



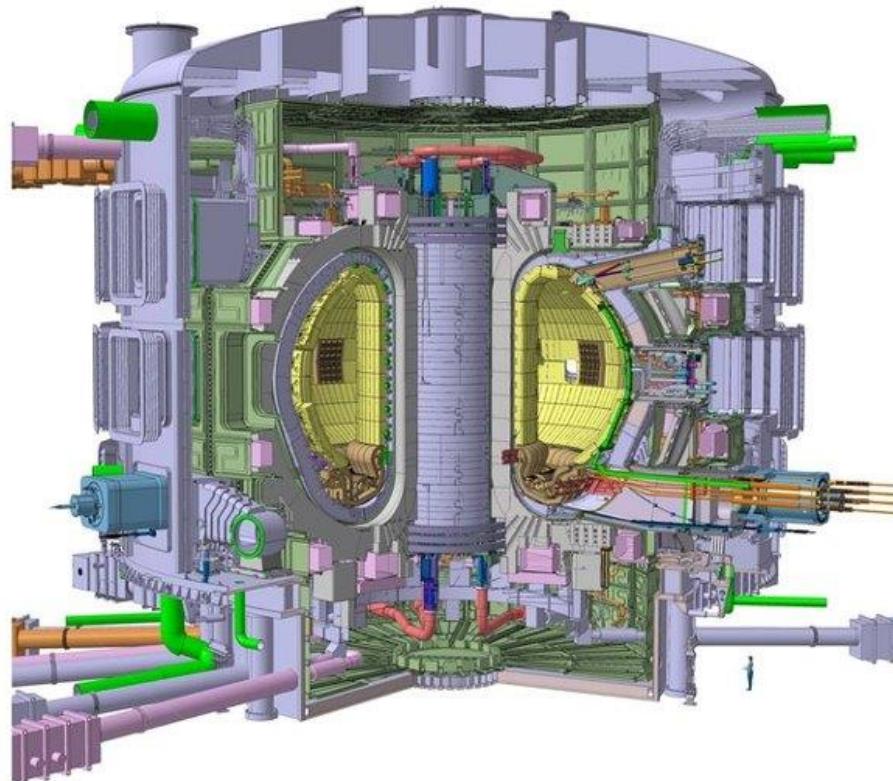
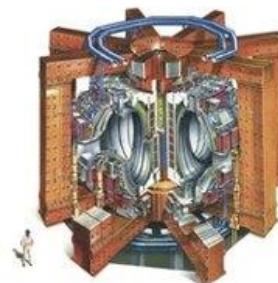
Fusion Technologies

IT, control, data acquisition

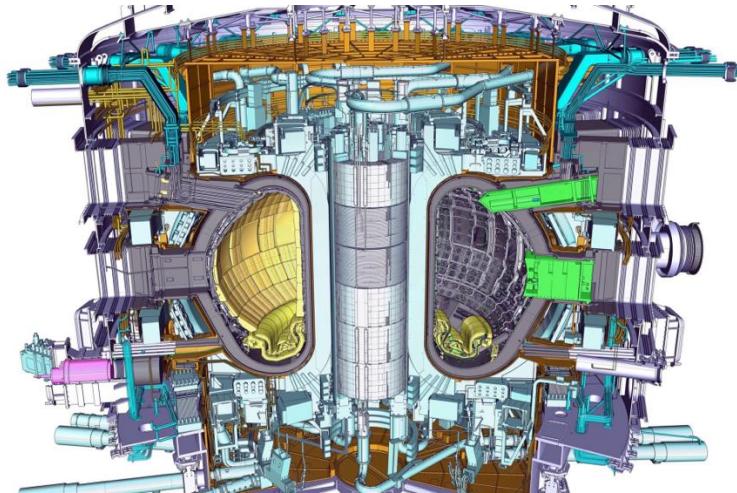
J. Decker

Special thanks to :
B. Duval, C. Galperti, SPC IT team

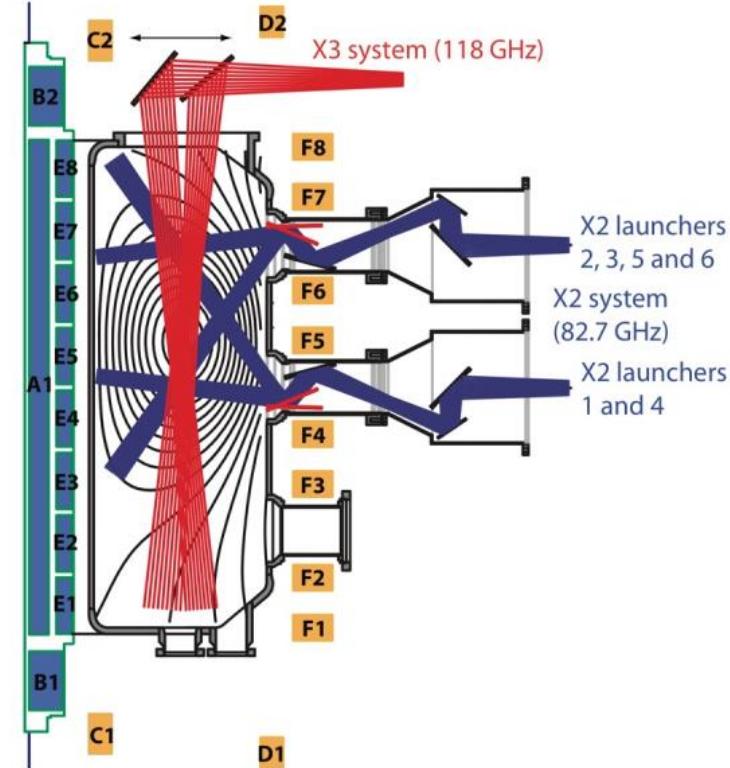
- Fusion experiments are complex interconnected systems
- A fusion plant requires safety protocols and 24/24 monitoring
- Magnetic fusion experiments operation mode is cyclic



- **Fusion experiments are digital systems**
 - Computers control (nearly) everything
 - Almost all measurements are voltages, currents, or (especially recently) images.
- **IT is central to all fusion activities**
 - Diagnostics
 - Actuators (heating, fueling, shaping, ...)
 - Real-time control
 - Data analysis
 - Theory

