O'REILLY®
Report

Redefining OS Selection

Transforming Your Linux Strategy into Competitive Advantage

Ned Bellavance & Chris Hayner

Red Hat
Enterprise Linux

Turn your OS choice into a competitive advantage

Foster rapid innovation with a modern platform. Learn how Red Hat Enterprise Linux delivers agility and control to adapt to evolving demands.

Explore more

Redefining OS Selection

Transforming Your Linux Strategy into Competitive Advantage

Ned Bellavance and Chris Hayner

O'REILLY®

Redefining OS Selection
by Ned Bellavance and Chris Hayner

Published by O'Reilly Media, Inc., 141 Stony Circle, Suite 195, Santa Rosa, CA 95401.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*http://oreilly.com*). For more information, contact our corporate/institutional sales department: 800-998-9938 or *corporate@oreilly.com*.

Acquisitions Editor: Megan Laddusaw
Development Editor: Gary O'Brien
Production Editor: Jonathon Owen
Copyeditor: Stephanie English
Cover Designer: Ellie Volckhausen
Interior Designer: David Futato
Interior Illustrator: Kate Dullea

February 2026: First Edition

Revision History for the First Edition
2026-02-13: First Release

979-8-341-66786-0

[LSI]

Table of Contents

Introduction

Historically, organizations have been slow to change their operating system (OS) deployment strategy. A company would run based on simple technical requirements and preferences and run multiple OSes only if it was absolutely necessary. Changes could be mandated by the technical requirements of a crucial application or forced upgrades from an OS reaching end of support. This led organizations to believe that OS selection was a settled issue—a technical necessity more than anything else. In truth, the speed at which the modern global technological landscape is evolving turns OS selection into a strategic imperative.

Nowadays, companies find themselves grappling with change that far outpaces prior experience. In a few short years, technologies such as containerization have made workflows more flexible, application deployment toolsets and strategies have changed radically, and technology such as AI and quantum computing have completely upended everything that was previously understood. As companies re-evaluate their technology needs, they must also adapt their OS selection strategy to meet these more turbulent times.

This report examines the drivers of this new strategic imperative. We draw upon empirical research including surveys and analyst reports to identify specific challenges facing modern IT organizations. As you will see, this change is driven by a number of strategic, and very modern, triggers. New applications, hardware, and platforms all align with a drive toward keeping up with innovation acceleration. Outside of innovation, organizations have to deal with an uncertain economic climate, sophisticated modern cybersecurity

threats, and geopolitical concerns such as data sovereignty and supply chain validation.

We will show how enterprises are redefining the concept of OS selection from one of operational necessity to one of strategic competitive advantage.

CHAPTER 1

The Strategic Shift in OS Selection

Requirements created by a modern IT environment are driving a major fundamental shift in enterprise OS infrastructure and app deployment strategy. The traditional benefits of an unchanging and unexamined OS ecosystem—simplified management and deployment, reduced need for training, etc.—are being rethought in the light of the modern speed of progress. Organizations are recognizing that rigid OS preferences are liabilities more than strengths; flexibility and rapid capability growth are more important than maintaining the status quo. Companies that don't consider this path risk being left behind by competitors that are more ready and able to be flexible. Additionally, the speed of the changing technology climate is further complicated by geopolitical factors that directly impact procurement decisions.

Major Themes and Changes

Several themes are now driving OS selection decisions in the enterprise. These are not isolated trends but rather interconnected forces reshaping how organizations approach their infrastructure strategy. We will talk about each of these in more detail later in this report.

Innovation acceleration is the dominant driver of OS selection decisions. Simply put: AI, automation, edge and cloud computing, IoT, and new custom hardware are evolving so fast that not everyone can keep up. Companies are concerned about technical debt as technology evolves faster than ever. Emerging technologies once required years to reach production readiness; now they can go from idea to

deployment in months. This is the primary driver for investigating OSes that enable businesses to compete.

Global uncertainty about supply chains has been on the rise for a number of years, prompting companies to look beyond their local environments to the supply chain as a whole that delivers the tools, software, and services they rely on. Many companies have identified single-vendor relationships, especially those with foreign nations, as a potential risk given the likelihood of supply chain disruption. These concerns have made vendor diversification both a risk mitigator and a strategic imperative.

