ISBN: 978-3-95450-252-3 ISSN: 2673-5350 doi: 10.18429/JACoW-IPAC2026-MOP7143

UPDATES ON THE SUSTAINABILITY IMPROVEMENTS OF THE KARLSRUHE RESEARCH ACCELERATOR

J. Gethmann*, E. Bruendermann, W. Mexner, T. Mexner, A.-S. Mueller, R. Ruprecht, B. Shen,
Karlsruhe Institute of Technology, Karlsruhe, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2026). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

Abstract

The thermal well system supporting the cooling plant of the Karlsruhe Research Accelerator (KARA) accelerator test facility at the Karlsruhe Institute of Technology (KIT) is in test operation after about one year of commissioning. To avoid any impact on the environment, follow the governmental regulations, and document different aspects and operation statistics, KIT developed a special robust and reliable data handling pipeline.

We describe the implementation of the data archiving strategy, as well as the experiences gained and statistics on, e. g. power reduction from the first year of commissioning of the thermal wells. Furthermore, our plans for improvement and extension of KARA efforts towards sustainable operation are presented.

INTRODUCTION

The Karlsruhe Research Accelerator (KARA) is operated by the Institute for Beam Physics and Technology (IBPT) at the Karlsruhe Institute of Technology (KIT). KARA is a 2.5 GeV electron synchrotron operated as a light source for up to 100 days a year in recent years and up to 50 days as a test facility for accelerator physics experiments. The latter are often carried out at energies of 1.3 GeV or 0.5 GeV.

Since last year, a new thermal well cooling plant is under commissioning which is described in more detail in [1]. The operation of the thermal well system must comply with legal regulations requiring continuous monitoring as well as monthly, annual, and multi-annual reports. To reduce operational complexity, all relevant limits, thresholds, and monitoring functionalities are integrated into a dedicated monitoring system. This allows cooling domain experts to document disturbances directly within the monitor interface, while automatically generated reports provide comprehensive documentation of the thermal well system's operation. On the other hand, the raw data is also interesting for accelerator physicists trying to optimize the energy efficiency of the overall system, so it should be accessible to them as well from within the internal machine networks. From an IT security perspective, it is desirable to segregate different domains and security-critical areas from one another in such a way that data can only flow out of security-critical areas, but no access to these areas is permitted. Finally, from an operational point of view, one wants to maintain as few additional systems as possible and make use of existing infrastructure.

After going into details of the data flow and archiving workflow, we present some results of the thermal well system with respect to our sustainability approaches and mention our ideas for future improvements therein.

* julian.gethmann@kit.edu

EXISTING INFRASTRUCTURE AT KIT

KIT operates a central long-term data archiving service for research data [2]. This service stores data for at least ten years and is accessible from within the whole KIT via FTP. We use this service to store our reports we are obliged to create every month and year.

We obtain additional data from a campus-wide network of monitoring wells used for groundwater monitoring. These data are used to see the temperature changes in groundwater beyond the thermal well system's monitoring.

Existing Infrastructure at IBPT

Our accelerator control systems are based on EPICS [3]. Data of selected process variables are saved in a Cassandra database via a dedicated input output controller (IOC) as described in [4]. These controls are running in our machine internal networks and are only directly accessible from within the same networks. Data from our internal networks are published towards our institute network in read-only mode via an HTTP file service to obtain a minimal attack surface. Besides the general machine networks, there are further dedicated networks for IoT devices, like for Programmable Logic Controllers (PLCs), technical process and building controls. From there data can only flow outwards to one EPICS server, which is collecting the data.

DEVELOPMENT OF THE DATA FLOW

The data flows are depicted in Fig. 1. Here, the darker KIT green represents services that were operated before the commissioning of the thermal wells, light blue are the services developed for the thermal wells, and light green that were developed for the reporting tasks. The boxes match the different network locations quite well, except for the service node which is also running in the EPICS network.

The EPICS IOC pulls its data from the PLC based thermal well control system via OPC-UA. Another IOC polls the first IOC and writes the minutely averaged data to the file server. The first IOC is also polled by our archiving system which writes the raw data to a Cassandra database that is also accessible from the institute network via a gateway providing an REST-API. The extra storage of the data on the file server/data store is done to provide the data in the already aggregated way — minutely data as requested by regulations and in a simple CSV format following the KISS principle of software development.

