DATA MANAGEMENT AT JET WITH A LOOK FORWARD TO ITER

J.W. Farthing*, T.Budd, A.J. Capel, N. Cook, A.M. Edwards, R. Felton, F.S. Griph, E.M. Jones, R.A. Layne, P.A. McCullen, D. McDonald, M.R. Wheatley, Euratom/UKAEA Fusion Association, Culham Science Centre, Abingdon, UK and JET-EFDA contributors#
M. Greenwald, Plasma Science and Fusion Center, MIT, Massachussets, USA
J.B. Lister, Association Euratom/Suisse, CRPP-EPFL, Lausanne, Switzerland
J. Vega, Asociación EURATOM/CIEMAT para Fusión, Spain

Abstract

The various databases and data management techniques used to configure the JET control systems are described by following the JET experiment pulse cycle of preparation, parameter setup and validation, countdown, real-time control, data collection, storage, analysis and access. Elements of the ITER-CODAC conceptual design are used to discuss where systems may be different on ITER.

INTRODUCTION

JET [1] is the largest Tokamak to date and has been operational since June 1983. It is located at Culham, Oxfordshire, UK and is operated by the United Kingdom Atomic Energy Authority (UKAEA) on behalf of the European Fusion Development Agreement (EFDA) [2]. JET has an ongoing role in performing ITER-relevant engineering and physics before ITER [3] goes operational in 2016.

JET CONTROL SYSTEMS ARCHITECTURE

The JET CODAS (COntrol and Data Acquisition Systems) architecture is hierarchical and modular in hardware and software at three levels [4]. Level-1, the central / supervisory level, provides a standard human-machine interface, data storage and analysis, central timing, central interlocks and safety systems, networking, experiment management and management of real-time control. Level-2, the subsystem level, provides conventional control, data acquisition, timing and interlocks. Each of the major plant systems, e.g. toroidal field (TF), poloidal field (PF), vacuum (VC), ion cyclotron radio frequency (RF), lower hybrid current drive (LH) etc correspond to a subsystem. There are 21 subsystems in total, including 10 that load-level ~80 diagnostics. The subsystems are designed to be autonomous in order to allow commissioning independent from the rest of JET. Level-3, the component level, consists of individual components of subsystems, e.g. vacuum pumps, power supplies, lasers etc.

The same architecture is applicable to the machine control and diagnostic subsystems. Machine control subsystems acquire comparable amounts of data to the diagnostics and the diagnostics are increasingly used in real-time control and protection - the plasma is the subject of the experiment.

Each subsystem is allocated a SUN Microsystems dual-processor Ultra-60 computer running Solaris-10. The computers run data-less, all subsystem data resides on a RAID server, in order to minimise swap times in the case of failure.

Most of the subsystems have attached CAMAC serial highway loops. These are interfaced via PCI-VME cards and VME-SHD cards and with a CAMAC device driver loaded into the Solaris kernel. Level-2/3 systems in VME running VxWorks and on PCs running Windows or Linux are gradually replacing the CAMAC. Three ATCA-based systems will be deployed in early-2008.

Many of the newer Ethernet-based systems, particularly those implemented by the EFDA Associate Laboratories, are interfaced to the Solaris subsystems using an HTTP-based protocol where the client runs on the Solaris host, and the server in the local controller [5]. The protocol defines the methods for parameter setup, system monitoring and bulk data read-out. These systems are referred to colloquially as “black boxes” and are similar to the ITER model of plant systems “delivered in-kind” together with their control and data acquisition systems. Typically all that is required is a network connection, a trigger, a synchronisation clock and the matching database definitions on each side of the interface.

SUBSYSTEM DATABASES

On each of the subsystems there are two hierarchical databases: the hardware and experiment trees. These are in-house developments dating from the early 1980s.

The hardware tree describes the I/O devices, the signal names, their calibration factors, the access control rules etc. A memory-resident live-database, known as the plant status image (PSI), is built from the tree and is used by all the conventional control room tools – synoptic displays, alarm handling, trend curves, set-points, level-2 control logic, state-machine definitions, component frameworks etc. All of the standard tools are data-driven, the more recently-developed ones using a CODAS configuration language based on RDF.