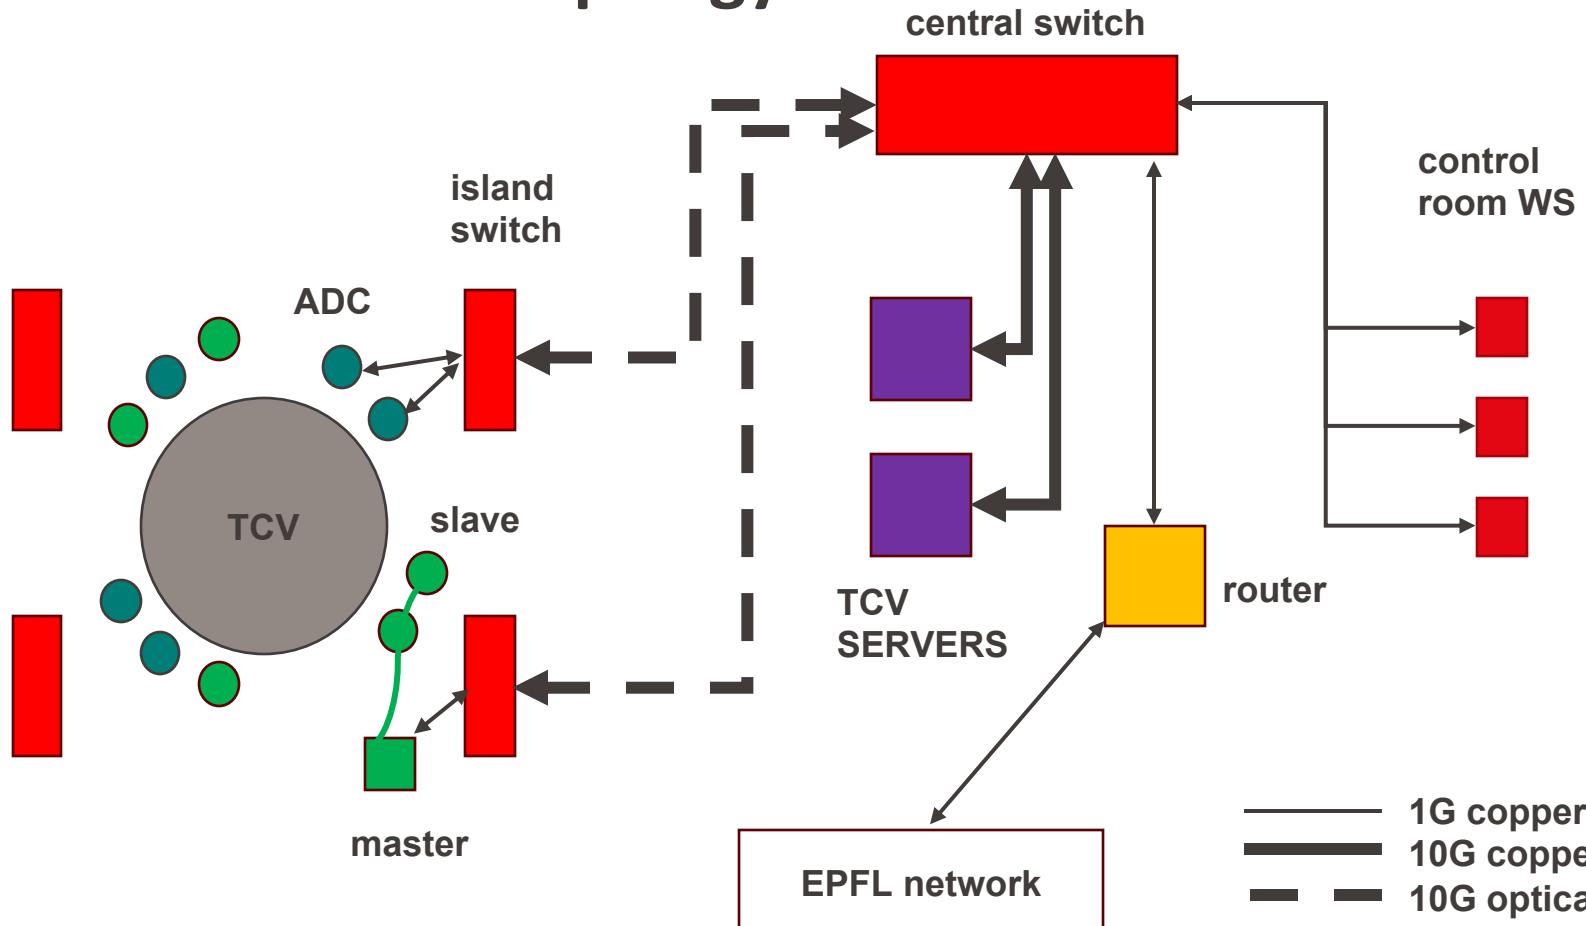


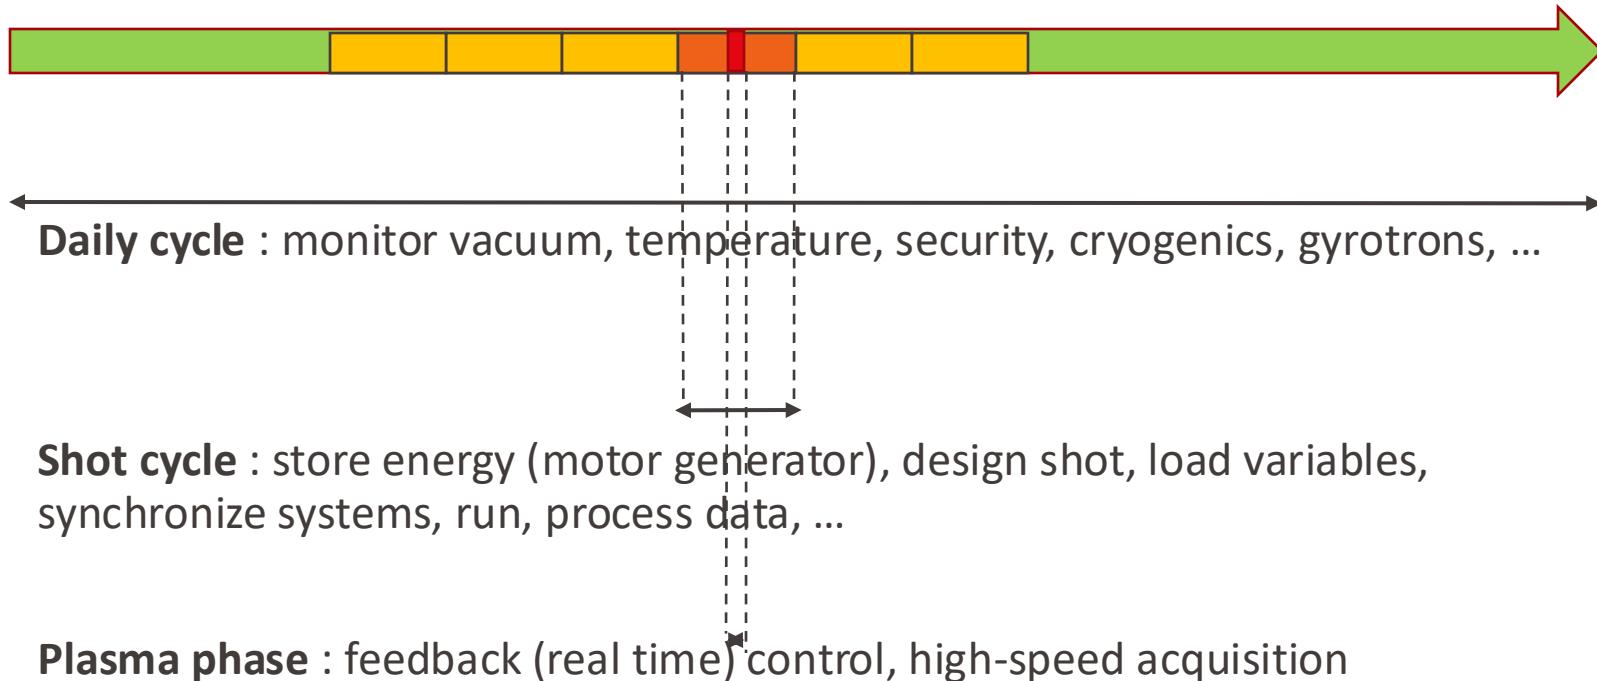
- Building a new diagnostic or actuator always requires advanced IT developments
 - Identify control variables and integrate in TCVCS (TCV control system)
 - Write drivers for equipment I/O
 - Program state machine and synchronization
 - Prepare structures for acquiring data
 - Setup signal digitization
 - Convert raw signals to physical quantities
 - Integrate in post-shot analysis



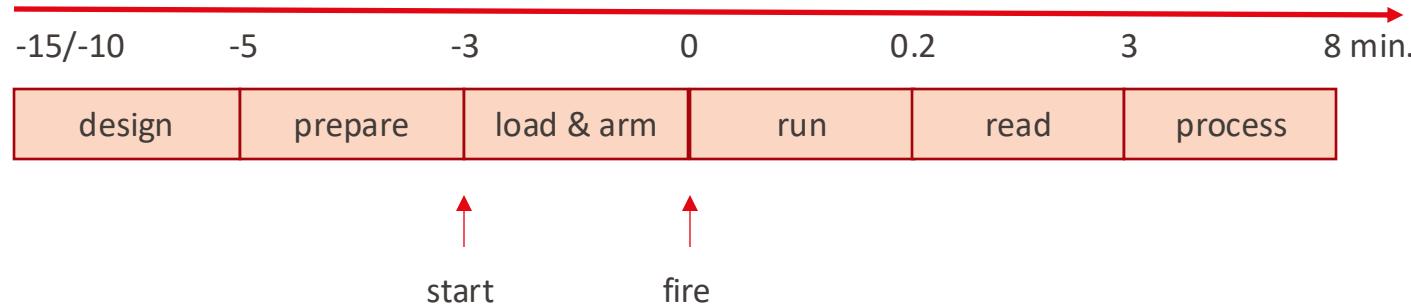
- Fusion experiments mix **hardware and software, physics and engineering**
-> **IT must integrate this diversity**
 - Hardware and software development are part of the research effort
- This talk focuses on TCV
 - 30 years of history, a remarkable mix of old and new technology
- Despite dozens of fusion research plants worldwide, there has been surprisingly little effort towards common IT solutions
- Control and data acquisition are increasingly intertwined
 - We need to save control data (for analysis, ML, and future shots)
 - We need to control data acquisition (setup, timing, ...)

- First Part
 - Fusion experiment timeline
 - Plant and plasma control
 - Real time control
 - Digitization & acquisition
- Second Part
 - Data processing
 - Data organization, data access, outreach
 - Data storage and backup
 - Fusion IT environment



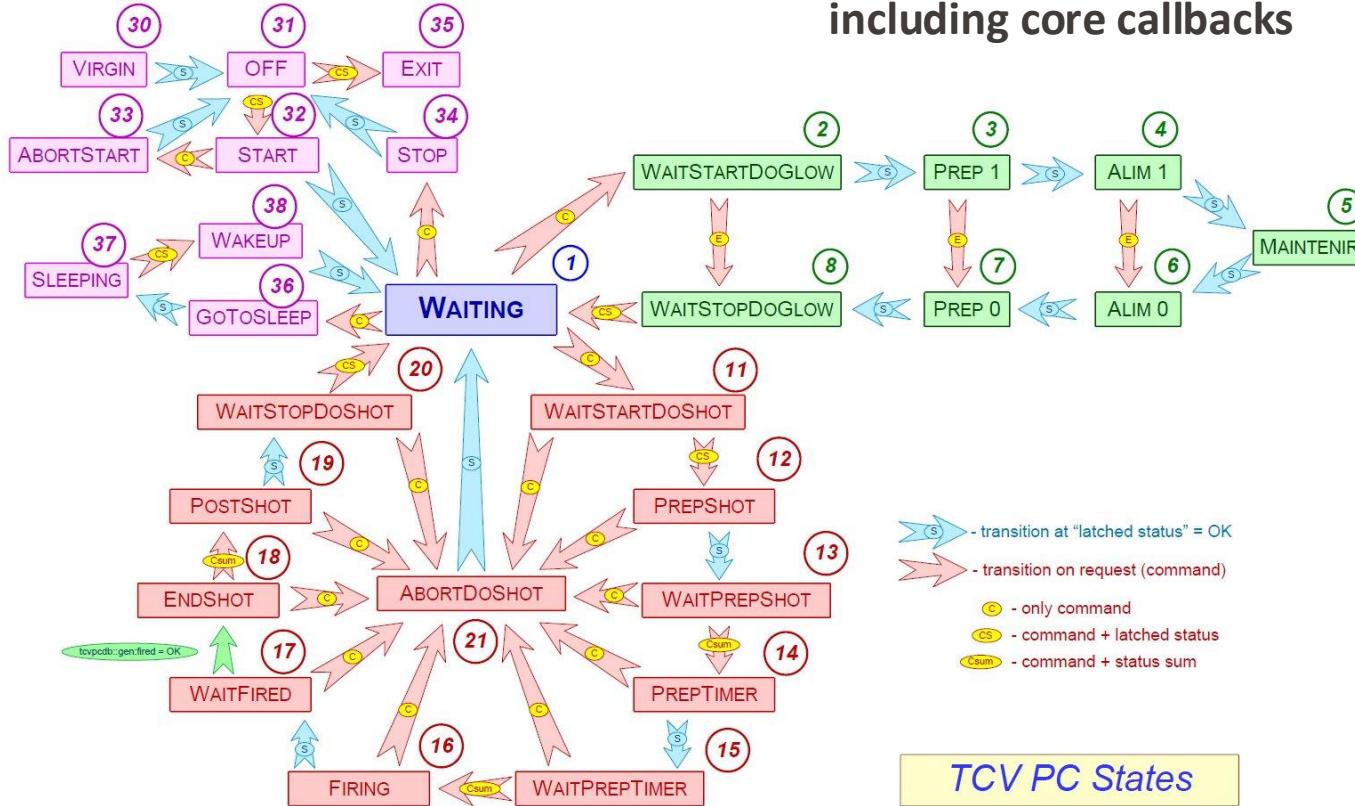


- Fusion experiments are based on shot cycles



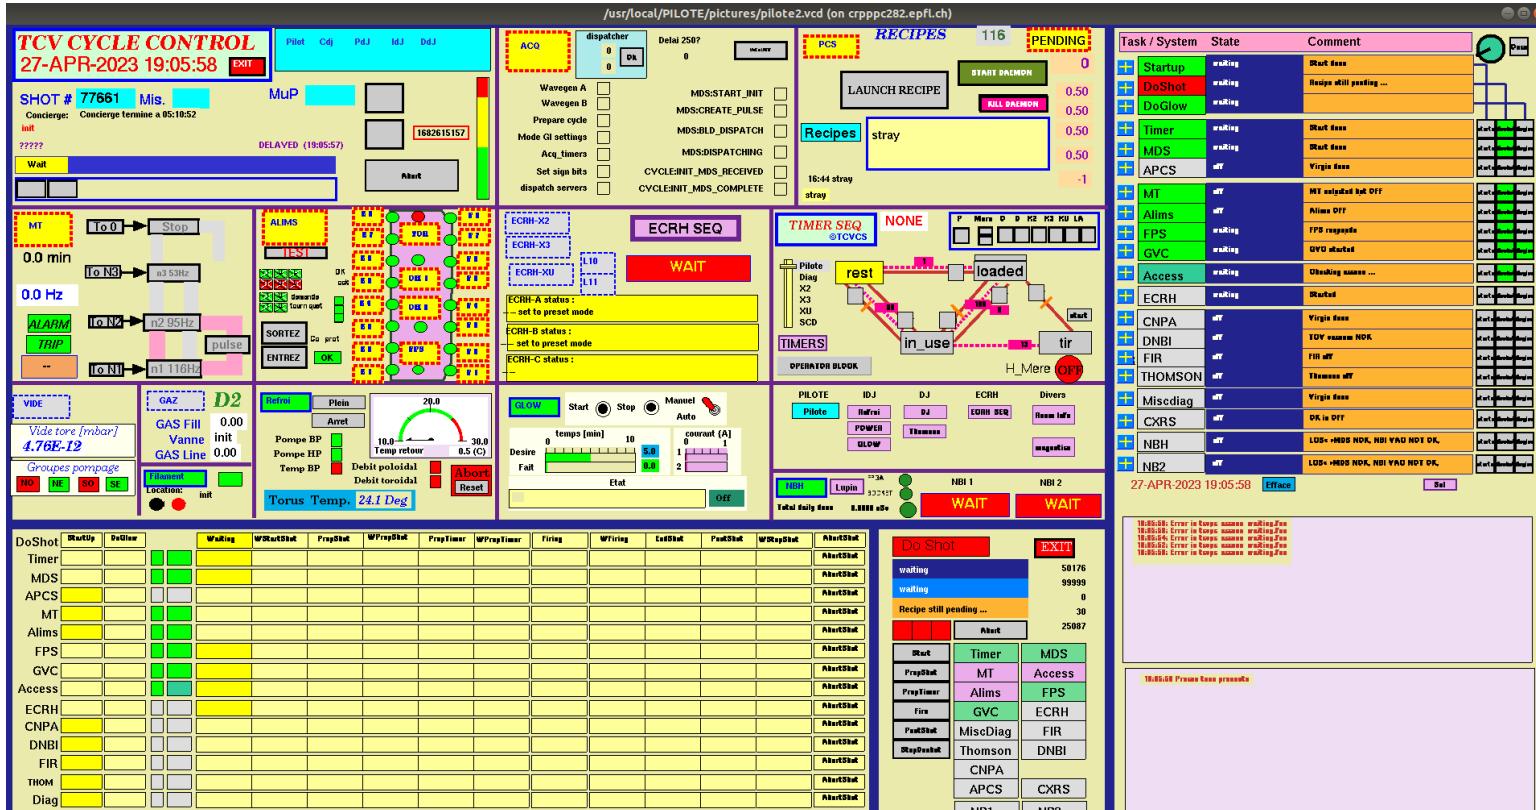
- design** (session leader) : plasma evolution, heating selection, diagnostics
- prepare** (pdj, ddj) : control waveforms, equilibrium prediction, timing, gains
- load & arm** (IT) : setup acquisition, load control variables into systems
- run** (pilot) : plasma discharge, feed-forward and feedback (real-time) control
- read** (IT) : transfer data from digitizers to server
- process** (IT + physicists) : analysis chain to calculate physical quantities

TCVPC is the TCV experiment state machine including core callbacks



TCV PC States

▪ Pilot main interface



- TCV live database : **VISTA CONTROL SYSTEMS**
- Two main channel type :
 - **IN**put : updated upon read or periodically
 - **OUT**put : callback upon update, includes use of “event”
- Multiple servers with remote access
- Multiple databases
- Permissions per channel or per database
- Tools :
 - **Vdraw** : control GUI
 - **Valarm** : graphical and sound alarm
 - **Vscript** : simple sequential language for “daemon”
 - **Vlogger** : storage of selected channels
 - **Vtrend** : graphical tool for visualizing logged data

- Example of out channels : TS gains
- Handler (callback) : C or fortran
- Access control : example

/vsys/crpppc282/db/TCV_PROTECTDB.access

DEFAULT: tcv_root=adefmnrx,
 \default=erx, tcv_idj=werx,
 \ech_oper=werx, \tcv_oper=werx

GAS_OPER ACCESS: tcv_root=adefmnrx,
 \default=erx, tcv_idj=werx,
 \tcv_oper=werx, gas_oper=werx

Note : channel-based access rights enabled

```
! sparam1      : i2c server IP adress
! hparam1      : i2c server port
! value        : initial value
! handler      : callback function
! automatic    : automatically write or read upon
channel access
```

```
DEFINE CRPPSRV23_I2C
•      sparam1 '10.27.128.209'
•      hparam1 50000
•      value    0
•      handler_c i2c_handler
•      ch_v3_function_style
•      automatic
```

```
! hfunction : i2c handler function selection
! lowlim 0  : low limit
! hilim 7   : high limit
```

```
DEFINE SETHFGAIN_I2C
•      integer out
•      hfunction 1512
•      lowlim 0 hilim 7
```

```
! hparam2      : i2c channel
! hparam3      : i2c address
```

```
$thomson:hf_gain:ch_060      SETHFGAIN_I2C      CRPPSRV2
3_I2C      hparam2 49      hparam3 40
```

- Alternative control framework :

EPICS

- selected by ITER
- used at SPC (gyrotron test bed)
- Multiple compatible protocols
- Requires I/O drivers

The following table contains a list of I/O devices for which EPICS device and/or driver support has been written, sorted by the primary I/O bus type. The table is also available sorted [by Manufacturer](#), [by Contact Name](#), or [by Link name](#) (you can also click on any active column title to switch views).

