

White Paper

Building an Open, Secure, and Flexible Edge Infrastructure with Red Hat Enterprise Linux

Sponsored by: Red Hat

Dave McCarthy
June 2021

IDC OPINION

Edge solutions are driving the next wave of innovation as organizations execute on digital transformation initiatives. Whether to automate operations, deliver rich customer experiences, or introduce new business models, CIOs want to take advantage of the benefits of cloud-native application design and software-defined infrastructure with the freedom to deploy anywhere.

This shift is part of an expanded definition of hybrid cloud. The concept has evolved from a duality of on-premises and public cloud resources to include edge locations such as factories, warehouses, and retail stores as well as infrastructure and applications deployed across multiple cloud providers. This flexibility in workload placement is essential to achieving the right balance of performance and security.

As these environments grow, there is an increased requirement for consistent management and interoperability. This is driving enterprise adoption of open source and, ultimately, an open hybrid cloud, where vendor lock-in is minimized, standards facilitate integration, and the open source community accelerates new capabilities.

At the foundation of the open hybrid cloud is the operating system (OS), which must be secure, flexible, and optimized for remote management. With the ability to create purpose-built OS images for small footprint devices, efficient over-the-air updates, and intelligent rollbacks, Red Hat Enterprise Linux has features specifically designed for edge deployments.

Red Hat Enterprise Linux is the cornerstone of the Red Hat Edge portfolio and a key element of the broader set of Red Hat products and solutions enabling the open hybrid cloud. IDC believes that Red Hat Enterprise Linux has all of the necessary functionality to ensure success at the edge now and into the future.

SITUATION OVERVIEW

Like the cloud before it, edge has emerged as a massive industry trend. According to IDC forecasts, the overall edge market is on track to reach \$240.6 billion in 2024, with a CAGR of over 15%. This is inclusive of the hardware, software, and services that encompass edge solutions.

The importance of edge has transcended IT departments and has caught the attention of the C-suite. In a recent IDC survey, 73% of enterprises viewed edge as a strategic investment and an additional 17% stated it is required for business operations.

In the simplest of terms, edge brings compute closer to the producers and consumers of data. As the growth of data becomes more distributed, it is necessary to distribute infrastructure and applications as well.

This distribution is causing a shift in IT investments. By 2023, over 50% of new enterprise IT infrastructure deployed will be at the edge rather than corporate datacenters, up from less than 10% today. By 2024, there will be an 800% increase in the number of apps at the edge. Despite this growth, it is important to note that edge is a complement to – not a replacement for – cloud infrastructure.

The edge has the potential to address limitations inherent in centralized architectures:

- **Latency:** The roundtrip time it takes for data to be transmitted to a remote datacenter and processed can be prohibitive for real-time applications that require quicker response times than the cloud can deliver.
- **Cost:** The proliferation of Internet of Things (IoT) and other connected equipment has led to an increase of data generated in remote locations, which can be expensive to transmit and store in a core datacenter or the cloud.
- **Security:** Corporate governance policies can limit the movement of data offsite, requiring local compute and storage.
- **Compliance:** Whether GDPR in Europe or the California Consumer Privacy Act in the United States, there continues to be more data sovereignty regulations dictating rules for where data can be stored.
- **Resiliency:** Distributing infrastructure and workloads to edge locations provides a level of business continuity in the event the network or cloud is unavailable.

Business Motivations and Benefits

The diversity in use cases can be seen in how enterprises are implementing distributed edge solutions. As illustrated in Figure 1, a majority of edge deployments are targeted at the modernization of existing infrastructure. This is a common scenario for organizations with remote office/branch office (ROBO) facilities. Older standalone servers and storage are being replaced with hyperconverged infrastructure that can host a combination of virtualized and container-based workloads.

Implementation of IoT systems is another business initiative linked to investment in new edge solutions. When these deployments were small in nature, IoT data would be transmitted to a central repository for analysis. Over time, cost and scalability challenges surfaced. It became clear there are many situations where it is beneficial to move the compute to where the data is than the other way around.

