

# Annual Linux market report for 2021



## Executive summary

### **The use cases that organizations have for Linux are changing as they shift their infrastructures to cloud environments.**

The last two years have seen an accelerated change with IT infrastructures, as organizations have launched digital transformation projects around new customer services and core operational services—to keep up with shifting workplaces and expectations.

Across both public cloud environments and physical infrastructures, Linux® has been at the center of digital change and has recognized that organizations adapt which Linux distribution they will use according to their specific project needs. Organizations on average are running more than two Linux distributions in their environments, and they shift between distributions based on the operating environment, the production stage, and the workloads types.

In 2021 Red Hat commissioned [Management Insight Technologies](#) to conduct a global survey of enterprises that deploy Linux® in their IT infrastructure. Management Insight Technologies conducted a series of phone interviews and web-based surveys, where respondents were asked a variety of questions about how their organizations were using Linux in the infrastructure and what drivers and considerations were used to select Linux distributions. These surveys were conducted in February and March 2021.

Respondents show similar flexibility in how they adopt other technologies, such as moving to a public cloud environment, using containers, and deploying devices at the network edge. New technologies are adopted fairly evenly and in a measured way with comparable growth in cloud and containers and steady incremental adoption for edge and Function as a Service (FaaS).

The major inflection point for organizations to change to a Linux distribution is the shift to a cloud environment. This could mean they are migrating from a different operating system (OS) or they are selecting a specific distribution for a specific reason, but the shift in environment provides an opportunity for them to choose a different infrastructure.

This annual report summarizes and shares some key findings about the current usage and future trends pertaining to enterprise Linux use. This report also covers four major areas for how organizations evaluate and use Linux within their infrastructures..

- ▶ General Linux deployments (including production stages and workloads)
- ▶ Security
- ▶ Public cloud
- ▶ Containers

More than 1,200 IT professionals in a variety of different roles and industries took part in this study. This research explores their current and planned Linux adoption, the patterns in adopting related technologies like containers, and the perspectives on major factors which influence these decisions.

## Highlights

### **Linux is the foundation of cloud infrastructures**

- ▶ On average, organizations are using 2.03 Linux distributions in their environment and they consider 3.05 distributions before making a decision.

- ▶ The top Linux distributions in use are Red Hat® Enterprise Linux (40%) and Ubuntu (40%).
- ▶ Security, reliability, and ease of management and maintenance are the top influences for choosing a Linux distribution.
- ▶ 58% of organizations have a hybrid cloud strategy, while another 25% have a multicloud strategy.

### **Adaptability and flexibility matter more than ease-of-access for applications**

- ▶ Only 19% of Linux applications are commercial or off-the-shelf software.
- ▶ 32% of applications are developed in-house.
- ▶ The release stage of an application has a huge influence on when organizations choose enterprise versions of Linux over community, with 78% of production workloads running on Red Hat Enterprise Linux in datacenters and 75% in cloud environments, but only 36% of test environments.

### **Standardization is the top way to manage complex environments**

- ▶ 52% of organizations use a single OS as part of their cloud management strategy.
- ▶ 43% of organizations use a single container image as part of their cloud strategy.
- ▶ 17% of organizations avoid using cloud provider tools to maintain a consistent management approach.

### **Portability is one of the major factors in how organizations use cloud computing**

- ▶ 58% of organizations migrated an application from a physical environment to a cloud service.
- ▶ 31% of organizations move workloads between public cloud environments.
- ▶ Over a quarter of organizations either repatriated a workload from cloud computing to on-premise (27%) or moved a cloud-native application to on-premise (30%).

### **Small-scale container deployments tend to mimic virtual machine deployments**

- ▶ 38% of respondents said that they were using containers similar to virtual machines (VMs), while another 38% said that using containers was easier for application migration than re-architecting.
- ▶ 55% of containers are deployed on-premise.
- ▶ 56% of containers host monolithic applications rather than microservices.
- ▶ 39% of respondents said that the main priority when adopting containers was developer productivity.

### **Methodology and firmographics**

1,204 IT professionals, across all major industries and from nine countries across all major geographic areas, took part in this research survey. Management Insight Technologies selected participants based on technology experience and knowledge, and all respondents had Linux experience (though only 484 were Red Hat Enterprise Linux customers) and some level of public cloud adoption.

Survey respondent firmographics

### Region



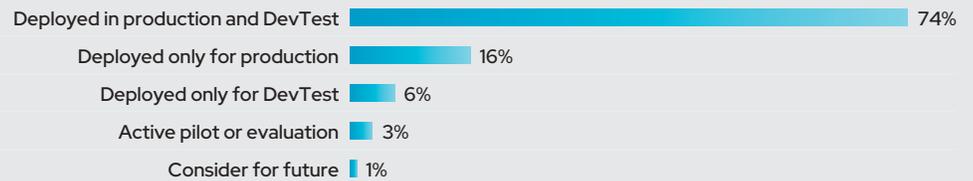
### Respondent type

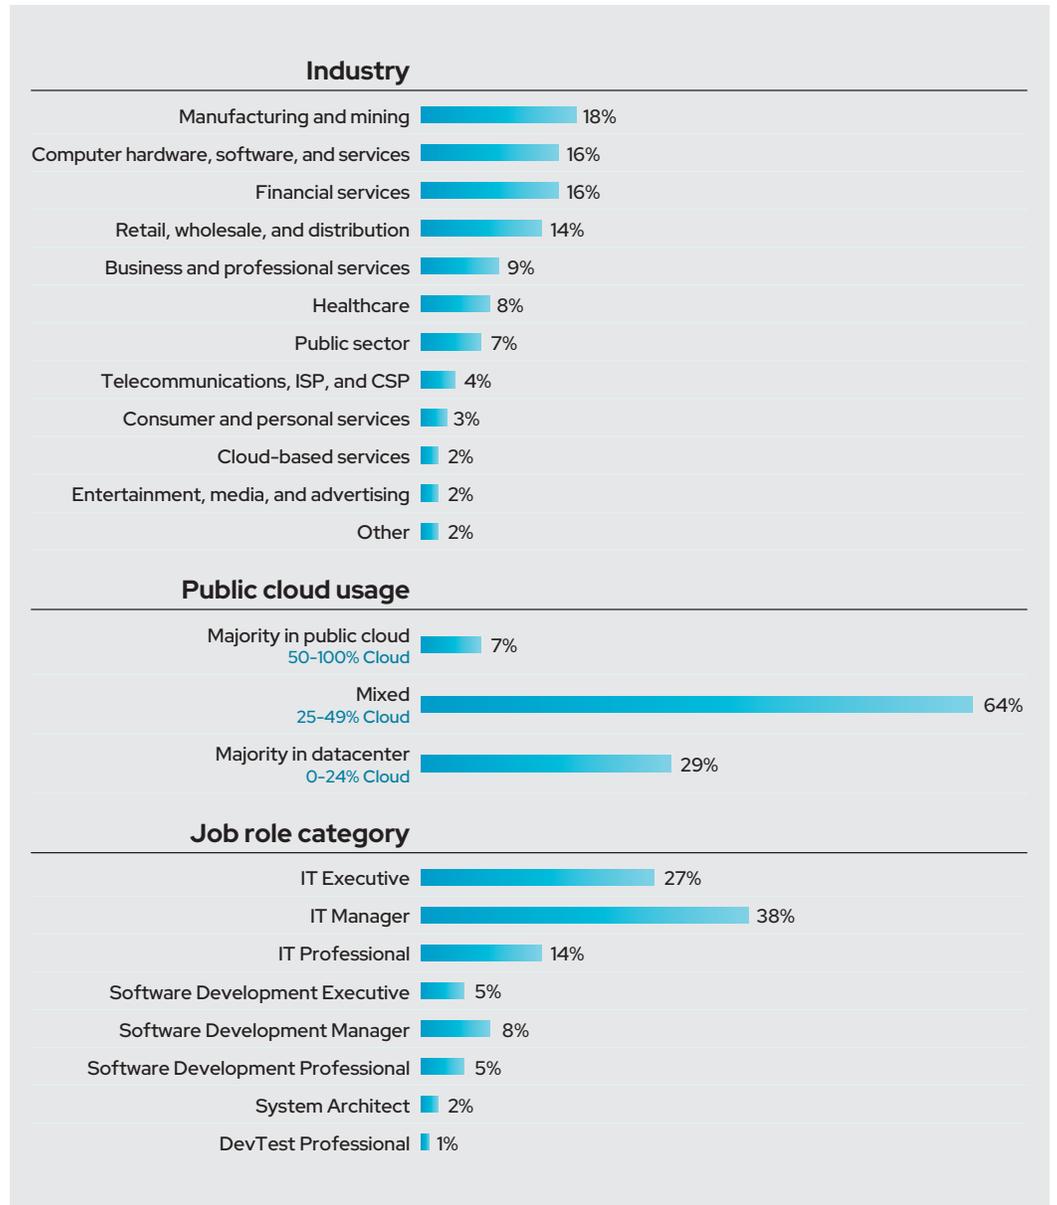


### Company size



### Linux OS status





## Selecting the right Linux distribution

### Organizations have a lot of options and they like to evaluate a mix of enterprise and community distributions.

There are a lot of different Linux distributions, and respondents consider a wide variety of options, but what they decide to use depends on the deployment environment and production stage.

Red Hat Enterprise Linux (57%) and Ubuntu (52%) are clear leaders with the majority of respondents considering them, but there is a pattern of evaluating other distributions because of the associated cloud service or the workloads which will run on it, such as Oracle Linux (26%) and Amazon Linux (20%).

An average of 3.05 distributions have been considered by organizations.

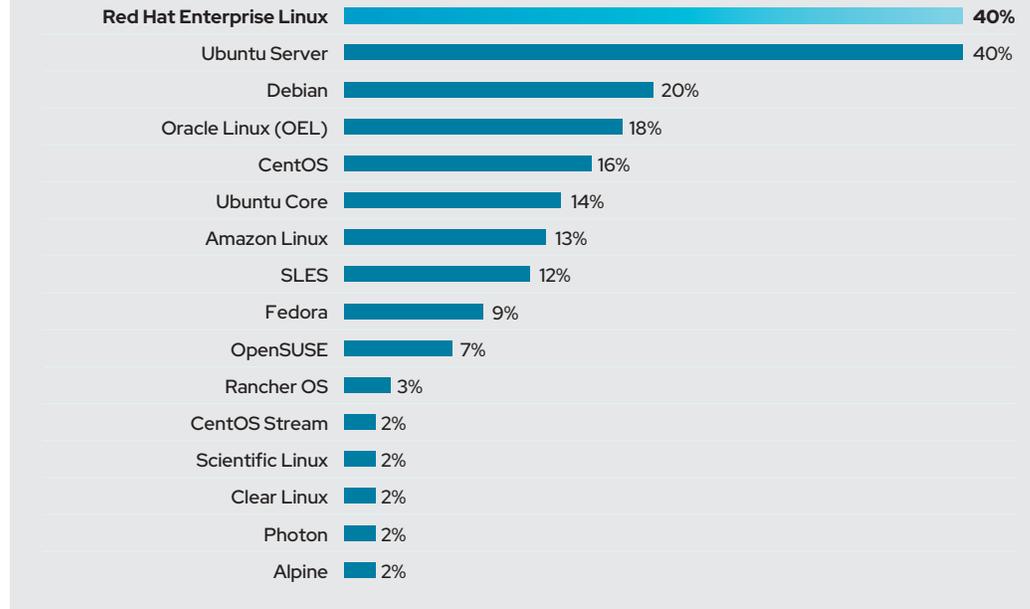
Figure 1. Linux distributions that organizations have considered using



The gap narrows once organizations choose what to deploy in their infrastructure. Red Hat Enterprise Linux and Ubuntu are the top choices (40% each). However, only an average of 2.03 distributions are deployed in infrastructures. This mirrored the respondents' preference to [standardize on a single OS](#) to manage their cloud environments.

