

15 YEARS OF THE J-PARC MAIN RING CONTROL SYSTEM OPERATION AND ITS FUTURE PLAN

Shuei Yamada*, KEK/J-PARC, Ibaraki, Japan

Abstract

The accelerator control system of the J-PARC MR started operation in 2008. Most of the components of the control computers, such as servers, disks, operation terminals, front-end computers and software, which were introduced during the construction phase, have went through one or two generational changes in the last 15 years. Alongside, the policies for the operation of control computers has changed. This paper reviews the renewal of those components and discusses the philosophy behind the configuration and operational policy. It is also discussed the approach to matters that did not exist at the beginning of the project, such as virtualization or Haber security.

INTRODUCTION

J-PARC (Japan Proton Accelerator Research Complex) is a high-intensity proton accelerator facility jointly planned, developed and operated by KEK (High Energy Accelerator Research Organisation) and JAEA (Japan Atomic Energy Agency). Construction began in 2001 and the operation started in 2007 [1]. The control system of the J-PARC accelerator consists of separate accelerator control systems for two different machine cycles. The accelerator control system for Linac (LI) and Rapid Cycle Synchrotron (RCS), whose repetition is 25 Hz, is managed by JAEA. While the control system for the Main Ring (MR) with 1.36 s or 4.24 s cycles, depending on the operation mode, is managed by KEK. These two control systems are built on a common infrastructure: timing system, network, storage system and EPICS [2], and work closely together to control the entire accelerator as a whole.

This paper discusses the 15-year operational history and future prospects of MR control system. The accelerator control system in the broadest sense also includes timing systems [3,4], personnel protection system [5], and machine protection system [6,7], but these are only mentioned in the references. The components of a networked distributed control system using EPICS, such as network, storage system, computing and software environment, will be focused on.

ACCELERATOR CONTROL NETWORK

Structure of the Network

The logical configuration of the J-PARC accelerator control LAN is shown in Fig. 1. The core switch is located in the Central Control Building (CCR) and is wired by optical fibre to the edge switches in the LI, RCS, Materials and Life Science Experimental Facility (MLF), and MR 3rd Power Supply Building (D3). Optical fibres are further wired from

D3 to the edge switches at D1, D2, Neutrino Experimental Facility (NU), and Hadron Experimental Facility (HD).

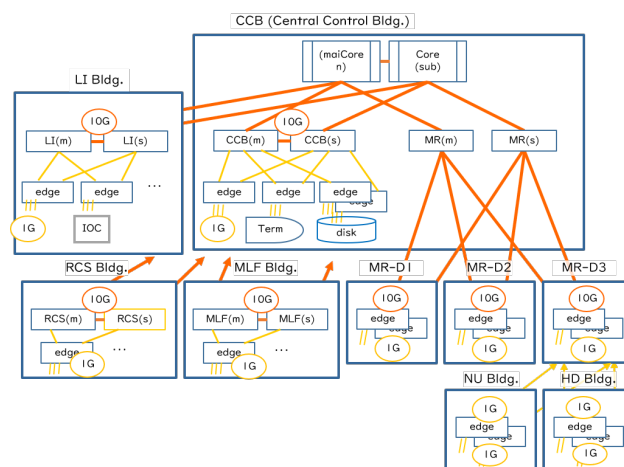


Figure 1: Logical Topology of the J-PARC Accelerator Network.

The core and edge switches are redundant and fails over from primary to the secondary system upon the failure. On-site maintenance services will promptly quickly identify the cause of the failure, take countermeasures and replace faulty components. VLAN divides the network into seven segments, namely, LI, RCS, MR, MLF, HD, NU, and CCR, minimizing the impact of network failures in one facility on other facilities.

The entire accelerator control network uses approximately 250 edge switches. LI and RCS are configured with edge switches right down to the end switch. MR uses a total of 12 edge switches for intermediate switches and SOHO switching hubs at the ends to reduce deployment and maintenance costs.

