ADMIN
Network & Security

Focus Guide

openEuler

Discovering openEuler

AI and the BooM Stack

A complete software stack tailored for AI

Automating Migration

Easily migrate your old CentOS systems to the modern openEuler

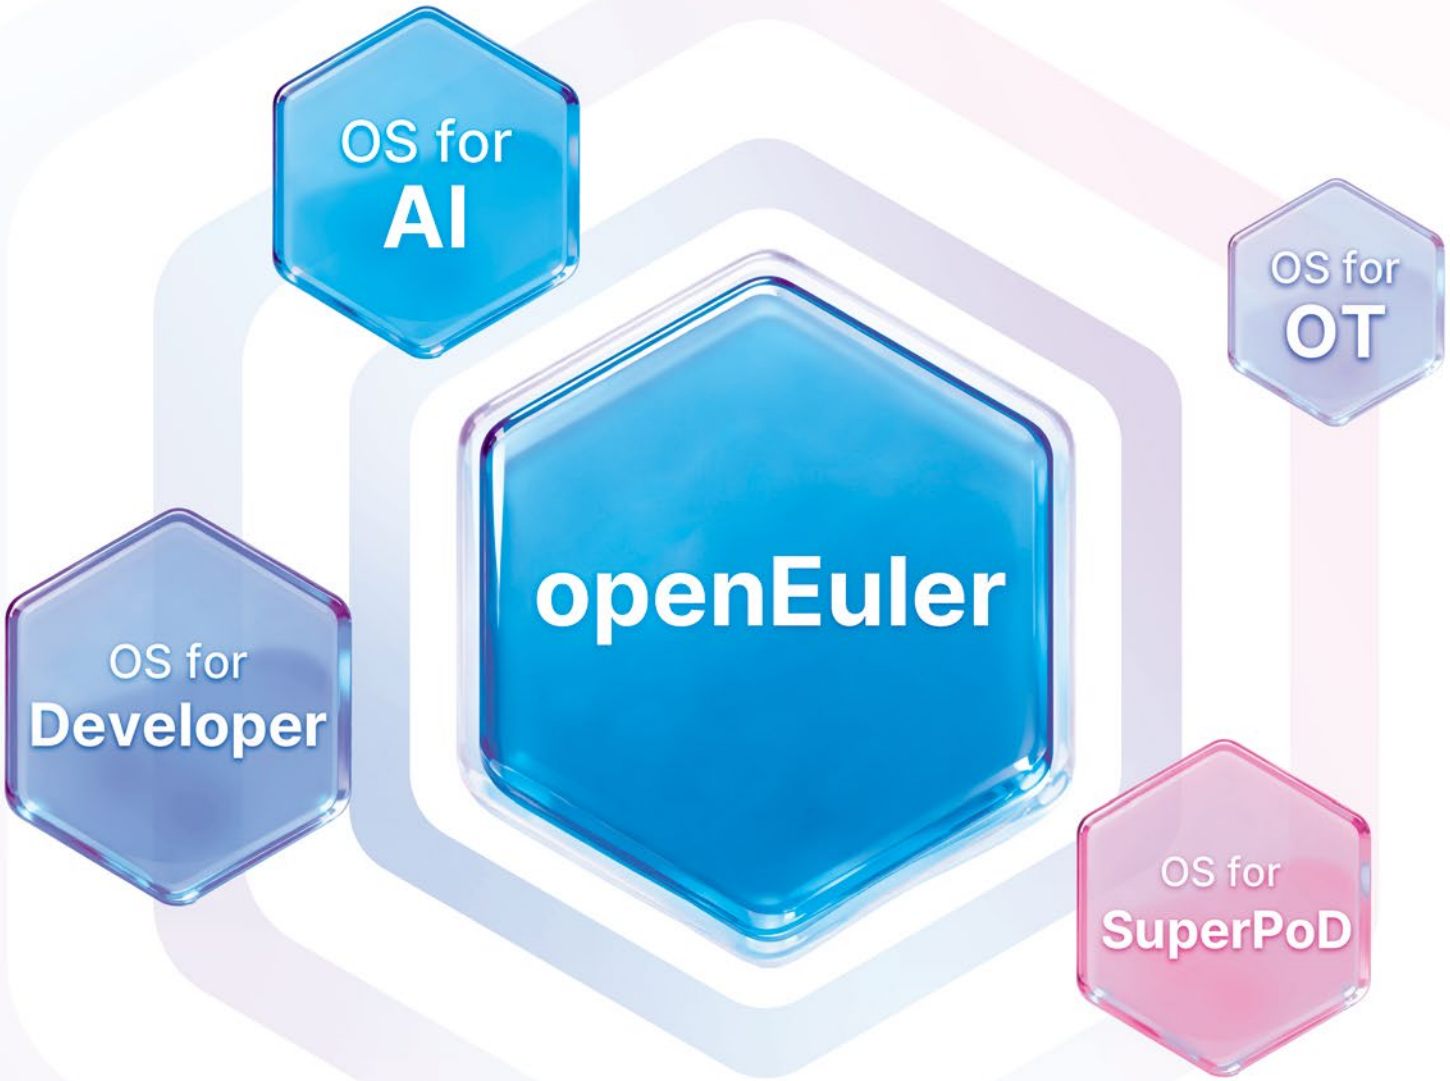
Interview

openEuler leaders and experts give an inside look at the project, its goals, and the path ahead

US$ 7.95

WWW.ADMIN-MAGAZINE.COM

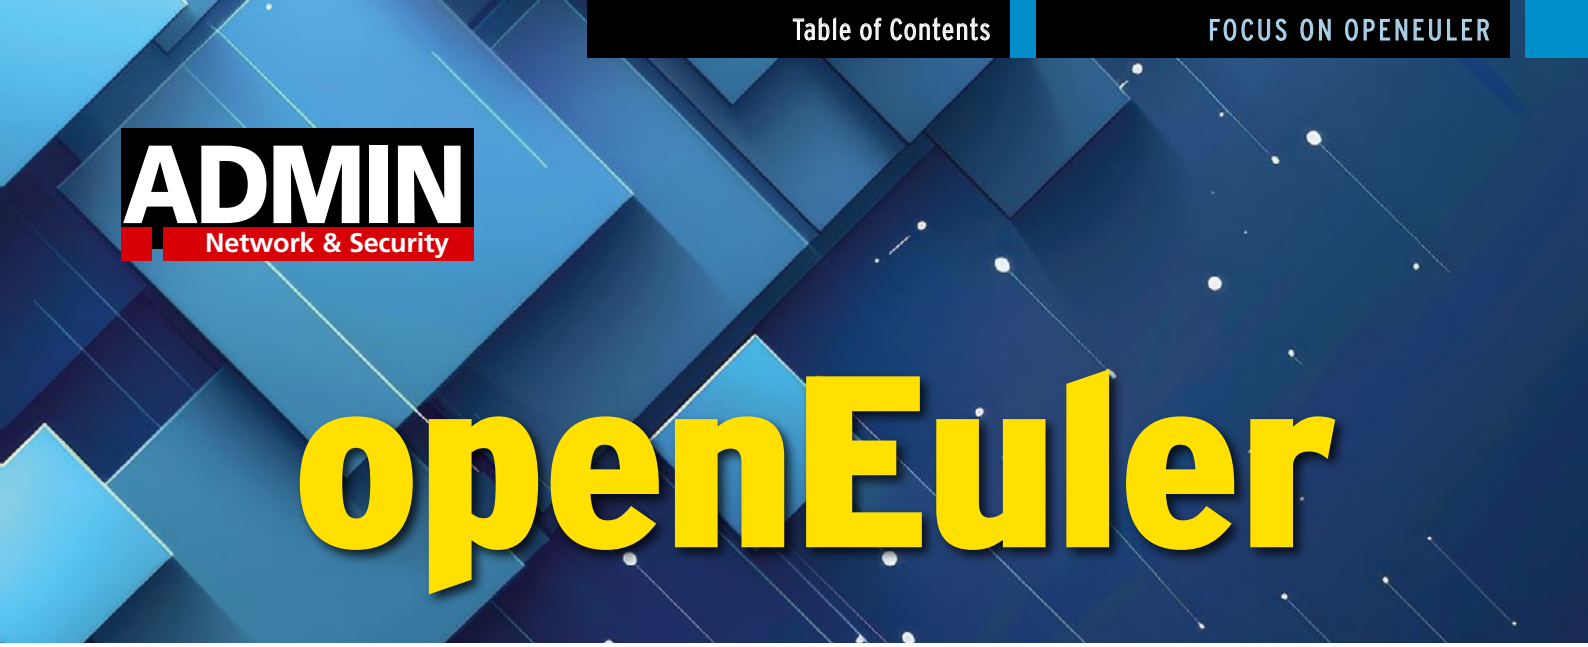
OPENATOM FOUNDATION | OpenEuler



Official Website

Rate on DistroWatch

contact@openeuler.io



Contact Info

Editor in Chief
Joe Casad, jcasad@linuxnewmedia.com

Associate Editor
Amy Pettle

Copy Editor
Aubrey Vaughn

Layout and Graphic Design
Dena Friesen, Lori White

Advertising
Brian Osborn, bosborn@linuxnewmedia.com

Marketing Communications
Gwen Clark, gclark@linuxnewmedia.com

Publisher
Brian Osborn

Customer Service / Subscription
Email: cs@linuxnewmedia.com
Phone: 1-785-856-3080

www.admin-magazine.com

While every care has been taken with the content of the magazine, the publishers cannot be held responsible for the accuracy of the information contained within it or any consequences arising from the use of it.

Copyright and Trademarks © 2025 Linux New Media USA, LLC

Cover Illustration © bravespirit, 123RF.com

No material may be reproduced in any form whatsoever in whole or in part without the written permission of the publishers. It is assumed that all correspondence sent, for example, letters, email, faxes, photographs, articles, drawings, are supplied for publication or license to third parties on a non-exclusive worldwide basis by Linux New Media unless otherwise stated in writing.

All brand or product names are trademarks of their respective owners. Contact us if we haven't credited your copyright; we will always correct any oversight.

Published by Linux New Media USA, LLC, 4840 Bob Billings Parkway, Ste 104, Lawrence, KS 66049, USA

Represented in Europe and other territories by: Sparkhaus Media GmbH, Bialasstr. 1a, 85625 Glonn, Germany.

Dear Readers:

openEuler has been a major player in Asia for years, but users in Europe and the Americas are only beginning to discover it. The massive openEuler project has one of the largest developer communities in the world and often ranks near the top for Linux kernel contributions. This focus guide is your introduction to the openEuler universe. We'll show you what's inside openEuler and study some special tools for migration and artificial intelligence. You'll also hear from some openEuler community leaders and experts on the project's goals and the path ahead.

Table of Contents

4 Discovering OpenAtom openEuler
The versatile Linux system known as OpenAtom openEuler is equally at home in the server room, cloud, or edge. With a huge developer community and strong corporate support, openEuler is ready to step into the foreground as a leading enterprise Linux.

10 x2openEuler Migration
Easily migrate your old CentOS systems to the modern openEuler with the x2openEuler migration toolset.

15 AI and the BooM Stack
openEuler prepares for the AI future with a complete software stack tailored for artificial intelligence.

19 Interview
We sat down with three experts from the openEuler community for an inside look at the project, its goals, and the path ahead.



Versatile Linux with eyes on the enterprise

Discovering OpenAtom openEuler

The versatile Linux system known as OpenAtom openEuler is equally at home in the server room, cloud, or edge. With a huge developer community and strong corporate support, openEuler is ready to step into the foreground as a leading enterprise Linux. By Joe Casad

The name openEuler [1] is coming up more frequently in Linux circles. This enterprise-ready Linux is a major presence in Asia, and the project has recently been connecting with more users in Europe and the Americas.

The openEuler project began as EulerOS [2], an in-house system maintained by Chinese mega-vendor Huawei. The Huawei developers needed a system that would serve a number of different roles, from server room to cloud to edge to embedded systems, and they wanted something that would adapt easily to a new generation of alternative chip architectures. Eventually, to facilitate collaboration with industry partners, Huawei contributed the code to the community as an independent project.

The openEuler project continues with the goal of providing a system that adapts easily to different architectures and use cases. The aim is to offer a system that allows the developer to "develop once and adapt to many scenarios."

OpenAtom

The OpenAtom Foundation [3] is a nonprofit organization focused on developing and promoting the global open source community. The foundation sees its mission as an effort to enhance industry-wide cooperation and empower open source across a broad range of sectors. Flagship projects include openEuler and the IoT-focused OpenHarmony, but OpenAtom also supports several other important projects, including the openKylin desktop OS, the OpenTenBase database, and the OpenLoong humanoid robotics initiative. The organization provides an array of services for supported projects, including open governance, communication, outreach, legal assistance, funding, and technical support.

