

## BESSY CONTROL SYSTEM ADMINISTRATION AND ANALYSIS TOOLS\*

R. Lange, T. Birke, D. Faulbaum, D. Herrendörfer, P. Stange, BESSY GmbH, Berlin, Germany  
M. Giacchini, INFN/LNL, Legnaro, Padova, Italy

### *Abstract*

The BESSY II storage ring has been in continuous operation for more than eight years. With BESSY's new projects and facilities requiring an increasing amount of attention, maintaining the BESSY II control system's high reliability and availability with less effort has become an important task. A set of professional non-commercial tools has been introduced. These tools are combined to detect and track errors, support system administration, and allow an efficient analysis and maintenance of the control system's hardware, software, and configuration.

The paper discusses problems and difficulties encountered, and presents the selected approaches and tools as well as the gained experiences.

### INTRODUCTION

Regular user operation of the BESSY II storage ring began in January 1999. Starting from the early stages of the BESSY II construction, during testing and commissioning of the accelerator, and through the more than 8 years of operation, its control system, which is based on the Experimental Physics and Industrial Control System (EPICS) toolkit and architecture, has provided a high level of reliability and availability [1,2].

### CHANGES OVER THE YEARS

#### *Knowledge Deterioration*

During the eight years of operation, a number of colleagues of the control system group have moved to new tasks at private companies, others were retired. For any given BESSY II control system application, the number of people with intimate knowledge of the application design and implementation has been shrinking. Some applications have to be considered as orphans, some are even completely forgotten.

#### *Hardware Deterioration*

After the initial phase of early drop-outs, the control system hardware was running reliably with a small number of failures for some years. Now, ageing electrical parts create an increasing problem. Large power supplies break down at a rising rate, any (scheduled or unplanned) power outage leads to a high number of device failures caused by failing small switching power supplies.

#### *Time Limitation*

Similar to most other accelerator facilities, BESSY is operated in 24/7 shifts. But – marking an important difference – the BESSY II operators and floor managers are not part of dedicated operations teams, as these tasks are distributed widely over most technical and physics

group members. Thus, all full-time employees of the control system group work operator and stand-by shifts, which limits their time for “regular” work. Single night shifts create an additional availability issue, as colleagues are not working regular hours for two consecutive days.

New projects, e.g. the design and implementation of the MLS control system [3], require an ever increasing amount of time and attention. Parts of the controls group have been highly involved in such activities, leaving only a minimal part of their time for BESSY II system maintenance tasks.

### THE MAINTENANCE CHALLENGE

Over the years, these effects were adding up towards a situation where a smaller number of people with significantly less time was facing an increasing number of problems. It was soon becoming clear that steps had to be taken to support the controls group maintenance procedures and minimize the time needed to detect and analyze problems. Within a year, a number of tools – all non-commercial – have been introduced to create support for the most important aspects.

### REVERSE ENGINEERING

Legacy applications are hard to change and debug, even more when the original author is not available anymore. While links within a medium size EPICS database can still be followed using a current database configuration tool, links from operator panels and correlations with configuration files of other applications are very hard to track.

#### *Irmis*

IRMIS is a collaborative effort between several EPICS sites to build a common relational database (RDB) schema and a set of tools to populate and search an RDB that contains information about the operational EPICS Input/Output controllers (IOCs) installed at that site. It allows a reverse engineering process of the complete existing EPICS installation, regardless of the particular applications that were used to create that installation. An intuitive web-based interface allows browsing the system representation, following links from the EPICS database records to other records, panels, tool configurations, and archiver entries.

BESSY runs a basic IRMIS installation that can be used by developers during problem analysis and solving.

### PROBLEM TRACKING

When a problem occurs, which can naturally happen any time on any day, the operator on shift may not know the responsible person within the appropriate support group. Even when the right support person is known, that person may be on vacation, working off hours due to

\*Funded by the Bundesministerium für Bildung und Forschung (BMBF) and the Land Berlin

night shifts, or busy with other tasks. A problem report should always reach enough people (in most cases the complete support group), to ensure a prompt and adequate response.

### *Trac*

Trac is an enhanced wiki and issue tracking system written and maintained by Edgewall Software, originally developed for software development projects. Its web-based, minimalistic approach is aimed to help its users in their development, while staying out of the way by imposing as little as possible on the users' established processes and policies.