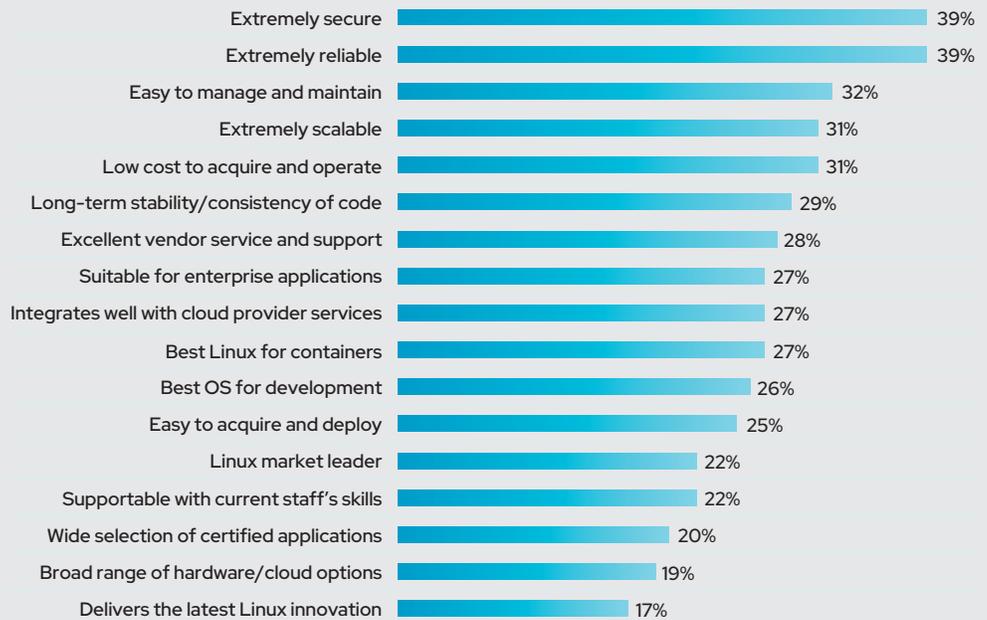
It should be noted that these percentages show whether an organization has a Linux server of that type in their infrastructure. It does not show the number of servers by deployment type.

Figure 2. Linux distributions in use



The choice of Linux distribution aligns tightly with usage, so there isn't one factor that dominates why respondents choose a distribution over another. The top selected factors include security and reliability, both at 39%. Despite the emphasis we hear from IT leaders that skills gaps are one of their top concerns, only 22% of respondents considered their current team skills an important factor in their Linux selection. The least considered factor selected by only about a sixth of respondents was technology innovation (not listed).

Figure 3. Most important factors in selecting a Linux distribution

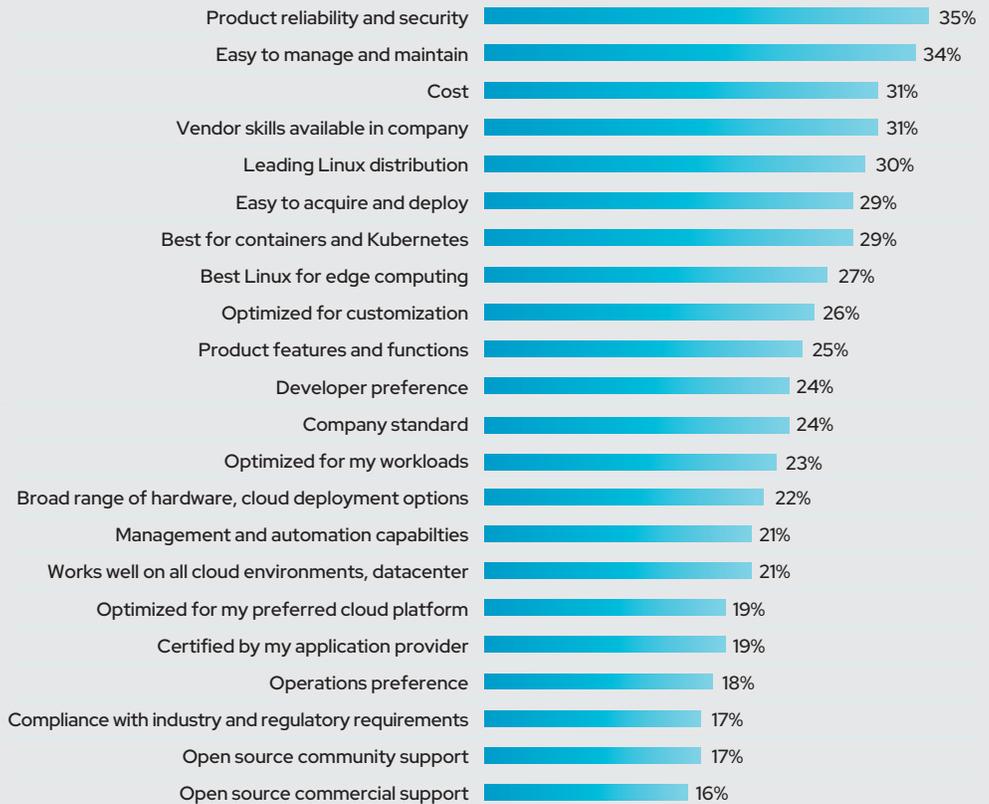


N=1204

**The most important factors for a Linux choice vary depending on how it will be used.**

There is a slight shift in priority for organizations which are heavily using cloud environments (defined as deploying more than 50% of their Linux instances on the public cloud). 39% of those respondents rate “integrating with cloud services” as the most important factor for selecting a vendor, as opposed to only 27% overall, and there is a similar emphasis on containers (34% versus 27%). Likewise, ease of maintenance and scalability are both significantly more important: 41% and 40% for organizations heavily using cloud environments versus 32% and 31% overall. Organizations that were heavily using cloud environments rated code stability and support as important significantly less often than the overall respondents.

Figure 4. Most important factors for the most recent Linux decision



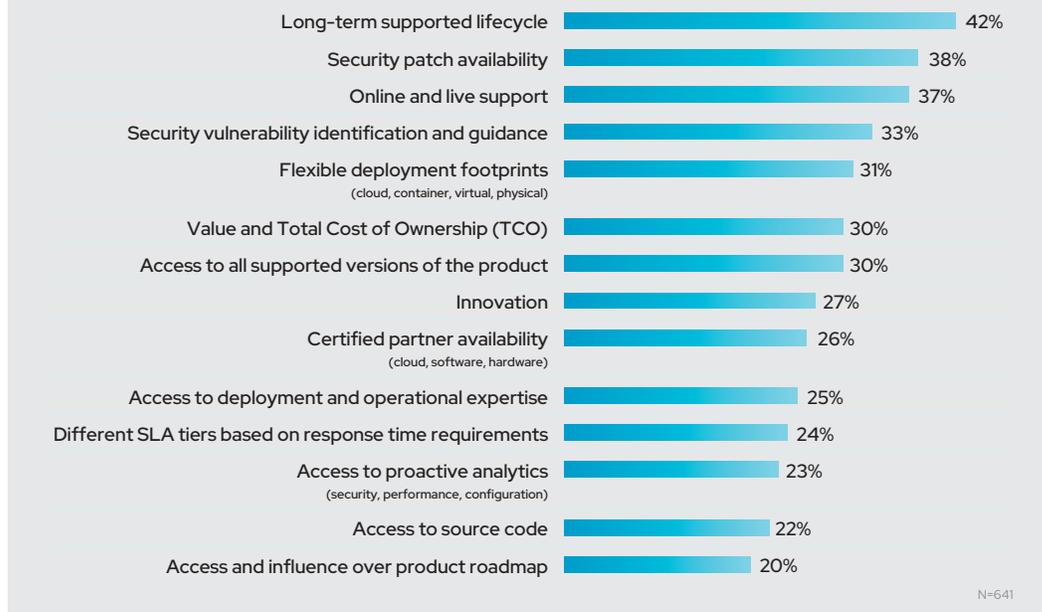
The relationship between stated attribute importance and provided preference was analyzed to determine the statistical drivers of Linux preference. With the results of this statistical analysis, some attributes strongly lean towards preference, although they may not be top of the list in stated importance.

Being a Linux market leader was the top statistical driver of preference for a Linux distribution, but respondents ranked it 13 out of 17 attributes in stated importance. This list includes top statistical drivers of preference in order of stated importance.

- ▶ Market leader
- ▶ Easy to manage and maintain
- ▶ Suitable for enterprise applications (otherwise ranked 8 out of 17)
- ▶ Best OS for development (ranked 11 out of 17 attributes)
- ▶ Extremely secure and extremely reliable (tied)

The factors that go into using a paid or enterprise Linux distribution differ from evaluating Linux. When there is a vendor relationship, the top factors become the overall OS life cycle (42%), the availability of security patches (38%), and the quality of support (37%).

Figure 5. Most important selection criteria for enterprise Linux distributions



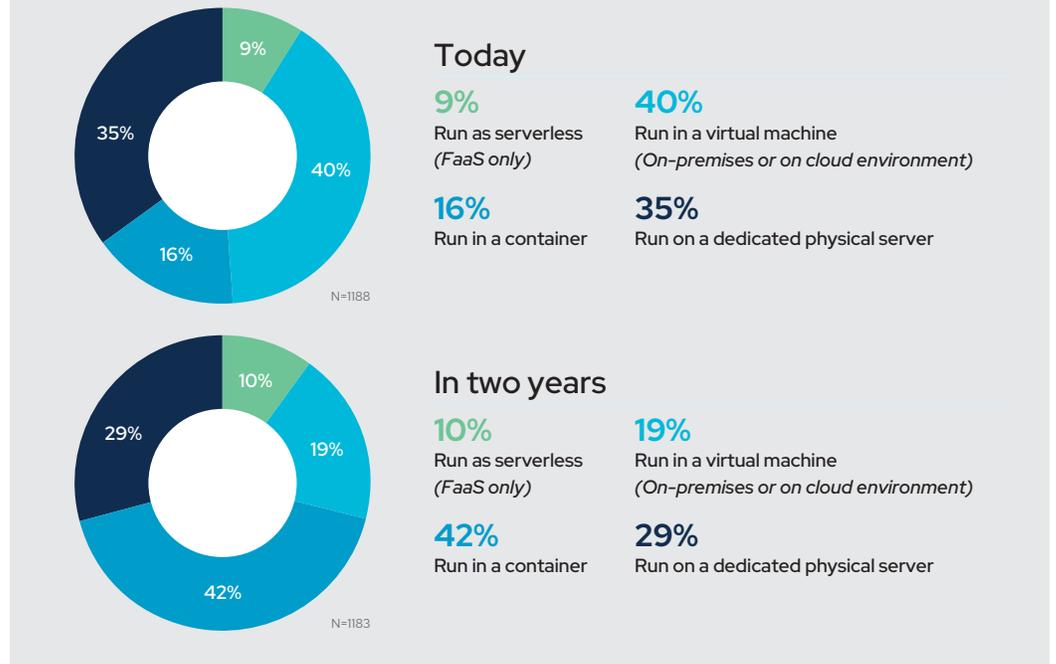
### How Linux workloads are deployed

There are several key aspects to workloads running on Linux systems: the type of deployment environment, the production stage of the environment, and the type of workload. All of those environmental factors shape how organizations select a Linux distribution and how they maintain that system.