OPENATOM FOUNDATION

Lead Image © videst, 123RF.com

The open source project officially launched in 2019, and since 2021 openEuler has been part of the OpenAtom Foundation (see the box entitled "OpenAtom"). If you're wondering about the name, the distro is named for the mathematician and scientist Leonhard Euler [4], who gave his name to the mathematical constant e (base of the natural logarithm) and was known for contributing to several other fields of math and science. The rigor and diversity of Euler's work is an inspiration for the openEuler team to stay flexible and keep a wide-angle view of the roles their distro can play in the ever-evolving IT space.

Table 1: Some of the community-certified distributions based on openEuler.

Partner	System
xFusion Digital Technologies	FusionOS
UnionTech Software	UOS
SUSE	SUSE Euler Linux
H3C Technology	H3Linux
Beijing Linx Software	Linxos
TurboLinux Inc.	TurboLinux Enterprise Server
Kylinsoft	Galaxy Kylin Advanced Server
CSG Digital Power	Pegasus Server
iSOFT	iSOFT Server OS

Introducing openEuler

The openEuler project has a very large base of contributors, users, and sponsors. The project's close collaboration with a wide range of vendors underscores the emphasis on versatility and portability. The openEuler community includes hardware companies, chip manufactures, OS vendors, service providers, ISVs, and research institutions. Some companies and organizations maintain their own Linux distros based on openEuler, which means that openEuler serves as a kind of meta-distro (similar to Debian) for many in the global Linux space (see Table 1).

The openEuler project sees itself as a full-featured alternative to Red Hat, Suse, Ubuntu, and other enterprise distros. That might sound surprising to those in the West who are less familiar with the project and its history, but a large global development community and close ties with the hardware industry put openEuler in a unique position for addressing many of the use-case challenges surrounding industrial Linux.

According to the openEuler project, more than 2,000 companies and institutions have joined the community, and more than 20,000 contributors have participated. Contributors include major international companies such as Intel, ARM, and AMD, as well as other organizations, such as the cloud-native NestOS project and the Institute for Software Chinese Academy of Sciences (ISCAS). Huawei, an openEuler strategic partner, has ranked as a top contributor in several recent Linux kernel releases (Figure 1). The breadth and depth of the openEuler project is quite impressive. According to Xiong Wei, Chairperson of the openEuler committee, more than 550 new projects have been incubated within the openEuler

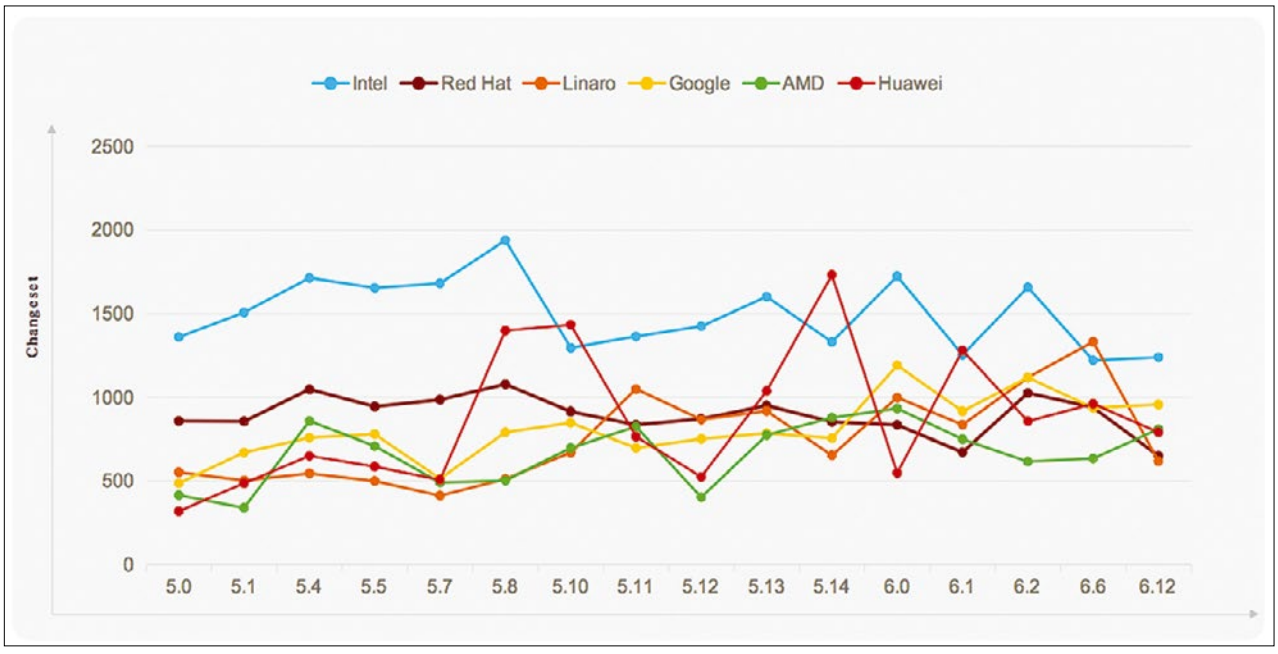


Figure 1: openEuler strategic partner Huawei is actively involved in Linux kernel development and frequently ranks near the top in code contributions.

community over the past six years. The developers believe their investment in AI, edge, embedded, cloud, and conventional enterprise settings puts them in a strong position to adapt to future digital and intelligent infrastructure (Figure 2). According to project leaders, openEuler's open collaboration, and its broad base of contributors, are its greatest strength.

Methods

openEuler puts out a new release every six months, and every two years they release a new Long-Term Support (LTS) edition. The LTS editions receive two years of comprehensive maintenance and support, and extended support is available for an additional two years. The emphasis of the LTS versions is on stability. The regular releases (dubbed "innovation releases") offer a chance for developers and users to try out the latest technologies.

The openEuler download page offers an indication of the project's versatility and scope (Figure 3). Choose an architecture (x86, aarch64, ARM32, LoongArch64, RISC-V) and choose a usage scenario (Server, Edge, Cloud, Embedded, or DevStation). Your selection will lead to other choices for hardware, image size, and other factors.

openEuler doesn't just accommodate this broad range of use cases – the community is actively engaged in developing and integrating tools for the server room, cloud, edge, and embedded space. The following sections highlight some of the advanced technologies built into openEuler.

AI

With its deep connections to the hardware industry and broad support within the developer community, openEuler is well positioned to serve as a platform for AI modeling and development. The openEuler project has made AI a priority and offers what they call "full stack" AI support. openEuler provides ready-made container images for the pytorch, tensorflow, and MindSpore AI frameworks. Other openEuler images offer support for the CUDA and CANN development platforms. You will also find images for installing Large Language Model (LLM) applications and toolchains, including Baichuan, ChatGLM, and iFLYTEK. openEuler's containerized approach provides a flexible and extensible ecosystem for optimizing AI development and workloads.

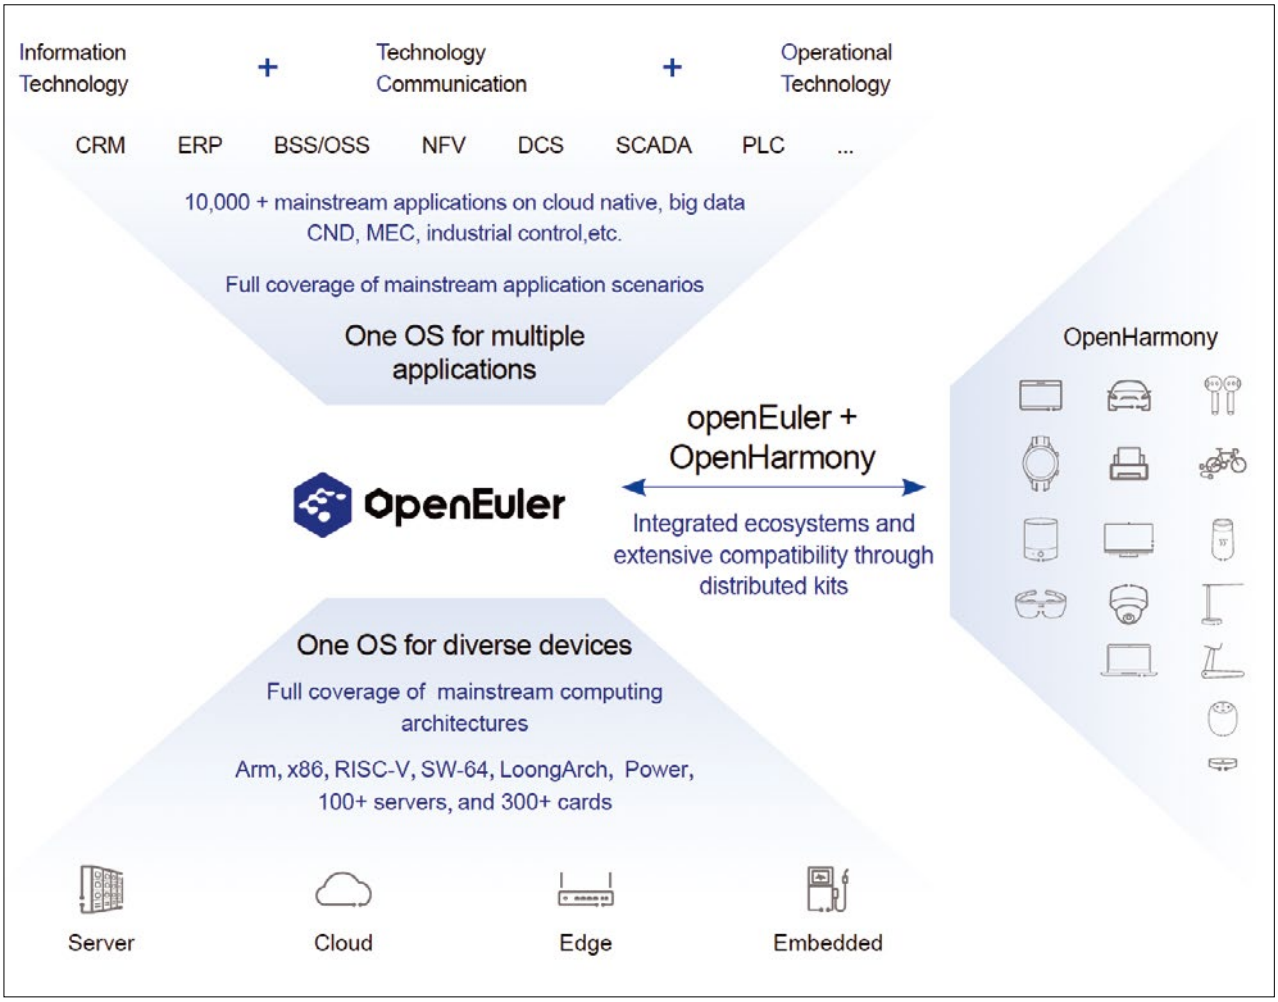


Figure 2: The openEuler project has the ambitious goal of supporting a broad range of components and frameworks across several hardware platforms and use cases.

openEuler also comes with the sysHAX LLM heterogeneous acceleration runtime, which enhances performance for transformer models such as deepSeek, Qwen, Baichuan, and Llama. SysHAX assigns inference tasks to CPUs to make optimum use of CPU resources in AI systems.

In addition to its ecosystem of AI frameworks and modeling tools, openEuler also offers an intelligent Q&A platform for streamlining code generation, troubleshooting, and Q&A. openEuler Intelligence uses multi-channel Retrieval-Augmented Generation (RAG) to bring AI techniques to AI development, giving developers greater control while improving efficiency. openEuler's AI development framework can help with:

- Data preprocessing and labeling – quickly generates Q&A data for AI model training.
- Document quality improvement – anonymizes, deduplicates, and standardizes documents for better information security and efficiency.
- Model building – fine-tunes to tailor a model for a specific task.

You can access openEuler Intelligence via the web or a shell interface.

Further down in the stack, the openEuler environment comes with a number of advanced components for supporting the data-intensive and compute-intensive systems that run AI applications (as described in the next section).

Server

The server room setting has undergone significant changes in recent years, and openEuler comes with several technologies designed to add efficiency and improve TCO. One example is DPUDirect, a cross-host collaboration framework that allows services to migrate between hosts for better flexibility and failover in a server room environment. DPUDirect responds to the expanding use of Data Processing Units (DPUs). DPUs, which are frequently used in AI and other data-intensive operations, combine features of a CPU and network interface to offload network-intensive tasks, improving performance in settings where complex computations are shared among multiple processors. DPUDirect facilitates offloading of the container or virtualization management plane, eliminating redundancy and improving efficiency in DPU-based environments.

openEuler also supports Enhanced Network File System (eNFS), an advanced version of NFS designed for Network Attached Storage (NAS) environments, as well as HPCRunner, a tuning assistant that reduces deployment costs and improves optimization efficiency.

EulerFS is a file system designed to work efficiently with non-volatile dual in-line memory modules (NVDIMMs). Although other file systems, such as the EXT4 file system, can work with NVDIMMs, heavy metadata management overhead can reduce the overall performance. EulerFS employs pointer-based dual directory views to reduce metadata synchronization overhead.