History of the Network

Equipment for the J-PARC accelerator control network has been supplied by Extreme, and all core and edge switches are replaced every 7–8 years. The control network started operation at LI in 2005 and was extended to MR in 2007. The bandwidth was 10 Gbps on the backbone and 1 Gbps between buildings. Network equipment was upgraded between 2011 and 2015, with bandwidth of 40 Gbps backbone and 10 Gbps between buildings. A second renewal of equipment has been underway since 2019. It will be completed in 2024, with a 100 Gbps backbone and 10 Gbps between buildings.

Network Security

The relationship between the J-PARC accelerator control network, the office network and the Internet is shown in Fig. 2.

* shuei@post.kek.jp

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI




Figure 2: Relationship between accelerator control network, J-PARC office network, and the Internet.

The control network is a private network independent of the office network and has various security measures to prevent entering malware and viruses.

The control network and the office network are not merely separated by firewalls, but direct communication between them is also prohibited. A DMZ is provided between them. Communication permitted are those between the control network and the DMZ, and between the DMZ and the office network. Network connections from the control network to the office or external network are restricted to those via HTTP proxy server or SSH server in the DMZ only. For http, only access to sites on the white list is permitted.

On the other hand, connections from the office network to the accelerator control network was initially restricted to HTTP and SSH to servers in the DMZ. In 2018, an EPICS Channel Access (CA) gateway was added to the DMZ, allowing read-only access from the office network [8]. SSH and CA connections are restricted by IP address on the office network side.

Anti-virus software is installed on the computers in the accelerator control network, depending on the level of risk. Furthermore, the terminals in the control network used by a large number of users are configured so that USB storage cannot be used. The use of older versions of operating systems no longer supported is also prohibited in principle.

STORAGE SYSTEM

Requirements for Storage

Reliability and scalability of storage systems are essential for long-term operation and high availability of accelerators.

MR selected a storage system manufactured by NetApp (or its OEM). It employs an operating system specialized for file servers, has redundant controllers, networks, and disks, and is able to expand its capacity by adding shelves and disks as needed. On-site maintenance services will promptly determine the cause of the failure, take countermeasures, and replace malfunctioning components.

History of Storage

An IBM N3600 with a capacity of 9 TB was introduced in 2008 and expanded to 28 TB in 2012. It was replaced by a

48 TB NetApp FAS220 in 2013 and expansion to 84 TB in 2014. The second replacement to a 192 TB Lenovo DM3000H took place in 2022.

None of them have experienced any serious problems, with exceptions of disk failure and replacement after about five years of operation. As of 2023, it is in operation as an NFS server for approximately 500 clients.

Linux Servers as Auxiliary Storage

The MR accelerator control system was designed to place all the applications programs for accelerator operation, data and each user's home directory on the storage system and share them via NFS.

However, rack-mounted file servers with a general-purpose operating system (Linux) have been deployed since 2021 for specific purposes. These file servers focus on disk capacity and affordability rather than reliability, I/O performance and scalability.

The longer accelerator operates, the amount of data increases (~10 TB/yr). However older data is less frequently referred to. Shot-by-shot waveform data acquired during MR operation will be moved from the storage system to the file server after one year. Two IBM x3630M3 with 24 TB of disks each were installed in 2012. They were replaced by an x3630M4 with 84 TB in 2016, followed by Lenovo SR550 with 216 TB in 2022.

Initially, the EPICS archiver was recording time series data to the storage system, with archive engine running on the blade server. The evaluation of the next archiver was started in 2016 on a x3630M3 surplus due to replacement. In 2017, a 112 TB Lenovo x3650M5 was installed and started as CPU and storage for the archiver [9].

SERVER COMPUTERS

Hardware of Servers

Initially, the control system for the MR accelerator adopted an architecture in which the application for accelerator operation was executed on a server computer and displayed on a terminal PC which is used as a X-terminal [10].

