IIR

Internet
Infrastructure
Review

Aug.2026
Vol. 70

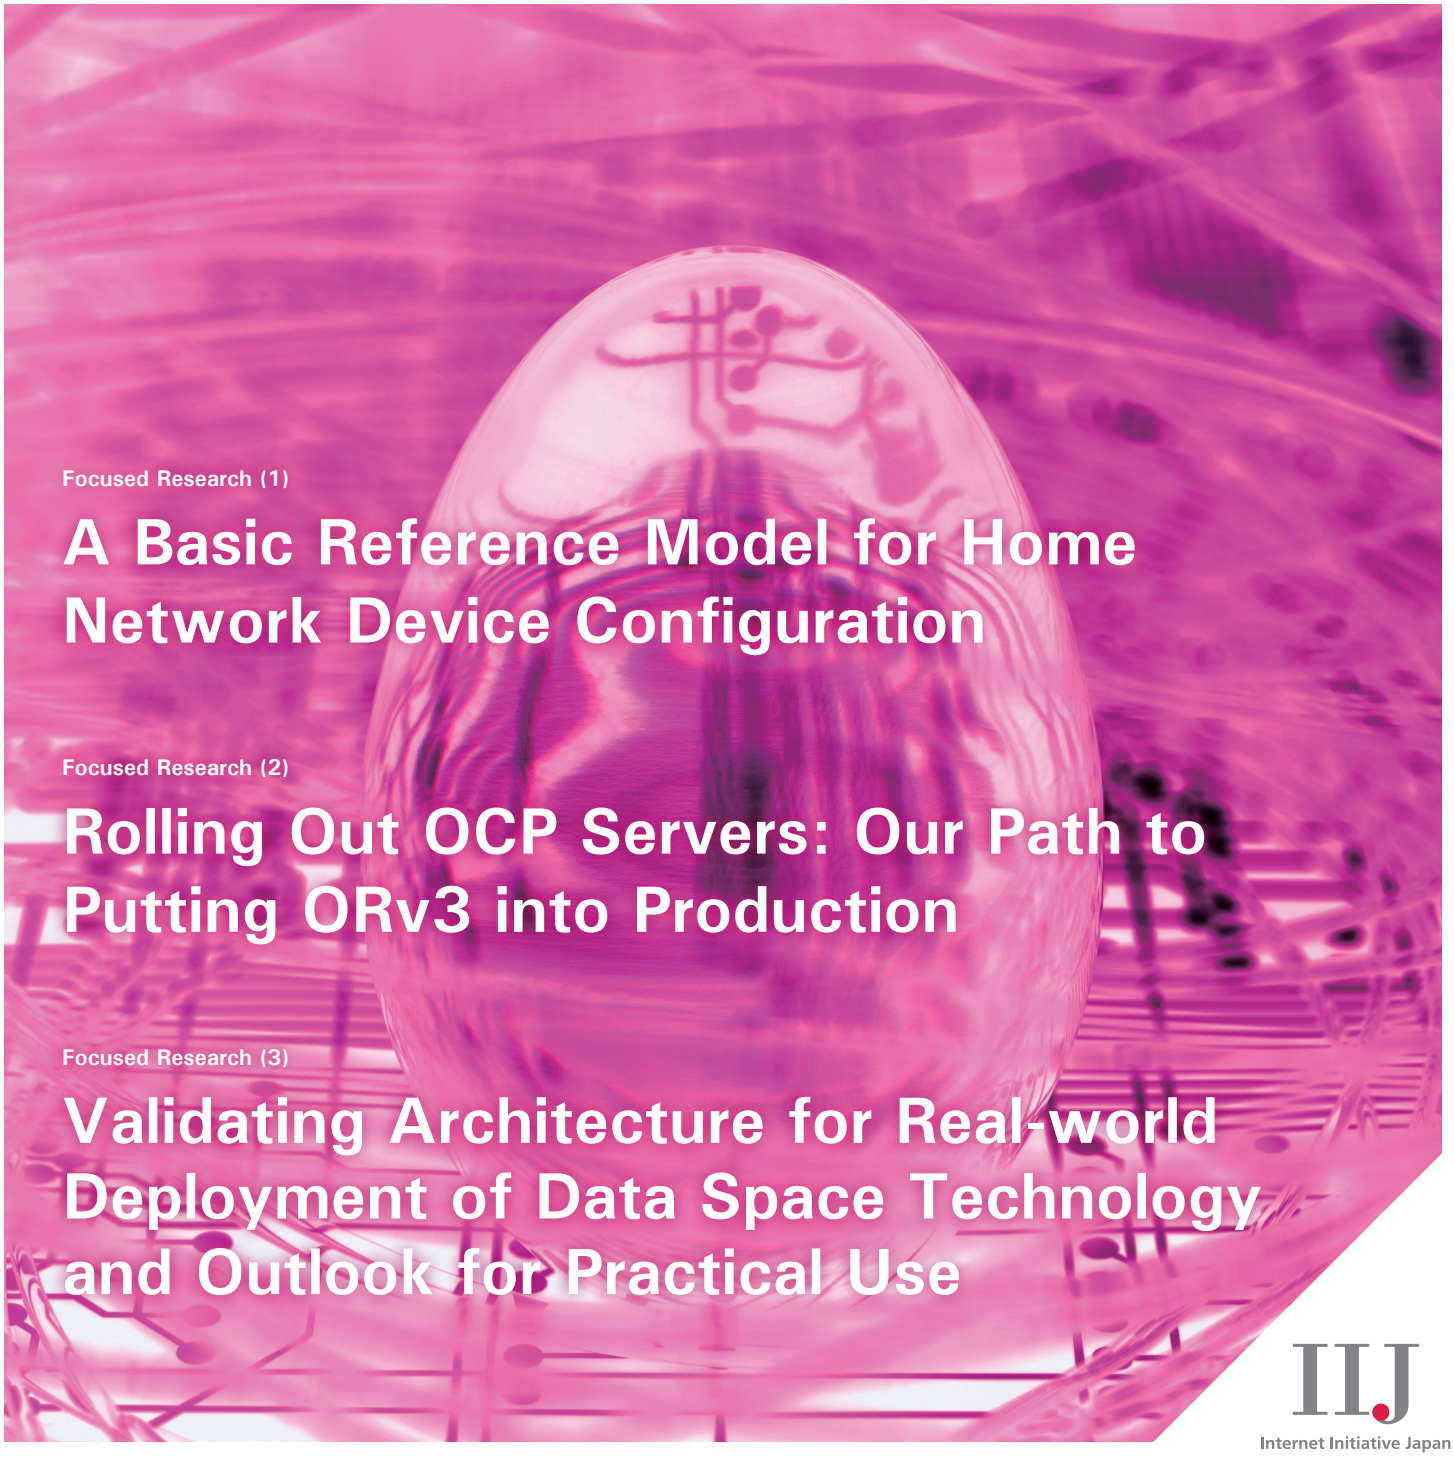

Focused Research (1)

A Basic Reference Model for Home Network Device Configuration

Focused Research (2)

Rolling Out OCP Servers: Our Path to Putting ORv3 into Production

Focused Research (3)

Validating Architecture for Real-world Deployment of Data Space Technology and Outlook for Practical Use

IIJ
Internet Initiative Japan

Internet Infrastructure Review

August 2026 Vol.70

Executive Summary ····· 3

1. Focused Research (1) ····· 4
1.1 Introduction ····· 4
1.2 The Need for Standardization ····· 4
1.3 Organizational Challenges in Home Network Standardization ····· 5
1.4 Role of IEC TC 100 TA 18 ····· 6
1.5 IEC 62608: Configuration Management for Multimedia Home Networks ····· 7
1.6 Background to Formulation of IEC 62608 and What's Next ····· 8

2. Focused Research (2) ····· 10
2.1 Introduction ····· 10
2.2 How IIJ's Thinking on Server Procurement Has Evolved ····· 10
2.3 Why We Were Able to Take On ORv2 ····· 10
2.4 The ORv2 System We Deployed and What We Learned Running It ····· 11
2.5 Why We Didn't Deploy ORv3 Right Away ····· 11
2.6 OCP Global Summit as a Turning Point ····· 11
2.7 IIJ's Key Considerations When Evaluating ORv3 ····· 12
2.8 The ORv3 Production Environment We Deployed in 2025 ····· 12
2.9 How We Made OCP Server Maintenance Work ····· 12
2.10 Why We Limit OCP Server Deployment ····· 12
2.11 Power Density Challenges in the SPR Generation and Planning the Next Configuration ··· 13
2.12 Conclusion ····· 13

