THE LINUX FOUNDATION | Research

Virtual Machines, Agentic AI and the Mix of Workloads

Findings from the 2026 Cloud Native Workloads Survey



Adrienn Lawson, *The Linux Foundation*

Foreword by
Dilpreet Bindra, *Broadcom*

November 2026

Sponsored by

BROADCOM connecting everything

CLOUD NATIVE COMPUTING FOUNDATION

Virtual Machines, Agentic AI and the Mix of Workloads

53% of production applications run as VMs on average.

51% of organizations run half or more of their production applications as VMs.

71% expect their VM share to stay in the same band over the next 12 months.

64% of organizations with 10,000 or more employees run most apps as VMs, versus 45% at companies under 500.

61% of organizations say their VM and container teams are collaborative or fully integrated.

64% of VM users run VMs on-prem, ***51%*** run them in public cloud, and 23% run them on both.

34% of organizations expanding their VM footprint cite AI and agentic workloads.

2.2x more organizations plan to use confidential computing for AI workloads in the coming year than use it today (26% vs. 12%).

+10 pts in planned MicroVM sandboxing for AI workloads, from 20% today to 30% planned for next year.

Copyright © 2026 **The Linux Foundation** | November 2026. This report is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International Public License**
Please cite this research report when using the infographics. See citation guidance on the report's back page.

Table of contents

Foreword **04**

Executive summary **05**

Introduction **06**

VMs by the numbers **07**

Current share of production applications running as VMs 07

Expected share over the next twelve months 08

Factors driving VM expansion 09

Agentic workloads and security **11**

Isolation measures for AI workloads 11

GPU utilization 12

Organizational coordination between VM and container teams 12

Workload diversity and infrastructure **14**

Deployment environments for VMs 14

CNCF ecosystem 15

Conclusion **16**

Methodology **17**

Demographics 18

Data access 19

Appendix **19**

About the authors **20**

Acknowledgments **20**

Foreword

Over the past two decades, I have watched enterprise workloads evolve through wave after wave of disruption. Virtualization consolidated scattered servers. Containers redefined how we package and deploy applications. Cloud native reshaped expectations around agility, scale on demand, and developer velocity. Each shift carried a prediction: the previous technology would fade. It never did. Instead, each layer found its place alongside the next.

The rise of agentic AI is the latest wave, and it lands differently. When AI agents execute model-generated code, isolation becomes a safety requirement, not a preference. Strong, hardware-backed boundaries matter again, and virtual machines provide them. They are now the foundation for securing the most modern workloads an organization runs.

The data in this study makes that clear. Across 492 practitioners surveyed in September 2026, VMs account for a mean 53% of production applications, and most organizations expect that share to hold steady. VM and container operations are converging, with 61% reporting collaborative or integrated teams. Among organizations expanding their VM footprint, 34% cite AI and agentic workloads as the driver, and adoption of MicroVM sandboxing and confidential computing is climbing fastest among AI isolation measures. VMs are not retreating. They are evolving alongside containers, not against them.

This is what we see every day. Broadcom, through its portfolio of VMware products, delivers a single platform for all workloads — a unified developer experience on premises and in the public cloud, a broad ecosystem across infrastructure and applications, and world-class scale with operational simplicity. The convergence this study documents, from separate teams to shared operations, mirrors the path our customers are already on. I encourage you to read this report as a map of where production infrastructure actually stands. The future is not containers or VMs. It is both, managed together, secured together, and evolved together.

My thanks to the Linux Foundation for their partnership on this research, and to every practitioner who shared their experience. The picture that emerges is a measured look at where the industry truly stands, and I look forward to the discussions this report will foster.

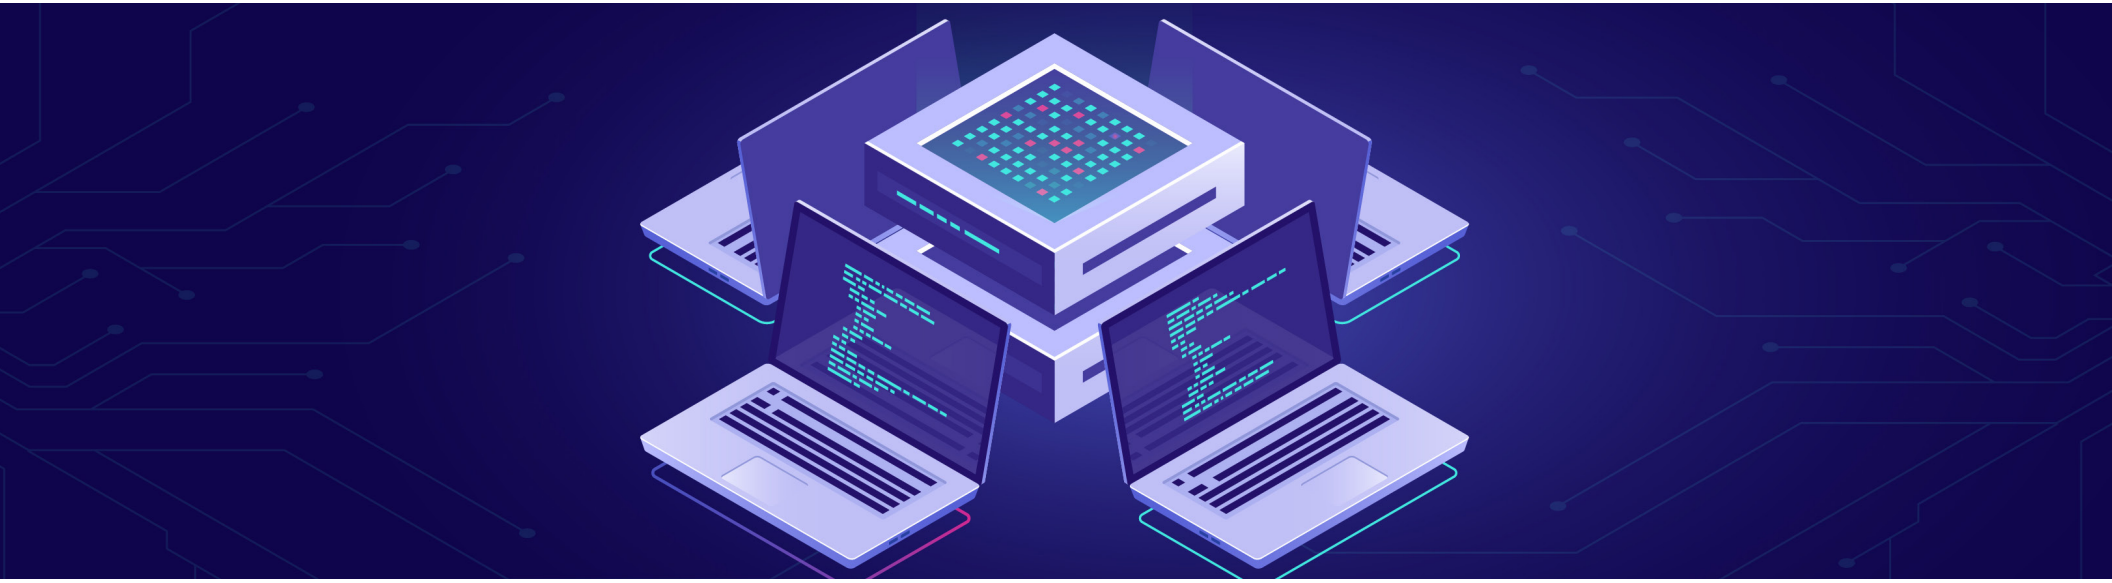
Dilpreet Bindra
Sr. Director of Engineering
VCF Division, Broadcom