An "Operations Trac" instance was configured with 20 different system components, which roughly map to the BESSY organization of support groups. In addition to personal accounts for all operators and support people, generic accounts were created (one for each component) to facilitate sending problem reports to groups.

Issue analysis, comments by others, and the process of problem fixing are easily documented using Trac's comprehensive and highly adaptable web interface. This led to a high acceptance of the Trac tool within operators and support groups.

### *Mailman*

Mailman is free software – written by Barry Warsaw and others as part of the GNU project – for managing electronic mail discussion and e-newsletter lists. Mailman is web-integrated, making it easy for users to manage their accounts and for list owners to administer their lists. It supports built-in archiving, bounce processing, content filtering, digest delivery, spam filters, and more.

In BESSY's Mailman installation, a mailing list for each of the mentioned generic Trac accounts was created. The support group leader acts as administrator of the group's mailing list. Any Trac report for a certain component is sent over this support group mailing list to the appropriate group of people, to reach the responsible and available persons at all times.

## CONSOLE ACCESS AND LOGGING

Console access to EPICS front-end computers (IOCs) is crucial. The console output contains error and diagnostic messages, most debugging tools need read and write access to the IOC console. This is similarly true for soft IOCs (EPICS IOCs running as a process on a server machine).

VME-based IOCs use serial console connections to a commercial terminal server device, so that access to the console is made through telnet or ssh. Soft IOCs, on the other hand, use the standard input and output of their server process.

A system is needed that allows uniform access and logging for both types of IOCs, so that connecting to a console does not need a-priori knowledge of the IOC's type and location or host.

### *Screen*

Screen is a full-screen window manager – also written as part of the GNU project – that multiplexes a physical terminal between several processes. All windows run their programs completely independent of each other. Programs continue to run when their window is currently not visible and even when the whole screen session is detached from the users terminal.

The soft IOC server machine was configured to start each IOC using a standard Linux initialization script under a named Screen session. A special user (named after the soft IOC) can log into the server and attach to the running session any time. Ssh access was set up to forward a ssh login using a certain key directly to the Screen session running the soft IOC.

### *Conserver*

Conserver is an application that allows multiple users to watch a serial console at the same time. It can log the data, allows users to take write-access of a console (one at a time), and provides a variety of features to accentuate that basic functionality. Conserver will log all serial traffic, so users can go back and review why something crashed or look at changes (if done on the console). Console logs can be tied into a monitoring system (by watching the created log files).

Three Conserver servers have been set up to control access to and log all IOC consoles in a uniform way. A user on any machine only needs to know the name of the IOC: the command 'console <iocname>' will connect to the desired console. All console output is logged and made available through a simple web server interface.

## NETWORK ANALYSIS

Network hardware failures and ethernet interface speed misconfigurations have been an increasing source of problems. BESSY's network group has introduced two systems to monitor and display network configuration, as well as traffic and error statistics.

### *NeDi*

NeDi is a lightweight network management framework, which is based on a scheduled discovery, a SQL back end and a web-based user interface. It retrieves all kind of configuration and statistical information off switches, and displays where nodes are connected to the network, traffic and errors for single or groups of nodes over time, and lots of other information.

When using compatible and correctly configured network hardware, NeDi's discovery crawlers will discover the network and its nodes automatically, thus only needing a truly minimal configuration: the IP of one switch. That way, NeDi keeps its databases up-to-date through all changes in topology and nodes, eliminating the need to keep its configuration manually in-sync with the real system.

### *Cacti*

Cacti is a complete network graphing solution based on RRDTool's data storage and graphing functionality. Cacti

provides a fast poller, advanced graph templating, multiple data acquisition methods, and user management features. It features an intuitive, easy to use interface, that makes it easy to configure and use Cacti for LAN-sized installations up to complex networks with hundreds of devices.

BESSY uses Cacti as a second choice for similar tasks as NeDi. It does not provide the extensive discovery features of NeDi, and is used with manual configuration for areas using hardware not compatible with the NeDi discovery crawlers.