3. Focused Research (3) ····· 14
3.1 Introduction ····· 14
3.2 Architecture for Realizing Data Spaces ····· 14
3.2.1 Trust-based Collaboration and the Preservation of Data Sovereignty ····· 14
3.2.2 Four-layer Architecture of a Data Space ····· 14
3.2.3 How Autonomous, Distributed Collaboration Works ····· 15
3.3 Design and Implementation of the PoC System ····· 15
3.3.1 Trust Model and Authorization Control in the Identity Layer (L3) ····· 15
3.3.2 Implementation of the Data Collaboration Platform and L2 Connector ····· 16
3.3.3 Integration of Backend Systems and the Data Layer (L1) ····· 16
3.3.4 Access Control and Future Extensibility ····· 16
3.4 Verifying Technical Effectiveness Through PoC Scenarios ····· 16
3.4.1 The Need to Strengthen Cyber Resilience Through Wide-Area Equipment Collaboration ···· 16
3.4.2 Improving Maintenance Processes via Attribute-Based Access Control ····· 16
3.4.3 ynamically Sharing Support Information Using a Chain of Trust ····· 17
3.4.4 Summary of PoC Results ····· 17
3.5 Technical Considerations and Future Outlook ····· 17
3.5.1 Technical Challenges for Real-World Deployment and Advancing the Trust Model ····· 17
3.5.2 Building a Data Ecosystem That Creates Value ····· 18
3.5.3 Prospects as General-Purpose Infrastructure and Convergence with AI ····· 18
3.6 Conclusion ····· 19

Executive Summary

IIJ recently announced its AI business strategy. Up to now, we have been building a track record of results through internal proof-of-concept work on AI utilization and the application of AI to some business operations. Having moved beyond that stage, we have now entered a phase in which AI is becoming embedded across the company as a foundation of our business, and it is in this light that we formulated and announced our company-wide AI strategy.

Since its founding, IIJ has fulfilled its role of providing social infrastructure centered on the Internet, with its business underpinned by the advanced technical and operational capabilities of its talented engineers. This is one of our great strengths and lies at the very core of the value we provide. Our aim is to use AI to convert the tacit knowledge of individuals, the experience and judgment of those working on the front lines, into explicit knowledge, build it into a corporate asset, and enable each and every employee to perform to their full potential. We plan to share details of specific initiatives as opportunities arise.

This issue of the IIR focuses on the evolution of core technologies underpinning the Internet, presenting IIJ's initiatives from three perspectives: home networks, server infrastructure, and data spaces.

Chapter 1 looks at the basic reference model for home networks based on international standard IEC 62608. With the increasing number and diversity of connected devices, home networks have reached the point at which we need to go beyond simply connecting individual devices and instead look at managing the network as a whole system. Moving beyond conventional device-to-device interoperability, the article lays out a framework for monitoring and managing the entire network in a unified manner, discusses the background to standardization, and lays out a future direction for home network management.

Chapter 2 discusses the evolution of IIJ's server infrastructure, looking at how IIJ came to deploy servers based on the Open Compute Project (OCP), the key considerations in its evaluation of Open Rack v3 (ORv3), and the configuration used in its production environment. In the course of moving from conventional enterprise servers to more scalable and efficient OCP-based configurations, IIJ's procurement model and design philosophy have also changed substantially. Building on the lessons learned from ORv2, and after repeated rounds of testing on factors such as power efficiency, rack utilization, and operability, IIJ moved ahead with the deployment of ORv3. As part of the infrastructure underpinning IIJ's services, OCP servers have become an important option in terms of both procurement and system building.

Chapter 3 examines the potential of decentralized data collaboration platforms through a proof of concept (PoC) on data space technology. Unlike conventional centralized platforms, a data space is a framework for secure data collaboration that preserves data sovereignty, a concept that has been championed primarily in Europe. The article explains the concept in concrete terms based on the results of IIJ's PoC. The PoC took the strengthening of cyber resilience in telecommunications infrastructure as its use case, confirming the effectiveness of attribute-based access control and policy-based data distribution. The article also sets out the challenges that lie ahead, including the standardization of semantics and the advancement of trust models, as well as the importance of building ecosystems founded on the distribution of data.

Guided by our mission of supporting society through technology, we will continue to advance the infrastructure technologies covered in this issue—networks, server infrastructure, and data—while also working to appropriately incorporate AI into our operations as we pursue the ongoing evolution of our services.

Junichi Shimagami

Mr. Shimagami is a Director and Executive Vice President and the CTO of IIJ. His interest in the Internet led to him joining IIJ in September 1996. After engaging in the design and construction of the A-Bone Asia region network spearheaded by IIJ, as well as IIJ's backbone network, he was put in charge of IIJ network services. Since 2015, he has been responsible for network, cloud, and security technology across the board as CTO. In April 2017, he became chairman of the Telecom Services Association of Japan's MVNO Council, stepping down from that post in May 2023. In June 2021, he also became a vice-chairman of the association.

A Basic Reference Model for Home Network Device Configuration

1.1 Introduction

Have you ever stopped and counted how many devices are connected to your home network? Beyond IT devices like PCs, smartphones, tablets, and printers, you may also find other devices connected to your home network such as game consoles and TVs, and even home appliances like rice cookers, air conditioners, and audio equipment that now come with Wi-Fi. It's commonplace for people to have two or more personal mobile devices these days, say a smartphone and a tablet, so in a family household, it's not at all uncommon for over 10 devices to be connected to the network.

Moreover, a much wider variety of devices can be found in the home than in the office. Most office IT equipment comes with standard tools for checking IP addresses and connectivity. But with TVs and other home appliances, it is not always easy to check the assigned IP address or MAC address, and the controls typically differ by manufacturer and model. On top of that, connectivity within the home involves a whole mix of things—common networking standards like wired Ethernet and Wi-Fi as well as dedicated audio cables and short-range links like Bluetooth.

And it's not necessarily the case that someone in any given household has some IT expertise. Many people reading this may have been asked to, say, fix a network problem at an aging parent's house and struggled to sort it out. Faced with the common refrain of "it won't connect" or "it doesn't work," there's currently no single way to determine whether the culprit is one particular device, a network misconfiguration, or a conflict with some nearby device. With there being few means of accurately assessing such a situation from afar, the result is that you generally have to go there in person to sort it out.

Not only is this state of affairs inconvenient, it can also affect the security of the entire network. There are real-world reports of vulnerabilities in home routers, webcams, and the like being exploited to conscript such devices into botnets involved in large-scale DDoS attacks. As home networks grow ever more important as everyday infrastructure, we need to move beyond the idea of individual devices merely being used to connect and instead think about how the system as a whole can be managed.

It is against this backdrop that the industry has long sought a common framework that would ensure interoperability and provide a systematic approach to the configuration management of multimedia devices in the home. This article looks at international standard IEC 62608-1 Ed. 2.0:2025, "Multimedia home network configuration – Basic reference model," developed within IEC TC 100 TA 18 (now handled by WG 31). We examine how it came to be formulated and the technical content it defines. I note that I have worked as a leader on this standardization effort since the first edition of IEC 62608.

1.2 The Need for Standardization

In an environment populated by a multifarious mix of devices like a home network, standardization is key to interoperability. There was a time when people preferred to buy all their appliances from a single manufacturer, but as the variety of products available grew more diverse, it became commonplace to mix and match products from different manufacturers. The cables and connectors that physically link devices have gradually been standardized. Even so, it was not uncommon for interoperability features to fail to work as expected when you connected, say, a TV and a content recorder from different manufacturers. Network connection methods and underlying assumptions also varied, and proprietary tweaks sometimes prevented data transfer between products from different manufacturers.

Standardization comes in several flavors, depending on how a standard comes about and what role it plays. International standardization broadly takes three forms.

1. De facto standards: Standards that take hold not through formal consensus but through broad market adoption. These are technologies developed by a particular company, or specifications set by a particular body (forums like the IETF or W3C), that became the industry standard through overwhelming market share or convenience. Many Internet protocols take this form, since "it works" is what matters most.

