The State of Self-Hosted Software 2025: A Survey by Replicated

Insights from industry professionals and Replicated customers on how teams deploy, monitor, and manage self-hosted software — plus where AI fits in.

REPLICATED

September 2025

What does it really take to deliver software into today's enterprise environments?

This report reveals how commercial software teams are evolving their deployment strategies across the Commercial Software Distribution Lifecycle — from supporting air-gapped environments to navigating Kubernetes complexity — and where they're investing next.

Inside, you'll find data-backed insights on installation experiences, support coverage, scaling plans, licensing models, and how vendors, including many Replicated customers, are aligning with the needs of security-conscious, infrastructure-savvy enterprise buyers. Plus, get a look at how teams are beginning to explore AI in deployment workflows — and where they believe it can make the biggest impact.

The Commercial Software Distribution Lifecycle

Inside the 2025 Software Deployment Survey

How modern software teams are adapting to complex customer environments — and where AI fits into the future of delivery.

This report offers a data-backed look into how commercial software vendors are evolving their delivery strategies for self-hosted, air-gapped, and BYOC environments. We interviewed almost 200 technical and operational leaders to understand how they deploy, manage, and support applications inside customer-controlled infrastructure — and how their practices are shifting to meet enterprise expectations.

What We Learned

Software vendors are adapting to complex deployment demands with a focus on visibility, flexibility, and enterprise-grade control. The need for scalable delivery, faster support, and smarter automation — including emerging AI use cases — is driving the next wave of investment.

You'll find insights on installation timelines, release velocity, licensing infrastructure, and reporting practices across SaaS, on-prem, and hybrid environments. We also dive into developer challenges with tools like Helm, common pain points in release management, and emerging interest in AI to accelerate scaling, testing, and automation workflows. Whether you're evaluating vendors or building your own delivery strategy, this report surfaces key trends shaping the future of enterprise software distribution.

Section 1:

Customer Base & Scaling

Context about the current self-hosted footprint and growth expectations.

REPLICATED

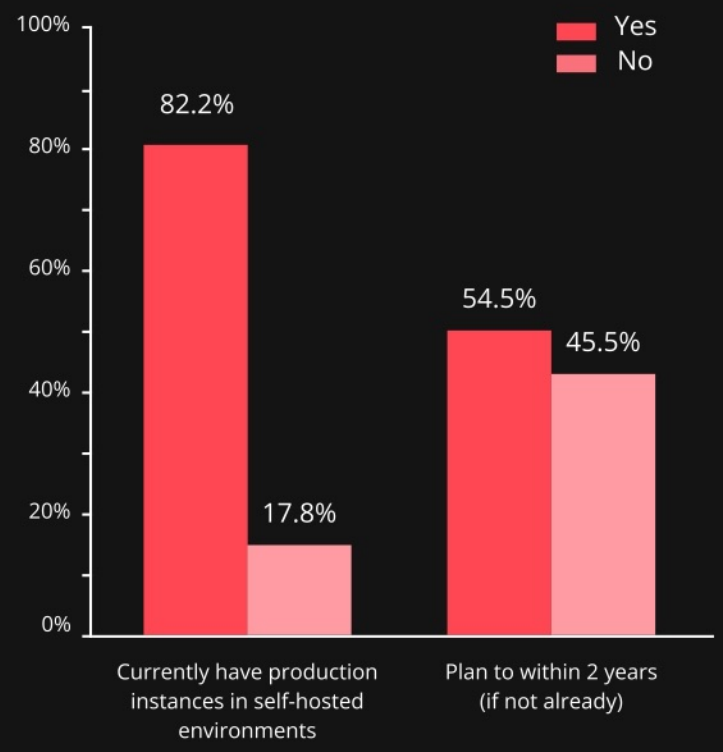

Question: Do you currently have production instances of your application(s) running in self-hosted customer-controlled environments? If no, do you plan to within the next 2 years?

Insight: 82% of vendors we surveyed already support production self-hosted deployments, and over half of those who don't plan to add them within two years — signaling long-term investment in enterprise-controlled environments.

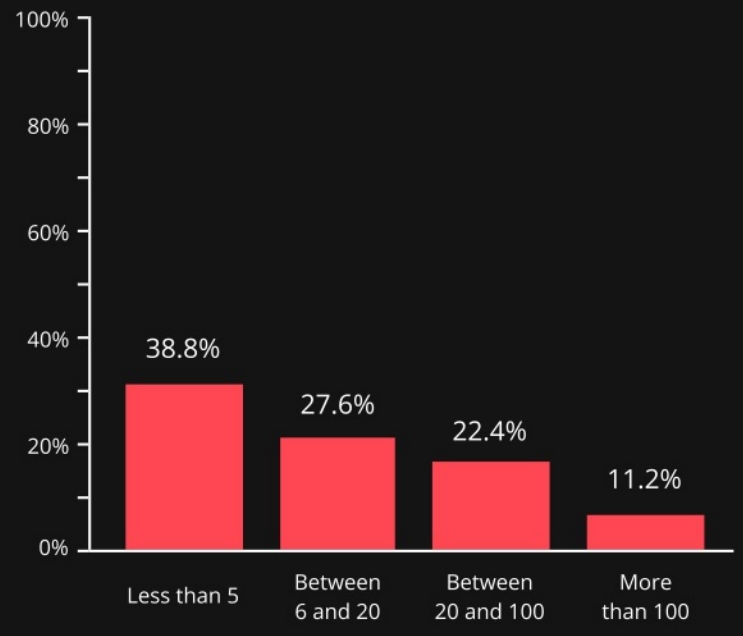

Question: How many active self-hosted customers do you have in the field?

Insight: Most companies offering self-hosted software support a small number of deployments. Our historical observations show that self-hosted deployments, though fewer in number, often serve larger or more deeply invested enterprise clients.

Question: How many self-hosted customers do you plan to scale to in 6 months?

Insight: Within 6 months, 42.9% of software vendors plan to expand their self-hosted offering, but not aggressively - staying under 25 new licenses.

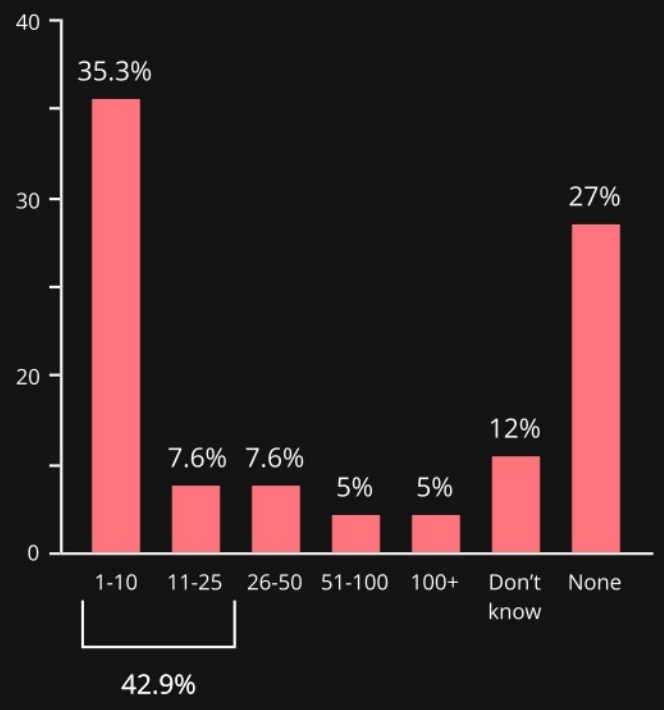

Question: How many self-hosted customers do you plan to scale to in 2 years?

Insight: Within two years, 52% of software vendors hope to scale their self-hosted offering substantially, by at least 25 new licenses. The other half expect small growth, or no growth at all.

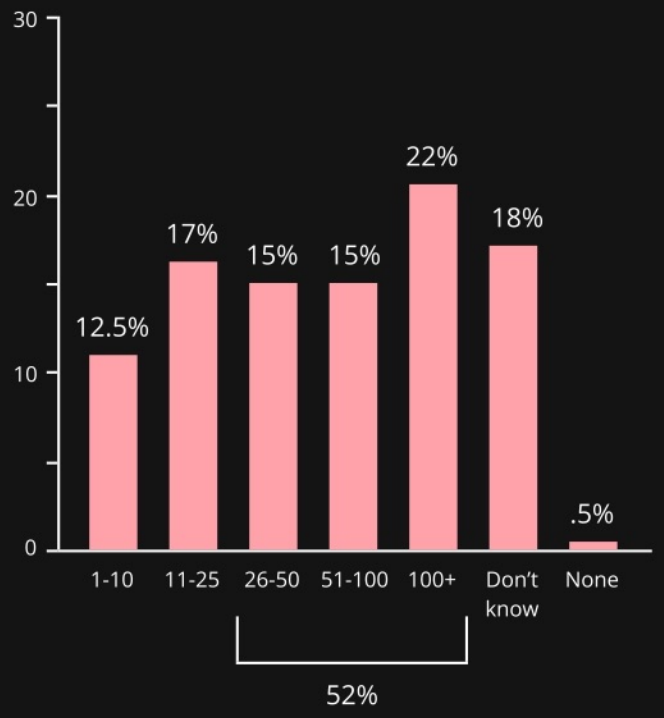

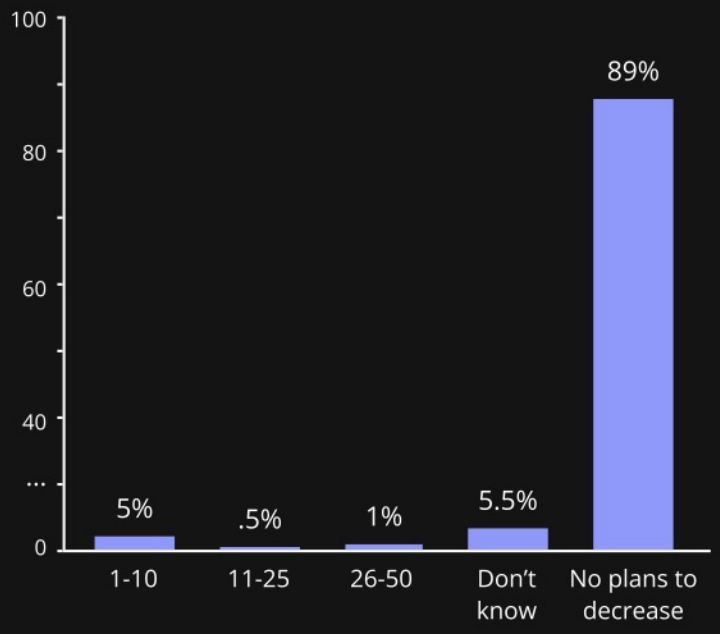

Question: If you don't plan to scale, how much do you plan to decrease your self-hosted customers by?

Insight: 89% of respondents indicated they have no plans to decrease their self-hosted customers. Self-hosted deployments continue to be an important part of their customer base, even if they don't anticipate growing substantially.

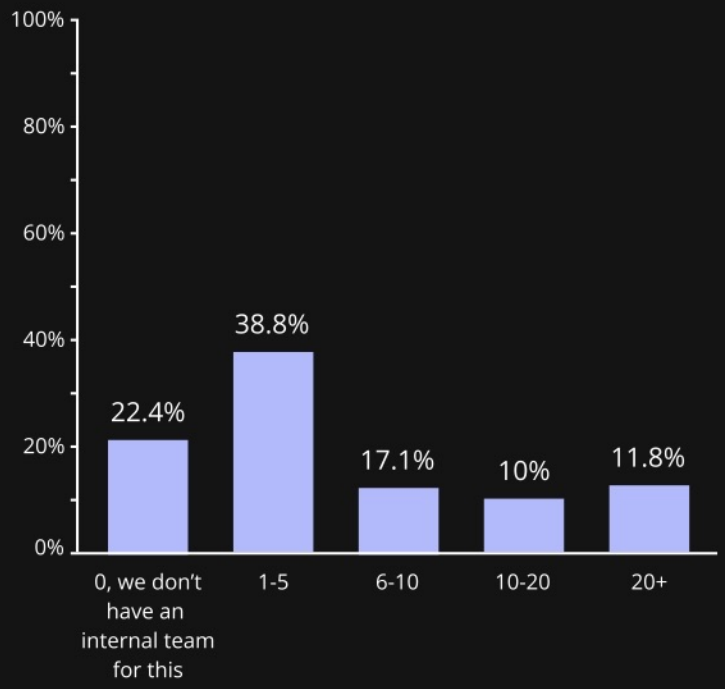

Question: How many people are on your internal development team, focused on supporting self-hosted deployments?

Insight: The largest percentage of organizations (38.8%) have small internal development teams of 1–5 people focused on supporting self-hosted deployments, while over one-fifth (22.4%) have no dedicated team at all. This indicates that support for self-hosted environments is often handled by lean or ad hoc resources rather than large, specialized teams.

Question: What categories does your application most belong?

Insight: The majority of applications fall into modern and technical domains, with AI (42.9%), Data (37.6%), and Developer tools (34.7%) leading the way. This suggests a strong concentration of respondents building software in cutting-edge or infrastructure-focused categories, while traditional sectors like Financial and Industrial remain less represented.

