ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA009

A HOMOGENOUS APPROACH TO CERN/VENDOR COLLABORATION PROJECTS FOR BUILDING OPC-UA SERVERS

Ben Farnham, Nick Ziogas, Fernando Varela Rodriguez CERN, Geneva, Switzerland

Abstract

Industrial power supplies deliver high and low voltage to a wide range of CERN's detector and accelerator components. These power supplies, sourced from external companies, are integrated into control systems via industry standard OPC servers. The servers are now being modernized. A key lesson learnt from running the previous generation of OPC servers is that vendor specific, black-box implementations can be costly in terms of support effort, particularly in diagnosing problems in large production-site deployments. This paper presents the projects producing the next generation of OPC servers; following an open, collaborative approach and a high degree of homogenization across the independent partners. The goal is to streamline development and support costs via code re-use and a template architecture. The collaborations aim to optimally combine CERN's OPC and production operations knowledge with each company's experience in integrating their hardware. This paper describes the considerations and constraints taken into account, including legal aspects, product commercialization and technical requirements to define a common collaborative approach across three hardware manufacturers.

PROJECT OUTLINE

The accelerator chain and experiments at CERN require a feed of consistent, stable and controllable high and low voltage. Where applicable, this requirement has been satisfied, by purchasing industrial power supply units from commercial companies. CERN has an incumbent investment of tens of millions of Swiss francs in a range of equipment from three Member States vendors: CAEN, ISEG and Wiener (ordered alphabetically). Each company designs, builds, sells and maintains their own models of industrial power supplies. This inhomogeneous mix of hardware is currently integrated in to CERN control systems via a middleware layer based on the industry standard OPC-DA [1] (DA stands for Data Access) protocol.

Due to obsolescence of the OPC-DA protocol (also latterly known as OPC Classic), this middleware layer is undergoing a significant migration, fromthe old COM based, MS Windows only, OPC-DA standard to the modern, platform agnostic OPC-UA [2] (UA stands for Unified Architecture) standard.

This paper outlines the process of collaborating with these three independent commercial suppliers, with assistance from CERN's Knowledge Transfer group (henceforth KT), to ensure that:

- Post-migration, the equipment installations at CERN continue to fulfil the required function in the laboratory.
- Opportunities for increasing homogeneity in the middleware integration software components and the procedures by which they are built are identified and acted upon.
- Maximum benefit is taken from the collective expertise of both CERN and the hardware providers: from CERN's side this includes experience of operational requirements, standards support and software engineering; from the commercial partners' sides this includes deep knowledge of their hardware and firmware and their systems integration experience.

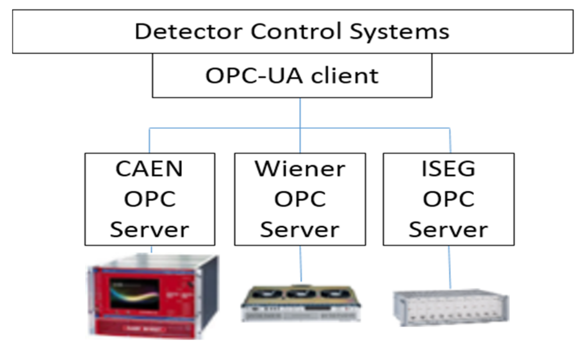

Figure 1: Hardware Controlled via OPC Middleware.

OPC-DA SERVER OPERATION: LESSONS LEARNED

Despite the similarity in names, OPC-UA is a fundamentally different protocol [3] from OPC-DA. This modernisation, moving from OPC-DA to OPC-UA, requires a complete rewrite of the OPC layer. This presents an opportunity to reflect on the development process and past 10 years of operation with the incumbent OPC-DA servers and apply the lessons learned to improving the development and operation of the forthcoming OPC-UA servers.

The development model followed previously for OPC-DA was for CERN to provide a requirements specification to each vendor who autonomously implemented and delivered a complete OPC-DA server component. Although this approach had the desired effect of incorporating expert hardware knowledge into the development process, it had the consequence of delivering three very different software components to CERN's control system chains, each component exposing its own unique behaviour for end-user installation, configuration and runtime operation. This inhomogeneity proved a costly model in terms of support.

