ISBN: 978-3-95450-168-7 doi:10.18429/JACoW-IPAC2015-MOPHA046

!CHAOS STATUS AND EVOLUTION

C. Bisegni, S. Caschera, G. Di Pirro, L. G. Foggetta, R. Gargana, E. Gioscio, D. Maselli, G. Mazzitelli, A. Michelotti, R. Orrù, S. Pioli, F. Spagnoli, A. Stecchi, M. Tota, INFN-LNF, Frascati (RM), Italy, and C. Di Giulio, INFN-ROMA2, Rome, Italy

Abstract

A synthesis of the "!CHAOS: a cloud of controls" project and its application to accelerators and large experiments will be presented. We will describe here how the !CHAOS project has evolved from a candidate for the SuperB accelerator control system to a facility for IT distributed infrastructures. !CHAOS is currently, not only suitable for accelerators and large High Energy Physics (HEP) experiments, but is also applicable to other contexts, such as social and industrial applications. Preliminary results achieved on an accelerator use case are discussed.

INTRODUCTION

!CHAOS project started as candidate of Distributed Control Systems (DCS) and Slow Control of the SuperB accelerator and detector [1], with the ambition to be a prototype of innovative DCS standards. Today, the project is supported by the MIUR (Italian Ministry of University and Research) and INFN to develop a national infrastructure prototype of high performing services devoted to devices and polyfunctional sensors distributed on LAN and WAN [2,3].

!CHAOS exploits the new high performance web technologies by strongly increasing control system performances and services, preserving scalability, redundancy and reliability, as well as the abstraction of processes at any level obtained by:

- key/value database (KVDB) implemented as distributed memory object caching systems (DOC).
- using non-relational databases, in a specific document database.
- optimizing and embedding high performance e/o standard inter process communication and handling (RPC/direct I-O/Events).
- abstracting data structure through binary data serialization (BSON)
- embedding COTS and, in general, open hardware, in order to minimize costs and integration time
- making "Controls as a services" to be responsive to new trends of IT technologies and the request of industries and society.

Moreover, the !CHAOS DCS architecture embeds by design the Data Acquisition (DAQ) topology and object data concepts which make the project able to handle Big Analog Data and their Variety, Volume, Velocity, Value and Visibility. The scope of the project is therefore to study and realize a prototype of a dynamic, on-demand Cloud-based infrastructure able to handle the "five V" for data collected from analog devices.

Finally, the project foresees applications on different use cases: accelerator DCS, building automation and environmental controls, integration in embedded devices like CRIO National Instruments, BPM Libera, Beagle Bone, etc… In March 2015 a second test [4] of !CHAOS accelerator DCS has been performed at the Frascati Beam Test Facility (BTF) [5], successfully replacing the original DCS for the monitoring and control of the beam transport magnets by integrating at the same time BTF DAQ functionalities.

ACCELERATOR USE CASE @ BTF

The BTF is a beam transfer line, part of the DAFNE complex, optimized for the production of single electron/positron in the rage 25-700 MeV mainly used for HEP detectors testing and calibration [6]. The high intensity LINAC beam, intercepted by a pulsed magnet, is attenuated by Cu target and a set of scrapers. Six quadrupoles, two 45^0 bending magnets and horizontal/vertical slits permits to select beam quality, the proper momentum and multiplicity down to single particle (Fig. 1).

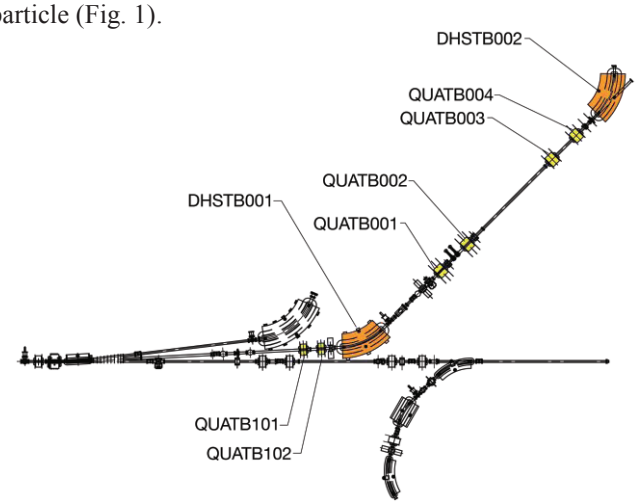

Figure 1: Beam Test Facility layout.

