WHITEPAPER

Overcoming Network Connectivity Challenges at the Edge

Practical strategies and solutions for building a more secure, resilient edge network

ZEDEDA

Table of Contents

Mastering Network Instability: Overcoming Connectivity Issues at the Edge

At the ever-expanding edge, where data processing and decision-making have migrated closer to the physical devices generating the information, network instability presents an imposing challenge. Unpredictable connectivity, intermittent links, packet loss, jitter, and fluctuating bandwidth can wreak havoc on the edge's promise of real-time insights and localized control. To truly harness the power of edge computing, developing strategies that mitigate these connectivity issues is imperative. This necessitates a deep understanding of the unique network conditions prevalent at the edge, along with the implementation of robust techniques to ensure seamless operation even in the face of network volatility. Network architects now face a complex landscape of:

- **Unpredictable connectivity:** Edge locations often operate in environments with intermittent or bandwidth-constrained network links. This requires architectures that transcend reliance on continuous connectivity to the cloud
- **Heterogeneous devices:** The explosion of devices at the edge, spanning sensors, cameras, industrial equipment, traffic infrastructure, drones, and other user endpoints introduces a dizzying array of protocols and connectivity requirements, including wired, wireless, and satellite connections.
- **Disparate, remote locations:** Disparate and remote locations strain edge computing networks due to limited or unreliable connectivity options, often requiring diverse networking protocols to link devices to central resources. This diversity in location and connectivity needs increases management complexity and raises the risk of bottlenecks or failures that hinder the real-time functionality expected from edge computing.
- **Security complexity:** Distributed edge networks exponentially increase the attack surface, demanding robust security strategies that encompass device identity, authentication, and data protection at multiple points.
- **Data sovereignty and regulatory compliance:** Edge deployments may span geographical boundaries, necessitating a deep understanding of where data resides and how it flows, ensuring alignment with data privacy regulations like GDPR.

The traditional centralized network model, optimized for cloud and data center environments, buckles under the unique demands of the edge. Realizing the full benefits of the edge demands a fundamental shift in our approach to networking.

The Challenge of Inbound Connectivity at the Edge

Conventional networking approaches often rely on the ability to initiate inbound connections to devices, allowing for remote management, troubleshooting, or software updates. However, this model becomes increasingly complex and risky when extended to the edge due to the following complexities:

- **Complexity:** Most common deployments rely on IPSec for security, which introduces a networking environment that is hard to manage, difficult to troubleshoot, and hard to configure.
- **Security vulnerabilities:** Every open inbound port on a device becomes a potential attack vector. Edge locations, often with less stringent physical security than data centers, increase exposure to unauthorized access attempts.
- **Lack of static IP addresses:** Many edge devices reside behind cellular networks or consumer-grade internet connections where dynamic IP addresses are the norm. Managing these shifting addresses at the scale required for edge deployments can be an operational nightmare.
- **NAT traversal complexities:** Network address translation is commonly used to conserve IP addresses, but it further complicates establishing inbound connectivity to specific devices behind a NAT gateway.

The Case for Outbound-Only Connectivity

To mitigate these challenges, edge network architectures should adopt an outbound-only "phone home" model. An outbound-focused edge networking approach involves:

- **Edge-device initiated connections:** Devices establish secure, persistent HTTPS connections to a central management platform or cloud service
- **Management via tunneling:** Once this outbound tunnel is established, authorized administrators can leverage reverse tunneling techniques to access devices for troubleshooting or configuration changes, eliminating the need for inbound ports to be open.
- **Centralized control:** The outbound connection model provides a clear central control point for policy enforcement, monitoring, and security updates.

Outbound-only connectivity offers a powerful strategy for securing and managing edge networks. The following are key considerations to successfully implement this approach:

- **Robust authentication**: Strong mutual authentication is critical to maintain the integrity of outbound connections and to prevent rogue devices from infiltrating the network.
- **Resilient communication protocols:** Protocols designed for intermittent connectivity, such as MQTT, ensure messages are queued and delivered when links are restored, promoting seamless offline operation.

- **Device provisioning:** Implementing secure mechanisms to bootstrap initial device enrollment and configuration without inbound access requires careful planning.

By embracing the proactive, controlled paradigm of outbound-only connectivity, you can significantly reduce the attack surface of edge networks while retaining the ability to manage and maintain a diverse array of devices, and enabling zero touch.

Diverse Connectivity and Intelligent Traffic Routing

Resilience at the edge demands diverse connectivity options. Ethernet, LTE, 5G, and even satellite links can be combined intelligently to mitigate the risk of single-point failures. Edge architecture solutions should incorporate intelligent traffic management features like weighted uplinks, which prioritize certain network connections based on performance or cost, and/or application-aware routing, which steers specific applications or data traffic types over optimal paths . Assigning weightings to different uplinks enables primary/backup configurations or load balancing based on traffic volume or cost considerations. By defining separate network instances and route-specific application traffic to designated uplinks, application-aware routing prioritizes critical traffic over less sensitive data flows, ensuring quality of service and optimal resource utilization.