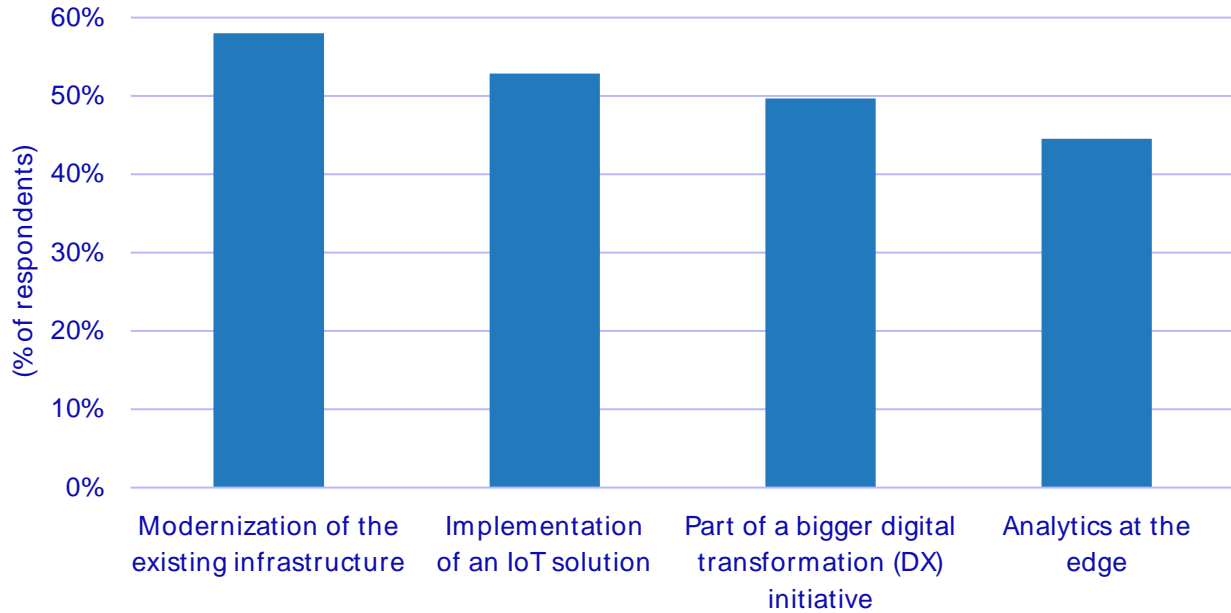
The industrial sector has some of the greatest potential for digital transformation. As IT and operational technology (OT) teams look to integrate systems, most of this activity is happening at the edge. This is part of a move from vertically integrated point solutions that can be costly to acquire and maintain to open and interoperable platforms that facilitate continuous innovation.

The edge is also playing an important role in analytics and data science initiatives. While machine learning and the creation artificial intelligence (AI) models benefit from the processing power of the cloud, the inference of new data against those models is best done at the edge. In these distributed computing scenarios, the emphasis on centralized management and security is paramount.

