




Conquering the Edge: Strategies for Kubernetes at the Edge and Beyond

A Guide to Secure, Scalable,
and Efficient Deployments



Table of Contents

| | |
|---|-----------|
| Unlocking the Potential of the Edge: What You'll Gain From This Guide | 3 |
| Beyond the Cloud: Computing on the Edge | 4 |
| An Edge Computing Revolution | 5 |
| The Containerization Advantage | 6 |
| Fundamentals of a Kubernetes Container Orchestration System | 7 |
| Benefits of Kubernetes in Data Centers and the Cloud | 8 |
| Emerging Trends in Kubernetes Adoption | 9 |
| Is Kubernetes the Missing Piece in Edge Computing? | 10 |
| Key Considerations for Kubernetes at the Edge | 11 |
| Practical Strategies and Solutions to Succeed at the Edge | 15 |
| How ZEDEDATA Simplifies Edge Orchestration and Management | 15 |
| ZEDEDATA Edge Kubernetes Service | 17 |
| ZEDEDATA Edge Application Services | 19 |
| Conclusion | 21 |
| Key Takeaways | 22 |
| Important Recommendations | 22 |



Unlocking the Potential of the Edge: What You'll Gain From This Guide

As edge computing rapidly evolves, its inherent complexities can challenge even the most seasoned architect. Decentralized deployments, constrained resources, and diverse device landscapes demand a new approach to edge orchestration and management. For many, container orchestrator Kubernetes is believed to be a first line of defense, particularly following success in data center and cloud environments. However, attempting to implement Kubernetes with the same approach at the edge is often characterized by costly trial and error, a necessity for numerous vendor relationships, and stop-gap measures, among other concerns.

This guide isn't your typical Kubernetes primer. We'll dissect the unique challenges of the edge itself, explore ways to determine if Kubernetes at the edge is the right choice for your application, unveil proven edge computing solutions, and equip you with proven strategies to conquer common challenges.


Beyond the Cloud: Computing on the Edge

The digital landscape is awash with data, a torrent generated by a multitude of interconnected devices, machines, sensors, and applications. This data deluge is particularly pronounced at the edge, where countless devices reside and interactions occur. According to International Data Corporation (IDC), over half of new enterprise IT infrastructure was deployed at the edge in 2023. And the sheer volume of devices connected to IoT systems is expected to reach 55.44 billion and generate roughly 73 zettabytes of data by 2025.

Traditionally, this vast trove of data would be hauled back to centralized cloud servers for processing. However, this approach is fraught with challenges at the edge. The operational costs associated with transferring massive amounts of data to the cloud are unsustainable, especially considering the potential short half-life of the data. Additionally, modern networks are designed asymmetrically, prioritizing the rapid download of data to end-user devices, not for uploading large volumes of data back to the cloud.

In many solutions, the only viable solution to unlock the value inherent in this edge data is to process it locally, at the source, where it's needed. Moving data back to the cloud for processing defeats the purpose of collecting it in the first place, and can render it obsolete by the time it reaches the cloud. To effectively harness this data, edge computing, a centralized approach to managing and orchestrating nodes, security, and applications, is essential.

In essence, edge computing represents a paradigm shift in the IT landscape, moving computation closer to the sources of data. This distributed approach enables real-time insights, faster response times, and improved decision-making, paving the way for a more efficient, responsive, and data-driven future.



An Edge Computing Revolution

According to Gartner, edge computing is part of a distributed computing topology where information processing is located close to the edge — the physical location where things and people connect with the networked digital world.

Once a nascent technology, edge computing has become a cornerstone of data processing and decision-making across diverse industries. From retail and hospitality to consumer electronics, manufacturing, renewable energy, and oil & gas, the landscape has experienced a quantum leap. The widespread adoption of smart devices, connected equipment, sensors, and applications has propelled a wave of data moving away from traditional cloud environments and data centers, closer to their source: **the edge**.

This exponential growth has ignited the need for a holistic edge computing solution. Such a solution would not only bring processing power closer to where data originates, enabling faster analysis and more informed decisions but also eliminate the cost burden associated with transferring massive data volumes to the cloud. This convergence of efficiency and cost-effectiveness paves the way for a new era of data-driven decision-making, empowering organizations to unlock unprecedented opportunities for innovation and customer experience.