An additional consequence of this approach was that each company delivered, initially anyway, black-box components whereby CERN had no access to the internal implementation code. In general, the CERN production environments contain large numbers of devices and, during data taking operations, these devices must all be controlled

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA009

and monitored simultaneously. This presents stringent requirements on the OPC layer in terms of effective handling of scale, load and concurrency. Unfortunately, however, hardware environments available to engineers developing and debugging the software at the company sites are often much smaller, and, therefore, insufficiently representative of a full production scale setup. Accordingly, defects observed in production related to scale, load and concurrency were often difficult for off-site engineers to recreate, diagnose and rectify. This resulted in long periods between faults reported in production and the corresponding fixes being made available. Over time this situation was regarded as untenable; source code access was eventually granted to CERN for all three OPC-DA servers. Note, however, even with the source code, support remained complex: the three servers were built without reference to one another, each, therefore, having its own unique internal architecture, implementation and developer-orientated diagnostic metrics.

A further consequence was that only the companies were able to provide server release versions. Users at CERN had to wait for server releases to be provided by the manufacturers, even for relatively trivial fixes.

A key observation from OPC-DA server operations was that a successful migration needs a mix of CERN's experience developing for and running large scale operations and the individual vendor's expert knowledge of the details of integrating hardware from their product catalogue; i.e. a collaborative approach was needed. Following on from the observed impedance resulting from the black-box mode of development and delivery, it was also clear that a collaboration down to the level of co-owned source code would be beneficial. Both middleware OPC layer related code and low level hardware access code must be accessible to all collaborating parties.

Another core lesson learned from developing and operating the OPC-DA servers was that homogenising the internal architecture and implementation was to be a top priority. Clearly the specifics of accessing the hardware are unique to each vendor, however, in broad terms, the OPC-UA servers all provide similar functionality, namely industrial standards based integration components to industrial power supplies.

TECHNICAL OUTLINE OF OPC-UA SERVERS

In line with the recommendation for homogeneity, the OPC-UA server implementations for CAEN, ISEG and Wiener hardware will adhere to the standard architecture depicted below in Figure 2.

OPC-UA has been identified as a strategic industrial standard middleware for CERN [4], and homogeneity and code re-use in OPC-UA servers was identified as a common requirement in a wider context than implementations for CAEN, ISEG and Wiener hardware. In response, CERN produced and maintains the quasar framework [5] – a cross platform, model driven, OPC-UA server development framework. The quasar framework provides commonly required functionality to developers straight out of the box. The server implementations for CAEN, Wiener and ISEG are based on quasar, delegating common tasks including schema based XML configuration, address space population, security, logging and publishing server meta-information to common code provided by the framework.

Using quasar tooling, the developer creates an XML defined model describing the required OPC-UA interface for each hardware class. Then, that model is fed in to internal framework scripts which use a collection of parameterized internal templates to generate the complete C++ code to support the OPC-UA client facing interface. The scripts also generate C++ code stubs which the server developer must complete themselves to handle device specific tasks including hardware connections, disconnections, reading data from and writing data to hardware and managing error states.

The unique details of how data is delivered to and retrieved from the hardware are encapsulated in a layer provided by the hardware manufacturers themselves – in terms of the generic architecture, this layer is referred to as the Hardware Access Layer (henceforth HAL). The different HAL implementations are generally provided in the form