Cloud

Cloud technology is increasingly important in the IT space. openEuler supports a number of components designed for delivering smoother and more efficient cloud operations. HybridSched is a full-stack solution for hybrid deployment of VMs, providing enhanced OpenStack cluster scheduling and kernel-mode basic resource isolation (Figure 4). Skylark, which is incorporated with HybridSched, is a Quality-of-Service (QoS) aware resource scheduler that improves utilization when high- and low-priority VMs are deployed together.

KubeOS is a lightweight operating system that supports containerization in cloud-native scenarios. You can manage your KubeOS environment through Kubernetes, including atomic upgrades and API-based operations for both containers and node OS components. openEuler also

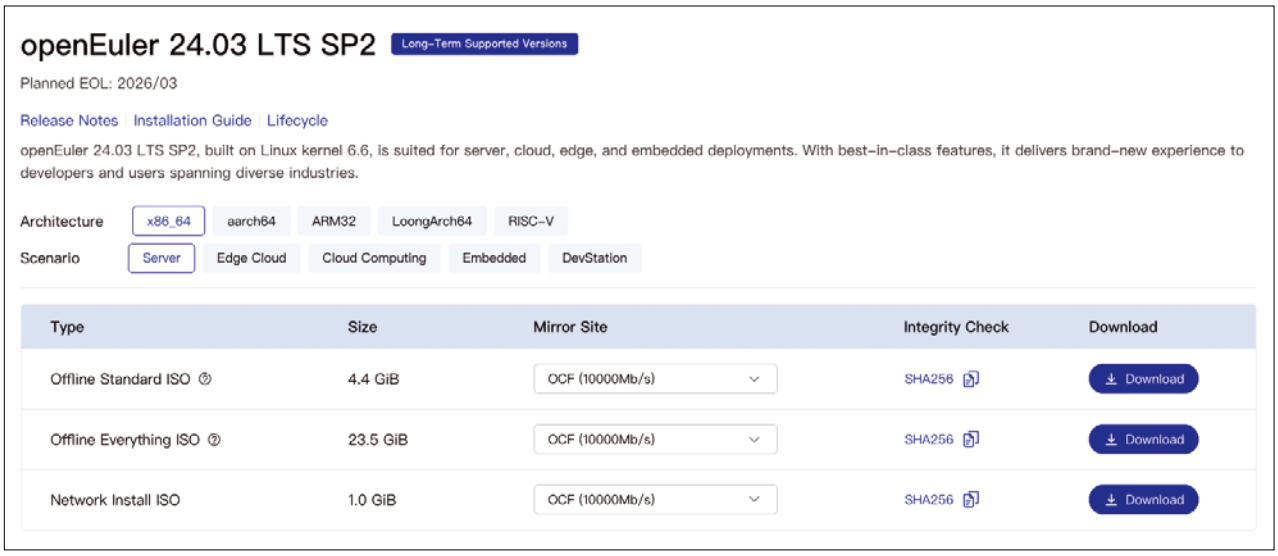


Figure 3: The openEuler download page.

supports Rubik, a hybrid container deployment engine that schedules and isolates resources to improve node resource utilization, ensuring prioritization for mission-critical services. NestOS is a cloud OS incubated in the openEuler community that offers secure management and integrates popular container engines such as Docker, Podman, and iSulad.

openEuler images are currently available directly within the AWS, Azure, and Alibaba cloud systems, and you can install openEuler as a community-provided image in the Google cloud.

Edge

Edge computing is an important topic for many organizations as the old paradigm of centralized processing gives way to exploding data volumes from self-driving cars, industrial sensors, IoT systems, and other smart devices. openEuler responds to the edge challenge by integrating KubeEdge, an open source management and provisioning tool incubated by the Cloud Native Computing Foundation. openEuler Edge integrates KubeEdge into a unified management platform (Figure 5) that streamlines AI deployments, implements service discovery and traffic forwarding, and improves processing and efficiency for edge devices.

Security Tools

Security is another important consideration, whether you manage a home network or a corporate colossus with 50,000 nodes. openEuler users can access all the powerful security components that populate the enterprise Linux space, and the openEuler developers highlight a few additional components that have received special attention.

The Integrity Management Architecture (IMA) is a mandatory access control subsystem built into Linux kernel 2.6 or later. IMA provides a deep level of file integrity protection, ensuring that boot files, system files, and other critical file assets are not subject to unauthorized change. IMA digest lists create a chain of trust that ensures processes operate without compromise.

SysMaster is an ultra-light, fault-monitoring, self-recovering replacement for the conventional init process that improves reliability and service availability without the complexity and long learning curve of systemd.

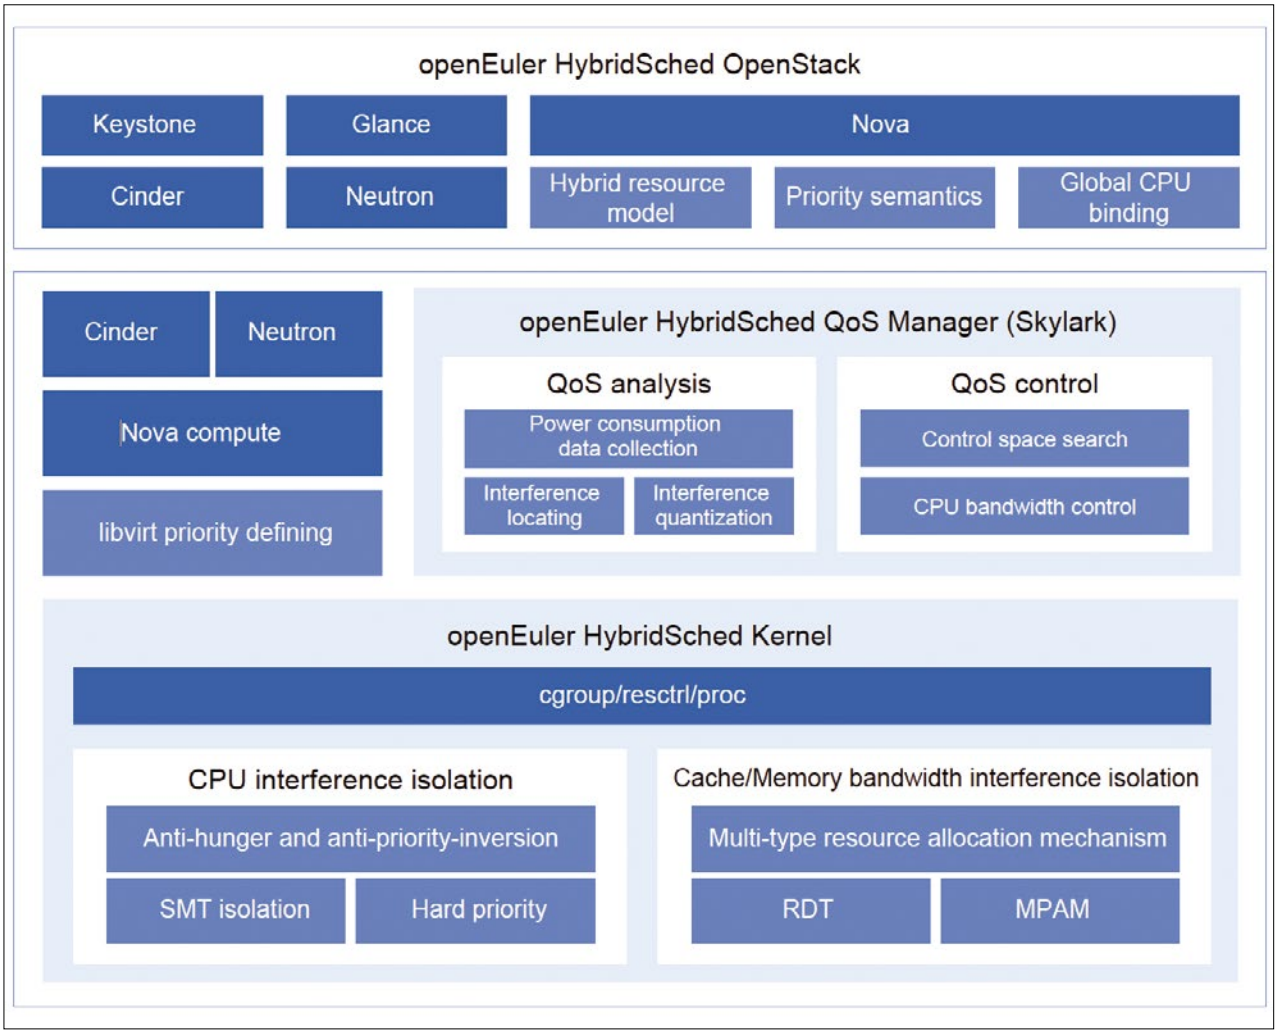


Figure 4: HybridSched supports OpenStack in hybrid environments.

openEuler also provides support for secGear, a secure application development kit, and secPaver, a tool for developing SELinux mandatory access control policies.

Desktop

When you first install openEuler, it lands you at the command prompt. The openEuler OS is used in a variety of server settings and is designed for efficiency, without assuming the need for a full desktop environment. If you are setting up openEuler on a user workstation, or if you prefer a GUI environment to the command prompt, the openEuler documentation describes the steps for installing a desktop environment. openEuler supports the Gnome, Xfce, DDE, and UKUI desktops.

Devices and Things

OpenHarmony is an OS used with embedded systems, smartphones, tablets, Internet of Things devices, and other systems. The OpenHarmony OS was also developed by Huawei and is maintained by the OpenAtom Foundation. openEuler is closely integrated with OpenHarmony [5] (as shown in Figure 2), offering seamless compatibility with a wide range of popular devices. The openEuler community also supports a number of other tools and systems for embedded development, as well as tools for interfacing with smartphones and other smart devices. The community even maintains a version of openEuler that runs on Raspberry Pi systems.

Conclusion

openEuler might seem like the "new kid" in some Western circles, but this expansive enterprise system is supported by thousands of developers. The openEuler environment is home to a wide range of tools for AI, edge, HPC, embedded, and conventional server room operations. Recent licensing changes at Red Hat offer an opening for other enterprise systems, and openEuler is gaining more attention as a leading alternative. Expect to hear more from this versatile Linux as new technologies such as AI and edge computing enter the foreground. ■■■

Info

[1] openEuler: https://www.openeuler.org/en/
[2] EulerOS: https://developer.huaweicloud.com/intl/en-us/euleros/index.html
[3] OpenAtom: https://www.openatom.org/
[4] Leonhard Euler: https://en.wikipedia.org/wiki/Leonhard_Euler
[5] OpenHarmony: https://en.wikipedia.org/wiki/OpenHarmony

Joe Casad

Joe Casad is the editor in chief of *Linux Magazine*.

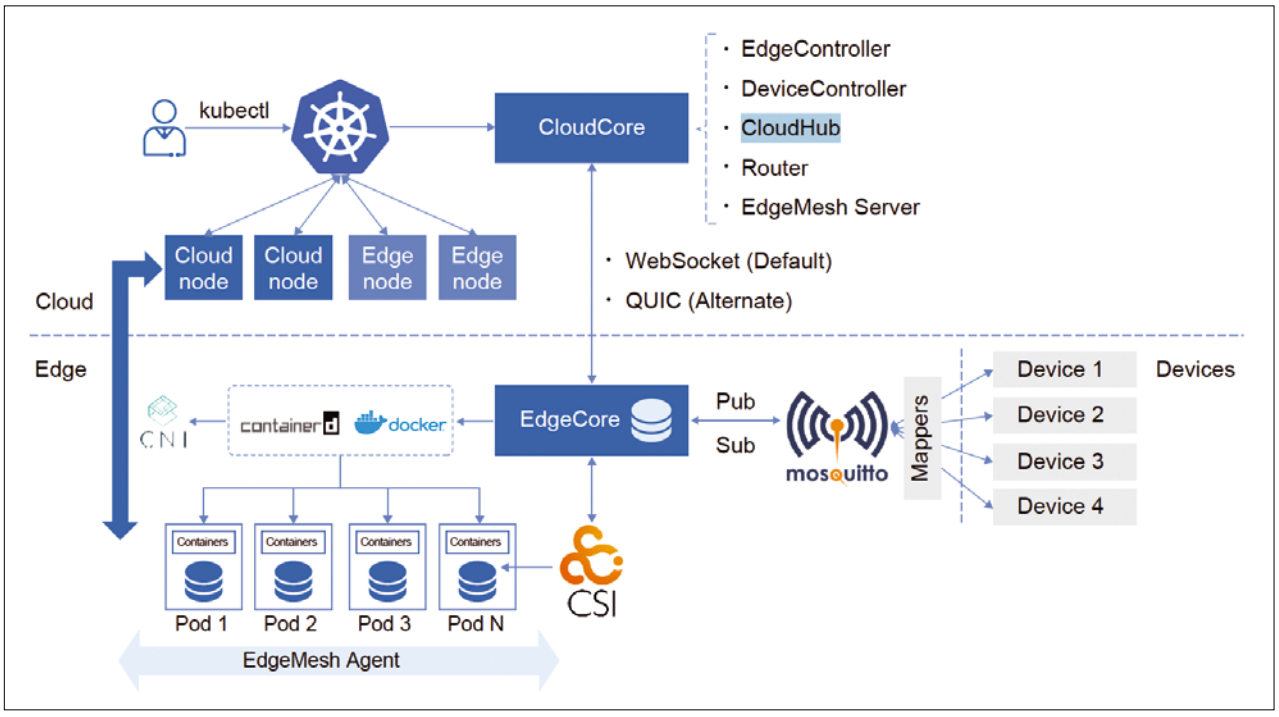


Figure 5: openEuler Edge supports unified management in edge computing scenarios.