Diverse Local Connectivity: Spanning Shop Floors and the Purdue Model

Edge environments, particularly in industrial settings, demand a mix of local connectivity protocols to support the diverse needs of devices, a complexity best understood through the lens of the Purdue Model, a conceptual framework for industrial control systems where physical processes intertwine with digital commands. This conceptual model, born from the collaboration between industry and academia, serves as a roadmap for understanding the intricate network architecture of industrial control systems. Envision the Purdue Model as a multi-layered network offering a structured approach to design, security and optimization, with each layer serving a distinct purpose in the flow of information. Level 0 through Level 4 comprises the physical process, intelligent devices, control systems, manufacturing operations systems, and business logistics systems, respectively.

Industrial ethernet protocols like EtherCAT and PROFINET provide deterministic, real-time communication for sensors, actuators, and PLCs at the lower levels of the Purdue Model. For less time-sensitive data, mobility, and flexibility within a local network, wireless options like Wi-Fi and Bluetooth Low Energy offer unique advantages, including ease of deployment and integration, reduced latency, bandwidth flexibility, remote access and control, and scalability. As devices may bridge layers (e.g., a sensor gateway aggregating data for upstream analysis), secure and reliable communication protocols also become essential. Traffic segmentation strategies, which typically divide a network into smaller, isolated segments to control the flow of information, play a vital role in maintaining the integrity of the Purdue Model's security zones. Strategies like virtual local area networks, firewalls, intrusion detection and prevention systems, data diodes, and proxies, for example, can help organizations strengthen the security of their ICS environments, reduce the risk of cyber attacks, and protect critical infrastructure.

Why Unreliable Connectivity is a Critical Edge Issue

Traditional data centers operate with a high degree of network reliability and uptime. The edge environment is radically different. Edge devices and applications may reside in remote locations, on vehicles, or in locations where consistent connectivity cannot be guaranteed, posing unique challenges.

While data centers might experience brief data packet loss, edge devices could face disconnections lasting minutes, hours, or even days. This necessitates designing applications that gracefully handle these situations. Edge applications are often linked to critical physical processes or systems. In contrast to a service disruption in a data center, a malfunction at the edge due to connectivity loss could have severe real-world consequences (e.g., a factory line stoppage). Applications must be "always-on" despite network conditions. Edge locations might also rely on satellite, cellular, or other potentially expensive connections. Applications need intelligence to utilize these resources in a cost-efficient manner, prioritizing critical data and relying on backup connections only as needed. And even when disconnected, edge devices and their data require robust security. Local encryption and secure offline operation become essential considerations.

The next section, Flexible Network Configuration for Edge Resilience, addresses some of the strategies edge architects should employ in response to these unique connectivity challenges.

Flexible Network Configuration for Edge Resilience

Edge architects need appropriate tools to tailor network configurations for specific applications while addressing the challenges of unreliable connectivity.

Consider these key approaches:

- Segment your edge network into multiple instances for improved security, performance, and manageability, assigning instances to different applications or tiers as needed.
- Support redundancy and security by allowing applications to have both internal and public static IP addresses.
- To ensure resilience during network outages, build offline functionality into your edge applications, leveraging local caching and data synchronization mechanisms.
- Where available, implement cellular networks as a failover mechanism for your primary edge connections.
- Finally, explore network virtualization and edge computing platforms to effectively implement these strategies, always prioritizing security alongside your configuration changes.

Network Segmentation: Building a resilient, flexible edge network

When connectivity is intermittent and resources constrained, adaptability is key. Network segmentation can be a powerful tool for addressing these challenges head on. By dividing your edge network into distinct instances, each catering to specific applications, you can create a flexible architecture capable of weathering the storms of unreliable connectivity.

This approach offers a multitude of benefits:

- Enhances security by isolating sensitive applications and data, minimizing the impact of potential breeches.
- Enables tailored network configurations, ensuring that each application receives the optimal resources and performance it requires
- Simplifies network management by providing clear boundaries and control points, enabling faster troubleshooting and reducing downtime.

Just as a well-designed ship has separate compartments to prevent flooding from compromising the entire vessel, network segmentation creates self-contained units that protect your edge infrastructure from cascading failures. It's a strategic approach that not only strengthens your network's connectivity but also empowers you to harness the full potential of your edge deployments.

Edge Challenges and Mitigation Strategies for Intermittent Connectivity

Unreliable edge connections can wreak havoc on applications designed for constant connectivity, introducing networking hurdles that traditional data centers simply don't face. Applications that were designed with constant connectivity in mind may fail when the network becomes unreliable. Edge applications need to be robust enough to handle extended periods of offline operation, ensuring continuity of critical processes. Inconsistent data and device state further complicates the picture, as synchronization across edge devices and cloud services is hindered by connectivity loss. Information about device state or application updates may be delayed, leading to a view of the system that's not in real-time. Finally, intermittent connectivity can also pose security risks. Devices isolated by network disruptions become vulnerable and need to implement strong local security measures even when they cannot communicate with a central management system.