FIGURE 1

Edge Solution Investment Priorities

Q. *What business initiatives are linked to investment in new edge solutions?*



Source: IDC, 2021

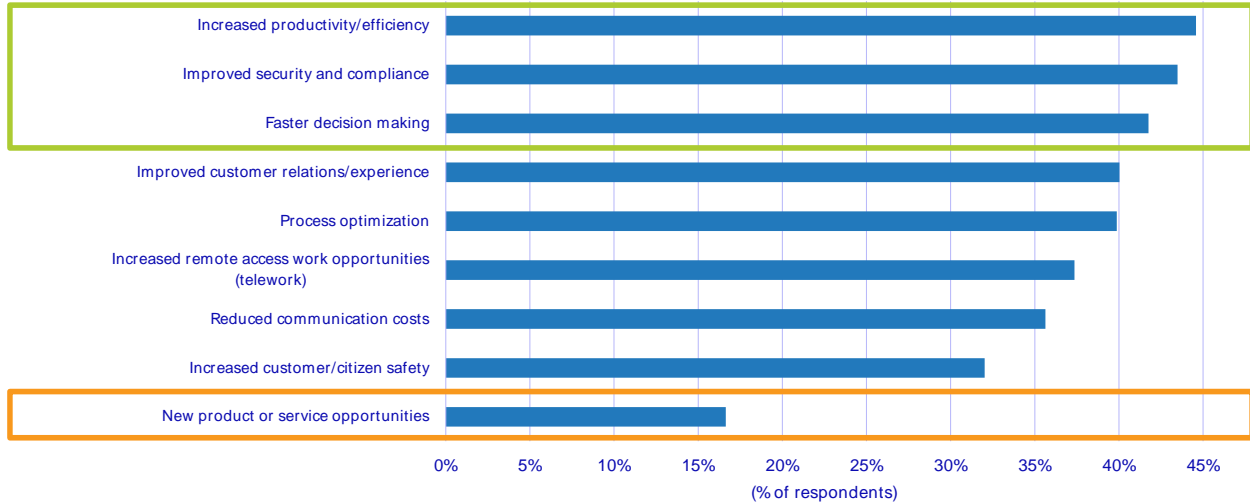
While the reasons for deploying edge solutions may vary, the benefits are consistent and can be thought of in both short-term and long-term perspectives. When surveyed, organizations recognize that placing workloads in optimal locations will lead to increased productivity and efficiency (see Figure 2).

By having full control of the environment, edge also improves security and compliance. These underpin perhaps the most important aspect of digital transformation: faster decision making in a world with ever-changing business requirements.

FIGURE 2

Business Benefits of Edge Computing

Q. What business benefits do you expect edge adds/will add to your organization?



Source: IDC, 2021

But there is also a long-term view. Once a flexible infrastructure is in place, it increases business agility. New products and services can be developed and launched faster. Cloud-native applications that are built on a microservices architecture with API-driven automation can be easily updated and scaled to meet user demand.

Deployment Considerations

While it is common to refer to edge as a singular location, the reality is quite the opposite. Edge deployment locations can be divided into two major categories.

Service Provider

A range of service providers are now offering edge solutions:

- **Cloud service providers:** Recognizing that hyperscale regions are not appropriate for all workloads, public cloud service providers have begun launching metro-level datacenters.
- **Communications service providers:** Often tied to the rollout of 5G networks, cloud and communications service providers are partnering for joint edge solutions.
- **Colocation service providers:** For organizations that need to address remote geographic regions, colocation datacenters can be used to host edge infrastructure.

Customer Premises

When lower latency or more control is needed, there are several customer premises options:

- **ROBO:** To support a distributed workforce, enterprises will deploy edge solutions in customer-owned facilities that often have little to no local IT staff.

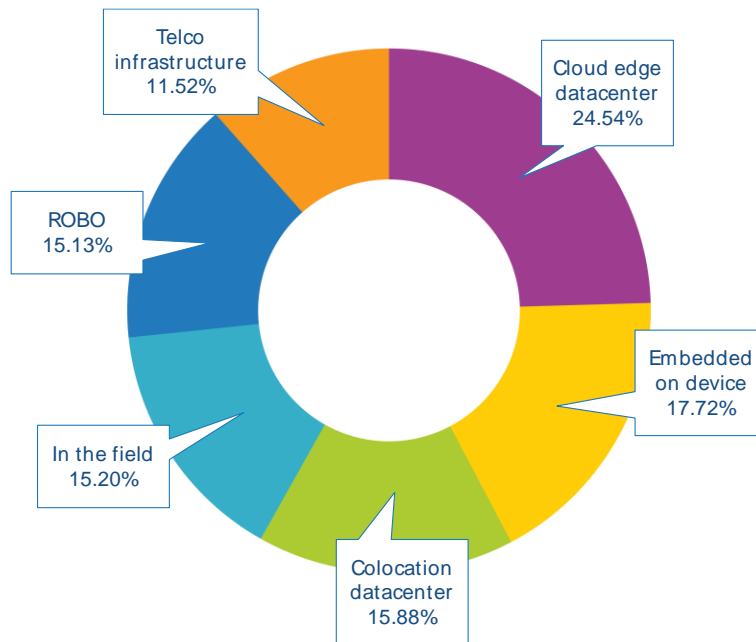
- **Industry field locations:** These are typically operations-oriented facilities such as factories, warehouses, and retail stores that may require ruggedization of equipment.
- **On device:** IoT-enabled equipment often has embedded compute capabilities that can process data and orchestrate a resulting action.

