# JCOP EXPERIENCE WITH A COMMERCIAL SCADA PRODUCT, PVSS

P. C. Burkimsher, CERN, Geneva, Switzerland

*Abstract*

In 1999, the four CERN Large Hadron Collider (LHC) experiments together chose a commercial Supervisory Control And Data Acquisition (SCADA) tool, PVSS, with which to construct their control systems. In this paper we review the functionality PVSS offers for building control systems. We show how PVSS has lived up to expectations (or not!) in its first four years of use at CERN and extrapolate to some important general considerations when adopting a commercial tool. This paper does not present a detailed technical performance report of PVSS but rather abstracts significant items of experience gained from using this particular SCADA tool. This generalized summary is relevant to potential users about to choose a commercial tool from any manufacturer. The paper considers the much-vaunted functional benefits of using commercial SCADA: Scalability, Extensibility, Ease of use, Stability, Cross-platform, Graphical User Interface, Archiving, Trending, Web-ability, Remote Access, Security, Alarms and finally Documentation. We identify the strengths and weaknesses encountered with PVSS in these areas. The paper further outlines our experience with the realities of dealing with a commercial company. Inbuilt, unspoken, even sub-conscious assumptions on the part of the SCADA manufacturer are brought into the spotlight and their important implications are described. We conclude on a positive note, referring to the recommendations of the External Review Committee which reinforced the Joint Controls Project's approach given the stringent financial imperatives under which the LHC experiments and CERN itself are currently operating.

## INTRODUCTION

The LHC experiments face a dilemma. Like the LHC accelerator team, they have to construct a control system for their equipment that will endure for up to 25 years. Unlike the accelerator controls group, the experiments are not staffed by career engineers. Rather, the work will be performed by people whose main interest is physics. How then to produce a quality control system to do the job? The experiments decided to build a control system using an industrial toolkit, thus avoiding the heavy investment of having to build the infrastructure themselves. Under the auspices of the Joint Controls Project (JCOP) [1] the four LHC experiments decided to use a common toolkit. Buying any piece of industrial software is rather like getting married. Although many aspects of your purchase are advertised as "component features", in practice they all arrive together in a single package. You purchase the good components along with the bad. Only afterwards do you find out which are really which. Rather like a dream wife (or husband) who fulfils all your criteria until you get to bed and find that (s)he snores. In this paper, we outline what really happened when the LHC experiments took a wife. Is this "marriage" working? Has PVSS lived up to expectations?

## THE OBJECTIVE CHOICE OF PRODUCT

Daneels and Salter gave a full-length description of SCADA in their paper "What is SCADA?" at ICALEPCS 1999 [2]. To quote: "As the name indicates, a SCADA system is not a full control system, but rather focuses on the supervisory level. As such it is a purely software package that is positioned on top of the hardware to which it is interfaced." Naturally, a SCADA toolkit is a toolkit for building SCADA systems. In a second paper, Daneels and Salter described how forty different products were evaluated against objective criteria drawn up by the experiments [3]. This lengthy procedure was designed to eliminate subjective bias from the final choice. Thus the products that were already familiar to several institutes were analysed with the same rigour as new products. According to these objective criteria, PVSS (Prozessvisualisierungs- und Steuerungs-System) from the Austrian company ETM [4] was chosen.

## EXPERIENCE TO DATE

Although PVSS was initially chosen by the LHC experiments, it has since been adopted more generally at CERN and used successfully in several fixed target experiments – COMPASS [5], HARP [6] and NA60 [7]. In addition to being used in the LHC test-beams, it has also been adopted by the cryogenic and vacuum groups in particle accelerator control. Approximately 150 people have followed a PVSS course and around 500 machines are currently licensed at CERN or in the participating institutes. We thus have a growing body of expertise and real experience in production systems. By the time the LHC proper comes on-stream, PVSS will be tried and trusted software.

## THE CRITERIA

The following paragraphs describe the principal objective criteria for which PVSS was chosen and show how the product has fared in practice.