In 2014 terminal PC was upgraded and it was changed that applications runs locally [11]. Subsequent sections discuss the evolution of the server hardware and its usage form, as well as the terminal. This section discusses the changes in the server hardware, followed by the transition of its usage form. Terminals are discussed in the following section.

The operational policies of server are: (1) leave rooms for server capacity, (2) when failed allocate load to other servers and keep the accelerator continue to operate, (3) install more servers if the load increases, and 4) form a on-site support contract.

The server system began operation with an IBM Blade-Center E (Fig. 3) enclosure populated with five HS20 blades in 2005. A second enclosure was installed in 2007 and more blades were added as needed. In 2015 the system was fully populated in two enclosures, making a total of 28 blades. When the HS23e was introduced in 2014, the power supply



(a) An enclosure populated with 14 blades.



(b) A single blade server.

Figure 3: Appearance of blade server system.

capacity of the enclosures were found to be insufficient. This forced a reduction to 22 blades, but increased the overall performance.

Enclosures have reached EoL during 2018–2019. Two 1U rack-mount servers (Lenovo SR530) were installed in 2019 to evaluate the transition from blade to rack-mount server. One of the enclosure failed in July 2020. Its repair was abandoned.

Eight SR530s were urgently arranged, as the number of available blades had been halved to 11.

At the same time the server system was temporarily restored with remaining blades, two existing SR530s, four fanless servers for IOCs, and one terminal PC. MR continued to operate on a temporarily system for six months until the arranged SR530 became available. In July 2021, all blade computers were decommissioned and replaced by 10 SR530s. Two additional SR530s were added in March 2022. Typical specifications for servers are shown in Table 1.

The advantage of blade-type servers is the low cost of model upgrades. During the 12 years of stable operation, seven models of four generations from HS20 to HS23e were in operation. On the other hand, disadvantages were also

Table 1: Typical Specifications of Servers

Model	CPU	Disk space
HS23e	Xeon E5-2470	48 GB
SR530	Xeon Gols 5215	96 GB

identified, such as requirements of large space for enclosures, power supply capacity limiting the number of blades available, and the cost of replacing the entire system. The foreseen scenario, where the lifespan of the enclosure determines the lifespan of the entire blade system, could also not be avoided.

Conversion to VM Hosts

When blade-type servers were first put into operation, individual blades were allocated individual functions, such as for running accelerator operation programs, DHCP servers and DNS servers [10]. In 2011 XEN virtual machines (VMs) as front-end computers were introduced on Scientific Linux (SL) 4 hosts.

In 2012 VMs were migrated KVM on SL6. At the same time multiple small servers were migrated to VMs and consolidated into one blade [12].

As discussed later, the replacement of terminal PCs in 2014 has reduced the load on the servers. Thus servers are being converted to VM hosts. As of 2023, approximately 100 VMs are in operation on five servers.

The introduction of virtual machines has enabled more effective use of computing resources and improved fault tolerance through live migration. It is also easy to load-balancing the IOC loads, launch IOCs for development and testing, and maintain the development environments for older OS versions. Container technologies such as Docker and LXC are also being evaluated, but have not been put into accelerator operation.

Servers for Specific Purposes

From 2010, servers for specific purposes have also been introduced. Servers for accelerator operation are not suitable for accelerator simulations. Servers with selected number of CPU cores, amount of memory, GPGPU, etc. are installed according to the characteristics of the simulation. The CA gateways and server for archivers described above would also included in this category.

TERMINAL COMPUTERS

About 20 HP Thin Client terminals with 2- or 4-display configurations were installed as terminal calculators in 2006–2007. For the simple operation and management of the large number of terminals, they were network-booted with SL4 and used like an X terminal [13]. Following the OS for servers and IOCs were updated to SL6, terminals were replaced by Intel NUCs in 2014 [11]. Their faster CPU and larger memory space enabled accelerator operation applications to be run on terminals, which were previously run on the servers. The appearance of the Thin Client and NUC

are shown in Fig. 4 and typical specifications are given in Table 2.