#### The way that Linux workloads are deployed is shifting—slowly.

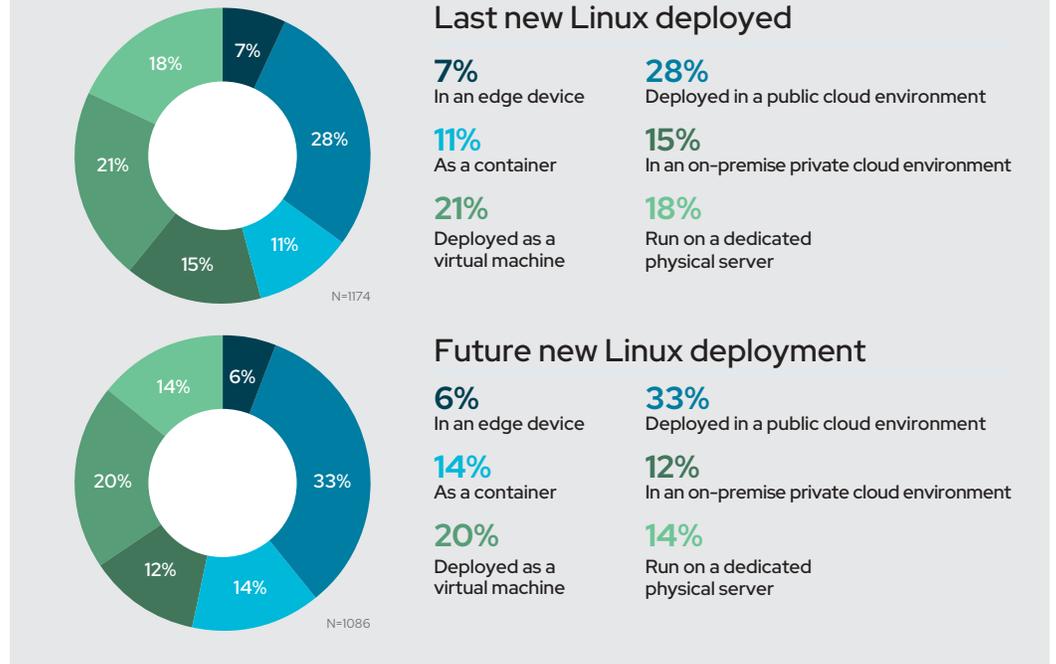
Although the overall change is minor, respondents are expecting a small drop in their physical systems, from 35% to 29% of systems over the next two years. Container environments are expected to grow the fastest, rising from 16% of systems to 19% of systems. Linux workloads running on virtual machines remain the most common deployment type, increasing slightly to 42%.

Figure 6. Current and expected allocation of Linux workload types



The shift toward using a public cloud environment seems more pronounced when looking more closely at the most recent deployment and the next planned deployment. Physical systems, private cloud environments, and virtual machine deployments are expected to drop by a cumulative 8%, while public cloud deployments are expected to increase to 33% (+5%) and container deployments to 14% (+3%).

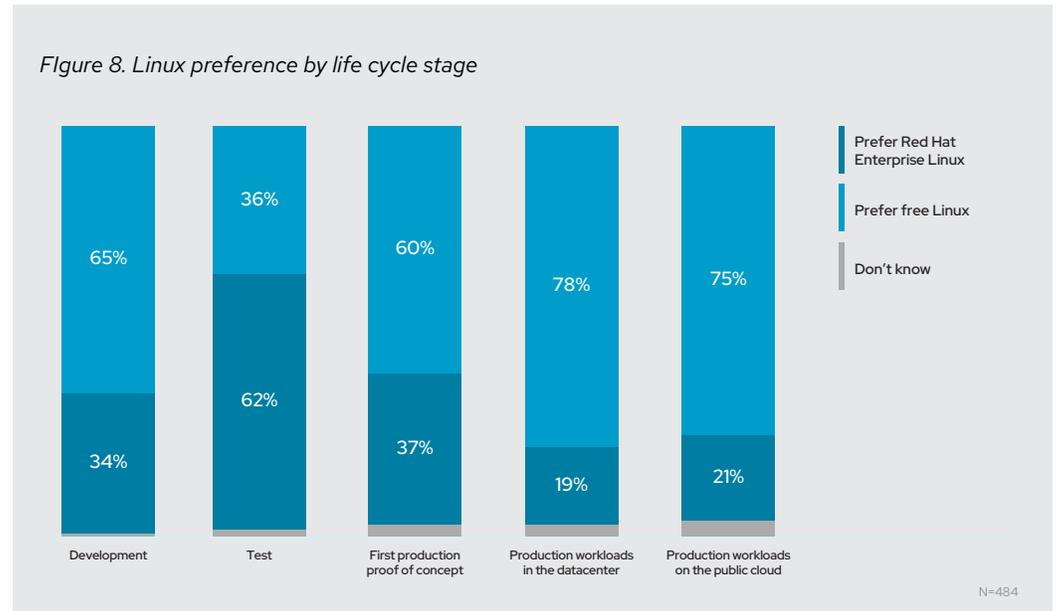
Figure 7. Last and next Linux deployment, by type



The production stage is a strong influence on whether organizations choose Red Hat Enterprise Linux or community versions of Linux.

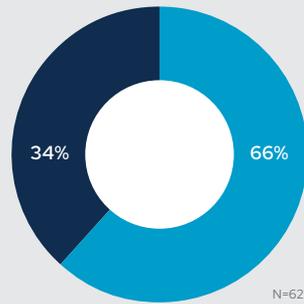
There is a strong preference for Red Hat Enterprise Linux for most environments, including development and proof of concept environments. The outlier is testing environments, which lean heavily toward free versions of Linux.

(Select one per row)



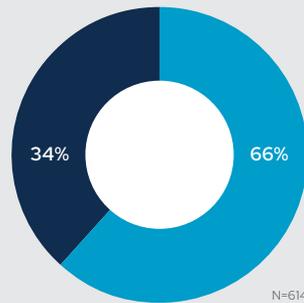
This seems to have more to do with the production stage of the workload, rather than directly with the environment, since the percentages of enterprise and community versions of Linux deployed in datacenters and public cloud are the same.

Figure 9. Percentages of enterprise and community Linux distributions in datacenters and public cloud



#### Linux workloads in datacenter

**66%** Paid **34%** Free

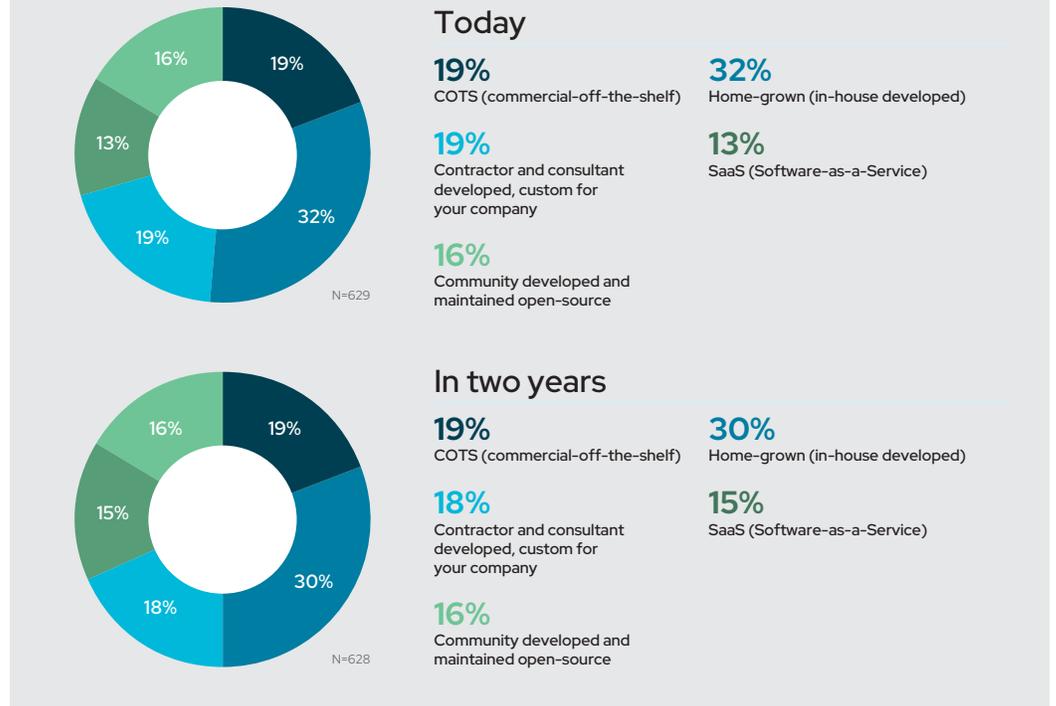


#### Linux workloads in a public cloud

**66%** Paid **34%** Free

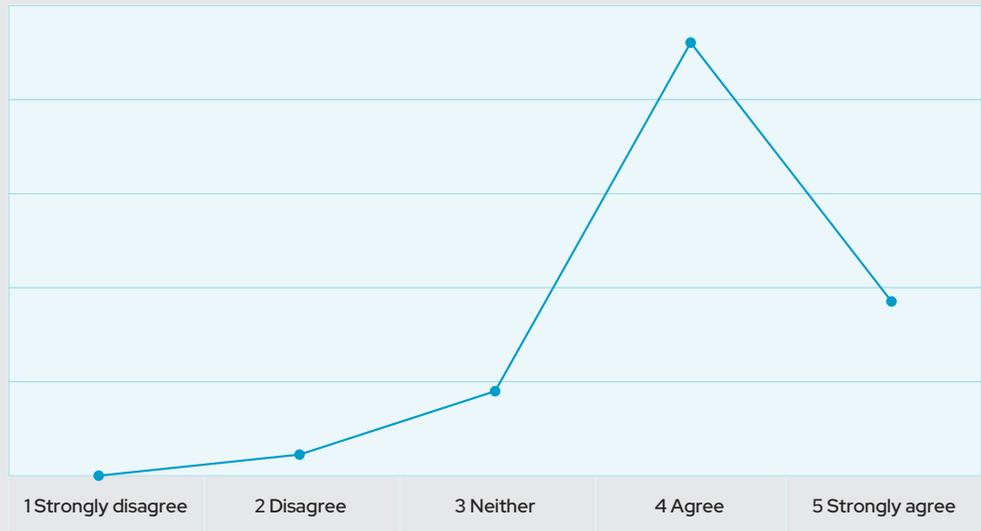
Respondents are expecting a minor change in what types of applications they are deploying or using within their infrastructures over the next two years. Right now, 51% of applications are developed within the organization or developed by a contractor and only about 19% of applications are commercial, off-the-shelf-software.

Figure 10. Breakdown of Linux workloads by type, current and projected



Even with commercial software, organizations don't just use the software as it is out-of-the-box. Respondents agreed that even commercial-off-the-shelf (COTS) software requires customization to be used in their infrastructures, with an average rating of 4.04 out of 5.

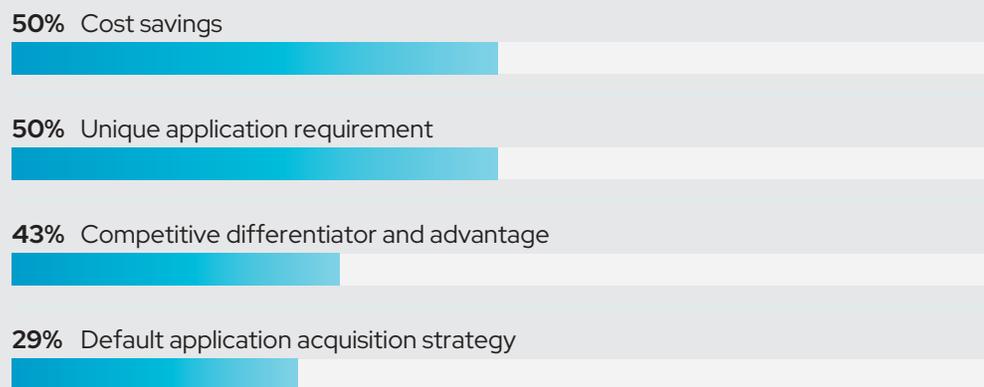
Figure 11. Rating for whether COTS requires customization



N=531

There was high agreement among the respondents that COTS requires customization, and a lot of organizations prefer to develop their own applications, either directly or through a contractor. The primary reasons relate to cost or because they have unique requirements not available in a COTS option (50% for both). Only 43% develop their own applications explicitly because it is a competitive advantage.