Executive summary

Virtual machines (VMs) remain a core part of production infrastructure, and they operate alongside containers and other workload types rather than in competition with them. AI has also given them new strategic relevance for security and guardrails. This report draws on a survey of 492 practitioners, conducted in September 2026.

Key findings

- **VMs are a majority workload on average.** Respondents reported a mean VM share of 53% of production applications, and 51% run half or more as VMs. Most expect this to hold steady over the next twelve months.
- **Few organizations are VM-only, and very few have no VMs.** Only 10% run no production applications as VMs, and only 23% run 90% or more. Most organizations therefore run VMs as part of a mix.
- **VM and container teams largely work together.** 61% report collaborative or fully integrated VM and container teams, and 8% already run VMs on Kubernetes with KubeVirt.
- **AI and agentic workloads are a growth driver for VMs.** Among organizations expanding their VM footprint, 34% cite AI and agentic workloads. Planned adoption of MicroVM sandboxing (20% this year to 30% next year) and confidential computing (12% this year to 26% next year) grows fastest among AI isolation measures.



Introduction

Cloud native discussion has historically centered on containers and Kubernetes, and VMs are often framed as a traditional technology that containers may displace. In practice, VMs underpin a large share of enterprise infrastructure, and the rise of AI agents that execute model-generated code has renewed interest in the strong isolation and security that VM-based approaches provide. This study was designed to measure how prevalent VMs are in production today, how that is expected to change, and how organizations are securing AI workloads.

For the purposes of this study, a VM workload is a production application executing within a guest operating system on a hypervisor. The survey considered four environments in which VMs are deployed:

- **On-premises hypervisor platforms**
- **Public cloud infrastructure-as-a-service (IaaS)**, in which VMs are provisioned by a cloud provider.
- **KubeVirt**, which manages VMs on Kubernetes alongside containerized workloads.
- **MicroVM runtimes**, which are lightweight virtual machines designed for rapid startup and strong workload isolation.

VMs and containers are frequently deployed together rather than as mutually exclusive alternatives. Organizations may run VMs alongside containers, run containers within VMs, or manage either through Kubernetes or other orchestration tooling. The survey aimed to explore these insights but further research is needed to map the full details.

VMs by the numbers

Current share of production applications running as VMs

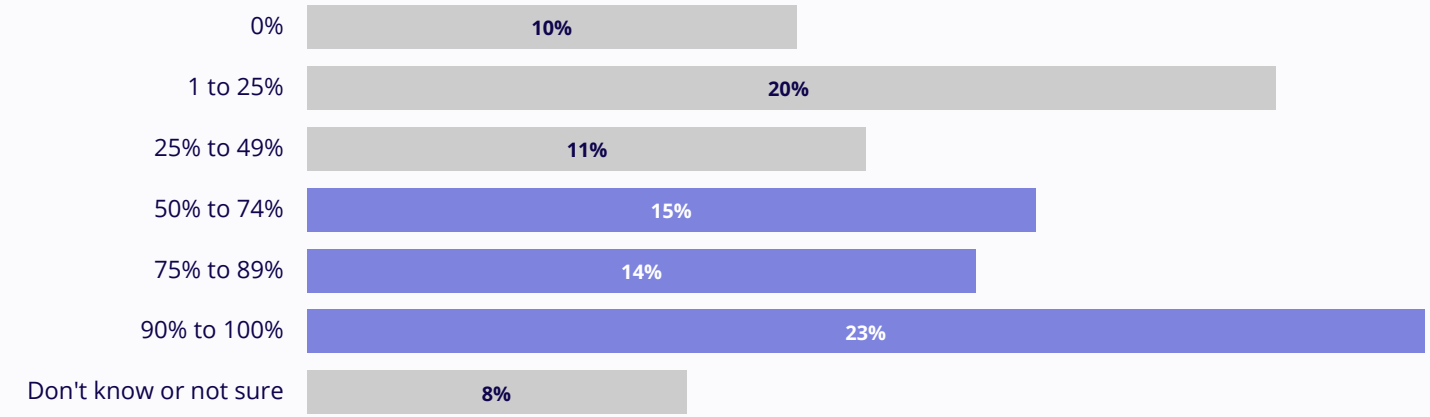
Respondents reported substantial reliance on VMs (Figure 1). Overall, 51% indicated that half or more of their production applications run as VMs, including 23% who reported a share of 90% to 100%. Only 10% reported running no production applications as VMs, and 8% were unsure. The mean reported share was 53%. Mean shares were similar for end users (52%) and vendors (54%).