2. De jure standards: Standards established by official standardization bodies based on international consensus. Standards from the IEC (International Electrotechnical Commission), which we discuss here, as well as ISO (International Organization for Standardization) and ITU (International Telecommunication Union), fall into this category. Delegates from each country gather to deliberate and formalize standards under set procedures.
3. Consortium standards: Standards developed by industry alliances (consortiums) formed for a specific purpose. Examples include the USB standards from the USB Implementers Forum (USB-IF) and Wi-Fi from the Wi-Fi Alliance, Bluetooth from the Bluetooth SIG, as well as UPnP (Universal Plug and Play) from the UPnP Forum for automatic device discovery and DLNA (Digital Living Network Alliance) for multimedia sharing.

IEC 62608, which we discuss here, is a de jure standard set by an official international body.

To date, UPnP has made device discovery and control easy, and DLNA has established a mechanism for sharing multimedia content on top of UPnP. IEC 62608-1 references these standards as well. But while they succeeded at making specific services work, they fell short when it came to managing the configuration of the entire network as a single system. There was no "reference model" to provide an overview of the devices combined in a home network in terms of what configuration information they hold and how they behave.

The purpose of IEC 62608 is to integrate existing individual protocols and provide a common framework for managing the whole from a higher level. Standardization at this layer can be expected to have the following benefits.

- Unified management in multi-vendor environments: Lets you inspect the status of devices across a home network through a common model, without having to rely on manufacturer-specific management tools.
- Remote monitoring and diagnosis: Thanks to standardized configuration management, lets you accurately assess conditions and troubleshoot quickly from external locations.
- Adaptability to technological change: By defining a "basic reference model" that does not depend on any specific communication protocol, continuity in management methods can be maintained even as new communication technologies emerge in the future.

Standardization like this benefits users, but it is yet to come to full fruition, and IEC 62608 cannot yet be called widely adopted. Getting there required more than technical discussion. We also had to tackle issues arising from the scope and sphere of influence of standards organizations in Japan and abroad.

1.3 Organizational Challenges in Home Network Standardization

Historically, communication standards and device (product) standards have been managed by different government agencies and different standards bodies. Home appliances in particular have long enjoyed a strong, independent global market and community as standalone devices. But with the spread of the Internet and the networking of devices in recent years, the boundaries between these domains (scopes) have been blurring.

So standardizing home network configuration management required, in parallel with the technical work, coordination stemming from the scopes of the standards organizations and government agencies involved. As Figure 1 illustrates, a home network straddles two different standardization worlds, with the home gateway sitting at the boundary. The scope of each of those worlds is as follows.

- Left side: IEC/JEITA (the world of devices and appliances). This covers everything inside the home gateway—household TVs, audio equipment, white goods, and so on. Internationally, this is part of the IEC's (International Electrotechnical Commission) purview. In Japan, it falls under the Ministry of Economy, Trade and Industry, with JEITA handling the work. The focus is mainly on "functionality as a product."
- Right side:ITU/TTC (the world of communications

and the Internet). This covers everything outside the home gateway—the communication path out to the Internet. Internationally, de jure standards are handled by ITU-T and de facto standards by the IETF (Internet Engineering Task Force). In Japan, this falls under the Ministry of Internal Affairs and Communications, with TTC handling the work.

Users are able to use the services they do because their devices in the home and the Internet outside connect seamlessly through the home gateway. The home gateway doesn't just link the home to the outside Internet, it also handles address translation (NAT) and firewall functionality, and in technical terms, it is designed as a network device based on Internet-side specifications. It gives devices in the home the means to connect to the network, but managing those devices was never part of the discussion.

In formulating IEC 62608, we had to keep tabs on developments in both of these worlds. So we built an information-sharing framework. With JEITA as our hub, we participated in TTC on the communications side, and in turn, TTC members also participated in JEITA. The findings we came to through this domestic collaboration were then submitted in the form of a Japanese proposal to IEC international meetings.

1.4 Role of IEC TC 100 TA 18

The IEC (International Electrotechnical Commission), where the international standard was proposed, is an international standardization body founded in 1906 that covers electrical engineering, electronics, and related technologies. With over 90 countries participating as full and associate members, it sets de jure standards spanning energy (power generation and transmission), electroacoustics, multimedia, and telecommunications, along with related terminology, symbols, compatibility, measurement, performance, and reliability.

The IEC's work is divided among Technical Committees (TCs) that handle a particular field. TC 100 (Audio, video and multimedia systems and equipment), which received this standard, specializes in audio, video, and multimedia systems and equipment. Beneath it sit Working Groups (WGs) responsible for specific technical areas.

WG 31, where this standard was discussed, defines its scope as "Multimedia home systems and applications for end-user networks."

WG 31's predecessor, TA 18, produced many international standards for technologies used in everyday life. The flagship example is USB. The USB specification itself is a consortium standard developed by the USB-IF (USB Implementers

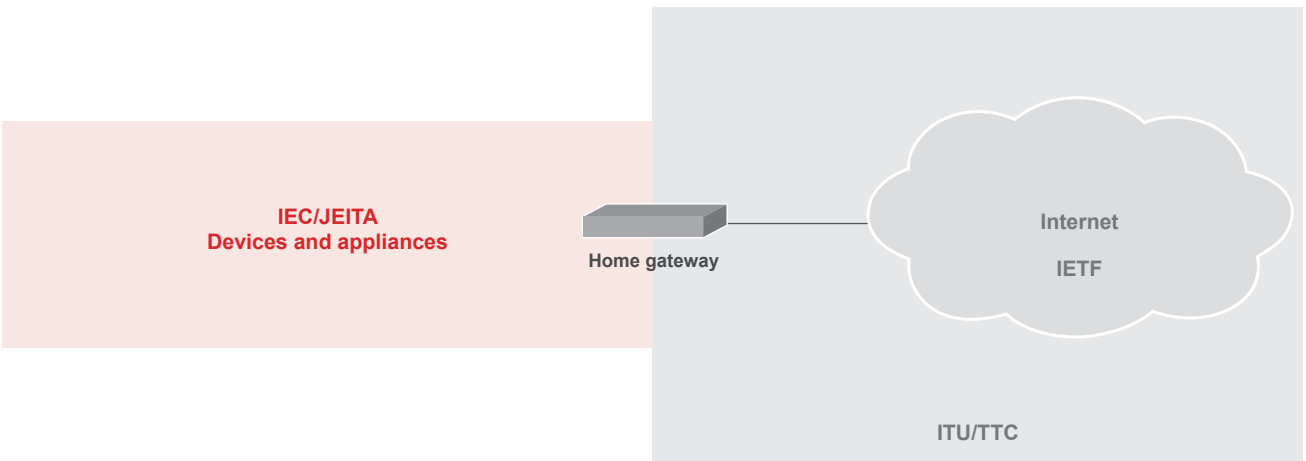

Figure 1: Domains of Standardization and Where This Standard Fits

Forum), but TA 18 turned it into an international standard, which it manages as the IEC 62680 series.

Recently, as part of its SDG-driven push to cut electronic waste, the European Union issued an amendment to the Radio Equipment Directive (Directive (EU) 2022/2380) unifying the charging port for certain electronic devices on USB-C. This mandates the USB-C charging interface for certain portable electronic devices. So instead of keeping a different charger and cable for every gadget, people can reuse the USB-C gear they already have. The technical reference standards cited are the IEC 62680 series of international standards.

IEC 62608 was also developed as an international standard for the logical management of home networks, with this kind of end-user connectivity in mind.

1.5 IEC 62608: Configuration Management for Multimedia Home Networks

IEC 62608 defines a "basic reference model" for managing a home network made up of diverse devices as one integrated system.

In this series, Part 1 (System Model) was published in 2014, and Part 2 (Operational Model) in 2017. But much time has passed since those first editions, and the world around home networks has changed dramatically with the rise of cloud services, the spread of wireless LAN, and so on.

As such, work to revise Part 1 began in 2021, and the latest edition, IEC 62608-1 Ed. 2.0:2025, was published in 2025. This section outlines the basic reference model, the background to the revision, and the technical features.

As Figure 2 shows, the standard's system model separates functions on the network into two logical entities: the side that manages, and the side that is managed.

- Configurator: The entity that keeps track of configuration information for the entire network, issues configuration instructions to, and checks the status of, each device. In Ed. 2.0, it is envisioned that this Configurator role will cover not only the home gateway but also cloud servers, smartphones, etc.
- Agent: Software that runs on a device, answers Configurator requests with the device's own attribute information, and changes its configuration as needed. Agents are intended to run on individual multimedia devices such as TVs, audio equipment, and white goods.