The experiment tree describes the pre- and post-pulse actions to be applied to the hardware. The link between the experiment tree and PSI is via the signal name.

The database files and the configuration files for the data-driven tools reside in a ZFS file system at the subsystem level named /jet/XX where XX=TF, PF etc.

*jonathan.farthing@jet.uk
#See Annex 1 of M.L.Watkins, “*Overview of JET Results*”, OV-1.3, Proceedings 21st IAEA Fusion Energy Conference, Chengdu, China, 16-22 October 2006.

LEVEL-1 SOFTWARE

The JET Level-1 software [6] is layered above the conventional level-2/3 control and data acquisition and provides cross-subsystem pulse preparation, setup and validation. It provides a single view of ~22,000 parameters. The same data-driven tool, known as the pulse schedule editor, with different views of the underlying data, is used by the Engineer-in-Charge (EiC), the Session Leader (SL), the Diagnostic Coordinator (DC), the power supplies operation engineer (PSOE), the pilots of the different heating and fuelling systems and many of the individual diagnostic responsible officers (ROs). Level-1 implements a separate memory-resident live-database that combines values from the PSI on the subsystems with values from the experiment trees to provide one single view of the main JET parameters. There are ~400 plug-in codes used to convert plasma requirements in terms of physics units into machine setup parameters and to validate the combinations of parameter values for consistency. Interfaces to simulation codes are also available. Included in the validation codes is the implementation of the pre- and post-pulse JET operating instructions (JOI). These are formal constraints on the JET operating space, defined by the Chief Engineer. Groups of parameter values, known as scenarios, provide recipes for particular plasma phases.

Increasingly the extra functionality of the level-1 tools is being used also at the subsystem level, thus replacing the use of some of the level-2 tools in normal JET operations.

THE JET EXPERIMENT PULSE CYCLE

JET is a pulsed machine. A single experiment is referred to as a pulse during which the plasma lasts for ~30-40s. The full pulse cycle from pulse preparation through to data analysis and hence preparation for the next pulse takes on average ~30 minutes, although technically can be completed in ~12 minutes. JET is operated in two shifts, known as sessions, from 06:30 to 14:45 and from 14:15 to 22:30. On average 25 pulses per day are executed. One of the challenges is to maximise the pulse repetition rate.

On ITER the plasma will last for up to 60 minutes.

Pulse preparation

Pulse preparation is performed using the pulse schedule editor. In advance of the session the SL prepares the planned pulse settings in a pulse schedule making use of scenarios, groups of parameters from other schedules, the settings stored for previous pulses and the normal editing facilities. The same validation codes used by the EiC at execution time can be used by the SL against settings in the pulse schedule. Electronic handover passes the required settings to the EiC, the PSOE and the pilots of the heating and fuelling systems.

Set-up

Parameter setup for the main plant parameters is performed by the EiC, for the detailed power supplies parameters by the PSOE and for the heating and fuelling systems by the various pilots. The diagnostics are setup by a combination of the DC and the individual ROs.

Validation

Once all of the parameters have been setup and checked the various ROs mark themselves as “ready”. The top-level cross-subsystem checks including the pre-pulse JOIs are then carried out by the EiC.

Countdown

The automatic part of the pulse cycle is controlled using a hierarchical state machine. All included components of a subsystem and all subsystems included in the experiment must be in a particular state before JET can move to that state. A NASA-like countdown from minus 2.5 minutes is initiated by the EiC once the JOIs have been checked. During the countdown the values from the experiment database are downloaded to the plant and the various components are requested to go “ready-for-pulse”. A hold-point at 30s is used for a final check that all subsystems are ready.