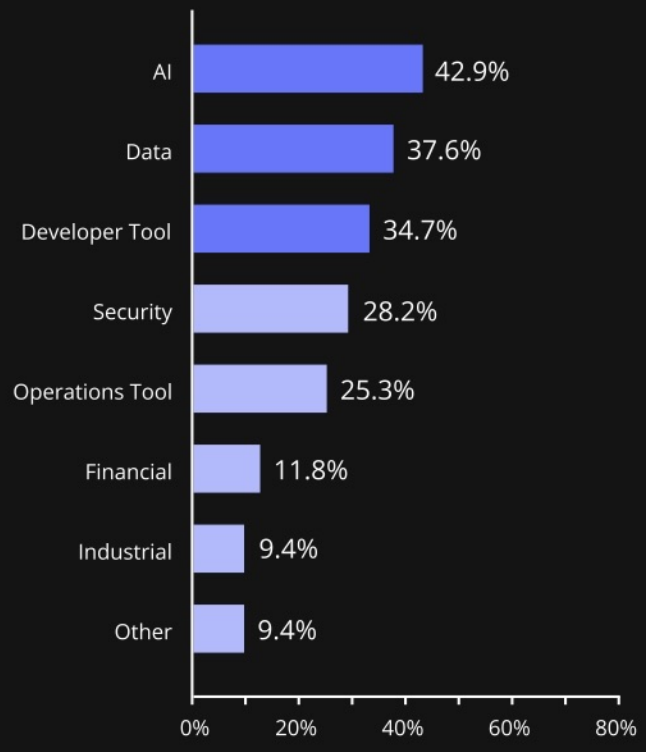

Section 2:

Deployment Models & Installation Experience

How applications are deployed and installed in customer environments.

REPLICATED

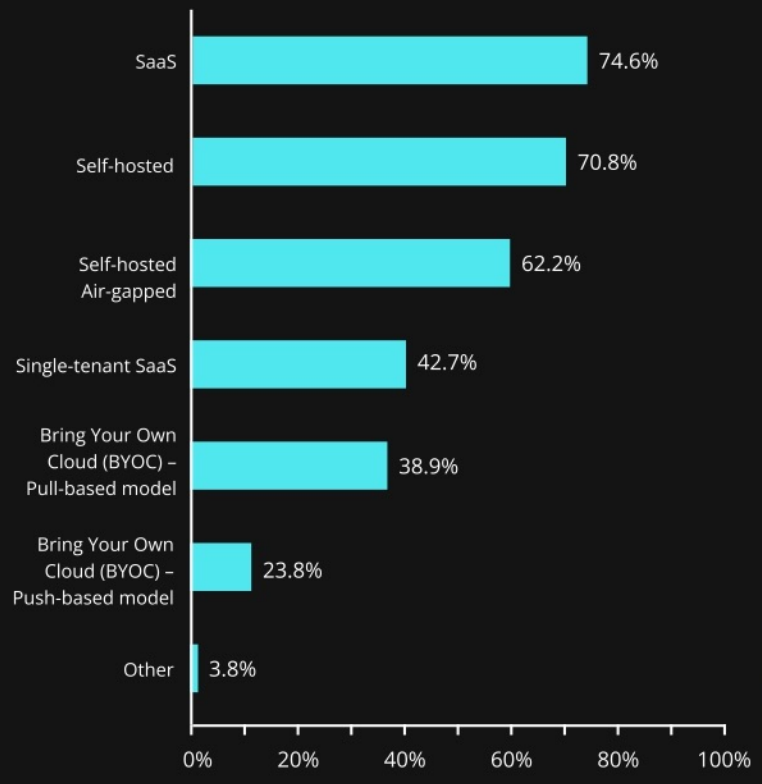

Question: What deployment strategies do you support for your application(s)?

Insight: Air gap and traditional self-hosted deployments continue to be the most popular methods to deploy into customer-managed environments, with BYOC gaining traction but still lagging far behind.

Question: What installation methods do you support?

Insight: Helm is by far the most popular tool for software installation among our respondents.

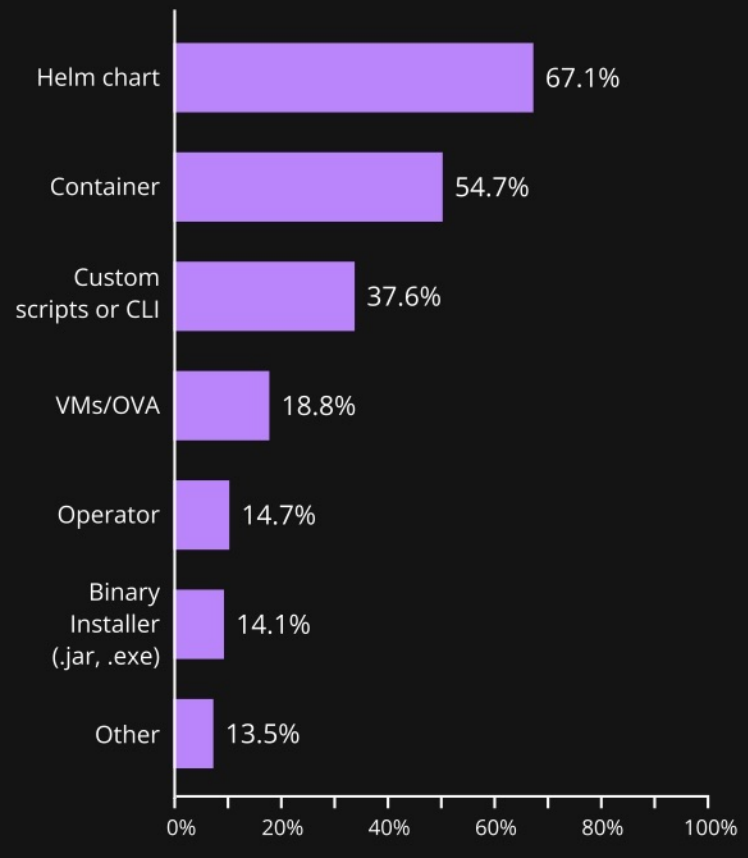

Question: Question: Do you provide tooling to help your customers provision or configure cloud resources as part of your application deployment?

Insight: Providing thorough, written documentation is the most popular way to help customers deploy self-hosted applications.

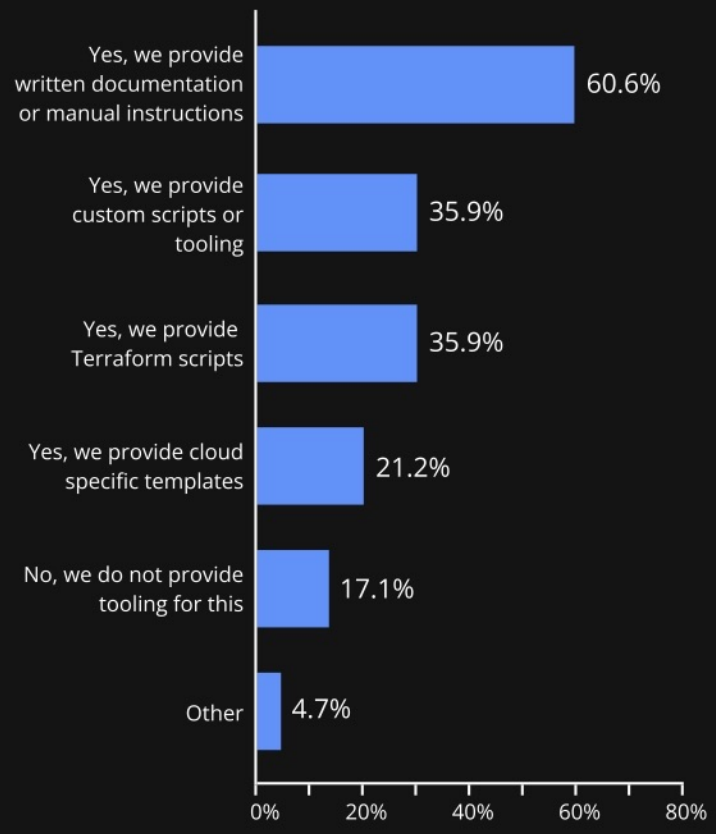

Question: How satisfied are your customers with the installation process for your application?

Insight: Most customers are reasonably satisfied — but not delighted. 65.9% of vendors report their customers are “Satisfied”, and only 12.4% report “Very Satisfied”. This means fewer than 1 in 8 vendors are delivering a truly delightful install experience.

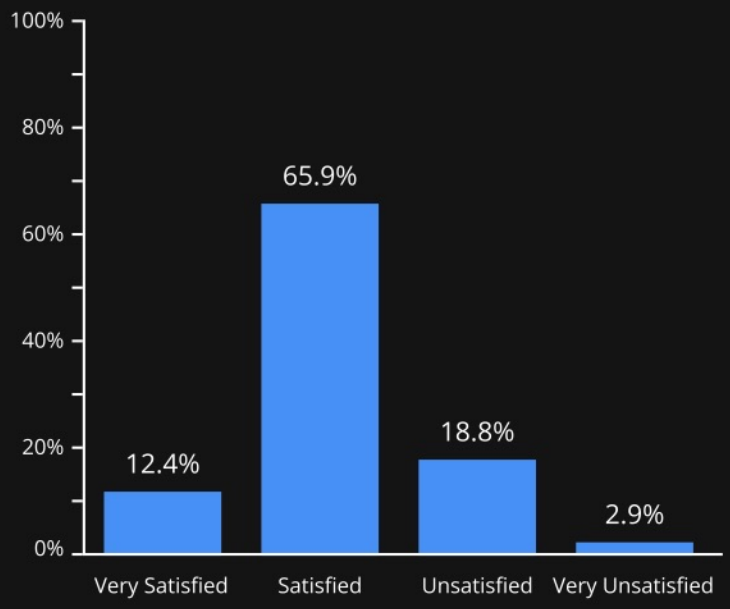

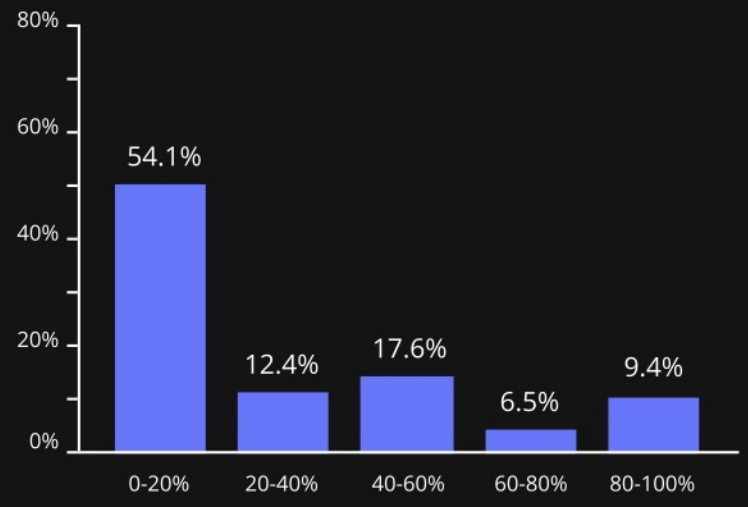

Question: What percent of your self-hosted customers are finished installing within 30 minutes of starting?

Insight: Most vendors have a slow or complex setup process. 54.1% of respondents said that less than 20% of their customers complete installation in under 30 minutes.

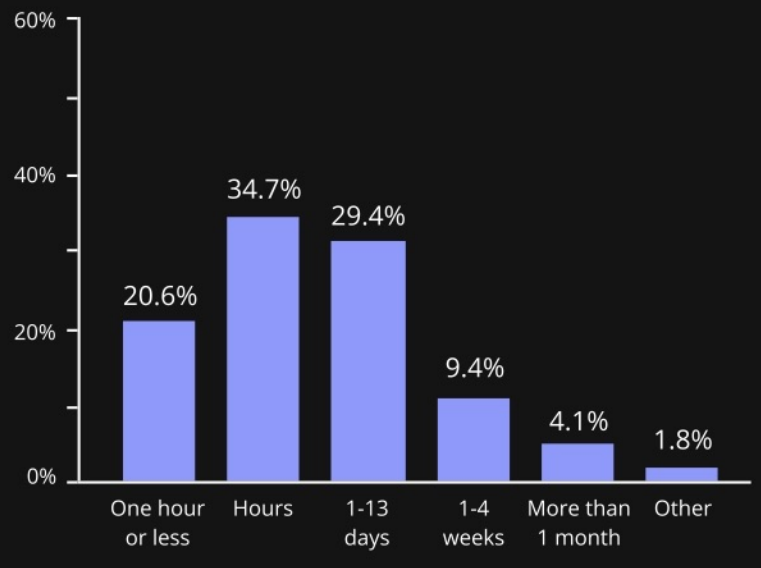

Question: On average, how long does a successful installation take?

Insight: Only one in five vendors report that a typical installation completes in under an hour, while nearly two-thirds say it takes multiple hours or even days. This highlights that "quick installs" are the exception — not the norm — in modern self-hosted software delivery.

