



Science and  
Technology  
Facilities Council



# EPAC Project Overview of the EPAC control system

Project	Doc Type	Work Package	Number	Version	Status
EPAC	RP	WP11	2	A	OFR

## Document Created at Issue A

Created By: *Chris Gregory*

CR Number: xx

Details: xx

# Contents

<b>1.</b>	<b>Introduction to EPAC .....</b>	<b>3</b>
<b>2.</b>	<b>Control system architecture.....</b>	<b>4</b>
2.1	EPICS .....	5
2.2	PLCs .....	5
2.3	Pulse ID and timing system .....	6
2.4	Detectors.....	6
2.5	Personnel safety system and interface to control system.....	6
2.6	Machine protection system .....	6
2.7	User interfaces .....	6
2.8	Motion control.....	7
2.8.1	Slow in/out positioners .....	7
2.8.2	Fast in/out positioners .....	7
2.8.3	Steppers, servos, and DC motors.....	7
2.8.4	Piezo actuators .....	7
<b>3.</b>	<b>Machine protection system .....</b>	<b>8</b>
3.1	Hard and soft stops .....	8
3.2	MPS example .....	8
<b>4.</b>	<b>Handling large volumes of data from machine diagnostics .....</b>	<b>9</b>
<b>5.</b>	<b>User interfaces .....</b>	<b>10</b>
5.1	Engineering interfaces.....	10
5.2	Overview displays.....	10
5.3	Task-based displays.....	11
5.4	PSS displays .....	11
5.5	Data archiving .....	11
5.6	Sequencer .....	11
<b>6.</b>	<b>The EPAC GUI .....</b>	<b>11</b>
<b>7.</b>	<b>Control system interface to the DAQ system .....</b>	<b>13</b>
7.1	The DAQ messaging system .....	13
7.2	The Pulse ID system .....	14
7.3	DAQ control interface .....	14
<b>8.</b>	<b>Security .....</b>	<b>15</b>
<b>9.</b>	<b>Deployment.....</b>	<b>16</b>
9.1	Provisioning Linux servers.....	16
9.2	Deploying IOCs to Linux servers .....	16
9.3	Provisioning Windows hosts and deploying IOCs .....	17
9.4	Future plans .....	17
	<b>Appendix A List of Acronyms .....</b>	<b>18</b>
	<b>Appendix B Team Structure .....</b>	<b>19</b>

# 1. Introduction to EPAC

The Extreme Photonics Applications Centre (EPAC) laser system will produce pulses with peak powers of around a petawatt at up to 10 Hz. This section presents a simplified version of the system to outline the main function of the control system in EPAC's operation.

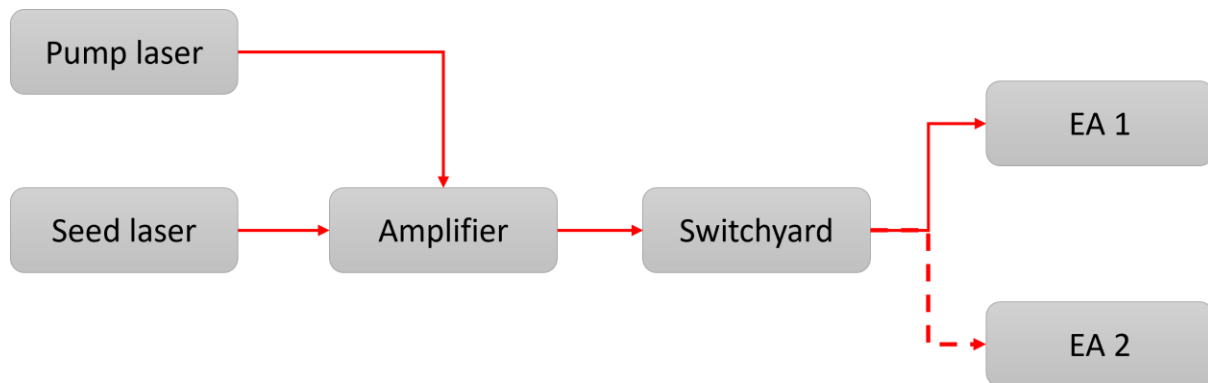


Figure 1: Simplified schematic of the EPAC subsystems

**Figure 1** shows a simplified schematic of the laser. The seed laser is the origin of each pulse delivered to the experimental areas (EAs). Here 30 fs laser pulses are generated at 10 Hz, before being 'stretched' (increasing the pulse duration) to 1 ns; they are then pre-amplified to 1 J. The pump laser uses cryogenically-cooled amplifiers to produce 5 ns, 150 J laser pulses at the same repetition rate as the seed.

The pump and the seed laser beams pass through separate interlock-controlled shutters to meet in the amplifier. For optimal operation of the amplifier, the pump laser beam needs to have its wavelength changed from the original infrared to green, which is done by propagating the beam through a non-linear optical element. After this wavelength conversion, the pump beam is directed towards a titanium-sapphire (Ti:Sa) crystal where its energy is stored. A short time (order 10s of ns) after this, the seed laser beam transits through the Ti:Sa, and the amplifier transfers the stored energy to it.

The full-energy beam goes to a switchyard to divert it to the correct EA; EPAC will only serve one EA at a time. Before entering the EA, the beam is recompressed back to a short pulse (30 fs). From here on, the laser is propagated in vacuum, via a focusing parabola, to the target that awaits in an interaction chamber. The interaction between the target and the sample is recorded by a range of detectors appropriate to the science under investigation. The primary detectors are typically area detectors.