Automating migration to openEuler

The Fast Migration Path

Easily migrate your old CentOS systems to the modern openEuler with the x2openEuler migration toolset. By Rubén Llorente

For years, CentOS was the distribution of choice for system administrators that needed to run a Red Hat Enterprise Linux (RHEL) compatible system, but could not afford a full support contract with Red Hat. CentOS, a fully compatible clone built from RHEL sources, functioned as a drop-in replacement for RHEL. CentOS promised an enterprise-grade Linux distribution that was available for anyone regardless of their budget.

Red Hat eventually became a sponsor for CentOS in 2014 **[1]**, in a movement the community found controversial at the time. The sponsorship agreement was, according to LWN.net **[2]**, a de facto acquisition of CentOS by Red Hat, since Red Hat took ownership of the CentOS trademarks and the CentOS governing board consisted of a majority of Red Hat employees.

In 2020, Red Hat unilaterally discontinued CentOS development, setting end-of-life (EOL) dates for the active versions much earlier than the community expected. According to Red Hat, CentOS users had no need to worry, because they could use the new CentOS Stream instead. Sadly, CentOS Stream is CentOS in name only. While the original CentOS was a RHEL drop-in replacement, CentOS Stream is an upstream development version of Red Hat, which made serious users wary of migrating.

The last CentOS supported release, CentOS 7, reached EOL on June 30, 2024. If you are still using CentOS today, you are running an unsupported operating system. The good news is that many alternatives have arisen in response to CentOS' demise, so system administrators have plenty of migration options.

One such option, openEuler from the OpenAtom foundation, offers an Enterprise Linux-like system for servers and cloud environments. It boasts a number of unique utilities, such as iSulad, a lightweight container runtime daemon; A-Tune, an AI-powered performance tuning tool for servers; and the x2openEuler migration tool. In this article, I explain how to use x2openEuler to migrate a small fleet of CentOS 7 servers to openEuler.

Your Migration Plan

Traditionally, upgrading systems from the RHEL family involved recreating your services on the new operating system, redeploying them from scratch, and discarding the old installations. For complex services with tight dependencies that cannot easily be carried over from your old CentOS to the new openEuler, this route might still be required. However, the vast majority of applications can be upgraded through different means.

The ideal, modern way of carrying out a migration is via in-place upgrades. When performing an in-place upgrade, you access each server that you want to migrate, replace the operating system that is running under your applications, and keep your production services

Photo by Chris Briggs on Unsplash

pretty much untouched. The advantage of this approach is that it is fast, requires less effort, and lets you reuse your old service configuration. As a rule of thumb, if the platform your production services run on is available in openEuler's repositories, an in-place migration is the way to go. For example, if your application runs on a modern OpenJDK, Perl, or Python environment, chances are an in-place migration will work free of fuss. To see if the platform your service runs on is available, check Repology **[3]** to verify that openEuler has the packages you need.

I strongly advise backing up your servers before you attempt to upgrade them. While the migration suite has a rollback functionality, you should not count on your migration system to safeguard your data. Ensure you have a sound backup before you proceed.

In-place migrations require the servers to reboot at certain points, therefore you should be prepared for service interruption during the process.

Your Migration Suite

The x2openEuler upgrade management tool needs to be installed on a machine in your network (preferably a server), be configured by an administrator, and then be instructed to perform a number of operations on your server fleet. x2openEuler operates by opening SSH sessions from the control machine to connect to your target servers. This design makes x2openEuler a fine option for migrating small- and medium-sized server fleets, because the administrator can install x2openEuler on a single control node, provide it with the all the target servers' addresses and SSH credentials, and have it migrate them all automatically. Servers can also be migrated in groups if you don't want to switch all your infrastructure in one go – you can also choose to migrate in small increments instead. To install x2openEuler, grab it from the openEuler download page **[4]**. x2openEuler is available as an RPM package for the x86_64 and AArch64 architectures. If your goal is to migrate an old server fleet composed of CentOS 7 systems,

Listing 1: Setting the MariaDB Root Password

```
SET PASSWORD FOR 'root'@'localhost' =
  PASSWORD('your_password_here');
FLUSH PRIVILEGES;
QUIT;
```

```
[root@linux service]# bash service_init.sh
Start the MariaDB custom configuration for the use of the x2openEuler tool.
Has MariaDB been installed on the server and can be used properly? [Y/n] (default: n) Y
Enter the password of the root user of the MariaDB again:
Set the password of the x2openEuler user for MariaDB:
If the selected database already exists, it will be overwritten.
Use default x2openEulerDb database? [Y/n] (default: Y) Y
MariaDB is configured successfully.
If authentication is enabled,
the SSH connection fails after the fingerprint of the machine changes.
Please confirm whether public key authentication is not required for SSH connection(y/n default: n):
```

Figure 1: The x2openEuler initialization process.



Figure 2: The migration suite uses a self-signed certificate. As a result, the web browser will throw a warning the first time you access it.

it's safe to assume you will be installing x2openEuler on a CentOS 7 control node. I recommend running x2openEuler from an instance devoted solely to this task that you can discard after the migration is complete. In order to install x2openEuler, just upload the RPM package to the machine and install it with `yum`:

```
# yum install ./x2openEuler-core*.rpm
```

The installation process pulls a surprising number of dependencies with it, including the MariaDB database (as well as other databases). You should set a root password for your database to make the initialization of x2openEuler smoother. Just start the database service with

```
# systemctl start mariadb
```

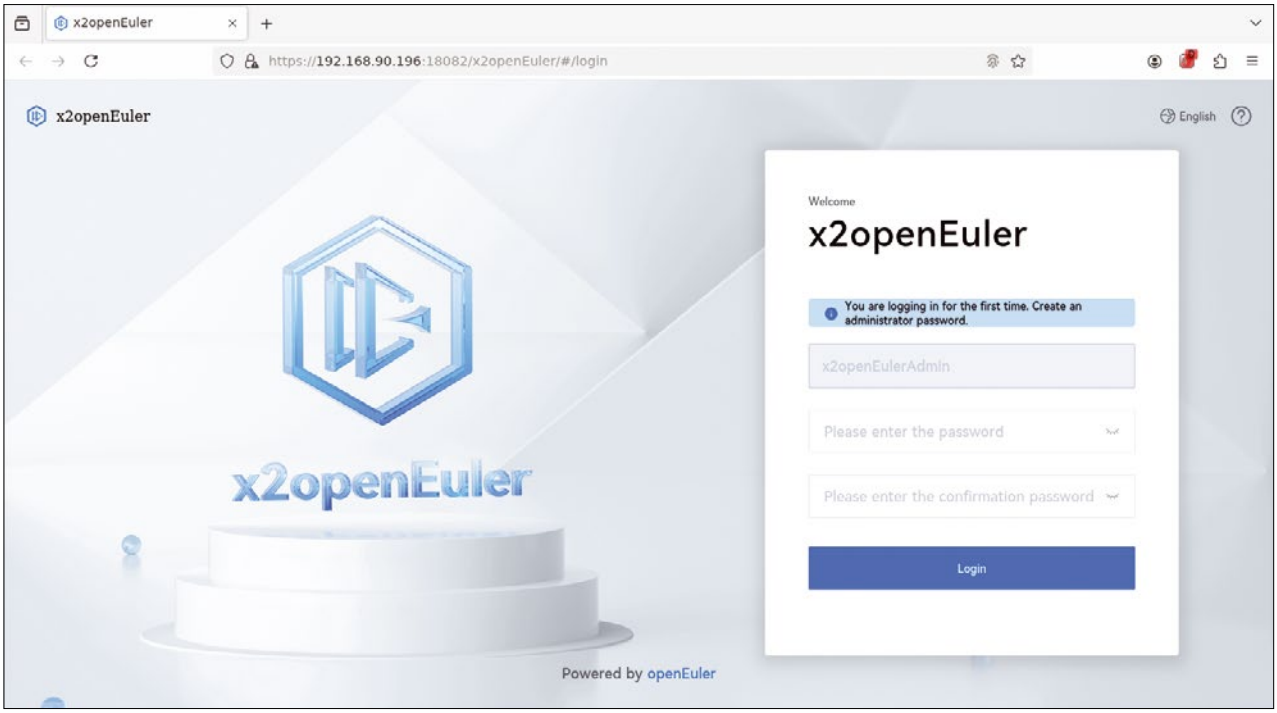


Figure 3: You must create an administrator user in order to log in for the first time. The web interface supports both the Chinese and English languages, which you can change from a menu at the top right of the page.

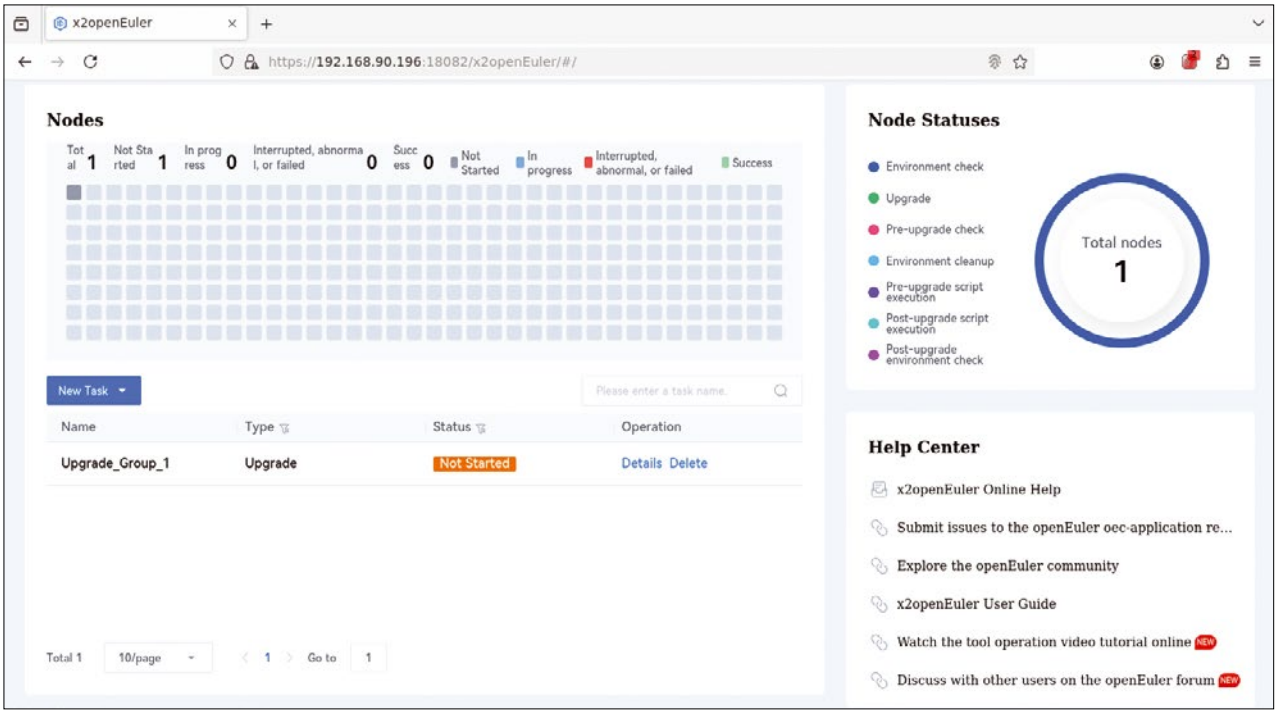


Figure 4: The task manager is the core of x2openEuler. Use the *New Task* menu to create a new *Upgrade* task.

Using the `mysql` command, launch a database prompt and set your desired password for the database root user. The root user here is the root user within the database context only – these actions don't modify your system root access credentials! See Listing 1 for the command sequence to set your database root password. Your password must feature a combination of lowercase, uppercase, and special characters; otherwise, it will be rejected.

Next, you need to initialize x2openEuler itself. Within your terminal session, move to `/usr/local/x2openEuler/portal/service/` and get started using the script `service_init.sh`. You will be asked a number of questions regarding your setup (Figure 1); feel free to use common-sense answers.

The initialization script will perform a number of tasks in order to enable and launch the service, including the generation of a self-signed TLS certificate for x2openEuler's web interface. At the end of the process, the server will be running, and you will be able to access the web interface by pointing a web browser to *https://$yourserver:18082/x2openEuler/#/login* (where *$yourserver* is the machine's IP address). Because x2openEuler uses a self-signed certificate, you should expect your web browser to throw a warning (Figure 2) before it lets you access the web interface (Figure 3). Also worth mentioning, the x2openEuler service will be vulnerable until an administrator logs in for the first time. Therefore, I recommend keeping the machine that runs the service behind a firewall with no connection to outside of your organization.