Section 3:

Architecture & Tooling

Technologies and packaging methods used in self-hosted deployments.

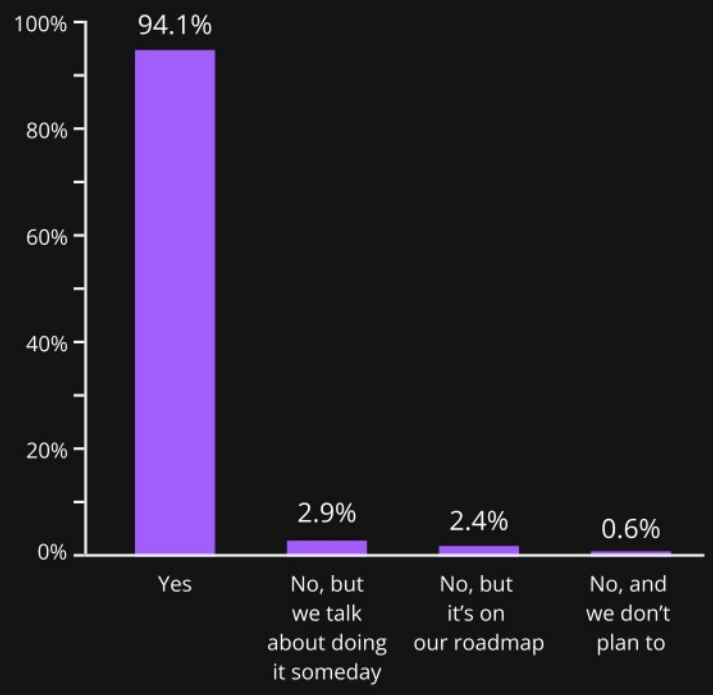

Question: Is your self-hosted application(s) containerized?

Insight: At this point, containerization of distributed applications is expected.

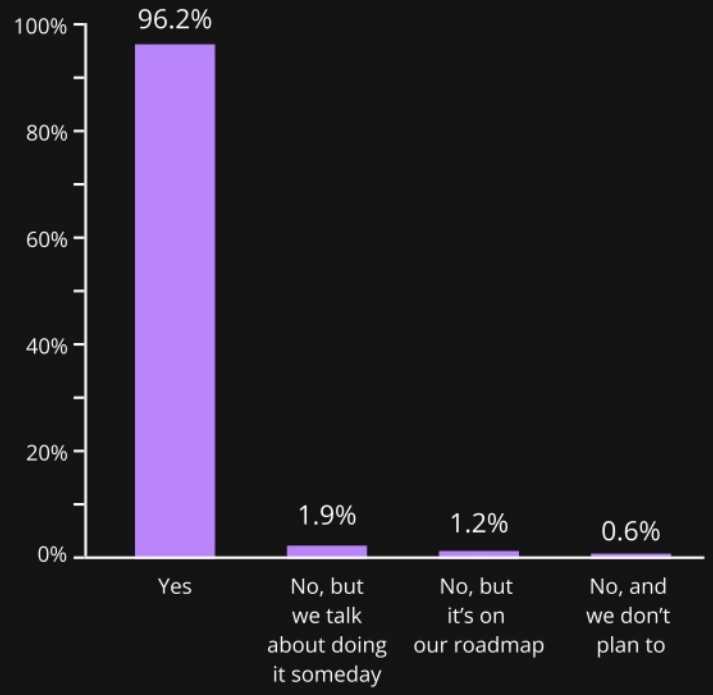

Question: If your self-hosted application(s) is containerized, do you deploy using Kubernetes?

Insight: At this point, Kubernetes is the standard for containerized self-hosted applications.

Question: Do you use Helm to package and deploy your self-hosted applications?

Insight: Helm is the clear standard for self-hosted deployments, yet a significant share of teams remain on the sidelines, signaling room for alternatives or complementary approaches.

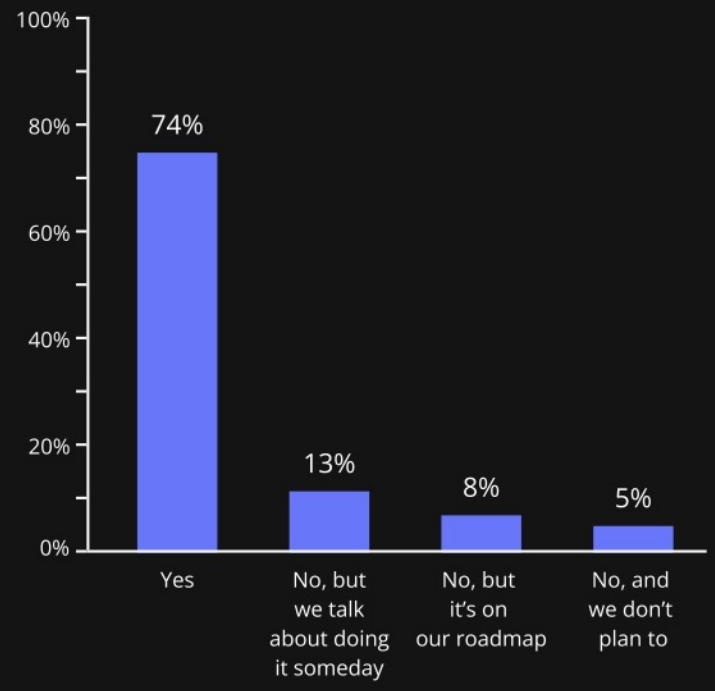

Question: How would you describe your team's skill level with Helm?

Insight: While most teams consider themselves competent or proficient with Helm, fewer than 3% identify as experts — suggesting broad adoption but limited deep specialization across the ecosystem.

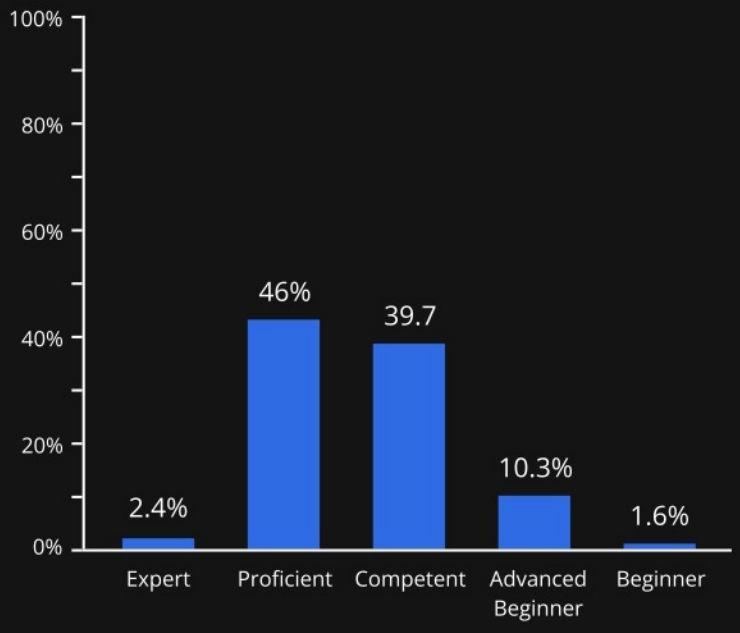

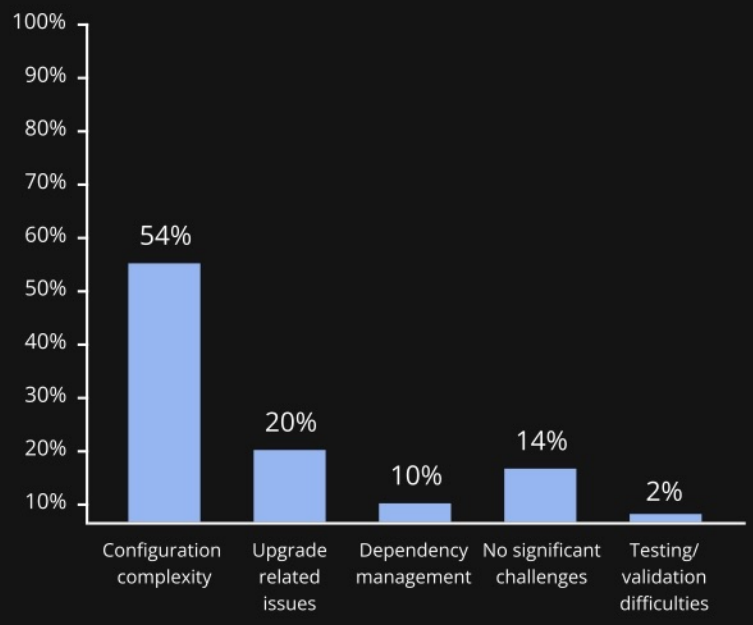

What is challenging about using Helm for self-hosted deployments?

Insight: While Helm is a powerful and widely adopted tool, its configuration model can become a challenge at scale, where managing complex values and templates often introduces significant overhead.

Section 4:

Testing & Quality Assurance

Validation across environments and configurations.

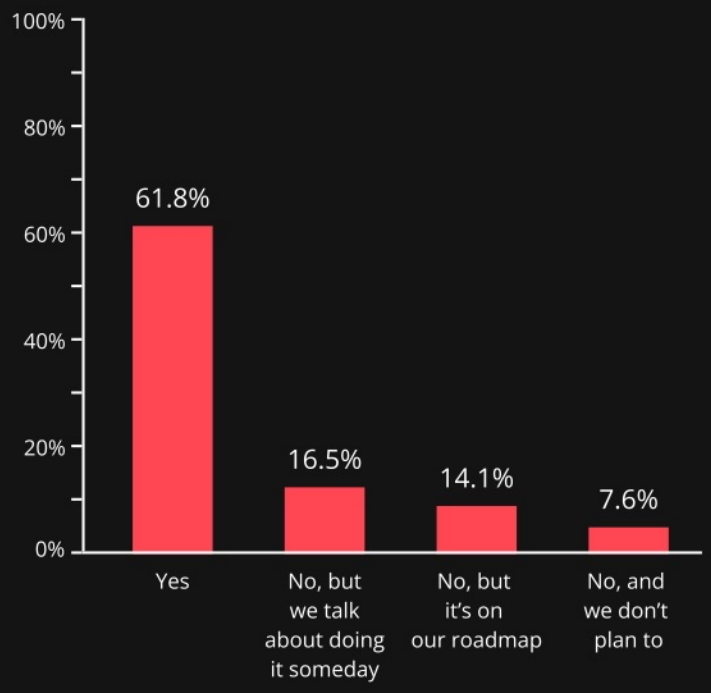

Question: Are you testing your self-hosted applications on customer-representative environments?

Insight: 61.8% of respondents say "Yes", they test in customer-representative environments. This is encouraging: it means most vendors understand that environment parity matters for reducing support issues and improving install success.

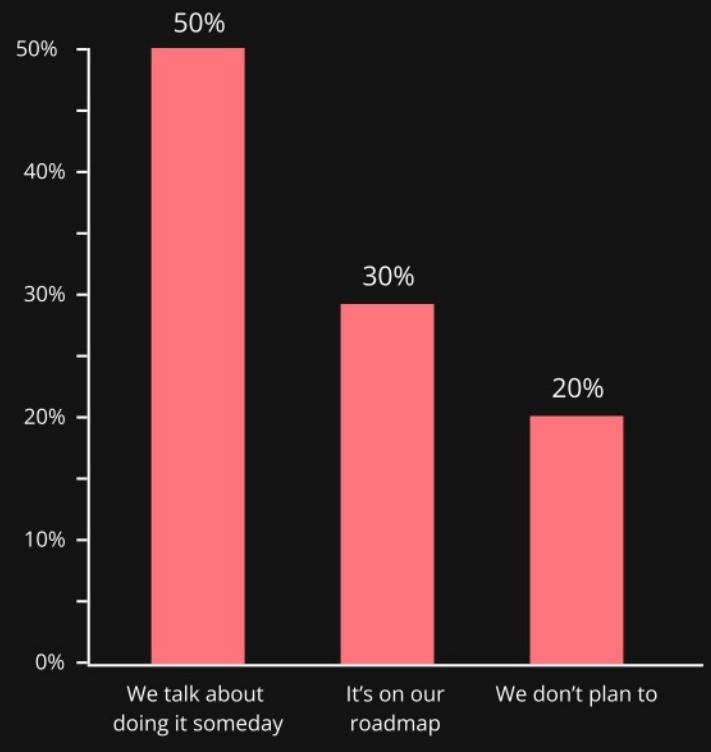

Question: If no testing is done, are there plans to introduce it and in what timeframe?

Insight: Among ISVs not currently testing on customer-representative environments, most are either considering it informally (50%) or have it on their roadmap (30%).

Question: If you are testing in customer-representative environments, what scenarios do you test?

Insight: 77.1% of respondents test for Kubernetes distribution compatibility (EKS, GKE, AKS, etc). These are among the most common deployment environments for modern enterprise customers, and testing here ensures basic compatibility across major providers.