Note that in this standard, Configurator and Agent do not refer to specific devices or products. They are defined purely as roles. So a single device can take on multiple

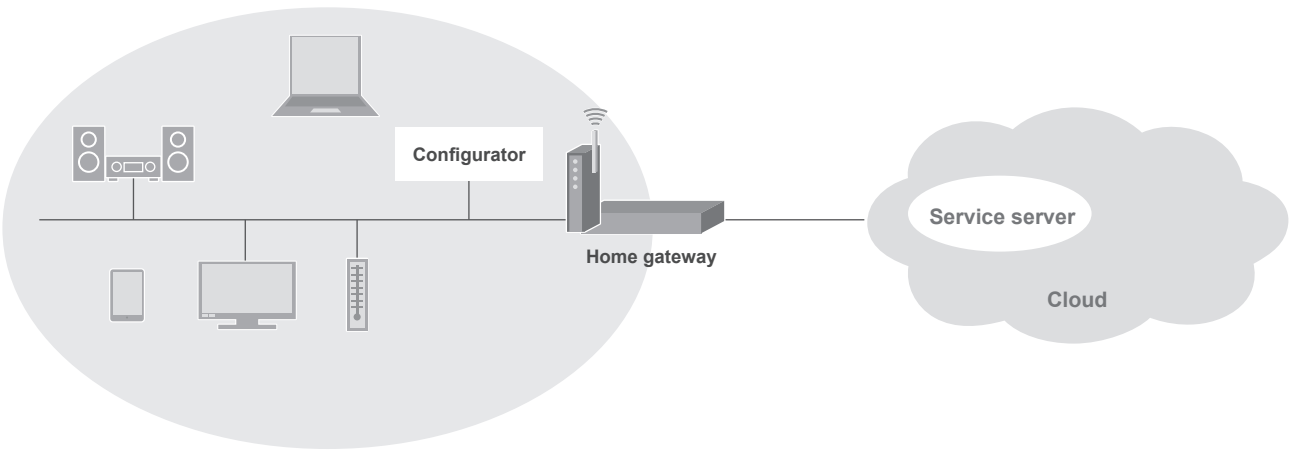

Figure 2: The IEC 62608 Reference Model

roles, and the delegation of roles can shift depending on how the system is used. This approach is intended to accommodate changes in devices, network environments, and usage patterns.

Now, when the first edition was published, wireless LAN was not all that widespread, so the standard covered wired networks only. Cloud-based data sharing was not commonplace either, so it went unmentioned. In those days, sharing data across home networks meant finding some way to connect into the home network itself.

The revision brings the standard in line with reality, with wireless LAN now in scope, and it introduces the concept of service servers designed for cloud use. Configuration management is thus no longer confined to the in-home network. Relationships with externally located functions can be covered too. The revision also addresses an issue in device–network relationships whereby it was easy to conflate the physical entity (the device) with the function (what it accomplishes); it does this by introducing a layered structure that makes the roles clear.

To illustrate specific applications, use cases were also added as an annex. Figure 3 shows an example of using the IEC 62608 reference model for automatic configuration of audio equipment in a studio. This provides a unique perspective that could only come from WG 31 (formerly TA 18) as a collection of multimedia and audio equipment experts and was created with ideas and advice from the group members.

1.6 Background to Formulation of IEC 62608 and What's Next

The idea behind IEC 62608 actually goes back to around 2002. At the time, XMLCONF (now NETCONF), a protocol for remotely configuring network devices, had been proposed at the IETF, the Internet standards body. When this was shared with the members of TC 100, the idea was put forward that home networks would surely need the same thing at some point.

But back then, when individual devices were each configured by hand, it was hard to convince people even within the IETF that XMLCONF was necessary, which made standardization slow going. It wasn't common for home appliances to connect to the network at the time either, and people generally didn't see why they would need to, so this was a time when adding networking features could actually hurt a product's sales. So while a handful of forward-looking members kept working on it, understanding was slow to develop on the TC 100 side as well.

More than 20 years later, both the Internet and the environment around home networks have been massively transformed. The Internet now has NETCONF standardization and automated operations management, while home networks have gained widespread wireless LAN, the arrival of the cloud, ubiquitous smartphones, and networked appliances. The landscape is utterly different. My sense is that we're now entering the phase in which home network management will be truly needed.

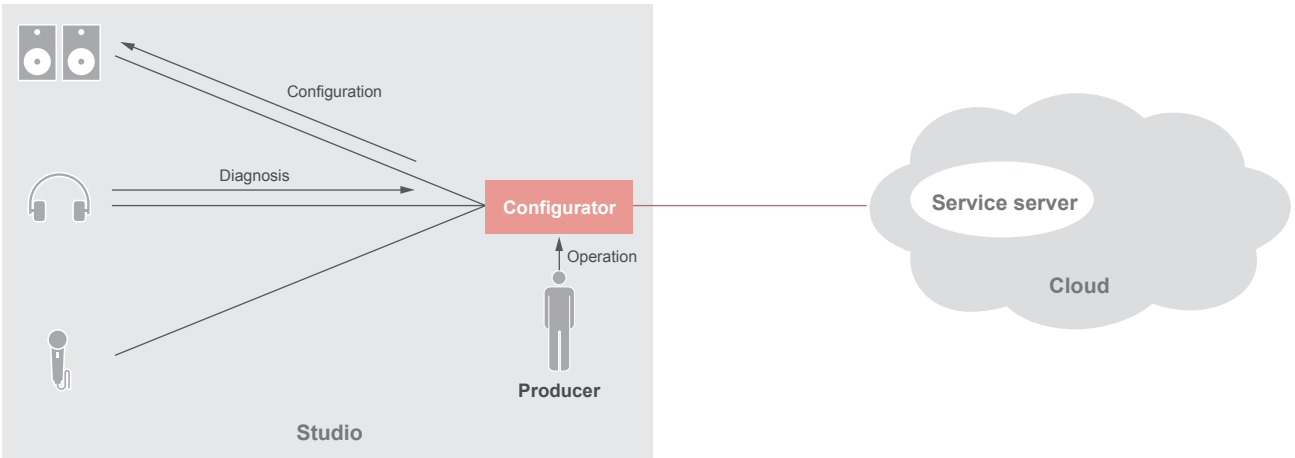

Figure 3: Studio Configuration Use Case

The revision of Part 1 of IEC 62608 is complete, but Part 2 (Operational Model) has also fallen behind the times and needs revising. Part 2 defines procedures for gathering information and applying configuration for home network management, and it corresponds closely to TTC TR-1062, "Customer support use cases for home network services."

So revising Part 2 creates the need to revise TR-1062 at the same time. Once again we'll be coordinating across two fields, but we share the same motivation, which is to make operations management easier. With the many people involved working together, I am confident we can move forward.

A new standard called Matter has recently appeared in the smart home world. It has brought a great variety of appliances, sensors, and even door locks onto the home network. Matter's connection procedures and applications are polished and intuitive, making it simple to operate appliances and view sensor data. Users may no longer even give a thought to the network at all.

Yet the more such devices can be connected without any thought for the network being needed, the harder it becomes to troubleshoot and respond in the absence of information when, inevitably, something won't connect or won't work. Network trouble, in particular, is rarely confined to a single device since it often arises from a confluence of multiple causes spanning software, configuration, and other factors. Having configuration information based on IEC 62608 available in such instances would no doubt aid in diagnosis and recovery. I believe this standard will serve as an unsung hero, quietly underpinning the reliability of home networks.

Finally, a word about international standardization. Companies that develop innovative technologies in-house while also leading the push to standardize them have been emerging of late. Standardizing can confer a market advantage in some situations. I feel that tracking standardization activity will be key when surveying new technology trends ahead.

Standards bodies differ not only in terms of creating de jure versus de facto standards. The way you are expected to participate in each also differs, and they each have their own procedures and culture. Looking beyond those differences, they also build cooperative relationships with one another through liaisons. We intend to keep working on this standard in cooperation with partners on many fronts.

Rei Atarashi
R&D Planning Group, Research Laboratory, IIJ
Aiming to realize a human-friendly Internet, Dr. Atarashi is engaged in research & development and standardization activities related to automated infrastructure designed to be easy for everyone to use.

