# THE CEBAF ELEMENT DATABASE AND RELATED OPERATIONAL SOFTWARE*

T. Larrieu, M. Joyce, M. Keesee, C. Slominski, R. Slominski, D. Turner
JLAB, Newport News, VA 23606, USA

*Abstract*

The newly commissioned 12GeV CEBAF accelerator relies on a flexible, scalable and comprehensive database to define the accelerator. This database, coined the CED (CEBAF Element Database) delivers the configuration for CEBAF operational tools, including hardware checkout, the downloadable optics model, control screens, and much more. The presentation will describe the flexible design of the CEBAF Element Database (CED), its features and assorted use case examples.

## BACKGROUND

The Jefferson Lab CEBAF accelerator is a superconducting recirculating linear accelerator capable of delivering continuous wave electron beams simultaneously to multiple experimental halls. Previously capable of delivering 6GeV, CEBAF recently underwent an extensive multi-year $338M upgrade project to double its energy capacity to 12GeV and add a fourth experimental hall (Fig. 1). The upgraded accelerator was commissioned successfully between November 2013 and May 2014. The CEBAF Element Database (CED) was an indispensable tool during the commissioning.

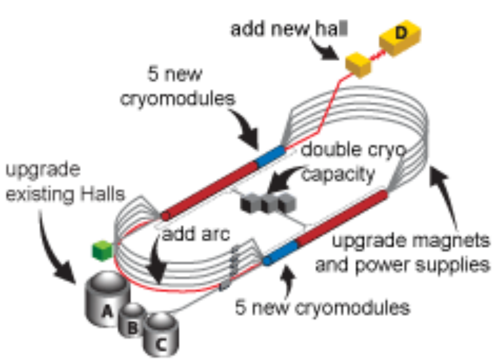


Figure 1: The CEBAF 12GeV Upgrade Project.

## DATABASE DESIGN

The CED design is based upon a modified Entity-Attribute-Value with Classes and Relationships (EAV/CR) data model [1]. Whereas in a traditional schema adding support for new accelerator hardware would involve adding additional tables and columns to the database, the EAV/CR data model employed by the CED is introspective – defining a new class of accelerator hardware in the CED simply involves adding rows to the already-existing metadata "catalog" tables. Once defined in the catalog, existing software is fully capable of interacting with the new entities after discovering their properties from the metadata tables.

A major benefit of the EAV/CR data model is that its static schema is well suited to integration with Oracle Workspace Manager [2] to provide timestamp-based versioning of table rows, named save points, and multiple independent workspaces (akin to branches in a software version control system).

* Notice: Authored by Jefferson Science Associates, LLC under U.S. DOE Contract No. DE-AC05-06OR23177. The U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce this manuscript for U.S. Government purposes.

## DATABASE IMPLEMENTATION

### *Layout*

To optimize performance for the different use cases, the CED database instance is distributed among three database users/schemas as shown in Figure 2. The operational schema stores the current machine configuration and is optimized for performance. The historical schema provides efficient storage and access to (read-only) historical save points (snapshots). And the development schema is used to create workspace branches where data can be edited and prepared before being promoted to the operational schema.

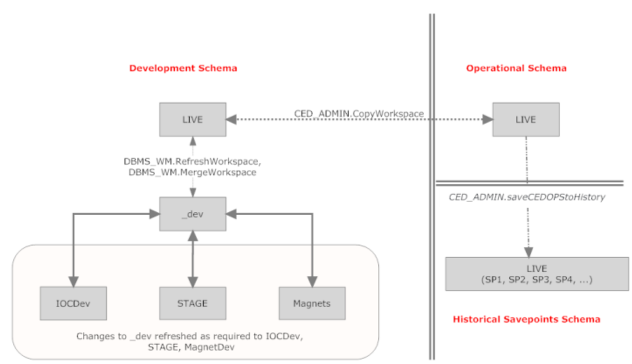


Figure 2: The CED three schema layout.