To combat these challenges, several key strategies are critical. Firstly, implementing a tiered connectivity system provides flexibility. This might involve using ethernet as the primary connection, LTE as a backup, and satellite for emergencies. With this approach, applications can intelligently select the most cost-effective and available option at any given time. Next, designing edge systems around the concept of eventual consistency is essential. This acknowledges that network disruptions can cause temporary delays in data updates and synchronization across devices and cloud services. Applications need to be built to tolerate some degree of "staleness" or inconsistency to maintain functionality. Furthermore, incorporating core offline functionality into applications ensures continuity of critical processes. This might involve caching data locally and queuing transactions to be sent when connectivity is restored. Finally, close monitoring of devices whose connection status is unknown becomes vital. Marking their health and state as "suspect" rather than simply offline provides a clearer picture of potential issues and allows for faster troubleshooting. By implementing these strategies, edge applications can become more resilient and adaptable in the face of unreliable network connections.

PVH: Optimizing Solar Farm Efficiency with Edge Management

PVH, a solar farm management company with a vast global network of installations, faced core challenges rooted in the remote and often unreliable nature of their edge environments. The company's solar farms located across 57 countries often rely on intermittent, low-bandwidth connections (such as satellite, LTE, and private networks), which greatly complicate remote software updates. Additionally, PVH's commitment to continuous innovation meant frequent software updates to deploy their latest solar panel optimization algorithms. With bandwidth being both a limited and expensive resource, they needed to minimize the amount of data transferred during these updates. Finally, ensuring safe deployment of updates was critical to prevent any accidental configuration changes that could disrupt solar farm operations.

To tackle these challenges, PVH leveraged ZEDEDA's edge orchestration and management platform. The platform's connectivity abstraction layer supports diverse connectivity types, which enabled PVH to deploy applications without having to tailor them to specific network conditions. Additionally ZEDEDA's expert guidance empowered PVH to streamline its application design, allowing for targeted software updates that maximized efficiency. A local operator console allowed PVH to manage and update its edge devices even during extended periods of offline operation. Importantly, a two-stage configuration lock process was implemented, minimizing the risk of errors and safeguarding the company's critical infrastructure.

As a result of these solutions, PVH now seamlessly deploys applications and updates regardless of location or connectivity status. Smaller updates translate directly to cost savings on their bandwidth-limited connections. The local console and configuration safeguards ensure reliable solar farm operation even in the most remote and disconnected environments. Most significantly, PVH can now accelerate their innovation cycle, rapidly testing and deploying their latest optimization algorithms to maximize the efficiency of their solar farms worldwide.

How ZEDEDA Meets Networking Needs at the Distributed Edge

ZEDEDA's architecture is designed to tackle the core challenges of edge computing: remote locations, unreliable connectivity, and bandwidth constraints. A key feature is its network-agnostic approach, allowing applications and updates to be deployed seamlessly across diverse network types (LTE, satellite, private networks, etc.), which eliminates the need to tailor software deployments to specific site conditions.

Furthermore, ZEDEDA empowers customers to manage edge devices and applications even when connectivity to a central system is lost. This is crucial in scenarios where intermittent connectivity is the norm. One way the ZEDEDA platform enables this is via the design of modular applications, composed of distinct, self-contained components, capable of receiving targeted updates. When an update is needed, only the specific component requiring a change is replaced, rather than the entire application, dramatically reducing the amount of data that needs to be transferred. This modular approach is particularly beneficial in low-bandwidth environments where large data transfers can be problematic or time-consuming.

Importantly, ZEDEDA balances flexibility with security. Customers can orchestrate their edge deployments from a centralized location while enforcing safeguards like two-stage configuration lock, which enables you to exclude an edge node from an automated configuration update and/or queue automatic updates rather than applying them across a fleet of edge nodes, and application snapshot and rollback, which enables you to create and save an application instance and then roll back to that instance as needed to recover from problems encountered during upgrades, data migration or normal operation. Both options minimize the risk of accidental changes that could disrupt operations across the edge infrastructure.

ZEDEDA's architecture and functionality is designed for high adaptability, making the platform suitable for a wide range of edge customers who demand resilient, efficient, and secure management of their distributed assets.

ZEDEDA's Architecture and Functionality for Secure Edge Connectivity

ZEDEDA takes a multi-layered approach to securing application networking and other forms of connectivity at the edge. At its core, the solution associates deployed applications with specific network ports and I/O resources. This restricts access by only allowing applications to utilize the ports and resources they are explicitly authorized for. All other ports remain disabled, reducing the potential attack surface. If a specific application requires special device drivers, which might not be fully trusted, the I/O device can be directly attached to the application itself. This eliminates the need to install potentially risky, unknown device drivers within the EVE-OS operating system.

ZEDEDA's Zero Trust security architecture further enhances security. With ZEDEDA, each application has a dedicated set of firewall rules and a "default deny" network configuration, which blocks all traffic by default unless explicitly permitted by a firewall rule assigned to an individual application–not just a single node or device. Additionally, each application has built-in visibility into potential security breaches through rule violation notifications.

Network segmentation, a core principle of Zero Trust security, is another powerful security feature offered by ZEDEDA. It allows applications to be connected to designated "network instances," providing granular control over network access. This means applications can be restricted to specific networks, such as a shop-floor network or an isolated "air gap" network, significantly enhancing security by preventing unauthorized access and lateral movement within the network.

ZEDEDA also integrates with public cloud providers' Virtual Private Cloud (VPC) offerings. By leveraging the cloud provider's IPsec configuration, a network instance can be designated as part of a VPC. This creates a secure, logically isolated environment within the broader network. The VPC implementation itself resides outside the deployed application, preventing the application from escaping the VPC or outside attackers from breaching its defenses.