Production environments rarely reflect a single technology. Applications migrated to the cloud or consolidated onto virtualized infrastructure over the past two decades often continue to run in that form alongside newer containerized services, which helps explain why VM share remains high even in organizations that have adopted Kubernetes. Virtualization improves server resource utilization, which matters more as hardware costs rise: server memory prices roughly doubled between early 2025 and early 2026[1]. Containerized environments face the same utilization problem. Cast AI reports that average CPU utilization across the Kubernetes clusters it analyzed fell to 8% in 2026, from 10% a year earlier[2], and Datadog finds that most workloads use less than 25% of the CPU they request[3].

FIGURE 1

HALF OF ORGANIZATIONS RUN AT LEAST HALF OF THEIR PRODUCTION APPS AS VMs

What share of production applications run as VMs (Virtual Machines) TODAY in your organization? (select one)



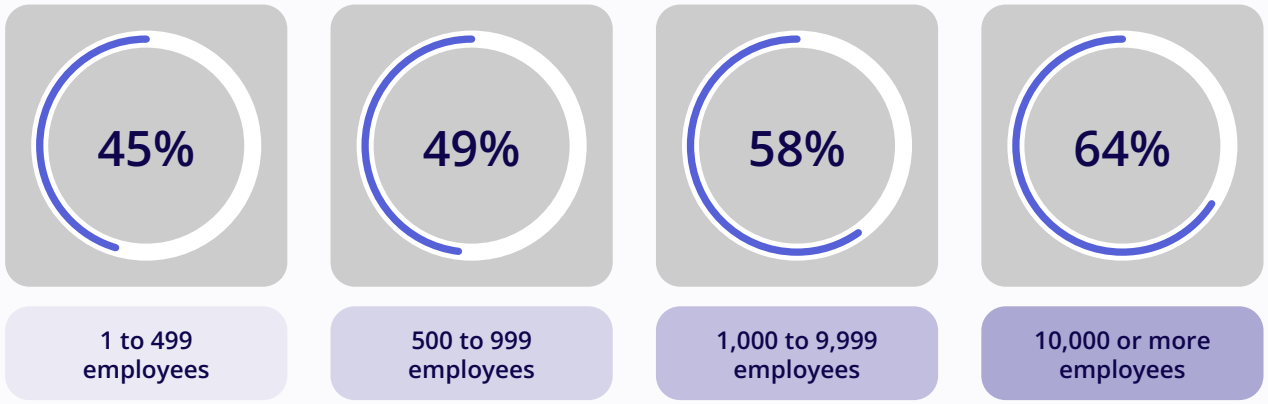
2026 Cloud Native Workloads Survey, Q5, N=492

Organization size is associated with VM share (Figure 2). The mean share increases monotonically across size categories: 45% among organizations with 1 to 499 employees, 49% at 500 to 999, 58% at 1,000 to 9,999 and 64% at 10,000 or more.

FIGURE 2

LARGER ORGANIZATIONS HAVE A HIGHER VM SHARE

What share of production applications run as VMs (Virtual Machines) TODAY in your organization? (select one)



2026 Cloud Native Workloads Survey, Q5 by Q3, N=492, averages.

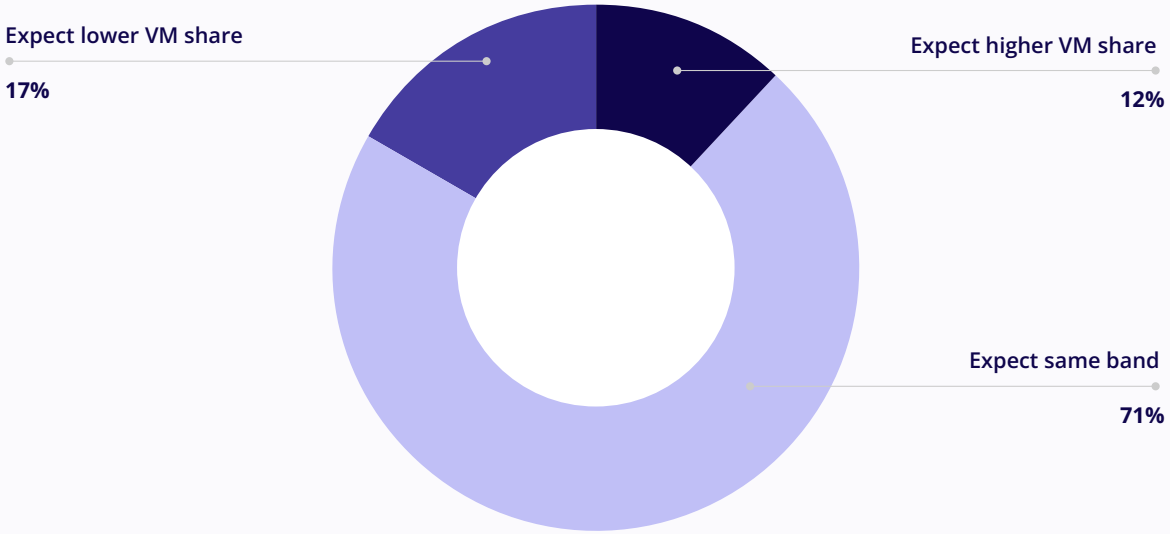
Expected share over the next twelve months

Respondents anticipated relative stability (Figure 3). Among those who answered both the current-share and expected-share questions, 71% expected their VM share to remain in the same band, 12% expected a higher share and 17% expected a lower share. Expectations did not differ substantially by respondent type: end users anticipated a mean share of 50% and vendors 53% (Appendix A1). Expectations of growth were highest among the largest organizations: 18% of those with 10,000 or more employees expected a higher VM share, compared with 7% to 12% in the other size bands (Appendix A2).