The Migration Targets

x2openEuler is a surprisingly complete tool. If you are operating within a serious environment, your first stop will probably be *Web Server Certificates* (accessible from the gear on the top-right corner of the administration dashboard). From here, you can replace the default self-signed certificate with one that is widely recognized by your organization's web browsers. Additional users can be added via the *User Management* section, in case you need to grant access to your coworkers.

Task management is x2openEuler's main feature (Figure 4). When you open the task management interface, you'll be presented with many things to do. If the only thing you need for the moment is to upgrade your server fleet, you can cut to the chase and hit *New Task* to create a new *Upgrade* task. Within the task creation section, you may add as many target servers as you like. You may either add them in a batch by uploading a Microsoft Excel file (a template file is provided) or you can add servers one by one. x2openEuler will manage the target servers over an SSH connection, and both password- and keyfile-based authentication are supported (Figure 5).

From the same menu where the target servers are defined, you can also configure a list of packages you want to be installed on the migrated system later, run custom scripts at different phases of the process, and perform a number of advanced actions.

After you have added all the servers you want to migrate, hitting *Ok* at the bottom of the task page will start the migration task. As a first step, it will begin performing preliminary environment tests. You will be taken to a status page that will show a global report of the ongoing testing. From here, you can check individual

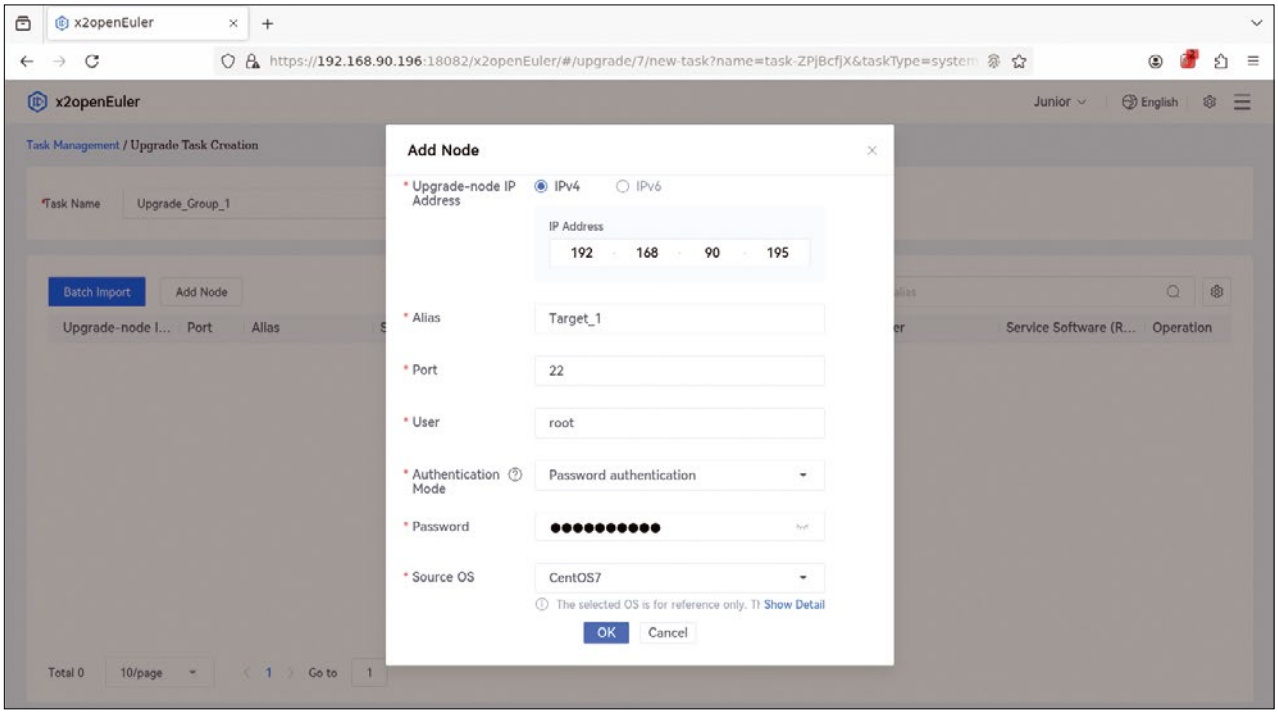


Figure 5: Servers can be added to your upgrade task one by one. If this method is too slow, you can also import a list via a Microsoft Excel file.

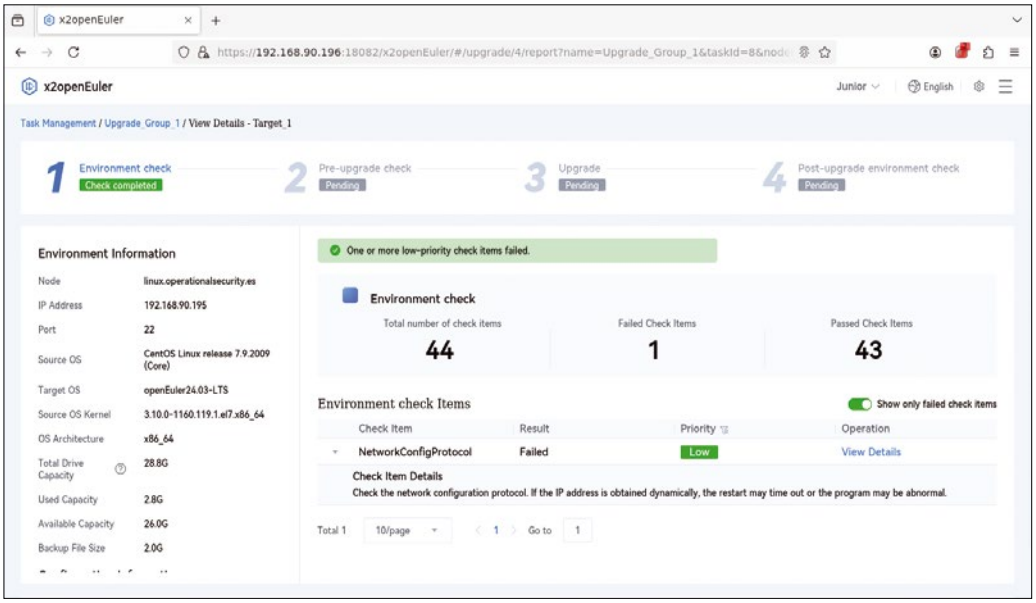


Figure 6: The detailed status for one of the servers: This migration task shows a minor failure because x2openEuler thinks this server might be using a dynamic IP address.

servers for an in-depth detailed view of their status (Figure 6).

The whole migration cycle for each server consists of four phases: an *Environment check*, a *Pre-upgrade check*, the actual *Upgrade*, and the *Post-upgrade environment check*. You can trigger these for each server from the *Task Management* tab or by calling a batch action that includes all the targeted servers (Figure 7).

Once the *Upgrade* phase is complete, you are given a final chance to rollback each server to the previous state. x2openEuler has the capability to rollback to the pre-migration stage, and you can go back using the *Rollback* button on the bottom right of the screen. If you are satisfied with the state of the server after the migration, you can move on to running the post-upgrade checks.

Conclusion

Migrating away from an obsolete CentOS installation doesn't have to be hard – it can actually be a quick process. You should expect a small server to be fully migrated in little more than an hour. If you're looking for a Enterprise Linux alternative that isn't focused on 1:1 compatibility with IBM's offering and that is actively developing their own cloud and server solutions, openEuler is worth considering, if only because their migration and upgrade tools are so powerful and convenient.

This is not to say that x2openEuler cannot be improved. Something that surprised me a bit, for example, was that users can't share migration tasks. x2openEuler users can only create tasks for themselves and can only execute their own tasks. While not a deal breaker, I think openEuler is missing the opportunity to make x2openEuler a true collaborative suite with granular permissions – useful, for example, if you want a junior system administrator to create the migration tasks and run the preliminary checks, but want a senior sys admin to trigger the actual upgrade. ■

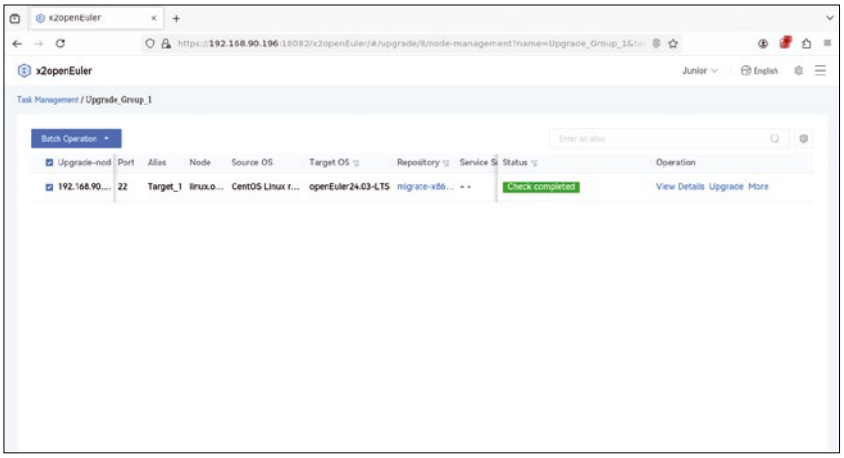


Figure 7: A task with a single target server: You can trigger the *Upgrade* phase by clicking on *Upgrade*. If a task contains more than one server, you can select *Batch Operation* to upgrade them all at once.

Info

[1] "CentOS Project joins forces with Red Hat": [https://lists.centos.org/pipermail/centos-announce/2014-January/020100.html]
[2] "CentOS and Red Hat team up": [https://lwn.net/Articles/579551/]
[3] Repology: [https://repology.org]
[4] x2openEuler download: [https://www.openeuler.org/en/migration/download/]

The Author



Rubén Llorente is a mechanical engineer whose job is to ensure that the security measures of the IT infrastructure of a small clinic are both legally compliant and safe. He is also an OpenBSD enthusiast.



openEuler: An OS for AI

Intelligent Stack

openEuler prepares for the AI future with a complete software stack tailored for artificial intelligence. By Joe Casad

It is no secret that tech companies, investors, governments, and futurists are all focused on artificial intelligence (AI). According to experts, the AI technologies emerging now could change the ways we live and work. And technology companies are leading the way, because they know if they don't embrace this intelligent future, they could be left behind. Despite all the attention to AI, most Linux systems look much the same as they always have – enterprise distributions tailored for a general, enterprise environment. A few offer variants for edge and cloud, but when it comes to AI, the common approach is to manually add extensions on the fly. Linux is the powerhouse behind the AI revolution, yet many of the adaptations for accommodating AI look like an afterthought. This approach can limit performance and reduce reliability, in addition to taking the time from experts who would prefer to be working on AI applications rather than tweaking the OS environment. Developers of the openEuler Linux system have taken a different approach: They envision a complete software stack tailored for the needs of AI – a whole system that is AI-ready out of the box.

The developers asked themselves:

- How do we ensure hardware compatibility amid diverse AI scenarios?
- How do we overcome installation difficulties in complex AI software stacks?
- How can the operating system support efficient and reliable AI?

These questions led the openEuler developers to envision a complete software stack that could serve as a universal foundation for AI installations. The Intelligence BooM stack lets you roll out a complete, AI-ready environment in as little as 15 minutes – tuned for performance and end-to-end security, and designed to work with an array of advanced AI-based hardware.

Intelligence BooM

Intelligence BooM is an open source reference implementation of an AI-software stack based on openEuler. The stack integrates seamlessly with mainstream open source solutions at every layer, offering both community implementations and openEuler-optimized alternatives.

Intelligence BooM (Figure 1) draws together a collection of advanced tools and services to provide a complete solution for AI. Joining these tools into a complete stack can save hours or

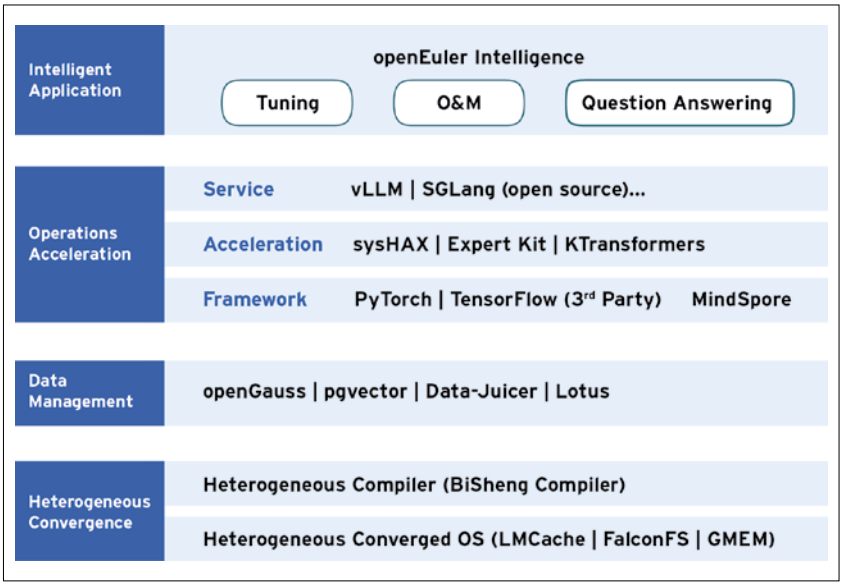


Figure 1: openEuler's Intelligence BooM stack – a complete environment built for AI.