The Containerization Advantage at the Edge

Containerization has long been a solution for organizations looking to seamlessly run multiple isolated applications on a device at the edge, and this deployment method is on a rapid incline. Gartner predicts that 80% of custom software running at the physical edge will be deployed in containers by 2028, a major increase from 10% in 2023. And a recent Cloud Native Computing Foundation (CNCF) survey indicated that 93% of organizations are employing or intend to employ containers in production, reinforcing the trend towards containerization in data centers, the cloud, and more recently at the edge.


The popularity of containerization is well earned, due in large part to revolutionizing the way applications are packaged. Containers enable businesses to bundle the various components of their applications with precise requirements and seamless execution. But managing a fleet of containers, often consisting of microservices, can become unwieldy. This is where Kubernetes has typically been introduced, particularly in data center and cloud environments.



of custom software running at the physical edge will be deployed in containers by 2028



of organizations are employing or intend to employ containers in production





Fundamentals of a Kubernetes Container Orchestration System

Kubernetes, or K8s for short, is an open source system that enables businesses to automate the deployment, scale, and management of containerized applications. Think of it as the conductor of a complex orchestra, where each container is a talented musician.

Kubernetes groups these containers into logical units, making them easier to manage and discover. This is like organizing the orchestra by sections, like strings, brass, and percussion. But unlike a traditional conductor, Kubernetes was crafted by Google's years of experience running production workloads, combined with the best practices of the broader tech community.

Kubernetes is portable, meaning it can run on any platform, and extensible, so it can be customized to specific business needs. It's also open source, which is one of the most notable characteristics of its appeal, so anyone from its massive and rapidly growing community can contribute to its ongoing evolution. As a result of this popularity, businesses of every type can find a wealth of services, support, and tools readily available to assist with Kubernetes implementation.

Benefits of Kubernetes in Data Centers and the Cloud

Since its debut in 2014, Kubernetes has been embraced by both open source developers and large enterprises, who have primarily used it to deploy applications within data centers or cloud infrastructure.

In these traditional environments, Kubernetes provides:



Agile scalability: Organizations can spin deployments up or down based on real-time demand, ensuring applications flex with the ebb and flow of usage.



Storage management: Kubernetes can easily handle the orchestration of both local and cloud-based storage solutions.



Service diversity: Kubernetes' open community enables access to a broad array of services available through DNS or IP addresses, eliminating the need for manual configuration and streamlining service interactions.



Container versioning: Kubernetes users can control which container versions they run, easily roll out new images or container resources, and replace outdated versions.



Load stabilization: Kubernetes orchestrates network traffic, ensuring seamless operation even during peak hours or under heavy loads.



Built-in security: Organizations can easily and securely update passwords, OAuth tokens, and SSH keys associated with specific container images to keep operating systems shielded from vulnerabilities.




Emerging Trends in Kubernetes Adoption

Now in use for more than 10 years, Kubernetes has become the ubiquitous open source platform for container orchestration and management, establishing itself as the de facto standard for deploying, maintaining, and scaling containerized applications in modern IT landscapes and across diverse environments. This widespread adoption is underscored by the CNCF Annual Survey 2022, which revealed that a record-high 96% of organizations are either using or evaluating Kubernetes, marking a significant leap from 83% in 2020 and 78% in 2019.

These numbers align with CNCF's estimation that more than 5.6 million in the developer community are currently using Kubernetes to build their applications.

Even with all its popularity, however, Gartner acknowledges that Kubernetes at the edge will necessitate the development of new tools and operational procedures, particularly when managing large fleets of clusters operating on resource-constrained hardware with limited local support. These challenges highlight the need for innovative solutions that address the unique demands of edge computing environments.



Is Kubernetes the Missing Piece in Edge Computing?

With the meteoric rise of containerization and a vast majority of new applications being containerized with Kubernetes in both cloud and centralized data center architectures, it's only logical that Kubernetes would emerge as the ideal solution for the distributed edge.

Leveraging Kubernetes for edge computing offers several compelling benefits, both from a technical and business standpoint. From a technical perspective, Kubernetes excels in orchestrating workloads across distributed systems within a cluster, making it well-suited for environments with multiple nodes or servers. However, adapting Kubernetes for use in geographically dispersed edge locations requires modifications to manage network constraints, reliability issues, heterogeneous hardware, and limited resources typical of edge computing.