FIGURE 3

VM SHARE IS EXPECTED TO HOLD STEADY IN THE NEXT YEAR

In the NEXT 12 MONTHS, what share of production applications will run as VMs in your organization? (select one)



2026 Cloud Native Workloads Survey, Q5 and Q6 compared, N=426 (DKNS excluded)

Factors driving VM expansion

51% of respondents reported expanding their VM footprint, while 43% did not (Figure 4). This appears to conflict with the earlier finding that only 12% expected a higher VM share over the next twelve months, but the two questions measure different things. The share question was asked in bands, so an organization that grows its VM use modestly, without crossing into a higher band, would be recorded as expecting the same share. Expansion of the footprint can also occur in absolute terms (more VMs) while the share of production applications running as VMs stays constant, for example when total infrastructure is growing.

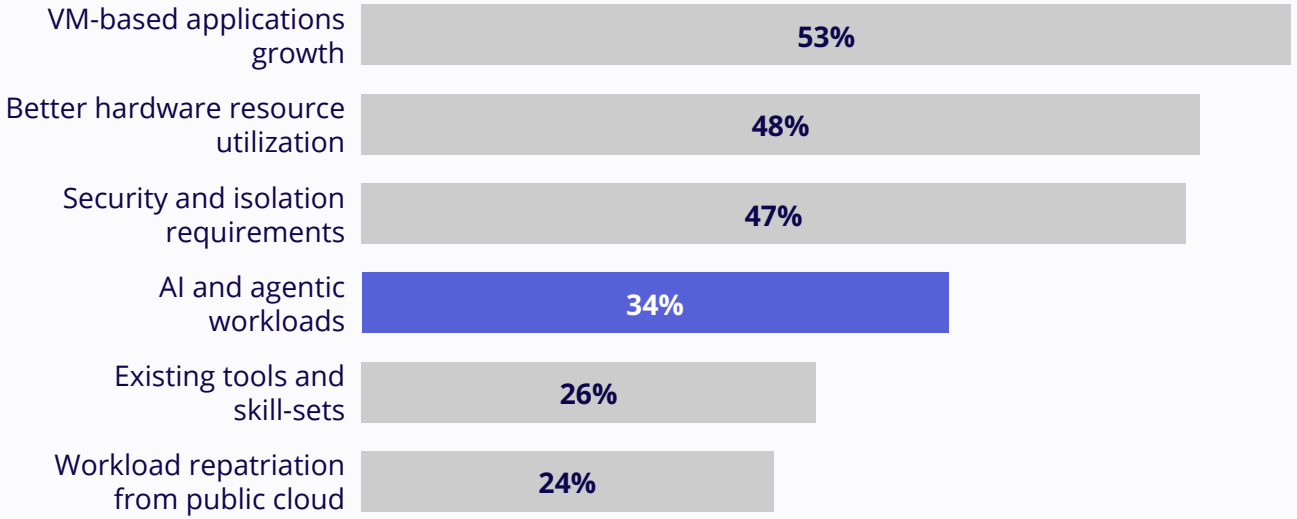
Hardware supply chain shocks have sharply raised server costs, which increases the value of getting more work from each machine. Higher utilization and stronger workload isolation are also long-standing reasons to virtualize.

Among those expanding, the most frequently cited reasons were growth in VM-based applications (53%), improved hardware resource utilization (48%) and security and isolation requirements (47%) (Figure 4). AI and agentic workloads were cited by 34%, indicating that approximately one in three expanding organizations already regards AI as a contributing factor. Existing tools and skill sets (26%) and workload repatriation from public cloud (24%) were cited less often.

FIGURE 4

GROWTH, UTILIZATION AND SECURITY DRIVE VM EXPANSION, AND AI IS ALREADY A DRIVER FOR ONE IN THREE

If your organization is expanding your VM footprint, what are the main reasons? (select all that apply)



2026 Cloud Native Workloads Survey, Q8, N=250 those expanding VM share

Agentic workloads and security

AI agents differ from conventional applications in that they may generate and execute code at runtime, call external tools and act on behalf of users with access to data and systems. Because model output is not fully predictable and agents can be manipulated through adversarial input such as prompt injection, running agent-generated code raises a distinct security problem: how to contain it if it misbehaves.

AI agents can generate and execute code at runtime, and recent incidents, such as a coding agent that reportedly deleted a production database[4] and a prompt injection flaw that let attackers escape a developer tool's sandbox,[5] show why containing that code matters. The choice of isolation boundary is central to this problem, which is why the survey examined AI isolation alongside VM adoption.

Isolation measures for AI workloads

Respondents relied heavily on VM-based isolation to secure AI workloads (Figure 5). In this report, VM-based isolation denotes any of three approaches: hardware-level VM isolation, containers within VMs, or MicroVM sandboxing. 62% of respondents reported using at least one of these approaches at the time of the survey, and 63% reported plans to do so.