Lead Image © antoinepham2903, 123RF.com

even days of configuration time – and potentially weeks of troubleshooting. The result is a whole new approach to the problem of AI in Linux: a complete stack that can serve as a native operating environment for AI.

As you will learn in this article, the Intelligence BooM stack doesn't just support efficient AI workloads; it also offers a comprehensive suite of development tools for building and implementing AI applications. openEuler's AI development environment is a core strength, supporting seamless and efficient AI development with language tools, training frameworks, retrieval-augmented generation (RAG) support, and other components for efficient AI coding.

Also included in the stack are AI-based tools for data management, as well as development frameworks, acceleration tools, and services for supporting AI applications. At the top of the stack is the innovative openEuler Intelligence system **[1]**, a toolkit for developing and implementing AI applications.

The openEuler Intelligence BooM stack is preconfigured and ready for developers who would rather get started on custom AI solutions without building the stack from scratch.

openEuler Intelligence

The top layer of the Intelligence BooM stack is populated with a set of AI tools collectively known as openEuler Intelligence. According to the developers, the goal of openEuler Intelligence is to let the user go "from ideas to AI applications in minutes with a user-friendly visual interface and low-code capabilities."

openEuler Intelligence (Figure 2) is a key feature of the Intelligence BooM environment, and it is the ingredient that makes openEuler an ideal platform for AI development. The openEuler Intelligence toolkit is preloaded with ready-to-use agents for OS development, deployment, performance tuning, and streamlining developer workflows.

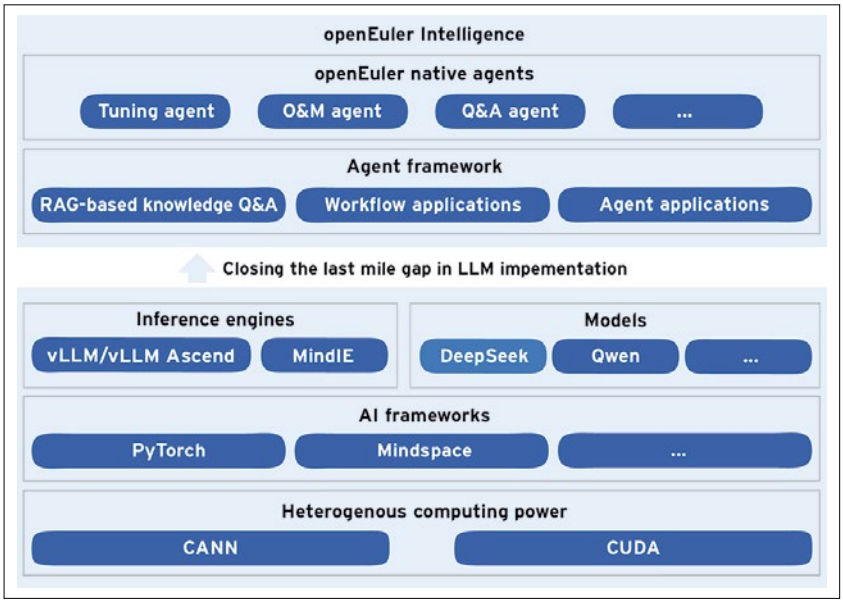


Figure 2: openEuler Intelligence and supporting components.

You can also develop your own agents with tools included in the openEuler Intelligence environment. The agent development framework encapsulates AI development capabilities, enabling developers to build custom agent applications with minimal overhead. The framework comes with visual drag-and-drop and multipath retrieval. Developers can access the underlying layers of the stack to efficiently build AI applications. You can upload local APIs through OpenAPI and then visually orchestrate workflows to create your own applications.

What Is GMEM?

New technologies like graphics processing units (GPUs), tensor processing units (TPUs), and field-programmable gate arrays (FPGAs) have complicated the task of memory allocation in today's heterogeneous systems. The good news is that accelerators have their own built-in memory. The challenge is that this memory is separated from the CPU, and conventional techniques for moving data between memory locations can introduce bottlenecks. The Heterogeneous Memory Management (HMM) feature built into Linux can hamper performance and relies on manual tuning. GMEM "combines the two independent address spaces into a unified virtual memory space" (Figure 3) **[2]**. The GMEM APIs allow for heterogeneous memory programming without memory migrations.

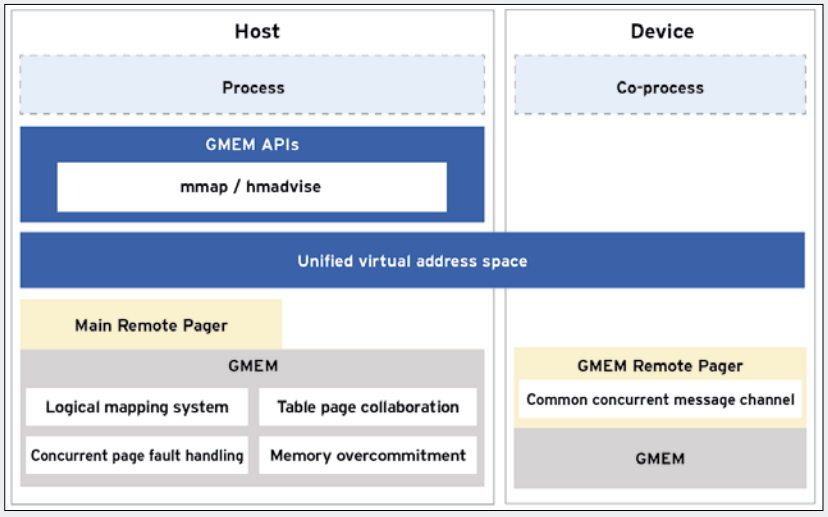


Figure 3: GMEM merges host and device memory into a unified virtual address space.

openEuler Intelligence also supports RAG, allowing you to parse document formats such as DOCX, PDF, and PPTX to build intelligent Q&A assistants around existing documentation.

openEuler at the Base

The openEuler operating system is tailored for heterogeneous convergence. You'll find support for a wide range of hardware environments, from x86 to ARM and RISC-V. openEuler also supports the new generation of AI-optimized chips. The Intelligence BooM stack enhances performance for AI workloads with:

- Coordinated scheduling of heterogeneous workloads
- Unified memory management for heterogeneous architectures
- Virtualization and pooling for heterogeneous computing power

A number of other components are built into the system for seamless AI integration. For instance, LMCache provides a key-value cache layer for memory optimization of large language models (LLMs). FalconFS is a high-performance distributed filesystem for AI workloads. Generalized Memory Management (GMEM) is a centralized management system for connecting heterogeneous memory in distributed environments (see the box entitled "What Is GMEM?").

Data Management

Like any enterprise-grade Linux system, openEuler supports a variety of tools for database management in cloud, edge, HPC, and data center environments. However, the openEuler developers are aware that AI is

Performance Matters

AI applications place significant demands on performance. The openEuler team has years of experience with tuning HPC systems, cloud environments, and server rooms, and they bring that expertise to the challenge of tuning systems for AI. openEuler includes a number of tools for improving and optimizing performance. For instance, A-Tune (Figure 4) is an AI-powered tuning engine that uses configuration data and performance metrics to automatically tune the system and then continues to iterate using a parameter search algorithm to achieve optimum performance.
Another tool, etmem (Figure 5), organizes memory into tiers, automatically routing hot data to high-speed DRAM and cold data to low-speed media. Hybrid Storage Acceleration Kit (HSAK) speeds up I/O for NVMe devices. Also included are custom versions of the GCC and LLVM compilers optimized for the openEuler environment.

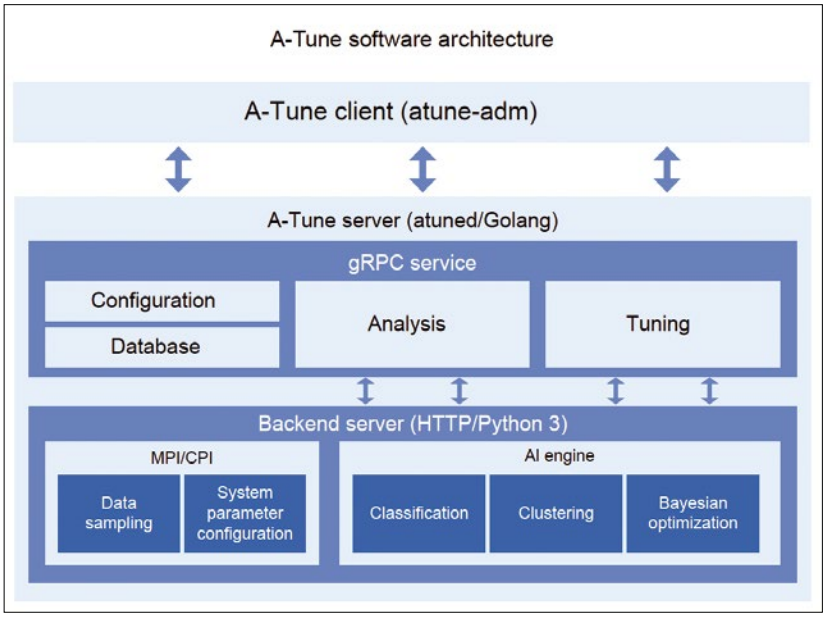


Figure 4: A-Tune software architecture.

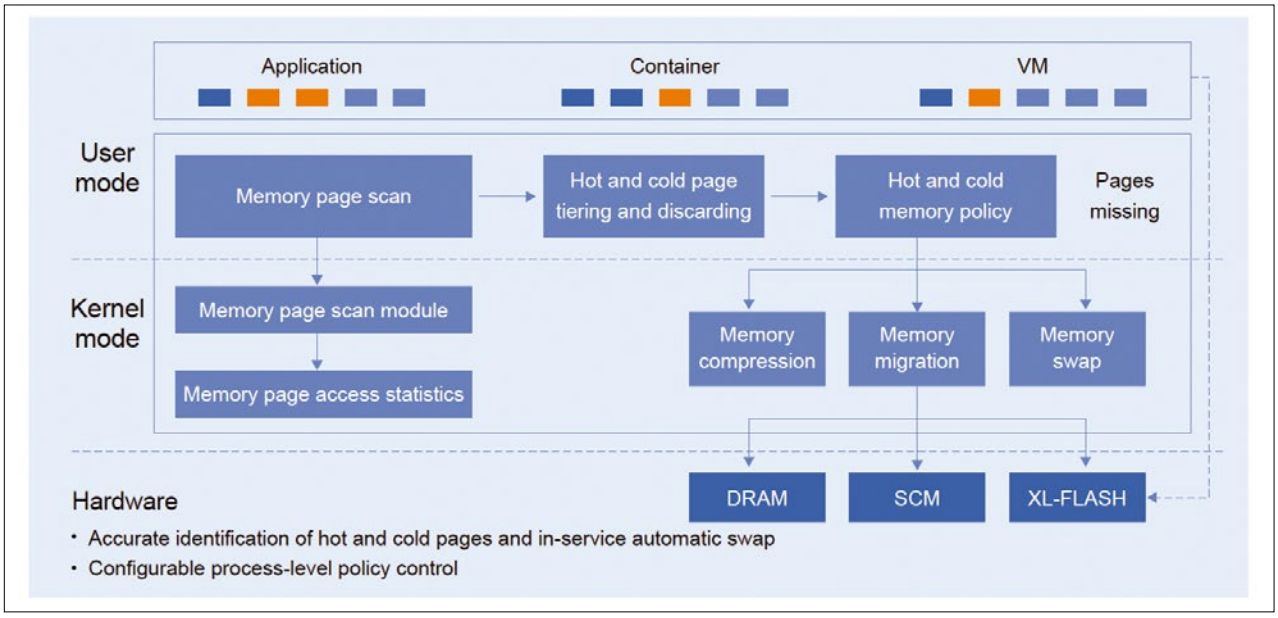


Figure 5: The etmem tool optimizes memory by saving frequently accessed data to faster media.

a special use with special needs for managing vast quantities of data. AI runs on data, and openEuler's BooM stack integrates a number of tools for meeting the challenge of managing data in AI environments. openEuler can work with openGauss, an open source database management system designed for high performance and high reliability. Also closely integrated with openEuler is pgvector, a PostgreSQL extension that supports the vector similarity search functions essential to many AI systems.
openEuler integrates another important component often used with LLMs and other AI applications: Data-Juicer, which is designed to efficiently generate data recipes used for training AI models.

Containerized Frameworks

openEuler supports a number of popular AI development frameworks in containerized form for easy deployment and integration. In a few clicks, you can prepare a complete development environment for:

- PyTorch
- TensorFlow
- MindSpore

Other images support the CUDA and CANN development environments. This modular, containerized design lets you easily integrate the development tools you need without the clutter of the tools you don't need – and if your needs change, you can easily adapt.