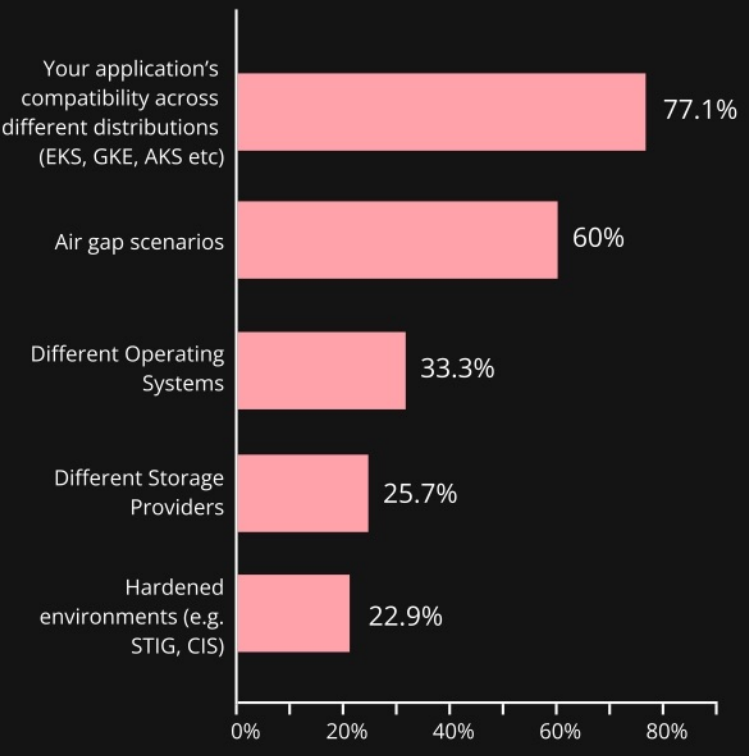

Question: If you test on different Kubernetes distributions, on which?

Insight: The majority of vendors focus testing on cloud-based Kubernetes distributions (EKS, GKE, AKS). These "big three" cloud Kubernetes offerings form the foundation of vendor testing, reflecting widespread enterprise adoption and cloud-native deployment patterns.

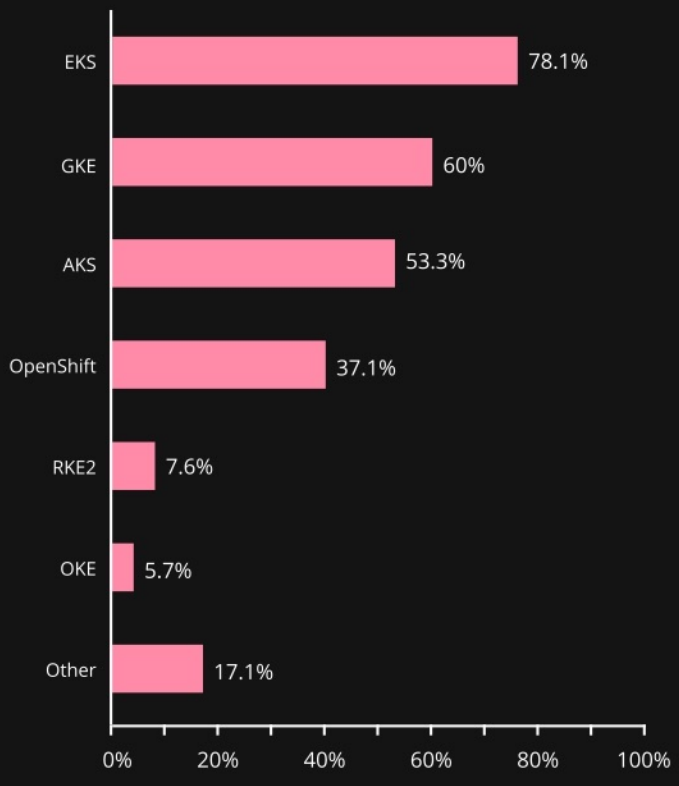

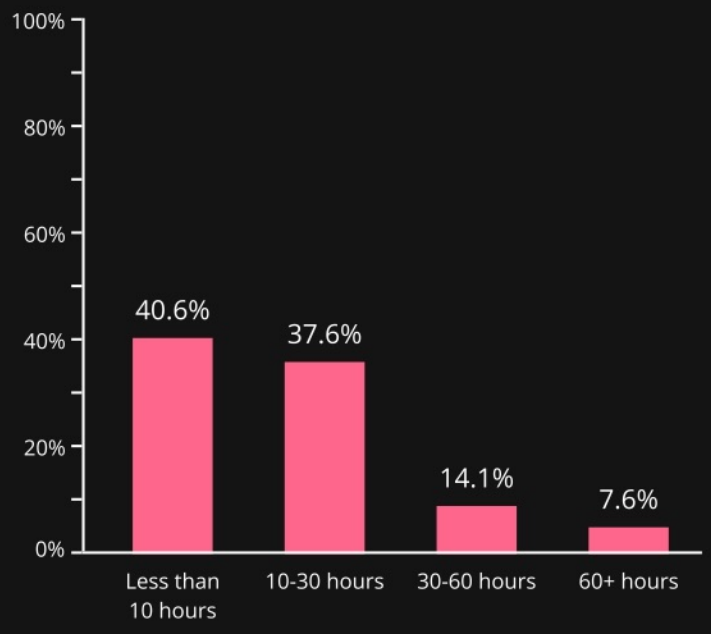

Question: How much development time does your team dedicate each month to writing and maintaining tests and test environments for your self-hosted applications?

Insight: Most independent software vendors (ISVs) dedicate less than 30 hours per month to writing and maintaining tests and test environments, with only a small minority investing more than 60 hours.

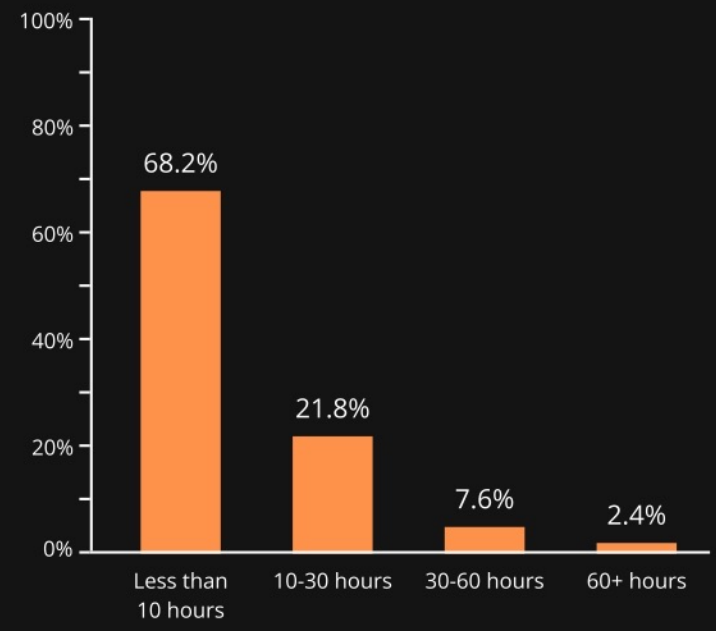

Question: How much time do you want to spend?

Insight: ISVs want to limit the amount of time they spend writing and maintaining tests (likely in favor of more automation or better tooling), with over two-thirds (68.2%) preferring to spend under 10 hours each month.

Question: How much money is your team spending each month on clusters dedicated to help test your application across different distributions and configurations?

Insight: Most ISVs (72.3%) are spending over $100 per month on clusters for the purpose of testing their app on multiple distros, suggesting that cross-environment testing is a meaningful part of their release process.

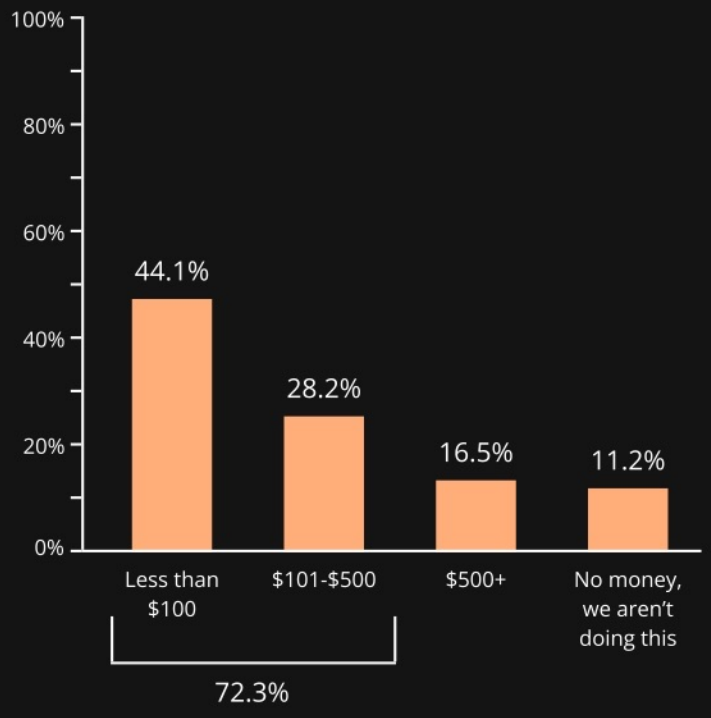

Question: How much money do you want to spend?

Insight: A large amount of respondents (35.3%) would prefer to spend less than $100 per month, highlighting a desire to reduce costs even among those currently spending on testing infrastructure.

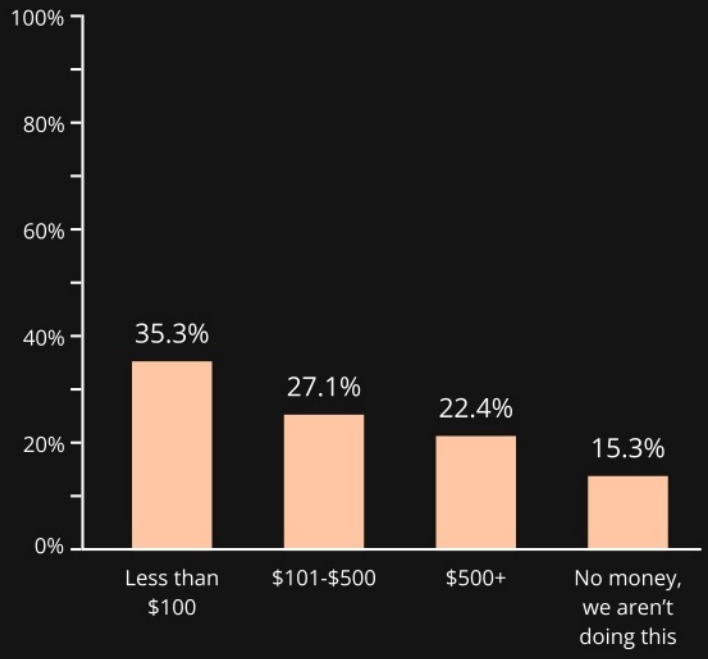

Section 5:

Release Management & Adoption

Release cadence and customer upgrade patterns.

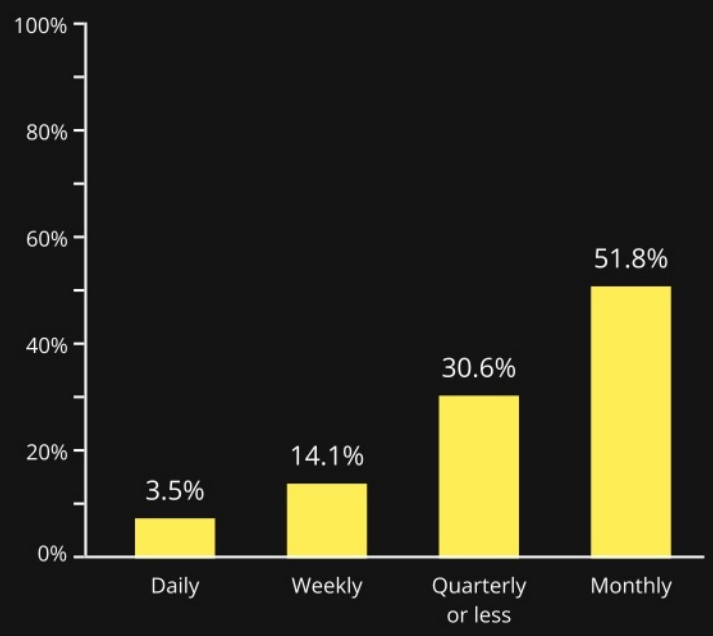

Question: How often do you release your application to self-hosted customers?

Insight: Longer release cycles (between monthly and quarterly) are the norm for self-hosted applications, with less than 20% of ISVs releasing weekly or daily.