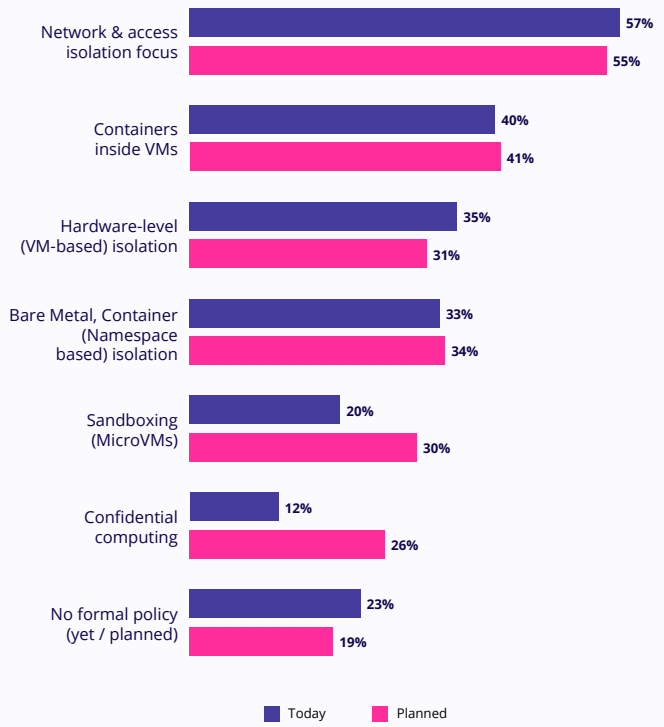
The largest projected changes involved two comparatively less common measures. Reported use of MicroVM sandboxing rose from 20% this year to a planned 30% in the next year, an increase of approximately 10 percentage points. Confidential computing rose from 12% to a planned 26%, more than doubling. Other measures were comparatively stable. Network and access isolation was the most frequently reported measure (57% today; 55% planned), followed by containers within VMs (40%; 41%).

Hardware-level VM isolation declined modestly (35%; 31%), and the share reporting no formal policy decreased from 23% to 19%.

FIGURE 5

SECURITY ISOLATION MEASURES FOR AI WORKLOADS

What security isolation measures does your organization have TODAY to secure AI workloads? (select all that apply)



2026 Cloud Native Workloads Survey, Q10 and Q11, (DKNS excluded: N=435 today, N=381 planned)

GPU utilization

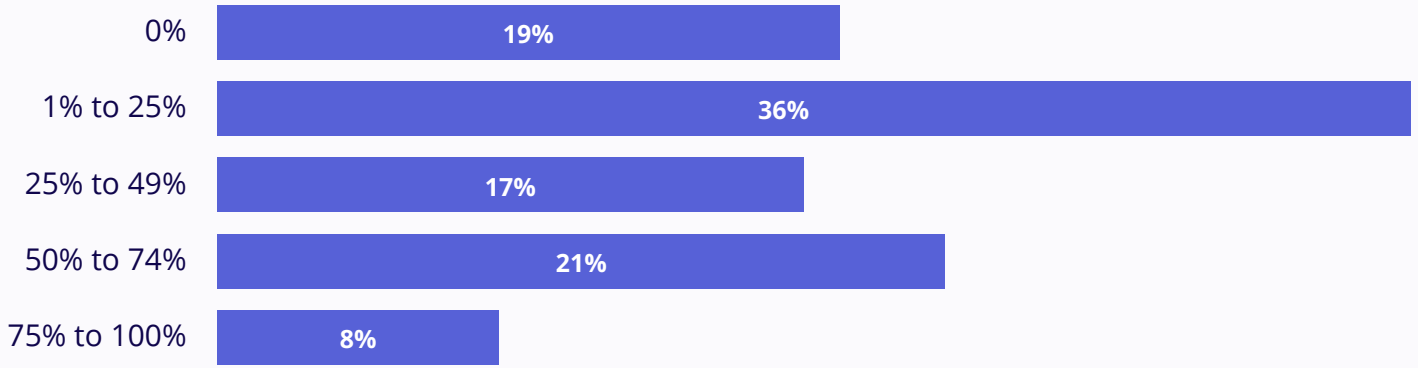
GPUs are the principal compute resource for training and running AI models, and they are expensive and, often, scarce. Sharing them efficiently is technically difficult. Such conditions can leave capacity idle, which makes utilization a possible indicator of infrastructure efficiency.

Among respondents operating on-premises GPUs, reported utilization was generally low (Figure 6). 54% reported average fleet utilization of 25% or lower, including 19% at 0%. A further 17% reported 25% to 49%, 21% reported 50% to 74% and 8% reported 75% or more. These figures exclude the 24% of respondents with no on-premises GPUs and the 17% who were unsure. GPU virtualization, which lets multiple workloads share a single physical GPU, is one approach that can raise utilization of this hardware.

FIGURE 6

MOST ON-PREM GPU FLEETS RUN AT A QUARTER OF CAPACITY OR LESS

What is the average utilization of your on-premises GPU fleet? (select one)



2026 Cloud Native Workloads Survey, Q12, (N=287 with GPUs)

Organizational coordination between VM and container teams

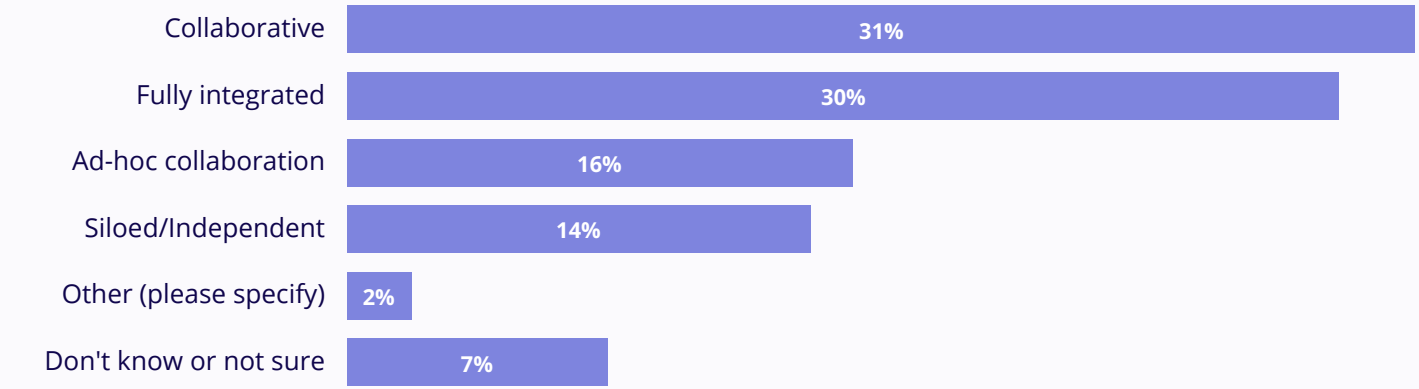
Historically, virtualization and container platforms were run by different groups with different tools, skills and operating models, and the degree of coordination between teams indicates how far that convergence has progressed.