Figure 4: Intel NUC (left) and Thin-client (right).

Table 2: Typical Specifications of Terminals

Model	CPU	Memory	Disk
HP t5720	Geode NX1500	512 MB	none
Intel NUC8i5BEH	Core i5-8259U	32 GB	SSD

About 20 HP Thin Client terminals with 2- or 4-display configurations were installed as terminal calculators in 2006–2007. For the simple operation and management of the large number of terminals, they were network-booted with SL4 and used like an X terminal [13]. Following the OS for servers and IOCs were updated to SL6, terminals were replaced by Intel NUCs in 2014 [11]. Their faster CPU and larger memory space enabled accelerator operation applications to be run on terminals, which were previously run on the servers. The appearance of the Thin Client and NUC are shown in Fig. 4 and typical specifications are given in Table 2.

NUCs are low-cost small PCs with no manufacturer maintenance service, but they are unexpectedly robust. 20 out of 30 machines installed in 2014–2015 are still in operation in

2023. New models are released annually with faster CPUs and larger memory available, but there is no significant difference from a software compatibility and management point of view. New models are tested from time to time and kept as spare units, in preparation for replacement upon failure or additional terminals. As of 2023, approximately 45 units are in operation with two-displays configuration.

FRONT-END COMPUTERS

Front-end computers in EPICS based control system are called I/O Controllers (IOCs). VME SBC selected as IOC in 2007 [14], followed by Linux ready PLC CPU (Yokogawa F3RP61) in 2008. In 2010, IOCs running on virtual machines (VIOCs) were introduced [12].

VME SBCs were robust and a good choice, but most of the control targets were found to be network-connected devices such as PLCs and oscilloscopes. VME crates were used as a mere power supply, with some exceptions described in later sections. As MR became more sophisticated, the number of IOCs also increased. Even the number of VME SBCs is small, a space for VME crate is required. VME SBCs share power supply when they are populated in one VME crate. These became operational bottleneck.

In 2014 a small fanless server (PiNON Sabataro Type-P) was introduced as an IOC [15]. The migration from VME SBCs to Sabataro will be completed in 2023.

The types and evolution of IOCs for MR operations are summarised in Table 3. The appearance of a typical IOC is shown in Fig. 5 and its specifications in Table 4.

OPERATING SYSTEM AND EPICS

The Basic Philosophy of Software Environment

From the operational aspect, it is preferable that servers, terminals and IOCs have the same version of OS and EPICS, and stay with them as long as possible. From the security, on the other hand, it is recommended to update the software as often as possible.

As a compromise between these conflicting demands, MR control systems keep using the same major version of the

Table 3: Evolution of Types and Numbers of IOCs

	2007	2008	2011	2013	2014	2015	2016	2021	2023
VME SBC	~80	~80	~90	~80	~80	~90	~90	55	37
microIOC	3	3	3	3	3	3	3	3	0
F3RP61		~10	~30	~40	~45	~45	~45	69	69
VIOC			a few	~30	~30	~30	~30	39	44
Saba-taro					a few	11	~30	50	80

Table 4: Typical Specifications of IOCs

Model	CPU	Memory	Disk
Sanritz SVA041 (VME SBC)	Celeron-M 600 MHz	512 MB	diskless
Yokogawa F3RP61	ppc603e 600 MHz	128 MB	CF card
PiNON Saba-taro Type-P	Celeron J1900	8 GB	SSD




Figure 5: Three form-factors of front-end computers.

operating system once installed, as long as it is supported. Patches and updates are applied as frequent as possible. In practice, operating system is patched (and rebooted if required) during long-term maintenance period of the accelerator, such as annual power outage in summer, New Year's holidays, and end of the fiscal year. Servers are only rebooted during annual power outages, in principle.