Acceleration

An accelerated system means faster learning and a shorter response time for AI queries. openEuler integrates several tools for accelerated operation in an AI context. The sysHAX LLM acceleration runtime enhances inference performance. sysHAX improves CPU throughput via NUMA-aware scheduling, parallelized matrix operations, and Scalable Vector Extension (SVE)-optimized inference operators. openEuler also provides the Expert Kit (EK) high-performance framework for scalable Mixture of Experts (MoE) LLM inference. (MoE is a machine learning technique where multiple expert neural networks, acting as learners, divide a problem space for more efficient learning. MoE layers are often used in Transformer models.)
Another powerful acceleration technology built into openEuler is KTransformers (pronounced "Quick Transformers"), an innovative tool developed at Tsinghua University. The Python-centric KTransformers framework offers kernel optimizations and parallel processing strategies. According to the KTransformers developers, "By implementing and injecting an optimized module with a single line of code, users gain access to a Transformers-compatible interface, RESTful APIs compliant with OpenAI and Ollama, and even a simplified ChatGPT-like web UI" **[3]**.

Conclusion

The Intelligence BooM reference stack is a departure point for visualizing the openEuler AI environment. The powerful tools of the openEuler application layer rest on a full stack of open source tools that is tailored and tuned for AI workloads. In addition to the tools described in this article, you will find drivers and software development kits (SDKs) to support out-of-the-box operations for a wide range of AI scenarios. See the box entitled "Performance Matters" for more on the performance tuning tools included with openEuler. The "openEuler Intelligence: A Case Study" box offers an example showing how openEuler Intelligence can bring the power of AI to solving real-world problems.
openEuler runs on a variety of hardware and is supported by a large community of users, including many AI specialists employed in both academia and industry. If you're interested in learning more about AI in openEuler, contact the openEuler support team at contact@openeuler.io. ■

Info

[1] openEuler Intelligence: [https://www.openeuler.org/en/projects/intelligence/]
[2] GMEM: [https://www.openeuler.org/en/blog/G11N-SIG/20240116-memory-management/20240116-memory-management.html]
[3] KTransformers: [https://kvcache-ai.github.io/ktransformers/]

Joe Casad
Joe Casad is the editor in chief of *Linux Magazine*.

openEuler Intelligence: A Case Study

The Intelligence BooM platform is designed to support custom solutions for practical problems. Sangfor Technologies, for example, recently implemented a patch management solution that leveraged the ready-to-use AI components built into openEuler Intelligence.
All enterprise Linux developers face a common problem. The Linux mainline kernel releases an update every 9-10 weeks, with each update containing 2,000-3,000 patches. Most enterprise organizations, however, don't want the disruption of a full kernel update every 10 weeks. Developers must therefore manually select the relevant patches and integrate them into production branches. This manual update process leads to inefficiencies and bottlenecks. One update can require up to 80 hours of a senior engineer's time, often leading to error rates exceeding 10 percent.
Sangfor wrote their patch management solution around the advanced components built into openEuler Intelligence. The solution leverages Qwen and other LLMs. Standardized templates generate reports, and a syntax analyzer translates code and patch submission into natural language suited for LLM comprehension.
The AI-based solution was deployed as an eight-pipeline workflow that reduced the analysis time for 100 patches from 200 minutes to 30 minutes. The direct patch application success rate rose from 90 percent to 98 percent, and the patch analysis speed improved 50 times. Overall, the solution reduced the workload for an update from one person-week to one to two hours.
Reducing the time for an update is particularly important in today's age of heightened security awareness. Sangfor's AI-based patch management solution reduces the response time for critical vulnerabilities from 14 days to less than 72 hours. By keeping an AI-based development environment close at hand, openEuler enabled a smarter solution to an age-old problem.

openEuler community leaders and experts tell about the journey

Up Close

We sat down with three experts from the openEuler community for an inside look at the project, its goals, and the path ahead.



Hu Xinwei joined Huawei in 2011 and is now a senior expert in ICT operating systems. Since 2020, he has served as the Chairperson of the openEuler Technical Committee and is responsible for the overall technical direction management of the openEuler Community.

Q Where did openEuler originate? Trace its history back to the very beginning. When and why did it start? What issues was it intended to solve?

A openEuler was initially developed by Huawei for internal use as a server operating system (OS), with the primary goal of addressing several internal needs that external commercial OSes could not meet.
One requirement was the widespread use of a diverse range of chips, encompassing various architectures and types. Huawei needs an OS that can support these different chip varieties. With the rise of heterogeneous computing and AI operations, there has been a Cambrian explosion in chip types, which has also significantly accelerated the iteration of traditional CPUs. However, the speed at which traditional commercial OSes support new chips is relatively lagging and slow. Therefore, there was an urgent need to build an OS platform capable of rapidly supporting these new processors.
Another requirement stemmed from Huawei's diverse business types, which include scenarios like servers and cloud computing, as well as edge computing and embedded devices. Typically, commercial OS vendors are categorized into server OS vendors and embedded OS vendors. Their software baselines often differ, making compatibility out of the question. However, with the convergence of device, edge, and cloud businesses, Huawei hoped to have a single OS that can support all these different scenarios. This would enable "develop once, adapt to many scenarios" for application development, significantly reducing development costs and accelerating development speed for its business units.
Another factor was that for large enterprises, there are many innovations closely tied to their business operations, many of which require modifications to the OS. AI has especially become an industry hotspot, leading to even more innovations in the OS. However, these changes are often difficult for commercial OS vendors to accept; thus, there is a pressing need for a model that supports these more aggressive innovations.
When Huawei and other companies discussed these issues, they found that these challenges were also faced by industry partners. Therefore, during the period of 2018-2019, Huawei worked with multiple industry partners for nearly a year and decided to contribute its internally used OS. Leveraging this OS as a blueprint, several companies jointly established the openEuler community. The community focuses on core technologies, such as all-scenario compatibility, multi-architecture adaptability, and AI. It aims to build a nonprofit, open, and innovative OS community that supports radical innovation and aligns with future technological trends.
The openEuler community was officially established in 2019. To ensure openEuler's long-term neutrality and sustainable development, Huawei and its community partners donated the project to the OpenAtom Foundation in November 2021. This donation included the full codebase, trademarks, community infrastructure, and related assets.
Since then, openEuler has been operated and governed by the community under the OpenAtom Foundation, attracting contributions from a growing number of enterprises, universities, and individuals. The project has evolved into a community-led, open-governance OS ecosystem, focused on building a vibrant and inclusive platform for digital infrastructure innovation.
To date, the community has been running for six years and has largely addressed the challenges faced by our companies. The community can proudly state that it now supports the most diverse range of chips and offers common baseline versions for edge, endpoint, server, and cloud. Additionally, the community is highly active in the field of AI with many innovative contributions.

Q Has the collaboration between openEuler and Huawei developers continued since the donation?

A Huawei has always been a staunch supporter and practitioner of open source. Before openEuler, Huawei participated in many international open source communities and made significant contributions. As I mentioned earlier, Huawei and its partners established the openEuler community with the principle that it should be open, inclusive, and nonprofit. After about two years of operation, the community's member organizations all agreed that, to further enhance open governance and attract broader industry participation, the best approach was to donate it to a foundation. This would ensure that a noncommercial entity manages the community, which is the ideal home for openEuler. It has proven that donating to the foundation was a crucial milestone and the right decision for the openEuler community. Since 2021, under the management of the foundation, the community has experienced rapid growth. Currently, more than 2,000 companies and institutions have joined the community, and over 20,000 contributors have participated, making it a large-scale community.
The donation had several goals. The developers wanted to create an open, transparent, and collaborative environment for OS development. They wanted to encourage more diverse contributors to participate, including developers from other vendors, academic institutions, and the open source community at large. They also wanted to promote shared innovation and ecosystem growth in digital infrastructure software.
Huawei remains one of many contributing members to openEuler, alongside other enterprises, universities, and individuals. Huawei remains an important contributor to the openEuler community. However, with more participants joining, Huawei's contribution ratio has been declining. Currently, over 60 percent of contributions come from companies and developers outside of Huawei. In the future, Huawei will continue to contribute to the openEuler community and hopes to work with more enterprises to further improve the openEuler community. Importantly, Huawei does not release proprietary commercial versions of openEuler. All contributions follow community-first, open source principles, ensuring that the technology remains accessible to all participants.

Q What is your vision for openEuler? What role do you envision it playing? Who is your target audience?

A We expect future OSes to evolve in two key dimensions:

- Intelligence-native design – Transitioning from POSIX-based software ecosystems to new paradigms integrating AI agents with traditional applications
- SuperPoD – Supporting emerging hardware architectures, especially SuperPoD (hyper-scale nodes) with tailored system capabilities

These transformations will reshape software ecosystems and developer collaboration.
I hope the openEuler community can become the most open and innovative OS community. I also hope that companies can build different business models based on the openEuler community. In short, I hope the openEuler community can be developer-friendly and also business-friendly.
We envision openEuler as the foundational OS for future digital and intelligent infrastructure, becoming the best development and runtime environment in the AI era. Beyond technology, we also want the openEuler community to be a vibrant hub for developers – a place where people who love open source, system innovation, and technical exploration can communicate, collaborate, and grow together.
Our target audience includes open source OS enthusiasts and developers, industry users building vertical applications on top of the OS base, independent OS vendors and solution vendors who want to develop their own competitive products, and global enterprises and partners looking for an open, reliable, and customizable platform. openEuler is not just about building an OS – it's about building a global, open source community for OS innovation.
As the Co-Chair of the openEuler Technical Committee, my primary responsibility is to organize and manage community processes. I also aim to attract more partners to join our community. Throughout this process, I am most eager to hear candid opinions and suggestions from all parties regarding the development of the openEuler community, which will help the management team to better guide the project.

Q What specialized tools or applications are included in the openEuler environment? What unique features does openEuler offer compared to its competitors?

A While the industry shows tremendous enthusiasm for large language models (LLMs) and AI – with many end-device OSes offering native AI services – no server OS has previously done this. For server OSes, providing native AI services is significantly more challenging: The model must be truly useful (not just novel) and cost-effective to operate.
openEuler is building the industry's first server OS with native AI capabilities. Its built-in model enables application tuning tools to outperform manual expert optimization – without requiring specialized hardware. This differentiates openEuler from other projects. Our sustained investment in practical AI applications has yielded tangible results, charting a clear path for intelligent OS development.
The community has been developing for six years and has accumulated a wealth of tools and innovative features. Currently, more than 550 new projects have incubated within the community. You can find the innovation white paper on the official community website, which lists over 50

projects covering areas such as networking, I/O, cloud-native, embedded systems, security, AI, and more. These projects come from contributions by various enterprises and organizations. I recommend checking them out; you're sure to find something that interests you. I do not believe that openEuler has any competitors, nor do we wish to emphasize our so-called uniqueness. We welcome everyone to use and integrate projects incubated by the openEuler community, and the openEuler community will continue to monitor industry developments and promptly incorporate new innovative technologies from the industry. We aim to create a platform for innovation and collaboration with openEuler, and this goal has remained unchanged.

Q What does openEuler excel at? What are you most proud of?

A I truly believe openEuler's greatest strength lies in its open collaboration. Building an OS community, especially one for enterprise applications, always comes with its share of challenges and even conflicts. For instance, chip manufacturers in our community often push for quicker version releases, as it makes supporting new chips much easier for them. On the flip side, OS vendors and end users usually prefer slower releases to avoid the higher costs that come with switching versions.
When we hit these kinds of roadblocks, the community brings all these companies together. We talk it out, negotiate, and work to find solutions that make sense for everyone. That's why I think what really sets openEuler apart is that its development truly comes down to discussions and coordination among all the participating enterprises and organizations.
What makes me most proud is how we've consistently kept our community diverse over the past few years. We have hardware companies, chip manufacturers, OS vendors, service providers, ISVs, end users, research institutions, and even universities all contributing. This mix means we get to hear from every angle, ensuring the community's development addresses a wide range of needs. Ultimately, this commitment to "listening" is the core reason behind openEuler's rapid growth.

Q What are you currently working on? Are there any new tools or applications in development that you are particularly excited about?

A I'm currently coordinating multiple open source projects to deliver an out-of-the-box AI full-stack solution called Intelligence BooM. This open source reference implementation enables users to experiment with digital and intelligent transformation using their private data – making AI truly applicable beyond just being a novelty.
Additionally, openEuler is investing in defining and developing future OSes for SuperPoDs. Our goal is to abstract and unify diverse high-speed interconnected technologies at the OS level, creating a new form of tightly coupled distributed OS.
openEuler will continue its exploration in intelligent OS development. In the era of Agentic AI, all OS tools and interfaces will become agent-enabled – an area where we are actively making progress.
I currently serve as the chairperson of the openEuler Committee, and my primary focus is on contemplating the openEuler development model for the next phase. From a technical perspective, AI is undoubtedly the most critical direction for the future development of openEuler. Our openEuler Intelligence platform provides businesses and developers with an out-of-the-box AI open source full-stack system. The community is also working on using AI agents to comprehensively transform openEuler. Our goal is to fundamentally change how developers and users interact with and develop OSes. Imagine being able to control complex enterprise-level OSes through natural language conversations. We hope to achieve this transformation in our next version next year.
Additionally, I am considering how to further increase the openness of the community. Currently, openEuler has become a popular OS community and platform in China, attracting many international companies. I'm focused on how to further reduce the barriers for global companies and communities to collaborate with other communities worldwide.