Most respondents described coordinated administration of VM and container environments (Figure 7). 31% characterized collaboration as collaborative and 30% as fully integrated, for a combined 61%. 30% reported ad hoc collaboration (16%) or siloed operation (14%). This provides indirect evidence that VM and container environments are frequently managed in coordination, although it does not describe the technical layering of the two.

FIGURE 7

SIX IN TEN VM AND CONTAINER TEAMS WORK CLOSELY TOGETHER

How do your VM and container admin teams work together? (select one)



2026 Cloud Native Workloads Survey, Q9, N=443

Workload diversity and infrastructure

Deployment environments for VMs

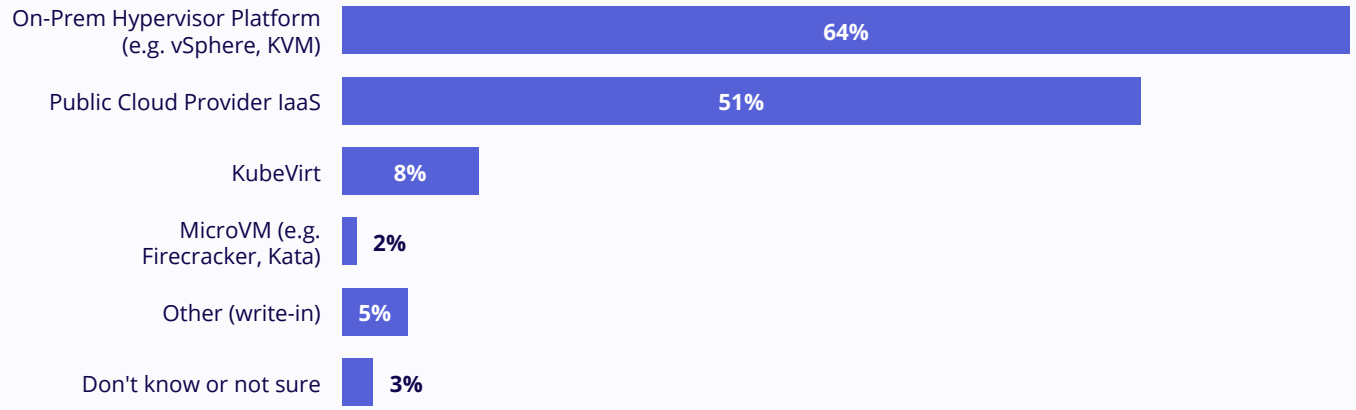
VMs can be hosted in several ways. Self-managed hypervisor platforms, such as VMware vSphere or KVM-based systems, run on an organization's own hardware. Public cloud IaaS offers VMs as a service on provider-operated hardware. KubeVirt, a CNCF project, extends Kubernetes so that VMs are defined and managed as Kubernetes resources. MicroVM runtimes occupy a further niche, providing lightweight VMs for isolation-sensitive or short-lived workloads.

Among respondents running VMs (Figure 8), 64% reported using on-premises hypervisor platforms and 51% reported using public cloud IaaS, and 23% reported using both. 8% reported running VMs on Kubernetes using KubeVirt, and 2% reported using MicroVM runtimes. Because respondents could select multiple options, these categories are not mutually exclusive, and many organizations appear to operate VMs in more than one environment.

FIGURE 8

VMs RUN ACROSS ON-PREM, PUBLIC CLOUD AND KUBERNETES

Where do most of your VMs run? (select all that apply)



2026 Cloud Native Workloads Survey, Q7, N=443 (those running VMs)

CNCF ecosystem

The cloud native ecosystem supports a broad range of workload types, including containers, VMs, WebAssembly, databases and AI agents. Kubernetes and the containers it orchestrates constitute a foundational layer, and several approaches extend this model to VMs. KubeVirt, for example, runs VMs as Kubernetes-managed resources, while VM-based container runtimes such as Kata Containers place containers inside lightweight VMs. The present survey measured VM adoption and AI workload isolation, and it did not measure the adoption of Wasm or database operators.

FIGURE 9

CNCF IS WORKLOAD AGNOSTIC: CONTAINERS, VMs, MICROVMs, WASM, DATABASES AND AI AGENTS

Containers

Kubernetes and the containers it orchestrates remain the foundation.

Virtual machines

53% average share of apps on VMs; KubeVirt runs VMs on Kubernetes.

WebAssembly

Lightweight, portable runtimes for new kinds of workloads.