Economic pressures have intensified, focusing on optimization and return on investment (ROI). ROI is not just financial—skills shortages are just as much of a concern as budget constraints, which complicates OS selection. In the past it was rare for a single admin to be an expert in many OSes. It was just not realistic for companies to have multiple high-level experts on staff—one of the reasons that a single-OS environment was the default. Thus, a demonstrable OS ease of use is a high priority.

Security concerns have redefined infrastructure requirements. More systems are being deployed faster than ever before, and thus, dramatically increasing the attack surface for companies. In addition, the new world of AI implementations, modern cyber threats, and regulatory compliance requires modern approaches.

Finally, *AI* has upended everything that we thought we knew about working with computers. The speed at which admins or developers can create new applications with the help of large language models (LLMs) is also unprecedented. Whether you're simply connecting to an LLM to ask questions, creating your own chatbot, or building sophisticated agentic AI tools, you need the best OS for the job.

The Importance of OS Selection

As we have seen, the choice of an OS has a lot more inputs than ever before. Poor OS choices can limit access to critical AI capabilities, hinder performance, and expose organizations to security risks or supply chain vulnerabilities. In contrast, selecting the proper OS can be a stepping stone into this high-paced, modern, productive world. Careful strategic selection of an OS can unlock competitive

advantage and digital flexibility, both essential to companies looking to capitalize on emerging technology opportunities.

Now let's take a look at each of the selection factors in a little more detail.

Innovation Acceleration

The relentless pace of advancement has fundamentally changed the standard model surrounding infrastructure decisions. This section looks at these accelerating factors and how companies are changing their approach to deployment strategies because of them.

Faster tech cycles

As the *McKinsey Technology Trends Outlook* report for 2025 states, "The global technology landscape is undergoing significant shifts, propelled by fast-moving innovations in technologies."[1] Simply put, developments in technology are measured at a pace that has never been seen before—and that pace will only continue to increase. For example, consider the rate of cloud adoption. In a large number of cases, companies started slowly, deploying dev/test environments in the cloud first and then graduating to full production workloads over a period of years. Then consider generative (GenAI) or agentic AI, which went from thought experiment to enterprise product in a matter of months.

In a traditional infrastructure strategy, companies utilize their existing OS choices and find a way to make the new technology work within them. The pace we are seeing in the modern global market is simply too fast for this to be practical. Development teams are increasingly experimenting with new platforms—searching for the ones that work best with the latest technologies—rather than trying to retrofit their old infrastructure to make it work. This is a strong driver for OS experimentation: if the current OS can't do what developers and engineers want it to do, they will begin looking for an OS that can.

1 McKinsey & Company, "McKinsey Technology Trends Outlook 2025," July 22, 2025, *https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-top-trends-in-tech*.

Targeted investment

OS reviews and changes are driven by new technologies more than anything else. This reflects the trend we discussed earlier: developers and admins seeking best-of-breed infrastructure rather than sticking with their in-house favorites. But that doesn't mean organizations are willing to spend unlimited resources; there is a strong desire to spend limited funds the right way, the first time. The choice of OS is more than just a decision around a single server, it's a part of a much larger platform, making the investment decision all the more important.

These investments are made to satisfy business needs. Technology advances, and organizations need to keep up with it as it evolves. For example, organizations are recognizing that AI capabilities require specialized infrastructure and development tools and that their standard OS may not be up to the task. Companies have been forced to build out massive farms of GPU-based servers just to keep up with the needs of AI tooling. Forward-thinking organizations recognize the pace of innovation and are adapting their infrastructure strategies to maximize its potential. Companies want to evolve fast enough to capture the maximum value of technology confidently, rather than just deploying it and hoping for the best.

Global Uncertainty

Global uncertainty has emerged as a factor in technology strategy. Geopolitical tensions leading to economic concerns have more of an impact on decision making than they have in years past. As such, organizations cannot look at their vendors in a technological vacuum—they must now evaluate their choices in terms of geopolitical considerations. The level of concern that organizations have about these uncertainties is higher than ever.[2] Organizations need to be concerned with tariffs and regulations, as well as export controls. The ongoing and uncertain nature of these geopolitical tensions threatens global supply chains, which could lead companies to keep their tech ecosystems within their own borders.