The related Soft Support database is available [here](#).

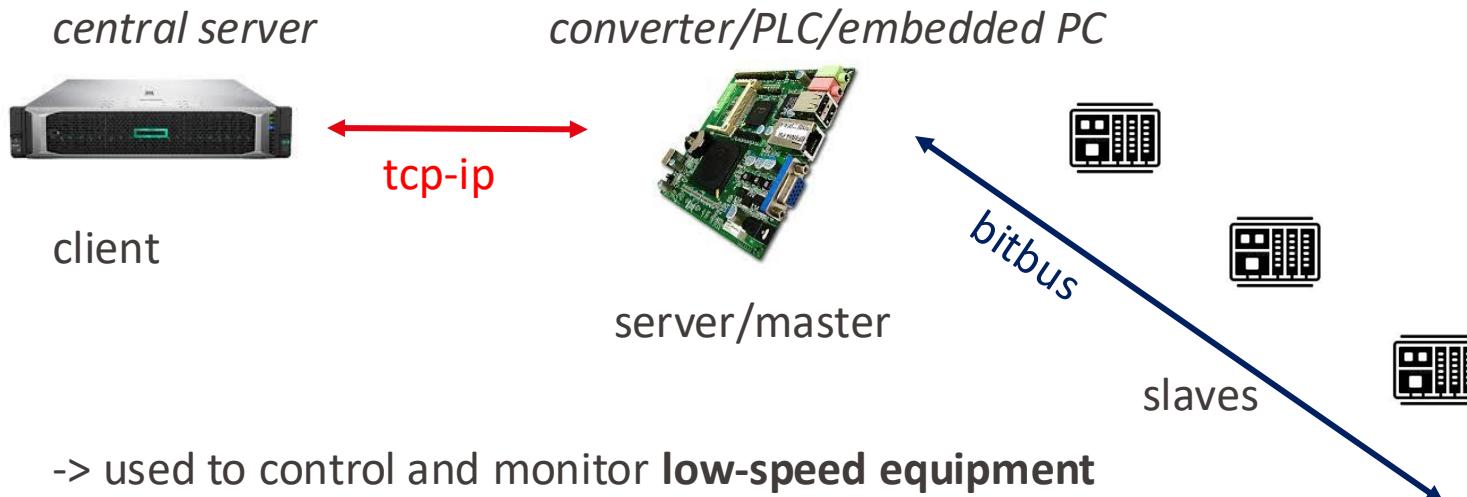
To submit a new entry to this list for your Hardware Support module, use [this form](#). Email corrections or questions about this page to [Andrew Johnson](#).

Contact email addresses have had an anti-spam filter applied. To mail anyone from the mailto: links below you will have to replace the strings "-AT-", "-DOT-" and "-DASH-" in the email address with an "@" , "." or "-" character respectively.

Contents

- [AB1746/1747](#)
- [AB1771](#)
- [AB1794](#)
- [ARM IOC \(Linux\)](#)
- [CAMAC](#)
- [CAN](#)
- [CAN \(iCS2\)](#)
- [cPCI \(devLib2\)](#)
- [CPCI/PXI](#)
- [EPP](#)
- [EtherCAT](#)
- [Ethernet](#)
- [Ethernet \(Asyn\)](#)
- [Ethernet \(HTTP\)](#)
- [Ethernet \(Modbus\)](#)
- [Ethernet \(Streams\)](#)
- [Ethernet SNMP](#)
- [Ethernet, USB, Firewire](#)
- [Ethernet-CAMAC \(Asyn\)](#)
- [Ethernet-GPIB](#)
- [Ethernet-GPIB \(Asyn\)](#)
- [Ethernet-Modicon](#)
- [Ethernet/NIM](#)
- [Firewire](#)
- [Firewire \(Asyn\)](#)
- [GPIB](#)
- [GPIB \(Asyn\)](#)
- [GPIB \(Streams\)](#)
- [Group3](#)
- [I2C \(Asyn\)](#)
- [IEEE1394 \(Asyn\)](#)
- [IOC \(Asyn\)](#)
- [IOC-Arduino \(Streams\)](#)
- [IOC-BPM](#)
- [IOC-CAN \(iCS2\)](#)
- [IOC-CPCI \(Linux\)](#)
- [IOC-DSO](#)
- [IOC-IPAC \(VxWorks\)](#)
- [IOC-PC/104 \(Linux\)](#)
- [IOC-RS-232 \(RTEMS\)](#)
- [IOC-Stepper \(RTEMS\)](#)
- [IPAC](#)
- [IPAC \(Asyn\)](#)
- [IPAC-CAN \(devLib\)](#)
- [IPAC-GPIB \(Asyn\)](#)
- [IPAC-RS-232 \(RTEMS\)](#)
- [IPAC-RS-232 \(VxWorks\)](#)
- [ISA-IPAC \(vxWorks\)](#)
- [Linux](#)
- [Modbus \(Asyn\)](#)
- [Modbus TCP](#)
- [Modbus/TCP \(Linux\)](#)
- [Modbus/TCP \(vxWorks\)](#)
- [PC/104 \(Linux\)](#)
- [PCI](#)
- [PCI \(Linux\)](#)
- [PCI, USB, GigE](#)
- [PCI-VME \(Linux\)](#)
- [PCI/CameraLink](#)
- [PCIe-IPAC](#)
- [PMC \(devLib2\)](#)
- [PXI](#)
- [RS-232](#)
- [RS-232 \(Asyn\)](#)
- [RS-232 \(Linux\)](#)
- [RS-232 \(ornlSerial\)](#)
- [RS-232 \(RTEMS\)](#)
- [RS-232 \(Streams\)](#)
- [RS-232 \(Windows\)](#)
- [RS-422](#)
- [RS-422 \(Streams\)](#)
- [RS-485](#)
- [RS-485 \(Streams\)](#)
- [RS232 \(Streams\)](#)
- [RS485 \(stream\)](#)
- [Streams](#)
- [TTL \(streams\)](#)
- [TTL \(Streams\)](#)
- [UIF Proprietary](#)
- [USB](#)
- [USB \(Asyn on Windows\)](#)
- [USB \(Asyn\)](#)
- [USB \(Streams\)](#)
- [USB_GigE, CameraLink](#)
- [USB_PCI, CameraLink](#)
- [VME](#)
- [VME \(Asyn\)](#)
- [VME \(devLib\)](#)
- [VME \(devLib2\)](#)
- [VME \(Linux\)](#)
- [VME-AB1771](#)
- [VME-Bitbus](#)
- [VME-CAMAC](#)
- [VME-ControlNet](#)
- [VME-GPIB](#)
- [VME-Group3](#)
- [VME-IPAC \(devLib\)](#)
- [VME-PC104 \(Linux\)](#)
- [VME-VXI](#)
- [VME64x-IPAC](#)
- [VXI](#)
- [VXI \(RTEMS via VXIGPIB\)](#)

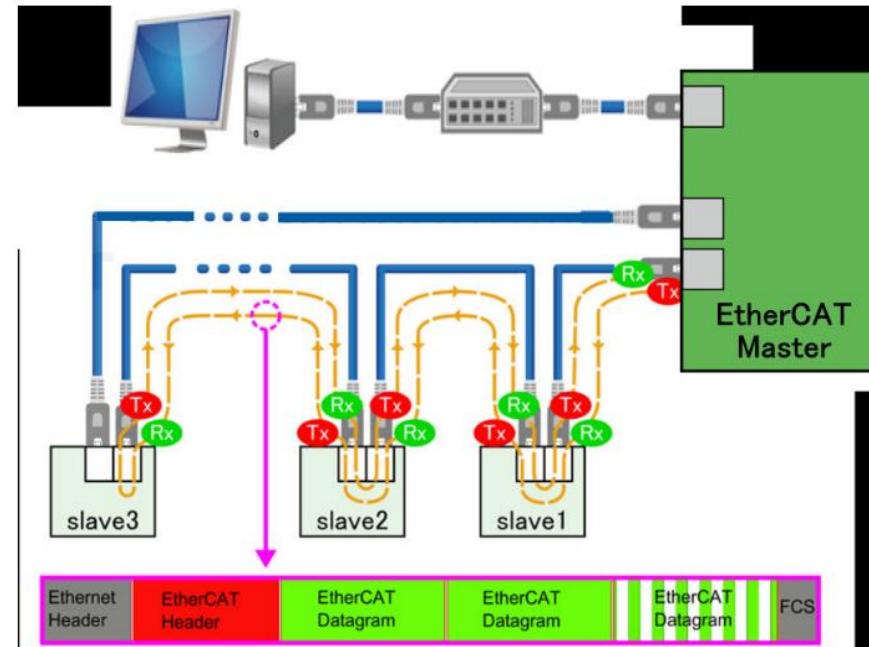
- TCV uses a combination of central and distributed control systems
- Central systems : device control from central computer via slave/master protocol (ex : BITBUS, RS232, MODBUS, I2C...)



-> used to control and monitor **low-speed equipment**
(pump control, power supplies, gains, timing setup, etc)

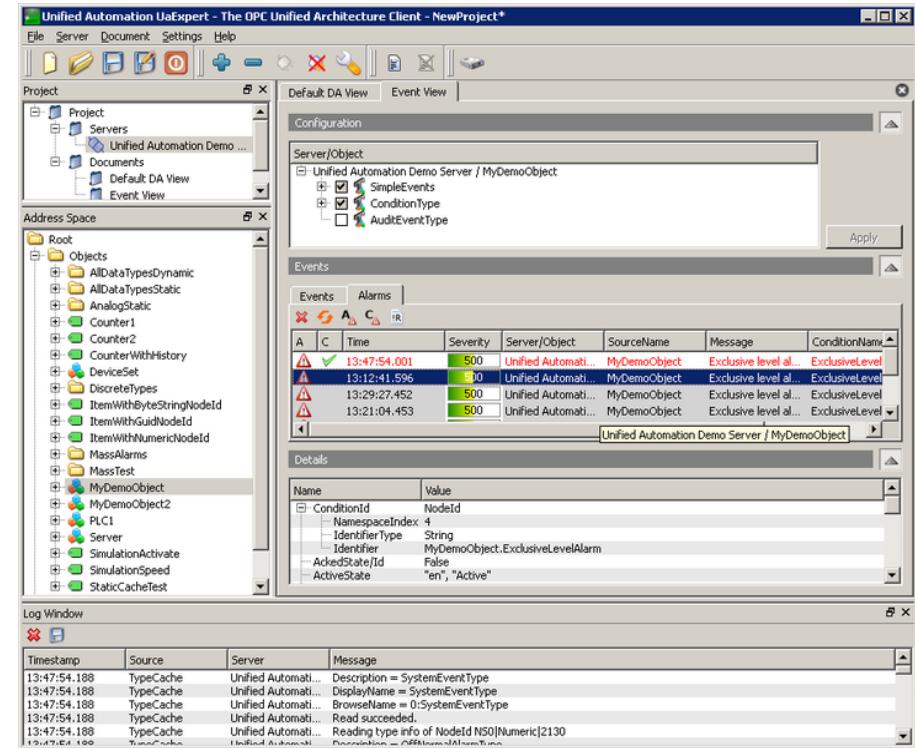
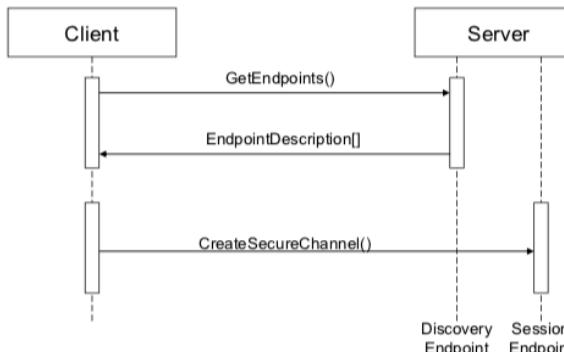
EtherCAT (Ethernet for Automation and Control Technology)

- Ethernet based control fieldbus.
- short data update times ($> 50 \mu\text{s}$)
- low communication jitter ($\leq 1 \mu\text{s}$)
- reduced hardware costs.



