ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

CONSOLIDATION AND REDESIGN OF CERN INDUSTRIAL CONTROLS FRAMEWORKS

P. Golonka[†], F. Varela-Rodriguez, CERN, Geneva, Switzerland

Abstract

The Industrial Controls Frameworks, JCOP [1, 2] and UNICOS [3, 4], have been employed to develop hundreds of critical controls applications in multiple domains, like the LHC Experiments Detector Control System, accelerator complex (cryogenics, powering, interlocks) and technical infrastructure, leading to an unprecedented level of homogeneity. These frameworks, used by around a thousand of developers worldwide, will now undergo a major consolidation and re-engineering effort to prepare them for the new challenges of the next 20 years in the High Luminosity (HL) LHC era, as well as to streamline their maintenance.

The paper presents the challenges that will be faced during this project due to the breadth of technological stack and large code-base contributed over two decades by numerous authors. Delivery of innovation induced by evolution of technologies and re-factoring of the ageing code must be done in a way that ensures backward-compatibility for existing systems. The vision and the current state of the frameworks is discussed, alongside the main deliverables planned in the medium term. Lessons learnt, optimizations of processes to make best use of the available resources and efforts towards open-source licensing of the frameworks are also presented.

INTRODUCTION

The development of the two CERN Industrial Controls Frameworks (JCOP and UNICOS) started in early 2000's. With the approach of the LHC era, the CERN Joint Controls Project [5] was set up in 1998 to streamline the development efforts for the controls of the LHC experiments and avoid duplicated efforts. The recommendation of the project, approved by the experiments, was to select a common set of industrial COTS components (at the hardware, middleware and software layer), with centralized effort to validate, enrich, enhance, and provide support. These components should be used as standardized and trusted building blocks. Key to the success of JCOP has been the strong partnership established with the commercial providers. At that time, high and low voltage power supplies and VME crate controllers were chosen with maintenance contract signed with vendors, as well as the OPC Data Access (DA) as the preferred middleware and PVSS SCADA, later rebranded to WinCC Open Architecture (OA) [6] as supervisory system. This careful selection/validation process involved around 10 man-years of the central team, yet it spared similar scattered efforts from each of the LHC experiment's controls team.

† Piotr.Golonka@CERN.CH

At the same time another initiative was started to develop the control system for cryogenics of the LHC [7], adopting the formal approach for process decomposition, modelling (IEC61512-1) and engineering that employed industrial components.

The project, called UNICOS, chose to standardize on Siemens and Schneider Programmable Logic Controllers (PLC) as well as the supervisory system and soon converged to use the same SCADA as selected by the JCOP project in 2001.

Both Frameworks are complementary in their approach and employ common components in their stack of technologies. Whereas the JCOP Framework provides guidelines and tools to implement flexible and complex control systems that comprise a large variety of equipment, like the detector control system of the LHC experiments, the UNICOS Framework defines a strict engineering process and provides tools to develop PLC-based control systems and allowing to generate parts through templates. Component-based approach used in both allow to compose functionality as required by the control system and extend it as necessary. UNICOS employs numerous components delivered by the JCOP Framework as shown in Fig. 1.

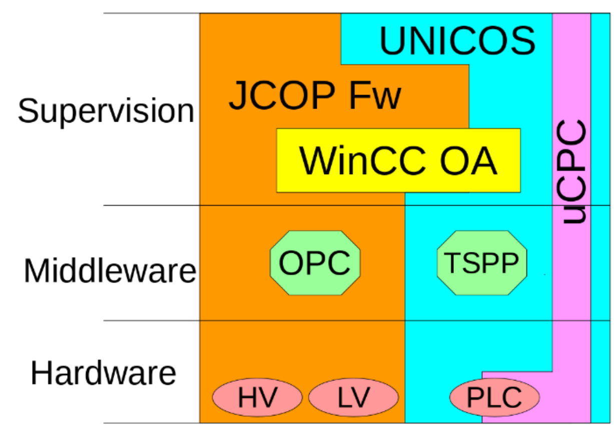

Figure 1: JCOP and UNICOS frameworks relation.

The JCOP framework is developed as a collaboration between the LHC Experiments and CERN industrial controls group, where each of the partners develops, maintains and supports generic components that are widely used across CERN. The UNICOS project is entirely developed by the industrial controls group and is provided to many groups at CERN to engineer their control systems.