For the EPICS base, the latest stable version at the time of deployment is selected and retained. EPICS patch levels may be updated if serious problems are found. All of these are tested and developed on spare computers and control targets and then put on production.

History of MR Control Software

SL4 and EPICS R3.14.7 were chosen as the basis for the MR control system in 2008. Standard EPICS GUI tools such as EDM, MEMDM and StripTool were used. In 2012 support for SL4 ended. Therefore servers migrated to SL6 and EPICS R3.14.12.3 in 2012, and IOCSs in 2013 [16]. Terminals were converted from thin clients to NUCs in 2014 due to difficulties in running SL6 [11]. Thus the software for servers, terminals and IOCs was unified with SL6 and EPICS R3.14.12.3. More modern but demanding software such as Control System Studio became available on terminals [17].

In 2019-2020, the software environment was migrated to CentOS7 and EPICS R3.15.5 to be prepared for the scheduled SL6 EoL in 2020. The NUCs purchased in 2019 were found not to work with SL6 expectedly. As it became necessary to get these into operation urgently, CentOS7 was chosen as the next OS, rather than wait for the release of CentOS8. In 2020, EPICS R3.15.5 was updated to R3.15.8 after it was found that the CPU could runaway. The EoL of CentOS8 was brought forward to 2022, earlier than originally planned, resulting in CentOS7 being the right choice.

CentOS7 EoL is planned for 2024. Migration to Alma Linux 9 and EPICS R7.0.7 started in 2023 and will be completed in 2024.

Exceptions to the Basic Philosophy

Many devices in MR are controlled through the network, but some exceptional ones are directly connected to the bus

of the IOC. These IOCs are difficult to upgrade the operating system.

The MR timing receiver is an in-house VME board. The IOC used for its control is VME-7700RC (GE) SBC running with SL4 and EPICS R3.14.7. The VME SBC and its operating system need to be updated, while device drivers and applications need to be significantly overhauled. Eventually, the decision was made for a migration to PLC module-type receivers in conjunction with the upgrade to next-generation J-PARC accelerator timing system [3, 4]. Of the 25 VME-type timing receiver IOCs, 23 are planned for migration to the PLC-type. The remaining are under consideration.

The ADC used for MR beam loss monitor is also an in-house VME board, controlled by a V7865 VME SBC (GE) with SL6 and EPICS R3.14.12. All 12 units will be upgraded to SVA061 (Sanritz) with CentOS7 and EPICS R3.15.8.

Yokogawa F3RP61 is a Linux-ready CPU module for Yokogawa PLCs, with PowerPC CPU and ELDK-based embedded Linux. It reads and writes PLC modules installed on the backplane and behaves as an IOC [18]. The F3RP71, its successor, has switched to ARM CPU and Yocto-based Linux. Approximately 70 F3RP61s are in operation in the MR control system. The transition to F3RP71 started in 2022.

SUMMARY

The accelerator control system of J-PARC MR over the past 15 years is overviewed. Each component has went through one or two generational changes. Some have simply been updated, but others have changed their way of operation. The background to this is discussed.

The new software and hardware are being tested on a spare unit and then deployed in production environment. This approach has worked well so far. The lifespan of control system components is 10 years for the longest and 2-3 years for the shortest, and tends to shorten year by year. It is expected that Linux and EPICS will continue to be used as the basis for control system software. Software and hardware are interdependent and need to be updated in tandem. Continuous R&D is vital to maintain the control system as a whole.