Pulse - real-time control

The pulse itself is triggered and synchronised by the central timing system and controlled by the various real-time feed forward and feed back systems. The JET real-time measurement and control systems [6] communicate over an ATM network with just a small number of residual analogue signal connections. Measurements from plant diagnostics are converted to physics values and made available to feed back controllers for the plasma control (shape, position and vertical stabilisation), the heating systems and the density. These controllers apply the required feedback by controlling the appropriate plant actuators. Ad-hoc feedback algorithms may be applied using the real-time central controller (RTCC) that has access to all the measurements and feed back controllers. Programming RTCC is another example of a data-driven application using the level-1 pulse schedule editor.

Real-time control is an ongoing growth area on JET. More physics diagnostics are being made available in real-time and more sophisticated control is being applied. Real-time protection, for example of the plasma-facing materials, is also becoming an important subject. Compile- and run-time checking of the ATM packet structures are being implemented using a configuration database.

Real-time control and protection will be fundamental to ITER operations.

Data Collection

Once the plasma is over the task is to collect the data from the plant systems as quickly as possible to make it available to the data analysis programs. The data is read from the CAMAC, VME and the PC systems and stored

as pulse files at the subsystem level. The pulse files are self-describing since they also include a copy of the experiment tree and the calibration factors from the PSI. Only the raw data is stored and is converted to calibrated engineering and physics units at data access time. Note also that it is not uncommon for diagnostics to be recalibrated following experimental campaigns. From the subsystems the files are transferred to the mass data store. The data collection time is dominated by the legacy CAMAC data. A multi-tier collection strategy controlled by flags in the experiment tree allows the most important data for the intershot analysis to be collected and transferred first.

Large data files from PCs and VME systems are now collected directly to the mass data store thus avoiding transfer to the subsystems as an intermediate step.

A separate system, not described here, is used for continuously recorded slow (<=0.25Hz) data.

Data Storage

The mass data store consists of a SUN Microsystems E4900 enterprise server with 8 UltraSparc IV dual-core processors with 64GB of memory running Solaris 10 and ZFS. The pulse files are initially stored on mirrored disk to allow the data analysis to start. In background three StorageTek tape copies are made, one remains in the tape silo and two are stored at geographically separated locations. The files are then compressed and copied from the mirrored disks to raided cache disks and then deleted from the mirrors. The aim is to keep all data online on the cache disks – budget permitting. If this is not possible a cost-based deletion strategy, function of file size and last access time, is used to free space. At present 49TB of usable cache disk is available and is 52% full. Files are automatically retrieved from tape when required.

Data Analysis

Six AMD 3400+ Linux nodes from the JET Analysis Cluster (JAC) running Fedora Core-6 are dedicated to intershot analysis. This is essentially a scheduling problem. A database models the dependencies of ~80 codes both on the raw data pulse files and on each other (processed data) and the codes are scheduled in parallel based on these dependencies. The database is also used for critical path analysis and the multi-tier data collection can be used to fine-tune the arrival of the raw data. Communication between the codes is via the writing of the processed data files although some work is ongoing to investigate the use of web services for this purpose. At present the most important parts of the intershot analysis are complete before the data collection is complete and hence do not impact on the pulse cycle time.

More detailed analysis typically takes place over the following weeks or even months and years and inspection of data access patterns shows that old (>5 years) data are read on a regular basis.

JET-like intershot analysis will not be required on ITER as the analysed data must be available in real-time. However by the time that ITER starts operating it may be possible to perform the more detailed analysis between pulses.

Data access

Client-server technology is used for both raw and processed data access with the servers running on the E4900. Access is via subsystem, signal name and pulse number. The raw data pulse file format is based on the structure of the experiment tree and dates from ~1981. The file headers hold the signal meta data. The meta data for the processed data is held in a MIMER RDBMS and the data in NetCDF files. A home-grown file system-based indexing system using ZFS is used to manage both the raw and processed data files.

Subsets of the raw and processed data are stored in a SAS database for reasons of statistical analysis. Access is client-server using SAS or SQL. We are in the process of migrating the underlying data storage from SAS to PostgreSQL.