The first controls applications at CERN based on the Industrial Controls Frameworks entered production in 2001. Today, the Industrial Controls Frameworks are in use in more than 600 critical controls applications to ensure the smooth running of the accelerators, detectors and technical infrastructure contributing to unprecedented performance of the research facilities. The flexibility of the frameworks

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

has allowed to drastically minimize the implementation time and effort to develop control systems for smaller experiments like NA62 [8] and ProtoDune [9].

Although a major effort was performed to maximize reuse, important differences exist between the two frameworks in terms of functionality, architecture (e.g. device models) and implementation that justify a further rationalization effort to streamline their integration.

In addition, the ambitious HL LHC project will require to keep many of these control systems operational for another 20 years, as well as to improve their performance and functionality. This, together with the natural ageing of the frameworks (like any other large scale and long lifetime software project), motivates a major consolidation effort where some of the building blocks will need a major re-engineering process.

The CERN Long Shutdown 2 represents a unique opportunity to launch the consolidation and redesign project since it offers enough time to test prototype implementations and to introduce major changes into the control systems.

THE CONSOLIDATION & REDESIGN PROJECT

Motivation

The main motivation for the consolidation and redesign of the CERN Industrial Controls Frameworks are to:

- prepare control systems for the challenges imposed by the HL LHC project and experiment upgrades;
- assure long-term maintenance while reducing effort;
- enable and integrate the use of new technologies.

For the first 15 years of operation developers resources needed to be shared, and had to prioritize the requirements needed to operate the LHC and its experiments with little time left to consolidate developments and minimize the maintenance effort of the frameworks.

Despite the fact that initial implementation of parts of the frameworks strictly followed specifications prepared in early phase of the projects, the unprecedented size and complexity of the control systems required dynamic evolution. Requirements had to be discovered or refined by applying prototypes to realistic preproduction systems. Rapid prototyping and development of new features were favoured. As result of this, various prototypes were continuously adapted to the users' needs until they became fully operational building blocks in production systems.

Despite the reduced amount of resources two consolidation attempts were made: a review of the JCOP Framework in the early years of the project and a major re-engineering of the UNICOS Continuous Process Control package in 2011 [10]. However, these reviews focussed on functionality rather than on maintainability.

Even though the industrial controls domain is rather conservative with respect to the adoption of new trends, the frameworks should respond to evolving needs by enabling available technologies such as containers, virtualization, Big Data, Industrial IoT and Programmable Automation Controllers (PAC), etc. Furthermore, technological changes led to the availability of new monitors with high-resolution and different form-factors. Similarly, standardization efforts such as unicode were widely adopted for computer systems. In addition, significant enhancements were introduced with every release of the SCADA, often providing natively the functionality which was previously delivered by the framework, e.g. advanced trending. The latest enhancements to the scripting language (object-orientation) allow to design interfaces with better abstraction and clarity, and develop more expressive code.

However, the requirement of quality and availability on the frameworks during the first years of LHC operation and reduced resources put on hold the integration of this new functionality and the adaptation of the frameworks to these technological changes.

Today the CERN Industrial Controls Frameworks span over all layers of the control system and comprise several million lines of code including: PLC code, a comprehensive library of scripts, tools and panels for WinCC OA SCADA developed using scripting language, as well as C++ used for OPC servers and WinCC OA integrations. In addition, Java tools for automatic generation of applications, Jython templates to generate PLC logic and numerous Python scripts are employed. Urgent needs to catch up on delayed adoption of enhancements or technology changes also motivate the need for major consolidation effort.

Due to the size of frameworks, diversity of technologies and complexity, mastering them present a steep learning curve. This has a major implication in our organisation with a high turnover where it is important to bring new developers up-to-speed as quickly as possible. This complexity also slows down the adoption of the frameworks by other research institutes, like GSI, ITER, or industry.

The Consolidation and Redesign project of the Industrial Controls Frameworks was endorsed by the governing bodies of the JCOP and UNICOS projects with the aim of securing reliable operation and evolutions for the next 20 years of the LHC. The project was launched in 2018 taking advantage of the LHC Long Shutdown 2 with the first steps focusing on creating an inventory of the components of the frameworks, identifying deprecated functionality and code, and the reconstruction of missing knowledge due to the loss of critical manpower over the years.

Lessons Learnt

The experience gained over 15 years of operation of industrial controls systems based on the frameworks are a vital input for the consolidation effort. Some of the lessons learnt include:

Involvement of all stakeholders in all phases of the project, continuous feedback on the status of the project, as well as a formal process to jointly define priorities with end

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

users are crucial to ensure transparency and share responsibility among all participating entities.