2 McKinsey & Company, "Navigating the New Geopolitical Uncertainty," January 16, 2025, *https://www.mckinsey.com/capabilities/geopolitics/our-insights/navigating-the-new-geopolitical-uncertainty*.

Tariffs and export controls

Tariff concerns and potential resultant currency shifts are changing the way companies look at purchasing across country borders. A long-term contract could look very different a year after signing it if volatility erupts, changing the currency dynamics between countries. This is a crucial issue that is currently hitting the entirety of IT strategy. Adding this economic dimension to vendor selection is a significant departure from previous thinking about IT purchasing decisions.

Digital sovereignty

Geopolitical pressures extend beyond just economics—there are regulatory requirements that differ from country to country, changing the dynamics between customers and vendors. This is most clearly seen in cloud computing, where regulatory concerns such as the General Data Protection Regulation (GDPR) are causing companies to clamor for "EU-only" clouds to guarantee that data will stay within the EU. This was after years of contention that caused a large number of companies to actively explore cloud vendors more conducive to their specific regulatory needs.

Economic Uncertainty

Economic pressures have always had an impact on IT. The speed of change, combined with global pressures, is making things even more complicated. There are limited funds to spend, and there are also skills shortages, which means that even if there is the option to fund a new administrator position, there might not be a worker to staff it.

Budget constraints have made cost optimization a primary factor in OS selection. Companies are increasingly focused on ROI and spending efficiency. The emphasis on cost optimization has implications for OS selection from both a financial outlay and ROI perspective.

Security and Data Breaches

Security of IT assets has been a primary concern for many years, and surveys such as Stack Overflow's 2025 Developer Survey clearly show that security, vulnerability analysis, and testing are top of mind

for developers. "Security and privacy concerns" ranked as the number one reason developers moved away from a given technology.

Supply chain attacks

Knowing where software comes from is only part of the picture; increasingly, companies need to know where the dependencies that software relies on come from as well. This can be a security problem in and of itself—if a minor package that is crucial to a major software offering is compromised, the organization is at risk. Perhaps the most famous example of this is the SolarWinds hack from 2021, which exposed untold numbers of SolarWinds customers to data breaches due to an insecure SolarWinds supply chain. The dependencies issue is still prevalent, however. Recently, the JavaScript registry npm suffered a supply chain attack that allowed packages assumed to be safe to be deployed on users' devices to mine cryptocurrency.

AI has proven to be both a security risk and a boon to productivity and business functions. AI-powered attacks have exploded in the past few years with no sign of slowing down. A recent report by the credit reporting agency Experian identifies AI as a threat significant enough that it will begin overtaking human error as the top cause of data breaches in 2026.[3]

Quantum encryption concerns

The emerging threat posed by quantum computing has been identified and researched for the past few years. The concern is that, due to the completely different approach quantum computers take to the problem of classical encryption, eventually (likely within the decade, but this timeframe is disputed), all existing encryption algorithms will be crackable in seconds. This will render all "classically" encrypted files and datastores completely vulnerable. Hackers know this and have been downloading as much as they can even though they cannot break the encryption yet. This security risk is referred to as "Harvest now, decrypt later."

3 Experian, *2026 Data Breach Industry Forecast*, December 5, 2025, *https://www.experian.com/thought-leadership/business/2026-data-breach-industry-forecast-report*.

Post-quantum cryptography (PQC) refers to the next generation of encryption designed to withstand the power of quantum computers. This type of cryptography is an emerging computer science research topic, and not all OSes are taking advantage of it yet. Encrypting technology with PQC algorithms means hackers will not be successful in harvesting now and decrypting later.

AI

AI has emerged as a dominant force driving a lot of technology buying decisions. This trend has been clear for some time, and there does not seem to be any decrease in momentum. Notably, AI is increasingly driving OS selection, as is the technology surrounding it.

Vibe coding

Thanks to AI, the modern use case of an OS has changed. Programming, data analytics, and business analysis are now accessible to people without specialized or technical backgrounds. To pick a single example, vibe coding is an AI-powered method of creating applications without the user necessarily having to be a programmer or application developer by trade. Increasingly, the market favors OSes that work with AI-assisted development tools that empower such behaviors.