## HOST AND SERVICE MONITORING

There are a lot of machines and both application and infrastructure services, which are important if not crucial for the BESSY II operation, but not very easy and obvious to monitor, including:

- Web servers (e.g. the Trac service, the operator help desk and instruction wiki servers)
- Mail services (Trac report distribution)
- Disk usage
- Server and IOC health and CPU load
- NFS file system availability
- DNS (name) and NTP (time) services
- Archiving engines
- Oracle database instances
- Channel Access gateways connecting the BESSY networks.

A framework is needed to regularly check these services and the machines they are hosted on.

### *Nagios*

Nagios is a host and service monitor designed to detect network and server problems as early as possible: before clients, end-users or managers do. It has been designed to run under Linux, but works under most Unix variants as well. The monitoring daemon runs intermittent checks on hosts and services specified using external "plug-ins", which return status information to Nagios. When problems are encountered, the daemon can send notifications in a variety of different ways (email, instant message, SMS, etc.). Current status information, historical logs, and reports can be accessed via a web browser.

BESSY started with a basic installation of Nagios, checking disk and health status of all relevant machines. Checks for the other services are added as needed, with the guideline of a new check always being added when the existing checks fail to detect something important.

A Nagios plug-in for Channel Access was developed at Legnaro, which allows for easy checks of IOC health and CPU load [4]. In combination with appropriate small EPICS databases, this plug-in also allows to check health of and access through Channel Access Gateways.

### *SEC*

SEC is a platform independent event correlation tool by Risto Vaarandi. It accepts input from regular files, named pipes, and standard input, and can thus be employed as an event correlator for any application that is able to write its

output events to a file stream. SEC is configured by rules, each rule specifying an event matching condition and an action list. Regular expressions, Perl subroutines, etc. are used for defining event matching conditions. SEC can produce output events by actions executing user-specified shell scripts or programs.

First tests with SEC in the BESSY networking group show that this tool will provide a means of watching the IOC console log files created by Conserver, to detect unsuccessful reboots as well as misconfigurations and hardware failures that print error messages on the IOC console.

## CONCLUSIONS

The newly introduced systems described in this paper have significantly helped the detection and analysis of problems occurring in the BESSY II control system. They provide operators, the control system group, and other support groups with useful and comprehensive tools and applications to efficiently help with detecting, reporting, debugging and solving problems.

## WEB LINKS

Irmis	<a href="http://www.aps.anl.gov/epics/irmis">www.aps.anl.gov/epics/irmis</a>
Trac	<a href="http://trac.edgewall.org">trac.edgewall.org</a>
Mailman	<a href="http://www.gnu.org/software/mailman">www.gnu.org/software/mailman</a>
Screen	<a href="http://www.gnu.org/software/screen">www.gnu.org/software/screen</a>
Conserver	<a href="http://www.conserver.com">www.conserver.com</a>
NeDi	<a href="http://www.nedi.ch">www.nedi.ch</a>
Cacti	<a href="http://cacti.net">cacti.net</a>
RRDTool	<a href="http://oss.oetiker.ch/rrdtool">oss.oetiker.ch/rrdtool</a>
Nagios	<a href="http://www.nagios.org">www.nagios.org</a>
SEC	<a href="http://www.estpak.ee/~risto/sec">www.estpak.ee/~risto/sec</a>

## REFERENCES

- [1] R. Bakker, J. Bergl, T. Birke, B. Kuner, B. Kuske, R. Lange, B. Martin, I. Müller, R. Müller, G. Pfeiffer, J. Rahn, H. Rüdiger, Rapidly Installable High Performance Control System Facilitates BESSY II Commissioning, Proc. EPAC 1998, 1676-1678 (1998).
- [2] J. Bergl, B. Kuner, R. Lange, I. Müller, R. Müller, G. Pfeiffer, J. Rahn, H. Rüdiger, Embedded Controllers, Field Bus, and a Modular IO Concept: Central Elements of BESSY II Controls, Proc. PAC 1997, 2493-2495 (1997).
- [3] R. Lange, T. Birke, R. Daum, S. Ehlert, D. Faulbaum, B. Franksen, R. Hartmann, B. Kuner, P. Laux, I. Müller, R. Müller, G. Pfeiffer, J. Rahn, H. Rüdiger, D. Thorn, Status of the MLS Control System, TPPB37, *this conference*.
- [4] M. Giacchini, R. Lange, N. Richter, LivEPICS: An EPICS Linux Live CD NAGIOS equipped, TPPA32, *this conference*.