The test consisted in monitoring and controlling the transfer-line magnets and the facility DAQ devoted to acquire the x-y scintillating fiber system profile detector [7] and the glass lead calorimeters monitoring and acquiring single particle position and energy.

The scope of the test was to exploit the !CHAOS framework in a real operating condition, test the IT infrastructure and the Graphical User Interface (GUI), verify the stability and capability of the information service – Meta Data Service (MDS) and storing – the !CHAOS Data Service (CDS).

ISBN: 978-3-95450-168-7 doi:10.18429/JACoW-IPAC2015-MOPHA046

One of the objectives of this second run in BTF was to add DAQ functionalities to the DCS. The first challenge was to overcome the problem related to the BTF HW configuration, in particular the machines that interface the DAQ system were VME controllers quite old and running Linux 2.4 kernel with maximum 512MB RAM. This poor configuration did not allow to run a full-featured Control Unit (CU). For this purpose, we have added in this test release of !CHAOS the support for old or bar metal devices that cannot run a full CU/UI image. We built a very light layer completely decupled from the core layers of !CHAOS, that provides REST APIs to push/pull live data in/out from the !CHAOS cloud based infrastructure.

This is a key functionality that allows to access to any HW device/WEB browser/Application the !CHAOS cloud based resources in a lightly and simple way.

Control Unit (CU)

The CU environment consisted of 2 Linux 3.0 beagleBoards and 1 Ubuntu 12.4 virtual machine, running !CHAOS CU and controlling via MOXA the quadrupoles and diploes (QUATB001-004, QUATB0101-102, DHSTB001-002). The mixed front-end devices configuration has been chosen in order to test first the !CHAOS framework on virtual environments that previously showed issues related to job synchronization and timing; then test the !CHAOS framework on low cost Single Board Computers.

A !CHAOS light CU, running on VME-VMIC7740 Linux 2.4.21 acquired the beam profile and the calorimeter through two QDCs CAENV792N and CAENV965. The data flow and the software handshake were controlled through a scaler SIS3800 and CAENV513 I/O Register. The !CHAOS light CU was able to push data in the IT infrastructure via REST API, ensuring a throughput up to 3KHz for the 60 Byte of data of the BTF DAQ, largely compatible with the 50Hz repetition rate of the DAFNE LINAC beam.

In order to speed-up the test/debug/development and problem troubleshooting during the 2 weeks of testing, most of the real BTF DCS have been replicated in a simulated configuration where low level device drivers were replaced with virtual ones. This allowed, for instance to develop and test new GUI interfaces without interfering with real devices or to quick find and isolate bugs that often we do not know exactly at which level of the SW stack they are.

Framework

A large amount of work has been done on the framework architecture, providing its implementation and adaptation to the cloud infrastructure service, the implementation of direct I/O communication channels among system nodes for large data throughput and the implementation of CDS alpha version. The CDS is a session-less !CHAOS Data Pack management service. It achieves two different jobs. The first aims to inspect the direct I/O header looking for the final destination of the datapack (Live, History or either). The second aims to indicize all staged datapacks while they are migrated (and accurately organized) to the data area. Every single datapack brings the information needed by the CDS to know how it works on it. Thanks to this 'session-less' behaviour the CDS is a good candidate for a cloud environment.

The test at BTF was mainly intended to exploit CDS services concerning the data storing and indexing. For two weeks all CDS and backend services have been worked without problems. In this test each CDS was configured to do either job in two separate threads (CDS can also be configured to do only one job, this can help to scale the acquisition or indexing of the job separately).