The details of the three schema layout are transparent to non-administrative users who need simply specify an optional workspace or savepoint name when interacting with the database.

### *Workspaces*

As a rule, users do not directly edit data in the operational schema. Instead, updates are prepared in a development workspace and merged into OPS upon request. A database administrator receiving a merge request will audit the proposed changes for validity and then promote the changes to the operational schema. Administrators make use of a web-based management tool to automate the auditing and merging process as well as the creation of new savepoints.

Most operational tools that utilize the CED default to using the OPS workspace, but allow the user to specify an alternate workspace. This capability allows an accelerator physicist, for example, to investigate changes to beam line layout or machine set points in a private sandbox. The official operational settings need not be altered in order to test tools against a what-if scenario.

### *API*

The generalized nature of the EAV/CR data model employed by CED requires a robust API to interpret the contents of the metadata tables and translate the abstract database storage into recognizable attributes useful to programs and users. The API is also responsible for enforcing much of the data validation and domain logic embodied in the metadata catalog.

For the CED, a shared library was written in C++ to be the sole intermediary for applications that will access its information. Native versions of the API are available not only to C++ programs, but also to scripts written in Perl, PHP, and Tcl. The script language versions are generated automatically from the original library via the open source SWIG (Simplified Wrapper and Interface Generator) tool [3,4].

On top of the core API, two general purpose user interfaces to the CED have also been built around the C++ library and its scripting language derivatives: a RESTful web interface and a full-featured command-line tool usable from JLAB Linux workstations.

### *CED Event Server*

The CED Event Server provides a means to link control system actions to database updates. Whenever rows are modified in the operational schema, a database trigger notifies the event server of the changes. Based upon rules in its configuration file, the event server can take actions including: telling long-running daemon processes to refresh their configurations, sending email to system owners, making logbook entries, or calling the ced2Epics utility to update process variables in the control system.

## HOT CHECKOUT

### *Readiness*

The immense scope of the work completed during the 18 months the lab was shut down necessitated a thorough system whereby all new, refurbished, and even supposedly unmodified systems could be verified as ready to be operated safely.

The CED components were used as the starting point for a web-based Hot Checkout (HCO) tool [5]. This tool tracked HCO readiness status on a daily basis in the months leading up commencement of beam operations. Geographic information from the CED granted users of the HCO tool an ability to determine readiness status region by region (Figure 3).


Figure 3: Web-based Hot Checkout Application.

### *CED-Linked Hardware Labels*

New hardware labelling was devised that utilizes the CED to provide detail that can't fit on the label itself. In addition to the key details such as name, weight, power supply, etc., the new labels contain a web QR code that can be scanned with camera-equipped tablets or smart phones to call a CED-linked web page. For example, the web page accessed from a Dipole label (see Figure 4) would allow the user to find out which other nearby hardware shares the same cooling water valves or power supply.


Figure 4: Magnet label with CED-linked QR code.

## SCREEN GENERATION

Screens used by CEBAF operators are built using the EPICS Control System Extensible Display Manager (EDM). For many large hardware systems, including magnets and vacuum, the process of building and maintaining screens by hand has been replaced by


ISBN 978-3-95450-168-7

programmatic screen generation. A large percentage of the screens used by CEBAF operators no longer exist as static files, but are generated on-the-fly using data from the CED. This process liberates EPICS programmers from the burden of creating screens for many types of new device instances and also ensures that the screens always match the current operational configuration.

In addition to the screens that are placed on menus for operators, a command line tool permits users to create task specific on-the-fly screens for many systems based on CED filters. For instance users may request screens that show non-contiguous regions or that group elements in a non-standard way (e.g. by power supply, by rack, by building, by controlling IOC, etc.).