Hardware acceleration

In addition, AI does still have traditional use cases that require extensive OS capabilities. Working on tasks such as inference, model training, or hosting an LLM-based chatbot requires the OS to interact effectively with GPUs and other custom hardware. Specialized chips allow for faster, larger, and more versatile AI capabilities, making hardware compatibility a critical selection factor. AI-driven automation and management functions will continue to tax the hardware resources available to an OS. Consequently, an OS's ability to interact with advanced hardware will have an impact on OS selection.

Conclusion

The market disruptions driving organizations to evaluate their OS choices present a strategic opportunity to explore new possibilities. Organizations are recognizing the unprecedented opportunity heralded by this global chaos and innovation explosion. They are being increasingly proactive in exploring new opportunities rather than trying to stick with existing platforms that may not be keeping up with the changes. This transformation of OS selection from a liability to an opportunity represents one of the most significant opportunities to enhance business operations. It is finding the best tool for the job in a storm of change, rather than simply doing things the way they have always been done.

CHAPTER 2

From Problems to Solutions

In Chapter 1, we identified several critical factors that drive OS selection: innovation acceleration, global uncertainty, economic pressures, security concerns, and the demands of AI. Recognizing these pressures is only the beginning, however. The challenge that faces IT leaders is how to respond to these pressures effectively—without creating new operational confusion or vulnerability. Successful organizations will recognize these issues and respond in a productive, competitive way. In this chapter, we will address each of these challenges and discuss how we can go from problems to solutions.

Accelerated Innovation Strategies

Innovation acceleration has become a driving force for IT leaders. This is not just from a need to keep up with the competition; thanks to concepts like first mover advantage, the organizations that can move fastest tend to capture the largest rewards. This has significant implications around OS selection strategies because the OS provides the base for the accelerated innovation an organization would like to capture. This kind of innovation will lead companies into uncharted waters, so it's important to be cautious and make decisions based on identifiable needs and benefits. Innovation acceleration does not mean making changes on a whim. A modern OS needs to empower experimentation while still providing security, performance, and reliability.

Thoughtful Experimentation at Speed

The acceleration of technological innovation has created market conditions where traditional development and evaluation methods are simply not fast enough. In the application development world this has been a trend for some time, going from a waterfall-based annual release cycle to Agile-based software development and DevOps-powered continuous delivery to today's hyperfast deployment models with multiple daily updates. This same trend is being seen across the technology landscape, with IT leaders recognizing that they need to follow suit to stay ahead of the curve.

This kind of innovation speed requires rapid cycles of developing and testing, all while maintaining a security baseline that doesn't put businesses at risk. Effective experimentation requires sophisticated sandbox environments that provide similar conditions to production without putting company data (or the OS hosting the application) at risk. This is the most efficient way to provide realistic testing environments that can be used to quickly assess performance under actual production conditions.

In a similar vein, applications need a fast, reliable pathway from testing to production. This should be a seamless process that allows applications that have been proven in testing to work reliably in production. This is not just about convenience—the reliability and repeatability mean that upgrades are just as efficient (and secure) as the initial test deployment.

Considering this, there is obvious value in maintaining a cautiously open mindset with OS selection processes. Being agile in a period of accelerated innovation is a clear strategic advantage. The OS ultimately hosts the applications in question and, with application requirements changing so quickly, organizations need to know that the OS will keep up. This nuanced approach has been borne out in studies and analysis by companies such as IDC and McKinsey. In June 2024, McKinsey stated, "The truth is that there's no one-size-fits-all approach. Instead, companies can align the pursuit of innovation with the need to maintain robust and dependable technological infrastructures."[1]

1 McKinsey, "Rethinking Conventional Wisdom: Future of Digital Tech Infrastructure," June 26, 2024, *https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/rethinking-conventional-wisdom-future-of-digital-tech-infrastructure*.

Still, tools and technologies must be selected carefully. Change for change's sake is a risk—you're adding complexity without an obvious benefit when something as fundamental as an OS is altered without oversight. The strongest evidence justifying such a change is that which shows clear advantages with regard to organizational capability. These advantages need to be demonstrable, not just running a new OS simply because it's new. To this point, organizations must be looking not only at current capabilities but also at an OS's future potential and stability.