Rolling Out OCP Servers: Our Path to Putting ORv3 into Production

2.1 Introduction

This article takes a look back at how we at Internet Initiative Japan Inc. (IIJ) came to view servers based on the Open Compute Project (OCP), and how, based on the lessons we drew from Open Rack v2 (ORv2), we came to put Open Rack v3 (ORv3)-generation servers into commercial use.

OCP is an effort to publish specifications for servers, racks, power supplies, and the like so that multiple manufacturers can build products to the same spec. Open Rack refers to the part of those specs that defines racks and power systems.

IIJ weighed ORv2 against a number of considerations before rolling it out in production: Was it suited to the power and rack specs in our own data centers? Would the hardware be available domestically on an ongoing basis? Could our operations teams handle it without trouble? Was the cost reasonable?

With ORv3, too, we waited until the full range of products was available and the path to deploying the hardware in Japan was clear, and we then started a proof-of-concept study (PoC) and moved to commercial deployment in 2025.

Below we walk through changes in our thinking on server procurement, what we learned from ORv2, and how evaluating ORv3 led to its commercial deployment in 2025.

2.2 How IIJ's Thinking on Server Procurement Has Evolved

During the mid-1990s, IIJ ran servers that combined PICMG CPU boards and I/O boards. Rather than using prebuilt servers, we procured custom servers made up of the components we needed.

Once the 2000s arrived, the options for PICMG-related products shrank while the number of servers we were procuring grew, so our old approach became hard to sustain in terms of both lead times and cost.

We then shifted to Sun's SPARC servers and enterprise x86 servers like NEC's Express5800 series and COMPAQ's ProLiant DL series. From around this time, we began to put more priority on ease of procurement, maintenance arrangements, and product completeness rather than the idea of optimizing individual components.

While it had become harder to sustain the approach of assembling servers to suit our own needs the way we had with PICMG, the need to think about the optimal server configurations for IIJ had by no means gone away. The reality was that, considering our procurement volumes and maintenance needs, we had little choice but to work within the specs that enterprise server vendors offered. Building servers to the specs we needed in-house was an option, but given the size of our operations, we could see that such an approach would not scale.

We reached a major turning point in 2009 with the launch of the IIJ GIO service. The initial rollout alone meant procuring roughly a thousand servers, and we had to design a system that held together as a whole, covering not just individual server specs but also power, racks, thermal design, and maintenance.

Subsequently, to handle the installation of over 5,000 units a year, we began bringing racks into the factory to mount the servers and do everything including cabling before then delivering complete racks. This made it possible to handle large rollouts while keeping on-site work time down.

We later learned about OCP servers, but at IIJ's scale of procurement, we were already able to keep costs low enough with enterprise servers, so we had no strong reason to opt for OCP right away. So at this point, we worked within our existing procurement model and put priority on finishing builds as fast as possible.

2.3 Why We Were Able to Take On ORv2

Around spring 2018, we received a proposal on ORv2 from Itochu Techno-Solutions Corporation (CTC), and this illuminated the potential for procuring the hardware more cheaply than conventional enterprise servers. So in September 2018 we visited an ODM in Taiwan with CTC to inspect its product lineup, development roadmap, and maintenance setup.

We also learned how the hyperscalers handle procurement. The motherboard and chassis do not account for all that much of a finished server's price, so for major components like CPU, memory, and SSDs, it often ends up being cheaper negotiating terms directly with the parts makers.

We had been introduced to a chipmaker's distributor right around that time, so during a trip to South Korea the following month, we also visited the company's headquarters there and worked out terms for memory and SSDs on the basis that the ODM would procure them on IIJ's terms. With another chipmaker, we ended up going through a distributor instead, but we were able to work out terms in the same way.

What makes this procurement model distinct is that you negotiate terms down to the component level on the premise of buying in volume and have the ODM source a configuration that fits your requirements. We believed this approach would work well at the scale of IIJ's operations and in the context of IIJ's data center operations, and that it would result in a reasonable cost structure.

2.4 The ORv2 System We Deployed and What We Learned Running It

The ORv2 system we deployed at our Shiroi Data Center Campus in 2019 accommodates 42 nodes per rack, and the centralized power supply draws single-phase 230V via four feeds. We based this configuration on an earlier deployment at a data center in Japan that we had toured around that time.

Data center floor space matters a lot in Japan, so IIJ focused on increasing rack utilization. It is not uncommon with typical enterprise server racks to limit the number of installed units to allow for peak power draw, but real workloads rarely use up that much power. As our racks are mostly for our own use, and because we had the ability to work with the user side to move hardware to other racks when needed, we avoided having too much design headroom for ORv2 and instead designed our system to fill the racks from top to bottom. We used a centralized power supply from Murata Manufacturing, and we had Nitto Kogyo Corporation build our racks to ORv2 spec at 700mm wide with front and rear doors. This fit the needs of IIJ's machine rooms while also leaving space for cabling inside the rack.

Running ORv2 in production also delivered clear results. We found that the configuration deployed at IIJ cut power consumption by roughly 30% versus off-the-shelf servers with similar specs. While there is no single factor behind this, within the scope of what we evaluated, the servers the ODM builds made it easy to choose the necessary and sufficient configuration for the task at hand, making it easy to keep power consumption down. OCP servers also use centralized power supplies, which raise power conversion efficiency, and we believe this made further power savings easier still. On the maintenance side, meanwhile, we found that a bit more work was involved in on-site fault isolation and replacement compared with pre-built servers, and that staff needed the skills to gauge the condition of individual components as they worked.

Some tasks, like CPU replacement, did mean more on-site work, but for management functionality, a serial console over IPMI was sufficient for us, so we encountered no major problems even with the simple management functionality provided by the ODM.

2.5 Why We Didn't Deploy ORv3 Right Away

The original plan was to keep using the same racks after deploying ORv2 and simply refresh the servers in them. In reality, though, socket form factor changes accompanying Intel's CPU generation updates, along with a market-wide shift in development resources toward AMD EPYC, meant that the range of ORv2-compatible products was gradually narrowing. It thus became harder to keep the ORv2 racks and simply refresh the servers.

Meanwhile, although the ORv3 spec had been published in 2020, the key supporting products, such as racks, centralized power supplies, and compatible servers, were not yet fully available at the time, so we were not in a position to make a decision on adopting it for our production environment. What mattered to IIJ was not that a standard existed, but whether we could procure the hardware within Japan on an ongoing basis, source maintenance parts in the event of failures, and handle everything comfortably within our data center operations. So when procuring new servers during this period, we chose enterprise rather than OCP servers. With ORv3, we decided to wait until a full range of products was available and the path to deploying the hardware in Japan was clear.

2.6 OCP Global Summit as a Turning Point

The turning point came at the OCP Global Summit held in San Jose in 2024. IIJ was looking for the ability to procure hardware within Japan on an ongoing basis, source maintenance parts, and comfortably deploy the equipment in its own data centers.

We had heard that Nitto Kogyo would be exhibiting ORv3 racks and Murata Manufacturing ORv3 centralized power supplies, so at that point we could already see the key building blocks coming together. Seeing the actual hardware on display at the summit helped us envision the prospect of actually assembling and evaluating the components in Japan, and we moved straight into a PoC upon returning home.

We borrowed a rack and centralized power supply of the same spec as those we had seen at the show, and we evaluated ORv3 servers from several manufacturers. We compared one-node-per-1OU and two-nodes-per-2OU configurations, ultimately adopting the latter for its cooling efficiency.

What we were looking at here was not just whether the servers held up on their own but whether we could handle them comfortably within our existing data center operations.

2.7 IIJ's Key Considerations When Evaluating ORv3

What we focused on most in our PoC was how the change in power delivery would affect real-world operations, whether the rack width and cabling would be viable, and whether we could put the servers into a production configuration as-is.

With ORv3, the voltage supplied to the servers changed from 12V to 48V. This was a natural change to make as power consumption climbed with CPUs, memory, and accelerators growing more powerful. The centralized power supply also no longer needs to sit in the middle of the rack, allowing a greater degree of design freedom.

For the PoC, we had Nitto Kogyo prepare a 600mm-wide rack with front and rear doors, but we found that plugging and unplugging cables and securing cable routes was a struggle. So we had them redesign the rack to 700mm wide, in line with ORv2.

The service outlets on Murata's centralized power supply disappeared when it moved to the ORv3 version. So for IIJ's setup, we needed a separate PDU to provide the power we previously drew from the service outlets. These decisions carried straight over into the production configuration described in the next section.