IT Infrastructure (IT)

The IT group finalised the analysis of the requirement of dynamic, on-demand Cloud-based infrastructure, and implemented two static infrastructures (development & testing) based on the virtual machine infrastructure. The test at BTF aimed to exploit the behaviour of the static infrastructure and its tree layers:

- Access policy via OpenVPN;
- !CHAOS services behaviour: CDS, MDS, CU and UIs;
- Back-end services scaling and capability: MongoDB (document database) for indexing, CouchBase (Key-value In-Memory-Cache DB) for live value, CEPH (Distributed object storage and Posix filesystem).

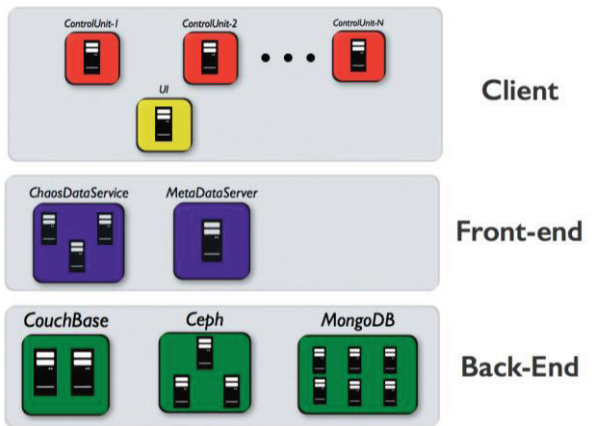

Figure 2: !CHAOS test & development IT infrastructure.

The !CHAOS test & development IT infrastructure (Fig. 2) is based on 26 nodes virtualized by KVM and managed by oVirt. All the VMs are distributed on 2 physical hosts, each equipped with two 6-core Intel(R) Xeon(R) CPU E5-2420 @ 1.90GHz connected each other by 10GbE. The servers storage resides on a Fibre Channel Storage Area Network capable of a cumulative peak of about 200MB/s. Essentially this is a test & development infrastructure and it is not designed for performance. The !CHAOS backend services are configured as follows. Indexing is achieved by a MongoDB cluster consisting of 3 shards each with 2 replicas and 1 arbiter. Live value resides on Couchbase 2-nodes cluster in replica mode. Historical storage lays on CEPH file system made by 2 OSD with replica-2 and 1 node for monitoring and

ISBN: 978-3-95450-168-7 doi:10.18429/JACoW-IPAC2015-MOPHA046

metadata, mounted as posix file system by FUSE. Infrastructure status and load has been monitored by one Ganglia server. All the hosts reside on a dedicated IP network in a dedicated VLAN accessible only through one OpenVPN server, a web portal has been developed to download OpenVPN configurations for users and external CUs that need to access the infrastructure.

During the BTF Run and other stress tests, have not been highlighted any overloads at each level of the infrastructure, at least while remaining below the available bandwidth of the 10 Gigabit Ethernet.

User Interface (UI)

Two user interfaces have been developed: one web-based and one with LabVIEW (Fig. 3). Both the interface access data via REST API, in order to test this functionality. An automatic slow feedback among beam energy and beam position has been implemented as simple application in LabVIEW.