Security-Focused Architecture

The pace of technological innovation does not just benefit legitimate businesses. Cybersecurity threats continue to evolve, leaving businesses vulnerable to attack if they do not keep their organizations secure. Market research such as CrowdStrike's *2025 Global Threat Report* clearly shows that cyber threats continue to evolve faster than cyber defenses. The modern threat landscape presents unprecedented, constantly evolving challenges that many traditional security frameworks and IT tools were simply not designed to address. Being ready to face these modern threats has become a strategic imperative.

This security architecture must be considered when looking at implementing a new OS. Organizations that don't incorporate security questions and verifications into their OS selection processes face a dual risk: the applications hosted on the system, and the security and reliability of the system itself. This is especially true at the time of this writing, as OS vendors are increasingly adding AI features to their OSes and allowing them a tremendous amount of access to perform system administration tasks in an automated manner.

Data security is essential

This leads to the main point: any OS needs to provide environments where applications can run securely and data can be stored securely. A company needs to be certain that anything that runs, whether it's a built-in component of an OS or an application that an OS is hosting, remains secure. The following three categories of data security have to be considered:

Data at rest
: This is the data that is stored for future use.

Data in motion
: This data is communicated in some way, be it across systems or within a system, from application to application.

Data in use
: This data is actively being processed by the computing platform.

Securing data at rest and data in motion is often solved by encryption. Securing data in use is more challenging. Modern OS architectures solve this problem by incorporating a Trusted Execution Environment (TEE). This is a section of the system that is limited to highly trusted software and accessible only through tightly controlled mechanisms. Anything that runs within a TEE is generally not accessible outside it. This can be both a hardware- and software-based solution that often extends beyond just one server and requires drivers and compatibility to be kept current in order not to lose capabilities. How this looks in practice varies by vendor and is outside the scope of this report, but confirming that applications can be run in such secure environments is essential.

AI workloads need more than protecting

While computational workloads of all kinds require the data protections and data sovereignty assurance requirements discussed previously, AI workloads are especially vulnerable for two additional reasons: they are new and data-intensive, and computationally complex.

The newness of AI workloads means that production workloads have not had time to be tested as thoroughly as standard workloads. Security is a collaborative effort, and the market will do a lot more testing and QA than any individual vendor can do on its own. This is true both for the product itself and the way that product is configured/used. In its *Tenable Cloud Security Risk Report 2025*, Tenable reported that cloud workloads supporting AI are more vulnerable than other workloads. Regardless of where the workload is hosted, AI's data-intensive nature, combined with persistent misconfigurations and vulnerabilities, demands a new level of diligence when it comes to cybersecurity.

Organizations need to consider AI-specific risks when evaluating an OS. Failure to do so means bringing in an unknown level of risk that could compromise current and future workloads.

Conclusion

Being intentional about OS decision making requires fundamental changes to how organizations evaluate technology. The challenges presented in this chapter provide the structure necessary to maximize competitive advantage by making the right choices. At the conclusion of this report, we'll look at a decision-making framework that encompasses all of these decisions. For now, the need for new thinking when it comes to OS deployment is clear—modern challenges require modern solutions. Organizations that embrace the challenge and think strategically will find themselves in a better position to utilize new technologies and maximize their competitive advantage.

CHAPTER 3
The Modern Linux Platform

The modern enterprise has, to a large extent, standardized server operations on Linux. This has been borne out over at least the past 10 years in numerous studies. To cite just one, *SQ Magazine* reported in August 2025 that 61.4% of large enterprises run at least one mission-critical workload on Linux and that 78.3% of internet-facing web servers run Linux.[1] This dominant position is the result of many years of growth—a trend that shows no signs of slowing.

In many cases, however, these statements never discuss individual distributions, leading many decision makers to believe that "all Linux is the same." This is not true; while all Linux distributions use the basic Linux kernel, the way the kernel is deployed and the tool and application ecosystem around it can be very, very different.