From a business standpoint, adopting Kubernetes at the edge offers a modern way to run modern software and containerized applications at the edge, accelerating time to value and innovation potential. Because Kubernetes has a vast and supportive community, this translates into long-term savings on development efforts for common features, ensuring the project's ongoing maintenance and security. Additionally, the widespread familiarity with Kubernetes among architects and developers simplifies the hiring and onboarding process.

If deployed and managed strategically, Kubernetes can play a critical role in unlocking the full potential of edge computing, harnessing both its technical capabilities and its distinct business advantages. But first operational technology and IT teams must be aware of and ideally resolve challenges that exist at the edge before attempting to implement Kubernetes. Second, they must be realistic about compromises that must be made to account for Kubernetes at the edge, as opposed to the familiar cloud or data center.

What is the distributed edge?

The "distributed edge" represents a sea change in computing, decentralizing processing power from remote data centers to a vast network closer to data sources. Unlike traditional cloud computing, where long distances between devices and central facilities result in latency, bandwidth strain, and single-point vulnerabilities, the distributed edge empowers a diverse network of local devices, edge servers, and small data centers to handle data on the spot.

This decentralized approach unlocks significant benefits, including faster response times for real-time applications like autonomous vehicles and reduced network congestion by processing data closer to its origin. However, there are inherent challenges at the distributed edge, including the homogeneity of hardware and software, limited computing footprint, geographic distribution, and expanded security needs, as well as the challenge of extending cloud-native development principles to the edge, integrating legacy systems and software, and the addressing the unique needs of operational technology systems.



Key Considerations for Kubernetes at the Edge

While processing and analyzing data via cloud-native applications and centralized data centers offers a familiar landscape, deploying and managing applications at the edge **presents a unique set of complexities and challenges.**

Kubernetes was designed for scale and flexibility. In the cloud, Kubernetes enables running thousands of containers in a cluster. For instance, an architect can easily manage three to five one-thousand-node clusters in the cloud. The edge, however, presents a different scenario. At the edge an architect might need to manage thousands of three- to five-node clusters, but there aren't any current tools that can manage this scenario.

Two options for managing Kubernetes at the edge include maintaining a manageable number of clusters and running multiple instances of a container orchestration platform, or implementing Kubernetes workflows in a non-Kubernetes environment like **EVE-OS**. The first option is ideal for users who intend to leverage core Kubernetes capabilities or are looking to manage a large number of containers at a limited number of sites. The second option may require a significant investment, but it takes advantage of existing Kubernetes-based workflows without implementing the Kubernetes runtime on the edge node. This option works well if you have a manageable number of containers, but if you're looking to manage more than 10, it would be advisable to go back to Kubernetes and rethink the

Key Considerations for Kubernetes at the Edge, Continued

Following are the most common challenges at the edge, some practical strategies to overcome them, and the considerations that must be explored before deploying Kubernetes an edge environment:

→ **Hardware, software, and infrastructure heterogeneity** – The ever-expanding diversity of workloads, combined with the sheer number of disparate systems, hardware, and software providers across today's distributed edge applications, creates a significant challenge: ensuring technology and resource compatibility while meeting desired performance standards. This complex ecosystem necessitates a strategic approach to edge infrastructure.

To address this head-on, consider adopting hyperconverged or hypervisor-based platforms. These platforms offer the agility and flexibility required to handle diverse workloads efficiently. Additionally, deploying open source solutions like EVE-OS unlocks a critical advantage: freedom from vendor lock-in and seamless interoperability across the entire open edge ecosystem. By embracing open source technologies, organizations can truly unlock the full potential of distributed edge computing.

→ **Smaller edge node footprints** – While off-the-shelf Kubernetes closely aligns with cloud principles like scalability, elasticity, flexibility, and resilience, Kubernetes deployment at the edge is problematic due to the inherently constrained nature of edge devices, which are limited in size, weight, power, and subsequently computing footprint and resources. This necessitates a tailored approach to Kubernetes at the edge. The cornerstone of this strategy lies in choosing the optimal Kubernetes distribution to align with your specific hardware and deployment needs.

While compact, open source container management solutions like K3s, Tanzu Tiny Stack, and OpenShift are attractive due to their low resource footprint, they fall short in addressing critical edge requirements like data sharing, communication, system interoperability, and elastic scaling.

Therefore, a more comprehensive approach requires partnering with an orchestration vendor or platform offering the flexibility to support a diverse range of Kubernetes distributions, including K3s, K8s, KubeEdge, MicroK8s, and others.