This comprehensive security approach safeguards applications from network-based attacks, protects the underlying EVE-OS infrastructure from compromised applications, and even shields the broader network and other edge nodes should a security breach occur on a single device.

Air-Gapped Environments: Challenges and mitigation strategies

Air-gapped environments are highly secure networks deliberately isolated from any external connections, including the public internet or even corporate intranets. This extreme security measure is common in industries handling sensitive data, critical infrastructure, or manufacturing processes subject to strict regulations (like medical device production). Edge computing often necessitates operating within air-gapped environments due to the remote or insecure nature of edge locations.

Managing and updating applications and devices within an air-gapped system poses a unique challenge, as traditional methods like using USB drives to individually update edge nodes introduces security risks and logistical inefficiencies. Customers with multiple air-gapped plants typically require solutions that safeguard against potential system failures and inherent data loss. These solutions should include full lifecycle application management to deploy, restart, and control applications throughout their lifecycle, even when the plant has no connectivity, and rapid disaster recovery to replace devices quickly in the event of device failure without compromising the air-gapped network.

Local Network Management: A solution for air-gapped management

Local network management, with local operator consoles, can address air-gap challenges by bridging the gap between centralized cloud-based management and disconnected edge environments. Operating alongside an edge orchestration and management platform, a local management solution should include:

- **Centralized configuration**: An administrator updates a device configuration in the orchestration tool, even though the device is currently air-gapped.
- **Secure export**: The updated configuration is exported in an encrypted format that can only be decrypted by the designated physical device using its unique Trusted Platform Module (TPM) for guaranteed security.
- **Flexible transport**: The encrypted configuration can be transported to the air-gapped site in any secure way (USB, SSH, etc.). This mirrors physical update methods while adding strong encryption.
- **Local application**: Within the air-gapped environment, the local management container receives the configuration data. The target device automatically fetches the updates from the local console, just as it would from the cloud. Additionally, local configurations can be scaled to multiple edge nodes.

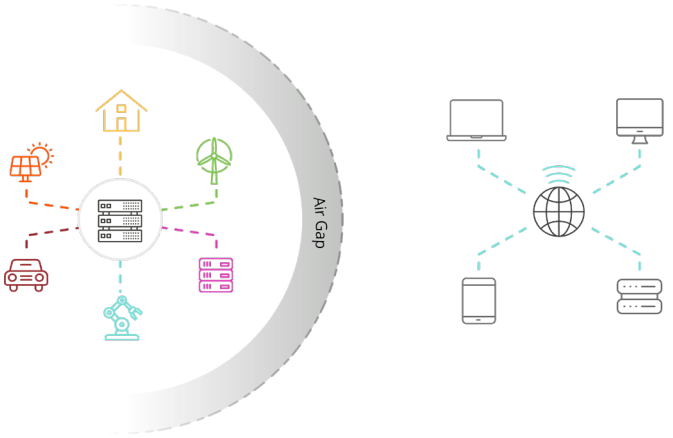

Use Case: Industrial Manufacturing with Air-Gapped Production Lines

Customer: Industrial manufacturer with a focus on regulated goods like medical devices, biological products, or other sensitive components.

Challenge: Manufacturing companies operate several plants with production lines that are completely isolated from external networks due to strict security and compliance requirements. Companies with these air-gapped production lines face unique challenges including application management and disaster recovery. Application management involves secure deployment, updating, and maintenance of applications on devices within the isolated network, all without compromising security or interrupting production. Disaster recovery protocols must be in place to allow for the rapid replacement of failed devices without risking the integrity of the air-gapped environment.

Solution: ZEDEDA Edge Sync, in combination with ZEDEDA's edge orchestration and management platform, provides a management and monitoring API that enables customers to build their own custom integrations, applications, and HMI interfaces leveraging ZEDEDA's API-based approach to enable:

Centralized management: All device configurations and application updates are managed from the central ZEDEDA orchestration platform, providing a single point of control even for air-gapped assets.

Secure configuration updates: When an update is needed, the configuration is exported in an encrypted format. Only the target device, using its TPM, can decrypt and apply the update. This ensures no unauthorized changes can be made.

Zero Connectivity: Designing for Offline Resilience

Prolonged or complete connectivity loss is inevitable in many edge scenarios. To ensure edge applications remain functional, the following offline capabilities should be core considerations in your architecture:

- **Local data caching and processing**: Edge applications require the ability to store data locally and continue processing it even when disconnected from the cloud. Smart caching strategies for constrained environments, like prioritization based on criticality, recency-based caching, least recently used (LRU) eviction, data expiration, and adaptive caching, as well as smart processing strategies such as data summarization, filtering and compression, and distributed processing are necessary, especially with constrained local storage on edge devices.
- **Offline functionality**: Design your applications with "offline-first" principles in mind. They should be able to provide core functionality even in the absence of connectivity. This may involve using locally cached data, or providing a degraded but still usable experience until the connection is restored, or designing applications with reduced online dependency.
- **Queue-based synchronization:** Implement a local storage queue to hold updates, data, and commands that occur during an offline period. Upon reconnection, this queue should be synced with the cloud systems using a reliable and efficient synchronization mechanism. Prioritize critical data over less time-sensitive updates to avoid unnecessary delays or bottlenecks when the connection is restored.
- **Conflict resolution:** Synchronization processes after offline periods could lead to data conflicts. Your application logic needs robust conflict resolution strategies. This could involve timestamps, user-based resolution, or application-specific rules (e.g., "last write wins") depending on the nature of your data.
- **Secure transport:** The encrypted configurations can be securely transported to the air-gapped production line using approved methods (USB, SSH, etc.).
- **Local deployment:** The local management solution, a lightweight container, is deployed within the air-gapped environment. Devices on the production line seamlessly pull updates from the local management tool, replicating the standard update process without requiring external connectivity.