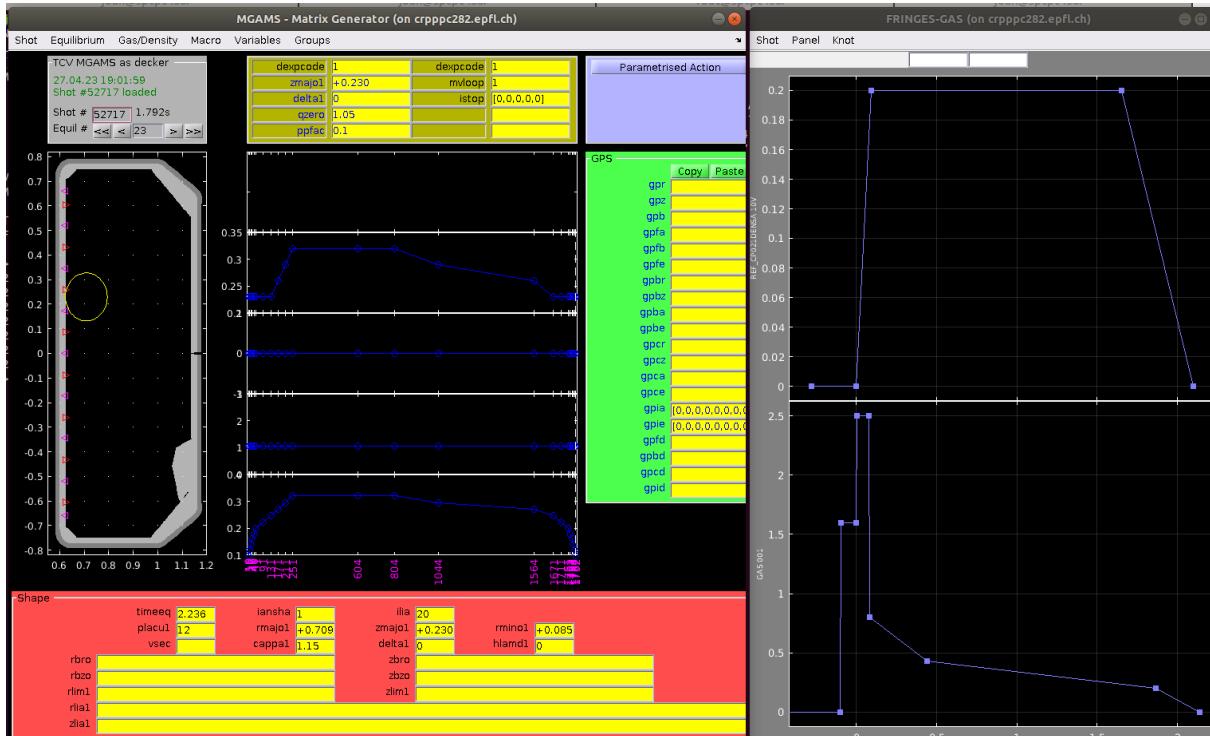
Distributed (decentralized) control : autonomous subsystems with own state machine

- Independent subsystem control PC
(i.e. **Labview** CODAC with NI PLCs)
- Advanced control protocols (**OPC-UA**)
 - Control channel list exposed by server
 - Secured communication
 - Self-describing variables (objects)

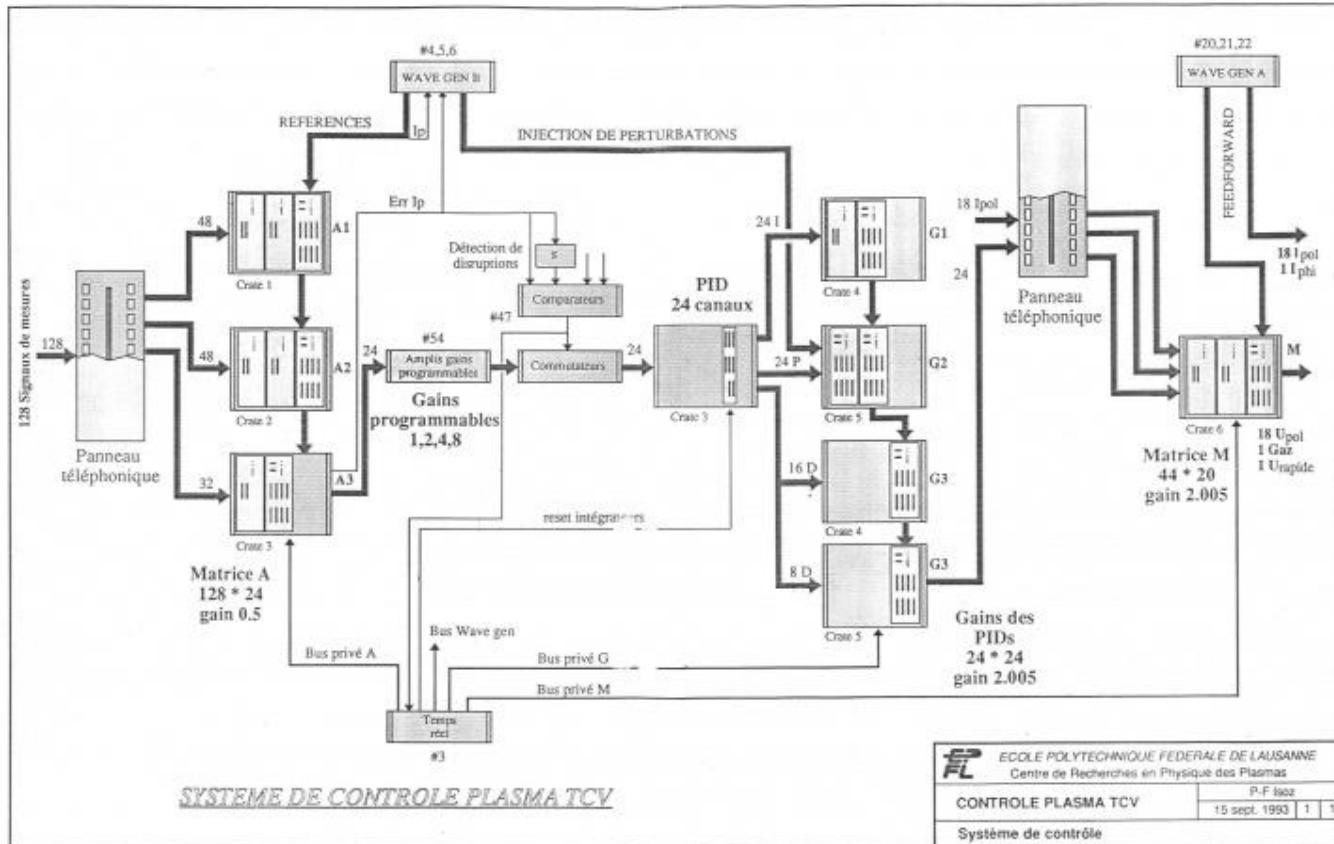


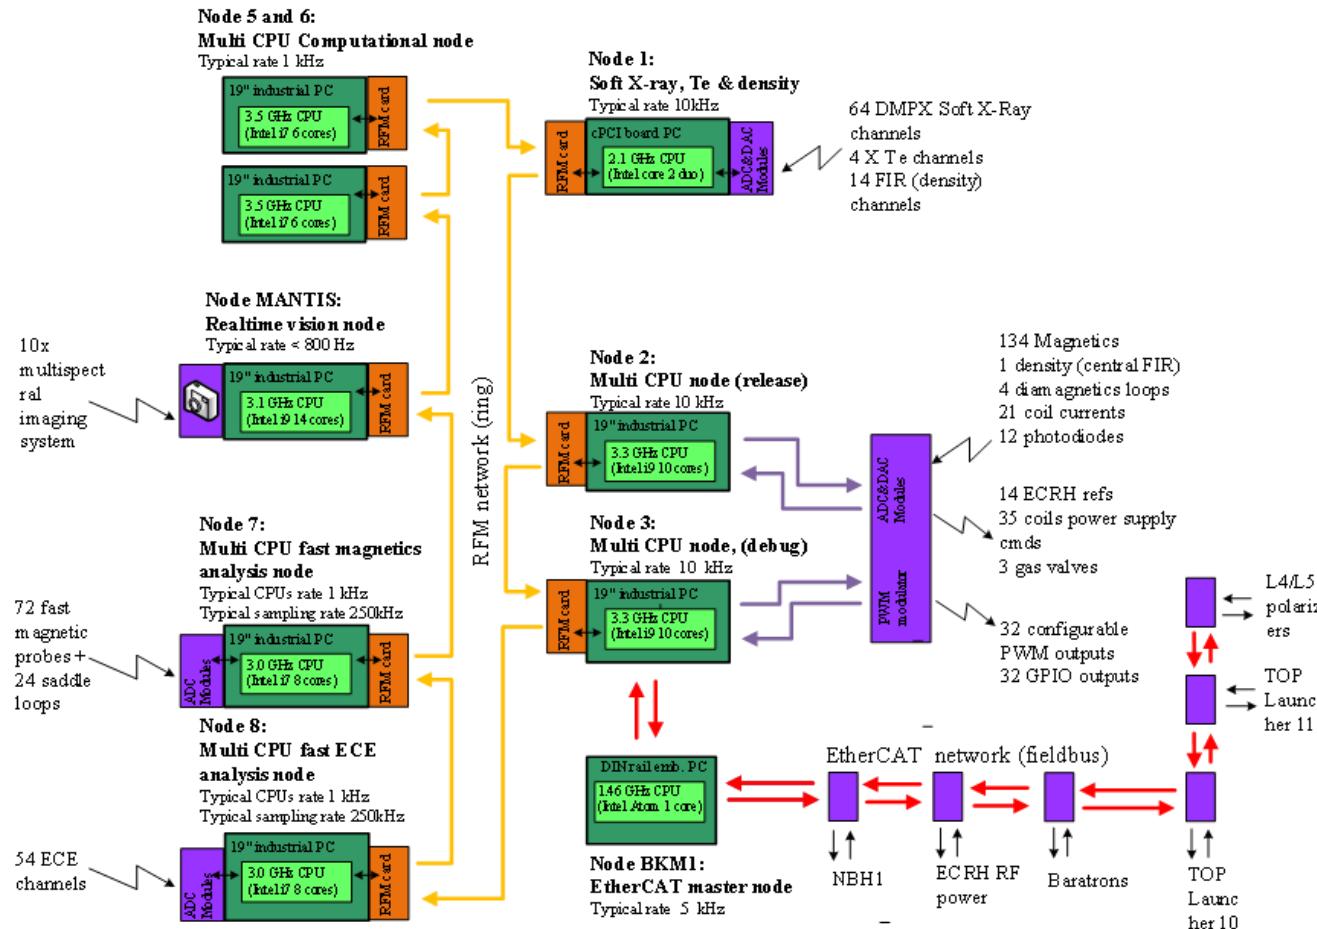
- Control a (limited) series of **actuators** through a combination of :
 - Feed-forward control waveforms
 - Feedback control by processing (real time) observations
- TCV may use thousands of observables to control a limited number of actuators:
 - Toroidal and OH coils (Bt and Ip)
 - Poloidal coils (plasma shaping)
 - Gyrotrons (EC heating and current drive)
 - Neutral beam injection systems (NBH and DNBI)
 - Fueling systems (gas valves + MGI)

- **Shot design interface : mgams**
- Secondary systems : echcs, nbhcs, fueling, etc