K3s - Created by Rancher Labs, K3s is a lightweight, certified Kubernetes distribution designed for production workloads across resource-restrained, remote locations or on IoT devices.

K8s - K8s is an abbreviation of "Kubernetes," which comes from counting the eight letters between the "K" and "S" in Kubernetes.

KubeEdge - KubeEdge is an open source system for extending native containerized application orchestration capabilities to hosts at the edge. Built upon Kubernetes, it provides fundamental infrastructure support for network, application deployment, and metadata synchronization between the cloud and the edge.

MicroK8s - MicroK8s is an open source system for automated deployment, scaling, and management of containerized applications. It provides the functionality of core Kubernetes components, in a small footprint, scalable from a single node to a high-availability production cluster.



Key Considerations for Kubernetes at the Edge, Continued

→ **Unreliable network connectivity and limited resources** – Unlike the predictable, always-on connectivity enjoyed by applications in the cloud or data centers, edge devices often operate in remote, harsh environments where network connectivity is unreliable or unavailable. What's more, many edge deployments are air-gapped, meaning they are completely cut off from the internet and have no connectivity to the cloud. This poses a significant challenge, as dispatching an IT technician to a remote location for every unexpected outage becomes impractical and expensive. Traditional Kubernetes relies on a persistent connection to online repositories for communication and container images, and so many off-the-shelf versions of Kubernetes are a poor fit for remote, isolated environments and air-gapped deployments.

For edge environments where continuous diagnostics, predictive maintenance, and outage management are critical, a centrally managed orchestration and management solution becomes indispensable. Such a solution ensures continued operation even during prolonged network and device outages, delivers reliable updates, and supports diverse communication hardware like satellite, LTE, and 5G. This not only ensures business continuity but also drastically reduces the need for costly field intervention, making edge deployments truly viable.

Key Considerations for Kubernetes at the Edge, Continued

→ **Data safety and security** – As more devices have moved away from tightly controlled centralized data centers and the cloud in favor of the edge, the security landscape has changed drastically and ushered in a new wave of security threats. Physical access to both devices and the data they hold has also introduced significant, new vulnerabilities that organizations must tackle to keep customer information safe. Lacking the security of data centers, edge devices in real-world public spaces like retail floors and factories introduce numerous vulnerabilities, including snatch-and-grab theft and the potential acquisition of vulnerable, proprietary data, as well as the potential for business interruption.

To counter this, robust security measures must extend beyond Kubernetes containers to include infrastructure-level protections for both hardware and software on these devices. Open source solutions, like EVE-OS, an infrastructure solution specifically designed for the distributed edge, address these concerns head-on. It safeguards against software and firmware attacks in the field, locks down physical ports and intellectual property on edge devices in case of theft, ensures security and consistency even with unreliable network connections, and facilitates large-scale application deployment and updates with limited bandwidth.

→ **Integrating legacy systems and software** – Organizations venturing into Kubernetes at the edge often worry about safeguarding the existing workloads they created in the cloud. Some believe they need to modernize all of their assets, hardware, and applications before implementing Kubernetes, while others attempt to rewrite applications and update technologies, leading to significant time investment.

A more suitable solution would be to partner with an edge Kubernetes service provider capable of supporting legacy containers and applications while providing a clear path to modernizing edge deployments along the way.

As you can see, Kubernetes at the edge is not a one-to-one relationship to Kubernetes in a data center or the cloud. A nuanced approach is required to secure and manage Kubernetes deployments at the edge. Recognizing that these challenges exist before embarking on Kubernetes or any other container deployment at the edge is the key to unlocking success. By proactively addressing these complex concerns, your organization can pave the way to securely and efficiently running containers at the edge.





Practical Strategies and Solutions to Succeed at the Edge

How ZEDEDATA Simplifies Edge Orchestration and Management

As the leader in orchestration and management at the distributed edge and a pioneer in the edge computing space since 2016, ZEDEDATA's solutions simplify edge deployment and management by making it open, secure, and effortless to manage. By leveraging cloud-native principles for configuration, management, and security, ZEDEDATA seamlessly replicates the agility and scalability customers have come to expect in the cloud at the edge.

How does ZEDEDATA help its customers achieve this straightforward simplicity?