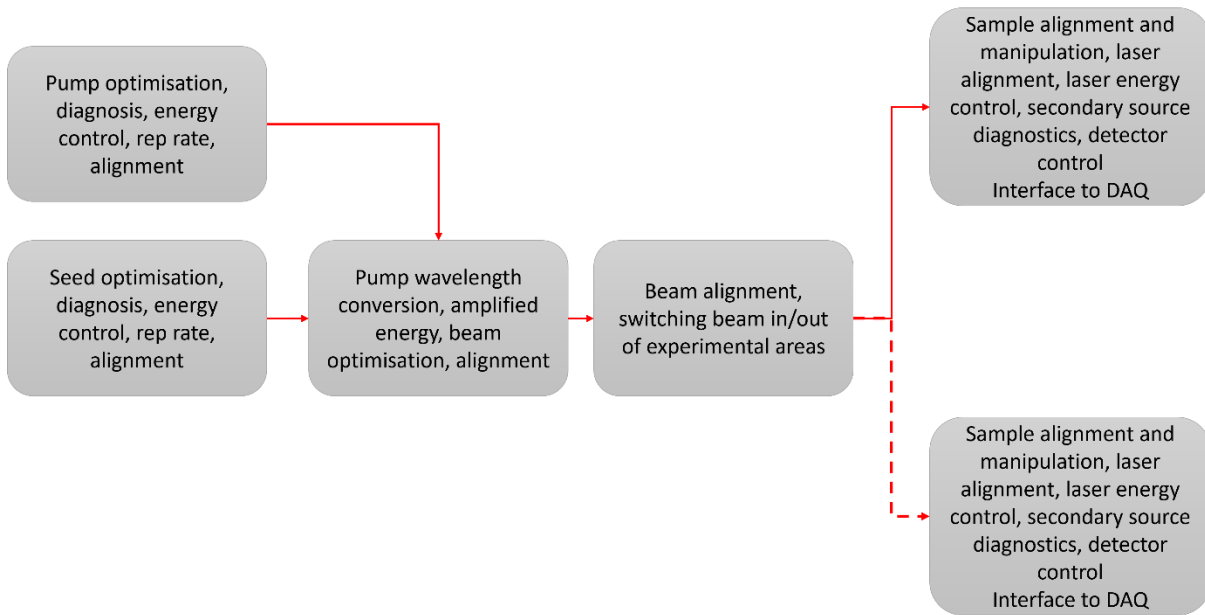


Figure 2: Primary responsibilities of the control system for in each subsystem shown in Figure 1

**Figure 2** shows the functions and responsibilities of the control system in each of the EPAC subsystems identified in Figure 1. Common to all areas is the alignment and manipulation of the laser pulses through actuated mirrors and motorised components. This is largely monitored in a non-blocking way with a suite of CCD cameras, which are always on and generating data. Beyond this, each area uses common systems as far as possible – for instance oscilloscopes or spectrometers – but all necessarily have specific custom devices that require software control.

The output of the laser system is configurable in several ways, such as repetition rate and pulse duration, and the output power can be changed by several orders of magnitude in step changes. There are different regimes – known as power modes – from which the end user can choose, with all but the highest power mode being used for the alignment of targets or detectors. Switching the power mode results in greatly different signal levels on many of the machine diagnostics, and the control system must compensate for this, for example by adjusting the gain on a detector or changing the degree of filtering in front of a camera.

## 2. Control system architecture

An overview of the control system architecture is shown in **Figure 3**. Systems are distributed over a network and communicate via the EPICS channel access protocol. Devices are attached to the network through input/output controllers (IOCs) that transfer commands and data.

The laser is inherently an event-driven system that runs at a maximum frequency of 10 Hz. Each laser pulse is associated with an electrical trigger from the timing system that is used to coordinate operations and synchronise data collection. More detail on the major components of the control system is provided below.