Expected Results: Manufacturing customers can expect uninterrupted operations as applications are managed and devices updated seamlessly within their air-gapped production environment. Enhanced security is a core benefit, with strong encryption and TPM-based verification preventing unauthorized updates and safeguarding the air-gapped network. In the case of a disaster, customers can expect fast recovery with minimal downtime due to the ability to quickly provision and deploy replacement devices. Additionally, customers can rely on ZEDEDA Edge Sync to aid in maintaining compliance with strict industry regulations that mandate air-gapped environments.

Scheduled Downtime: Tailoring solutions for disconnection

While both scheduled downtime and offline/zero connectivity scenarios involve periods without service, the way we recommend approaching them and the solutions we implement differ significantly. Scheduled downtime offers the advantage of being planned in advance, allowing for:

- **User communication**: Informing users about the upcoming downtime and estimated duration helps manage expectations and minimize disruption. Scheduling maintenance during off-peak hours further reduces impact.
- **Controlled timing:** You have control over the timing of downtime, enabling coordination across dependent systems and ensuring maintenance doesn't clash with critical operations.
- **Focus on mitigation:** The primary goal is to minimize the impact of downtime. Techniques like rolling updates, blue/green deployments, and feature flagging, a technique used to control the gradual rollout of new features in an application, aim to provide continued service or controlled rollouts during the update window.

How Offline Scenarios Differ from Scheduled Downtime

Unlike scheduled downtime, intermittent offline scenarios are unpredictable and can potentially last longer than anticipated. Here, the focus shifts towards resilient applications and user experience. Designing applications that function even when disconnected is paramount. This involves incorporating strategies like offline data storage, queuing mechanisms to hold updates until reconnection, and synchronization procedures to ensure data consistency when connectivity resumes. During offline periods, users need to understand limitations and be informed about how the application handles disconnected states.

The core difference between scheduled downtime and intermittent offline scenarios lies in the type of planning and focus required. Scheduled downtime is proactive, allowing you to choose the timing and minimize disruption, while offline scenarios necessitate a reactive approach, requiring applications to be inherently resilient to unforeseen connectivity loss. Scheduled downtime solutions aim to mitigate the impact of a controlled event, while offline solutions focus on enabling functionality despite the lack of connection. Lastly, scheduled downtime typically has a predictable scope (affected users and services), while offline scenarios can have a wider and more variable impact, depending on the duration and affected systems.

ZEDEDA Edge Sync: Secure, seamless, local edge management and monitoring

ZEDEDA Edge Sync addresses a critical challenge in managing edge deployments: unreliable or completely absent cloud connectivity. It provides a robust and secure local management solution embedded within your own network through a management and monitoring API that enables customers to build their own custom integrations, applications, and HMI interfaces leveraging ZEDEDA's API-based approach.ZEDEDA Edge Sync is specifically designed to empower architects working with air-gapped environments, ultra-secure deployments, or projects where network connections are inconsistent due to the nature of the deployment site. Even when disconnected from the cloud, ZEDEDA Edge Sync ensures continuous monitoring and management of edge nodes, allowing for configuration updates and essential status checks.

ZEDEDA Edge Sync delivers a cost-effective approach by eliminating the need for specialized hardware. It operates on standard systems, reducing deployment complexity and lowering the total cost of ownership (TCO). It ensures a single source of truth for device configuration. When cloud connectivity is restored, ZEDEDA Edge Sync synchronizes with the cloud for seamless, up-to-date management. This approach guarantees operational continuity within the edge deployment, regardless of temporary internet outages.

For archiects currently managing air-gapped environments and contemplating a future cloud migration, ZEDEDA Edge Sync serves as a valuable bridge solution. It delivers secure local management capabilities while providing a controlled and risk-mitigated path to the cloud.

In essence, ZEDEDA Edge Sync is an indispensable tool for architects seeking a reliable network connectivity solution in the complex world of edge computing. From ultra-secure deployments to projects located in areas with limited internet access, ZEDEDA Edge Sync ensures the manageability and integrity of your edge nodes.