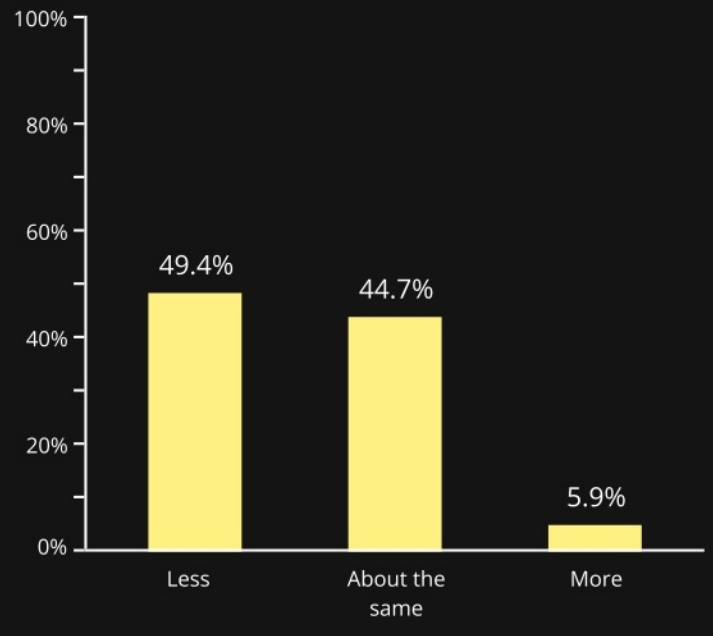

Question: Do you deliver to your self-hosted customer more or less often than your SaaS or managed offerings?

Insight: About half of vendors release their self-hosted offering less frequently than their SaaS or vendor-managed product.

Question: What percentage of customers are on your latest release version?

Insight: While most vendors have at least 50% of customers on the latest version, only about one-fifth have reached 90% adoption for the latest version. Additionally, 10% of vendors report having no way of knowing which versions their customers are on, pointing to a lack of visibility into customer instances.

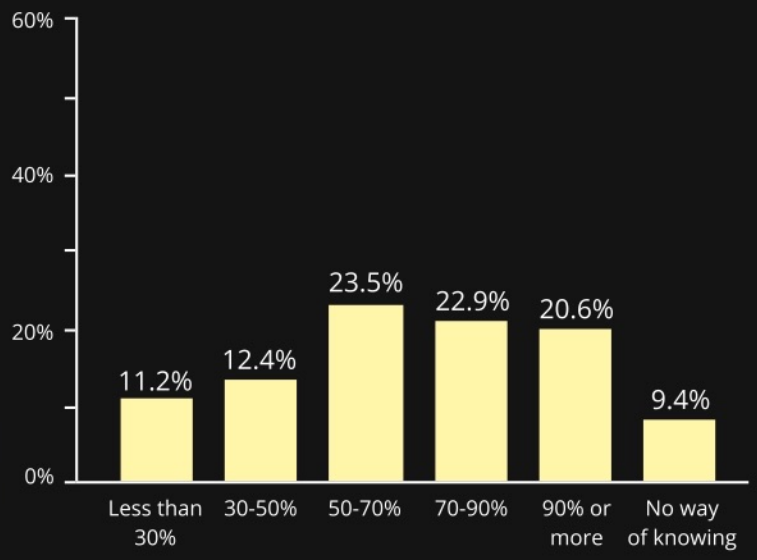

Question: On average, how long after the release of a new version do your customers upgrade?

Insight: Most ISVs report that customer upgrades lag weeks or months behind a new release, with over 40% saying it usually takes months. This type of delay can impact support, docs, security, new feature development, and more.

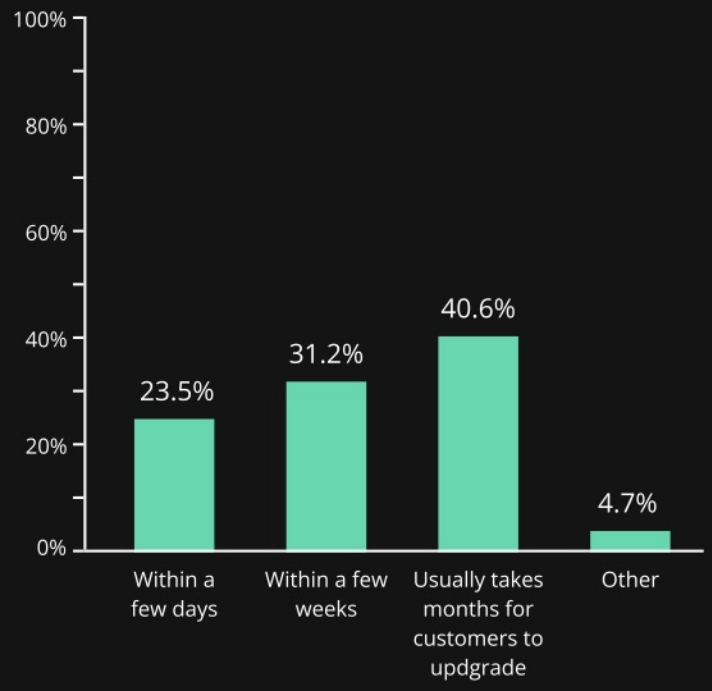

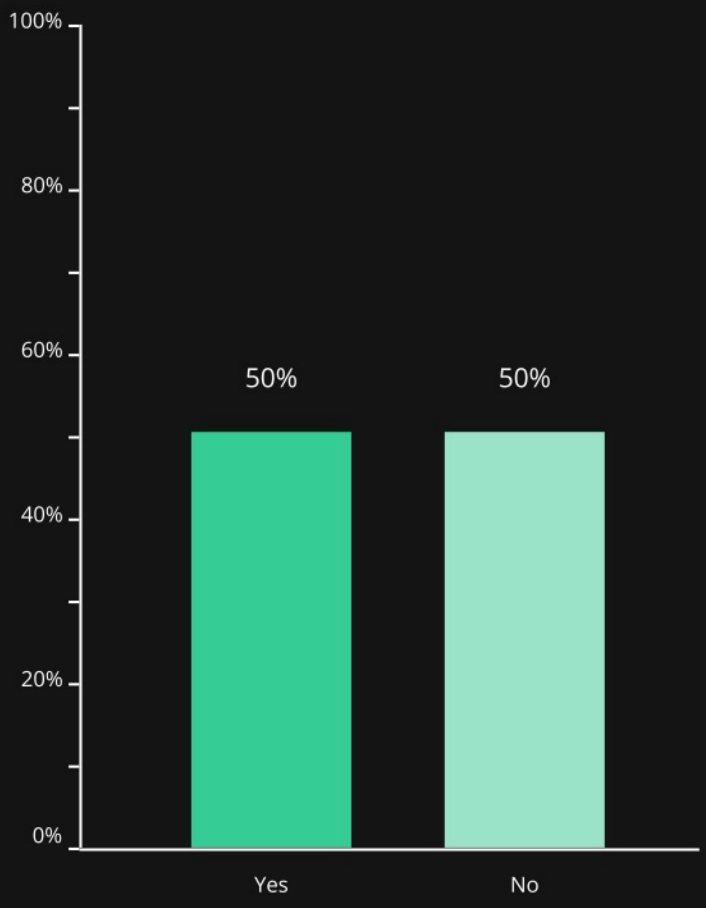

Question: Are you happy with how quickly your customers adopt your new releases?

Insight: This 50/50 split suggests that customer release adoption speed is a significant pain point across the industry. The vendors who are dissatisfied likely face challenges with:

- Support burden from customers running multiple versions
- Security risks from outdated deployments
- Difficulty rolling out new features and improvements
- Increased complexity in maintaining backward compatibility

The equal split indicates this is a universal challenge rather than something that affects only a subset of vendors, making it a critical area for improvement in commercial software distribution strategies.

Section 6:

Licensing & Entitlements

Mechanisms for managing customer access and entitlements.

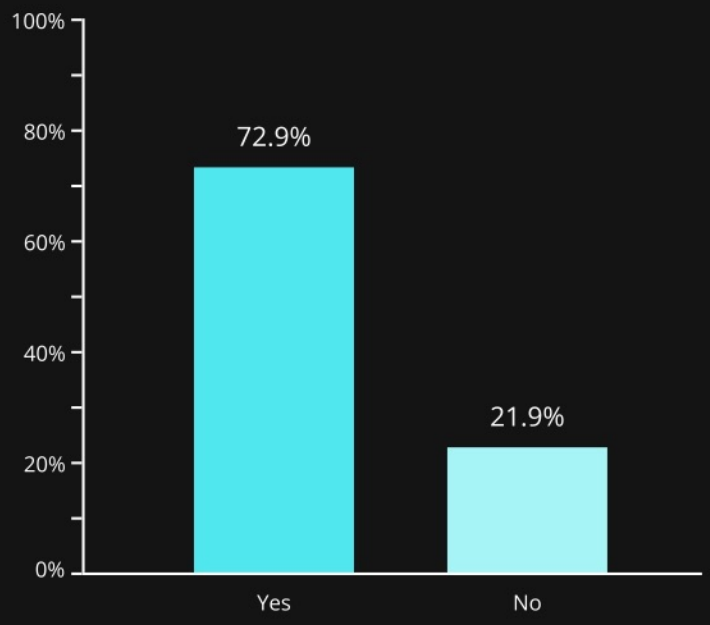

Question: Do you have a licensing mechanism to manage custom licenses, entitlements, etc?

Insight: Most vendors have licensing in place, suggesting that licensing is becoming a more important aspect of the Commercial Software Distribution Lifecycle. This is a shift from the beginnings of the Open Source and open core eras, when licensing was generally based on trust.

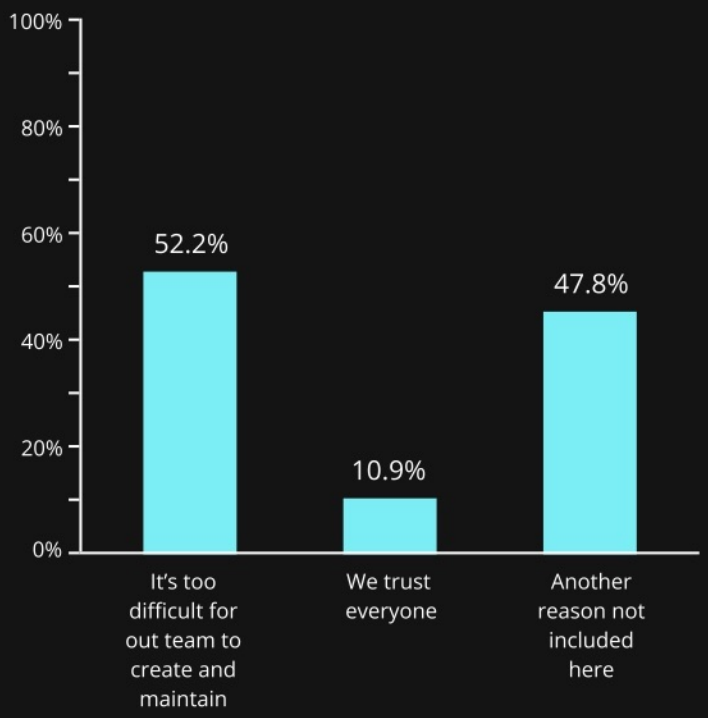

Question: If you don't have licensing, why not?

Insight: Taken together, complexity and trust represent the majority of responses as to why a vendor doesn't have software licensing in place. Given that such a small minority use a licensing mechanism, it's likely that simpler implementations would increase the use of licensing.

Question: If you don't have licensing, do you plan to introduce license/entitlements?

Insight: Most vendors without licensing intend to introduce it, but inertia is strong within this group. Only one in four have no intent to implement it. Licensing is valued, but it's not being prioritized over delivering features and supporting customers. This further argues for simpler solutions that can help make licensing nearly universal.

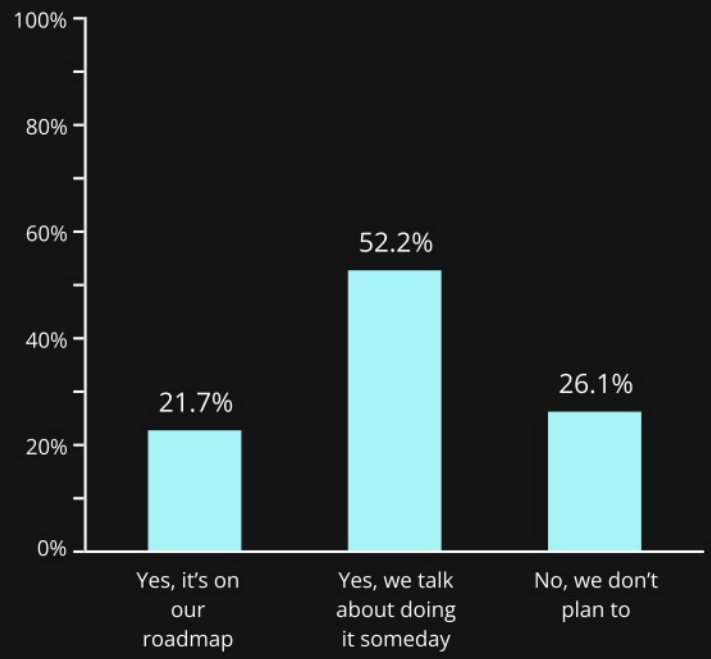

Question: Is your license server integrated with cloud marketplaces?

Insight: Only 20% of vendors integrate their licensing with cloud marketplaces. This represents a potential miss on a massive revenue opportunity: procurement through marketplaces is generally lower friction and often taps into pre-approved budgets. Vendors who do integrate into marketplaces can benefit from a significant market opportunity.