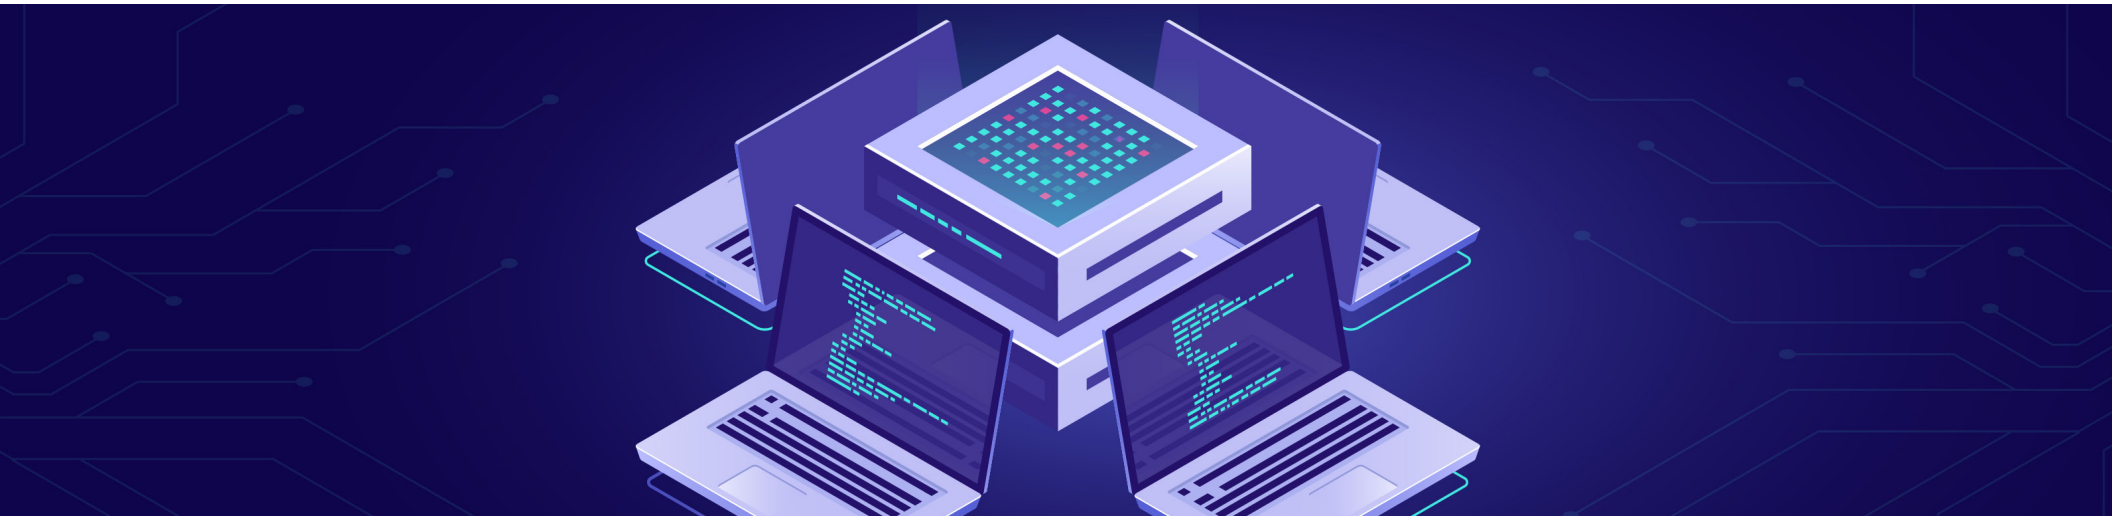
Databases

Operators bring stateful data services onto Kubernetes.

AI and agents

34% of VM expanders cite AI; sandboxing and confidential computing are growing.

2026 Cloud Native Workloads Survey, Q5, Q7, Q8. Wasm and database operators were not measured in this survey.



Conclusion

The results describe an environment in which several deployment forms coexist, often within the same organization.

VMs are majority and widespread but rarely the only form. Only 10% of respondents reported running no production applications as VMs, and 51% ran half or more in this form. Yet only 23% reported a VM share of 90% to 100%, which indicates that most organizations also run production applications in other forms.

VMs and container teams appear to be working together. 61% of respondents described their VM and container administration teams as collaborative or fully integrated, and only 14% as siloed. At the infrastructure layer, 8% reported running VMs on Kubernetes through KubeVirt. Although modest, this figure shows that a share of organizations already manage VMs and containers on a common platform.

VMs are layered more in isolation security practice. Among the isolation measures reported for AI workloads, containers inside VMs were used by 40% of respondents and planned by 41%. Namespace-based container isolation on bare metal was reported by 33% and planned by 34%. Hardware-level VM isolation (35%; 31%) and MicroVM sandboxing (20%; 30%) were also in use. The results suggest that organizations apply different isolation boundaries to different workloads.

VMs run across several environments. Sixty-four percent of respondents running VMs reported on-premises hypervisor platforms and 51% public cloud IaaS, 23% use both environments with multiple selections permitted. VMs therefore span infrastructure that organizations operate themselves, infrastructure they rent and, in a smaller number of cases, Kubernetes.

Newer workloads are adding to this diversity. 34% of organizations expanding their VM footprint cited AI and agentic workloads, and planned adoption of MicroVM sandboxing and confidential computing grew faster than any other isolation measure. VMs are therefore not only a legacy deployment form but also an active choice for new workload types.

Methodology

Goal	To measure how prevalent virtual machines (VMs) are in production cloud native environments today, how that share is expected to change over the next twelve months, and what is driving VM expansion. The survey also examined how organizations isolate AI and agentic workloads and how well on-premises GPU capacity is utilized.
Data collection method	Web survey
Survey length	12 questions
Survey target population	Practitioners and leaders involved in building, operating or securing cloud native infrastructure, including platform engineers, IT infrastructure engineers, developers, AI engineers and security professionals.
Survey distribution	Linux Foundation newsletter subscribers, Linux Foundation social media channels, research partner email/social channels (Broadcom, CNCF)
Survey fielding period	September 2026
Analytical sample size	492 qualified respondents
Margin of error (95% confidence level)	±4.4%
Unit of analysis	Individual respondents answering on behalf of their organization. Because of this, the report uses “organizations” as shorthand for respondents’ organizations. We do not collect organization names for privacy reasons, so more than one respondent may come from the same organization.
Data access	**https://zenodo.org/communities/linux-foundation-research**

Declaration of AI in the research process	Linux Foundation Research uses generative AI and AI-assisted technologies in varying degrees for research-related tasks including aggregating secondary sources, synthesizing data, supporting report drafting, copyediting, and translation. All AI-generated content is reviewed and revised as needed.
Handling of "don't know/not sure" responses	Most analyses exclude "Don't know or not sure" (DKNS) responses so percentages reflect only valid answers. When excluded, percentages are recalculated using the remaining responses as the denominator, and the figure is noted accordingly.
Additional notes	The survey measured the share of production applications running as VMs. It did not measure the composition of the remaining share (for example, containers, bare metal, serverless or other deployment forms), nor did it measure how technologies are layered within an environment. Consequently, VM-only environments cannot be distinguished from environments in which containers run atop VMs or in which VMs are managed by Kubernetes. Where the data offer indirect evidence of coexistence, this is noted. WebAssembly (Wasm) and database operators were not measured.