What an organization needs to do is find key differentiators in the OSes they are evaluating. Many modern OSes have evolved cutting-edge features, but some handle these features differently than others, and some features are simply not available everywhere. Some of the most valuable features to look for include modern secure deployment models, AI-integrated system tools, fleet management tooling, enterprise-grade security, and expanded support and partner certifications. In addition, the capabilities of an OS need to be considered in light of the hosted workload requirements that will be running on it—best-of-breed capabilities that match modern requirements.

1 Robert A. Lee, "Linux Statistics 2025: Desktop, Server, Cloud & Community Trends," *SQ Magazine,* updated November 18, 2025, *https://sqmagazine.co.uk/linux-statistics*.

Revolutionizing OS Deployment

Deploying OSes has not changed that much over the years. The traditional deployment involved an installer of some kind layering the kernel and other tools onto the hard disk, followed by the applications installed by users. The advent of package management was a huge leap forward for application management and system administration because it meant users no longer had to compile code from scratch every time they installed or updated an application.

A modern Linux distribution should provide a deployment mode where systems are essentially immutable containers that can be swapped in and out as needed. You can download a new image, and a simple reboot is all it takes to either upgrade (or downgrade) whatever component of the system you're changing. This provides security, reliability, and confidence that all systems in the fleet are running the expected versions of all binaries and libraries.

You can manage the created images in the same way as application containers (or images), but on an even smaller scale—binaries and even the OS kernel itself. This is a major improvement over the former package-based update model. While packages are a superior deployment model to compiling from source, they have their own issues with compatibility, package requirements, and potential version conflicts. A container-based model removes these concerns because everything is contained in the image. Versions of anything in the container will remain consistent, greatly lessening the concerns many system administrators have regarding configuration drift. Additionally, the image installation is far faster and easier—simply download the image, apply it, and reboot.

Images also offer security benefits because they are immutable. This means that a bad actor or malware would not be able to modify the image in place, nor would they be able to replace it. Because these images are managed in much the same way as application containers, you can look at the image with standard container tools such as Podman or many other common container platforms. This familiarity with the tooling will enable administrators already familiar with application containers to deploy OS images quickly and with confidence.

Finally, the images you create can be used anywhere. It doesn't matter if you are deploying to physical, virtualized, cloud, or edge

environments. The image will be deployed and run in the same fashion everywhere. Images will also be easier to manage from beginning to end. Figure 3-1 shows a sample image mode process across three simple steps: building images, deploying images, and managing images after deployment.

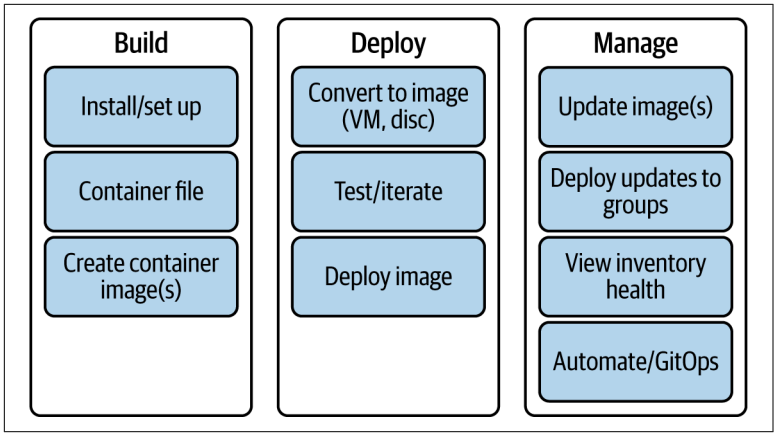

Figure 3-1. A sample containerized (image mode) OS deployment model

AI-Assisted System Administration

Many application creators are beginning to incorporate AI directly into system administration tools such as terminal emulators (which provide access to the command line). These are exceptionally valuable because they allow an admin to query LLMs of their choice without opening a browser window or separate application. The next logical step is for OS manufacturers to do the same, incorporate AI directly into the OS itself. Then the AI assistant can be called from the command line, regardless of the terminal that an end user chooses. This capability will greatly benefit admins of all skill levels, but junior administrators in particular will benefit from the ability to ask an AI assistant questions directly in the command line to re-evaluate on the fly and ensure they're doing the right thing in the right way.

