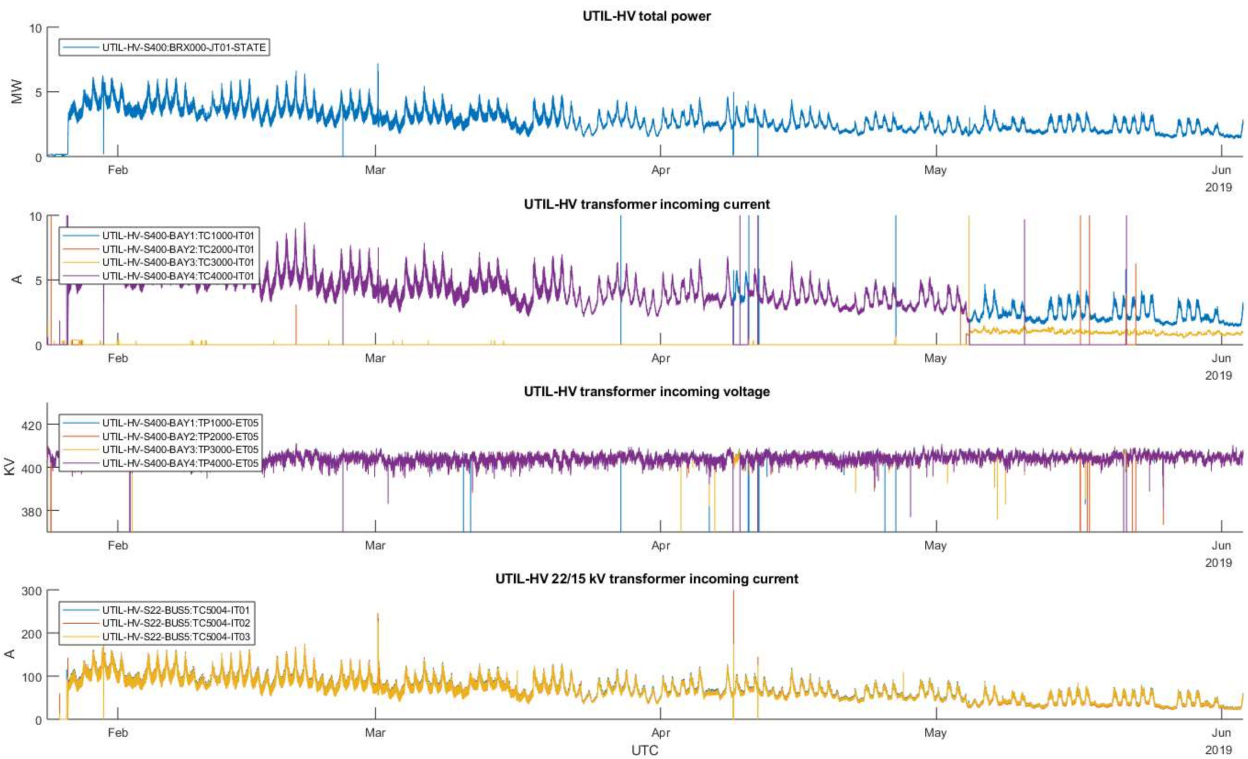

Real data from Electrical system
~30K variables
~more than 1 year data
~100GB of data

Python Interface

courtesy from MAST (diagnostics data)