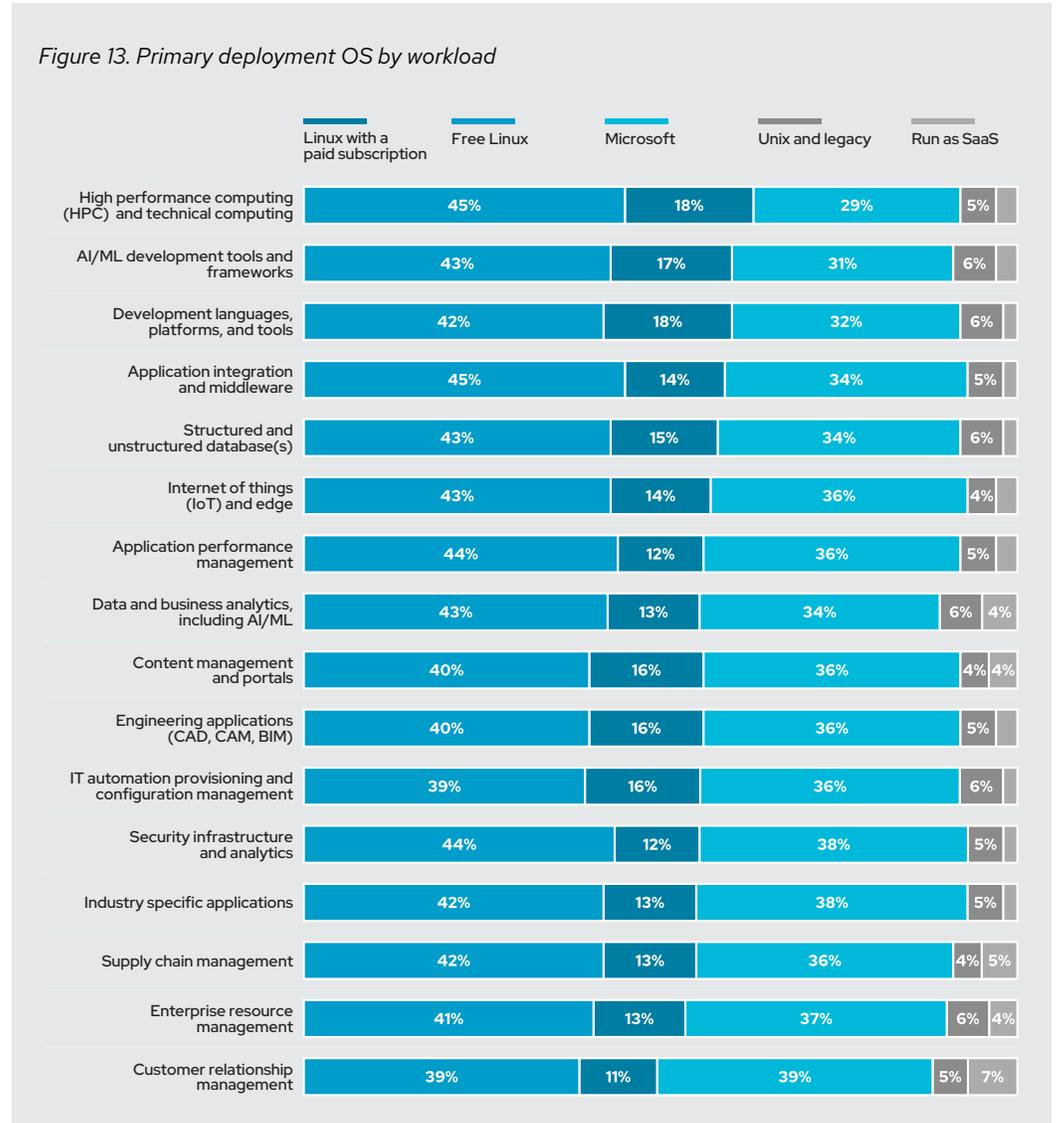
Figure 12. Reasons for developing software in-house or through a contractor



N=803

Enterprise Linux is generally the preferred OS for key business-critical applications, especially ones that benefit from a powerful performance or from standard protocols. The top workloads are high performance computing (45%), integration (45%), application performance (44%), and security analytics (44%).

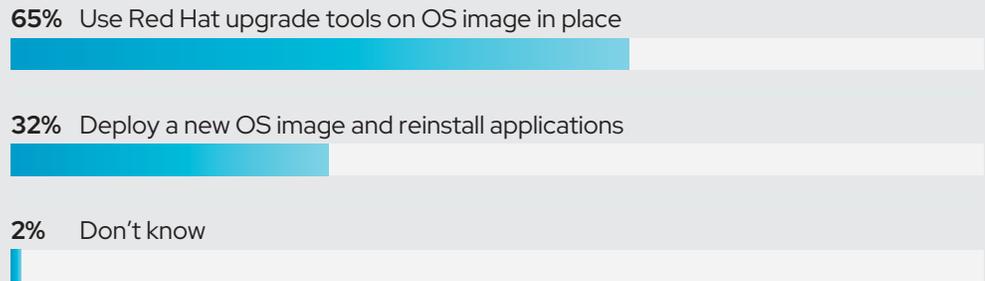
Figure 13. Primary deployment OS by workload



**Red Hat Enterprise Linux instances are stable, with administrators choosing to migrate rather than redeploy for upgrades.**

Overall, respondents prefer to upgrade systems rather than redeploy and this pattern holds true both for users, which were heavily using cloud environments and primarily in the datacenter.

Figure 14. Approaches to updating a Red Hat Enterprise Linux instance



N=484

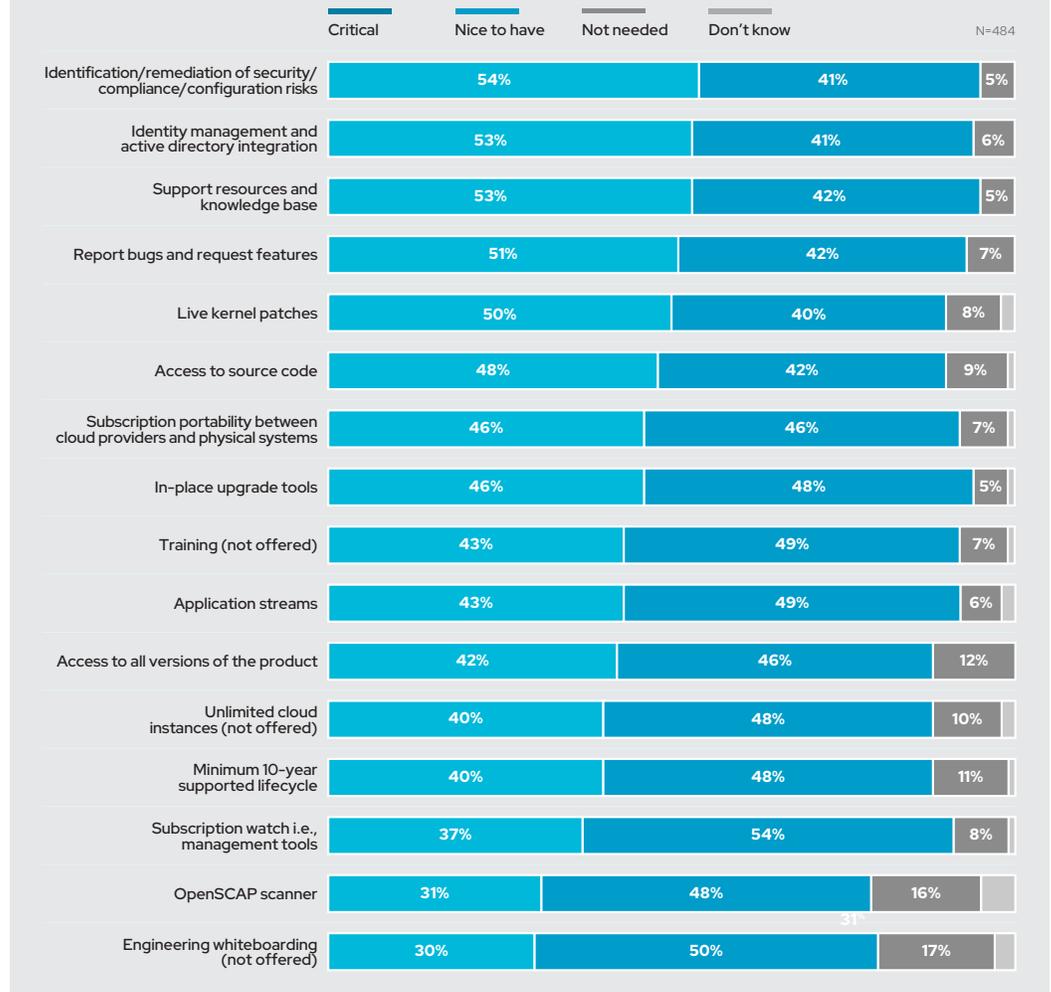
## Security

Features related to security are highly valued by customers.

When asking Red Hat Enterprise Linux customers, specifically, they rated security related features as the most critical benefits they get from their subscription, with three of the top 5 features related to security.

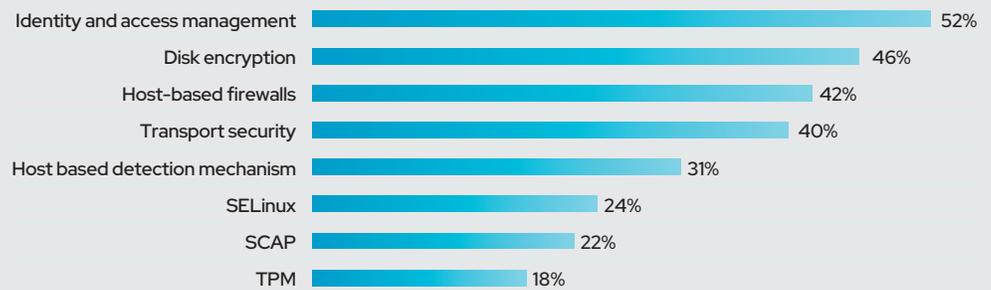
- ▶ Proactive analytics / identifying configuration, compliance, and security risks (Red Hat Insights) (54%)
- ▶ Identity management and Active Directory integration (53%)
- ▶ Live kernel patching (50%)

Figure 15. Ranking the most important Red Hat Enterprise Linux subscription features



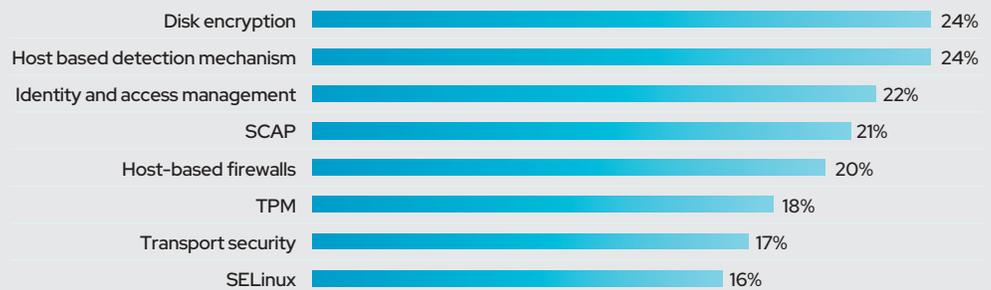
Looking more deeply at security features that are in use within infrastructures, over half of Red Hat Enterprise Linux customers (52%) are using their identity management application. Disk encryption is the second most used security technology (46%), while the other three features in the top 5 all relate to network and access security.

Figure 16. Red Hat Enterprise Linux Security features in use by customers



Those respondents who are not using a particular security feature in Red Hat Enterprise Linux were asked if they addressed that security area with another tool, and the answer seems to be not really. No security area stood out for third-party tools.