Like all AI tools, command-line assistants should be used with care. It may be the case that you'll need to refine your question to a command-line assistant to obtain an accurate answer, especially if you're asking the assistant to generate code, such as a shell script.

When in doubt, those same junior admins would be wise to ask a senior colleague whether the answers given by the AI assistant are going to do what they think they're going to do.

The tool would be more powerful than simply understanding things like basic bash scripts—command-line assistants can offer convenient tips on how to configure applications and services, as well as utilizing OS features to troubleshoot. Crucially, these assistants would be specifically trained on both Linux in general as well as the distribution's finer points and specific needs and requirements. This fine-tuning is well known in the AI world: a custom model versus a generic LLM trained exclusively on the right kinds of documentation and built to minimize hallucinations makes this kind of tool a significant differentiator in built-in tooling.

A command-line assistant would be especially valuable for tools such as Security-Enhanced Linux (SELinux) and the built-in firewall. These are both very important to system operation, but they are complex and their configurations are changed only infrequently, in most cases. A system administrator will find that configuring these tools will be a lot faster and easier when a command-line assistant is utilized to help. In addition, the AI tool will have a lot of vetted examples to compare to, allowing it to have a good sense of best practice and be able to call out misconfigurations that might not be obvious even to a seasoned administrator.

Finally, command-line assistants can be utilized to help understand the structure of the tools being used. For example, if you're trying to understand how modern host-based firewalls are implemented, you can simply ask the assistant. It will query relevant documentation resources for best practices and examples based on what you've submitted, and it will formulate a plain-English answer that supplements the configuration or source code changes to help you understand what it's recommending and why.

Fleet Management Tooling

The two tools we've discussed are powerful and greatly simplify the management of systems on an individual system level. The next tool would take things to a higher level and help users proactively manage their entire deployed Linux environment. Such a tool would have deep connections into everything that an OS manufacturer offers, meaning that the entire platform can be managed from one

place, simplifying management and auditing needs. Ideally, it would also include a planning feature that helps organizations understand their deployed application lifecycles and how these affect the systems running in their environment. These integrations would allow for security and configuration validations across the entire organization in real time.

Roadmap features would provide detailed information about upcoming app and OS releases so organizations can plan upgrades based on anticipated new features rather than reacting once those new features are released.

Fleet management tools can also be used to keep systems up-to-date. Coordinated patching (whether automated or scheduled and confirmed as successful), feature/software updates or rollouts, and centralized configuration management make it so that managing one thousand servers is as straightforward as managing a single server.

Enterprise-Grade Security

An OS designed with security in mind should include more than the security benefits such as immutable images and configuration management discussed previously.

One of the most discussed future attacks revolves around post-quantum cryptography. As we discussed in a previous chapter, PQC provides organizations with quantum-resistant algorithms for keys, encryption, and digital signatures. This helps future-proof data and applications and is an essential feature to look for when comparing OS distributions.

An enterprise grade OS platform will also provide a robust set of supply chain security features. This is not restricted to just the OS—it is a platform play that helps empower developers and organizations looking to deploy in a DevSecOps manner. The simplest example of this is a trusted location to host known good application dependencies, but this can go much deeper. An ideal DevSecOps platform would provide development teams with an end-to-end security pipeline for all their applications and dependencies.

The combination of PQC protections, automated security tooling for admins and developers, and proactive software management creates a layered defense that stays current and addresses the modern

sophisticated threat landscape. Crucially, this security is provided by the tools and OS the organization uses, making it a full-featured platform play rather than forcing organizations to rely on extensive external third-party security investments.

Expanded Hardware and Software Support and Certification

As powerful as an OS platform is, organizations will still need external compatibility to operate. This is true both for the hardware the OS runs on and for the third-party software that will be leveraged in the running system itself.

In the modern era, hardware compatibility is a fast-moving market. AI hardware in particular is changing extremely quickly, requiring OS drivers and compatibility to keep up. It's important to ensure that your selected distribution excels at this, offering up-to-date AI libraries, drivers, and applications as well as hardware optimizations for manufacturers such as NVIDIA, Intel, and AMD. In addition to these three major players, however, an OS manufacturer must have a robust third-party ecosystem providing the latest functionality to its end users in a speedy, secure, and reliable manner.