2.8 The ORv3 Production Environment We Deployed in 2025

With the PoC done, we began building out ORv3 at IIJ in July 2025. Our design accommodates 34 nodes per rack at 13kVA, with power supplied via two three-phase feeds. With ORv3, we again worked on the premise of racks being filled from top to bottom, so we kept rack utilization high while supporting higher power density than in the ORv2 era.

The ToR switch we adopted for this configuration is made by Juniper and is not ORv3 compliant. It thus needs a separate single-phase 230V feed, so we drew single-phase power off the two three-phase feeds coming into the rack. We had Harting make custom cables for this, and we connected a single-phase PDU downstream to supply power to the ToR switches. Our design balances the phases across the entire facility rather than within each individual rack.

2.9 How We Made OCP Server Maintenance Work

With OCP servers, it's not uncommon for recovery efforts to involve isolating faults at the component level, as opposed to the equipment being maintained as an indivisible finished product. Working within the ODM's warranty, IIJ handles initial on-site fault isolation and parts replacement in-house, while CTC takes care of returning failed parts and coordinating with vendors. Failed parts are sent to a repair center in Taiwan as needed.

This differs from how maintenance works with enterprise server vendors, but with this clear delineation of roles, we have made OCP server maintenance work in an actual production environment at IIJ.

2.10 Why We Limit OCP Server Deployment

At IIJ, we limit OCP server deployments to data centers we operate ourselves, such as the Matsue Data Center Park and the Shiroi Data Center Campus. This is because deploying OCP servers and Open Rack requires detailed advance coordination with the data center side on power systems, rack specs, and installation conditions. Also, while total rack power consumption is not pinned to its ceiling under real workloads, densely packed racks do produce localized hot spots. So to take advantage of OCP servers' high rack utilization, we had to work with

the data center side, including on cooling and placement conditions, to figure out just how far we could take the build-out. In some cases this means locally lowering cooling setpoints or adjusting airflow direction.

This equipment is not something you can install anywhere you want under the same conditions like you would with typical 19-inch rack servers. So it is more efficient to deploy at a decent scale than in small, one-off installations. This ties back to something IIJ has been aware of since around 2009: procuring in bulk increases the degree of design freedom.

2.11 Power Density Challenges in the SPR Generation and Planning the Next Configuration

While the 2025 commercial rollout was underway, IIJ was also working on its next configuration. The 4th Gen Intel Xeon Scalable Processor (Sapphire Rapids-SP, or SPR) generation, which IIJ had been evaluating since autumn 2024, came with performance gains but also increased power consumption, so per-rack power density and power delivery design mattered more than the performance of any single server.

So the question was not how the performance of individual CPUs compares but how to make the rack function properly as a whole as power density climbs.

For the CPUs we will adopt in 2026, we have chosen Intel Xeon 6 (Granite Rapids, or GNR) and 5th Gen AMD EPYC (Turin). Turin leads on power efficiency, but given CPU generation compatibility, we also needed GNR as a live-migration target from our current Intel Xeon servers.

For the chassis, we selected the 2U2Node format, which delivered good power efficiency in the SPR generation, for the GNR generation as well. Our decision was based not just on CPU performance in isolation but also on whether the configuration would work properly given our per-rack node count, power requirements, and the characteristics of our services.

2.12 Conclusion

In this article, we have taken a look back at how IIJ took on ORv2 and tested ODM-based procurement and whole-rack configurations, including centralized power supplies, in real operations. We then walked through how we built on the lessons from ORv2 and moved on to the ORv3 PoC and commercial deployment once the key building blocks were in place and the prospect of being able to handle the hardware comfortably under our operating conditions had emerged.

Along the way, we had to change how we work, not just the way we chose servers but also our approach to procurement and system building. On the procurement side, rather than simply comparing products from enterprise server vendors, we increasingly had to think about configurations and pricing with ODMs in mind. On the system building side, we also had to make the whole system work—not just the individual servers but the racks, power delivery, cabling, and maintenance operations as well.

The end result is that adopting OCP servers has broadened IIJ's options for procurement and system building. It has become clear, however, that one needs to be selective about where OCP servers are deployed and the conditions under which they are operated. Looking ahead, and with those constraints in mind, we hope to continue taking advantage of OCP servers as one of the options available to us at IIJ.

Masahiro Takahata

Manager, System Infrastructure Engineering Section, Infrastructure Engineering Department, Infrastructure Engineering Division, Network Services Business Unit, IIJ

Since joining IIJ in 2001, Mr. Takahata has been involved in the design and construction of server infrastructure. In 2008, he served as a core team member in the launch of NHN, the infrastructure underlying IIJ's services, and in 2009 he was responsible for building the infrastructure for IIJ GIO. He subsequently focused on operating the NHN infrastructure, and in 2011 he carried out the deployment of server infrastructure facilities in the containerized data center at Matsue Data Center Park. In 2019, upon completion of IIJ's Shiroi Data Center Campus, he worked on the rollout of ORv2 OCP servers, and since 2025 he has been working on the rollout of ORv3 OCP servers.

Validating Architecture for Real-world Deployment of Data Space Technology and Outlook for Practical Use

3.1 Introduction

As our digital society evolves, the conventional centralized platform model, under which vast amounts of data are collected in one place for management and use, is increasingly being complemented by a new approach that is drawing attention: the "data space." A data space enables autonomous, decentralized interconnection while preserving data sovereignty. This growing interest is driven by a fundamental need: although the ability to use data across organizational boundaries is now crucial, organizations also need to eliminate by design such risks as the leakage of trade secrets and the use of data for unintended purposes, thereby reconciling trust and control in the way data is shared.

These developments originated in Europe, where frameworks for trust-based data collaboration have been put forward through pioneering initiatives such as Gaia-X and International Data Spaces (IDS). In Japan as well, the drive toward real-world adoption is gaining speed, as efforts advance to establish the reference architecture known as ODS-RAM through collaboration among industry, government, and academia.

To verify the technical effectiveness of this new framework for data collaboration, we conducted a proof of concept (PoC) in April 2026. Our primary objective was to establish foundational technologies for enabling secure data collaboration in distributed environments while preserving data sovereignty. Specifically, we built an interconnection environment that combined an ODS-compliant data transfer module with the IIJ Cloud Data Platform and validated its effectiveness by integrating authorization mechanisms such as Relationship-Based Access Control (ReBAC). Building on the results of this work, we also examined the technical feasibility of extending this approach to domains requiring more stringent governance, with a view to future integration with advanced foundational technologies for data protection currently being developed as part of the K Program (NEDO)[*1].

This article presents the basic architecture of the data space technology we examined in the PoC, its specific implementation, and the PoC scenarios we designed with a view to enhancing the cyber resilience of telecommunications carriers, together with the findings obtained.

3.2 Architecture for Realizing Data Spaces

What defines a data space in technical terms is not the centralized accumulation of data, but rather collaboration among participants based on autonomous trust. This section outlines the core architecture of a data space and explains the role of each component.

3.2.1 Trust-based Collaboration and the Preservation of Data Sovereignty

The biggest feature of a data space is that it provides data sovereignty, under which data providers retain control over their data. In conventional data-sharing models, once you entrusted data to a platform, you had to delegate control over its use to that platform. In contrast, a data space introduces policy-based access control to establish a mechanism for usage control, whereby providers define the terms of use and recipients are required to comply with them through technical enforcement.

This is supported by a trust framework that extends across organizational boundaries. ODS-RAM places particular emphasis on the principle of Trust by Design (trust is not assumed but established through system design). Under this approach, even participants that do not have a fixed pre-existing relationship can initiate secure data exchange by proving their identities through verifiable credentials. Note, however, that an authenticated entity is not necessarily a trusted one. What matters is that the parties agree on the level of trust required for the relevant risks and intended use, and then combine measures such as verification, digital signatures, and trusted timestamps to build a trust structure that is fit for purpose.

3.2.2 Four-layer Architecture of a Data Space

In our PoC, we implemented a minimal configuration, referring to the layered structure defined by ODS-RAM, Japan's standard reference model. Conceptually, the architecture of a data space consists of the four layers described below. Because the layers are mutually independent and loosely coupled, the architecture allows the necessary layers to be

*1 Key and Advanced Technology R&D through Cross Community Collaboration Program (NEDO), Data no hogo to ryutsu no jidoka gijutsu no kenkyu kaihatsu [Research and development of automation technology for data protection and distribution] (https://www.iij.ad.jp/news/pressrelease/2023/1108.html, in Japanese).

implemented selectively according to the maturity of the domain and use case.