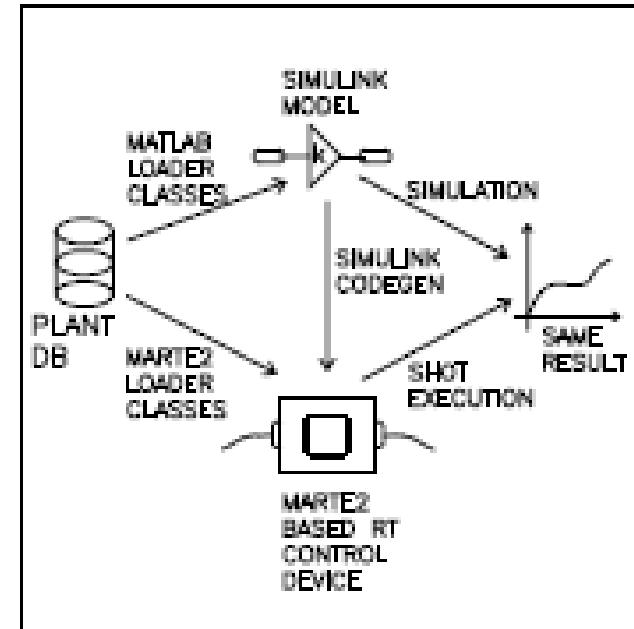


Analog real-time control

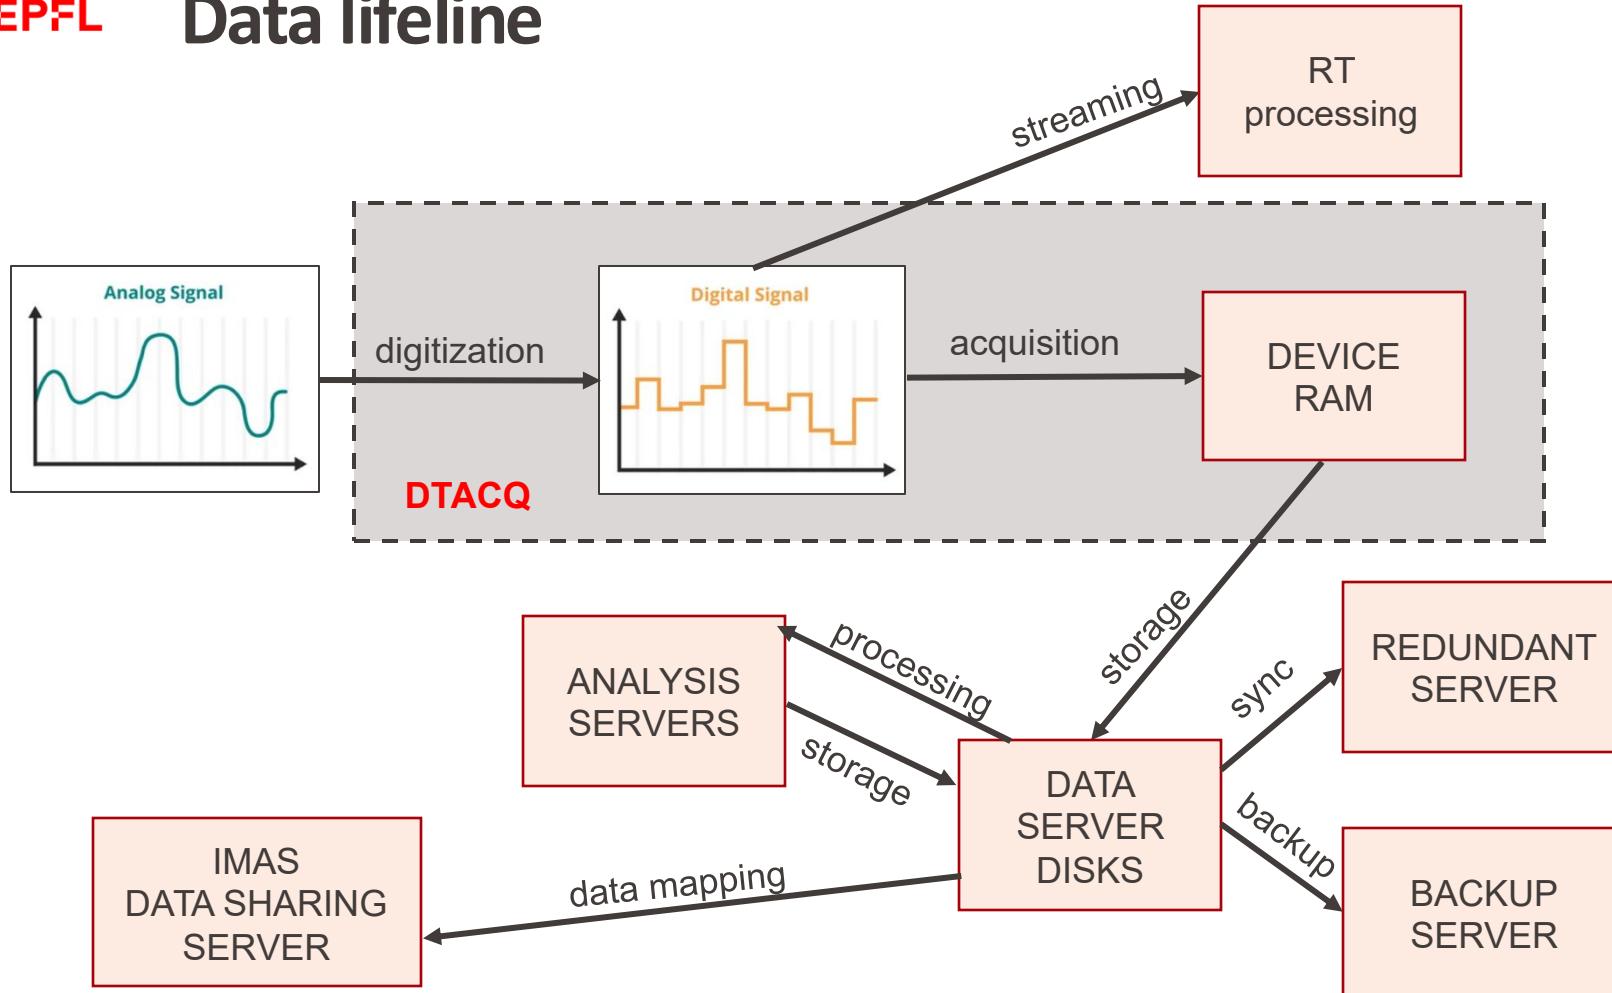




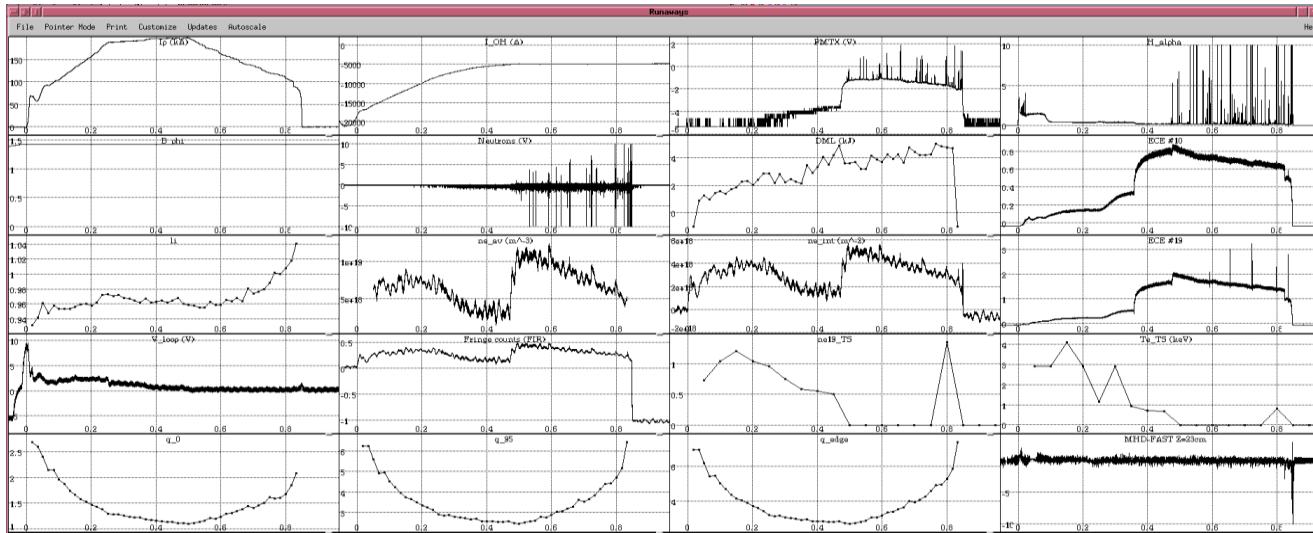
- Increasingly complex control algorithms
 - Equilibrium reconstruction (shaping)
-> real time **liuqe**
 - ECRH beam tracing
-> real time **torbeam**
 - Transport and current diffusion
-> real time **RAPTOR**



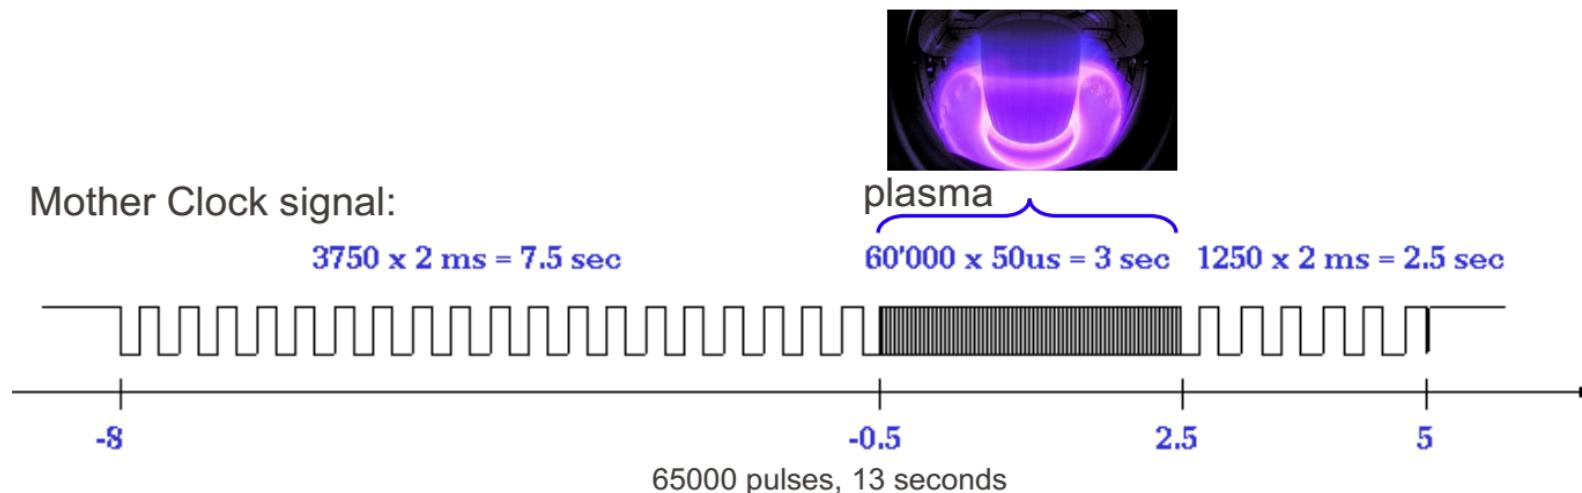
- Global platform for RT simulation and compilation using MATLAB/Simulink
- Global RT execution platform using MARTE2
- Deployment of AI-ML techniques
 - > complex events on very short timescales (disruption mitigation and avoidance)



- Causal analysis is essential in experimental plasma physics
- Tools dedicated to multi-signal time evolution comparison (i.e. dwscope, jScope)
- Ideally, synchronization precision must at least match acquisition rate
- For a long time, TCV sync precision was less than one microsecond
(but better for signals acquired on the same ACQ unit)
- Better synchronization enables cross-correlation analysis from different diag



- Main time source = “Mother Clock” (MC):
 - Bitbus slave generates mother clock signal and 1 MHz
 - 65000 clock ticks limited by 16 bit counter on MC signal
 - Slow (2 ms) period before and after the plasma phase
 - Fast (50 us) period for refined timing during plasma phase



- Excellent precision : < 1 ns !
- Fiber = no need for isolation



- D-TACQs (ACQ2106) shipped with WR
- Trigger times are programmed at endpoint level
- Inter-operable with legacy system
 - 1MHz can be generated locally = synchro

- Program acquisition modules (mdsplus Init action)
 - Model tree (enable actions, set # of samples and frequency)
 - Setup timers (WR or legacy)
- Arm acquisition modules
- Trig -> Run
- After shot, either :
 - Push data from dtacq
 - Send event for pulling data from central servers
 - Wait for postshot phase and pull from (mdsplus Store Action)
- Recent feature : share acquisition with RT systems
 - Instead of splitting signals, combine streaming and writing into RAM
 - Still some work to do... magnetic signals acquired 4 times...