Conclusion

An OS's security, ecosystem, and feature set work together to address the challenge identified at the beginning of this report, namely, exploring other OSes is a strategic imperative. The major features identified in this chapter make an OS stand apart from its competitors, both in capability and manageability. The technology advances combined with the strong partner programs represent strategically significant differentiators that cannot be overlooked when choosing a new OS platform.

CHAPTER 4
Next Steps

OS selection is a foundational decision for organizations, with implications for competitive position and internal infrastructure management. Navigating the pressures of innovation acceleration, global uncertainty, economic constraints, modern security needs, and AI requires a platform specifically built to address these challenges in an integrated manner.

This selection is not a decision to be taken lightly. There are four key steps of evaluation that an organization should consider when making decisions around future OSes and platforms:

- Identify the pressures that are most significant to your organization.
- Map your workloads to capability requirements.
- Evaluate overall value to the business.
- Validate supply chain security and long-term viability.

Identify the Most Significant Pressures

Organizations will discover that different teams face different pressure combinations: AI product teams will prioritize innovation and hardware requirements, while regulated industries will have a higher focus on security and compliance capabilities. Organizations should rate each pressure domain (innovation acceleration, global uncertainty, economic constraints, security needs, and AI) based on

current and expected future importance and use that list to build out lists of "must-haves" versus "nice to haves" when it comes to future evaluations.

Map Your Workloads to Capability Requirements

One factor that will drive any requirements is the workloads that are most important to an organization. A workload analysis will reveal whether your organization benefits from workload-specific optimization or whether a larger standardization approach would make more sense. This benefits organizations of all sizes as the number and importance of workloads are not necessarily related to company size.

Evaluate Business Value

Once workloads have been analyzed, an organization will have a firm understanding of server count, employee count (and skill levels), and total workloads under management. These will help to assess direct costs such as OS licensing, support, and infrastructure needs. It will also help estimate indirect costs such as the level of in-house expertise, delays in workload deployment due to unanticipated constraints, and missed business opportunities due to infrastructure inflexibility. Once this data is gathered, companies can make more informed strategic decisions that will enhance operational efficiency and increase competitive advantage.

Validate Supply Chain Security and Long-Term Viability

An organization making an OS or platform decision needs to know that the OS or platform in question will be around for the long haul. Understanding a company's history is important but so is knowing how it vets its applications, the depth and breadth of its partner portfolios, and its experience and capabilities regarding critical aspects like regulatory compliance as well as cutting-edge features such as AI. Organizations expect workload deployments to have multiyear lifecycles; it is essential to ensure the platform can handle the workload now and in the future.

Conclusion

The key to success in this process is understanding that infrastructure decisions create dependencies. Choices made today either enable organizational success or constrain capabilities and delay projects for years to come. Organizations that approach an OS selection decision strategically, with clear requirements and an evaluation framework in mind, position themselves to pursue innovation aggressively, successfully navigate uncertainty, and optimize both productivity and security now and into the future.

About the Authors

Ned Bellavance is an IT professional and technical educator with more than 20 years of experience in the field. He's been a helpdesk operator, systems administrator, cloud architect, and product manager. Most recently, Ned runs Ned in the Cloud LLC where he develops courses, runs multiple podcasts, writes books, and creates original content for technology vendors. Ned has been a Microsoft MVP since 2017 and a HashiCorp Ambassador since 2020. He has three central pillars: embrace discomfort, plan to fail, and be kind.

Chris Hayner is a seasoned IT professional with decades of experience spanning operating systems, infrastructure, cloud computing, and cybersecurity. His career began in the datacenter, where he managed a diverse range of systems, from mainframes to AlphaServers to white-box x86 servers. From here, his scope expanded to include virtualization, cloud technologies, and overarching cybersecurity and IT strategies. For the past 15 years, Chris has worked in the consulting realm, serving as a subject matter expert (SME), architect, and analyst. In these roles, he has helped hundreds of organizations bridge the gap between business objectives and IT solutions, driving innovation and operational success. In addition to holding a CISSP certification and numerous vendor and industry credentials, Chris earned an MBA from Temple University, equipping him with a unique blend of technical expertise and business acumen.