■ **L1: Data Layer**

This layer addresses usage control, protection against data tampering, and data quality. A key requirement is that the actual data remain within the provider's own infrastructure, allowing the provider to exercise usage control at its own discretion.

■ **L2: Transaction Layer**

This layer governs the entire transaction process between the provider and the consumer in a manner that is independent of data format and transfer method. In addition to physical transfer, it handles routing, automated policy enforcement, and the technical enforcement of agreed terms. These functions are performed by connectors deployed at the edge of each participant's environment in a distributed design that is fundamentally different from a centralized hub-based architecture.

■ **L3: Identity Layer**

This layer handles participant authentication and authorization. It provides credentials in a verifiable form and applies access control based on the terms of use set by the data provider. ODS-RAM breaks identity down into three elements: verification of real-world existence, authentication, and authorization. And it treats successful technical authentication and the decision to trust an entity as separate issues.

■ **L4: Semantics Layer**

This layer deals with addressability and semantics. By making metadata accessible, it enables interoperability through discovery, interpretation, and use of data across organizations.

3.2.3 How Autonomous, Distributed Collaboration Works

When the layers described above operate together, they enable high-trust data collaboration through the following process, without relying on a centralized platform.

■ **Dynamic Binding of Identity and Policy**

At L2 (the transaction layer), the usage policies defined by the data provider in L1 (the data layer) are dynamically matched against the participant attributes proven through L3 (the identity layer). As a result, even in the absence of a prior agreement with a given counterparty, access can be authorized securely based solely on real-time attribute verification.

■ **Enforcing Policy at the Edge**

A core feature of a data space is that the connector located at the point of data egress (the edge) performs the L2 role and physically enforces policy at that point. By making authorization decisions and applying controls before data leaves the provider's administrative domain, it prevents the provider from losing control over the data and ensures data sovereignty by technically enforcing terms of use even after the data has been handed over.

When these mechanisms work together, participants can manage their data autonomously while forming a "network of trust" in which they share information only with the parties they need to, only when they need to.

3.3 Design and Implementation of the PoC System

To implement the architecture described in Section 3.2 for our PoC, we built a demonstration environment by interconnecting an external ODS-RAM-compliant data space infrastructure service with our own service assets. For this phase, we adopted a minimal configuration aimed at rapidly assessing technical viability, so of the architecture's four layers, the implementation focused on L1 (the data layer), L2 (the transaction layer), and L3 (the identity layer). L4 (the semantics layer), which is responsible for defining data semantics, was therefore excluded from the scope of this PoC. Cross-cutting design elements such as governance, security, and trust (Trust by Design) were considered only at the level of basic design guidance. Below we describe the implementation structure in detail, including how each layer's functions were mapped to specific components.

3.3.1 Trust Model and Authorization Control in the Identity Layer (L3)

In implementing the identity layer (L3), we used an external infrastructure service as a trust anchor for managing the root of trust. For the authentication process in this PoC, we adopted OpenID Connect (OIDC) using Keycloak and built a mechanism for verifying the legitimacy of participating entities. For authorization control, we implemented Relationship-Based Access Control (ReBAC) across L2 and L3, enabling dynamic authorization checks based on inter-organizational relationships and privileges. By linking authorization information with an external infrastructure service compliant with industry-standard specifications, we verified that this configuration can support future interoperability across multiple domains.

3.3.2 Implementation of the Data Collaboration Platform and L2 Connector

For the connector function at the core of the transaction layer (L2), we adopted a data transfer module developed in accordance with the ODS L2 specification. Complementing this, our cloud-based data collaboration platform (the IIJ Cloud Data Platform) acts as an intermediary between the L1 data source and the L2 transfer module, providing advanced hub functionality to facilitate seamless data collaboration. Their respective roles are as follows.

■ Policy Enforcement

The L2 data transfer module evaluates, in real time, authorization information from the L3 infrastructure against the provider's usage policies. By ensuring that data is sent only to authorized recipients and only under the specified conditions, it serves as a technical barrier for maintaining the provider's data sovereignty.

■ Interface Abstraction and Push-Notification Handling

The IIJ Cloud Data Platform abstracts away interface differences across individual L1 data sources and gives the L2 module a unified HTTP API for data access. And to complement the inherently pull-based communication model of the data space, it also handles push notifications that instruct the counterparty to retrieve data, thereby enabling a practical data collaboration process.

3.3.3 Integration of Backend Systems and the Data Layer (L1)

For the data layer (L1), anticipating future integration with operational management systems, we used a dataset consisting of a subset of equipment information. A key design point in our implementation is that it dynamically extracts and provides only the data required through the data collaboration platform described above, without directly modifying the existing business systems. With this mechanism, we confirmed that our approach can quickly and securely make static data stored in existing systems available and usable as dynamic assets within the data space, while maintaining confidentiality.

3.3.4 Access Control and Future Extensibility

At the boundary with the data space, we implemented an access control mechanism that combines the L2 transfer module with the L3 authorization mechanism (ReBAC). This mechanism allows only requests authorized by the L3 infrastructure to pass, providing robust protection against unintended disclosure to third parties. Based on the fit-for-purpose trust structure described in Section 3.2.1, it ensures data trustworthiness by combining encryption of communication channels with strict verification of authorization information.

The L3 authorization mechanism we adopted is general-purpose enough that it can also be referenced by components other than the L2 modules. As such, in the future we expect it to enable flexible integration with IIJ's other services, as well as with the advanced foundational technologies for data protection being developed as part of Japan's K Program (NEDO). By further advancing the outcomes of this research and integrating technologies that more effectively achieve both data confidentiality and usage with policy-control and access-control technologies that govern control across organizational boundaries, it may be possible to develop this into a data distribution platform that combines security and flexibility even in domains where stricter governance is required.

3.4 Verifying Technical Effectiveness Through PoC Scenarios

To verify the value that data space technology can deliver, we conducted a PoC assuming a wide-area telecommunications carrier equipment data space.

3.4.1 The Need to Strengthen Cyber Resilience Through Wide-Area Equipment Collaboration

As attacks on telecommunications infrastructure grow more sophisticated, it is becoming difficult for any single operator alone to take steps to manage risks across the entire supply chain. In particular, when it comes to understanding vulnerabilities in equipment spanning organizational boundaries and the sharing of maintenance status, information silos and delays in information sharing pose structural challenges. In this PoC, we aimed to solve these challenges through autonomous information linkage.

3.4.2 Improving Maintenance Processes via Attribute-Based Access Control

In the first scenario, we focused on identifying equipment subject to maintenance updates between a telecommunications carrier and its partners.

Ensuring that End of Service Life (EOSL) is not overlooked is essential not only for avoiding risk, but also as a basis for making timely decisions on strategic system refreshes that ensure business continuity and adaptation to the latest technology.

In the PoC, we confirmed a process for selectively disclosing and obtaining data within the scope authorized based on attributes (credentials). This confirmed the practical utility of accelerating the identification of equipment requiring updates and enabling planned capital investment.

3.4.3 ynamically Sharing Support Information Using a Chain of Trust

In the second scenario, we tested a process for reliably delivering vendor-issued vulnerability information to operations personnel.

Here, we built a flow in which those who need the information obtain it at the time they need it via the data collaboration platform, with information authenticity being guaranteed. We confirmed that this eliminates the degradation and delays that come with passing information through human hands, minimizing lead time in the initial response to an incident.

3.4.4 Summary of PoC Results

Through this series of scenarios, we confirmed that it is indeed possible to implement systems that provide both confidentiality and transparency (controlled transparency), such that, based on attributes, data is disclosed only to those meant to see it, while preserving data sovereignty. The main results of our PoC are as follows.

■ Autonomous Trust Relationships Shorten Collaboration Lead Times

Sharing equipment data across organizations has traditionally carried heavy coordination costs, such as those involved in signing individual NDAs and building dedicated networks. In this PoC, we combined attribute-based authentication on the L3 platform with policy enforcement by L2 connectors to create a flow that pulls in data from trusted parties instantly and securely. This provides a foundation for rapid decision-making across organizational boundaries in emergencies such as cyberattacks.