Figure 17. Third-party security tools used by customers

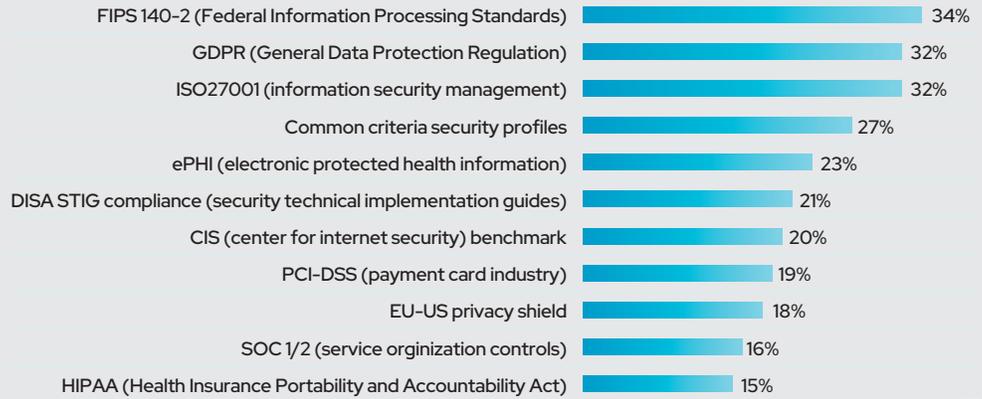


**While organizations have to deal with a variety of security regulations, government-defined standards are more common than industry ones.**

Only a small percentage of respondents said that complying with regulations was the most important factor in their Linux deployment requirements. By far, those respondents have to manage compliance with government regulations over industry regulations. The two top compliance standards are Federal Information Processing Standards (FIPS) 140-2 (34%) and General Data Protection Regulation (GDPR) at 32%. By comparison, the next industry regulation was electronic protected health information (ePHI) at only 23% and Payment Card Industry Data Security Standard (PCI-DSS) at 19%.

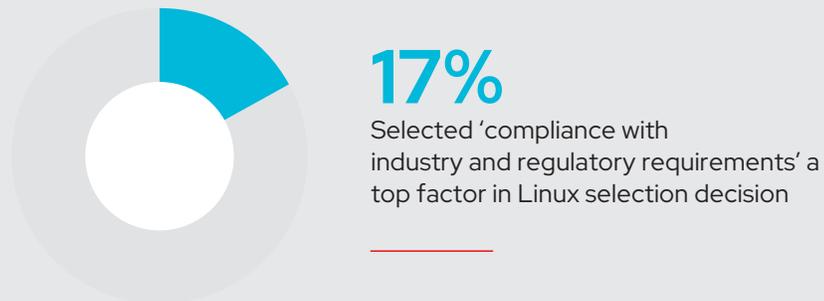
On average, respondents had to comply with 2.59 government or industry regulations on their most recent Linux project.

Figure 18. Top industry or government security regulations to comply with for the last deployment



Overall, while compliance is a necessary part of their Linux deployments, only one in six respondents said that regulatory compliance was a major factor in selecting a Linux distribution.

Figure 19. Percentage of respondents who cited regulatory compliance as the most important factor in their most recent Linux decision



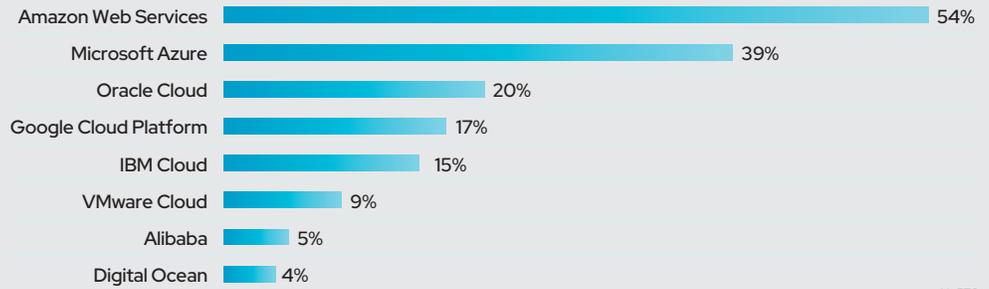
## Public cloud

### Linux distributions cross all types of environments.

Respondents are using a consistent distribution of cloud and on-premise environments that are seen across the industry. The top cloud providers are Amazon (used by 54% of respondents), Azure, and Oracle, with an average of 1.63 public clouds per organization.

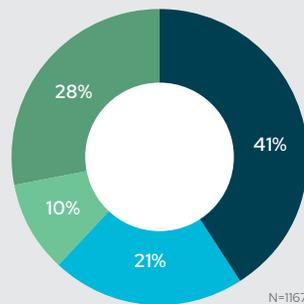
When breaking down the overall environments, Linux systems are deployed across on-premise and hosted physical systems (51% combined), with another 21% deployed on-premise in a private cloud. Only 28% of Linux workloads for this study were deployed on a public cloud.

Figure 20. Public cloud providers used for Linux workloads



N=570

Figure 21. Linux systems by infrastructure environment



N=1167

### Linux instances

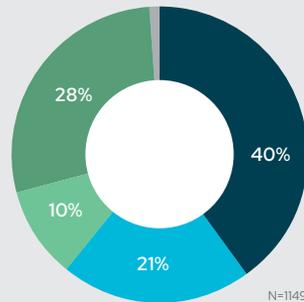
**41%** On-premise, traditional

**21%** On-premise, cloud (e.g., private cloud)

**10%** Off-premise, traditional (e.g., hosted)

**28%** Off-premise, cloud (e.g., public cloud)

**0%** Others including, edge



N=1149

### Windows instances

**40%** On-premise, traditional

**21%** On-premise, cloud (e.g., private cloud)

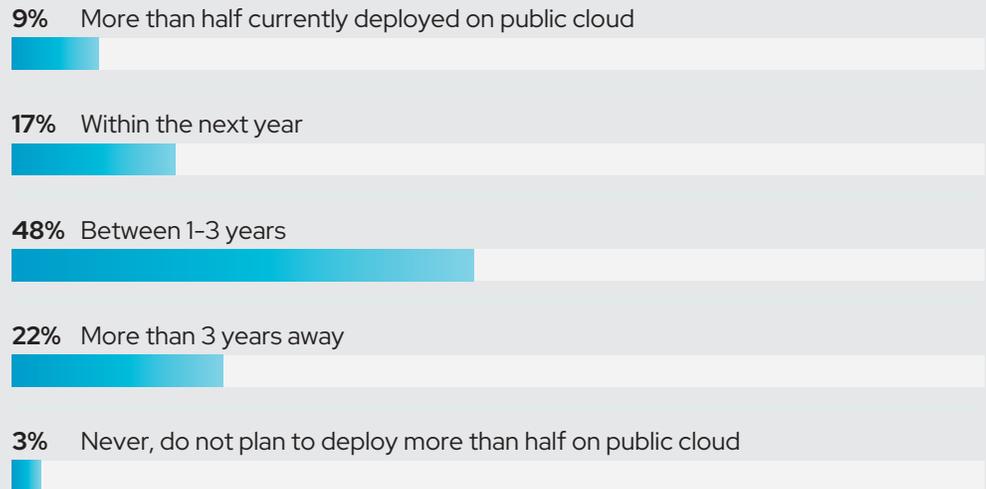
**10%** Off-premise, traditional (e.g., hosted)

**28%** Off-premise, cloud (e.g., public cloud)

**1%** Others including, edge

A full 70% of respondents expect it will take at least one year to deploy over half of their workloads on a public cloud; only 9% already have over half of their workloads on a public cloud.

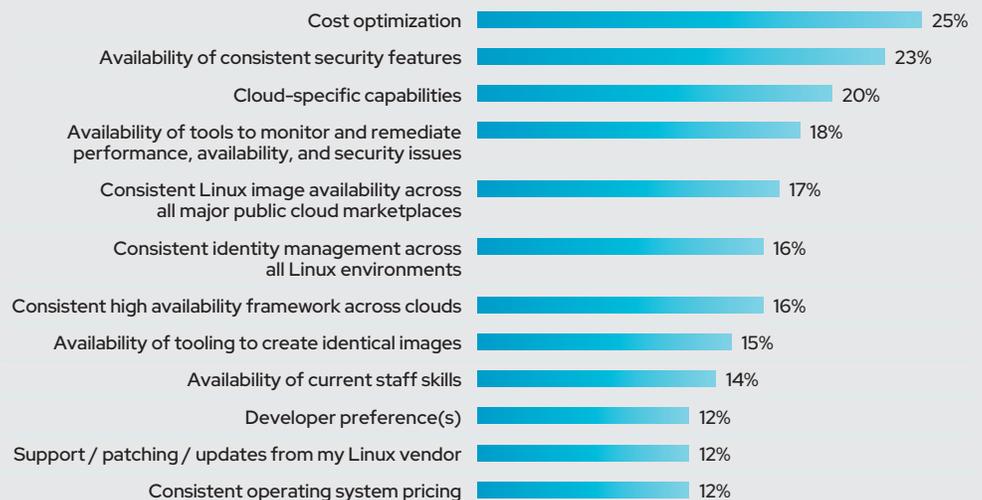
Figure 22. Timeline to deploy 50% of Linux workloads on public cloud



N=570

As part of selecting a Linux distribution for a project, developers are more influential than other roles (27%) in selecting a community Linux distribution. However, when deploying Linux in multiple cloud environments, developer preference was minimal (ranked 10 out of 12 factors). When there are multiple

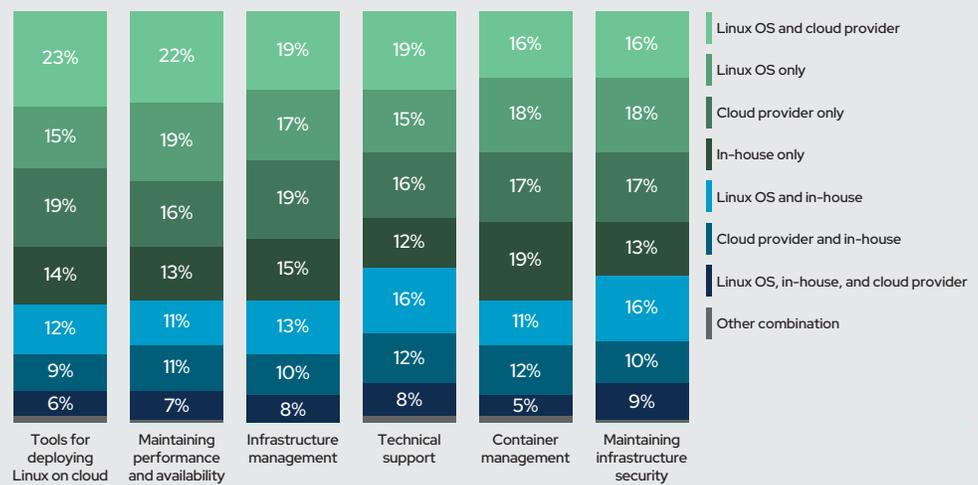
Figure 23. Factors that influence the Linux selection when deploying in multiple clouds



N=568

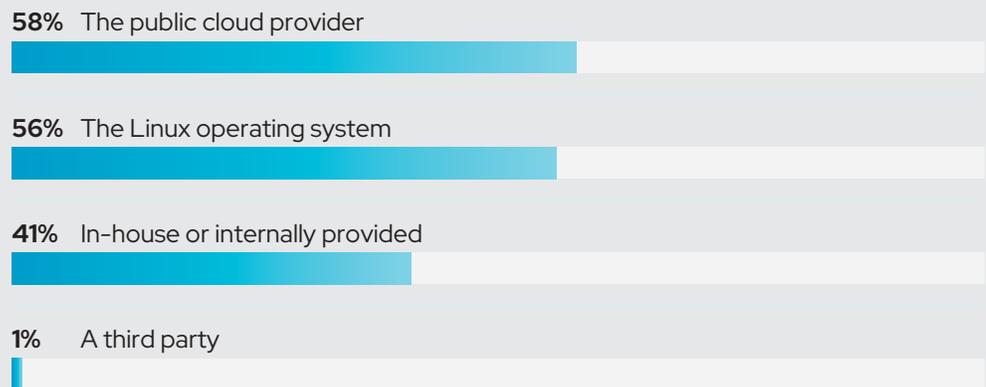
Moving to cloud environments shifts how respondents want to interact with software vendors and where they want to get their software and tools. Respondents were almost evenly split between using the Linux vendor, cloud provider, third party, or in-house development for key software assets like security, management, containers, and support. On average, respondents are using over 1.5 different service providers for each area.