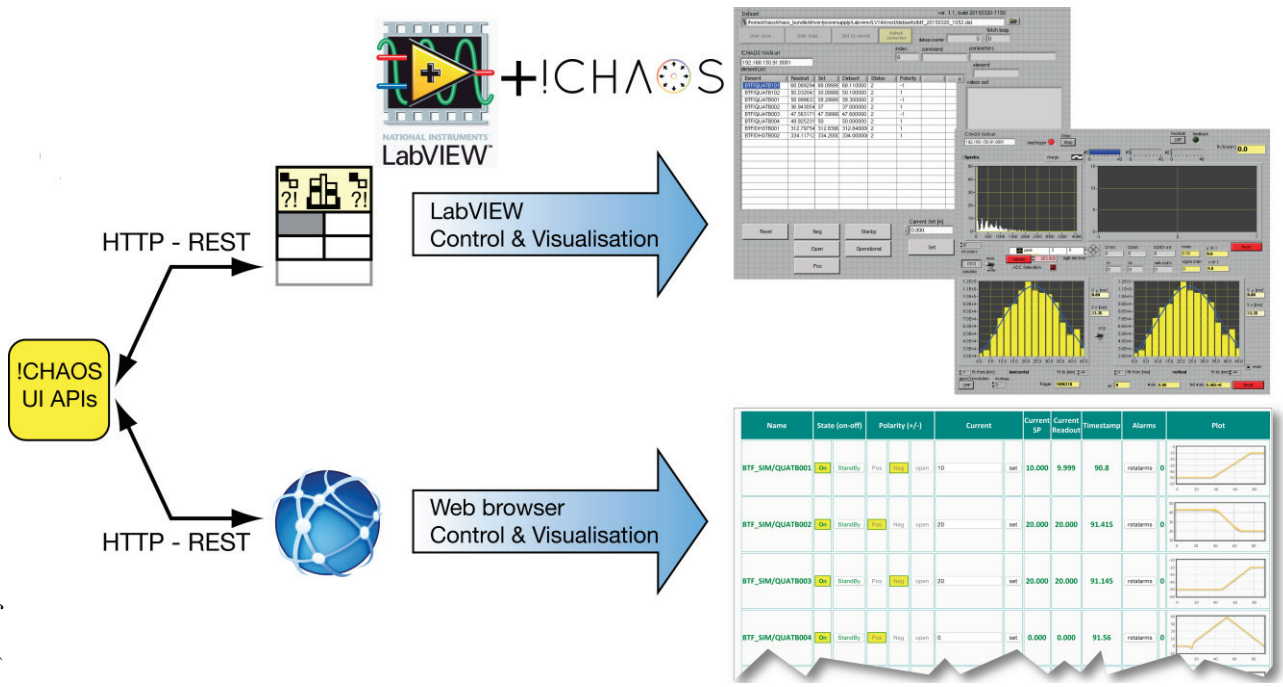

Figure 3: User interface tools.

CONCLUSION

Tests run on the accelerator enabled us to fully verify the functionality of the framework on a real case, as accelerator complex. Several preliminary stress tests of the system have been performed without showing limitations in terms of throughput and rate needed for managing the accelerator. During these tests the only limitation emerged was related to the Ethernet bandwidth (10 Gigabit/s).

Currently, the !CHAOS team is developing an independent benchmark to be able to qualify the framework and infrastructure performances.

Al the information concerning the project are available at http://chaos.infn.it/.

ACKNOWLEDGMENT

We thank the Computing Services of the LNF.

REFERENCES

[1] G. Mazzitelli et al, "High Performance WEB Applications for Particle Accelerator Control Sysytems", Proceedings of IPAC2011, San Sebastián, Spain; http://www.JACoW.org

[2] F. Antonucci et al, "!CHAOS: a cloud of controls - MIUR project proposal", INFN - 14-15/LNF; https://www.lnf.infn.it/sis/preprint/

[3] L. Catani et al, "Introducing a new paradigm for accelerators and large experimental apparatus control systems", Phys. Rev. ST Accel. Beams 15, 112804 – Published 29 November 2012.

[4] C. Bisegni et al, "First operational experience with the !CHAOS framework", Proceedings of PCAPAC 2014, Karlsruhe, Germany; http://www.JACoW.org

[5] G. Mazzitelli et al, "Commissioning of the DAFNE beam test facility", Nuclear Instruments and Methods in Physics Research A 515 (2003) 524–542.

[6] B. Buonuomo et al, "Performance and Upgrade of the DAFNE Beam Test Facility (BTF)" Nuclear Science, IEEE Transactions on (Volume:52 , Issue: 4), 2005.

[7] B. Buonuomo et al, "Diagnostics and Upgrade of the DAFNE Beam Test Facility (BTF)" Nuclear Physics B - Proceedings Supplements Volume 150, January 2006, Pages 362–365.