How ZEDEDA and ZEDEDA Edge Sync Protect Disconnected Edge Deployments

The edge of the network is characterized by many of the same challenges that can impact traditional cloud-based management approaches, not the least of which is a constant threat to security. When distributed deployments are likely to experience unreliable connectivity, planned outages, or be entirely air-gapped for security reasons, maintaining operational continuity, control, and robust security protocols is critical. ZEDEDA's orchestration and management platform already addresses the security of edge deployments through a multi-faceted approach based on Zero Trust principles, ensuring that no device or user is inherently trusted and every access request is verified, regardless of its origin. ZEDEDA's security practices include secure boot and measured boot methods to ensure software integrity, remote device attestation for identity verification, secure remote management, data encryption, policy-based access control, and threat detection and response. These measures collectively strengthen ZEDEDA's Zero Trust security model, protecting edge devices and data.

ZEDEDA Edge Sync further addresses network connectivity challenges by providing a local, secure management solution tailored for these unique edge environments.

Edge device security is of paramount importance for numerous reasons: 1) edge devices are often deployed in less secure locations, making them vulnerable targets for theft or tampering; 2) sensitive intellectual property, data or proprietary software may be stored on edge devices, necessitating robust protection; and 3) compromised edge devices can lead to service outages or malfunctions, directly impacting customers or critical operational processes.

ZEDEDA Edge Sync prioritizes security by incorporating several key features:

- **Air-gapped support**: Ideal for deployments requiring complete network isolation for maximum protection against external threats.
- **Zero trust architecture**: Minimizes the attack surface by adopting a "never trust, always verify" approach to authentication and authorization.
- **Seamless cloud integration**: Ensures that the cloud-based CUSTOMER controller remains the authoritative source for deployment configuration, reducing conflicts and maintaining alignment.
- **Security posture management**: Enables centralized tracking and management of security settings, patches, and vulnerabilities across the fleet.

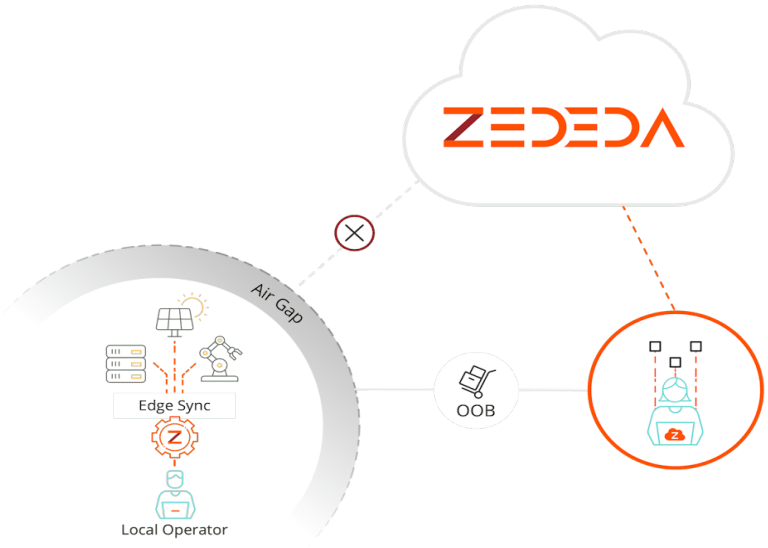

ZEDEDA Edge Sync: Key features and benefits

Centralized management with offline operation: Enables the centralized orchestration of devices even in air-gapped environments. Users interact with the ZEDEDA platform as normal, while local operator management ensures operational continuity regardless of connectivity.

Operational resilience in air-gapped environments: Enables deployment of new applications, updates, and configuration changes within air-gapped environments, guaranteeing uninterrupted critical operations.

Fleet management: Eliminates the need for multiple local management solutions on-site. A single operator console can manage and monitor an entire fleet of devices, streamlining operations.

Data synchronization: Designed to gather monitoring and health metrics from devices on-site. These insights can then be synchronized from the individual nodes with the central ZEDEDA orchestration and management platform when connectivity is restored.

Flexible data stores: Enables application and associated image storing, including (container images, virtual machine disks, etc.) even in air-gapped environments.

Evolutionary path to the cloud: Manages edge devices with the flexibility to easily transition to a cloud-based model in the future, mitigating risk and enabling a seamless evolutionary path.

Fully Customizable: ZEDEDA Edge Sync provides a management and monitoring API that enables customers to build their own custom integrations, applications, and HMI interfaces leveraging ZEDEDA's API-based approach.

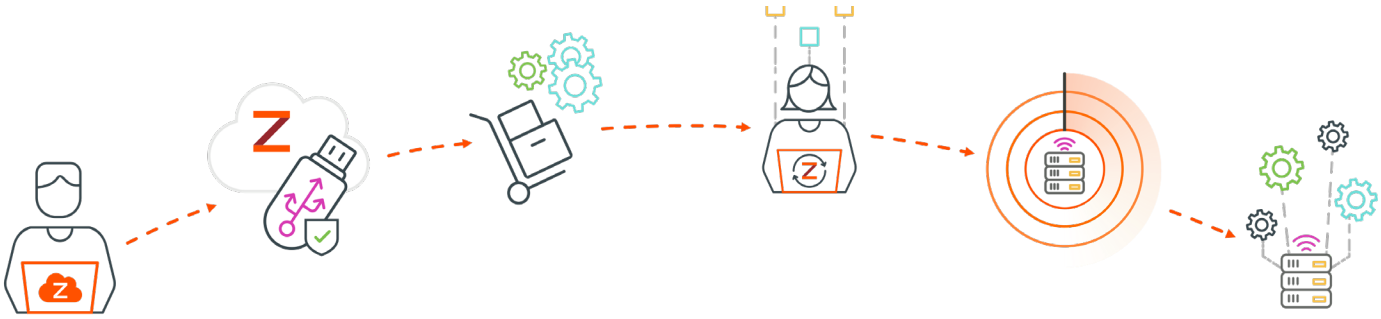

ZEDEDA Edge Sync Use Cases

ZEDEDA Edge Sync resolves three major cases:

- Planned intermittent connectivity loss to the ZEDEDA cloud controller
- Complete air-gapped solutions
- Scenarios where WAN connectivity interruptions occur

ZEDEDA Edge Sync enables fleet-level monitoring and configuration, ensuring service continuity and seamless operations. The local management solutions also ensures that the users can continue to manage their fleets, including configuration updates and monitoring, even when the connectivity to ZEDEDA is severed.

Scaling Connectivity Solutions for Thousands of Devices

When managing a single edge device, connectivity concerns are relatively straightforward. However, scaling solutions to encompass thousands of devices introduces new complexities in network traffic, data management, and operational costs. ZEDEDA experts offer the following key strategies to consider:

- **Minimize network traffic through polling intervals:** Instead of continuously pushing updates to each device, configure devices to "call home" and fetch necessary configuration changes. This reduces network strain significantly. ZEDEDA allows for customizable polling frequencies (e.g., hourly instead of every minute), striking a balance between rapid updates and conserving bandwidth.
- **Prioritize necessary data:** Not all metrics are equally crucial. Reduce network traffic and storage requirements by allowing granular control over which metrics are reported and how frequently. Prioritizing essential data helps streamline operations at scale.
- **Cost optimization**: With thousands of devices, the cost of connectivity becomes a major factor. Carefully evaluate the cost-benefit ratio of options like cellular SIM cards versus satellite connectivity, factoring in your specific geographical deployment and bandwidth requirements.
- **Management Platforms:** At scale, a robust centralized management platform becomes essential. Look for solutions that allow you to efficiently provision, update, monitor, and troubleshoot large fleets of devices from one location, streamlining your operations.

Additional Considerations for Large-Scale Deployments

Imagine deploying ZEDEDA to manage hundreds of sensors across multiple manufacturing facilities. By configuring devices to poll for software updates on an hourly basis and selectively reporting temperature and vibration metrics, you can minimize network traffic while ensuring critical data is available for real-time monitoring. This strategy, along with a well-designed management platform, lets you effectively operate a large-scale edge deployment. Consider the following:

- **Network architecture**: Assess whether existing network infrastructure can handle the additional traffic from thousands of devices. You may need to implement bandwidth upgrades or network segmentation strategies.
- **Data handling:** Develop solutions for storing, processing, and analyzing the increased volume of data generated by numerous edge devices. Cloud-based storage and analytics tools may be necessary.
- **Security:** Robust security measures are non-negotiable. Large deployments can present a broader attack surface. Implement strong encryption, access controls, and intrusion detection systems to maintain the integrity of your network.

It's important to remember that every deployment is unique. Consult with experts to develop a tailored scaling strategy that aligns with your business objectives and technical constraints.

Optimizing Edge Performance in Bandwidth-Constrained and High Latency Environments

Edge environments often operate under myriad unpredictable network conditions, including limited, or low, bandwidth. Low bandwidth presents a significant bottleneck for edge deployments, where devices often rely on constrained or unreliable network connections.

This limited data throughput hinders critical operations like software updates, remote management, and the transmission of telemetry data or AI model outputs. High latency, on the other hand, is all about delay in the delivery of data, which not only severely hampers the responsiveness of edge applications but also can disrupt real-time interactions or control flows often required at the edge. In scenarios requiring real-time decision-making, control loops, or time-sensitive interactions, latency introduces unacceptable delays that can degrade the performance of distributed systems or lead to incorrect actions based on stale data.

The combination of low bandwidth and high latency poses unique design challenges for edge architects. Traditional cloud-centric models, where data is constantly transmitted to a centralized location for processing, become inefficient or infeasible.

To mitigate these constraints and deliver optimal performance in these scenarios, edge architectures must prioritize:

- **Local data processing**: Deploying compute capabilities at the edge allows for pre-processing, filtering, and analysis of data directly on devices. This reduces the volume of data that needs to traverse the network, conserving bandwidth and mitigating latency effects.
- **Push vs. pull models:** Proactive, scheduled updates and management directives from a central server to edge devices can better utilize low-bandwidth scenarios compared to constant polling from the edge. This allows for intelligent batching of data during off-peak network periods.
- **Traffic prioritization and QoS:** Implementing Quality of Service mechanisms ensures that critical management traffic, real-time control signals, and time-sensitive data flows take precedence over lower-priority traffic.

Network Design Best Practices for the Edge: A solution designed for expected and unexpected disconnectedness

Network outages and air-gapping can wreak havoc on traditional solutions designed for constant connectivity. The very design of network solutions must prioritize offline functionality and secure local management. Tools must provide the ability to monitor, configure, and update edge deployments even when the connection to central cloud management is severed. This ensures uninterrupted service delivery and reduces the risk of downtime due to network disruptions. To successfully tackle these challenges, it's crucial to adopt specialized network design best practices that prioritize security, resilience, flexibility, and ease of management.