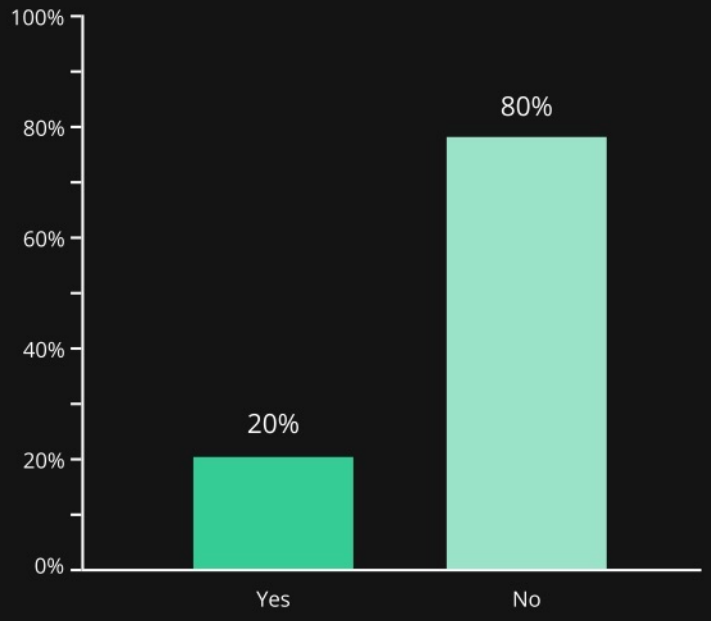

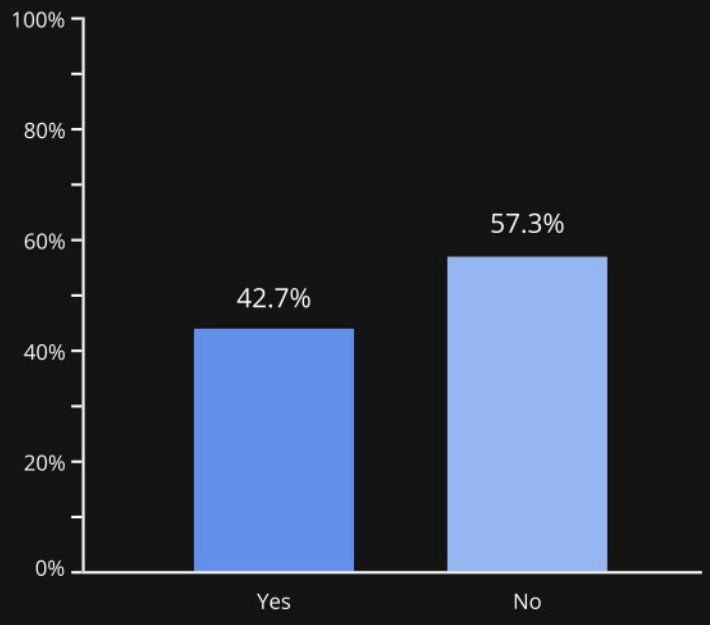

Question: If Yes, is your license server integrated with your CRM tools (Salesforce, Hubspot, etc)?

Insight: Even though nearly three-quarters of vendors implement licensing, about 60% of those do not integrate their licensing with CRM systems. This can lead to poor customer experiences like delays in license activation and mismatches between contracts and entitlements. Vendors without tight integration are at a competitive disadvantage when it comes to scale, customer support, and insights into how their products are being used.

Section 7:

Reporting & Telemetry

Data collected from deployments and how it is used.

Question: What types of reporting are you gathering about your application?

Insight: This data is dominated by operational metrics, both high-level measures like uptime and lower-level details like disk and memory usage (though those are skewed toward Saas). Product insights are collected from on-prem instances at a significantly lower rate than they are for SaaS, suggesting those insights are highly valued but can be difficult to gather for on-premise delivery.

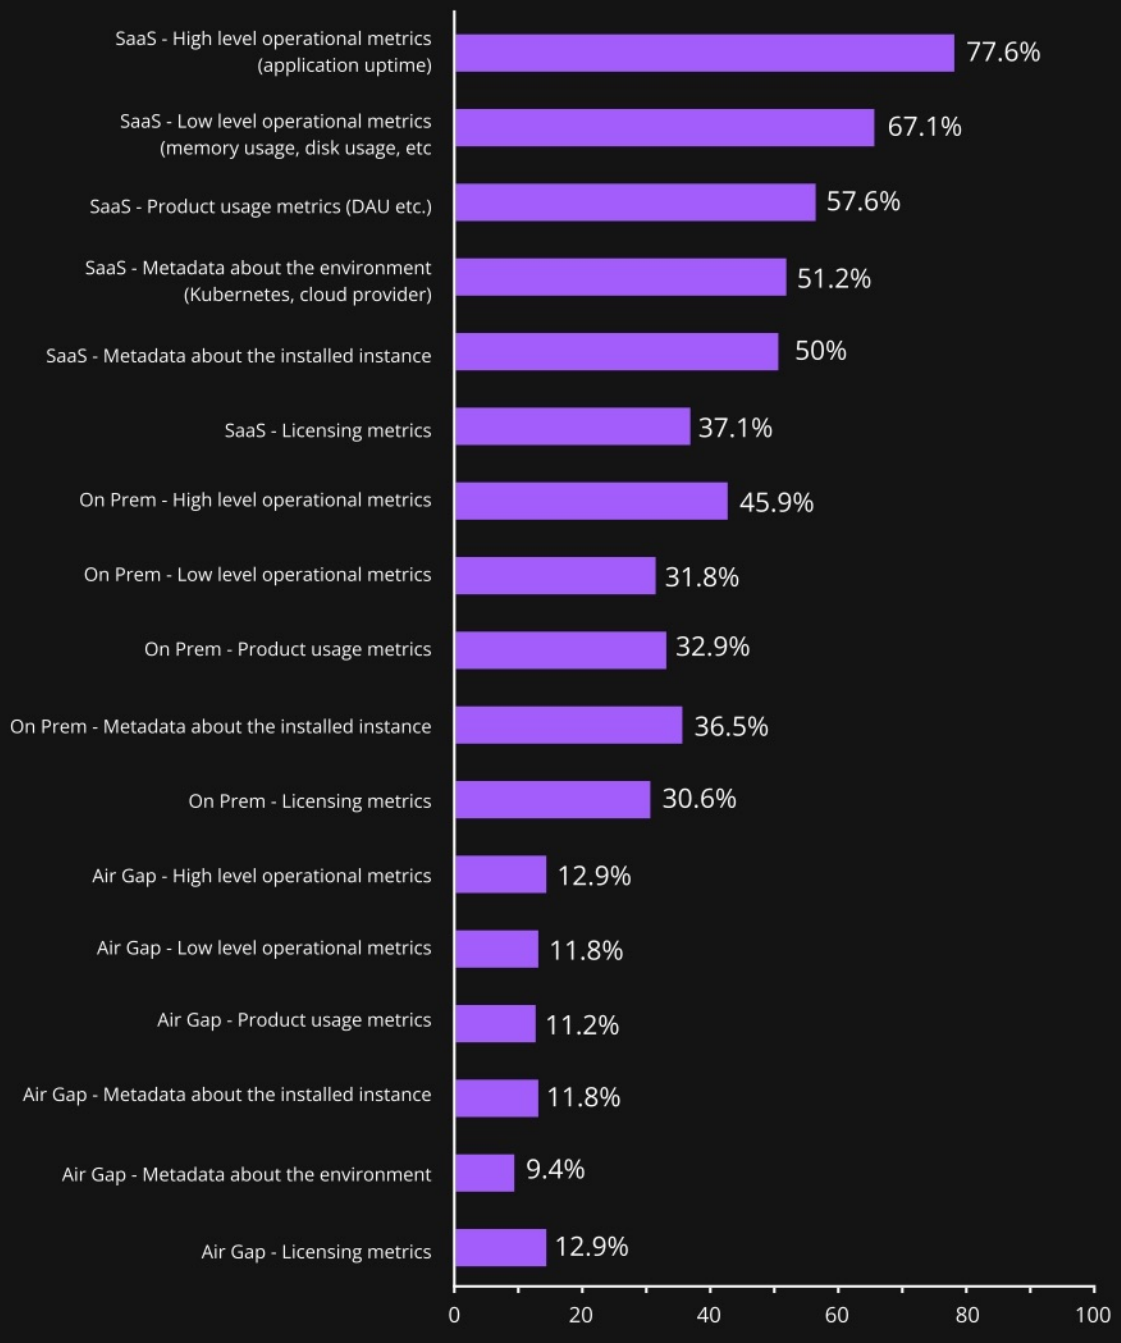

Question: If you could add reporting of one metric to your self-hosted deployments, which metric would you add?

Insight: A full 36% of respondents indicated that they either don't track additional metrics or feel they already have what they need—suggesting either a lack of formalized reporting strategy or satisfaction with defaults. With most other responses focused on product usage, vendors are increasingly invested in understanding customer engagement, feature adoption, and user behavior—even in on-prem environments.

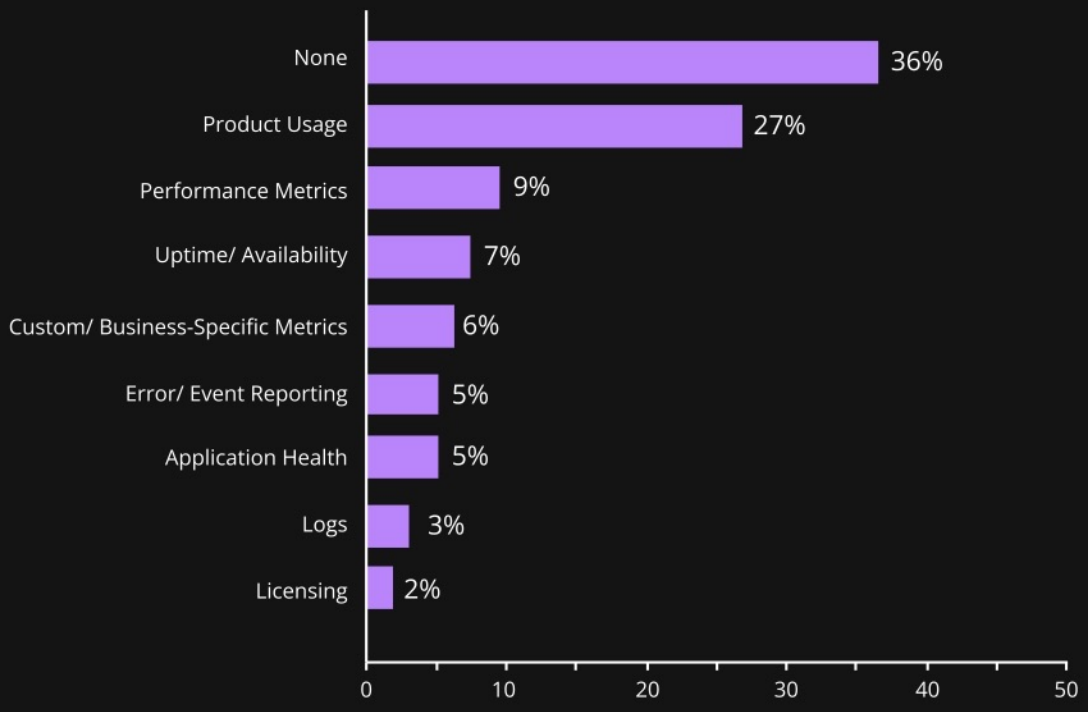

Question: Who uses the reporting gathered by your self-hosted deployments?

Insight: Engineering teams are extremely data-heavy, but may not be taking into account the business context of the data they gather. Likewise, product teams are using a lot of data, but the data they most value are being collected at much lower rates. Lastly, sales and marketing aren't taking advantage of the data available to them and missing significant opportunities to understand their customers and maximize their revenue.

This may also suggest a "data hoarding" problem (probably not intentional) where engineering teams build in data collecting for their needs without considering what other teams may want and need.