Figure 24. Where organizations are obtaining tools and services for Linux administration



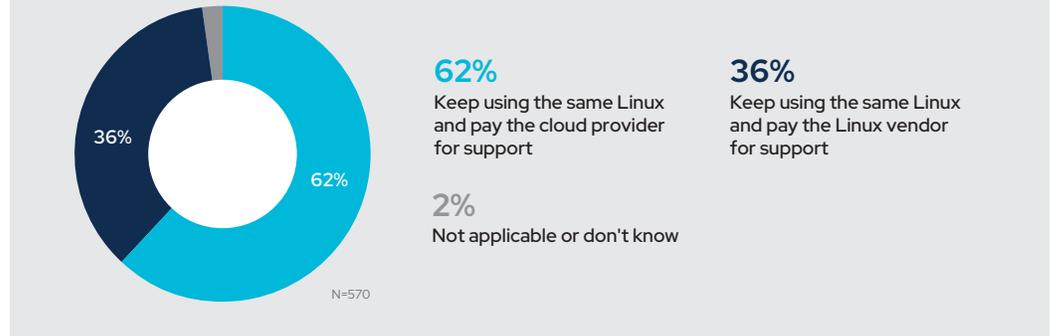
This is a less than ideal situation for most organizations. Respondents indicated that their preferred relationship would be to obtain Linux OS through their cloud provider and pay them for support, switching the relationship from the vendor directly to their cloud provider.

Figure 25. Where organizations obtain key services or tools (simplified)



Average number of sources used\*: 1.58 | N=570

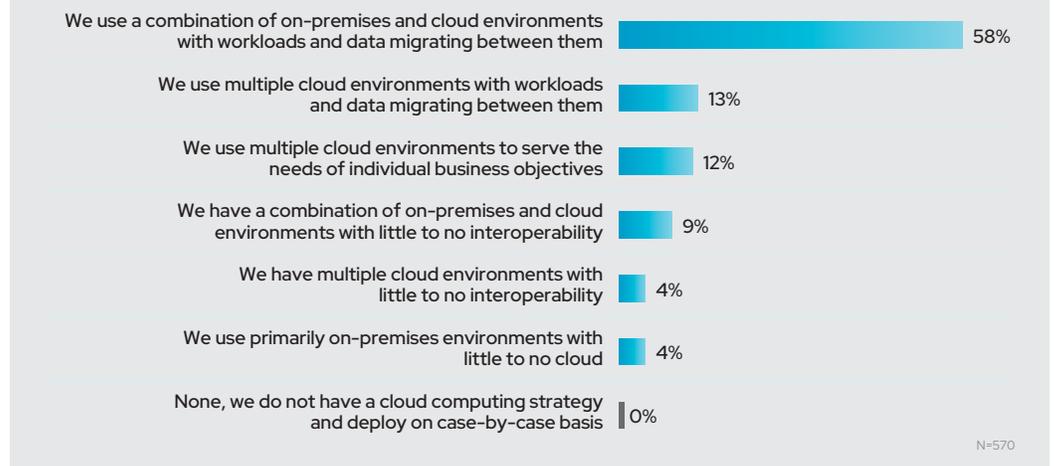
Figure 26. Preferred way to purchase support for cloud-based Linux deployments



**Where workloads reside is surprisingly stable.**

By far, the majority of respondents (58%) have a hybrid cloud strategy, with active integration and migration between physical and cloud environments. A quarter of respondents are using a multicloud strategy, and still expect some level of interoperability between environments. Only 13% of respondents (combined) have little or no interoperability between their cloud environments or hybrid cloud and on-premise environments.

Figure 27. Current cloud strategy



Interoperability and portability are critical considerations for cloud environments because organizations are frequently moving workloads, most frequently from a datacenter to a cloud environment (58%), but also moving between multiple cloud environments (31%). Only 2% of respondents had neither moved a workload between environments nor deployed a new workload.

Figure 28. Actions taken in the past 12 months



N=570

Overall, organizations expect to have similar patterns for migrations and new deployments, with three exceptions: a significant increase in new cloud deployments (from 40% to 48%) and significant decreases in migrating cloud-native workloads into a datacenter (from 30% to 23%) and repatriating workloads (from 27% to 21%).

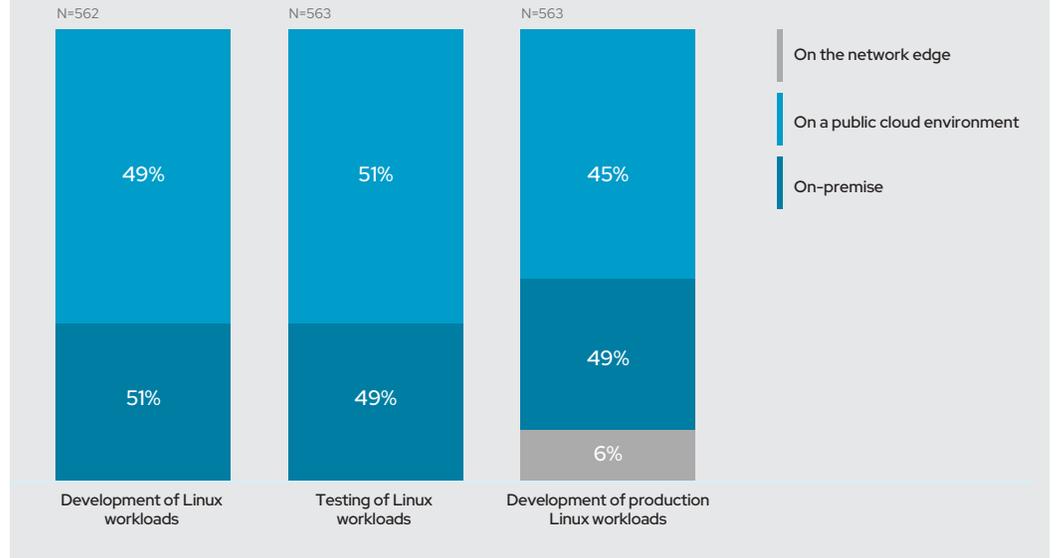
Figure 29. Actions expected in two years



N=570

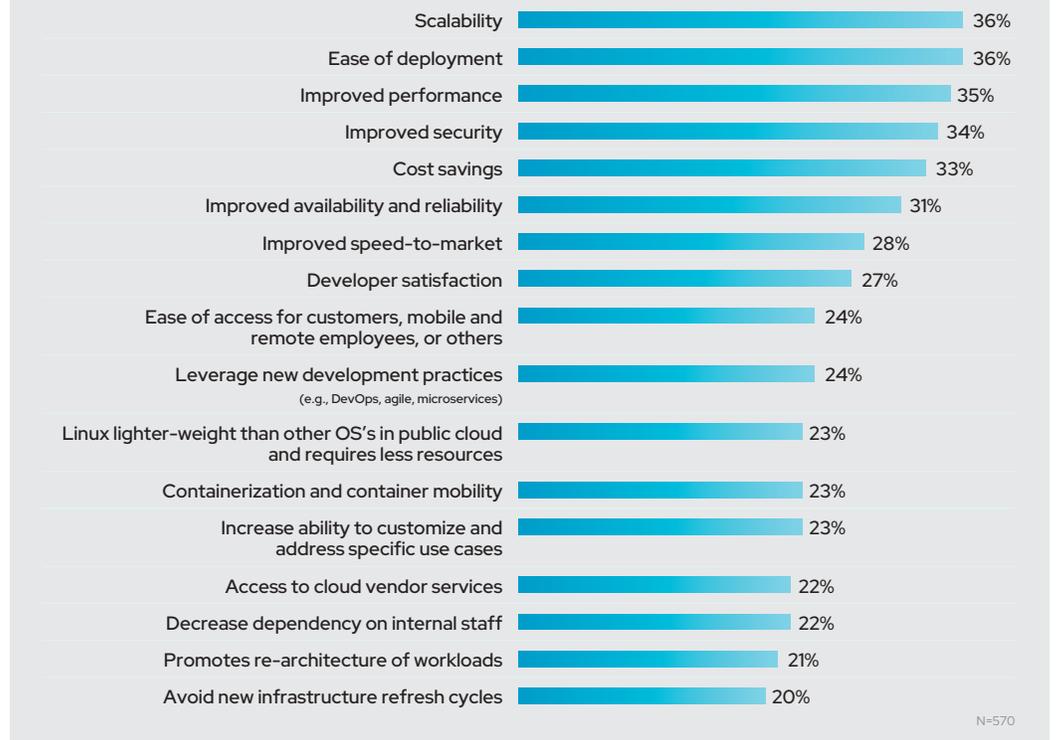
Although workloads are frequently moving between environments and new workloads are being deployed, the overall allotment of workloads between datacenters and cloud environments stays consistent across workload life cycle stages. The only slight difference is production environments, where 6% of workloads are using edge computing, while neither development nor testing environments have edge components.

Figure 30. Linux workloads by production life cycle stage



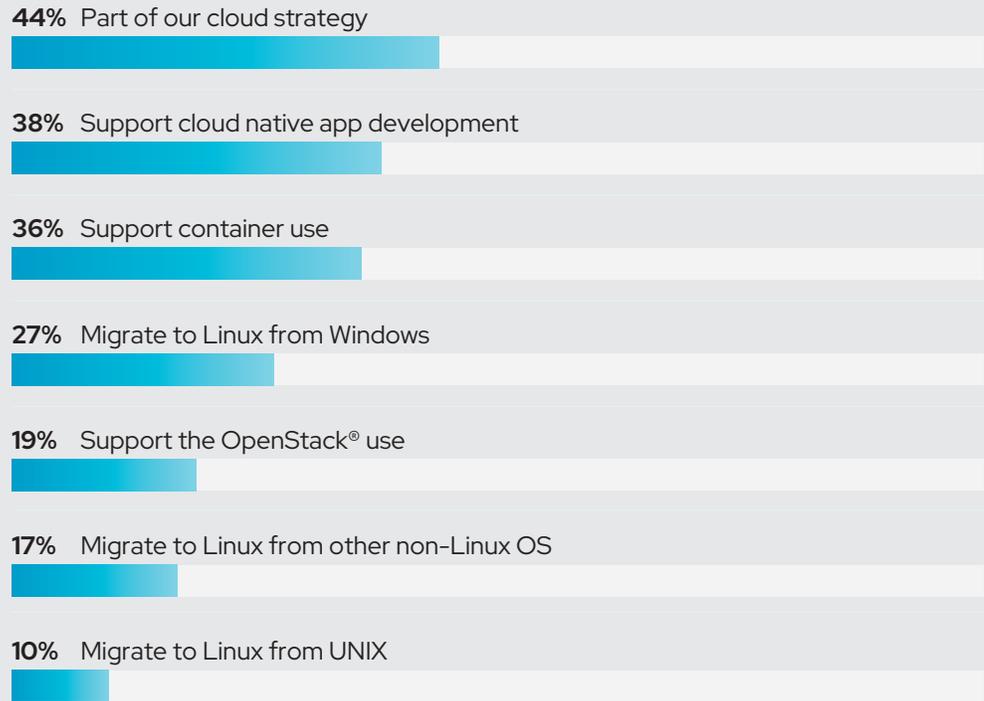
The top benefits that organizations expect to see when deploying Linux in a cloud environment are increased scalability and ease of deployment, tied at 36%. Rounding out the top five benefits are improved performance, security, and cost. Interestingly, cost savings is only the fifth most important benefit for using cloud generally, but cost optimization is the primary factor in selecting a Linux distribution for multicloud deployments.