- **Prioritize security**: In a world where edge devices are often more vulnerable, security must be the bedrock of your network design. Embrace zero-trust models that assume no inherent trust within the network, implementing robust encryption, strict authentication, and continuous threat monitoring. Air-gapped environments, in particular, demand extra protection due to their physical isolation.
- **Plan for disconnection:** Acknowledge that intermittent connectivity or planned outages are a reality at the edge. Design your network with resilience in mind, incorporating redundancy, failover mechanisms, and local management tools like ZEDEDA Edge Sync. These tools ensure operational continuity and seamless management even when direct cloud connectivity is unavailable, without sacrificing security.
- **Flexibility and scalability**: Edge deployments evolve rapidly, demanding flexible network designs that can seamlessly adapt to changing requirements. Leverage software-defined networking, containerization, and zero-touch deployment solutions to achieve flexibility and scalability while keeping your TCO in check.
- **Ease of management:** Edge environments often lack specialized IT resources. Your network design should focus on simplicity, customization, and intuitive interfaces to make management accessible even without deep networking expertise. Solutions that emphasize ease of use are crucial for successful and efficient operations.

Solving Network Connectivity Challenges for Remotely Managed Devices

Remote devices, often integral to edge infrastructure, face a unique challenge: unreliable or spotty network connectivity that severely hampers essential communication, monitoring and management. ZEDEDA Edge Sync, along with ZEDEDA's edge orchestration and management platform, deliver a solution designed to address these specific network hurdles.

Firstly, ZEDEDA ensures devices maintain autonomous operation even if the connection to the central management system drops. Applications continue to function, and data is stored locally until connectivity is restored. In situations constrained by low bandwidth, ZEDEDA intelligently prioritizes critical updates and management traffic. It also cleverly compresses data and schedules transfers for times when the network is less congested. Additionally, ZEDEDA's support for various connectivity options (cellular, Wi-Fi, wired Ethernet) means devices can leverage whatever network is available, or whichever offers the most cost-effective connectivity option, maximizing their uptime. Finally, ZEDEDA streamlines the deployment of new devices through zero-touch provisioning, further reducing the need for on-site technicians even when connectivity is unreliable.

To elevate your remote device management strategy, consider these best practices:

- **Centralized management**: Opt for a management platform (cloud-based or on-premises) that provides a unified view of all remote devices, irrespective of their location.s
- **Security**: Prioritize robust security measures for your remote devices to defend against cyberattacks and unauthorized access. Enforce strong authentication, data encryption, and regular patching of vulnerabilities.
- **Monitoring**: Proactively monitor the health and performance of your remote devices. Employ tools that provide real-time alerts for issues ranging from hardware failures and security breaches to unexpected downtime.
- **Resilience**: Architect your remote deployment with redundancy at its core. Consider backup communication channels, failover systems, or distributed setups for mission-critical applications.
- **Documentation:** Maintain comprehensive documentation of your remote device infrastructure—network configurations, software versions, and maintenance logs are all crucial in streamlining troubleshooting.

By harnessing the power of ZEDEDA's capabilities and following these best practices, you'll conquer the challenges inherent to remote management, establishing a reliable, secure, and efficient edge computing environment.

Key Takeaways

- **Evaluate Existing Networks:** Analyze your current infrastructure's ability to handle increased traffic and identify potential areas where upgrades or network segmentation might be needed to support a larger edge rollout.
- **Assess Connectivity Options:** Explore the availability of diverse connectivity methods (Ethernet, LTE, 5G, satellite) at your edge locations. Determine if application-aware routing can be implemented for intelligent traffic prioritization.
- **Investigate 'Outbound-Only' Strategies:** Research how this approach might reduce the security risks associated with inbound traffic at the edge.
- **Design for Resilience:** Architect your applications with offline functionality, local caching, and queue-based synchronization in mind to handle intermittent connectivity.
- **Explore Air-Gapped Solutions:** If relevant to your use case, learn more about offerings like ZEDEDA Edge Sync designed for securely managing isolated networks.
- **Optimize Data Handling:** Research cloud-based and edge-based solutions for efficiently storing, processing, and analyzing the increased volume of data at the edge.
- **Re-examine Security**: Prioritize robust encryption, strong access controls, and intrusion detection systems explicitly tailored to your expanding edge environment.
- **Plan for Multi-network Integration**: Develop strategies for how to effectively merge and manage disparate network types (wired, cellular, satellite, etc.) to create a robust and adaptable edge infrastructure.
- **Consider Data Sovereignty & Compliance:** Map data flows and storage locations, ensuring they comply with relevant regulations like GDPR, particularly in geographically dispersed environments.
- **Create Documentation:** Emphasize thorough documentation of network configurations, software versions, and maintenance logs throughout the planning and rollout phases.
- **Engage Experts**: Consult with specialists in edge computing and network connectivity optimization like ZEDEDA to develop a comprehensive scaling plan that aligns with your unique business objectives

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, while increasing visibility, security and control.

ZEDEDA ensures extensibility and flexibility by leveraging a partner ecosystem, and EVE-OS, open-source Linux-based edge operating system.