EPAC Control System Design Review Committee Report

1. Introduction

A Control System Design Review of the Extreme Photonics Applications Center (EPAC) was conducted on 24-25 April 2024 by an independent committee with a charge provided by Prof. Rajeev Pattathil, head of Novel Accelerator Science and Applications. The controls team presented the plans for the EPAC control system and the committee was asked to provide feedback and recommendations based on the presented material.

Committee members:

Timo Korhonen, European Spallation Source ERIC (*Chair*) <timo.korhonen@ess.eu>
Ralph Lange, ITER Organization <ralph.lange@gmx.de>
Graham Cox, STFC UKRI <graham.cox@stfc.ac.uk>
Tyler Johnson, SLAC National Accelerator Laboratory <tjohnson@slac.stanford.edu>

The committee members have agreed to be contacted with specific questions.

2. Charge Questions to the Committee

In general, the committee is happy with the status of the control system development and congratulates the project team for the excellent work done, including the comprehensive presentations. The committee did not have any major concerns regarding the feasibility and robustness of the plan. However, there are some recommendations for the team to consider, which in the opinion of the committee would support the project by increasing the control system potential and making it more future-proof and robust in production.

1 - Overall control system architecture

Please comment on the proposed architecture/development processes and potential challenges in integration.

In general, the design follows a typical model for EPICS-based systems. The committee endorses the proposed approach and thinks it is on a good track.

2 - Appropriateness of the control system for enabling smooth user operations

Is it well thought-out with user operations in mind?

Regarding user operations, it seems that there is still a need for clarification of details of the different operation workflows. In general, the plan seems to be adequate.

How functional is the automation system?

The proposed system can, as far as the committee can judge, support automation of any required tasks with moderate effort.

Is it flexible enough to adapt to changing user requirements? How easy is it to make changes?

While unexpected changes will always be a strain on the controls team, the committee considers the setup flexible enough to accommodate changes as they appear. Adapting to changing requirements is as easy as can be reasonably expected.

3 - Machine Protection System

Please comment on the state of the Machine Protection System and its suitability for EPAC operations.

The committee is pleased to see the effort of implementing the machine protection logic in the lowest appropriate layer (PLC) and making it independent of the control system. This is universally accepted and the recommended practice for safety and protection systems.

In general, functional separation is a good principle to follow and we encourage the team to go even further, separating the MPS and vacuum controls that in the present state are running on the same platform.

4 - Graphical User Interfaces

Please comment on the user-friendliness of the GUIs. Does developing our own GUIs expose us to more risks?

Given the history of CLF and its experience in producing high-power lasers, the committee is convinced that the user-friendliness of the presented GUI solution is appropriate and meets UX requirements.

Are they resilient to future versions of EPICS?

EPICS Channel Access (CA) and pvAccess (PVA) interfaces and standard data types are quite stable. Typically, client software can be expected to work with new versions. This, however, may not apply when using custom structures, as the EPICS core modules cannot guarantee backward compatibility in all cases.

With CA being slowly phased out and EPICS users shifting to PVA, any tool that only supports CA is not fully future-proof.

5 - Interface with DAQ

Is the plan for data acquisition appropriate? Can it cater to large data volumes? Please comment on our plans for the PulseID system.

The plan for data acquisition seems to be sufficient for the initial planned 10 Hz repetition rate.

The proposed Channel-Access-based PulseID system should be able to generally keep up with this rate. However, one cannot rule out occasional drops and difficulties in aligning data to particular laser pulses. Depending on the severity of the consequences of such occasional drops, it might be advised to consider countermeasures to make the system more robust. (See the more detailed section below.)

The committee recommends further clarification of this matter and following up with appropriate measures.

6 - EPICS Gateway

Is our use of the EPICS gateway appropriate? Would it pose a limitation to bandwidth?

The CA Gateway does not handle a mixture of scalar and waveform/image data well. Any big data item (image) blocks other data during the time it is handled by the Gateway. This effect can be mitigated with a strategy that properly separates scalar and waveform PV traffic to separate Gateway instances, but this is a non-trivial task and requires careful planning.