### *Scalability*

The issue of scalability is probably the most fundamental item in the list of LHC experiment requirements. Many commercial systems as well as independently developed systems use "tag" technology, i.e. a flat namespace of process variables and control parameters. PVSS was specifically chosen because it is device-oriented, meaning that the namespace is structured. Consequently one can create and manipulate complex devices (such as high voltage power supplies) at a level of abstraction that is manageable. This has already allowed CERN to develop systems more quickly, and with fewer errors, using PVSS than would have been possible using a tag based system.

Some SCADA tools that were reviewed had built-in limits, e.g. 64,000 maximum items. It is very important to discover if such limits exist when choosing a SCADA product as often the limits are designed-in and cannot be increased. The PVSS design is open in this respect and we are currently not aware of any software limits that will cause problems using the present generation of 32-bit hardware in the detector control systems.

Size matters. Beyond the limits of a single system, even a device-oriented system, the LHC experiments need to be able to expand further. Some SCADA tools are very centralized, but PVSS is not. The PVSS architecture is one of communicating processes and can take advantage of multi-CPU systems. Those systems may in turn be further distributed across many machines, producing a federation of cooperating systems. During investigations at CERN, an unspoken assumption of the manufacturer was exposed. ETM had assumed that a distributed system would always be fully interconnected, that is all member systems would have direct connections to all other members. At CERN, a “large scale” system could have ~100 component subsystems spread across even more machines and the question was raised as to whether PVSS would scale sufficiently. CERN has performed tests in this area to confirm that the ETM implementation does indeed scale to this level [8]. One hundred and twenty-one systems were successfully connected in a tree of valency 3 with 5 levels (Fig. 1).

Every system (node) in the tree was directly connected to all of its ancestor nodes higher up the tree. Thus the top-most system had a direct connection to all other nodes in the tree. To ensure a realistic test, some of the systems in the configuration ran Windows and some ran Linux. The cross-platform nature of PVSS and the connectivity capability for a large number of systems was amply demonstrated.

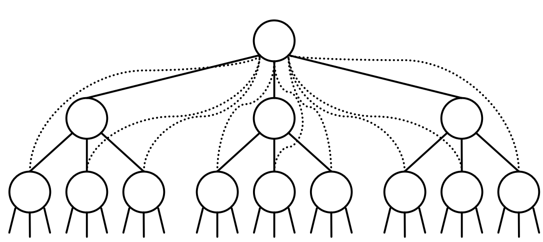

Fig 1. Distributed system test showing 3 of the 5 levels.

Another unspoken assumption on behalf of the company initially concerned the static nature of resultant systems. Before CERN chose PVSS, ETM’s customers had all had relatively static systems, meaning that the number and range of devices being controlled was defined at the outset and very rarely changed. This is not the case in a High Energy Physics (HEP) experiment where equipment changes at least annually and sometimes more frequently! The technical consequence of this mindset was that, in the original PVSS, all systems had to have an identical, local copy of all data type definitions. Any change in any system’s type definitions would involve a shutdown and resynchronization of all participating systems. From the outset ETM agreed to address this issue and have done so. The upgraded scheme now means that PVSS is ideally suited to a physics experiment built from subsystems where those subsystems are both introduced to the cluster, and then later modified, in an independent manner.

## *Extensibility and Ease of use*

An extremely important criterion was that systems built with the toolkit should not be constrained by it. In other words, it should be possible to extend the toolkit in any areas where the toolkit as delivered was less than ideal for HEP. Another criterion was that it should be easy to use. These two requirements often conflict, of course. The more powerful a tool is, the more training will be needed to use it effectively. Some users have required the sophistication and have been pleased to be able to exploit it. Other users wanting to perform more trivial tasks have felt that PVSS was complicated. Under JCOP, the experiments and IT-CO (the controls group in IT division at CERN) were able to collaborate and assist in this regard. They identified certain HEP-specific functionality that was not initially provided in the product. Functionality that was going to be needed by all four experiments was then earmarked for implementation in a modular package that became known as the Framework [9]. There are many Framework components, but typical examples are:

- CAEN crate functionality. These high voltage crates are widely used in HEP, in particular across the LHC experiments. Drag-and-drop CAEN crate functionality was of great interest to all concerned.
- Finite State Machine (FSM) functionality. The ability to model a detector’s run-time state is a vital feature required in all four experiments.