Magnetic Fusion **acquisition rate** examples

- <1 Hz : monitoring
- 20-60 Hz : pulsed lasers (thomson scattering)
- 1 kHz : diags with large integration period
- 1 MHz : diags with short integration period
- 10 MHz : pulse shape processing
- 10 GHz : pulse reflectometry

-> different acquisition technologies required

Corresponding **synchronization technology**

- 1 – 10 ms precision : NTP (network time protocol, synchronisation)
- 1 μ s : TCV MHz clock and legacy trig system
- 0.1 – 1 μ s : PTP (precision time protocol, synchronisation + syntonisation)
- < 1 ns : White Rabbit (CERN), uses dedicated network hardware

- **Constraints**
 - Capture relevant time evolution
 - Relevant amplitude variations
 - Avoid saturation : gain controls (DDJ)
- **Digitization specifications**
 - Rate
 - Duration
 - Resolution
 - Dynamic range
 - Differential or not
 - Connectors
 - Streaming capabilities ?



CAMAC

DTACQ



Title: SPC Fast Acq Catalogue
 Date: 23/08/2024
 Author: L. Simons

Type	Company	Model Name	Stand Alone?	Max Chans	BW [MHz]	Acq rate [MHz]	Prec [bits]	Acq Time [s]	Sum of Price
ADC	Acqiris	Acqiris SA248P	n	1	2500	8000	14	0,500	- €
	Acquisys	ATS9364	n	2	1000	3000	12	0,171	- €
	CAEN	CAEN2730	y/n	32	250	500	14	0,156	- €
	D-tacq	ACQ482ELF	y	16	40	80	14	0,500	- €
	GUZIK	GUZIK ADP7000 series 10-bit	n	4	6500	16000	10	1,600	100 000,00 €
		VR8042	n	2	4000	10000	12	2,570	35 000,00 €
	Keysight	M5200A	n	4	2000	4800	12	1,200	- €
	NI	NI-5186	n	2	5000	6250	8	0,080	- €
		NI-9775	n	4	13,9	20	14	0,105	4 310,00 €
		PXIe-4481	n	6	10	1,25	24	0,000	14 100,00 €
	Red Pitaya	SIGNALlab 250-12 v1.2a	y	2	60	250	12	0,700	2 367,97 €
	Spectrum Instrumentation	M4x.22xx	n	4	1250	1250	8	0,800	- €
		M5i.33xx	n	2	4700	5000	12	0,400	- €
	Teledyne	ADQ14	y/n	4	1200	2000	14	0,537	- €
		ADQ30	y/n	1	1000	1000	12	4,295	- €
		ADQ32	y/n	2	1000	2500	12	0,859	9 260,00 €
		ADQ33	y/n	2	1000	1000	12	2,147	5 960,00 €
		ADQ35	y/n	2	2500	5000	12	0,429	13 750,00 €
		ADQ36	y/n	4	2500	2500	12	0,429	- €
		ADQ7/DC	y/n	2	3000	5000	14	0,200	20 690,00 €
	Vitrek	EON Express	n	2	1750	3000	12	1,333	16 750,00 €
Osc				36	5500	7437,5	23	0,635	68 970,76 €
PHA				3	50	162,5	26	0,000	11 965,00 €

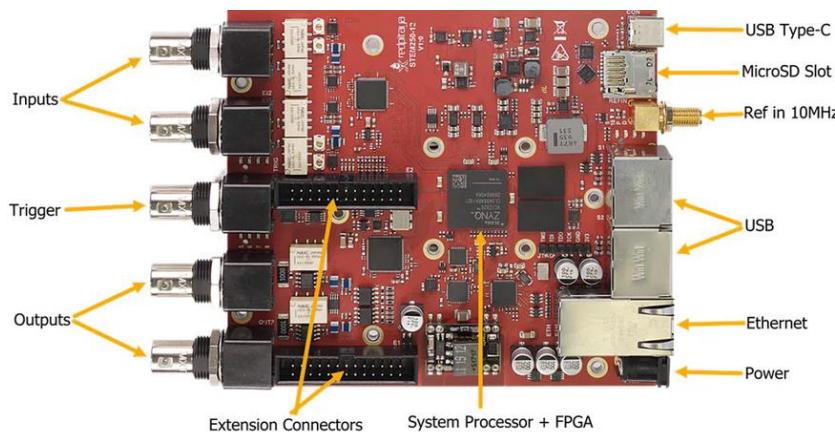
Courtesy of L. Simons

Examples of ultrafast Acquisition Modules

RED PITAYA 250-12



BOARD LAYOUT

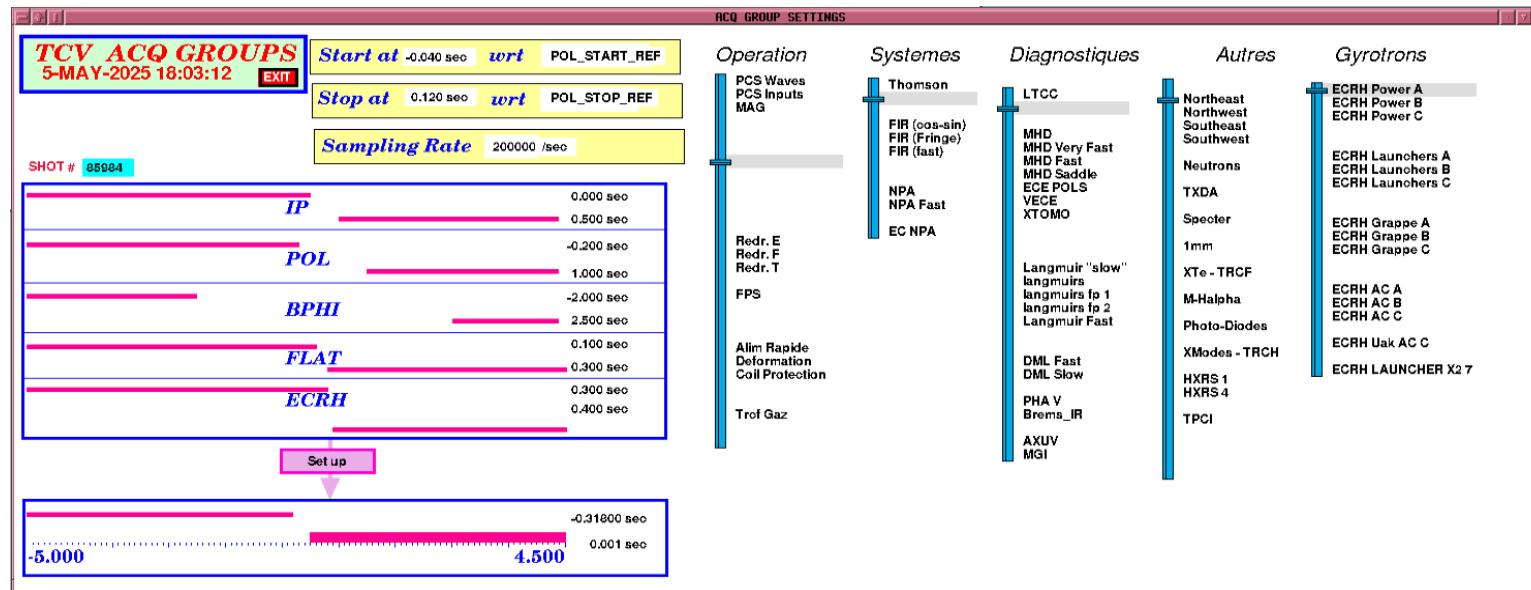


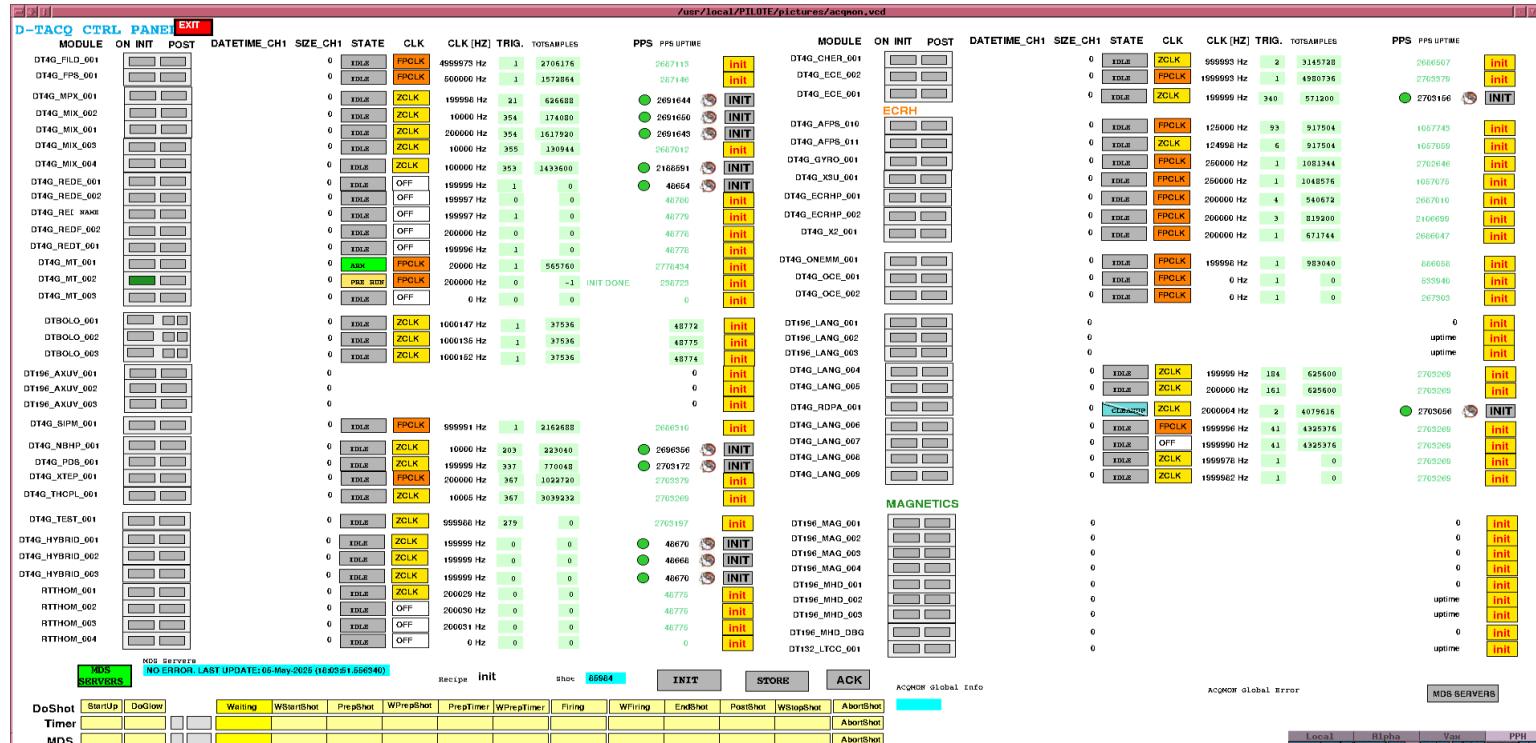
TELEDYNE ADQ35



Acquisition timers programming

- The number of acquired samples might be limited by board RAM
 - Compromise between sampling rate and acquisition window duration
- Acquiring useless data is detrimental (storage, I/O, backup, etc)
 - Acquire only relevant diagnostics during relevant shot phases
- Acquisition windows and sampling rate can be set in ACQ GROUPS or directly in VISTA





... to acquire **everything** for **as long as possible** at the **highest possible rate**, and keep it **forever**

- Imagine 50'000 channels, acquired on 16 bit resolution, for 10s, at 1MHz

-> this is **1 TB of data for one shot ! 40 TB in one day...**

- Can the network carry all this data between shots?
- Can we hold the data in one place?
- Can we process this data between shots?
- Can we back up the data overnight?