Hou Jian, Vice President of Server R&D Center, KylinSoft. As a Technical Committee Member of the openEuler Community and Maintainer of the openEuler High Availability Special Interest Group (HA SIG), Hou Jian serves as KylinSoft's technical lead for open source community contributions to openEuler.

Q I understand that you have personally launched six SIGs within the openEuler community, including the HA SIG and the Kubernetes SIG. What were the key technical gaps in the ecosystem that prompted you to set up these groups? How do you coordinate cross-SIG collaboration, given the diversity of these focus areas?

A These six SIGs were established in the early days of the openEuler community. We created SIGs focused on our technologies – both to address community gaps and to demonstrate

KylinSoft's technical strengthens. Two examples illustrate our rationale: The Ultimate Kylin User Interface (UKUI) SIG was formed when we noticed openEuler's initial release lacked a GUI interface, coinciding with KylinSoft's desire to promote its self-developed UKUI display software. Today, the UKUI SIG has been adapted for all openEuler releases. Similarly, the HA SIG addressed the essential need for high availability in a server OS. KylinSoft planned to develop its Kylin-HA cluster software while incubating a community version through the SIG – attracting more developers to reduce costs while driving commercial sales to the HA solution.

Q KylinSoft participates in 70+ SIGs with 332 contributors. What governance best practices have you developed to maintain contribution quality at such a scale?

A KylinSoft's internal organizational structure operates similarly to openEuler's SIG model – both organize teams by technical domains to incubate technologies. For complex technical fields requiring cross-team collaboration, workgroups are established. For example, openEuler's AI Joint workgroup coordinates SIGs across the entire stack from kernel and compilers to application software.

Q For openEuler 24.03 LTS, KylinSoft led technical planning. What were the most critical kernel/distributed storage optimizations for production environments?

A One important feature is Multipath TCP (MPTCP). MPTCP provides users with a more powerful capability for multipath utilization. By breaking free from the constraints of traditional single-path TCP, MPTCP supports intelligent multiplexing of redundant channel resources, achieving a qualitative leap in data transfer rates – where the effective rate equals the sum of all available channel bandwidths. This significantly enhances both network efficiency and reliability. Currently, MPTCP has demonstrated outstanding performance in path selection, traffic distribution, and handover for WiFi and cellular networks. Moreover, MPTCP was designed with full backward compatibility to traditional TCP, ensuring seamless integration and smooth migration in existing network environments.

Another development is system compatibility detection and assurance. To address compatibility issues between legacy business applications and the NeoKylin V11 system during the upgrade from NeoKylin V10, this technology focuses on compatibility diagnosis and system-level assurance mechanisms for applications in OS upgrade scenarios.

The solution promotes standardized external interfaces for the OS, maintains ecosystem compatibility across version iterations, and facilitates seamless application-compatible upgrades.

Application scenario-based optimization technology provides comprehensive real-time system metric collection, automated bottleneck detection and analysis, expert-level tuning, and AI-powered optimization. Addressing the high labor/time costs and incomplete data acquisition of traditional multi-scenario optimization methods, it enables automated performance tuning for different application scenarios. It captures real-time system performance data and hotspots while visually presenting bottlenecks. With a single command, users can fully collect server metrics – including CPU utilization, NUMA status, network traffic, process-level graphs, and system resource allocation. The solution identifies bottlenecks while providing tuning recommendations and scenario-specific parameter suggestions.

Q Your teams contribute to 8+ domains, from compilers to AI. How do you prioritize resources between foundational (e.g., kernel) and emerging (e.g., AI) technologies?

A Our server OS runs mission-critical customer applications, making stability and compatibility of core components (kernel, toolchain, virtualization) paramount. This shapes our prioritized technology approach: foundational capabilities first, emerging technologies in phased adoption. Our basic strategy is 50 percent foundational + 30 percent enhancement innovation + 20 percent exploratory R&D.

Q With 24,000+ PRs contributed, how does KylinSoft align its enterprise product roadmap with community-driven development?

Before product planning, we first seek open source prototypes. If a suitable open source prototype exists, we develop our commercial product based on it while actively contributing to the prototype's technical evolution. If no prototype is available, we create one as an open source project for cutting-edge technology exploration and then develop the commercial product with a focus on usability, professionalism, and stability.

Q What training mechanisms does KylinSoft use to transition traditional Linux engineers to cloud-native contributors?

A KylinSoft employs Linux Foundation engineers with core competencies in kernel development, toolchain construction, and debugging (primarily using C). Their development follows traditional version/feature-based models with relatively slower iteration cycles. For cloud-native capabilities, required skills shift to cluster debugging and microservices (using Golang/Python), requiring agile methodologies with faster release cadences.

We support internal training programs to strengthen cloud-native fundamentals. We also find that community project immersion enables Linux engineers to rapidly apply existing skills in cloud-native contexts through open source contributions.

Q As China's core OS ecosystem expands internationally, what

lessons do you bring with you that could benefit global open source governance?

A The openEuler ecosystem is not enterprise-led with community assistance, but rather inversely designed: The community is the main body, with enterprises as contributors.
The community owns technical roadmaps, principles, and governance rules, with the Technical Committee (TC) making directional decisions. Enterprises participate as contributing organizations through SIGs, collaborating equally with other vendors. All submissions undergo public review; internal enterprise development processes do not override community workflows. openEuler's governance is neither a "disorganized community" nor "centrally controlled project management." The SIGs operate independently, determining their own goals, pace, and implementation. The TC sets annual community release goals.



Jiang Biao is the Chief Expert of Cloud Computing & Computing Power at China Telecom Group and Director of Infrastructure Software Center under eSurfing Cloud Infrastructure Division.

Q How do you define the strategic synergy between eSurfing Cloud's computing power goals and openEuler's community-driven development model?

A eSurfing Cloud and openEuler share a highly aligned vision in advancing the autonomous and intelligent transformation of digital infrastructure. openEuler delivers an agile and technical foundation, and eSurfing Cloud focuses on scenario-driven implementation. Together, we drive the evolution of China's fundamental software ecosystem.

Q What were the key technical/operational hurdles in adapting a cloud-native OS for telecom workloads, and how were they mitigated?

A During the adaptation of cloud-native OS for telecom workloads, eSurfing Cloud addressed four key challenges: inefficient image evolution, transmission security risks, multi-team resource conflicts, and delayed user feedback. We targeted measures across four dimensions:

- Image building: Logical/data separation architecture
- Security: digital signature verification
- Deployment isolation: MTC isolation
- Operational experience: webhook validation

These systematic approaches enabled stable NestOS deployment in production environments.

Q With 600+ PRs contributed, how does eSurfing Cloud measure the ROI of open source participation while ensuring sustainable engineering investment? What specific kernel features or optimizations in openEuler have been driven by real-world production issues from telecom or cloud providers?

A The mutual empowerment that comes from open source participation and internal development leads to a very productive ecosystem. Open source engagement not only accelerates project timelines and elevates product quality but also directly feeds back into project pipelines while simultaneously advancing technology, talent, and ecosystem development.
In the Linux kernel domain, eSurfing Cloud has contributed to the openEuler community through vulnerability fixes, performance optimizations, and MPTCP protocol improvements. These contributions exemplify how operational insights drive upstream innovation, reinforcing the bidirectional value flow between enterprise needs and open source evolution.

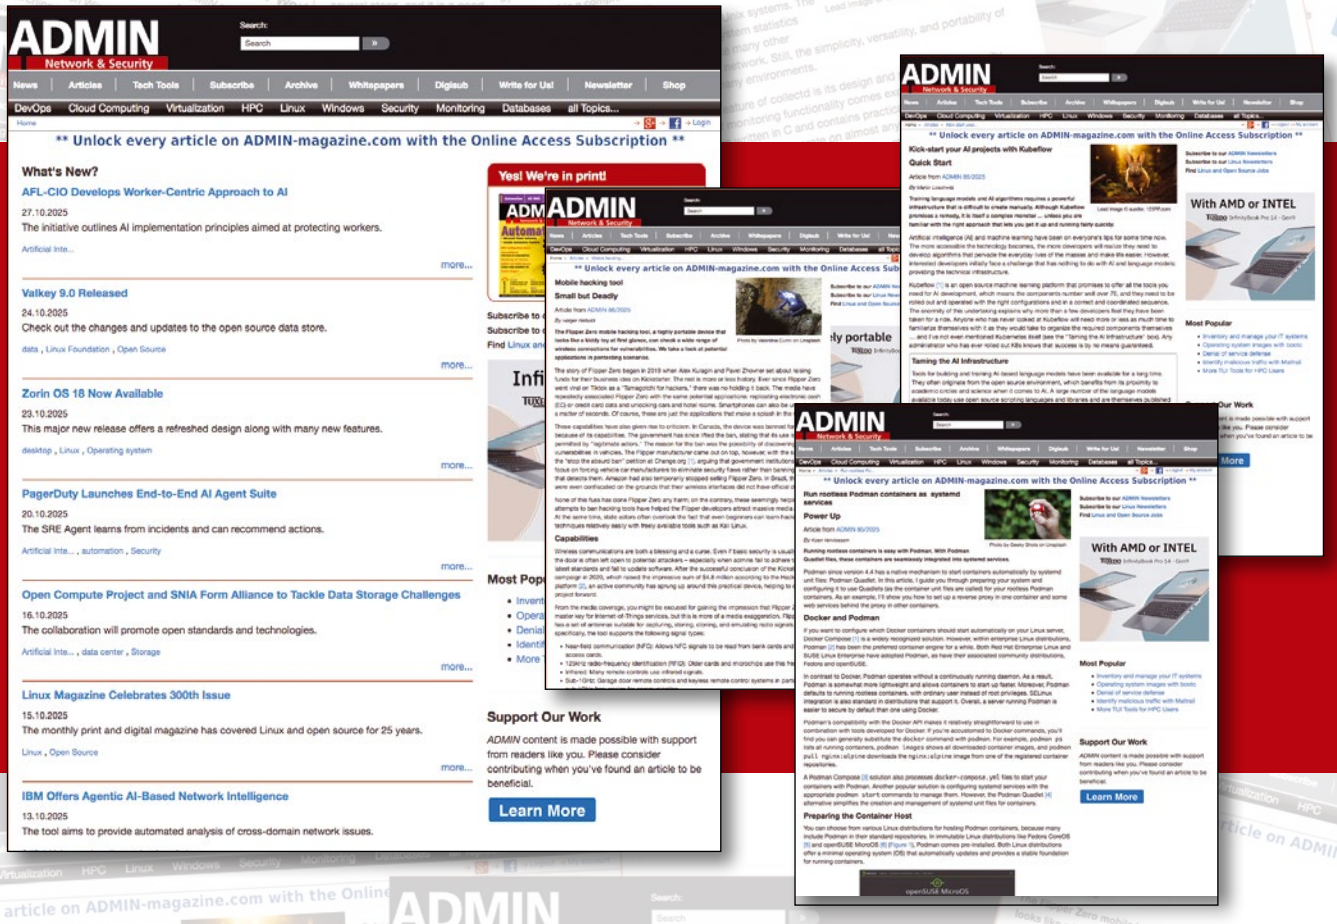
Q openEuler aims to be a "full-stack" OS for cloud-native and edge scenarios. How does its architecture address the emerging needs of hybrid cloud environments?

A As an OS designed for cloud-native and edge computing, openEuler drives innovation through its full-stack architecture. CTyunOS, built upon openEuler, further enhances cloud-native capabilities.
In terms of architectural design, openEuler adopts a layered, decoupled, and modular approach, enabling independent optimization of the kernel, container runtime (iSulad), and scheduler (Volcano) to meet the dynamic requirements of hybrid cloud and heterogeneous environments. openEuler's hot patching mechanism enables interruption-free upgrades on edge nodes, ensuring business continuity.
Regarding cloud-native efficiency, openEuler natively integrates Kubernetes and lightweight container technologies and leverages lightweight virtualization to achieve unified orchestration of virtual machines and containers.
For telecom-grade scenarios, CTyunOS further optimizes real-time performance based on the openEuler community version and embeds national cryptographic algorithms to meet the deterministic computing demands of telecom operators. ■

Get the Complete ADMIN Archive

▶ADMIN-magazine.com◀

This monthly subscription gives you access to the full ADMIN website, including thousands of technical articles, news, features, and online content back to 2010.





Activate your access today, and get access to the complete ADMIN archive!

Customers in Europe and the United Kingdom (EUR/GBP)

Customers in all other countries (USD)