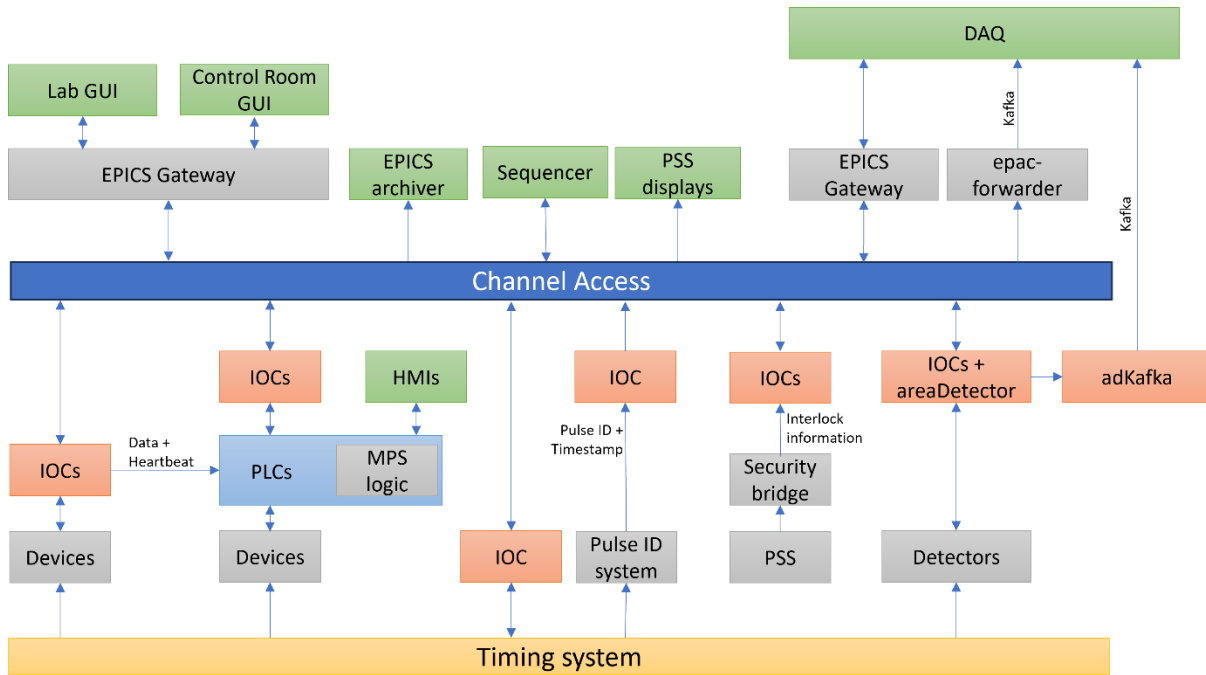


Figure 3: Overview of the EPAC control system architecture, showing the links to the data acquisition system (DAQ), the personnel safety system (PSS), and the machine protection system (MPS)

## 2.1 EPICS

EPICS (see <https://epics.anl.gov/>) is an open-source distributed control system framework used in many Big Physics experimental facilities around the world. It has been the control system toolkit of choice for lasers developed at the CLF for around a decade, due to its flexibility and ease of use, coupled with its use in facilities close to the CLF, both geographically (ISIS, Diamond) and in terms of a working relationship (Daresbury). Although the use of EPICS is not yet widespread in the laser community, it has also been adopted by ELI, one of the largest laser developments in the world.

The CLF’s EPICS systems run mostly on Ubuntu 22 servers, using EPICS v7. At present only the Channel Access message protocol is used, but pvAccess may be considered in the future if the need arises. There are a small number of devices that require Windows only software, and these use IOCs running on Windows 11 PCs.

In addition to IOCs from the EPICS toolkit, Phoebus will be used for some engineering screens, the archiver for machine data storage (except for images), the sequencer for automation, and the EPICS gateway for location- and role-based access control.

## 2.2 PLCs

Programmable logic controllers (PLCs) control many of the moving parts in EPAC as well as the vacuum systems. They have an s7nodave interface to EPICS for integration into the control system, but also human-machine interfaces (HMIs) for local control.

## 2.3 Pulse ID and timing system

A Greenfield Technology timing system provides a range of frequencies to drive the facility. The primary operating frequency of the EPAC facility will be 10 Hz, but it will also be able to run at 1 Hz and shot-on-demand.

The timing system synchronises all the laser sources and diagnostics. It is important to be able to associate the data from all detectors and diagnostics with each laser pulse. To do this, each pulse is given a unique ID and a timestamp, with this information distributed via Channel Access to all devices. The back up approach is to synchronise all devices with Network Time Protocol (NTP) and compare data timestamps with the primary timestamp given to the pulse by the Pulse ID system. More detail can be found in **Section 7.2**.

## 2.4 Detectors

EPAC has around two hundred cameras that collect data every laser pulse to provide information on alignment and beam quality. It also has other detectors used to diagnose the temporal shape of the pulse, the pulse energy, and other key metrics. These diagnostics may be areaDetectors, but may produce waveforms or scalar values. Almost all areaDetectors are GigE, resulting in large volumes of data on the network that need to be dealt with.

## 2.5 Personnel safety system and interface to control system

The personnel safety system (PSS) is a read-only input to the EPAC control system. The data from the PSS is used to relay information about the hazard status of individual laser areas and of the overall facility. This information is used to:

- display the laser hazards present in the laser areas.
- provide an overview of all the interlocked systems that are operational in the facility
- provide a 10-second warning when EPAC is about to change power mode. This warning is used by the control system to ensure that any necessary changes are made beforehand.

## 2.6 Machine protection system

The machine protection system (MPS) is responsible for preventing damage to equipment in the EPAC facility through non-standard operation. The risk may come from the laser beam itself. For example, damage in an optical component can imprint intensity modulations on the laser that can go on to damage more components downstream: these modulations must be diagnosed and the MPS must prevent further propagation through the laser path. Other MPS rules relate to engineering services. For example, if two vacuum chambers have a large difference in pressure, the MPS must prevent the gate valve between them from opening.

The MPS is described further in **Section 3**.

## 2.7 User interfaces

The different user interfaces within the EPAC control system are shown in **Figure 3** as green boxes. They are each described in **Section 5**.

## 2.8 Motion control

Motion control in EPAC is generally straightforward. There are no complex movements to calculate, and most motion falls into one of the categories described below.

### 2.8.1 *Slow in/out positioners*

These are used, for instance, to insert and remove alignment aids into the laser path. They are usually moved by pressurised air under the control of a PLC. The motion completes in around a second, and these components are only moved in and out when the laser beam is at low energy.