ISBN: 978-3-95450-252-3 ISSN: 2673-5350 doi: 10.18429/JACoW-IPAC2026-MOP7143

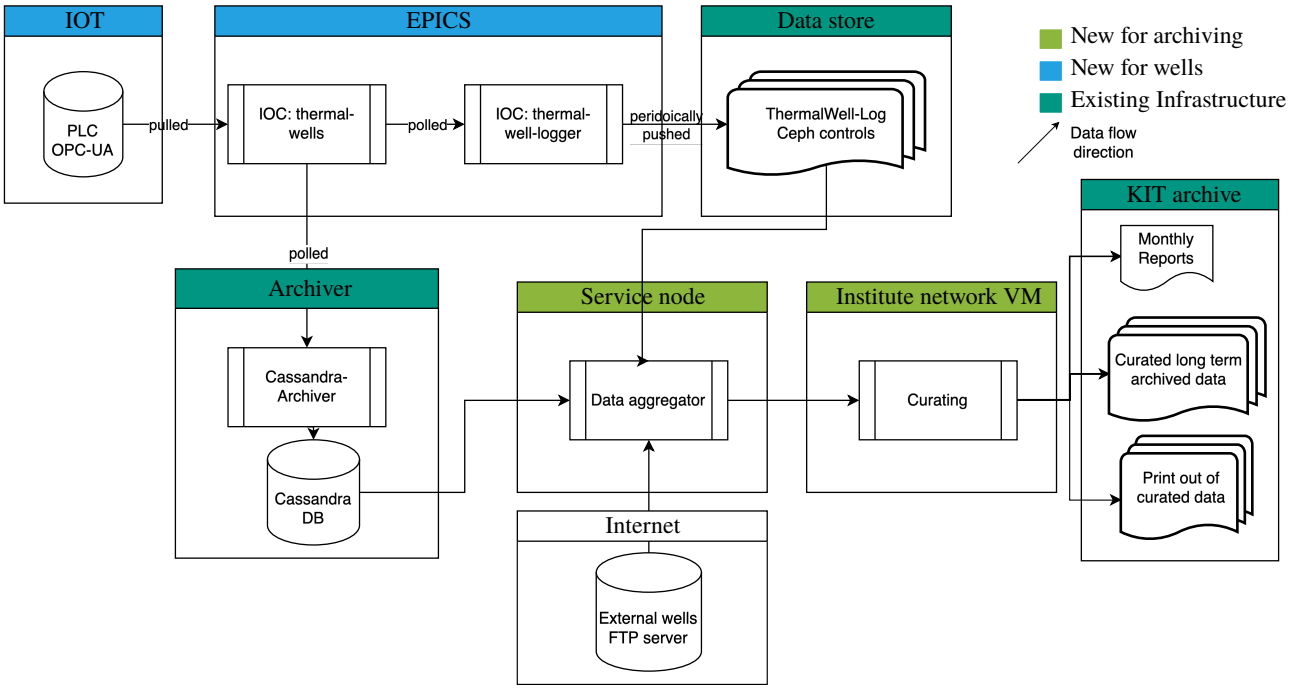


Figure 1: Data flow from the thermal wells' PLC and external wells through our existing data acquisition to the newly created aggregating and curating stages. Each box is one service in a different network.

For parts that require a bit more logic and flexibility, a Python program was written to aggregate the data from the data store, do error flagging, fill up missing data from the archive in case the IOC thermal-well-logger IOC was offline and warn if some limits are reached or data is not available. In case the data cannot be archived in the Cassandra database (e .g. because the EPICS infrastructure needs to be maintained or the network between the PLC and EPICS is unavailable, because of maintenance or unexpected incidents), it is also possible to manually fetch the data required by the authorities from the PLC for up to five days and insert them into the data aggregator. Lastly, the aggregator downloads temperature and groundwater level monitoring wells operated by campus groundwater monitoring services. That data serves as reference and cross-check data for trends outside of our own facility's surrounding. That is the only point where the data flow from a less trust-worthy network to a more-trust-worthy network.

For curating the data and creating a monthly report, we use a Python based visual interface to inspect the data. The curating includes a validity check of the data and commenting on special events which are logged in an electronic logbook. Trends also have to be commented on.