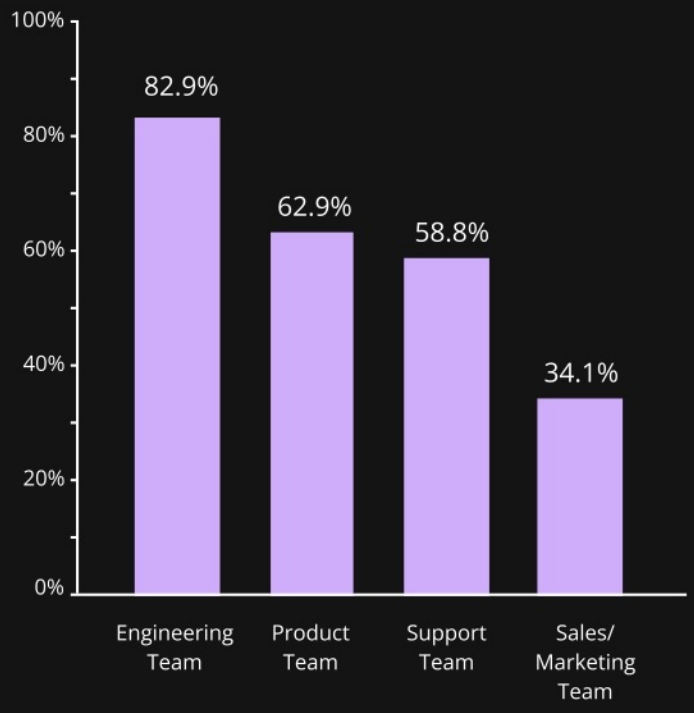

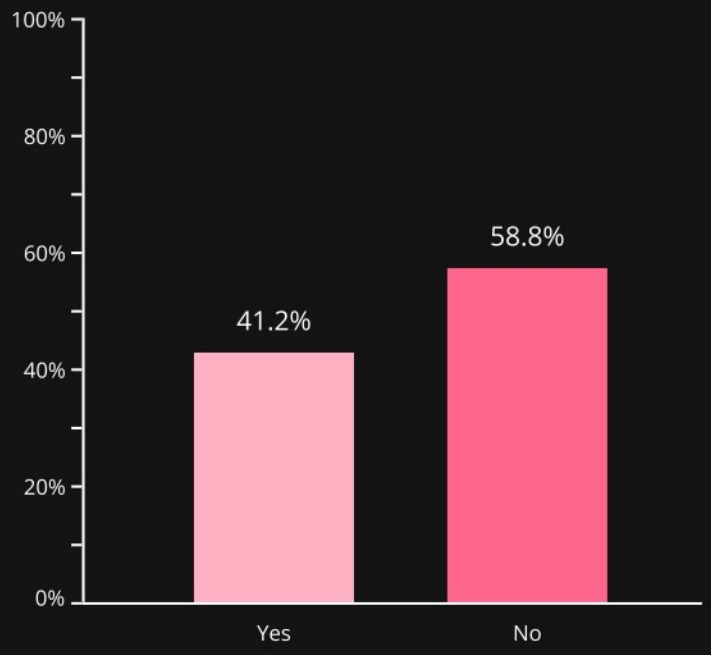

Question: Of the metrics that you report on, do you surface any insights to your customers?

Insight: This answer suggests the majority of vendors have an approach to data collection and sharing that is focused on their needs rather than their customers. Coupled with the high-level of engineering-focused metrics and usage, this suggests a sense that the data collected is not considered valuable to customers. Collecting more of the data that product and GTM teams clamor for also opens the door to providing significant value to customers with data.

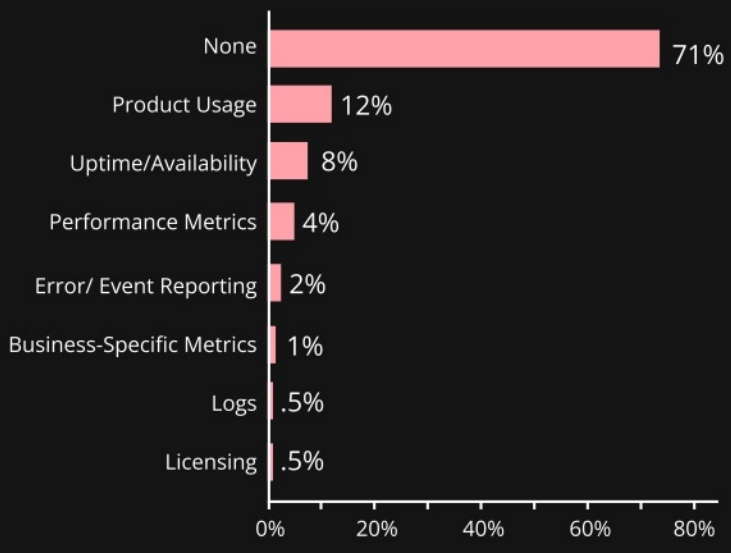

Question: If you do surface metrics to customers, which do you surface?

Insight: Over 70% of respondents marked "None"—suggesting that reporting is often inward-facing, with limited investment in surfacing metrics directly to end users. Among vendors that do surface data, product usage insights lead the way, indicating a desire to show value, adoption, or ROI to customers—especially in usage-based models.

Question: Do your customers need to opt-in in order for you to report on application instances running in their environments?

Insight: Vendors respect customers' sensitivity around data collection, with over 63% of vendors choosing some type of opt-in approach. They also show sophisticated thinking about privacy and security by treating air-gapped customers differently. Transparency about data collection can differentiate you competitively and increase trust with customers.

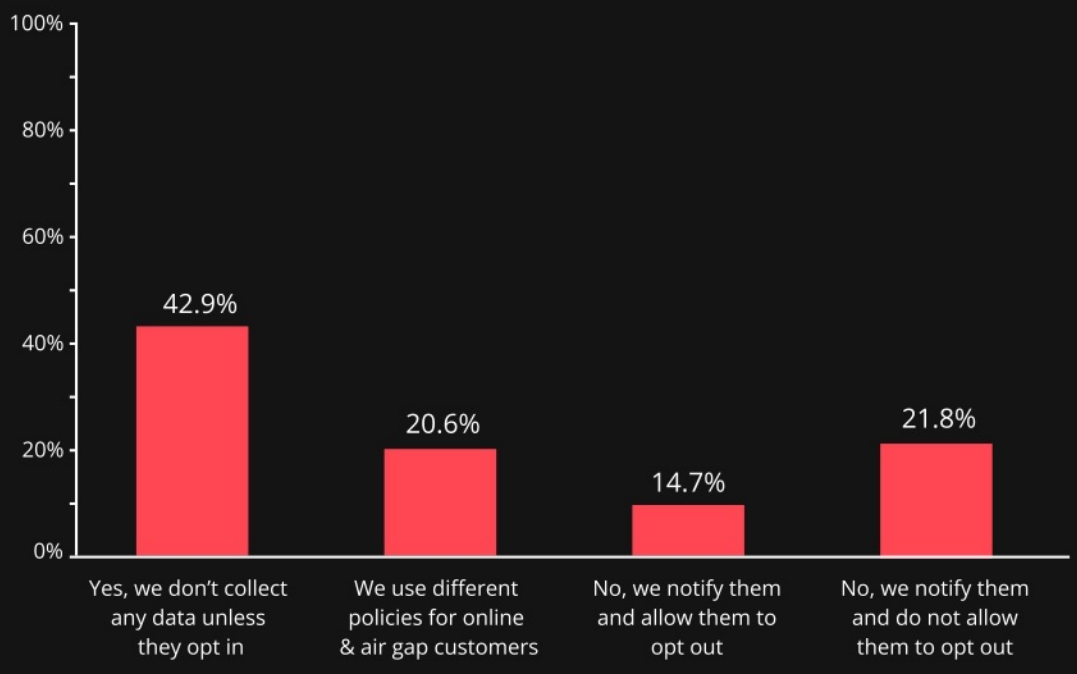

Section 8:

Support & Customer Experience

Support structures, severity handling, and customer satisfaction.

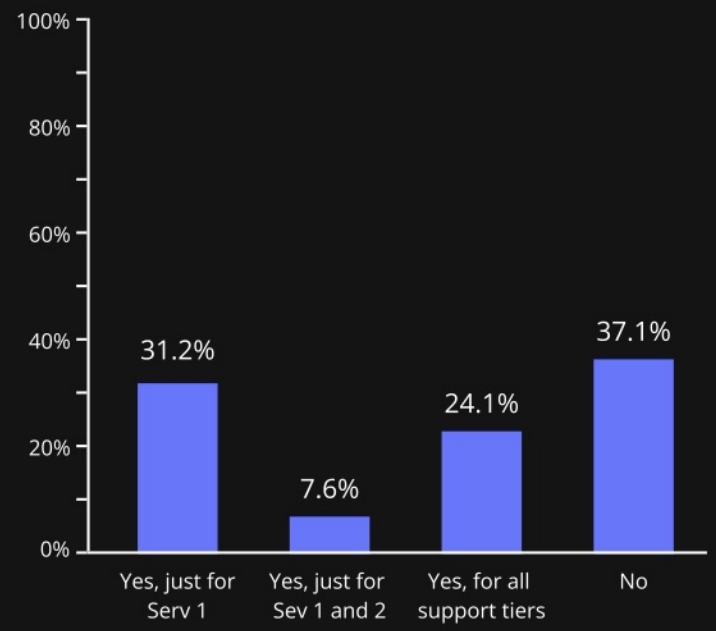

Question: Do you offer 24/7 support coverage?

Insight: Most vendors do not offer 24/7 support coverage, and those that do are more likely to offer a tiered approach where only the most severe issues receive 24/7 responses. This suggests companies are striking a sophisticated balance of cost considerations and customer needs.

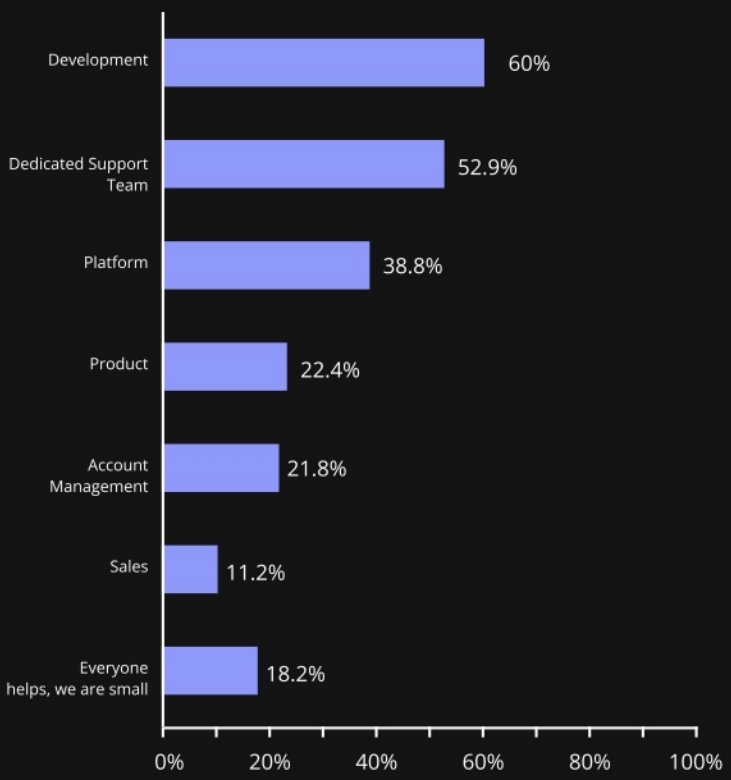

Question: Which of your internal teams support your self-hosted customers when issues arise?

Insight: A majority of organizations rely on technical teams to support self-hosted customers during issues, with 60% involving Development and 53% having a dedicated support team. This highlights a strong dependency on engineering resources, while customer-facing functions like Sales and Account Management play a more limited role in incident response.

Question: How does your support team gather the right information to resolve customer issues?

Insight: Support teams most commonly rely on ticket content (75.3%), log collection (70%), and Slack or email exchanges (69.4%) to gather information for resolving customer issues. More direct methods like calls and environment metadata are used less frequently, indicating a preference for asynchronous and documented communication channels.

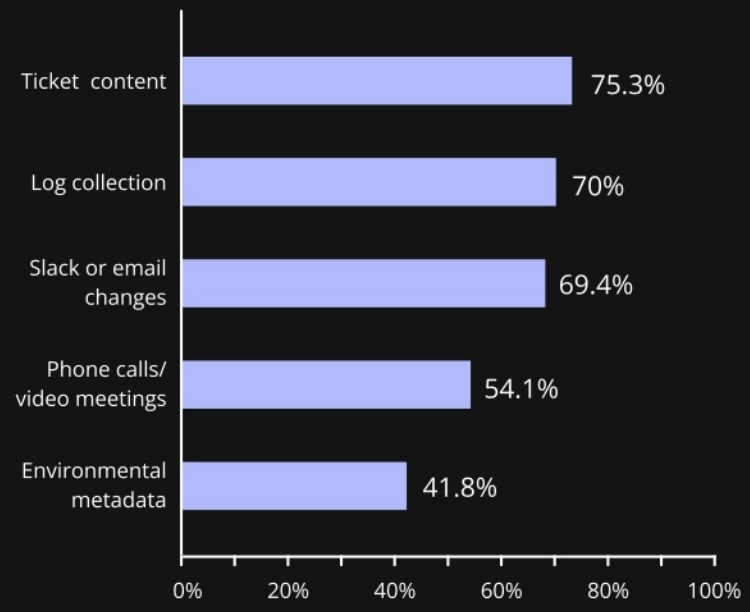

Question: How long is the median time it takes to resolve a Sev-1 issue?

Insight: Most organizations are able to resolve Sev-1 issues quickly, with nearly two-thirds (65.9%) reporting resolution within 24 hours. Only a small minority (7.1%) take more than 48 hours, suggesting strong prioritization and responsiveness for critical incidents.