Figure 31. Primary benefits expected when deploying Linux on a public cloud



The most common action that organizations have taken within the last 12 months or expect to take over the next two years is to migrate existing on-premise applications into a cloud environment. Most of those migrations are strategic choices as part of their cloud infrastructure strategy or their application strategy. Fewer respondents migrated from one OS to another as part of moving to a cloud environment; the biggest subset was moving from Windows on-premise to Linux in a cloud environment at 27%.

Figure 32. Reasons for most recent on-premise to cloud migration



N=331

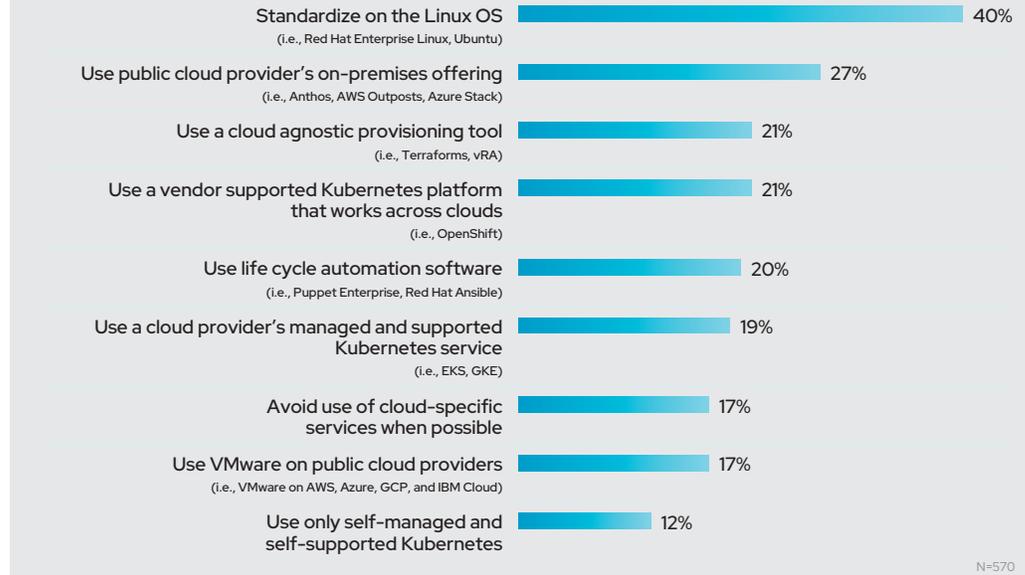
### **Consistency is the central principle for cloud portability.**

Because workloads are constantly moving between environments in our respondents' infrastructures, part of managing applications is trying to define a consistent approach to portability. Interestingly, respondents had slightly different approaches overall depending on whether the migration was from on-premise to a cloud environment or between multiple cloud environments.

For both types of migrations, the number one preferred strategy is using a consistent OS, with some nuances in their approaches to each environment.

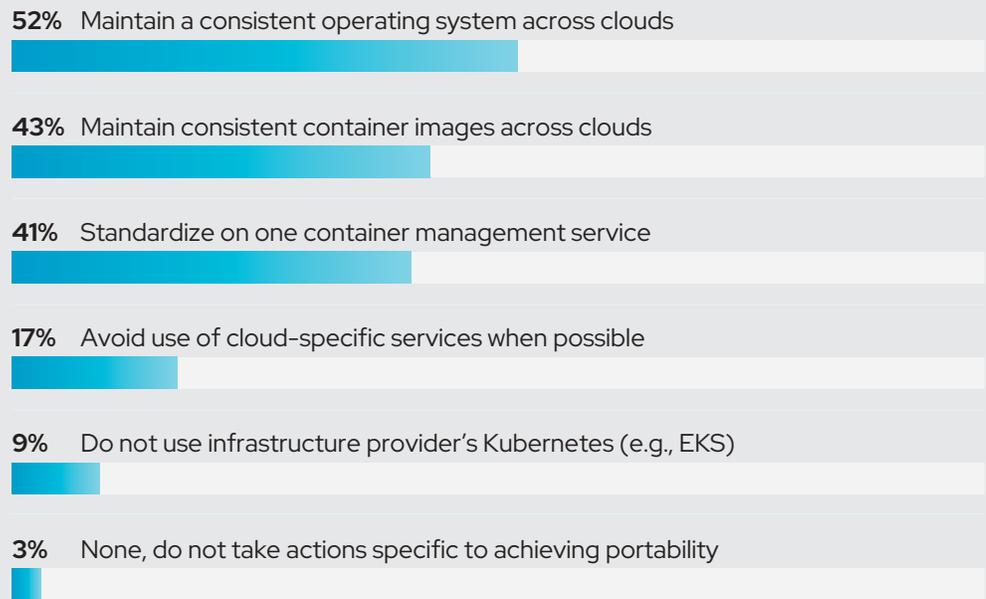
Standardizing on an OS is the preferred approach to migrate workloads by 40% of respondents, the top option. However, respondents also are willing to refactor an existing application or use extensive deployment pipelines for on-premises-to-cloud migrations. Only 17% of respondents avoid using cloud-specific tools to maintain portability.

Figure 33. Actions to make datacenter to cloud migrations easier



However, for cloud-to-cloud migrations, standardizing on an OS was preferred by a clear majority of respondents (52%). Another 43% were committed to maintaining consistent container images across cloud environments. As with on-premise migrations, a significant percentage (41%) were willing to use container management services for deploying workloads, but only 17% would avoid cloud provider-specific tools to maintain portability.

Figure 34. Actions to make cloud-to-cloud migrations easier



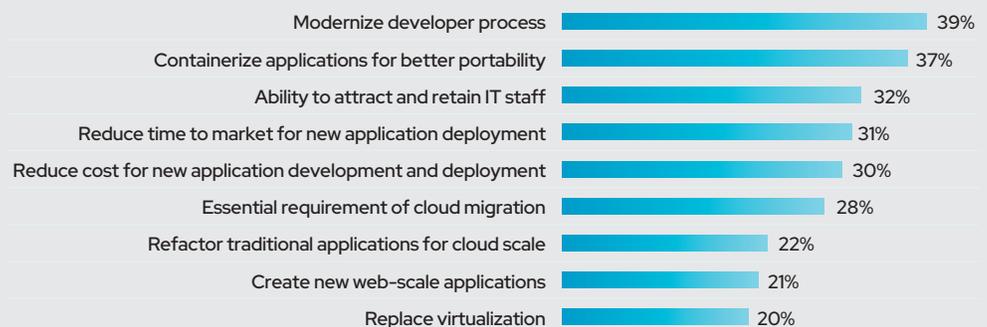
N=570

## Containers

**The top container goal is for developer productivity—but the goals reflect a similar adoption pattern as cloud.**

The top reason that organizations are adopting containers is to modernize the developer process (39%), but a close second is application portability (37%). Application portability is a similar priority with cloud adoption, since workload migrations from the datacenter to a cloud environment are the most common. Interestingly, portability is much more important to IT executives than it is to developers.

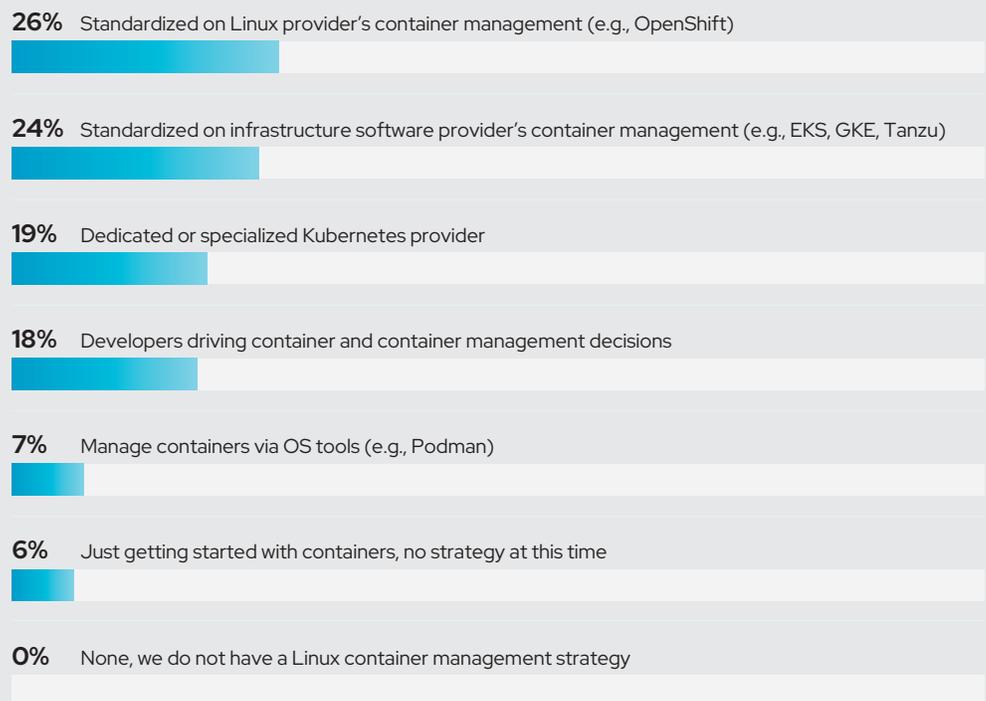
Figure 35. Top goals for adopting containers



N=547

Much like cloud adoption, standardization is the preferred way to manage the container environment, although unlike with a cloud environment, that standardization is on a container management system rather than the container image.

Figure 36. Strategy for managing containers



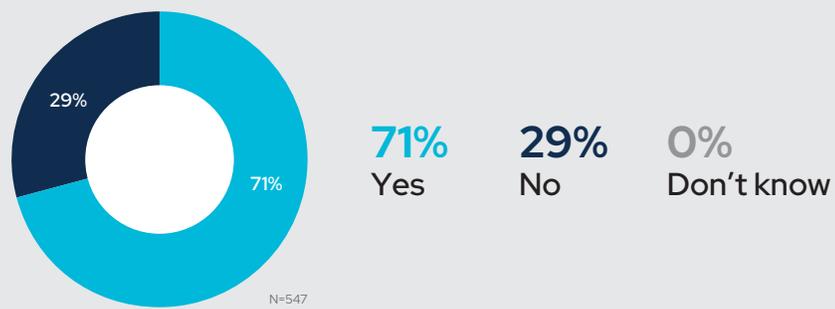
N=547

The approach to container management is segmented, and perceptions of how organizations should manage containers are skewed. Over half, 54%, of respondents are using a Kubernetes-based platform to manage containers (either from a container vendor or from a service provider), but an overwhelming majority (71%) say that Kubernetes is a requirement for deploying containers.