By centralizing edge management and orchestration, ZEDEDATA can ensure consistent operations across edge nodes, enhance visibility, security, and control over distributed edge infrastructure and applications, and ultimately reduce costs associated with trial and error and fragmentary solutions.

Our holistic approach provides customers with the most flexible, universal, and open solution for their distributed edge orchestration needs, made even more powerful through our integrations with leading technology partners.

CORE PRINCIPLES

ZEDEDA's solution is architected with three core principles in mind:



Zero Limits — As edge computing grows rapidly, orchestration solutions must scale to manage vast fleets of edge compute nodes. ZEDEDA's future-proof architecture seamlessly handles geographically dispersed edge deployments, providing remote visibility and control at scale.

Leveraging the open source EVE-OS as a bare-metal foundation, ZEDEDA liberates architects from vendor lock-in, enabling them to build solutions on a diverse mix of hardware and software. This flexibility extends to supporting any combination of Docker containers, Kubernetes clusters, and virtual machines on any edge node while connecting to any cloud or server configuration.

ZEDEDA empowers organizations to transition from legacy applications to modern cloud-native development while preserving customer options for evolving needs.



Zero Touch — ZEDEDA simplifies distributed edge computing deployments, empowering field teams to manage hardware and applications anywhere, without on-site IT expertise. Our built-in app marketplace provides unparalleled out-of-the-box access from custom innovations to leading partner offerings. Additionally, customers can white-label our cloud service and curate their own private app marketplace.

ZEDEDA empowers architects, developers, system integrators, and distributors to effortlessly scale distributed edge computing projects for end customers. Our solution automates everything, from deploying hardware and installing software to managing the entire lifecycle, further simplifying all phases of edge computing and empowering end customers to succeed.



Zero Trust — ZEDEDA's Zero Trust security architecture offers a layered approach to safeguarding edge computing deployments, from silicon to cloud. This sophisticated framework leverages cutting-edge features such as hardware-based root of trust, measured boot, and remote attestation to ensure the integrity and provenance of edge devices. Data is encrypted both at rest and in motion, while I/O port blocking prevents unauthorized access and tampering in the field. A distributed firewall further enhances security by governing data flow from edge to cloud with granular, policy-based controls. This intelligent firewall segments field assets from upstream networks and data systems, creating an isolated micro-perimeter that safeguards critical operations. ZEDEDA's unique ability to deliver this comprehensive security architecture at the distributed edge sets it apart as a truly innovative solution.



ZEDEDA Edge Kubernetes Service

ZEDEDA Edge Kubernetes Service is an industry-first, fully managed Kubernetes service for the distributed edge to address the increasing desire to deploy Kubernetes to run containerized workloads and applications at the edge. This unique service includes a Kubernetes runtime that is curated, managed and supported by ZEDEDA, as well as integrations with industry-leading orchestrators.

ZEDEDA Edge Kubernetes Service is a managed Kubernetes infrastructure service that enables organizations to deploy, manage, and modernize their edge deployments effortlessly. This allows organizations to remotely deploy Kubernetes infrastructure at the distributed edge cost efficiently, with minimal configuration, and complete security from day one. ZEDEDA Edge Kubernetes Service makes it easy to deploy Kubernetes at the edge while providing organizations with a clear path to modernizing their edge deployments while still supporting legacy containers, virtual machines and applications. ZEDEDA Edge Kubernetes Service supports the customer modernization journey from legacy deployments, to application containerization, all the way to Kubernetes at the edge, providing a flexible future-proof service.

There is no one-size-fits-all solution for deploying and managing Kubernetes at the edge so ZEDEDA has partnered with industry-leading Kubernetes providers to ensure organizations are able to support existing investments and corporate standards across all deployments. ZEDEDA's partnerships and integrations with industry-leading orchestrators, such as IBM Edge Application Manager, Rafay, Red Hat OpenShift, SUSE Rancher, and VMware Tanzu provide a robust solution for the modern edge landscape.



[Learn More](#) about the benefits of ZEDEDA Edge Kubernetes Service.

FEATURED USE CASE

How ZEDED A is Helping One of the World's Largest Auto Manufacturers Manage the Largest Edge Kubernetes Project in the World

In July 2023, the global car manufacturer embarked on what has become the largest edge project in the world, specifically the deployment of new edge compute servers throughout all of its global dealerships. To assist in achieving its goals, the auto manufacturer selected ZEDEDA to modernize its edge infrastructure and provide a centralized, encrypted operating system, global orchestration platform, and software delivery system that would be compliant with UNECE R155 cybersecurity system management standards.