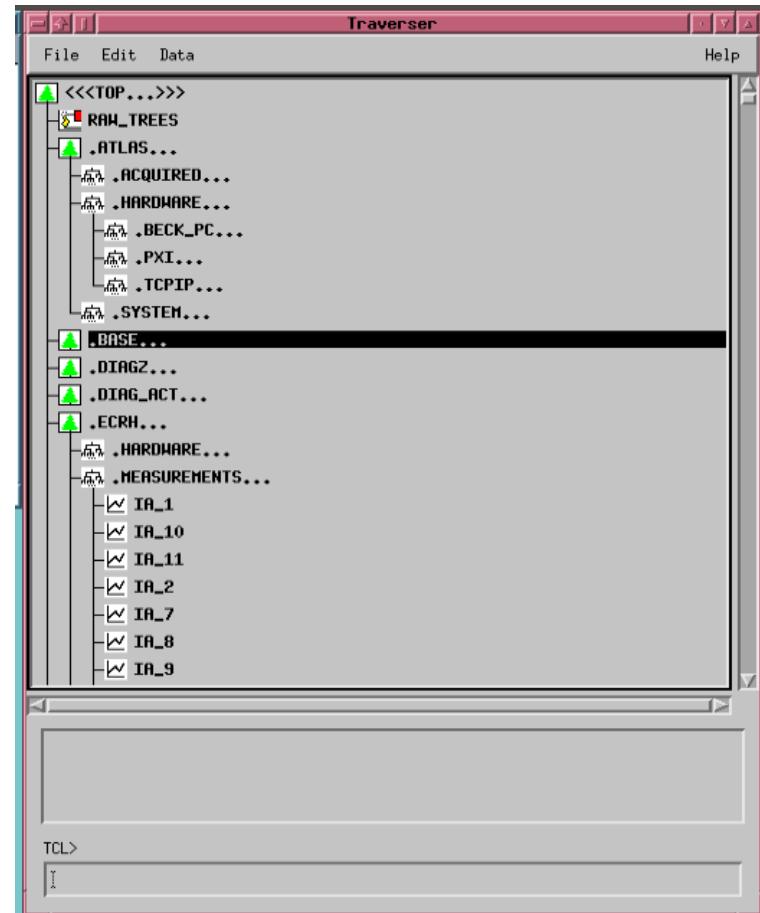
- Examples **good practices** :

- train of pulses -> keep peak amplitude and timestamp?
- Oscillating signals -> FFT and keep strongest/most meaningful features?

- FPGAs can help (process data between digitization and acquisition)
- Electronicians can help : pre-ACQ treatment of analog signals

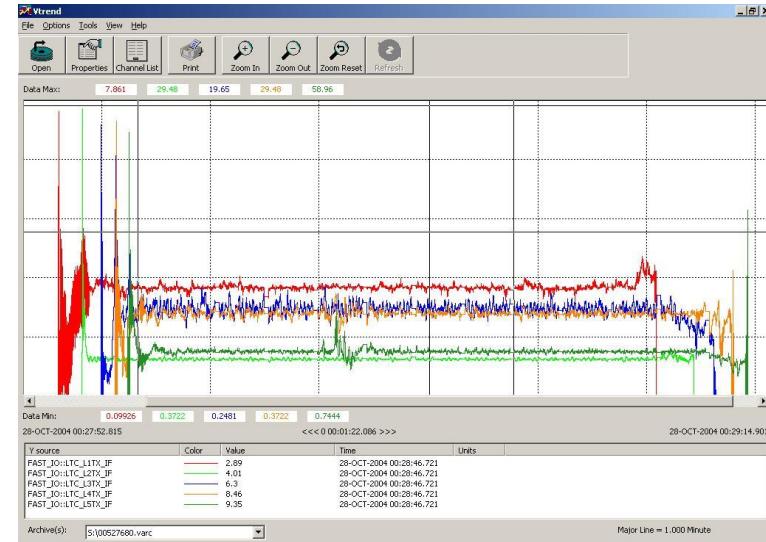
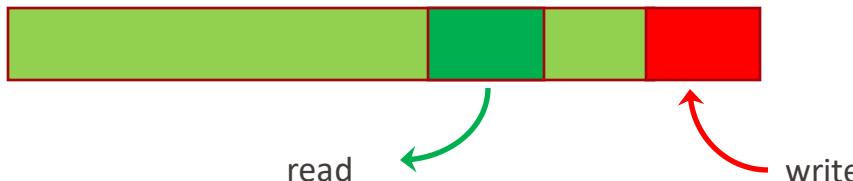
- First Part
 - Fusion experiment timeline
 - Plant and plasma control
 - Real time control
 - Digitization & acquisition
- Second Part
 - Mdsplus : data storage and organization
 - Data processing
 - Data access, outreach
 - Storage and backup
 - Fusion IT environment

- Large number of time traces (signals) with different timestamps
- Shot-based acquisition
- **MDSPlus** : data storage in shot-based files
- Intrinsic language : TCL/TDI
- Built-in objects (signals, data with units)
- Built-in basic data processing
- On the fly data compression
- Video storage poses different challenges
-> favor single frame or full movie access?



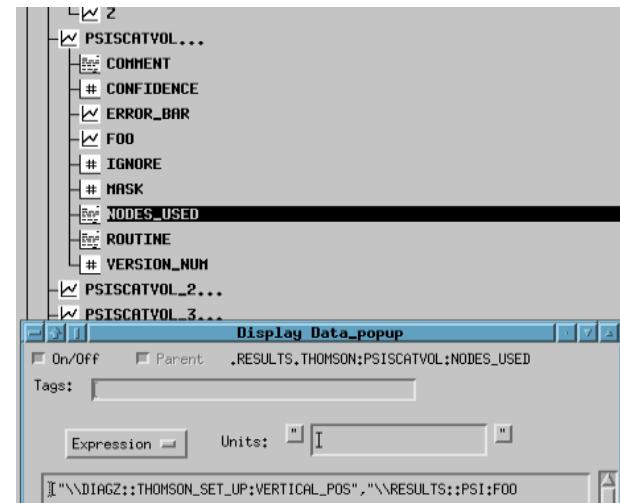
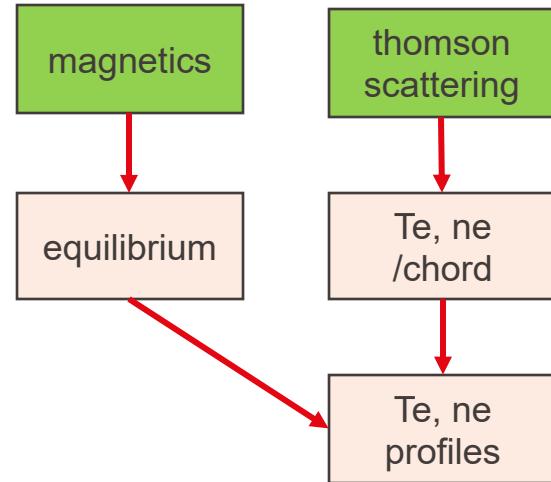
- **Challenge** : let authorized people in, keep everyone else out
- One powerful protocol : **SSH** (Secure SHell)
 - Ensures both authentication and encryption
 - Tunable based on hosts, users, groups, etc.
 - Passwordless key-pair authentication
 - X11 forwarding
 - Tunneling, proxy jumping
 - Command execution, filtering
- Several virtual desktop systems (NX, VNC) use ssh (x2go, nomachine, thinlinc, ...)
- **MDSPlus remote data access** is based on a IP protocol : **mdsip**
 - Possibility to use in thin or thick client
 - Can be encapsulated in SSH

- Read and store parameters beyond shot cycle. Ex : vessel pressure
- Heterogeneous data sources
- Data storage : shot -> day
- Solutions :
 - Vista Vtrend
 - MDSplus segmented data



- Note : future experiments will combine high-frequency acquisition (physics does not slow down) with long durations... how will they manage? Who knows....

- Data analyses are interdependent
- One analysis can depend upon a number of raw data samples and anterior analyses
 - analysis chain
- Enforcing consistent causality means updates must be followed downstream
- TCV solutions : **anasrv**
 - calculations upon request
 - calculation of upstream nodes if missing
 - recalculation of downstream node if outdated
- Thinking of a forward solution



- **Provenance** is metadata tracking the analysis chain to a given dataset with the possibility to reconstruct the dataset identically
- **Essential provenance metadata** includes :
 - Electronic treatment tunable parameters (gains)
 - Acquisition setup parameters (timing, precision, calibration, conversion)
 - Raw and processed data chain of dependences
 - Analysis codes versions
 - Analysis codes tunable parameters
- Metadata also include :
 - Data physical description
 - Data samples for **shot searching**, grouping, etc
-> average, min/max, steps, ramps, coarse timegrid....

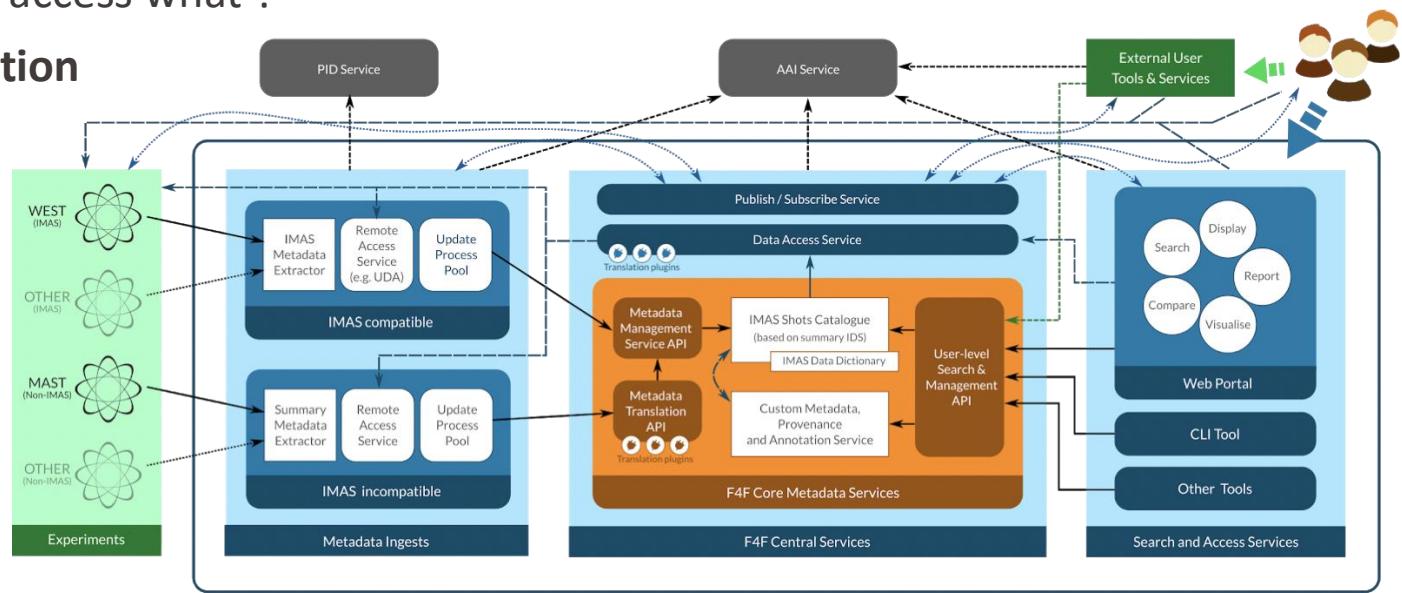
How many different TCV data subsampling systems do we have at SPC?