MDSplus – Model Data System

MDSplus is the de-facto standard way of accessing fusion experimental data. It was first developed by MIT, CNR-RFX Padova and LANL in ~1987 [8]. It ".. allows all data from an experiment or simulation code to be stored into a single, self-descriptive, hierarchical structure" [9]. JET provides an MDSplus glue layer to the pre-existing raw and processed data servers and in this way also provides remote data access (RDA) using the MDSplus API. RDA is used as a complementary technique to that of remote computer access (RCA) depending on user preference, the software tools available at the different sites, the amount of data to be transferred and remote IT security rules. Note that JET policy states that all processed data must be stored centrally at JET so that it can be used by all collaborators.

MDSplus provides the basic functional requirements of the ITER data access system from the users' perspective. Extensions have been proposed to cover continuous data acquisition as part of the ITER conceptual design.

Data Volumes

Since the first JET pulse in 1983 the raw data collected per pulse has roughly followed a Moore's Law-like doubling every 2 years. Today we collect up to ~10GB per pulse and the total data collected over ~70,000 pulses amounts to ~35TB. Enhancements to JET in 2007, 2008 and 2009 should result in ~60GB per pulse being collected by 2010. This is small by comparison to HEP experiments but still represents a challenge to maintain the pulse repetition rate, data access times and data security.

JET SUPPORT DATABASES

In addition to the control systems databases JET supports a large number of heterogeneous databases covering all aspects of the project including: financial systems, planning, maintenance schedules, long-term

experiment planning, machine configuration, rostering, publications, OS account management, management systems and QA, electronics database, document management etc. These are not all as well integrated as one would wish for with the inherent risks of duplicated, out-of-date and inconsistent data. An ongoing task is to merge and integrate these databases where possible.

WHERE WILL ITER BE DIFFERENT?

To compare ITER with JET some elements of the ITER-CODAC (COntrol, Data Access and Communication) conceptual design [10] can be examined.

ITER will be long-pulse / pseudo-continuous

The pulsed and continuous data acquisition systems on ITER will be merged. All data sources, including internal CODAC metrics, can be considered as streaming data at different rates. A single data access layer will provide access to all stored data. A single signal naming scheme as part of a wider plant naming scheme, based on URIs (RFC 3986), managed name spaces and Qnames has been proposed. Data will be accessed via absolute time although pulse numbers and relative times to pulses and segments within pulses will still be available as shortcuts. However data retrieval methods should tend towards physical criteria rather than the old pulse number and time interval methods.

The pseudo-continuous data implies the need for novel data mining and data classification techniques. One such technique currently being developed at JET involves encoding signals according to a discrete set of alphabetical values according to their gradients. The encoded values are stored in a PostgreSQL RDBMS and the full power of SQL can be used to search for similarly shaped data [11].

All project data can have an associated time

There is a requirement to replay history for operator training and for investigations into incidents. This applies to software versions and configuration data as well as the plant data. There will be a concept of future as well as past for experiment planning purposes. Full data provenance should be provided from the published papers back to the ADC boards, firmware, software, cabling changes etc.

Merge slow controls with plasma controls

Just as the slow and pulsed data acquisition and access systems will be merged so will the slow, e.g. vessel conditioning, and plasma controls. A single schedule editor and real-time scheduler should be able to be used for both sets of requirements. It is just a question of scheduling tasks / actions with different time scales. This can be extended to long- and medium-term experiment planning, maintenance scheduling, shutdown planning, automatic, or semi-automatic, commissioning procedures etc.

Real-time control, analysis and protection

Real-time control, analysis and protection will be the norm on ITER as is already the general trend on fusion devices today.

Data Volumes and Data Rates

Rough estimates based on the ITER project requirements and on extrapolations from JET suggest that ITER will collect <10PB/year in 2016, small by comparison to LHC, and that the maximum data rate will be <10GB/s. Note that there is always a tendency to underestimate data volumes, for example the original JET design documents from 1977 estimated the raw data collected on JET to the end of 1990 as 20GB. This turned out to be an underestimate by a factor of 8.5.