### 2.8.2 *Fast in/out positioners*

For some applications, the laser must be stopped or diverted when pulsed at high energy. This may be to divert the laser to a beam dump, or to select a single pulse for propagation to an experimental area. When a laser pulse arrives, the positioner must be either fully in or fully out: a partially obstructed, high-energy beam may cause damage to the optics in the laser system. A Linmot<sup>1</sup> linear slide and controller is used for this application, which can open or close in less than 100 ms. The slide is programmed to move in or out when it receives an electrical trigger, and these triggers are synchronised to the laser pulses. The control system has supervisory control.

### 2.8.3 *Steppers, servos, and DC motors*

A variety of stepper, servo, and DC motors are used in EPAC for applications that include sample and target positioning, fine adjustment of optical path length, and positioning a parabola to achieve optimal focus of the laser. ACS<sup>2</sup> motion controllers and the EPICS motor record support module are used for these motors.

### 2.8.4 *Piezo actuators*

Two types of piezo actuator are used for motion control: Newport picomotors<sup>3</sup> and piezoelectric stacks with a mirror mounted on them.

Newport picomotors are stick/slip linear actuators that perform slow movements to change the angle of mirrors and other components to align the beam. Movement can be made by operators, or in a software-controlled closed-loop that uses cameras to measure the beam position and moves the actuators to align it.

Piezoelectric stacks can make small-angle adjustments to the mirror mounted on them at kHz frequencies, and are used to counteract vibrations in the facility to keep the laser focal spot steady. The system uses a continuous pilot laser that samples vibrations through the optical path. Vibrations are measured using a position-sensitive detector and used as feedback to the closed-loop controller. The controller is a National Instruments cRIO<sup>4</sup> and the feedback loop is written on the FPGA of the device. The EPICS control system has supervisory control.

---

<sup>1</sup> <https://linmot.com/products/linear-guides-linear-modules/>

<sup>2</sup> <https://www.acsmotioncontrol.com/>

<sup>3</sup> <https://www.newport.com/f/picomotor-piezo-linear-actuators>

<sup>4</sup> <https://www.ni.com/en/shop/compactrio.html>

### 3. Machine protection system

The MPS will operate on a range of timescales, the shortest of which is the time between two consecutive laser pulses (100 ms). To meet this requirement the logic for the MPS is defined as far as possible in PLCs. The action to make the system safe in the event of a breach is ideally an output directly from a PLC to a piece of hardware, for example to close a shutter.

Optimally, the inputs to the PLCs are also hardware signals, but in some cases this is not possible, and the only solution is derived from software in an IOC. In such a case, the IOC will connect directly to the PLC via modbus and provide two pieces of information: the value of the safety condition, and a heartbeat that oscillates high/low at 1 Hz. The PLC can use this heartbeat to determine if the IOC is up and responsive, and act accordingly.

#### 3.1 Hard and soft stops

Two transitions to safe states - a hard stop and a soft stop - have been defined for each of the subsystems of EPAC (as shown in **Figure 1**) which the MPS can enact in the event of a breach. These are known as hard stop and soft stop.

A soft stop is a safe state that is quick to recover from, and typically involves inserting a beam shutter or changing the timing system to minimise the output energy. A hard stop is more intrusive – such as cutting the power to a laser system – and more time is required to recover from such actions.

For each MPS rule, the possible damage to the system is evaluated, and the appropriate safe state is chosen.

#### 3.2 MPS example

**Figure 4** shows an example of how the MPS will work. At the bottom of the image, the EPAC laser is propagating from right to left. It is focused by a lens into a vacuum tube, then recollimated on its exit from the chamber by a second lens, before being directed by a mirror to the rest of the system. A small portion of the light leaks through the mirror and is used to check the beam parameters are in specification. The leakage beam may be too intense for the diagnostic camera and so a filter wheel – a cassette that holds different levels of attenuation at a series of indexed positions – is placed between the camera and the beam to reduce the intensity to acceptable levels. The intensity of the laser beam changes with the EPAC power mode. Changing and monitoring the power mode is carried out by the PSS.

This system presents two risks that need to be handled by the MPS. Firstly, the filter wheel may be set to the wrong index, which could result in a signal intensity that is too high for the detector and could cause damage. Secondly, the laser beam must only be focused when the vacuum level is below a given threshold. If the pressure is too high, the laser beam will cause breakdown in the gas, leading to loss of energy and a reduction in beam quality.






Figure 4: The vacuum system inputs to the MPS are direct hardware inputs to the PLC.

Some components, such as filter wheels, are only readable through software. In this example, the software sends the data and a heartbeat signal as proof of correct operation.

The index of the filter wheel can only be monitored through software. The IOC takes the EPAC power mode as an input and selects the correct filter to provide the required level of attenuation. It confirms that the filter wheel is then at the correct index, and if not reports an error. The IOC provides a Boolean value indicating the status of this error to the PLC, along with a heartbeat that oscillates at 1 Hz. If the PLC input shows an error, it sends an output to close the shutter to prevent damage to the camera. The PLC also monitors that the IOC is up and responsive through its heartbeat signal. If the heartbeat signal is absent for a period greater than the timeout period, then the PLC will close the shutter. These two cases are shown as flow diagrams (a) and (b), respectively, in the PLC inset.

The vacuum level is monitored directly by a PLC, which stores the threshold value internally. If the threshold is breached, the PLC can take action to put a shutter into the beam to prevent propagation and risk to the machine. This is shown in flow diagram (c) in the PLC inset.