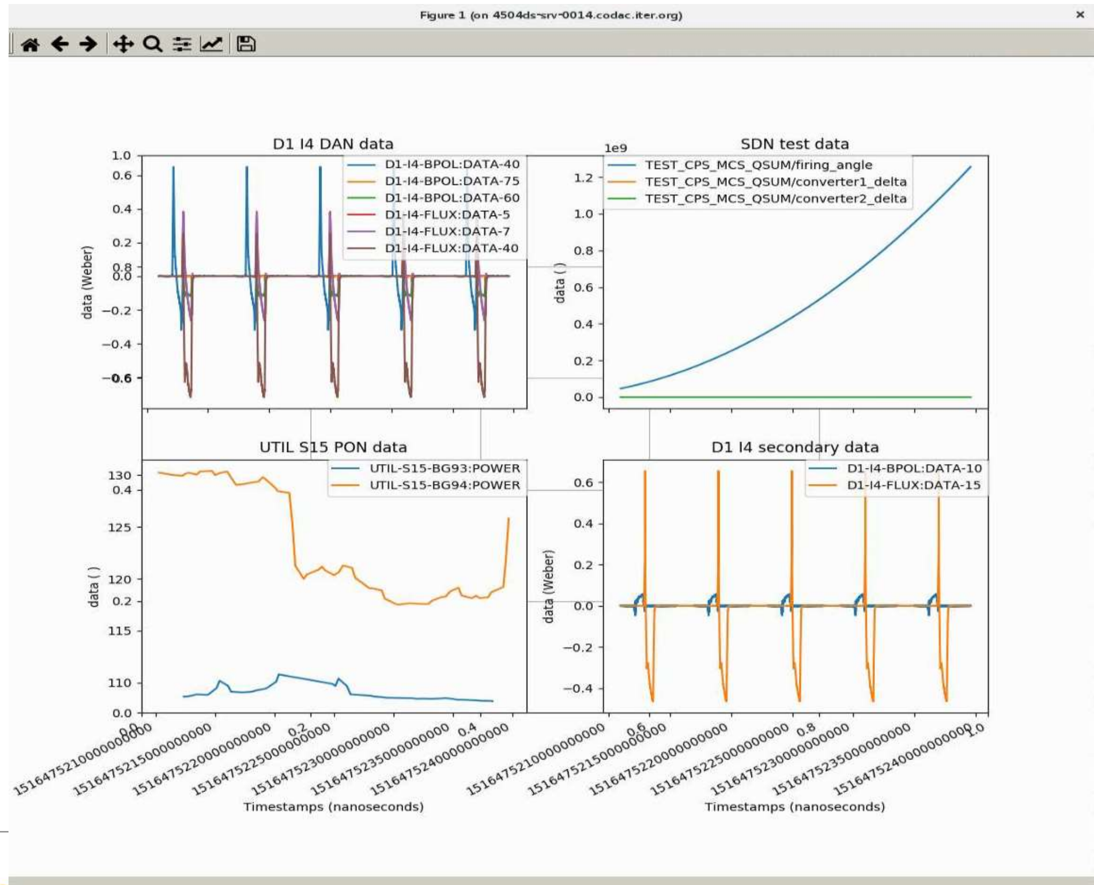

C++/Qt based tool

C++/Qt based tool

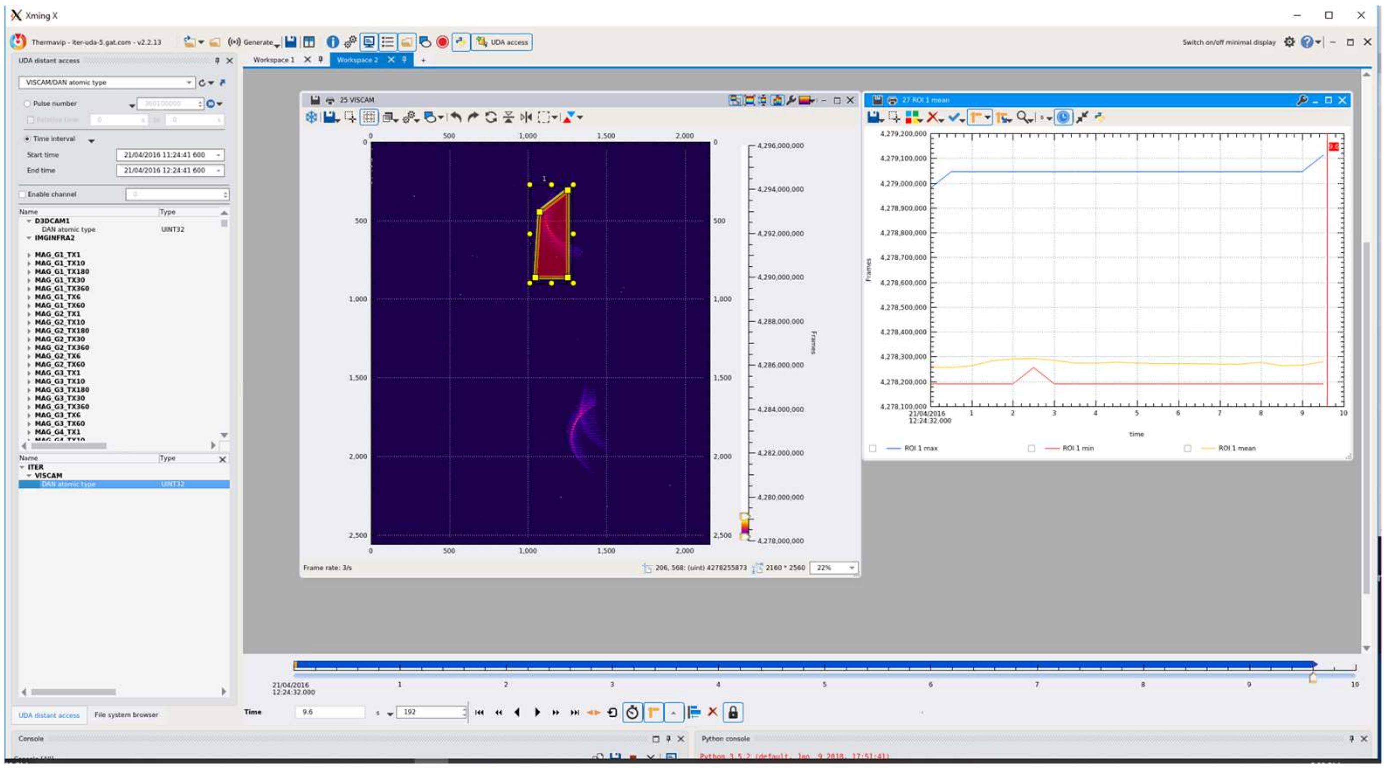

Summary

- Good progress, waiting for diagnostics to make more tests, but still a lot to do…
- Scalability tests for UDA on-going
- Overall satisfaction with SWMR
- Good support from HDF group
- Next steps:
  - Evaluation of HSDS (object store)
  - Better integration with physics workflow
- How the scientific community sees use of cloud and object storage?