On the other hand, the new generation pvAccess Gateway (part of p4p [3]) handles mixed loads of scalar and array data items well and generally performs much better than the CA Gateway. If described performance issues are encountered, it is worthwhile to consider that as an option.

In general, gateways can pose a limitation to bandwidth and should only be used to

- Limit access to a different subnet as being read-only, possibly allowing a limited set of writeable channels.
- Protect IOCs from being overloaded by too many CA/PVA subscriptions.

7 - Future Proofing

If we are to implement feedback loops, can the current DAQ-control system interface support that?

The current control system architecture appears to be able to support the level of automation that was described. The use of the EPICS State Notation Language was discussed for simpler routines such as beam switching, and the use of the Bluesky/Ophyd python framework was discussed for running more complicated experiment automation such as tomography scans. Beam pointing and centring, another use case for automation that was mentioned, could be accomplished by either method.

However, it is important to note that the Bluesky/Ophyd Python framework does not currently support pvAccess.

8 - Delivery

Please comment on the size and expertise of the existing team to deliver the EPAC control system.

The committee has no special concerns about the expertise of the team. However, the team size is somewhat of a concern and the committee recommends adding resources whenever possible.

During the construction period, contractors can help to bring in know-how and resources as needed, but for steady-state operation, it is important to retain the necessary skills in-house. Only this will guarantee the control system's long-term maintainability and robustness against unavailability of key resources.

9 - Anything else the panel thinks we are missing

The committee did not notice any particular omissions in the plan. All important areas seem to be considered and covered as far as possible.

4. Detailed Findings, Comments and Recommendations

Presentations

Introduction

The committee appreciated the introductory presentations given by Prof. John Collier, Cristina Hernandez-Gomez and Rajeev Pattathil, which provided a comprehensive overview of the project and the science it will enable. This introduction prepared the ground and set the context for the following technical talks.

The EPAC Control System – Overview

Chris Gregory presented the primary responsibilities and an architectural overview of the EPAC control system.

Network Architecture, Security and Handling of High-Volume Data Traffic

Michael Chikwanda presented a Network overview, bandwidth requirements as well as key elements and aspects of the design.

Machine Protection System (MPS)

Rafael Sarasola presented the requirements, technology, architecture and development workflows of the MPS.

Development Process and Deployment Overview

Becky Harding presented the controls team, tools and principles of the development process, and the deployment workflow that has been developed.

Control System GUI

Aoun Muhammad presented the requirements, architecture, design and tools used for the EPAC GUI application.

Data Management Overview

Stephen Dann presented the scope, overview, architecture and framework used for the measurement data management.

Bluesky – Experiment Orchestrator

Ulrik Pedersen presented Bluesky and the concept of its application for EPAC Laser-synchronised Bluesky tomography scans.

Findings

The EPAC Control System – Overview

1. The presented design follows a standard model for EPICS-based systems.
2. Laser systems and experimental systems are largely independent and separated, with separate control rooms.
3. Setting up an experiment at EPAC is a time-consuming multi-stage process that requires two-way communication between laser and experiment.
4. User operation is foreseen for 40 hours weekly (one daily shift, Mon-Fri).
5. EPAC repetition rate will be 10Hz initially. (Ideas to go to 100Hz in an upgrade.)
6. The timing system is implemented in hardware (i.e., no software interfaces at the receiving end.) The pulse ID is distributed from one IOC over Channel Access to all other IOCs (using TSEL).
7. Mainly to support alignment, a total of ~200 cameras will have to be integrated.
8. Six network domains (seed laser, pump laser, laser control room, experimental area, experimental control room, DAQ) are interconnected by CA Gateways.

9. Four energy modes are switched/controlled by the PSS (as they have occupational safety consequences).
10. The switchyard features a moveable mirror setup that allows beam delivery to one of the experimental areas. Switching between target areas is slow.
11. The controls team consists of 10 colleagues in the UK (6 exclusively for EPAC), one in Poland and seven in India. The average experience is 4 years (median 2 years).
12. The by-subsystem schedule foresees the three laser-related subsystems (seed laser, pump laser, amplifier) to be integrated in 2024. Experimental areas (requirements not complete yet) are scheduled to follow by mid-2025.
13. The infrastructure for deployment is working.
14. The Drive System infrastructure is done for the low-level part and still needs the user-facing GUI (scheduled by the end of 2024).
15. The Machine Protection System requirements are not complete yet, the design principles have been proven on a different system.