■ Making Data an Asset While Leveraging Existing Systems

We showed that the existing production asset, the Unified Operations Management service (UOM), can be adapted for use in a data space solely through on-demand extraction via a connector, without major modifications. For infrastructure operators with vast legacy assets, this is a realistic and highly effective way to move toward the latest in distributed data collaboration while protecting existing investments.

■ Establishing a Starting Point for Improving Resilience Across the Entire Infrastructure

Configuration and vulnerability information that had been siloed within individual operators can be loosely linked through the data space to the extent needed, opening a path to visualizing risk across the entire supply chain in near real time. This presents a new form of data collaboration, one that connects together the measures taken by individual operators and raises the cyber resilience of social infrastructure as a whole.

3.5 Technical Considerations and Future Outlook

While our PoC confirmed the effectiveness of data spaces, it also revealed the challenges and prospects involved in implementing them as wide-area social infrastructure and building a sustainable ecosystem around them.

3.5.1 Technical Challenges for Real-World Deployment and Advancing the Trust Model

To improve practical utility further, we first need to solve technical challenges around the dynamics of data distribution and sharing of semantics.

The first challenge is implementing an event-driven delivery mechanism that improves the timeliness of information. In addition to the current on-demand data extraction, we need a mechanism whereby specific occurrences (events) trigger connectors to autonomously push information out to the relevant organizations.

This would be useful not only for incident response tasks, such as detecting equipment anomalies or discovering vulnerabilities, but in any use case in which information needs to be shared without delay, such as changes in inventory levels or environmental sensors exceeding thresholds. It would broaden the scope of what data spaces can do, from static collaboration, where you actively go and fetch information, to dynamic collaboration, whereby the information you need arrives at just the right time.

The second challenge is standardizing semantics to automate the interpretation of data across organizations, which was outside the scope of this PoC. As a data space grows,

differences between the vocabularies independently defined by each participant and cases in which different organizations assign different IDs to the same entity end up becoming structural obstacles to collaboration. Effective approaches here are dynamic mapping based on a common ontology for the former, and the assigning of global identifiers using DIDs (Decentralized Identifiers) for the latter. What matters in both cases is that participants can achieve interoperability without changing their own existing definitions or IDs. Combined with the event-driven delivery described above, this becomes a foundation that supports both the distribution of data and its interpretation and identification as a unified whole.

To undergird these information dynamics and shared semantics, we must progressively advance the trust model and build the institutional frameworks that support it. In our PoC, we adopted a community-based model that uses a specific authentication infrastructure as its trust anchor. The next technical milestone is to build on this and extend it into a more autonomous, decentralized trust model based on VCs (verifiable credentials).

A key part of this process will be mechanisms that go beyond simply verifying technical signatures to dynamically determine who has vouched for a given credential and against what criteria. Specifically, we will need a function like the Clearing House proposed in Gaia-X, which verifies participant conformance and records and audits transaction evidence, together with a federated catalog function that catalogs data assets and policies and shares them among participants.

In such mechanisms, defining governance as in who serves as the trust anchor (the root of trust) goes hand in hand with technical implementation. By creating an environment in which both providers and consumers of data can assess and reference dynamically changing trust levels under a common framework, we will be able to maintain trustworthiness even as the number of participating organizations grows, while dramatically reducing interconnection costs.

3.5.2 Building a Data Ecosystem That Creates Value

If data spaces are to become a sustainable framework rather than ending as one-off demonstrations, they must evolve from mere venues for exchange into ecosystems that turn the information flowing through them directly into real business value. To make this happen, it will be crucial to design a structure that facilitates independent participation by not only those who provide and obtain data but also third parties who create value on top of it (application developers, analytics service providers, and the like).

As European precedents show, one effective step is to build marketplaces that offer not just the data itself but also analytics capabilities (services) as a package for processing and analyzing the data. The trouble with data provision and use alone is that participants struggle to see any direct benefit, so it is hard to generate motivation that outweighs the cost of connecting. Catena-X tackles this problem by limiting the applications and services listed on its marketplace to certified offerings only, giving participants an environment they can use with confidence and without the burden of vetting vendors. This lowers the barrier to connecting, and it also drives network externalities, such that as the number of participants increases, the data become more diverse and the analytics become more accurate.

Standardized on-ramps like KITs (Knowledge, Implementation, Tooling) are also an important way to encourage third parties to enter the ecosystem. By packaging the technical specifications, implementation guidelines, and tools needed for each use case, they let application developers join the ecosystem in a standards-compliant way without having to start from scratch with their own research and individual negotiations.

And to support the healthy growth of this ecosystem, guaranteeing the neutrality of the underlying technology stack is also crucial. Building on open source software (OSS) components that do not depend on any particular vendor, and on open standard specifications that anyone can conform to, ensures fair connection opportunities for participating organizations. Ensuring this neutrality and transparency encourages diverse stakeholders to join with confidence, broadens the pool of those creating value on the platform, and drives the ecosystem's self-reinforcing growth. Offering small and midsize organizations a low-cost option for connecting via managed services is another essential element for broadening the ecosystem base.

3.5.3 Prospects as General-Purpose Infrastructure and Convergence with AI

Just as the Internet once removed the barriers to communication and transformed the structure of society, data spaces have the potential to drive the evolution of data usage from

a special undertaking into a standard operational foundation. For this to happen, we will need to see progress on technical standardization as well as widespread awareness of the real benefits that participants gain by connecting to a data space.

As a step toward real-world deployment, a realistic approach is to start with closed, community-based trust within a specific industry and among parties with shared objectives, building trust relationships grounded in operational rules. From there, gradually shifting and expanding to general-purpose trust based on machine verification using verifiable credentials (VCs) should ensure interconnectivity between different ecosystems, such that data spaces grow into a truly open data distribution platform.

This evolution of the trust infrastructure will further accelerate deep convergence with AI technology. With large language models (LLMs), for example, we can now envision the prospect of building environments in which a question posed in natural language safely and automatically retrieves and integrates the necessary information, without the need to think about complex policy settings or data mappings. And if general-purpose trust provides machine-verifiable guarantees of data provenance and rights, the trustworthiness of AI-generated output will also rise dramatically.

Meanwhile, data spaces will also be a vital foundation for AI, because an environment in which high-quality data of clear authenticity flows across organizations will have a direct impact on improving the training accuracy of and ensuring governance over AI models. Once data flowing across organizational boundaries becomes the norm, this will surely create fertile ground for the creation of new value, not just in terms of strengthening cyber resilience but in all sorts of areas, including carbon neutrality and supply chain optimization.

3.6 Conclusion

Our PoC was an attempt to put the data space concept into practice in telecommunications infrastructure, a domain that demands high reliability, as a step toward realizing next-generation data collaboration. The insights we have gained go beyond the challenges of any specific industry and point toward the sort of flexible form that data distribution should take going forward.

Our aim is to bring about a society in which data is used safely and efficiently in a way that is trusted, unfettered by organizational boundaries. To achieve this, we will deploy advanced technologies such as AI to simplify implementation, opening the door to a future in which the new infrastructure of the data space sees widespread adoption and ongoing development as an open service freely available to everyone.

At IIJ, we will continue to work with stakeholders on efforts to bring about a new data-driven society, in terms of both technology and services.

Yasushi Toriumi

Business Development Office, Enterprise Sales Unit, IIJ
Since joining IIJ, Mr. Toriumi has worked primarily as a project manager in the field of system integration, launching solution businesses aimed at expanding the integration business. He currently serves as a technology strategist in the Business Development Office, pursuing initiatives to bring about a data-driven society, mainly through the development and utilization of data distribution platforms.

August 2026 Vol.70

Internet Initiative Japan

About Internet Initiative Japan Inc. (IIJ)

IIJ was established in 1992, mainly by a group of engineers who had been involved in research and development activities related to the Internet, under the concept of promoting the widespread use of the Internet in Japan.
IIJ currently operates one of the largest Internet backbones in Japan, manages Internet infrastructures, and provides comprehensive high-quality system environments (including Internet access, systems integration, and outsourcing services, etc.) to high-end business users including the government and other public offices and financial institutions.
In addition, IIJ actively shares knowledge accumulated through service development and Internet backbone operation, and is making efforts to expand the Internet used as a social infrastructure.

IIJ-MKTG020-0068

Internet Initiative Japan Inc.

Address: Iidabashi Grand Bloom, 2-10-2 Fujimi, Chiyoda-ku,
Tokyo 102-0071, Japan
Email: info@iij.ad.jp URL: https://www.iij.ad.jp/en/