Strong partnership with industrial providers is key to the success of the project, as the requirements of HEP facilities may drastically differ from those of industry, e.g. scale, integration with external systems. The partnership must ensure the adherence of the companies to the project, to respect its milestones and, ideally, to influence the evolution of the products. Strong and mutually beneficial industry-research partnership also protect from the undesired effects of technology lock-in when the *development partner* relationship is established instead of *supplier-consumer*. An example of this is the excellent relation built over many years between CERN and Siemens/ETM which led to multiple joint efforts to strengthen the capabilities of the WinCC OA SCADA system. Another example is the collaboration [10] established with the power supply vendors (CAEN, iSEG and Wiener) to work together on the implementation of new OPC Unified Architecture (UA) servers.

Validation of components of the control systems and the stability of the technologies chosen for their implementation, as well as long-term support and availability of spares are of vital importance. The blooming of attractive new technologies, programming languages and tools with an uncertain future could result in decision leading to problems in the long term.

Stable core team in organisations with a high turnover of people like CERN, a stable core team is fundamental to ensure the knowledge preservation and sharing of both, the frameworks themselves and their underlying technologies, as well as the domain specific knowledge. An example is the CERN BE-ICS group providing central maintenance and development for numerous packages, and the central controls teams of the experiments with practical experience in the use of these components at their setups.

Avoid "black-box" components Nowadays, there is a tendency to include third-party software modules that provide an immediate benefit, e.g. shorter implementation time, without insight of the product. Experience shows that in the long run, some of these may require complicate troubleshooting and significant maintenance effort. An example of this was our previous experience with OPC DA servers implemented exclusively by the power supply vendors that motivated the joint effort to develop the new OPC UA servers in collaboration with vendors.

Commercial components may be cost-effective in the long run Although commercial solutions like WinCC OA imply an initial licensing cost, the package provides solid foundations to build industrial control systems and may significantly reduce the development and maintenance effort since they provide much of the functionality required (e.g. scalability, distribution, drivers, trending, archiving, alarming, notifications) and allow developers to focus their efforts onto the necessary business logic to solve a specific problem. We are deeply convinced of applicability of this strategy to implement control systems for research.

Constraints and Challenges

Given the size and complexity of the Industrial Controls Frameworks and their vast deployment in more than 600 critical controls applications, embarking on a major consolidation and redesign effort is a significant challenge. Changes must be done in way that minimize the impact and efforts to upgrade already existing systems. Backwards compatibility must be ensured whenever possible and incompatibilities and their impact must be studied and documented. Migration tools and support must also be provided as necessary.

Project Management

Considering the constraints, resources and its scope, the management for the project is clearly a challenge on its own and calls for optimizations. Previously in the frameworks a single person was responsible for the development and maintenance of a set of components. As a consequence, *knowledge silos* appeared, with responsibility for a significant number of components in the hands of few senior staff members. Experts' overload was effectively preventing any evolution of the components as their experience was also in high demand to resolve issues in production systems, leaving little time for knowledge transfer and proper supervision.

To break this classical deadlock situation, an *agile attitude* was adopted, with elements of its widely known methodologies. We aim at involving junior team members to contribute to development of all components, which requires spending some resources for knowledge transfer and building a high-trust culture. Challenging long-existing architectures and technical solutions and their participation in code reviews already showed benefits and allowed to increase the uniformity of the code.

Adopting *team responsibility* combined with technical means (branched development, automated build process, change-review and acceptance process) encourage everyone to contribute the necessary changes wherever needed. Complete re-engineering of the build-automation, release-control and quality-assurance infrastructure (automated tests) was an essential step towards *continuous-integration* and already paid-off for the invested efforts. In the future, we hope to embrace the atmosphere of welcoming contributed changes, e.g. by the operation/maintenance teams posting bug fixes for issues in production systems, or even larger parts of code in a style known from *community* projects.

Our current experience shows that applying standard *agile* frameworks such as Scrum or Kanban directly is not realistic on the project-wide scope: we experiment with the adoption of some of concepts such as time-boxing (sprints), Kanban boards, stand-up meetings. However, instability of development team, and the need to share time with other activities, such as support for production systems allows us to apply these at a limited scope.

In the early phase of the project we considered an approach with a team of developers dedicated exclusively to the consolidation and redesign project, working on a

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

separated code branch. This idea was abandoned due to the limited resources and also the possible divergence of the "consolidation" branch from the one used in operation, which would consequently increase the risks of deployment in production systems. Instead, the work progresses through "sprints", in small increments, merging the consolidation efforts with necessary bug-fixing.