A part of the extensibility of PVSS comes from its programming language, called CTRL (pronounced “control”). CTRL is very close indeed to the programming language ‘C’. An application developer who knows ‘C’ can start using CTRL very quickly. Yes, he has to learn the programming model and the functions provided in the Application Programmer’s Interface (API), but at least he does not have to learn a whole new programming language at the same time. This was not the case for other SCADA products on the market and has been an appreciated feature of PVSS.

### *Stability*

Whatever tools matched the functionality requirements on paper, a very important consideration was that of stability. The initial detailed technical evaluations looked carefully at this aspect. Over the last 4 years the PVSS kernel has continued to work very reliably with users reporting kernel stability above their expectations. Over this period, PVSS has been extended in many areas, frequently at CERN's explicit request. Not surprisingly, these new additions have contained some bugs; nevertheless corresponding patches have been provided promptly. In general, ETM brings out a significant new version annually, with patches as requested in between. CERN feels that this frequency is about right, given the rate at which the product is currently being developed.

### *Cross-platform*

A clear requirement from the HEP community was to have a cross-platform Windows and Linux system. Several very good products were excluded from adoption because of their Windows-only implementation. One must be careful about defining what "cross-platform" means. To some it means "I want the best of Windows and the best of Linux". More cynical folk accept that cross-platform usually means "I'll get the lowest common denominator between Windows and Linux". PVSS actually manages to achieve the better of these options. PVSS does not merely allow you to develop systems that will run on Windows or Linux. We have configurations that show that a truly mixed platform is possible - we have examples of Windows systems inter-working with Linux systems. The experiments are thus able to take advantage of the better performance observed under Linux as well as the more comfortable user interface experienced under Windows.

### *Graphical User Interface flexibility*

When PVSS was first delivered, its development environment looked like Linux, irrespective of the platform. A Windows look-and-feel development environment was already under production and was in fact delivered shortly afterwards. Fundamentally, panels developed under either "look-and-feel" are PVSS panels that can run on either platform. Inevitably, however, these panels run a restricted subset of widgets available on both platforms. Nevertheless, Windows users can incorporate ActiveX components into their panels and thus have access to the rich and varied interfaces available using these items. Of course, such panels will not at present run properly on Linux. It is still an open issue as to whether the experiments will formally decide to sanction the use of such platform-specific components. In practice, as it looks as if all systems will be mixed-platform, there is no reason not to go this way. PVSS is amenable to whatever the experiments decide.

### *Archiving*

PVSS offers sophisticated archiving tools. Complex value archive arrangements can be configured via panels. 'C' language programming is not required for this! Nevertheless, we have observed that the sheer range of options has overwhelmed some users. PVSS cannot, though, be criticised for being restrictive.

There have been concerns expressed regarding access to archived data. ETM had assumed that access would always be through the PVSS API. Hence the file format behind the API was not released. This "black box" approach left ETM free to optimise the files for best performance, but meant that some level of PVSS had to be installed on any machine wanting to read the archives. At CERN, where the archives need to be read by independent programs, possibly many years after they were first recorded, this was a matter of concern. The company has responded positively to these worries and is making an interface to record archive data in an Oracle database. It has furthermore addressed an issue that made it inconvenient to access archive data from any machine other than that where it was initially recorded.

### *Trending*

The purpose of archiving data is usually to view it again afterwards in value-over-time plots known as trends. PVSS provides basic trending facilities, but there have been many philosophical discussions as physicists' expectations are very high. PVSS is not primarily a data presentation package in the sense of PAW [10] and achieving the right balance is an open discussion. Nevertheless, the company has responded favourably to our requests and, for example, have upgraded their trend charts with logarithmic axes capability. Another common request is to have multiple plots on a single panel. IT-CO, on behalf of JCOP, has taken advantage of the extensible nature of PVSS and, working within its existing features, has successfully provided this additional functionality in an easy-to-use form to end-users. This feature is another of the facilities provided in the JCOP Framework.