Figure 2 shows you one view of the GUI, which displays relevant parameters, including the temperature of the extraction wells, the absorbing wells, and the corresponding temperature differences. One can mark areas of interest and highlight them for comments or correct obviously incorrect data points like unit changes during installation of the system. On the right hand side you have free text forms to comment on your marks and on the general trends. The marks can



Figure 2: Screenshot of the editor for curating the temperature measurements. One can mark certain areas in the plots, modify them transparently, and write some comments. When finished a PDF report is being generated.

be either for one particular parameter, like temperatures, or for all of them as sometimes only certain parts might be affected by issues, e. g. failure of one sensor, and sometimes the whole system, e. g. the afore mentioned network interrupts, or EPICS maintenance. The GUI was developed taking into account feedback from end users and will be optimized for simplicity of use in the upcoming evaluation phase. The end user does not have to be an expert with the EPICS control system, the different networks, or the IT infrastructure, and can focus on their task with one portable entry point and a dedicated GUI. The final step of the curating process involves writing the PDF report to KIT's archive and keeping one hard copy of it.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2026). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-252-3 ISSN: 2673-5350 doi: 10.18429/JACoW-IPAC2026-MOP7143

© Content from this work may be used under the terms of the CC BY 4.0 licence (© 2026). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

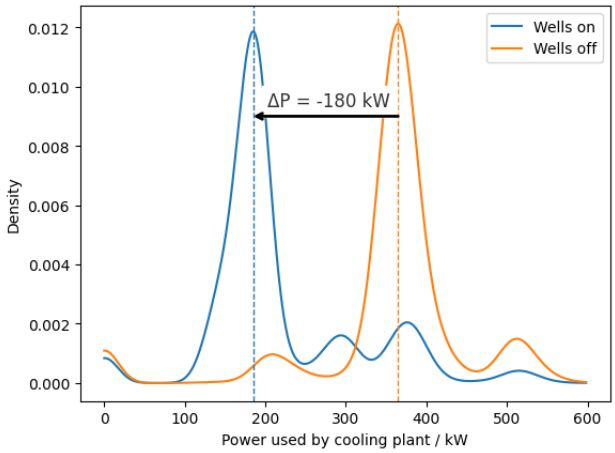


Figure 3: Kernel density estimation (KDE) of the quarter hourly power of the cooling plant with the thermal wells operating and not cooling.

OPERATIONAL EXPERIENCE AND DATA ANALYSIS

We operate the thermal well system for several months stable with only one of the three cooling machines as the main cooling system for our accelerator. However, we also had to deal with usual issues of new systems and we could especially improve the software systems to be more resilient and user friendly. We also tested several issues that might occur during normal operation like different error states of the PLC, wells, and network interrupts. During that time of commissioning, we could already save electrical energy for cooling. Figure 3 shows the Kernel density estimations data for the two cases of the wells, with and without cooling. The case "cooling" is here defined as water flowing into the dissipating wells, regardless of the amount nor how much it was heated up in the three heat exchangers. Also only the case of light source operation as typical for a so-called user operation is taken into consideration by filtering for beam energies of more than 2 GeV. The difference in the two cases is clearly visible as a 180 kW decrease which is a one by 49 % of the total amount. As the wells were on for about 1800 hours this yields 320 MWh of energy savings or about 200 MWh for half a year.

EXTENSIONS AND IMPROVEMENT IDEAS

Our next plan is to use some of the warmer water from our thermal well system as the supply water for the heat pump at a large, nearby cleanroom facility. This will allow us not only to cool our accelerator but also to reduce the amount of energy required to operate the air conditioning in the cleanrooms. Another possibility is to incorporate the thermal well system into our main building's existing cooling network.

Besides the thermal wells, we are installing two 932 kWh batteries in our accelerator complex to buffer the 540 kWp power of our solar power plant and enable autonomous operation of the wells also during the nights. Also this enables us to stabilize the grid by peak shaping.

SUMMARY

After years of successfully optimizing the power consumption by means of turning devices off when possible, replacing inefficient ones, installing photovoltaic power plants and already saving up to 1/3 in some cases, we could further decrease consumption by the installation of thermal wells. We implemented a data aggregation system for our new thermal well system that collects data from different data sources in a secure manner and enables the operations manager to curate the collected data and create monthly and yearly reports. The first year of commissioning already showed savings. A reduction of power consumption by 180 kW are realistic. We plan to make use of our waste heat and get more independent from the power grid with a battery system.

REFERENCES

[1] J. Gethmann, E. Bründermann, W. Mexner, A.-S. Mueller, R. Ruprecht, and B. Shen, "Utilization of renewable energies for sustainable accelerator operation at KIT", in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025, pp. 2670–2672. doi:10.18429/JACoW-IPAC2025-THPB096

[2] Research Data Archive, `https://www.rda.kit.edu/english/index.php`, visited May 2026

[3] The Experimental Physics and Industrial Control System, `https://epics-controls.org/`, visited May 2026

[4] S. Marsching, "Scalable Archiving with the Cassandra Archiver for CSS", in *Proc. ICALEPCS'13*, San Francisco, CA, USA, Oct. 2013, paper TUPPC004, pp. 554–557.