## 4. Handling large volumes of data from machine diagnostics

In laser systems, most diagnostics for monitoring and optimising beam performance are cameras. The mirrors in the laser beam path reflect around 99.9% of the incident laser light, transmitting the rest. The transmitted portion can be directly imaged onto a camera to diagnose beam quality, alignment, and energy. An analyser of some kind positioned before the camera can also be used to check for other metrics, such as pulse duration, optics damage, and spectral phase.

There are around 200 diagnostic cameras around the EPAC facility. They all observe the small, transmitted portions of the laser and do not need to obstruct useable beam to record data. This means that they all run simultaneously and produce data with every laser pulse.

Collecting data from the cameras and displaying it to operators is a crucial design consideration for the EPAC network infrastructure. The most common diagnostic cameras have 2k x 2k, 16-bit pixels. Running at 10 Hz, these cameras each stream data at rate of 640 Mb/s.

Servers with a 40 Gb fibre connection are used to manage the cameras, with one server for every 25 cameras, giving peak data rates of ~19 Gb/s. Testing has shown that the CPU load on the server is the limiting factor to increasing the number of cameras per server, not the network bandwidth. In the current configuration, both are within acceptable limits.

Streaming this data to client computers poses a further difficulty. EPAC's operator interfaces run on Windows 11 PCs with 10 Gb network interface cards. A typical screen will show a maximum of 8 - 10 cameras at once, which is manageable; however, the option to stream any camera in EPAC is required. Two options are being considered to address this:

1. Produce a second, low-resolution, image in the IOC that is used for screens that view, say, more than two cameras at once.
2. Within the operator application, disconnect from any camera image that is not currently being viewed.

Option 1 places extra demand on other parts of the network, as two images from each camera would be streamed rather than one, so Option 2 is preferred. The GUI framework being used for operator screens (described in **Section 6**) has the capability to know which images are being displayed and can disconnect the other data streams, but this option has yet to be implemented and tested.

## 5. User interfaces

This section describes the interfaces with two kinds of user: humans who interact with the control system to achieve some task, and other software components that read from, and/or write to, the control system. The interface to the DAQ system is omitted, as it is covered in its own section (**Section 7**).

### 5.1 Engineering interfaces

These graphical user interfaces are designed to enable expert users to control all the functionality of a given device or set of devices. These interfaces will be used, for example, to set calibrations, optimise performance of critical components, or set safe operating limits. They may be provided in software on a computer screen, or through an HMI depending on the device under control.

### 5.2 Overview displays

Overview displays will contain sufficient information from each subsystem to let operators know that EPAC is operating correctly, and provide information about the overall system status and performance. They will give quick and easy feedback on any problems, and indicate which subsystem is at fault.

### 5.3 Task-based displays

Task-based displays will aggregate the necessary data and control for operators to achieve routine tasks without needing to open multiple windows. An example would be optimisation of the amplifier output energy, which requires adjusting the temporal overlap of the pump and seed laser pulses through manipulation of a delay generator to maximise the energy. An aggregated screen might show the delay generator delay setting, the current output energy of the amplifier, and oscilloscope traces of the pump and seed pulse for a quick visualisation of their relative timing for rough adjustment.

### 5.4 PSS displays

Information from the PSS is interpreted by the control system and used to display safety information to operators. For example, at the entrance to each laser area the display shows which lasers are on, from which the users know what protective eyewear to wear to enter safely.

### 5.5 Data archiving

Machine data in the form of scalar values, string, and waveforms will be stored in the EPICS archiver appliance. The stored parameters, sample frequency, and retention duration will be set by the EPAC data policy and is out of the scope of this document. Image data will be handled by the DAQ system.

### 5.6 Sequencer

Some of the procedures needed to setup and configure EPAC are repetitive and can be modelled as simple state machines. These are good candidates for automation, where the control system can verify pre-conditions, arrange a series of actions, and check for successful completion without further operator intervention. The State Notation Language (SNL) from the EPICS toolkit will be used to automate these procedures. To begin with, the number of SNL programs will be small, but as the project moves through the commissioning phase and more is learnt about operating the facility, this list is expected to grow.

## 6. The EPAC GUI

The EPAC GUI is a .NET application, written in C# and using the Blazor framework. It can be deployed either as a web-based or a desktop application: currently it is deployed to the desktop. The communication layer is a .NET wrapper around the standard Channel Access library from the EPICS toolkit. The GUI is not yet fully developed, but working prototypes are available in the pump and front-end sections of the facility.

The GUI design allows for a variety of features, such as:

- The ability to toggle between 'normal' and 'goggle' mode, the latter using a colour palette that is more clearly visible through laser-protection eyewear
- Login authentication based on Azure, which can link to the CLF's centrally managed fed-id system.
- Widgets that can be dragged or hidden, to allow users to customise the screen to their needs and save their preferences.