### *Web-ability and Remote Access*

One of the original requirements was that control systems built with the tool adopted by JCOP should be accessible using a Web browser, and hence accessible from remote locations. Only one product that was reviewed used Web pages as its main GUI, and this product was not selected for other reasons. All other Web-capable products assumed that you would implement panels twice, once for the "local" GUI and again for the Web version, always for reasons of performance. PVSS initially fell into this category and has always provided the necessary hooks. It is nevertheless highly undesirable to implement and maintain two sets of panels this way. ETM have had long discussions with CERN and their other customers and are currently proposing a scheme whereby the proprietary GUI gets downloaded when you open a Web page. This solution does restrict the functionality to Web browsers running on Windows or Linux on Intel systems, but has several advantages: the user does not need to pre-install PVSS; he is presented with a familiar screen and the implementation is the standard PVSS GUI, offering speed benefits over other, usually interpreted, solutions. Again the company has worked hard with CERN and their other customers to find an acceptable compromise. There remain a host of social and security questions concerning who should be allowed to do what and from where, but PVSS is technically able to support whatever is decided upon.

### *Security*

The topic of remote access immediately raises questions about security and protection from accidental access or even malevolent intruders. To date, worldwide, almost all control systems that have been built have been assumed to run on dedicated isolated networks. There has been no need for specific security protection. Suddenly, an Internet connection is becoming a must, whether to download updates or even just to reset the clock via NTP. Secure links between the different parts of the control system, foreseen in the LHC experiment requirements of four years ago, are now urgently needed. The same need is surfacing more generally in industry at large. The popular magazine New Scientist has carried several articles [11, 12] recently on the dangers posed by ignoring this topic.

### *Alarm Handling*

An integral part of a SCADA system is its alarm handling capability. PVSS has an extremely flexible alarm definition scheme. It does not come with a predefined alarm model, but allows you to configure your own. This flexibility does mean that first-time users have to learn a little more and become familiar with their experiment's alarm model before they can define their alarms, but this is the only way to offer the flexibility demanded.

Can the alarm system scale? We refer again to the model of communicating independent subsystems. Alarms are dealt with locally. PVSS offers a concept of summary alarms, known as "group alerts" which can be propagated up a hierarchy of systems.

The LHC Alarm Service Project at CERN chose not to use PVSS for their alarm handling [13]. Why? In outline, there were 3 reasons:

- Their alarms were not datapoint alarms in the PVSS sense
- They explicitly wanted to centralize incoming alarm states rather than use PVSS's distributed model.
- They needed alarm grouping and reduction functionality not natively provided in PVSS and preferred to use a different solution rather than extend PVSS in this way.

In short, their application did not fit naturally into a SCADA solution. On the contrary, the PVSS capabilities do map well onto the LHC experiment sub-system architecture and it is turning out to be very appropriate for this.

### *Documentation*

If PVSS has been weak in any area, it is in the area of documentation. PVSS has been historically used by control system experts working directly for ETM, or sometimes for third-parties, implementing control system applications. The documentation that was provided was reference documentation only, with no introductory User Guide. This certainly caused problems, as not all users at outlying institutes were able to attend the courses held at CERN. A reference manual is usually insufficient to get started. Secondly, in making a contract with CERN, ETM was making its first steps into the English language market. Unfortunately not all the documentation was originally available in English. Although CERN was able to negotiate an attractive financial agreement with the company, the IT-CO helpline at CERN fielded many calls that would have been avoided with better documentation at the outset. Although the documentation is not yet perfect, each new release brings a significant improvement. ETM is demonstrably working on this aspect.

## HUMAN ASPECTS

### *The users*

The users of PVSS are not technical attributes of any software product, but they are a key element in the success of its adoption. The human element should never be overlooked when discussing technical systems. The PVSS users from the LHC experiments are physicists who are not interested in becoming controls engineers. "I don't want to be a PVSS expert, I just want to control my XYZ" was an often heard refrain. There was clearly a lack of communication, as many newcomers to PVSS thought that PVSS was either a ready-made control system, or that, by using it, a control system could be created with a few simple mouse clicks. This expectation was clearly at odds with reality.