Simulation data volumes will be significant

Simulations of ITER plasmas will be performed in advance of each pulse, and maybe also in near real-time applying correction factors as the pulse proceeds. 1 minute of burning plasma is predicted to produce 1TB of simulation data.

Archive all raw data or only novel data

The cost of archiving all the raw data is small by comparison to the ITER investment and so should always be archived [12].

Internet Age

Unlike when JET was designed there are now a wide variety of commercial and open source solutions available addressing all aspects of data management. Very little should be developed in-house. The trick will be to select the most enduring technologies among those available. Success will rely both on good judgement and a degree of luck.

SINGLE PROJECT-WIDE DATABASE

ITER should aim for a single logical database to last the project lifetime. It should cover all aspects of the project including: document management, CAD, structured data, plant naming schemes, signal naming schemes, software, planning and IT. It should cover the project lifetime: construction, operations, maintenance, decommissioning. It should cover all areas of the project: administration, engineering, physics, IT. It should merge the long-term experiment planning with operations data, data storage, data analysis and publications. The control systems data should be derived from this single logical database. All data should have an associated time, history should be kept and hence full provenance of published data should be available. Management systems: processes, quality, risk management, obsolescence management should also be included.

The ITER-CODAC conceptual design [10] proposes that the plant systems “delivered in-kind” should be self-describing via schema provided by CODAC. The data to be provided includes the cubicle layouts, the modules and

their history, the wiring, signals, firmware, software, commands to the local controllers etc. It is essential that these data be incorporated into the single, logical site-wide database at system integration time.

SUMMARY

The JET hierarchical and modular architecture and the hierarchical databases and hence the raw pulse data file formats developed in ~1981 have stood the test of time. The processed data file formats have changed, but the underlying structures are hidden from the end-users by using a well-defined API and client-server technology. The level-1 software has been layered on top of the level-2 (subsystem) control to provide JET-wide control and increased functionality.

In addition to the control system databases JET has a large number of heterogeneous support databases developed at different times, by different groups and using different technologies. Integration between them is not always ideal.

ITER is starting with a clean sheet. The aim should be to provide an all-embracing, project-wide database to last the project lifetime. The control systems data should be derived from this database.

REFERENCES

[1] http://www.jet.efda.org
[2] http://www.efda.org
[3] http://www.iter.org
[4] H. van der Beken, et al., CODAS: The JET Control and Data Acquisition Systems", Fusion Technology 11(1), pp. 120-137
[5] C.Hogben and F.S.Griph, "Interfacing to JET Plant Equipment using the HTTP protocol", http://www.iop.org/Jet/fulltext/EFDR02004.PDF
[6] H. van der Beken, B.J. Green, C.A. Steed, J.W. Farthing, P.A. McCullen and J.A. How, "Level-1 software at JET – a global tool for physics operation", Fusion Engineering, 1989. Proceedings, IEEE Thirteenth Symposium on Fusion Engineering.
[7] R. Felton, T.Budd and F. Sartori, "Real Time Measurement and Control at JET – Status 2007", ICALEPCS 2007
[8] J.A. Stillerman, T.W. Fredian, K.A. Klare and G. Manduchi, Rev. Sci. Instrum. 68, 939 (1997)
[9] http://www.mdsplus.org
[10] J.B. Lister, J.W. Farthing, I. Yonekawa and M. Greenwald, "Status of the ITER CODAC Conceptual Design", ICALEPCS 2007
[11] J. Vega et al., "Data mining techniques fpr fast retrieval of similar waveforms in Fusion Massive databases", 6th IAEA TM on Control, Data Acquisition and Remote Participation in Fusion Research, Inuyama, 2007, to be published in Fusion Engineering and Design
[12] J.B. Lister, B.P. Duval, J. W. Farthing, T.J. Fredian, M. Greenwald, J. How, X. Llobet, F. Saint-Laurent,W. Spears, J.A. Stillerman, "The ITER Project and its Data Handling Requirements", ICALEPCS 2003