Figure 37. Tools to manage Linux containers

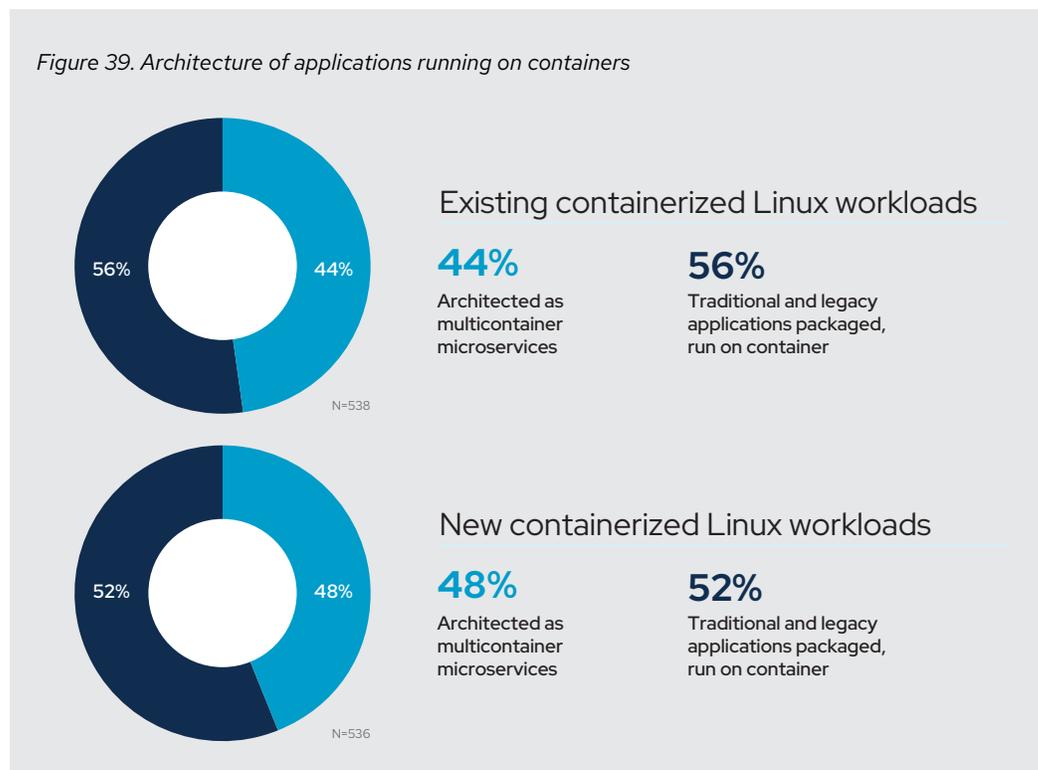
Vendor Kubernetes product		Region			
		NA	EMEA	APAC	LATAM
Red Hat OpenShift	31%	25%	39%	39%	32%
Open source and community Kubernetes	17%	10%	22%	30%	9%
VMware Tanzu	10%	8%	13%	15%	6%
Rancher	9%	7%	13%	7%	13%
Pivotal Container Service (PKS)	9%	5%	6%	8%	31%
Heptio	2%	1%	1%	5%	2%
<b>Non-Kubernetes management</b>					
Docker Enterprise	33%	48%	15%	14%	36%
Amazon Elastic Container Service (ECS)	18%	17%	19%	16%	22%
<b>Cloud vendor Kubernetes service</b>					
Azure Kubernetes Service (AKS)	21%	19%	26%	18%	26%
Amazon Elastic Container Service for Kubernetes (EKS)	21%	21%	20%	21%	23%
Google Kubernetes Engine (GKE)	12%	12%	7%	12%	19%
<b>Built-in tool</b>					
Container management tool built into OS	5%	4%	3%	8%	10%
Podman	3%	3%	2%	4%	8%
None of the above, we self-manage	1%	1%	0%	0%	2%
	N=547	N=259	N=119	N=116	N=53

Figure 38. Is Kubernetes a requirement for deploying containers



Despite the popular association of containers with microservices, the majority of workloads running on containers are architected as monoliths. This is true both for existing container applications (56%) and for upcoming applications that will be developed on containers (52%).

Figure 39. Architecture of applications running on containers



Existing applications are commonly converted to containers (simply moved as a monolithic application), and this is a trend that the majority of respondents see continuing for over a year or more.

Figure 40. Expected timelines to port applications to containers

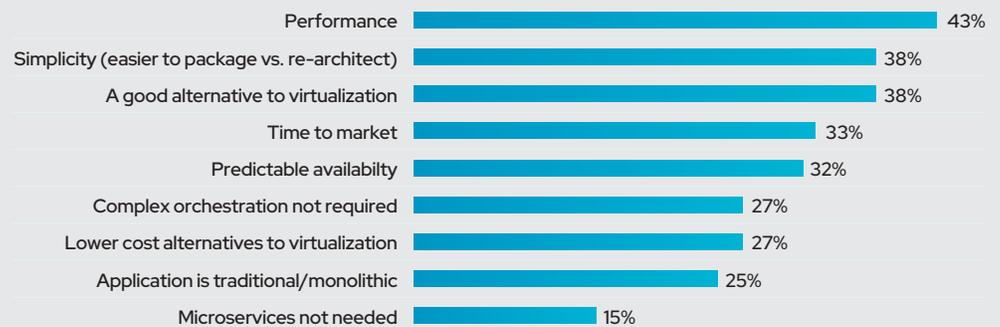


N=539

**A major appeal of using containers is simplicity.**

When organizations create or migrate applications in a single container, major reasons are performance and simplicity. Simplicity could mean the architecture itself (38%), an alternative to virtualization (38%), or application performance (43%). Other considerations like speed, availability, and cost ranked lower.

Figure 41. Reasons for using a single container architecture

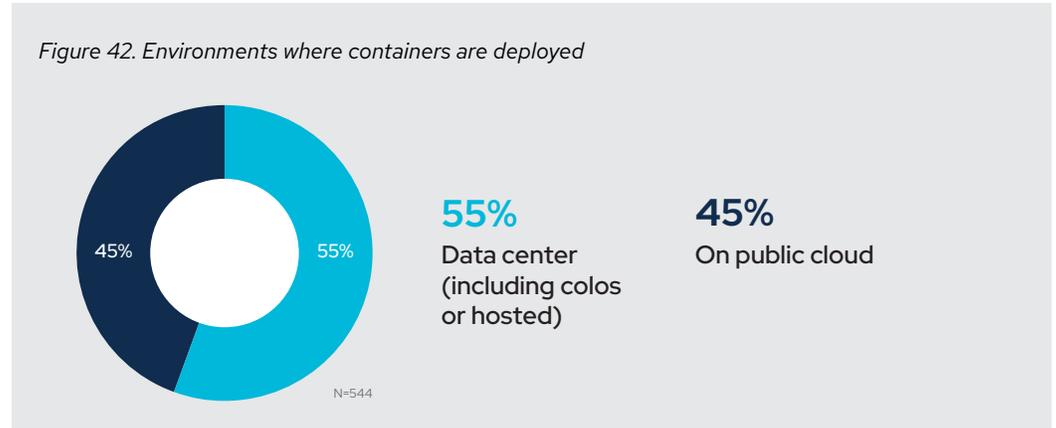


N=539

**Using containers doesn't necessarily mean moving to a cloud environment and moving to a cloud environment doesn't necessarily mean using containers.**

As with monolithic versus microservices architectures, there is an even divide between containers hosted in a datacenter or in a cloud environment—with a slight edge to datacenters (55%).

Figure 42. Environments where containers are deployed



That ratio is reversed when asking whether containerization is required for cloud environments, with 53% of respondents saying that they are required. The overwhelming majority of cloud applications are currently deployed in virtual machines (60%), a number which is not expected to move much in the next two years (58%).

Figure 43. Is containerization required for public cloud environments?

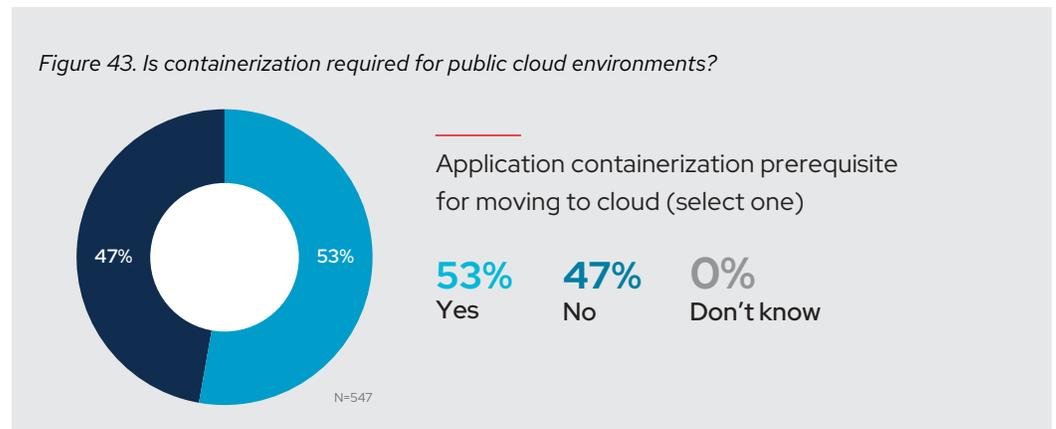
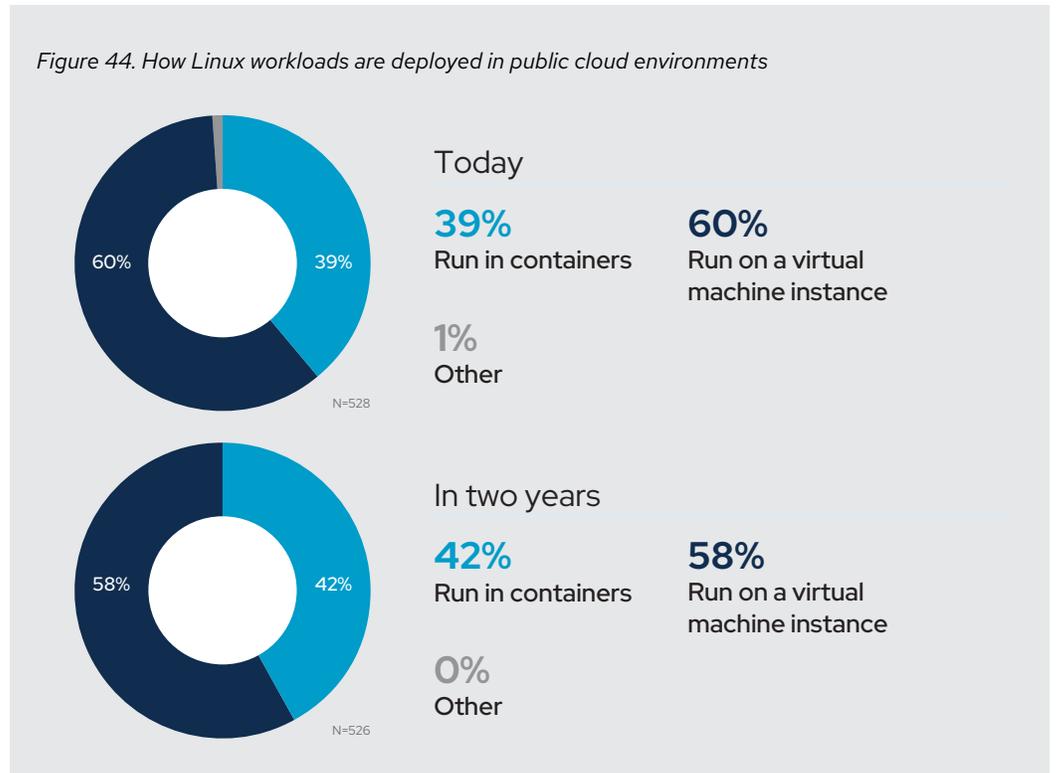


Figure 44. How Linux workloads are deployed in public cloud environments



### About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

#### North America

1 888 REDHAT1  
www.redhat.com

#### Europe, Middle East, and Africa

00800 7334 2835  
europe@redhat.com

#### Asia Pacific

+65 6490 4200  
apac@redhat.com

#### Latin America

+54 11 4329 7300  
info-latam@redhat.com

f facebook.com/redhatinc  
@RedHat  
in linkedin.com/company/red-hat

redhat.com  
#F31489\_1222

Copyright © 2022 Red Hat, Inc. Red Hat, the Red Hat logo, and Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation in the United States and other countries, and are used with the OpenStack Foundation's permission.