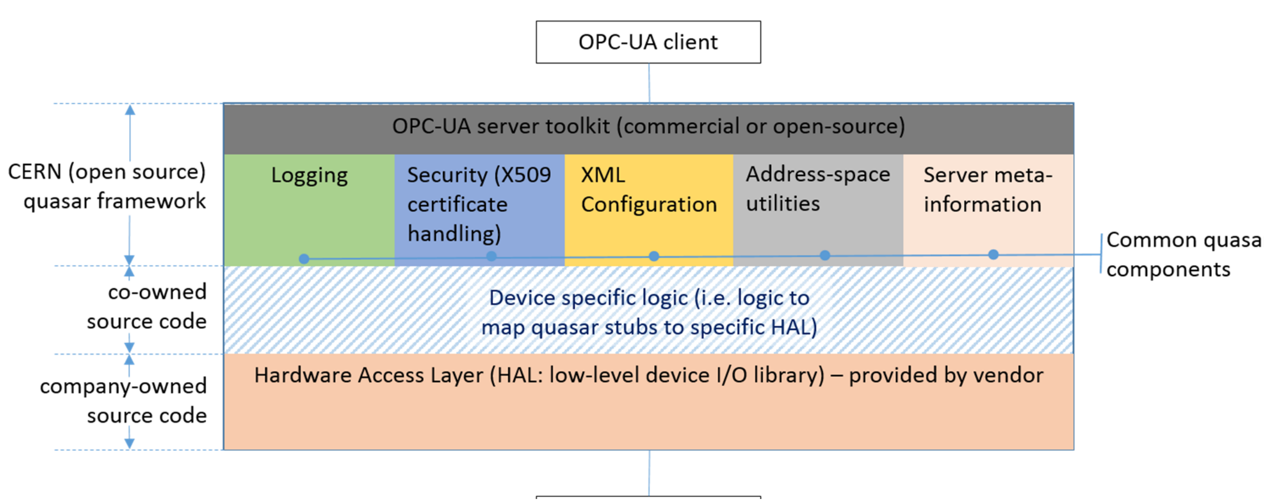

Figure 2: Components of the Internal Template Architecture for Industrial Power Supply OPC-UA Servers.

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA009

of a library with a clearly defined and well documented programmatic interface.

In general, the division of labour is that CERN is largely responsible for delivering and maintaining the implementations for the quasar generated stub handlers, whereas the manufacturers are largely responsible for delivering and maintaining the HALs. Manufacturer expertise, however, is required for tuning the stub handler code for optimal usage of the HAL API and resulting operations in device firmware (e.g. to provide expert advice on which device registers may be slow to respond; requests for such values therefore, may be less suitable for time critical code segments etc.)

In keeping with the shared-source-code guideline, both CERN and the manufacturers have full access to the source code in the OPC-UA layer (i.e. the quasar model and the user-defined code in the filled-in stubs) and the source code in the HAL. The source code for each of the three OPC-UA server implementation is shared via Git projects hosted on a cloud-based repository, with CERN and the manufacturers each having a single representative with full project administration rights.

COMMERCIAL PARTNER ORIENTATED OUTLINE

The requirement for homogeneity extends beyond technical considerations: A common model for setting up and managing the collaborative projects was also sought. Clearly, these projects involve transfer of intellectual property created at CERN into a commercial environment. Expert assistance in this field was provided by CERN's KT group. KT's mission is to promote and transfer technology and expertise developed at CERN, aiming to accelerate innovation and maximise the global positive impact of CERN on society. Each of the manufacturers is an independent commercial entity from a CERN member state, with a worldwide customer base. The terms and conditions of the projects and the final product (an industrial standards based integration component) must be beneficial for CERN's laboratory operations and a similarly motivating prospect from the companies' commercial perspectives, for example allowing them to innovate novel ways of satisfying customer requirements or to address new markets.

Open-source Licensing for the Quasar Framework

An initial task to address before starting any collaborations with the external companies was to make the quasar framework available outside CERN to assist any OPC-UA server development project, i.e. for quasar to become a fully open-source project. The goal, defined through discussion with the quasar development team, was to open-source quasar in such a way that the framework itself should maintain a strong 'copy left' open source licensing approach, motivated by the desire to permit users to freely run, study and develop with quasar and to encourage contributions from external users to help strengthen and grow the project. The framework itself should not become the subject of commercialisation. However, so as not to dissuade any parties from industry from using the framework to build servers, it must also be possible to apply proprietary licensing (including commercial licensing) to servers built using quasar.

With these goals in mind, a scheme was devised whereby quasar was licensed with a mixture of BSD licensing (for the core scripts) and LGPLv3 licensing (for the XML and code generation templates). This licensing approach provides the openness and free-for-use desired for the quasar framework itself, whilst allowing quasar users the freedom to license the products they build as they wish.

With the framework on which the CAEN, ISEG and Wiener servers were to be built satisfactorily open sourced, the foundations were prepared for describing a common model of the project with each company.

The Common Project Model

Working together with KT, CERN's common requirements of each collaboration projects were summarised as follows:

- A commitment from each company to provide a HAL implementation with a documented interface and to support the HAL for up to 15 years.
- CERN access to the HAL source code.
- A commitment to ensure the HAL remains compatible with the OPC-UA server for 15 years.
- Hardware expert opinion on optimal HAL usage in the OPC-UA layer quasar stub implementation code.
- Free use of the OPC-UA server for CERN and collaborating institutes.
- Up to 5 days of on-site hardware expert assistance during the initial rollout of the OPC-UA servers (if requested).

Accordingly, CERN's commitments and contributions to the collaboration projects were summarised as follows:

- Development and long-term maintenance of the quasar framework.
- Development and long-term maintenance of the implementation code to fill the quasar generated stub handlers.
- Access to CERN's large scale production environments (during testing) to validate HALs.
- User feedback for continuous product improvement.