Network Architecture, Security and Handling of High-Volume Data Traffic

16. The network design is guided by scalability, efficiency, reliability, security and simplicity.
17. The network will use hub-and-spoke topology with two hubs on separate floors for redundancy.
18. The network will use single-mode fibre for consistent high speed.
19. Timing, Interlocks and CCTV networks will use physically separate infrastructure. Corporate services, Profinet and Science networks will use shared infrastructure.
20. The chosen camera type (AVT Mako G234) uses ~400 Mbps bandwidth on its 1 Gbps interface.
    Integrating 25 cameras, the IOC server requires ~10 Gbps camera data bandwidth (10 Gbps interface).
    A section with 40 cameras requires ~16 Gbps bandwidth (25 Gbps interface).
    All 200 cameras require ~80 Gbps bandwidth (100 Gbps interface).
21. The core layer switches (Aruba 8325) provide 48x 1/10/25 GbE to connect to access layer switches and 8x 40/100 GbE for the backbone.
    The server farm switches (Aruba 6300) provide 24x 1/2.5/5/10 GbE and 4x 1/10/25/50 GbE uplinks.
    The access layer switches (Aruba 6100) provide 48x 1 GbE and up to 740W PoE (for cameras), with 4x 10 GbE uplinks.
22. Aruba's Virtual Switching Extension (VSX) technology will be used to ensure high availability and reliability.
23. Logical segmentation will be achieved through VLAN configurations.
24. Network monitoring will be based on Grafana.

Machine Protection System (MPS)

25. The Machine Protection System concept is being migrated from a software-based system running on a server to a PLC-based system that is independent of the control system for its operation.
26. The MPS PLC (Siemens S7-1513) will communicate (unidirectionally) with EPICS to display the status, but the logic will be independent of control system actions.
27. The EPICS-PLC communication is currently based on S7nodave, with plans to move to Modbus/TCP.

28. The MPS PLC also hosts vacuum system functionality at the moment.
29. MPS PLC programming uses a modular approach, automated testing and a version control system.
30. Notably the beam switchyard, and possibly some other components are still to be added to the protection system.

Development Process and Deployment Overview

31. The development process follows Agile methodology, with two teams (seed laser & pump laser) that both contain UK and Indian staff, doing synchronised two-week sprints.
32. GitHub Project boards are used to manage and track the development process.
33. Deployment is using Ansible to move sources onto remote hosts (production servers running IOCs), then compile locally on those hosts.
34. EPICS Base and support modules use a separate playbook each.
35. IOCs use a separate playbook and repository for each IOC.
36. There are plans to deploy IOCs on Windows and to evaluate container use.

Control System GUI

37. The EPAC GUI needs to answer specific challenges (accessibility, e.g., when wearing laser goggles) and provides advanced graphical features (dashboards, themes, etc.) as well as user authentication/authorisation and hidden/custom scripting. The framework ensures testability and maintainability.
38. Development is done using the .NET framework on Windows, using the C# language and the Blazor framework (to create modern interactive UIs using HTML, CSS, JS and C#).
39. EPICS communication is implemented as a C# wrapper around the original Com.dll and ca.dll libraries from EPICS Base. (The native C# library EpicsSharp is not used.)
40. The widgets follow the model-view-viewmodel architecture pattern to provide a clean separation between presentation, presentation logic and business logic.
41. The same set of libraries is used for desktop applications and web applications.
42. Authentication/authorization is based on Azure Active Directory (AD).
43. Deployment works currently (binary) through GitHub Release. An installer is planned, which would provide automatic updates.