Reflecting on feedback from key users as well as on our own experience we recognized that transparency and predictability in project planning and progress also needed improvement. Even though issue-tracking was an integral part of the project since its beginning, some procedures for delivery planning and prioritization were not clear. This resulted in a very backlog of pending requests building up (incidents, bug reports, feature requests, some of which not relevant anymore). Some user requests were never tracked or were accepted for implementation without proper prioritization, blurring the planning further. To address these issues, we formalized the workflow for requests for both projects. In addition, following the proven method for request approval already present in the JCOP Project, the structure for the management of the UNICOS project was changed and two new bodies were implemented, the Technical Committee and the Advisory Board, taking the responsibility for defining and adjusting priorities as well as making recommendations to the project management. Systematic clean-up of the issue-tracker followed allowing to build a clear high-level overview of goals and priorities for the project. This ongoing activity of reviewing existing issues is essential for the next steps of the consolidation and redesign project.

Expected Deliverables

The Consolidation and Redesign project will run for the next years due to the need of sharing resources with other concurrent projects. Project deliverables will be staged according to the priority defined jointly with the stakeholders, considering the impact of the changes, implementation cost and possible timing for deployment, e.g. changes that break backwards-compatibility can only be deployed during long shutdowns. Reflecting the broad scope of the project, the deliverables span from clean-up, deprecation of code and functions, code refactoring, quality assurance, new build and release process to enhancements and new features as well as adopting new trends and technologies. In this chapter, we present the most important of these, which are already in progress or planned for execution

OPC-UA Migration With the new OPC UA specification released in 2008, the major limitations (such as complexity, hard requirement on the MS Windows platform, complex security setup) were overcome. Soon after OPC DA was phased out and the need for migration reaffirmed by rapid adoption of OPC UA.

As presented in [11], CERN observed the need for migration as an opportunity to collaborate with hardware vendors on development of the new OPC UA servers contributing in-house experience, coordination as well as developing some of software components. Following the agreement signed with CAEN, Wiener and ISEG, OPC UA servers have been developed. At this moment the migration from the OPC DA to OPC UA is in progress, after having been successfully tested in realistic large-scale setups. With some of the installations counting thousands of high voltage channels, and tens of thousands of OPC items, the supervision layer of control systems requires a massive migration/reconfiguration effort. Adequate automated tools were prepared within the JCOP Framework to make this task effortless for the users. It is expected that the CERN-wide migration will complete before the end of the Long Shutdown 2, i.e. in 2020. Even though the OPC UA server codes are not formally a part of the frameworks, they will be available to all interested users.

CANbus Gateways and CAN Common Module (CCM) CANbus remains the proven fieldbus used by geographically-spread systems for large detector control systems. Their physical interface to computer equipment used to be provided by the PCI interface cards, and CAN-to-USB interfaces. In recent years, CAN-to-Ethernet interfaces have become an attractive option since they remove the need for a direct connection between the CANbus and the computer, thus enabling redundancy and virtualization. Confronted with the lack of standardized API to access CAN hardware (on Windows), CCM was developed allowing for common programming interface to be used for a variety of gateways, on Windows and Linux. The CCM is a key component of the Quasar framework [12] for the OPC UA servers development.

Changes to the Framework motivated by the evolution of the CERN controls infrastructure The Long Shutdown 2 defines time limits on some elements of CERN controls infrastructure, namely in the accelerator sector. Deprecation of the RDA2-based CMW (Common Middleware) [13] used to communicate with the accelerator control systems, as well as the CERN Accelerator Logging System (CALS) [14], with appropriate functional replacements required development efforts in the frameworks, as well as orchestration of the migration process with the service providers. Efforts related to the migration to the new NXCALS system are described in [15].

Replacement of the ageing Windows Terminal Server infrastructure with sessions provided by Linux machines, and made available on *any* desktop machine/platform, with appropriate security and ergonomics is being worked on in collaboration with the CERN central IT support and the accelerator control groups. Numerous technologies, such as Xpra [16] are being validated and prototypes are presently tested. Progressive deployment to pilot users is envisaged once the decision on the new technology is made.

New S7+ Driver for the Siemens S7-1500 PLCs As reported in [17], progress is being made for proper integration of this new family of PLCs. Enabling new block optimizations and integration with engineering tools should be achieved, while assuring optimal performance and reliability of communication. This new PLC series could be interfaced to WinCC OA in two ways: the native

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

S7+ WinCC OA driver or via an OPC UA server embedded in the hardware modules. The final choice will be based on the performance, stability and on the determinism of the communication. The adoption of the new driver would require modifications at the SCADA level to adapt to the new addressing schema. However, even with the promising S7+ communication, several issues need to still be addressed [17]

Modernization of the HMI The original human-machine interface panels for production applications based on both the JCOP and UNICOS frameworks were designed and developed over a decade ago with numerous functional constraints enforced by shortcomings of the technology at that time. Thousands of panels and synoptic views were developed with fixed-size to match the resolution of 19” screens present in control rooms, due to lack of layout-management features in WinCC OA.