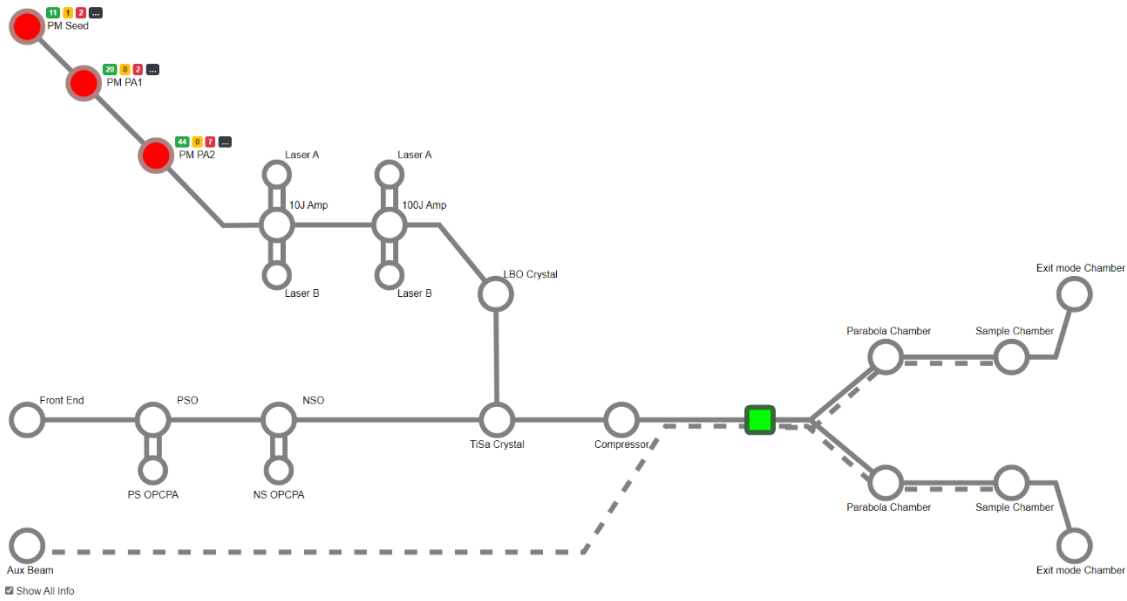


Figure 5: The overview screen for EPAC

An example of an overview display is shown in **Figure 5**. This is designed to be viewed from the laser control room. Each subsystem is represented as node, arranged according to its position in the laser path. After commissioning, a quick view of the subsystem status will be given by colour coding a node – red for error, yellow for warning, and green for OK – with the small boxes to the side of a node showing the number of components in that subsystem in each of the three states. Clicking on the node will bring up a new window with an explanation of any problems.

Figure 6: An example summary screen of the pump laser subsystem. Panels can be moved and/or minimised according to user preferences.

The GUI shown in **Figure 6** provides an intermediate level of detail and is designed for daily use by operators to perform routine adjustments. The screen provides access to control and visualisation of all device parameters that would be expected to change at least daily.

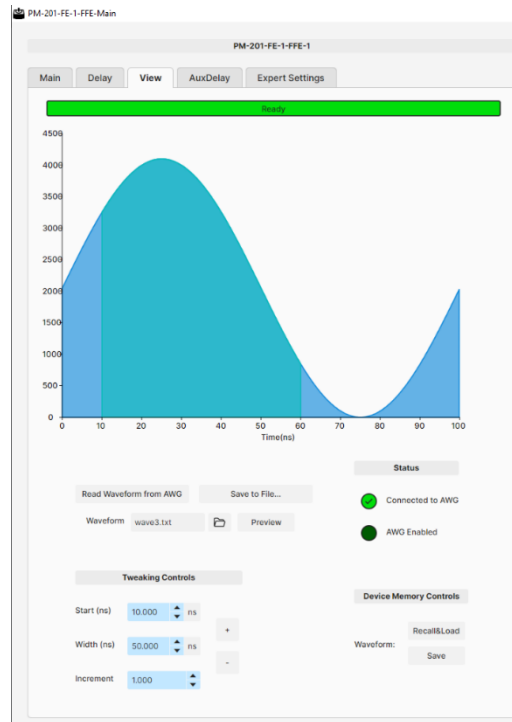


Figure 7: Example of an engineering interface will access to all device parameters

Engineering interfaces are designed for expert use and expose all the available commands and settings of a device. An example of this is shown in **Figure 7**.

## 7. Control system interface to the DAQ system

The EPAC DAQ system is responsible for deciding what data and metadata are stored; transporting those data from their collection point to visualisation, analysis, and storage platforms; and controlling some parameters of the machine in response to experimental requests that come from the end user. An example of the last of these would be to rotate a stage from  $0^\circ$  to  $180^\circ$  in steps of  $1^\circ$  and to capture an image at each angle. The control system provides the mechanism to move the stage (writing to a process variable (PV) over Channel Access) and the DAQ makes the request.

The control system has the following responsibilities to the DAQ:

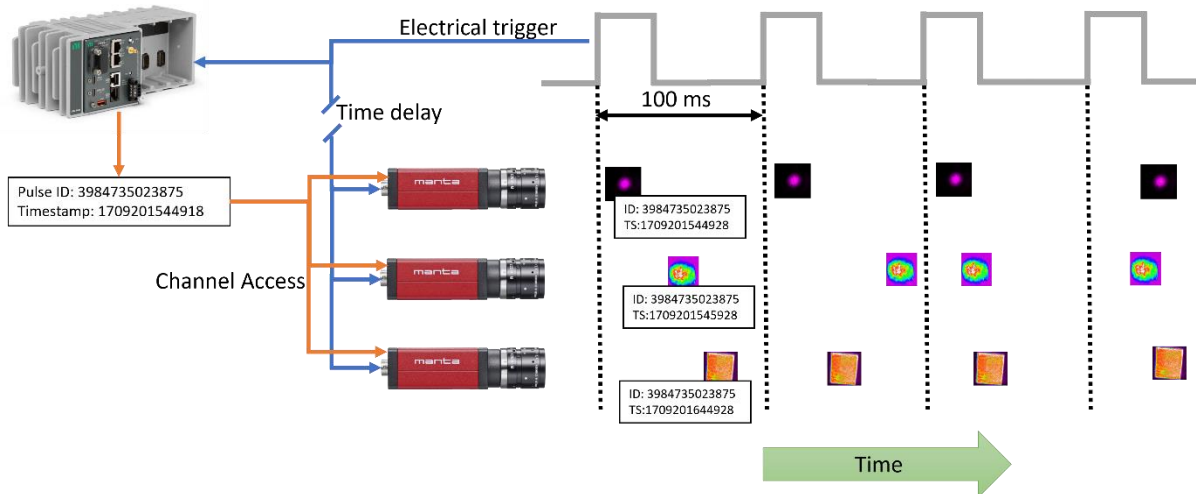
- Acquire data from detectors and pass them to the DAQ system
- Provide a way to associate data from all detectors with the laser pulse that generated them (the Pulse ID system)
- Provide an interface through which the DAQ system can instruct the machine to take data and change the operating parameters.