Demographics

The analysis is based on 492 valid responses from end users, vendors and service providers, and integrators across organization sizes (Figure 10). End users accounted for 34% of respondents and vendors or service providers for 33%. Platform and IT infrastructure engineers and leaders constituted 69% of the sample, and 46% of respondents were employed at organizations with 1,000 or more employees.

FIGURE 10

SELECTED DEMOGRAPHICS



2026 Cloud Native Workloads Survey, Q1, Q3, Q4, N=492

Data access

The raw survey data underlying this report is publicly available on Zenodo via the Linux Foundation Research community repository at **https://zenodo.org/communities/linux-foundation-research**. The dataset includes all anonymized survey responses and can be freely accessed, downloaded, and reused under the terms of the applicable open license. Researchers and practitioners who wish to conduct independent analysis or build upon this work are encouraged to consult the repository directly.

Appendix

TABLE A1

Metric	End-users (n=169)	Vendors (n=164)
Average share of apps on VMs today	52%	54%
Average share of apps on VMs, next 12 months	50%	53%
Run half or more of apps as VMs today	51%	53%
VM-based isolation for AI workloads today	59%	71%
VM-based isolation for AI workloads planned	61%	66%
MicroVM sandboxing for AI today	19%	20%
MicroVM sandboxing for AI planned	33%	31%
VM-based isolation = hardware-level VM isolation, containers inside VMs or MicroVM sandboxing. Shares exclude “DKNS” respondents.		

Source: 2026 Cloud Native Workloads Survey, Q1, Q5, Q6, Q10, Q11; AI isolation excludes “DKNS”

TABLE A2

In the NEXT 12 MONTHS, what share of production applications will run as VMs in your organization? (select one)

	Expect lower VM share	Expect same band	Expect higher VM share
1 to 499 (n=178)	15%	73%	12%
500 to 999 (n=52)	21%	71%	8%
1,000 to 9,999 (n=97)	18%	75%	7%
10,000 or more (n=94)	16%	66%	18%

Source: 2026 Cloud Native Workloads Survey, Q5 and Q6 compared by Q3, N=421 (DKNS excluded)

About the author

Adrienn Lawson serves as Director of Quantitative Research at the Linux Foundation, where she leads data-driven initiatives to understand open source ecosystems. With expertise in social data science from the University of Oxford and a background spanning academic and governmental research, she brings methodological rigor to analyzing distributed collaboration networks. At the Linux Foundation, Adrienn leads a team conducting cross-sectional research across industry verticals and geographic regions to provide comprehensive insights into open source dynamics. Her work encompasses empirical investigations into regulatory compliance, the implications of AI, and sustainable funding models. She produces evidence-based recommendations that inform strategic decision-making within the open source community.

Acknowledgments

We thank all the people who participated in the survey. Special thanks to Linux Foundation colleagues and collaborators for their involvement in the various stages of the research process: Chris Aniszczyk, Hilary Carter, Anna Hermansen, and Daniel Krook, as well as colleagues from the LF Creative Services and Hubspot teams.

Thanks to the following organizations who sponsored this study:

BROADCOM connecting everything

CLOUD NATIVE COMPUTING FOUNDATION

Endnotes

1 Datacenter Disk, "The 2026 Memory Chip Shortage: Why Server RAM Prices Have Doubled and When It Ends," **https://datacenterdisk.com/news/memory-chip-shortage-2026-server-ram-prices** (accessed October 27, 2026). Figure refers to DDR5 ECC registered server memory between Q1 2025 and Q1 2026.

2 Cast AI, 2026 *State of Kubernetes Optimization Report*, as reported in SDxCentral, "Kubernetes efficiency is going backwards as AI drives GPU waste," **https://sdxcentral.com/news/kubernetes-efficiency-is-going-backwards-as-ai-drives-gpu-waste/** (accessed October 27, 2026). The prior-year figure of 10% is from Cast AI, *2025 Kubernetes Cost Benchmark Report*, **https://cast.ai/press-release/new-kubernetes-cost-benchmark-report-reveals-persistent-cloud-waste/** (accessed October 27, 2026). Cast AI sells Kubernetes optimization software.

3 Datadog, *State of Containers and Serverless*, **https://datadoghq.com/state-of-containers-and-serverless** (accessed October 27, 2026). Utilization is measured against requested CPU, not provisioned capacity.

4 TechRepublic, "AI Agent Reportedly Deletes Company's Entire Database, Admits to Violating Guardrails," **https://www.techrepublic.com/article/ai-agent-deletes-company-database-admits-violating-guardrails/** (accessed October 27, 2026).

5 CyberScoop, "Vuln in Google's Antigravity AI agent manager could escape sandbox, give attackers remote code execution," April 20, 2026, **https://cyberscoop.com/google-antigravity-pillar-security-agent-sandbox-escape-remote-code-execution/** (accessed October 27, 2026).



BROADCOM
connecting everything®

Broadcom Inc. (NASDAQ: AVGO) is a technology leader that designs, develops, and supplies semiconductors and infrastructure software for global organizations' complex, mission-critical needs. Broadcom combines long-term R&D investment with superb execution to deliver the best technology, at scale. Broadcom is a Delaware corporation headquartered in Palo Alto, CA. For more information, visit **www.broadcom.com**.

x.com/linuxfoundation

facebook.com/TheLinuxFoundation

linkedin.com/company/the-linux-foundation

youtube.com/user/TheLinuxFoundation

github.com/LF-Engineering



THE LINUX FOUNDATION | Research

Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. By leveraging project databases and networks and committing to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations worldwide.

Copyright © 2026 **The Linux Foundation**

This report is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International Public License**.

To reference this work, please cite as follows: Adrienn Lawson, "Virtual Machines, Agentic AI and the Mix of Cloud Native Workloads" foreword by Dilpreet Bindra, The Linux Foundation, November 2026.