An ongoing project to renovate consoles in the CERN Control Centre foresees that screens are replaced with ones with much higher resolution and different form-factor. A requirement was put up by the operation team that a new HMI is developed such that the existing synoptic panels could be displayed in an ergonomic way on the new screens, allowing to zoom the existing panels and making most optimal use of the available screen size, which calls for responsive design and the possibility for the user to be able to “compose” the layout of the screen.

Over the past years numerous significant enhancements have been implemented in WinCC OA, including dynamic layout-management for UI elements within a window, style-sheets, support for responsive design and touch interfaces, new or improved trending and visualization widgets, etc. After evaluation of these new technical means, and a prototype validated with users, a project of a new HMI (main operator window) for the UNICOS framework was started. It will deliver the first version, focusing on the most common needs of the operators, by the end of 2019. Subsequent versions will complete the functionality so it will become a replacement of the currently used fixed size HMI.

In parallel, efforts are undertaken to review and to consolidate other frequently used panels to enhance their usability and in particular, to make them “resizable” by applying layout management. Tools such as trending profit most from these enhancements, improving the ergonomics and responding to long-standing operator requests. It was decided that the trending tool would undergo major refactoring and the invested effort would not only clarify and stabilize the existing functionality but it would also introduce the new features that became available in WinCC OA throughout the years in response to CERN requests.

Even though good practices for ergonomic HMI design prescribe conservative adoption of visual enhancements, we also look at modernization of the user-experience, following contemporary style guides. The style-sheets technology available in the SCADA product allows for easy customization at runtime. In addition, enabling gestures such as drag-and-drop known commonly from computer interfaces or touch displays are also considered.

New Alarm Screen The alarm screen is considered as one of the most important elements of any supervision system and hence should provide adequate functionality and be ergonomic. Throughout the years, a layer of enhancements was delivered by the frameworks on top of the standard WinCC OA alarm screen. However, limits were reached that prevent further functional enhancements and scalability/performance improvements calling for a completely new approach.

To probe the functionality aspects, a proof-of-concept was rapidly developed in JavaScript, and integrated with WinCC OA using the WebView widget. New use cases such as powerful filtering or responsive UI design were explored, yet its performance was not adequate for large systems. The requirements were formally documented and presented together with the prototype to the WinCC OA users’ community as well as ETM’s technical management, and the project for a *Next Generation Alarm and Event Screen* (NGAES) for WinCC OA was started. The choice of the architecture and technologies to be employed was agreed with WinCC OA experts at ETM, with the aim of possible hand-over of the project to the company for further maintenance.

By employing C++ and advanced data-modelling and filtering techniques provided by the Qt framework, as opposed to scripting, the NGAES allow for tens of thousands of alarms to be displayed and filtered with great performance in either live or historical mode. Free-text or complex logical expressions on meta-data may be employed for search and filtering and the layout is fully configurable. Dedicated optimizations allow to use it optimally also on “faceplate” panels that need to report on alarms for a handful of selected devices (aggressive filtering).

The project is in an advanced stage of development, having delivered first functionality previews in Spring 2019. The first version that could be used alongside the current alarm screen should be ready by mid 2020.

The architecture and components of NGAES are developed with generic use-cases in mind to be applied whenever a live-display with powerful filtering/grouping of tabular data is needed. Therefore, at the second stage of the project, a new Event Screen will be implemented to report on live and historical value changes in large distributed systems. The completion of the second stage is planned for 2020 and should respond to long-standing demands for missing functionality expressed by our framework users’ community. This topic will be subject of a dedicated paper.

Next Generation Archiver Improvements to the archiving system of the CERN Industrial Controls applications, were reported in numerous occasions at this conference [18] and showcase the excellent partnership in product development between CERN and ETM/Siemens, within the CERN openlab project [19].

Most recently, we reported on the shared development of the *WinCC OA Next Generation Archiver* (NGA) [20] providing a flexible and extendable archiving architecture whereby new data-storage technologies may be easily

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

integrated through a standardized interface and the mechanism of plug-ins. Archiving to many back-ends working in parallel, with various data-decimation policies enable a completely new set of use cases. While ETM develops a backend based on InfluxDB [21], CERN focuses on the development of an Oracle backend to assure maximum compatibility with the RDB Archived used currently in all CERN production systems.