### 7.1 The DAQ messaging system

The DAQ system uses the Apache Kafka messaging system to move data. The data collected by IOCs is forwarded to Kafka through one of two methods. In the case of areaDetectors, the adKafka plugin is used. For all other types of data, the epac-forwarder reads the value over Channel Access and passes it to Kafka. The epac-forwarder is part of the DAQ system not the control system, and so is not discussed here.

## 7.2 The Pulse ID system

The Pulse ID system will be a National Instruments cRIO with digital input/output (IO) modules. The concept is shown in **Figure 8**.



*Figure 8: A schematic of the Pulse ID system, showing the cRIO in the top left corner. The cRIO digitises the trigger signals, generates a timestamp and ID, and sends them to device IOCs in advance of data acquisition.*

For every laser pulse, there will be a coincident electrical trigger pulse from the EPAC timing system in the form of a 5 V transistor-transistor logic (TTL) signal. These triggers will be digitised by the cRIO, which in response will:

- Generate a timestamp when the trigger was received
- Increase the value of the EPAC Pulse ID by 1
- Distribute the new value of the Pulse ID over Channel Access.

The trigger signal to the Pulse ID system will be 'early': it will arrive before the associated laser pulse has been generated. With a 100 ms interval between laser pulses (with many parts of the laser operating at 10 Hz), it should be possible for IDs to be sent far enough in advance for them reach all IOCs before the associated data acquisition has taken place. The Pulse ID can then be associated with the data at acquisition time, before forwarding to Kafka. This approach is shown for the first set of images in **Figure 8**, which all have the same pulse ID despite being recorded at different times.

If this system fails and the IOCs do not receive the new Pulse ID in good time, the second option is to associate data by comparing timestamps. Each of the computers running IOCs in the EPAC control network will be synchronised to an NTP server. The timestamps of each piece of data can then be compared after acquisition and placed into 100 ms buckets that represent each pulse. This data association can be carried out at any time after collection, but the disadvantage here is that Pulse IDs are not assigned before entering the DAQ system.

## 7.3 DAQ control interface

The final responsibility of the control system is to provide the DAQ system with access to the laser and experimental controls. This will be provided through Channel Access to PVs through an EPICS gateway, which will ensure that only parameters agreed by the machine operators can be controlled by the DAQ.



## 8. Security

The networking infrastructure in EPAC has layers of security to reduce the risk of uncontrolled access to the system:

- VLANs segment the network into functional areas, preventing control of systems from unauthorised areas of the building and reducing the scope for potential attacks
- Access control ensures that only authorised users can access specific resources: switchport port-security will be implemented to limit switchport access by restricting the number of MAC addresses allowed, preventing unauthorised devices from connecting
- The EPAC network sits behind the STFC RAL site firewall, which should prevent malicious packets from outside EPAC from reaching their intended destinations.

VLANs divide the EPAC network into several subnets. Referring to **Figure 1**, VLANs are used to place all the devices in the pump laser on their own subnet. The same is true for devices in the seed laser, the amplifier, and each of the experimental areas. In conjunction with access control, and appropriate use of the EPICS gateway, this will prevent users from accidentally changing hazards in areas that they do not own. For example, operators in the pump laser should not have control over any devices in the seed laser, and vice versa, whereas operators in the amplifier may need some control over both laser subsystems, as they act as inputs to that subsystem. The available control parameters, and the degree of control, will be managed through VLANs and a gateway.

In addition to this, segmenting the network has the advantage of isolating separate subsystems, reducing the risk of unauthorised access or lateral movement by attackers. It also brings the more straightforward benefit of knowing where a device is located, by identifying the subnet on which it sits.

To supplement VLANs, access control lists on switches will filter traffic based on predefined rules. These rules will be configured to permit or deny traffic based on source or destination IP addresses, protocols, and ports; for example, to only allow access to the EPAC control network from a fixed list of IP addresses. This action will further block malicious traffic and prevent unauthorised access to sensitive resources.

The outer layer of security is the centrally-managed STFC RAL site firewall. This is the first line of defence against external malicious actors.

There are two key parts of the EPAC network that require special attention. Firstly, the PSS network is almost completely physically separated from the rest of the infrastructure, the single exception being read-only access for a named server on the control network through a Pilz SecurityBridge<sup>5</sup>. This is an industrial firewall developed in accordance with IEC 62443-4-1.