Data Management Overview

44. The estimated data volumes are 1-5 GB per second and a few PB per year.
45. Data is collected from both “facility diagnostics” and “experimental detectors”.
46. Archived data will be provided to both facility staff and external users.
47. Management of archived data is considered out-of-scope for controls.
48. Apache Kafka acts as a “data broker”: EPICS data is sent to Kafka, and retrieved from Kafka to be used for analysis, plotting and saving to files.
49. Each Kafka message is a single measurement, including metadata. Each topic contains data from one source.
50. Camera integration uses the areaDetector plugin that converts NDArrays to Kafka messages.
51. Other EPICS data uses the EPICS-Kafka forwarder from ESS as well as a waveform variant developed in-house.
52. “Event building”: mapping archive data and metadata to pulse IDs (siloing) using timestamps and a single PV that has the nominal time for each pulseID.

53. The pulseID logic in cRIO technology counts the TTL trigger pulses from the timing system and maps an NTP-synchronised clock timestamp.

Bluesky – Experiment Orchestrator

54. It is proposed to use Bluesky to orchestrate the experiments at EPAC, including the interaction/communication with the lasers.
55. To investigate and demonstrate the EPAC use case "laser-synchronised tomography scan", a two-step simulation/emulation environment was set up, including hardware "laser" (LED), camera, sample stage & motor controller and trigger box as well as software containers for EPICS IOCs, Bluesky and a database (not Kafka).
56. While many parts of Bluesky/Ophyd can be used as-is, further development is necessary in a few areas:
    - Kafka data broker integration
    - pvAccess integration
57. There have been difficulties running the demo simulation/emulation at a 10 Hz rate, caused by the preliminary hardware and software stack used in the setup.

Comments

The EPAC Control System – Overview

1. Five years after its release, the pvAccess protocol (PVA) is mature.[1] The committee suggests any new project use PVA where possible and rely on Channel Access mainly as a fallback.
2. The CA Gateway application has serious shortcomings, especially when large arrays (images) are forwarded. CA Gateway instances need considerable effort to set up and maintain. The committee would suggest an architecture that tries to minimise the number of CA Gateways needed. Control rooms should be directly connected to the systems they control.
3. PVA's standard image container (NTNDArray [2]) is especially well-suited for transporting images across the network. The 2nd generation PVA Gateway [3] does not show the disadvantages of the CA Gateway mentioned above.
4. The high number of cameras will be a main integration challenge, as the requirements with respect to network bandwidth, image processing CPU power and archiving are demanding.
5. The high number of cameras will be producing data streams of serious bandwidth. The committee would suggest carefully analysing which image data is needed in the control room at which operation step to avoid unnecessarily sending high bandwidth streams.
   A similar analysis needs to be conducted with respect to archiving of camera data.
6. Decimated image streams (lower resolution, lower update rate) for general control room use should be prepared in the camera integration IOCs and preferably used in control room GUIs.

Network Architecture, Security and Handling of High-Volume Data Traffic

7. The committee appreciates the robust and appropriate design. The chosen topology and technologies seem adequate, the hardware providing the required bandwidth even for worst-case scenarios.
8. The shown design seems to suggest the physical and logical separation of the connections between cameras and camera IOCs from the IOCs "uplink" (EPICS) network. The committee endorses such a separation.

Machine Protection System (MPS)

9. The development is progressing in the right direction; separating protection functions from controls is the recommended practice. This is endorsed and continuation work in this direction is important.
10. The approach to PLC programming seems good and appropriate. Notably version control is important and traditionally rarely used in this field.
11. Sharing the same PLC for MPS and vacuum creates unnecessary dependencies that can compromise both functions and also make development, testing and verification more complex.
12. The committee suggests evaluating OPC UA as the communication protocol for integration of Siemens PLCs, which provides a much better interface than Modbus and is supported by many COTS controllers from different vendors.
    E.g., some NI compactRIO crates offer an OPC UA server for integration. (cRIO being used in MPS and in the timing/pulseID system.)