The project approaches the milestone of deployment of its first version in the detector control system of the ALICE experiment [22]. Through the development of a custom “ALICE O2” backend, the NGA will not only be used to archive the data but also at the same time assure the fundamental data link between the detector control system and the brand new O2 system [23].

The InfluxDB backend is undergoing functionality and scalability tests. It enables an interesting way to create web dashboards using solutions such as Grafana [24]. In addition, archiving to the InfluxDB backend with local storage, when combined with parallel archiving to Oracle for safety, may provide much quicker access to short-term data, for small or medium-sized project. This conjecture and its limits are being validated now.

Efforts are also made to validate the technology for the future high-performance backend, allowing for low-latency access to large sets of time-series data, which, at the same time, will enable effective data analytics. Technologies such as Apache Kudu [25] or TimescaleDB [26] are being considered. A native read-write backend to the NXCALS system is also envisaged, which will significantly simplify the architecture.

The NGA will become the main archiver system for the ALICE Detector Control System as of 2020. We also plan its progressive deployment in other domains to gain more experience and to build further confidence in the new solution.

Clean-up and consolidation efforts Review and then removal or deprecation of unused functionality was one of the first goals for the project. For a number of unused components (some very complex) it was agreed with the users to stop the support and maintenance.

Using static code-analysis techniques a tool was created to analyse the existing code base of the frameworks, as well as copies of production systems, in search of unused functions. For every function reported by the tool, code-analysis was performed by developers, to either deprecate the function, or document the reasons for which it should be kept. The efforts completed in 2019 allowed to remove a few thousands of lines of dead code and also a significant amount of lost knowledge was reconstructed in the process of code reviews.

On a number of occasions refactoring of code was applied, with careful testing of compatibility of existing interfaces through unit-testing. The use of object-oriented extensions of the WinCC OA scripting language allowed to make code much clearer, and reduce copy-and-paste patterns, leading to reduction of up to 50% of the code.

The structure of components was cleaned up to align them with the conventions, simplify the structure and to optimize the resources (e.g. change the library-loading policies). The risk of potential incompatibilities that arose was discussed and approval was given by the JCOP Coordination Board to proceed. The changes will become effective in the next release of framework, which will be used to upgrade of all control systems at CERN in 2020.

The activities related to consolidation and clean-up are far from being complete. A systematic review of all libraries and components, starting from the most important ones will require a couple of years to complete.

Quality, reliability, robustness Certain concepts of DevOps[27] are particularly applicable to industrial controls systems when used at scale. A spectacular growth in the number of systems to maintain, e.g. during SCADA upgrades with limited service time slots, made us realize the importance of trust in the reliability of delivered components and the process of their validation in all layers of frameworks and realistic setups. From the developer perspective, early detection/reporting of software defects and realistic system tests to easily reproduce problems, makes the fixes much easier to be provided than applying a hotfix in production system.

The *Vertical Slice* project provides a testing/validation platform available for developers allowing them to contribute unit and validation tests. It consists of real hardware (PLCs and power supplies) or sophisticated emulators [28], a number of middleware layers, and a supervision system with data archiving. The platform is rebuilt daily from the most recent versions of framework components, and stores the results of tests for future reference and comparison. Daily reports are available and displayed on a status screen as shown in Fig. 2.

Figure 2: Monitor exposing the results of *Vertical Slice Test Platform* motivates the developers, increases work visibility and provides early reporting of faults.

A solid build automation and release management infrastructure is a major contributor to overall quality. Each commit leads to the complete framework release, ready to download and undergo testing: automatic unit, nightly or manual. The involvement of developers in the definition of Continuous Integration workflows and their technical implementation was essential to the success.

To increase test coverage for framework components the catalogue of unit-tests and integration-tests is currently being extended with new features being systematically developed with the corresponding test cases. Even though we recognize the usefulness of strict unit-testing approach,

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

at this stage we set the priority on integration and system tests due to their larger coverage.

Improving the documentation is frequently raised as a priority task, in particular, by new users of the framework. We plan to address this in stages: firstly, we decided to simplify the effort required to create or update components help: online reference documentation is generated automatically by the CI infrastructure as a step in the build process. This online help integrates smoothly with the engineering tools. In a second stage, the documentation for public framework interfaces will be reviewed, completed and updated. Ultimately, user-guides and user-reference chapters for documentation will be completed giving meaningful contexts and examples.