With CERN's core requirements identified, there followed a series of meetings with each company, the project outline and requirements were communicated, discussed and negotiated. KT then drew up three sets of documents, each set comprising a pair of documents per vendor. The two documents were: a collaboration agreement encapsulating the common requirements above in addition to the needs and requests specific to each company and a joint ownership agreement defining the ownership and licensing terms of the specific vendor server, built using quasar and company code and expertise. The ultimate aim is to have three pairs of co-signed (CERN/company) agreements. These agreements form a clear and contractual background to the collaborative effort of engineers in CERN and the

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA009

companies involved in producing the technical product but also to provide a clear ownership structure for further exploitation by the companies for the respective technical product.

Each company agreed in principal that having a modern standards based integration component to offer existing and future customers would be desirable, however, each also raised concerns during the discussion and negotiation phase, common points included: the need to constrain the cost of in-house engineering development time and long-term support commitments; how engineering cost expended during development could be recouped, and indeed how to capitalise further on these efforts over the longer term; each company was also concerned that a software proxy to their hardware should uphold the respective reputations for quality and performance associated with their brand.

These points and their respective resolutions were discussed and written in to the two collaboration agreement documents as follows:

- The collaboration agreement document explicitly describes the project goal and details the contributions each party must make, including finite development and testing effort and long-term support exposure.
- The joint ownership agreement document describes the rights each party has as regards licensing the resulting OPC-UA server implementation. CERN has perpetual, royalty free rights and is free to grant similar rights to third parties for the execution of the CERN mission. Each company may license the server as they see fit with other third parties, including licensing options for commercially exploitation.

CURRENT STATUS AND TIMELINE

The three processes, regarding reaching a collaboration agreement and implementing a technical solution are at different stages: at time of writing the current status is:

CAEN

Collaboration agreement documents agreed and signed by CAEN and CERN.

OPC-UA server in advanced beta stage for Linux, pre-production testing carried out in the ATLAS experiment look promising, no major defects observed so far. No MS Windows version available yet.

ISEG

Collaboration agreement documents agreed and signed by ISEG and CERN.

OPC-UA server in advanced beta stage for Linux, pre-production testing carried out in the ALICE and ATLAS experiments look promising, no major defects observed so far. No MS Windows version available yet.

Wiener

Collaboration agreement in the process of negotiating acceptable levels of development and support commitment between Wiener and CERN.

OPC-UA server a proof-of-concept Wiener HAL interface and implementation developed by CERN has been provided to Wiener for review. A skeletal implementation based on this HAL interface is in development for MS Windows, however, no further effort is likely to be expended before the collaboration agreement is finalised.

CONCLUSION

Drawing on key lessons from operational and support experience with the old OPC-DA servers, the projects to deliver the new OPC-UA server implementations have been built with a focus on homogeneity. This homogeneity extends to both the technical aspect and the process of collaborating with commercial partners. KT's assistance in creating legal and contractual foundations for these generic processes has been critical. During the current development phases and the subsequent operation and maintenance phase, this commonality is expected to reduce the complexity of the support task and provide an improved service to the end-users through:

- A well-defined collaborative project approach, homogenous across all companies.
- A consistent end-user interface for installation, configuration and operation.
- Server implementations based on template architecture and common code re-use, with expert device integration knowledge encapsulated in HALs.
- Full-stack source code transparency for efficient on-site fault diagnosis.
- Reduced duration of fix to release cycles.

Furthermore, as supported by the licensing scheme, other research institutes have already approached CERN to express an interest in using these OPC-UA servers to integrate commercial power supplies in to their control systems. Feedback, and source code contributions, gathered from these users will be valuable to help harden and genericise the servers to better handle a wide range of equipment configurations and diverse operational constraints.

REFERENCES

[1] The OPC Foundation, "OPC Classic", http://opcfoundation.org/opc-classic/

[2] The OPC Foundation, "OPC Unified Architecture", http://opcfoundation.org/opc-ua/

[3] B. Farnham, R. Barillère, "MIGRATION FROM OPC-DA TO OPC-UA", in *Proc. 13th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'11)*, Grenoble, France, Oct. 2011.

[4] P. Nikiel, B. Farnham, S. Franz, S. Schlenker, H. Boterenbrood, V. Filimonov "OPC Unified Architecture within the Control System of the ATLAS Experiment", in *Proc. 14th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'13)*, San Francisco, CA, USA, Oct. 2013.

[5] S. Schlenker, D. Abalo Miron, B. Farnham, V. Filimonov, P. Nikiel, C. Soare "quasar - A Generic Framework for Rapid Development of OPC UA Servers", in *Proc. 13th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015.