Secondly, EPAC will be used to perform experiments that require enhanced security clearance. In such cases, sensitive data will be contained on an entirely physically separated network, and will be stored on servers in a room with strict access control and surrounded by a Faraday cage.

---

<sup>5</sup> <https://www.pilz.com/en-GB/products/networks/fieldbus-and-ethernet-systems/securitybridge>

## 9. Deployment

The CLF uses Ansible<sup>6</sup> to both provision the Linux servers with EPICS base and support modules, and deploy IOCs, and this approach is also being taken in EPAC. Ansible is installed and run on a control node and can perform actions on multiple remote hosts by providing lists of their IP addresses in a configuration file. IOCs are run using procServ<sup>7</sup> and installed as services on the Linux servers.

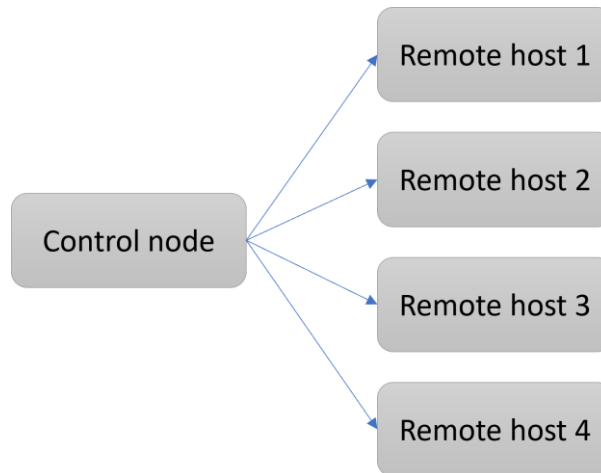


Figure 9: Schematic of Ansible deployment from the control node to servers

### 9.1 Provisioning Linux servers

Ansible playbooks allow the user to specify the version of EPICS base and each support module to be installed. For each of the specified hosts, the playbook then downloads the necessary sources, builds them on the host, and sets the environment variables to the site configuration.

### 9.2 Deploying IOCs to Linux servers

Each IOC has its own GitHub repository and a playbook for deployment. The playbooks are stored in a separate repository. Each playbook is a YAML file that specifies a user-friendly name for the IOC, the repository to pull from, and the commit to deploy, and supplies any macros that are needed by the IOC at run time, such as device IP addresses or serial ports. An excerpt from a playbook is shown in **Figure 10**.

Once run, the playbook:

- Creates a directory on the host for the IOC and sets permissions correctly.
- Clones the IOC to the control node, makes any modifications specified in the playbook, then zips it and copies it to the remote host
- Unzips and makes the IOC; if a version of the IOC already exists, it is stopped and deleted first
- Generates a service file, copies it to the correct folder, and starts the service
- Updates an IOC information list on the host that contains the installed IOCs.

<sup>6</sup> <https://www.ansible.com/>

<sup>7</sup> <https://github.com/ralphlange/procServ/>



```
---
- name: [redacted] Description
  hosts: [redacted] Server to deploy IOC on
  vars_files:
    - config.yml

  roles:
    - role: ioc
      repository_name: [redacted] Name of IOC repo
      ioc_id: [redacted] Name for this instance
      sha: [redacted] Branch name to use
      vars:
        params:
          SHORT_NM: PM-201-FE-1-PA-1
          ID: 10
          PORT: PA1
          SERIAL_PORT: ttyUSB0
    } [redacted] Extra
    } [redacted] params
    } [redacted] (varies by
    } [redacted] device)
```

Figure 10: A section of an Ansible playbook to deploy and IOC

### 9.3 Provisioning Windows hosts and deploying IOCs

There are some IOCs that require a Windows host. At present, EPICS and support modules are installed manually, and IOCs are deployed manually. The use of Ansible to automate this is being investigated.

### 9.4 Future plans

The existing deployment infrastructure works reliably, but could be improved. It is expected that there will be a move towards containerisation of IOCs by the end of the EPAC project, for example by using [epics-containers](https://github.com/epics-containers)<sup>8</sup> or a similar scheme, and this is being investigated now.

---

<sup>8</sup> <https://github.com/epics-containers>

## Appendix A List of Acronyms

CLF	- Central Laser Facility
DAQ	- Data acquisition
EA	- Experimental area
ELI	- Extreme Light Infrastructure
EPAC	- Extreme photonics application centre
EPICS	- Experimental physics and industrial control system
ESS	- European Spallation Source
FGPA	- Field Programmable Gate Array
GigE	- Gigabit ethernet
HMI	- Human machine interface
IOC	- Input-output controller
MPS	- Machine protection system
PLC	- Programmable logic controller
PSS	- Personnel safety system
PV	- Process Variable
RAL	- Rutherford Appleton Laboratory
SNL	- State Notation Language
STFC	- Science & Technology Facilities Council
Ti:Sa	- Titanium sapphire
TTL	- Transistor-transistor logic
VLAN	- Virtual Local Area Network

## Appendix B Team Structure

The CLF software control team has eight members. There are three senior engineers, three mid-level engineers, and two junior engineers. EPAC also has a dedicated network engineer, supported by the digital infrastructure team for the whole organisation. The CLF software team works in partnership with a team of software developers at TIFR, Hyderabad, India. They are funded by an international collaboration grant, and 100% of their work time is allocated to the EPAC project. This collaboration has been going for over six years now, with regular visits of members of the India team to the CLF for training and to help with commissioning software.