OPEN SOURCE LICENSING

Over the past few years we observe a growing interest in the use of the frameworks beyond CERN. In addition to a few research institutes (GSI, ITER) already making use of CERN frameworks, other labs as well as industrial companies have also expressed an interest in using them. In addition, various systems, such as cryogenics or detector cooling, based on CERN blueprints are assembled out of CERN by third party sub-constructors. For their operation, these facilities will require controls that are based on the CERN frameworks.

Up to present, the use and copyright for the frameworks was restricted to CERN, its associated institutes and permission was granted explicitly. Considering the effort for preparation of individual bilateral agreements with every new interested party, we prefer a globally applicable solution for licensing the framework, preferably on the non-restrictive open source license.

With assistance of the CERN Knowledge Transfer group, we are in the course of clarifying the legal aspects, such as IP clearance from all contributors. On the technical side, numerous efforts mentioned in the previous chapters, such as clean-up, documentation, code and version management are required for effective open-sourcing of the code.

The decision to choose open source is aligned with the *CERN Policy on IP Management* announced this year, and particularly with recommendations from the Open Source License Task Force. Assuming enough interest is expressed, a community around the project may emerge with needs to structure the contributions and releases, assure expertise and maintain technical knowledge. Balancing the needs and resources of CERN, and setting up the project management model will need to be addressed, yet the steps undertaken already pave the way to such extension.

We hope to complete the work on the legal aspects of the framework towards the end of the year, and gradually release subsequent components with open source release; a complete functional framework release on an open source license should be possible in 2020. Details will be announced to the research and industrial communities.

The scope for the open source license release of the framework does not extend to its necessary commercial components (hardware and software). However strong partnership with vendors described in the previous chapter allows to effectively avoid the effects of the vendor lock-in.

SUMMARY AND OUTLOOK

The list of other high-level features and enhancement to the project is long (performance optimisation, secure and scalable web access, event replay, data-driven engineering, integration of data analytics, IIoT to name a few) as well as dynamic evolution of industrial automation (Industry 4.0). The priorities of these in the context of the consolidation and redesign project need to be considered and well balanced over the next years to assure their completion.

Once the most important consolidation tasks are completed for the frameworks, a number of other ambitious tasks is awaiting, such as the consolidation of the device models to provide extensible and reusable tools working across both the JCOP and the UNICOS frameworks or to streamline their integration. We would also like to explore possible synergies with other major control frameworks and their ecosystems, following prior proof-of-concepts such as those described in [29].

Last but not least, we believe in the vision of *DevOps* [27] that may to some degree be applicable to certain domains of Industrial Controls. Solid, reliable and functional frameworks are a corner stone to trigger the transition.

ACKNOWLEDGEMENTS

The authors would like to thank to all Contributors to the project, and namely the members of the JCOP Coordination Board, as well as the former JCOP and UNICOS Framework project leaders, Wayne Salter, Manuel Gonzalez-Berges, Philippe Gayet, Enrique Blanco-Vinuela whose vision, ideas as well as their conceptual and technical contributions made the project reach its goals.

REFERENCES

[1] O. Holme, M. Gonzalez-Berges, P. Golonka, S. Schmeling, “The JCOP Framework”, in *Proc. 10th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'05)*, Geneva, Switzerland, Oct. 2005, paper O3_005.

[2] M. Gonzalez-Berges, R. Barillere and F. Bernard “Frameworking: a collaborative approach to control systems development”, in *Proc. 10th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'05)*, Geneva, Switzerland, Oct. 2005, paper O1.049-6.

[3] E. Blanco, F. B. Bernard, P. Gayet, and H. Milcent, “UNICOS: An Open Framework”, in *Proc. 12th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'09)*, Kobe, Japan, Oct. 2009, paper THD003, pp. 910-912.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEDPL04

[4] R. Barillere, "UNICOS a Framework to Build Industry Like Control Systems Principles Methodology", in *Proc. 10th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'05)*, Geneva, Switzerland, Oct. 2005, paper I3_002.

[5] D. Myers, "The LHC Experiments' Joint Controls Project, JCOP", in *Proc. 7th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'99)*, Trieste, Italy, Oct. 1999, paper MA1O06.

[6] Siemens Simatic WinCC Open Architecture, by ETM, `https://www.winccoa.com/`

[7] P. Gayet and E. Blanco, "LHC Cryogenics Control System: Integration of the Industrial Controls (UNICOS) and Front-End Software Architecture (FESA) Applications", in *Proc. 11th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'07)*, Oak Ridge, TN, USA, Oct. 2007, paper WOAA03, pp. 281-283.