Figure 3 supports the notion that edge is anything but one size fits all and highlights the importance of the open hybrid cloud concept. No proprietary platform can support the unique needs of creating a holistic infrastructure that spans this range of deployment types, and the prospect of integrating multiple systems together can be costly and time consuming.

FIGURE 3

Edge Infrastructure Deployment Trends

Q. *Where is your organization's edge infrastructure deployed?*



n = 637

Source: IDC, 2021

IDC believes that success at the edge requires the utmost flexibility while maintaining a common control plane that unifies management of the overall life cycle of the solution. Given the remote nature of these systems, edge often requires a longer support life than infrastructure in the datacenter.

There is also a need to support multiple application types. While most organizations plan on cloud-native development practices for new applications, they must coexist with traditional and virtualized ones as well.

While keeping data onsite can be a benefit, edge locations can often have less physical security controls. This limitation combined with the need for remote management brings unique challenges. It is a much different workflow managing 1,000 servers in one location versus one server in 1,000 locations.

This is why automation is so important for scalability in large fleet deployments, including down to the OS level. It can often be an overlooked component of the solution, but being intentional about selecting an OS that is designed to support the complexities of the edge is critical.

CHALLENGES AND OPPORTUNITIES

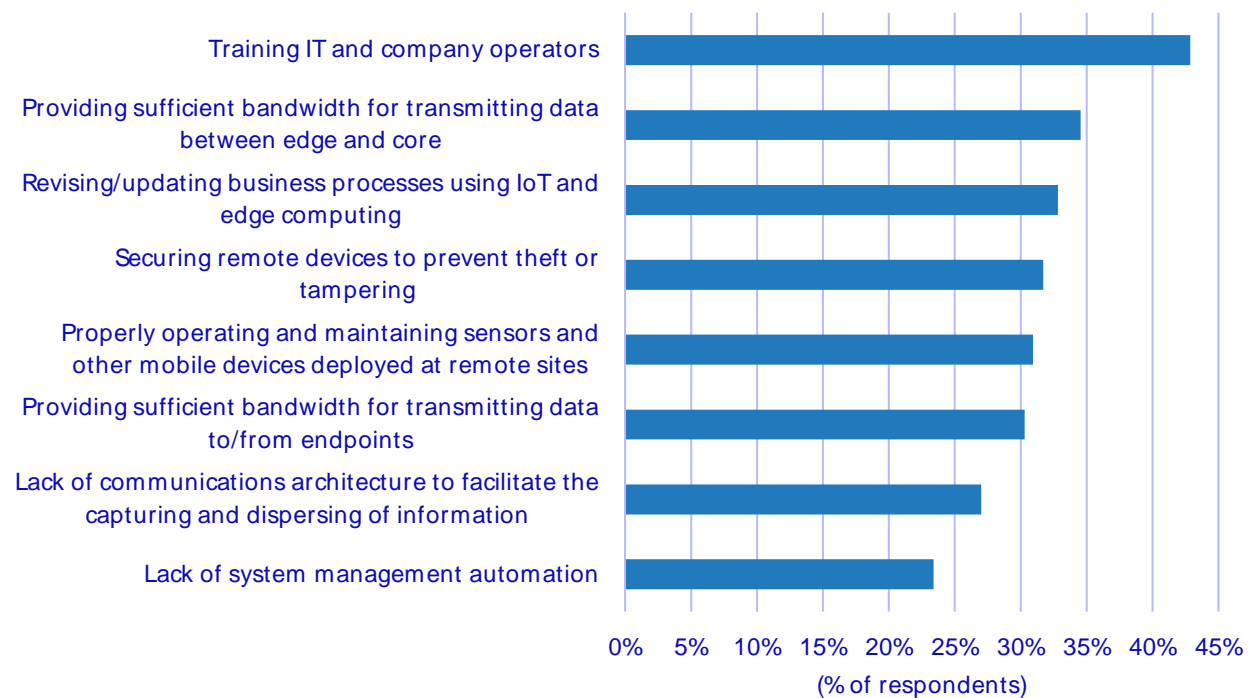
Despite the potential benefits and alignment with business objectives, organizations are still challenged in deploying an edge strategy. For some, provisioning the correct amount of connectivity and bandwidth is the issue. For others, it is the burden of updating processes and procedures to stay consistent with corporate standards (see Figure 4).