- TCV : trial indices, derived trees
- Essential when carrying out multiple shot analyses
- Must decide on default version (standard, latest, intelligent)
- Capacity to browse among versions -> annotations
- Complexity increases exponentially with the number of interdependent analyses

-> So far, a satisfactory global data versioning system is still missing at TCV

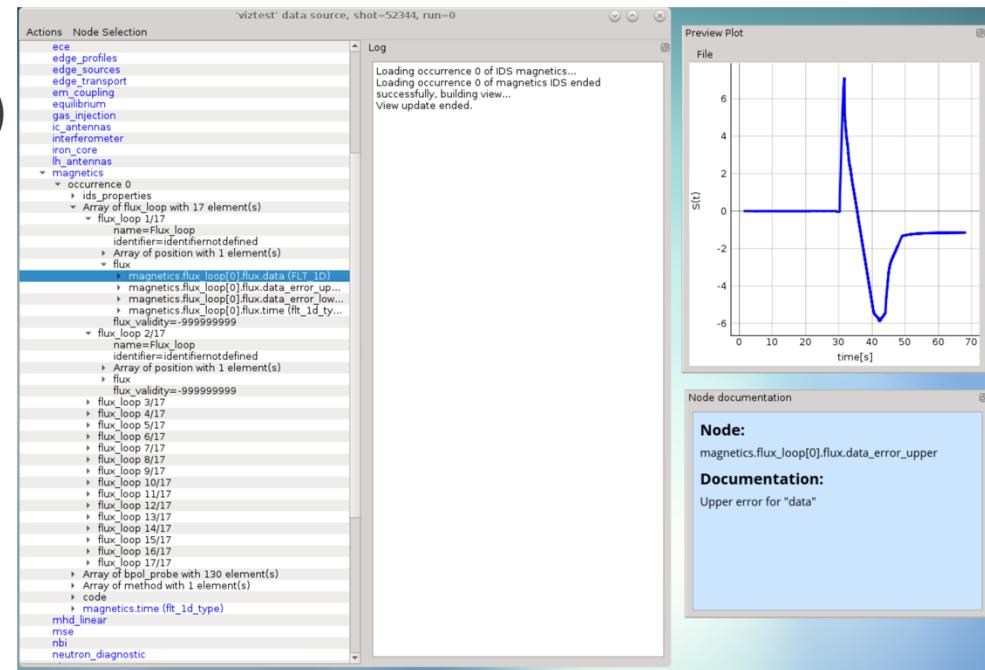
- Public funding imposes open data
- Delay to protect original work
- Meaningful data sharing is quite complex
- A&A : who can access what ?
- **Data identification & definition**
 - Dictionary
 - Metadata
 - Ingestion
 - Format
 - Search
 - Access

FAIR4FUSION
-> Eurofusion DMP



- Integrated Modelling and Analysis suite (IMAS)
 - Interconnected software for ITER modelling
 - Modular selection of components

IMASViz



- Interface Data Structures (IDSS)
 - Organised as tree structures
 - Describe both experimental and modelling data
 - Work for any device
 - Describe plant systems and physics object
 - Self-describing objects
 - Information on provenance

IMAS Data Model 3.34.0

Heating systems

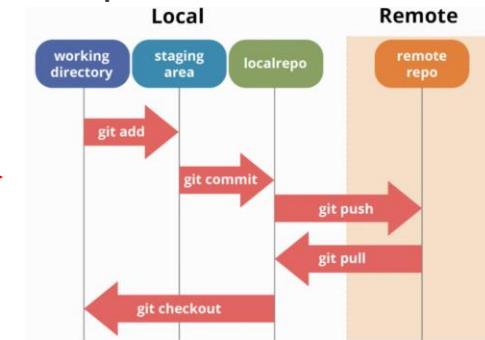
Diagnostics

amns_data	disruption	iron_core	reflectometer_profile
barometry	distribution_sources	langmuir_probes	refractometer
bolometer	distributions	lh_antennas	sawteeth
bremsstrahlung_visible	divertors	magnetics	soft_x_rays
calorimetry	ec_launchers	mhd	spectrometer_mass
camera_ir	ece	mhd_linear	spectrometer_uv
camera_visible	edge_profiles	mse	spectrometer_visible
charge_exchange	edge_sources	nbi	spectrometer_x_ray_crystal
coils_non_axisymmetric	edge_transport	neutron_diagnostic	summary
controllers	em_coupling	ntms	temporary
core_instant_changes	equilibrium	pellets	thomson_scattering
core_profiles	gas_injection	pf_active	tf
core_sources	gas_pumping	pf_passive	transport_solver_numerics
core_transport	gyrokinetics	polarimeter	turbulence
cryostat	hard_x_rays	pulse_schedule	wall
dataset_description	ic_antennas	radiation	waves
dataset_fair	interferometer	real_time_data	workflow

- The tokamak WEST has selected IMAS for experimental data storage organisation

- **Findable** : capacity to search through datasets with meaningful criterias/filters
- **Accessible** : A&A solution, API
- **Interoperable** : Data definition, format
- **Reproducible** : provenance, versioning
- Are we there yet...?
 - Findable? Logbook, alma database, IMAS summary IDS
 - Accessible? Remote access servers, LDAP A&A
 - Interoperable? Eurofusion DMP effort (IMAS) (3 years from now)
 - Reproducible? Not systematically

- Simultaneous code development by several people (and even several teams) can be a strong accelerator, but requires organisation and discipline
- An essential requirement is **code versioning**
 - centralised systems : CVS, SVN
 - distributed systems : GIT (gitlab, github)



- Verification, moderation, merging, and documentation of large codes are essential tasks that :
 - Require a significant effort
 - Fall on people with a deep understanding of the code (core developer)
 - Are generally not provisioned by labs nor rewarded by the community

- SPC servers are attacked constantly
- When one protection fails, attack rates explodes
- Password strength matters

- 2FA is effective
- Some essentials
 - Set up firewall
 - Enable selinux
 - No root access with password (use SSH key pairs or su from user account)
 - Enable fail2ban (delay or block following unsuccessful attempts)

Number of Characters	Numbers Only	Lowercase Letters	Upper and Lowercase Letters	Numbers, Upper and Lowercase Letters	Numbers, Upper and Lowercase Letters, Symbols
4	Instantly	Instantly	Instantly	Instantly	Instantly
5	Instantly	Instantly	Instantly	Instantly	Instantly
6	Instantly	Instantly	Instantly	Instantly	Instantly
7	Instantly	Instantly	2 secs	7 secs	31 secs
8	Instantly	Instantly	2 mins	7 mins	39 mins
9	Instantly	10 secs	1 hour	7 hours	2 days
10	Instantly	4 mins	3 days	3 weeks	5 months
11	Instantly	2 hours	5 months	3 years	34 years
12	2 secs	2 days	24 years	200 years	3k years
13	19 secs	2 months	1k years	12k years	202k years
14	3 mins	4 years	64k years	750k years	16m years
15	32 mins	100 years	3m years	46m years	1bn years
16	5 hours	3k years	173m years	3bn years	92bn years
17	2 days	69k years	9bn years	179bn years	7tn years
18	3 weeks	2m years	467bn years	11tn years	438tn years



- OS selection : **rocky linux** (centOS, fedora, ...). Compromise stability vs update
- Filesystems : ext4, **xfs**, zfs
- Partitioning : **GPT** : boot (EFI), space management (/home, /var/log)
- Storage arrangement : RAID 0/10 (speed), RAID 5/6 (volume)
- Invest in **RAM** : Swapping kills performance, FS caching is a powerfull feature
- Software : **minimalist approach** (faster, less maintenance, reduced risk of conflict)
- Redundancy of power supplies, network interfaces (bonding)
- Use UPS (Uninterrupted Power Suply)
- ILO/IDRAC : onboard chip accessible if server is down. Hardware survey, console
- Monitoring : **Zabbix** , Nagios
- Backup

ZABBIX Monitoring Inventory Reports Configuration Administration

Dashboard Problems Overview Web Latest data Graphs Screens Maps Discovery Services

Global view

All dashboards / Global view [Edit dashboard](#)     

System information

Parameter	Value	Details
Zabbix server is running	Yes	localhost:10051
Number of hosts (enabled/disabled/templates)	135	46 / 0 / 89
Number of items (enabled/disabled/not supported)	3540	3005 / 481 / 54
Number of triggers (enabled/disabled/problem/ok)	1736	1053 / 683 [24 / 1029]
Number of users (online)	5	2
Required server performance, new values per second	37.03	

Problems by severity

Host group ▲	Disaster	High	Average	Warning	Information	Not classified
Linux servers	1	3	15	3		
Templates/Operating systems		1				
Virtual machines		1				

Problems

Time ▾	Info	Host	Problem • Severity	Duration	Ack	Actions	Tags
2023-04-30 01:00:40	April	TSM server	RAID array device MD md127 is in check sync action on TSM server	2d 8h 27m	No		
2023-03-11 04:01:35	2023	spcpc126	Free disk space is less than 20% on volume /	1m 22d 4h	No		
2022-12-09 09:44:39	December	spcsrv7	Zabbix agent on spcsrv7 is unreachable for 5 minutes	4m 23d 22h	No	4	
2022-12-08 15:19:18	2022-11-18 19:02:12	icvm0046.epfl.ch	Zabbix agent on icvm0046.epfl.ch is unreachable for 5 minutes	4m 24d 17h	No	4	
2022-11-03 09:51:01	November	spcsrv8	Free disk space is less than 20% on volume /data	5m 14d 13h	No		
2022-10-29 04:31:10	October	spcsrv6	Free disk space is less than 20% on volume /data-vboxadmin	5m 29d 22h	No		
2022-06-26 18:08:30	2022-06-26 18:08:30	tcvdata	Free disk space is less than 20% on volume /Terra16	6m 5d 4h	No		
		spcipc1	Zabbix agent on spcipc1 is unreachable for 5 minutes	10m 9d 15h	No	2	

Local



Favourite maps

No maps added.

Favourite graphs

No graphs added.

technology	speed	capacity	CHF/ TB
SD	25 MB/s	< 0.25 TB	200
USB 3.0	80 MB/s	< 4 TB	50
HDD	150 MB/s	< 20 TB	15
SATA SSD	500 MB/s	< 4 TB	50
NVMe SSD	2000 MB/s	< 4 TB	50

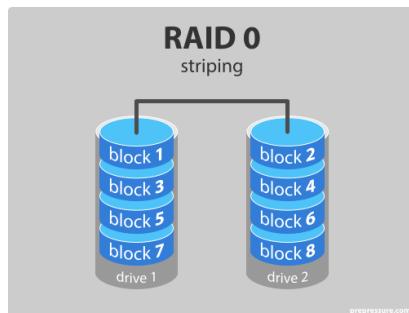
Compare with network speed :

- 100 Mbs : 11 MB/s
- 1 Gbs : 110 MB/s
- 10 Gbs : 1100 MB/s
- Fastest storage at TCV (SPR diag) : 4 x raid0 NVMe : 8GB/s (~RAM speed)

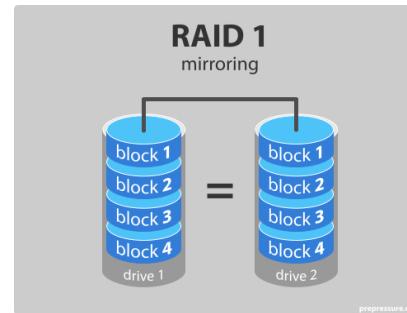
RAID array

Redundant Array of Independent Disks

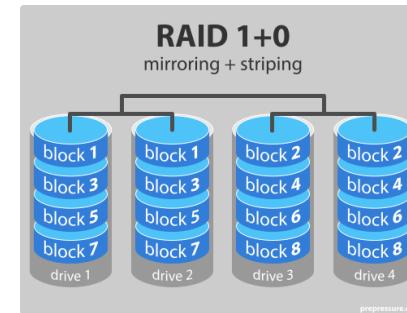
RAID0 : stripping



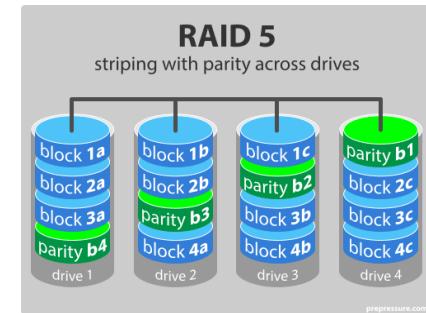
RAID1 : mirror



RAID10 : combine



RAID5 : parity



Fast but risky

Safe but slow

Fast and safe
but expensive

Fast (read)
and safe

- **To not back up data is an act of carelessness or... faith ?**
- Important (and/or voluminous) data should be backed up **at least twice**
 - If possible, **in different buildings**
 - If possible, using different technology
- Different backup types
 - Copies – archives – differential – incremental
 - **System restorable** or not (DRM protection)
 - Drive, partition, file backups
 - Redundant systems (**replication**)
- Some remarks :
 - **Restore operations must be tested** (including bootable full system restores)
 - If the main drive is encrypted by malware, will the backup be impacted ?
 - **Think of all possibilities**

- **IBM spectrum protect**
 - Secure (the client asks the server what to do, the server cannot access the client)
 - Multi platform incremental backup
 - Monitoring, alerts, ...
 - Replication : seamless synchronization with second server
 - Deduplication : mark identical binary data blocks and save space
 - Strategies : number of versions kept, delay to keep outdated objects
- **Rsync** : an amazing tool (coming out of a Ph.D. thesis) : copies, backups, etc
- Atempo Lina : official EPFL solution

SPC IT team

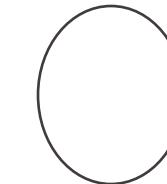
- **SPC wiki**
<https://spcwiki.epfl.ch/>
- **Redmine** knowledge base (ticketing system)
<https://spcsrv18.epfl.ch/redmine/>
- **EPFL Gitlab**
<https://gitlab.epfl.ch/>
- **SPC user web pages**
<https://crppwww.epfl.ch/~<username>/>
<https://crpplocal.epfl.ch/~<username>/>
- IAEA meeting on Control, Data Acquisition and Remote Participation (every two years)



Christos
Paraskevopoulos



Joan
Decker



Pierre
Etienne



Mateusz
Gospodarczyk



Luke Simons



Cenk Yildiz