Development Process and Deployment Overview

13. The selection of Ansible as a configuration management and deployment tool is appreciated. It seems an appropriate and capable solution for a system this size.
14. Compiling all control system software from sources on each server is very unusual. Especially if all hosts are under configuration control, there is no obvious need to recompile separately on each.
    Having the complete sources and toolchain on each machine opens up the room for uncontrolled local modifications.
15. Using a freely variable set of versions for EPICS Support modules may be appropriate and have flexibility advantages during development.
    As soon as commissioning starts, the used set of EPICS Modules needs to be under strict version control. Many installations are using fixed sets of Base and Support modules, creating their own "releases", to limit the number of combinations that need to be supported.
16. "Build and Deployment" was a focus topic at EPICS Collaboration Meetings in 2019 [4] and 2022 [5]. The materials from these meetings will provide good input when reiterating the requirements for deployment.
17. The chosen granularity of "one repo per IOC" is on the fine end of the spectrum. Considering systems like the camera integration that contain many identical IOCs, it might be more appropriate to choose "one repo per application", where the application may contain multiple IOCs running the same software. This matches the standard EPICS application structure.
18. Running IOCs in Linux containers adds a lot of complexity. The committee would suggest conducting a cost-benefit analysis before committing to such an approach.
19. The committee has not seen plans and workflows related to development, testing, staging and production. Who will decide about upgrades? Who will be able to deploy? What testing is required? How are upgrades and deployments documented? Again, some of these aspects may not be critical now, but they will be as soon as commissioning starts.

Control System GUI

20. The committee has concerns about the chosen technology.
    While there are certain advantages in the proposed solution, it is not obvious to the

committee that they are sufficient to justify the development effort, especially considering the risk of technology obsolescence and the foreseeable effort for long-term support.

21. Dependency on (as we understand) a single developer introduces high risks in case this resource becomes unavailable.
22. Given the experience of the controls team with laser facilities, the committee is confident that an appropriate level of user-friendliness will be achieved.
23. Having to write HTML to create the screens severely limits the number of contributors who can and/or are willing to make the effort. Other EPICS laboratories have had good experiences with operators creating screens using graphical editors.
24. The chosen programming language C# has no libraries for the pvAccess protocol (PVA). To take advantage of PVA's capabilities, especially for image data, further development will be necessary and should be factored into the cost and resource plan.

Data Management Overview

25. The EPICS-Kafka interface is based on the ESS ADPluginKafka. The committee appreciates this and encourages EPAC to continue the collaboration. New features and enhancements can be shared for mutual benefit.
26. The EPICS-Kafka forwarder should be based on pvAccess and use Normative Type containers - the structures in the NT specification [6] should be mapped to Kafka data schemas. (That mapping could be a shared community standard.)

Bluesky – Experiment Orchestrator

27. The committee appreciates the approach using a collaborative framework. Bluesky seems to be a reasonable, appropriate choice.
28. Starting with an (albeit limited) full-loop demonstration setup seems a good idea to involve stakeholders early on.
29. The committee fully supports expanding the collaboration with close partners (notably Diamond and ISIS) on this approach.

Recommendations

1. The committee recommends adopting pvAccess (PVA) and Normative Type structures as the standard EPICS protocol. Its advantages (especially regarding the transport of images and overall future-proofness) outweigh its risks. Use of Channel Access should be restricted to being the fallback where PVA is unavailable or not possible.
2. In the opinion of the committee, the EPAC-specific GUIs can be used where they are appropriate. The committee recommends conducting a thorough evaluation of the cost vs. benefit balance and preparing a fallback strategy to avoid full dependency on the custom tools. A possible fallback strategy could be using community tools (e.g., CS-Studio Phoebus) for a part of the GUIs and being prepared to migrate custom GUIs to the community tool as required.
3. The committee recommends considering adding a minimal (MRF) event timing system to make the pulse ID system more solid and give the possibility to add useful functionality on top. The cost for a minimal system is moderate and well justified by the added determinism, capability to handle the foreseen higher (100 Hz) repetition rates and the additional functionality.

4. Concerning the Machine Protection System, the committee recommends taking a more rigorous approach to analysing the risks. The beam switchyard appears to have a high-risk potential. Developing a good understanding of the associated hazards and a robust fail-safe protection strategy will be important.
5. The committee recommends separating the MPS and vacuum control functionality and allocating each of them their own PLC resources. The added cost is minimal and will be more than offset by the resulting independence of the systems and their maintainability.