However, the most identified reason is a lack of training. This assumes that edge systems are managed separately from other IT infrastructure, requiring a different set of tools and knowledge in how to use them.

FIGURE 4

Edge Strategy Challenges

Q. What challenges is your organization facing with running and/or implementing an edge strategy?



Source: IDC, 2021

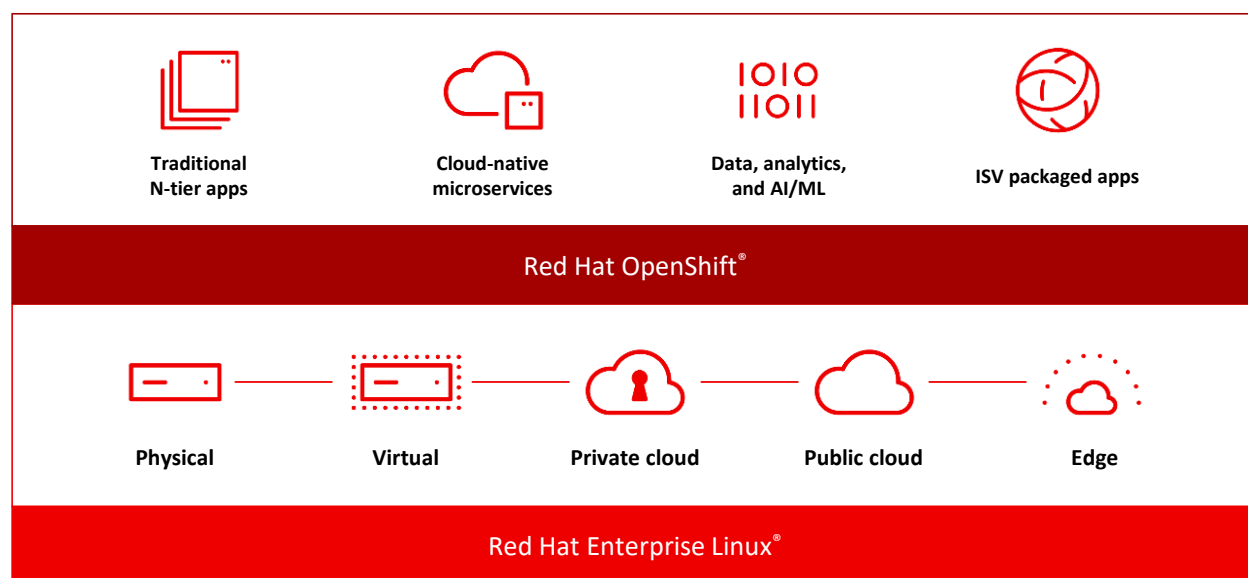
This is the opportunity for the open hybrid cloud where management, security, and automation are consistent regardless of deployment location, size, or workload.