The ZEDEDA edge orchestration platform will give the automaker a single dashboard view into the status of all its deployments and health of all its devices. It will also enable the auto manufacturer to carry out remote maintenance while eliminating the need for additional manpower in the field.

ZEDEDA's centralized operating system, orchestration platform, and software delivery system will also help the automobile manufacturer solve several additional technical and logistical issues, namely:

- **Removing vendor lock-in** with open source technology;
- **Reducing the complexity of disparate ICT hardware** with hardware abstraction and virtualization; and
- **Eliminating threats from hostile environments** with zero trust security architecture and advanced firewalls, including inclusion detection and prevention.

ZEDEDA's orchestration offering also extends beyond the technology with ZEDEDA Edge Kubernetes Service, ZEDEDA's latest edge-as-a-service offering. Designed with ease of use and rapid responsiveness in mind, ZEDEDA Edge Kubernetes Service will enable ZEDEDA to eliminate troubleshooting ambiguities that come with employing multiple vendors, provide remote diagnostics to address problems in the field, and work directly with the automobile manufacturer's vendors to solve problems that were once out of reach.

[Read the full case study.](#)





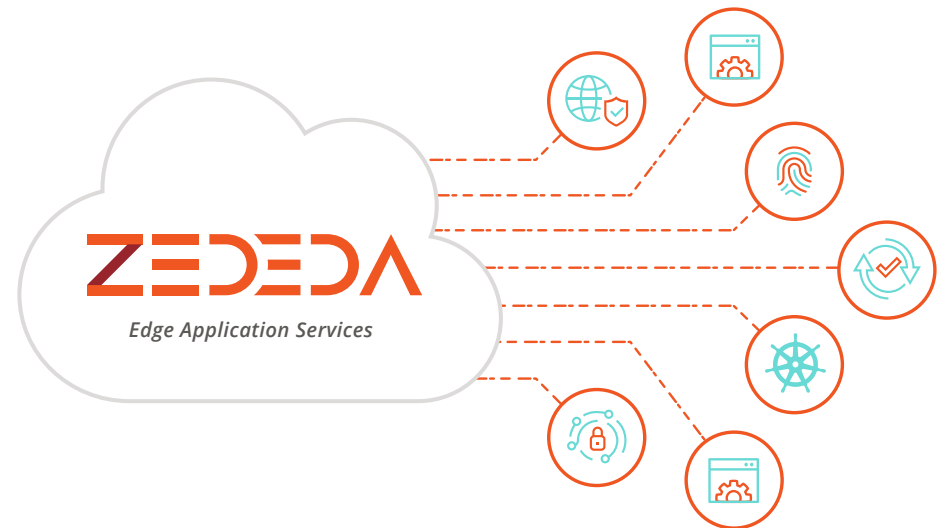
ZEDEDA Edge Application Services

ZEDEDA Edge Kubernetes Service is part of a growing number of ZEDEDA Edge Application Services that enable ZEDEDA customers to easily manage, configure, and control their edge applications within the robust ZEDEDA ecosystem. ZEDEDA Edge Application Services are distributed, cloud-native services designed to simplify the security and remote management of edge infrastructure and applications at scale.

ZEDEDA Edge Application Services are built on ZEDEDA's edge management and orchestration platform, which is delivered as a service and powered by EVE-OS.



Learn more about ZEDEDA's Edge Application Services suite, including [ZEDEDA Edge Access](#) and [ZEDEDA Edge Kubernetes](#).



Edge Development Pain Points: How EVE-OS can help

Edge computing has become crucial for transforming how organizations build and run software. But this evolving frontier comes with unique challenges for developers and architects unlike any encountered in the familiar landscape of data centers and the cloud. At the edge, curveballs like unsecured environments, hardware and software diversity, network connectivity woes, and difficulty managing edge infrastructure scalability, among others, have brought about the need for new tools capable of handling challenges that data center and cloud tools are ill-equipped to cure. EVE-OS offers an edge-native remedy that addresses the unique security, networking and manageability issues of the edge, while abstracting away hardware diversity.