Detailed Assessment of the pulseID Mechanism

The plan for data acquisition seems to be sufficient for the initial planned 10 Hz repetition rate. The proposed Channel-Access-based PulseID system should be able to generally keep up with this rate. However, one cannot rule out occasional drops and difficulties in aligning data to particular laser pulses.

Depending on the severity of the consequences of such occasional drops, it might be advised to consider countermeasures to make the system more robust. At a minimum, consideration should be given to the mechanisms to determine when an acquired dataset cannot be reliably aligned with a particular laser pulse. The committee recommends further clarification of this matter and following up with appropriate measures.

The main contributors to indeterminism in the proposed PulseID system would be:

1. Accurately tracking the laser triggers and reliably incrementing/persisting the aggregated PulseID number in the EPICS control system.
2. Reliably distributing this PulseID from the counting/tracking system to the interested data acquisition IOCs. These IOCs are envisaged to be mainly camera acquisition systems with many connected GigE cameras via AreaDetector.

The CLF proposal of using an FPGA-based system to reliably count/track trigger pulses in a deterministic manner will mainly address item 1 above. However, consideration must be made of the type of system selected, as the PulseID must ultimately be transferred into an EPICS IOC for transmission via Channel Access protocol. The software interface between the FPGA and CPU in the counting/tracking system must also be considered when making the selection of an appropriate architecture. The integrated system must be capable of reliably incrementing and transferring the PulseID from the FPGA to the CPU and onwards via Channel Access within the available 100 ms window between pulses. Software interfaces based on periodic polling may not be able to reliably achieve this sustained rate. The panel would recommend a direct memory register access method between CPU and FPGA using, for example, the EPICS regDev driver.

Item 2 can only reliably be addressed with a deterministic PulseID system. With the proposed system for 10 Hz operation, reliability would likely be improved by considering measures such as protecting the transmission of PulseID via Channel Access over the Ethernet network from latency introduced by a busy network. The isolation of this PulseID transmission onto a second private intra-IOC network would improve the reliability of PulseID transmission.

One option worth considering is to augment the timing system with a set of MRF event system hardware. With that addition, PulseID transmission and reception up to (and even beyond) a possible 100 Hz operation would be deterministic. A minimal system for PulseID counting and transmission would include a master event system, comprising an IOC and

MRF Event Generator/Master, and individual MRF Event Receivers added to each host of data acquisition IOCs interested in deterministically receiving PulseIDs.

There are resources available to help set up an MRF system and the EPICS integration.

An MRF event system would also open up several other opportunities. These include:

- Direct integration of the external timing triggers into the EPICS control system, allowing trigger events to be used to drive software events. This also adds the ability to generate software events that can be distributed across the system and converted into electrical triggers.
- Ability to distribute master oscillator-locked triggers and/or clocks to data acquisition hardware.
- EPICS-integrated software monitoring of the event system to monitor health and diagnose faults.
- Ability to integrate electrical triggers into the system, which can be distributed over the event system and emitted as electrical signals from any system with an Event Receiver. This can complement an existing trigger system or act as a future replacement.
- Option for delay compensation, to correct timing drifts across fibre connections and to reduce the need to accommodate for fibre installation asymmetry.
- Ability to send a limited set of deterministic data across the event system, e.g., machine mode/setup changes.
- Locking of events/triggers to MO signal and/or mains 50Hz or another external reference signal.

5. References

[1] Five Years of EPICS 7…, ICALEPCS2023 https://accelconf.web.cern.ch/icalepcs2023/papers/th1bco01.pdf

[2] NTNDArray Normative Type (specification), https://docs.epics-controls.org/en/latest/pv-access/Normative-Types-Specification.html#ntndarray

[3] P4P, https://mdavidsaver.github.io/p4p/

[4] Build and Deployment Session at the 2019 EPICS Collaboration Meeting, https://indico.cern.ch/event/766611/timetable/#20190606.detailed

[5] Build and Deployment Workshop at the 2022 EPICS Collaboration Meeting, https://indico.cern.ch/event/1173788/sessions/459213/#20220920

[6] Normative Types specification, https://docs.epics-controls.org/en/latest/pv-access/Normative-Types-Specification.html