RED HAT EDGE

Red Hat is addressing edge as part of its open hybrid cloud strategy with a foundation built on Red Hat Enterprise Linux. The company has chosen a portfolio approach to edge solutions leveraging familiar tools and patterns, making it easier to deploy and manage (see Figure 5).

FIGURE 5

Red Hat Open Hybrid Cloud Platform



Source: Red Hat, 2021

Instead of creating an edge-specific OS, Red Hat has built edge functionality directly into Red Hat Enterprise Linux that includes:

- **Quick image generation:** Easily create purpose-built OS images optimized for the architectural challenges inherent at edge.
- **Small footprint edge OS:** Scale down the size of Red Hat Enterprise Linux to support memory-constrained edge servers and IoT gateways.
- **Remote device update mirroring:** Staged and applied image updates occur at the next reboot or power cycle, ensuring minimal downtime.
- **Efficient over-the-air updates:** Updates transfer significantly less data and are ideal for remote sites with limited or intermittent connectivity.
- **Intelligent rollbacks:** Application-specific health checks detect conflicts and automatically revert an OS update, preventing downtime.

In addition, Red Hat Enterprise Linux has features that are important regardless of deployment location. Red Hat recognizes that enterprises deploying open source software in production need assurances that the platform has clearly defined security and support over the life cycle of the OS. Red Hat Enterprise Linux delivers 10 years of support, which can be continued even further with an

extended life-cycle support contract. Red Hat Enterprise Linux subscription licensing provides full access to future releases while letting the customer choose when to upgrade.

Red Hat Enterprise Linux bridges the gap between traditional and cloud-native workloads. This protects investments in existing applications while enabling the move to containers. Containers enable better portability, isolation, and workload flexibility on static edge servers and devices. Cloud-native containers facilitate life-cycle independence between the application and underlying OS.

The Red Hat Edge portfolio also includes Red Hat OpenShift – an enterprise orchestration platform that extends Kubernetes capabilities to edge sites with consistent operations and application life-cycle management at scale using flexible topology options to support low-latency workloads. Deployment models include single-node edge servers with combined control and worker nodes, distributed configurations with remote worker nodes, and three-node clusters that offer high availability in a small footprint. The ability for Red Hat Enterprise Linux to be deployed as a container provides applications with the same level of stability and security as a base image. This results in a smoother collaboration across operations and developments teams as they scale edge deployments on top of OpenShift.

Considering no one vendor provides all of the technologies needed for modern IT operations, it can be challenging to find the right hardware, software, and cloud platforms; ensure interoperability and stability between those components; and optimize performance across multivendor environments. Red Hat partners with leading hardware, software, and cloud vendors to provide more choice, innovation, and stability.

Through testing, certification, and collaborative engineering, Red Hat works with partners to deliver multivendor technologies and solutions that meet an organization's requirements for innovation, availability, performance, and security. Red Hat fosters a large and certified partner ecosystem – including partners such as Microsoft, SAP, and major cloud providers – so you can choose the products and platforms that best fit your needs, knowing they will work reliably with Red Hat Enterprise Linux.

CONCLUSION

Organizations are looking to digital transformation as a method of increasing agility, creating differentiation, and achieving sustainable competitive advantage. The concept of an open hybrid cloud, which provides the freedom to innovate across datacenter, edge, and cloud infrastructure, is an essential element of this strategy.

Red Hat Enterprise Linux provides organizations with an OS that not only handles the complexity that can come with edge computing but also serves as a key building block in Red Hat's portfolio. Red Hat Enterprise Linux helps organizations reduce operational costs while preserving stability and interoperability with core systems, allowing them to maximize the value of their in-house skills.

Red Hat is giving its customers the foundation needed to implement the future of digital infrastructure with the consistency, security, and flexibility to scale required to process large amounts of data that drive real-time decisions.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2021 IDC. Reproduction without written permission is completely forbidden.