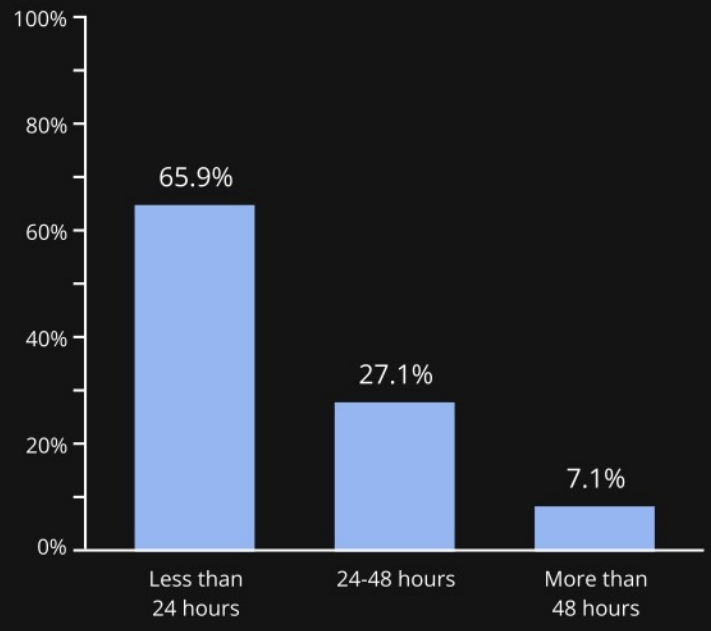

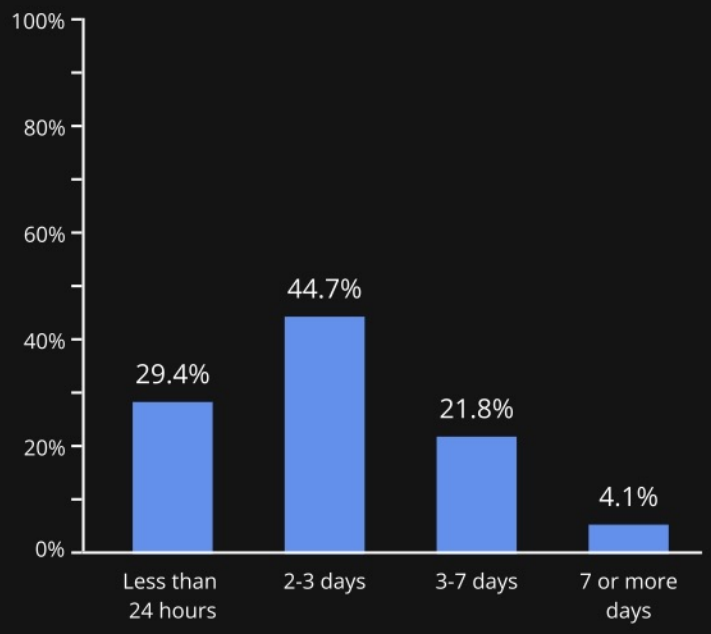

Question: How long is the median time it takes to resolve a Sev-2 issue?

Insight: The majority of organizations resolve Sev-2 issues in a 2–3 day window (44.7%), while nearly 30% are able to resolve them in under 24 hours, indicating some prioritization overlap with Sev-1. However, a meaningful portion (25.9%) require 3 or more days, reflecting the broader range of acceptable resolution times for less critical issues.

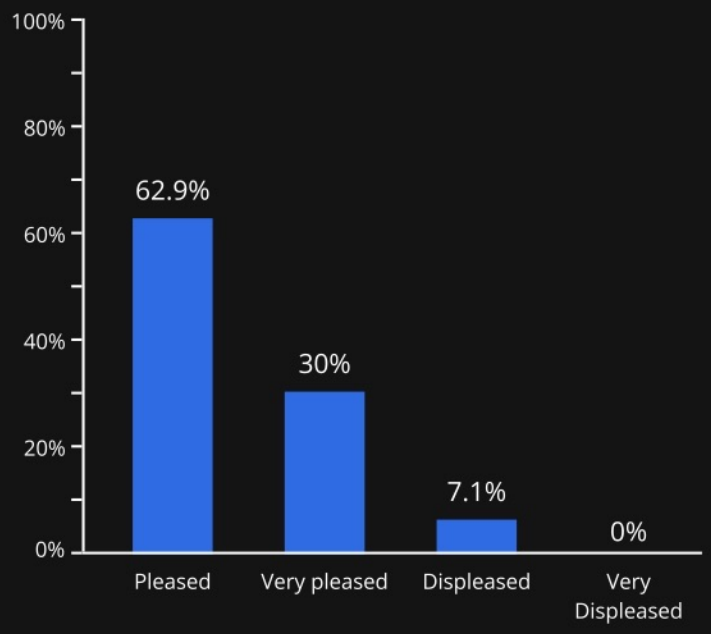

Question: How pleased are your customers with your response time on support issues?

Insight: Customer sentiment toward support response times is overwhelmingly positive, with nearly 93% reporting they are pleased or very pleased, and zero respondents indicating they are very displeased. This suggests that current response practices are effectively meeting or exceeding expectations for the vast majority of customers.

Question: How pleased are your customers with your mean time to resolution (MTTR) on support issues?

Insight: Customer satisfaction with mean time to resolution (MTTR) is strong overall, with nearly 89% expressing they are pleased or very pleased. While slightly lower than satisfaction with initial response time, the absence of "very displeased" responses signals that resolution timelines remain broadly acceptable.

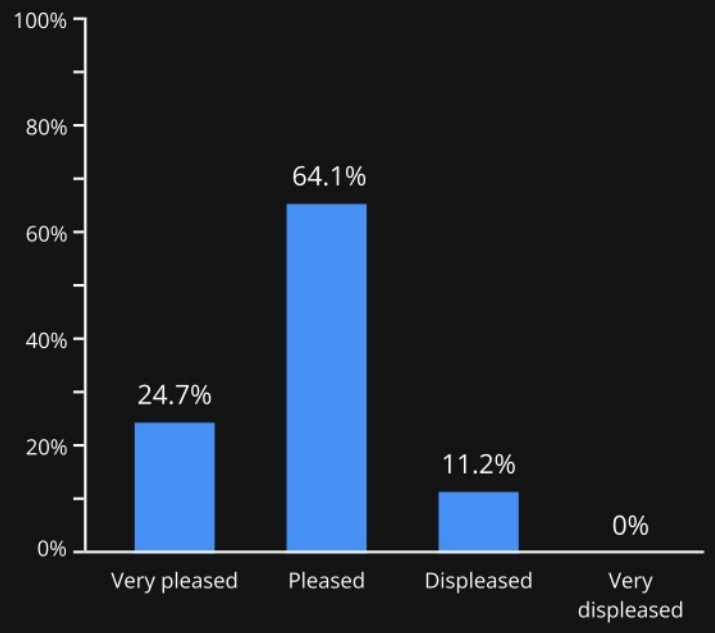

Question: How big is your organization?

Insight: Most respondents come from small to mid-sized organizations, with 68.3% reporting under 250 employees. Only 9.4% represent companies with over 1,000 employees, suggesting that insights from this data are most reflective of support practices in leaner or scaling teams.

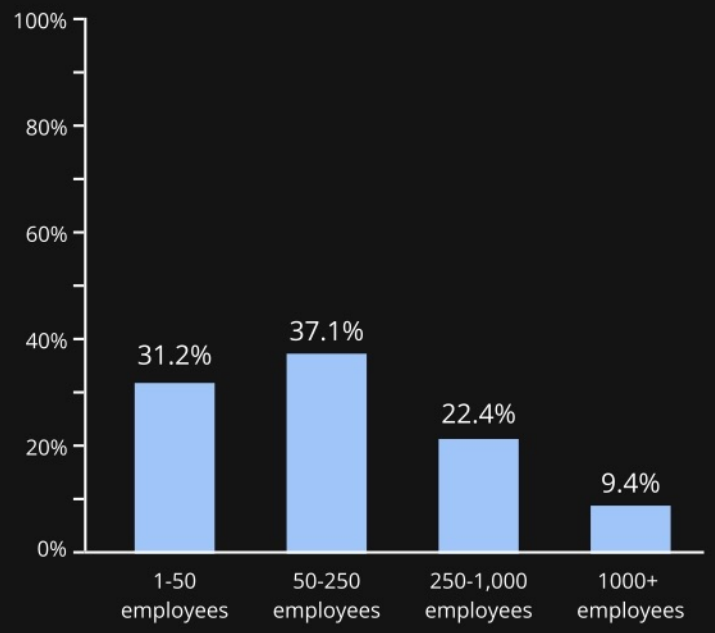

Section 9:

AI in Deployment & Development

Current and potential use of AI in self-hosted workflows.

REPLICATED

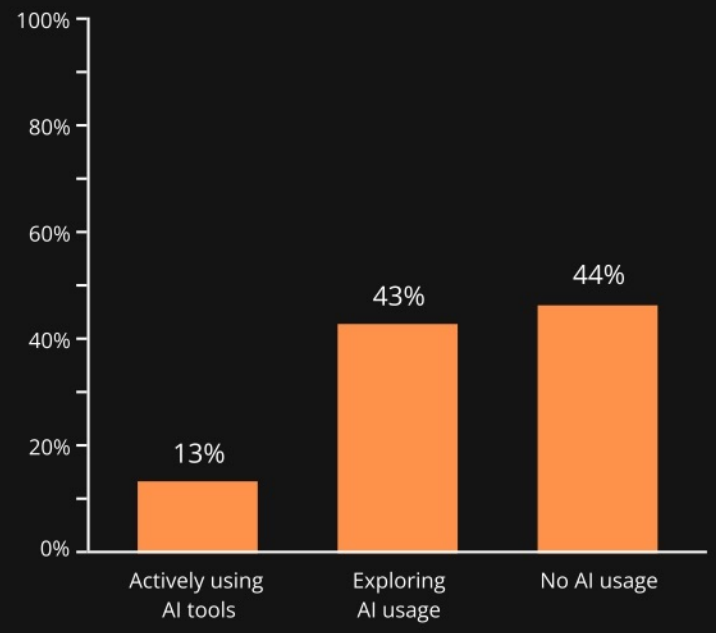

Question: Do you use any AI tools in your self-hosted deployment and development process?

Insight: While 13% of respondents are actively using AI tools in their daily jobs, a much larger group — 42% — are in exploration or evaluation mode. This indicates high interest, but also hesitation or uncertainty about how and where to apply AI productively.

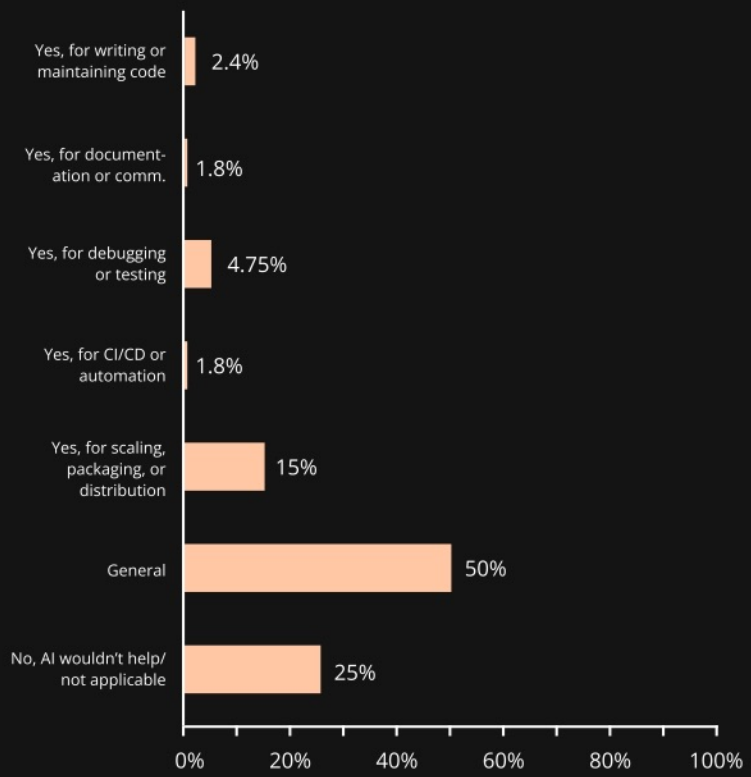

Question: Are there any tasks involved with building and distributing your application(s) where you are not using AI tools today, but think AI tools could be helpful?

Insight: Most teams are still exploring where AI fits into their workflows, but early signals point to strong potential in deployment and delivery automation. 15% of respondents mentioned scaling, packaging, distribution, or deployment tasks as promising areas.

Conclusion

The 2025 Software Deployment Survey reveals a market in transition — one where vendors are balancing traditional self-hosted models with modern tooling, stronger automation, and growing support for enterprise-controlled environments. Kubernetes, containerization, and Helm have become table stakes, while investments in testing, licensing, and customer support show a clear shift toward enterprise maturity.

Looking ahead, the opportunity is clear: teams that streamline installation, surface meaningful insights, and embrace automation — including AI-powered workflows — will be best positioned to meet enterprise expectations at scale. While challenges remain, the path forward is taking shape: smarter infrastructure, faster delivery, and better experiences for customers and vendors alike.