Developed within the Linux Foundation's LF Edge consortium, EVE-OS is a lightweight fully open source, Linux-based operating system tailor-made for the distributed edge. Its flexible architecture [addresses common edge issues with the following solutions:](#)

- **Protection against firmware and software vulnerabilities** and attacks in the field, even in untrustworthy environments
- **Predictable and secure environmental conditions** for edge applications, even when network connections are unreliable
- **Deployment and software updates** across diverse edge devices at scale, even with limited or inconsistent bandwidth







Conclusion







Due to the growth of edge devices and data, there is a need for a comprehensive edge computing solution that moves data processing and analysis closer to endpoints and removes the cost associated with transferring large amounts of data to the cloud. For many, Kubernetes can provide a powerful platform for deploying, maintaining, and scaling applications at the edge. However, there are several complexities to consider when deploying Kubernetes clusters on edge infrastructure, such as the heterogeneous nature of hardware, software, and skill sets, limited computing footprint, geographic distribution, and unique security needs. It's important to be aware of these challenges before attempting to deploy Kubernetes at the edge and to seek support where needed to avoid the dreaded cost of trial-and-error and quick-fix solutions.


In addition to a proven edge management and orchestration platform, ZEVEDA Edge Kubernetes Service draws on numerous customer projects to address the challenges of deploying and managing Kubernetes at the edge. It provides a curated, managed, and supported Kubernetes runtime, as well as integrations with industry-leading orchestrators, making it easy to deploy Kubernetes at the edge while providing a clear path to modernizing future edge deployments.

Key Takeaways

-  **Edge computing is a rapidly evolving paradigm** that brings computation and data storage closer to the sources of data, enabling faster analysis, improved decision-making, and enhanced efficiency.
-  **Kubernetes, an open source container orchestration system**, has emerged as a powerful tool for managing containerized applications at the edge.
-  **Deploying Kubernetes at the edge necessitates addressing unique challenges**, including hardware and software heterogeneity, limited resources, unreliable network connectivity, and data security concerns.
-  **ZEDEDA Edge Kubernetes Service offers a comprehensive solution** for deploying, managing, and modernizing edge deployments, providing a simplified and secure approach to Kubernetes at the edge.

Important Recommendations

-  **Carefully evaluate the specific needs and requirements** of your edge computing environment before implementing Kubernetes.
-  **Consider adopting hyperconverged or hypervisor-based platforms** to address hardware and software heterogeneity at the edge.
-  **Utilize compact, open source container management solutions** like K3s, Tanzu Tiny Stack, or OpenShift to address resource constraints at the edge.
-  **Deploying open source solutions** like EVE-OS offers freedom from vendor lock-in and seamless interoperability across the entire open edge ecosystem.
-  **Implement a centrally managed orchestration and management solution**, like ZEDEDA, to ensure continuous operation and outage management in unreliable network conditions.
-  **Leverage ZEDEDA Edge Kubernetes Service** to simplify Kubernetes deployment, management, and modernization at the edge.



Resources for Additional Learning

[*ZEDEDA Edge Kubernetes Service: A Game Changer for Kubernetes Edge Deployments*](#)

[*Kubernetes on the Edge: Getting Started With Kubeedge and Kubernetes for Edge Computing*](#)

[*Kubernetes as a Service for the Distributed Edge*](#)

[*Linux Foundation LF Edge - Project EVE*](#)

[*Simplifying Kubernetes Deployments for the Distributed Edge*](#)

[*Seamlessly Deploy Kubernetes at the Edge*](#)

[*EVE-OS: An Open Operating System for the Edge With ZEDEDA*](#)

[*How to Avoid the Pitfalls of Deploying Kubernetes at the Edge*](#)

ZEDEDA



CONTACT@ZEDEDA.COM

About ZEDEDA

ZEDEDA makes edge computing effortless, open, and intrinsically secure – extending the cloud experience to the edge. ZEDEDA reduces the cost of managing and orchestrating distributed edge infrastructure and applications, while increasing visibility, security and control. ZEDEDA delivers a distributed, cloud-native edge management and orchestration solution, simplifying the security and remote management of edge infrastructure and applications at scale.

ZEDEDA ensures extensibility and flexibility by utilizing an open partner ecosystem with a robust app marketplace and leveraging an open architecture built on EVE-OS, from the Linux Foundation. EVE-OS is a lightweight, open-source Linux-based edge operating system. ZEDEDA delivers instant time to value, has thousands of nodes under management and is backed by world-class investors with teams in the United States, Germany and India. For more information, visit www.zededa.com.