[8] P. Golonka *et al.*, "Detector and Run Control Systems for the NA62 Fixed-Target Experiment at CERN", in *Proc. 15th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015, paper MOPGF020.

[9] A. Scarpelli, "ProtoDUNE and a Dual-phase LArTPC", Feb. 2019, `https://arxiv.org/pdf/1902.04780.pdf`

[10] E. Blanco Vinuela *et al.*, "UNICOS Evolution: CPC Version 6", in *Proc. 13th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'11)*, Grenoble, France, Oct. 2011, paper WEPKS006, pp. 786-789.

[11] B. Farnham, F. Varela, N. Ziogas, "A homogenous approach to CERN/vendor collaboration projects for building OPC-UA servers", in *Proc. 16th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, paper THPHA009.

[12] P. Nikiel, P. Moschovakos, B. Farnham, and S. Schlenker,, "Quasar, the Full-Stack OPC-UA Software Ecosystem for Controls Middleware of LHC Experiments" , presented at the 17th Int. ICALEPCS Conf. (ICALEPCS'19), New York, USA, October 2019, paper MOPHA100, this conference.

[13] K. Kostro, J. Andersson, F. Di Maio, S. Jensen, and N. Trofimov, "The Controls Middleware (CMW) at CERN - Status and Usage", in *Proc. 9th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'03)*, Gyeongju, Korea, Oct. 2003, paper WE201.

[14] C. Roderick, L. Burdzanowski, and G. Kruk, "The CERN Accelerator Logging Service - 10 Years in Operation: A Look at the Past, Present, and Future", in *Proc. 14th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'13)*, San Francisco, CA, USA, Oct. 2013, paper TUPPC028, pp. 612-614.

[15] I. Prieto and M. Sobieszek, "Big data archiving from Oracle to Hadoop", presented at the 17th Int. ICALEPCS Conf. (ICALEPCS'19), New York, USA, October 2019, paper MOPHA117, this conference

[16] P. Wulles, "Remote Forwarding of Human-Machine Interfaces for Industrial Controls.", CERN Internal note CERN-STUDENTS-Note-2018-015, 2018, `https://cds.cern.ch/record/2633643/files/ShortReport.pdf`

[17] J. O. Ortola Vidal and M. Vazquez Muniz, "Integration of New Siemens S7-1500 PLC Family in UNICOS-CPC: Engineering Challenges and Performance Evaluation", presented at the 17th Int. ICALEPCS Conf. (ICALECPS'19), New York, USA, October 2019, paper WEPHA114, this conference

[18] P. Golonka, M. Gonzalez-Berges, J. Hofer, and A. Voitier, "Database Archiving System for Supervision Systems at CERN: a Successful Upgrade Story", in *Proc. 15th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015, pp. 129-132. `doi:10.18429/JACoW-ICALEPCS2015-MOPGF021`

[19] CERN openlab project, `https://openlab.cern/`

[20] P. Golonka, M. Gonzalez-Berges, J. Guzik, and R. Kulaga, "Future Archiver for CERN SCADA Systems", in *Proc. 16th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 1442-1445. `doi:10.18429/JACoW-ICALEPCS2017-THPHA037`

[21] InfluxDB, `https://www.influxdata.com/`

[22] F. Carena et al., "The ALICE Experiment Control System", in *Proc. ICALEPCS'05*, Geneva, Switzerland, Oct. 2005, paper MO4A.3-7O

[22] The ALICE Collaboration, "Technical Design Report for the Upgrade of the Online–Offline Computing System", CERN-LHCC-2015-006.

[24] Grafana, `https://grafana.com`

[25] Apache Kudu, https://kudu.apache.org/

[26] TimescaleDB, `https://www.timescale.com/`

[27] DevOps: a set of software development practices, `https://en.wikipedia.org/wiki/DevOps`

[28] M. Ludwig *et al.*, "Developing and Validating OPC UA based Industrial Controls for Power Supplies at CERN", in *Proc. 12th Int. Workshop on Emerging Technologies and Scientific Facilities Controls (PCaPAC2018)*, Hsinchu, Taiwan, `doi:10.18429/JACoW-PCaPAC2018-WEP04`

[29] M. Ritzert, E. Blanco Vinuela, M. Ostrega, and L. Zwalinski, "UNICOS Framework and EPICS: A Possible Integration", in *Proc. 16th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 915-917. `doi:10.18429/JACoW-ICALEPCS2017-TUPHA201`