## BEAM APPLICATIONS

### *elegant Download Tool (eDT)*

The ced2elegant and elegant Download Tool (eDT) [6] software tools were developed at JLab as the mechanism whereby the design set points stored in the CED can be used to perform accelerator set up. First, the ced2elegant utility is executed to retrieve design magnet setpoint values and geographic coordinates from the CED and produce a valid elegant [7] lattice. Then eDT is run to calculate new energy-specific settings by invoking the elegant accelerator modelling software to compute the momentum at each element in the lattice. At the end of the process, an output file is produced that can be used to load the set points directly into the control system. This process is faster and more robust than the previous method of scaling an historical machine save and relying upon operator tuning in order to achieve a stable orbit.

### *quadscan Utility (qsUtility)*

The qsUtility application includes data-gathering as well as analytical capabilities. The tool gathers emittance measurement data by altering the focussing strength of one or more quadrupole magnets in discrete steps while performing harp measurements using a rapid zigzag stroke technique. The details of each harp and matching quad that qsUtility uses to gather data are stored in the CED. In a manner similar to eDT, qsUtility can retrieve design parameters from the CED to create an input lattice and then execute Elegant to calculate desired quadrupole settings for a segment of the accelerator at a particular energy.

### *Linac Energy Management (LEM)*

The Linac Energy Management (LEM) tool (Figure 5) is used to calculate and optimize the energy gradient profile for the CEBAF accelerating RF cavities. From the CED, LEM retrieves configuration details of each cavity. The retrieved data for each cavity includes physical attributes such as its Q0 and also operational details such as whether it has been bypassed or is currently power-limited. To bypass or limit a cavity, an operator modifies the CED and allows the CED Event Server to propagate that value to the control system.

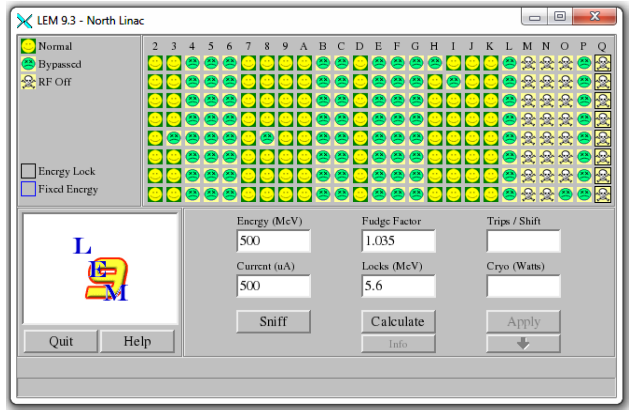


Figure 5: LEM retrieves cavity data from the CED.

## SUMMARY

The CED has proven its utility during installation and commissioning of the CEBAF 12GeV upgrade project. Its flexibility and the growing suite of tools that have been built upon it ensure that it will play a central role in efficient and safe future operation of the upgraded accelerator.

## REFERENCES

[1] P.M. Nadkarni, L. Marenco, R. Chen, E. Skoufos, G. Shepherd, P. Miller, "Organization of Heterogeneous Scientific Data Using the EAV/CR Representation," J. American Medical Informatics Association **6**(6), 478, PMID 10579606 (1999).

[2] http://www.oracle.com/technetwork/database/ twp-appdev-workspace-manager-11g-128289.pdf

[3] D.M. Beazley, "SWIG : An Easy to Use Tool For Integrating Scripting Languages with C and C++," Proc. of the Fourth Annual Tcl/Tk workshop. Monterey, CA (1996).

[4] http://www.swig.org/

[5] K. Baggett, "New and Improved: The JLab (state-of-the-art) HCO System," WAO2014, Workshop on Accelerator Operations, Mainz, Germany (2014).

[6] D. Turner, "eDT and Model-based Configuration of 12 GeV CEBAF," these proceedings, Proc. of IPAC2015, Richmond, Virginia (2015).

[7] M. Borland, "elegant: A Flexible SDDS-Compliant Code for Accelerator Simulation," Advanced Photon Source LS-287, September 2000.