### *Relations with the Company*

IT-CO relations with the company have been, and remain, excellent. As the previous paragraphs have shown, the company listens and responds. Within its capabilities, it tries to respond positively. CERN is not ETM's only customer and sometimes priorities differ, but quarterly meetings keep both sides informed. CERN currently has a first-line support contract with ETM which provides more direct contact with the developers. An on-site support engineer provides CERN with immediate quantifiable technical assistance and also feeds back less quantifiable "vibes", such as mood and urgency to the company. The on-site person is available for consultation anywhere on the CERN site and IT-CO has daily email contact with the office in Austria. Furthermore, ETM developers (not just managers) have visited CERN, at our invitation, on more than one occasion. The provision of support from a commercial company is more than just a telephone number or an online helpdesk. ETM have become involved in ascertaining CERN's real needs. We have presented them with sometimes lengthy lists of feature requests and the company has made many changes, some of them major, as a result.

We have been fortunate in having a partner such as ETM. Although the original product met the basic criteria, the evolved product fulfils the requirements more fully and more conveniently for the user. The global lesson to draw is that when you choose a software product you are not just making a technical choice. Clearly, the basic technical criteria must fit, but the right choice of company can have a profound influence on the outcome of your project. Does the company listen? A company that is too big and remote may not be very responsive. A company that is too small may not have the resources to respond, or may one day become extinct. If the sexy software is the bride, the company is her family. You get married not just to the bride, but also to her family. If you contemplate marriage, you should take a long, careful look at your prospective mother-in-law! Fortunately we have seen that the close relationship between CERN and ETM has resulted in real, tangible benefits to our end-users.

## THE FUTURE

Future enhancements currently under discussion with ETM include histogramming, authentication, efficient alarm handling in very large systems and further improvements in documentation. Other non-PVSS specific issues which need to be followed up include managing component-ware inter-dependencies and the administration of large clusters of systems.

## SUMMARY AND CONCLUSION

This paper has reviewed the experience at CERN using the SCADA product PVSS. This established product from a medium sized company has already been used in several production systems at CERN providing continuity in their development environment and leverage to free-up personnel to work on their end-user applications. Although no product is perfect, excellent relations with the PVSS manufacturer have meant that the company has been actively and successfully involved in addressing with us all of the important issues encountered at CERN.

The JCOP activity at CERN was the subject of an independent external review in Spring 2003. The External Review Committee, including members from outside institutes, unanimously endorsed the JCOP approach, the choice of ETM as supplier and the PVSS product itself.

Successful marriages may lack romance. But a real working relationship endures.

## REFERENCES


[1] CERN Joint Controls Project, JCOP http://cern.ch/itco/Projects-Services/JCOP/welcome.html

[2] "What is SCADA?" A. Daneels, W. Salter, CERN, Geneva, Switzerland http://www.elettra.trieste.it/ICALEPCS99/proceedings/papers/mc1i01.pdf

[3] "Selection and Evaluation of Commercial SCADA Systems for the Controls of the CERN LHC Experiments." A. Daneels, W. Salter, CERN, Geneva, Switzerland http://www.elettra.trieste.it/ICALEPCS99/proceedings/papers/ta2o01.pdf

[4] ETM Corporation http://www.etm.at/english/index.htm

[5] COMPASS Experiment, CERN http://wwwcompass.cern.ch/

[6] HARP Experiment, CERN http://cern.ch/harp

[7] NA60 Experiment, CERN http://cern.ch/na60

[8] "Scaling Up PVSS", CERN http://cern.ch/itcobe/Projects/ScalingUpPVSS/welcome.html

[9] The JCOP Framework Project, CERN http://cern.ch/itcobe/Projects/Framework/welcome.html

[10] Physics Analysis Workstation (PAW), CERN http://cern.ch/paw

[11] "Red alert on the e-war front." Duncan Graham-Rowe, New Scientist, 5 July 2003

[12] "Electricity grids left wide open to hackers", Duncan Graham-Rowe, New Scientist, 30 August 2003

[13] The LHC Alarm Service Project, CERN http://cern.ch/ab-co-tech-committee/PROJECTS/LHC ALARMS/AB_CO_TC_LASER_7_11_02.pdf