REFERENCES

- [1] N. Kamikubota *et al.*, “J-PARC Control toward Future Reliable Operation”, in *Proc. ICALEPCS'11*, Grenoble, France, Oct. 2011, pp. 378-381, paper MOPMS026.
- [2] EPICS - Experimental Physics and Industrial Control System, <https://epics-controls.org/>
- [3] N. Kamikubota *et al.*, “Ten-year operation and experienced troubles of J-PARC MR timing”, in *Proc. PASJ'19*, Kyoto, Japan, Jul.-Aug. 2019, paper THOI07.
- [4] F. Tamura *et al.*, “Next generation timing system for J-PARC”, in *Proc. PASJ'19*, Kyoto, Japan, Jul.-Aug. 2019, pp. 149-152, paper THOI08.
- [5] N. Kikuzawa *et al.*, “Present Status of Personnel Protection System at J-PARC”, in *Proc. PASJ'19*, Kyoto, Japan, Jul.-Aug. 2019, pp. 877-880, paper FRPH006.
- [6] T. Kimura *et al.*, “Performance Evaluation of MR-MPS and Development Plan of New MR-MPS for J-PARC”, in *Proc. PASJ'17*, Sapporo, Japan, Aug. 2017, pp. 1148-1150, paper WEP101.
- [7] H. Takahashi *et al.*, “Update of MPS Modules for J-PARC LINAC and RCS (2)”, in *Proc. PASJ'21*, Takasaki, Japan, Aug. 2021, pp. 914-917, paper THP038.
- [8] S. Yamada, “Real-time and Detailed Provision of J-PARC Accelerator Operation Information from the Accelerator Control LAN to the Office LAN”, in *Proc. PCaPAC'18*, Hsinchu, Taiwan, Oct. 2018, pp. 167-169.
doi : 10.18429/JACoW-PCaPAC2018-THP04
- [9] S. Yamada *et al.*, “Deployment of archiver appliance at J-PARC main ring”, in *Proc. PASJ'17*, Sapporo, Japan, Aug. 2017, pp. 1144-1147, paper WEP100.
- [10] N. Kamikubota *et al.*, “Computer Environment for J-PARC MR Operation”, in *Proc. PASJ'10*, Himeji, Japan, Aug. 2010, pp. 690-692, paper WEPS116.
- [11] S. Yamada, “Renovation of PC-based Console System for J-PARC Main Ring”, in *Proc. PCaPAC'14*, Karlsruhe, Germany, Oct. 2014, pp. 81-83, paper WPO021.
- [12] N. Kamikubota *et al.*, “Experience of Virtual Machines in J-PARC MR Control”, in *Proc. ICALEPCS'13*, San Francisco, CA, USA, Oct. 2017, pp. 417-419, paper MOPPC131.
- [13] S. Yoshida *et al.*, “Console System Using Thin Client for the J-PARC Accelerators”, in *Proc. ICALEPCS'07*, Knoxville, TN, USA, Oct. 2007, pp. 383-384, paper WPPA33.
- [14] N. Kamikubota *et al.*, “Operation Experience and Migration of I/O Controllers for J-PARC Main Ring”, in *Proc. PCaPAC'16*, Campinas, Brazil, Oct. 2016, pp. 101-104.
doi : 10.18429/JACoW-PCaPAC2016-THPOPR009
- [15] S. Yamada, “Deployment of a Tiny Fan-Less Server as IOC in J-PARC Main Ring”, in *Proc. PASJ'16*, Chiba, Japan, Aug. 2016, pp. 634-636, paper MOP092.
- [16] S. Yamada, “Upgrade of software toolkits for EPICS Input Output Controllers in J-PARC Main Ring”, in *Proc. PASJ'13*, Nagoya, Japan, Aug. 2013, pp. 1106-1108, paper SUP080.
- [17] S. Yamada *et al.*, “Deployment of Control System Studio at J-PARC Main Ring”, in *Proc. PASJ'15*, Tsuruga, Japan, Aug. 2015, pp. 767-769, paper WEP103.
- [18] J. Odagiri *et al.*, “Development of Embedded EPICS on F3RP61-2L”, in *Proc. PASJ'08*, Higashihiroshima, Japan, Aug. 2005, pp. 240-242, paper FO05.