Connection
we solve IT

MODERN INFRASTRUCTURE AND MULTICLOUD SOLUTIONS

Accelerating Edge Rollouts with Automated and Scalable Solutions

E-BOOK

Table of Contents

Introduction: Edge Computing as a Business Imperative

As enterprises strive to modernize IT infrastructure and deliver more responsive services, edge computing is emerging as a critical enabler. By decentralizing data processing to locations closer to where data is generated, edge computing supports real-time decision-making, improves operational efficiency, and unlocks innovative new services powered by AI and IoT. In fact, as highlighted in recent business strategy research, edge computing is now tightly aligned with achieving key business objectives such as agility, resilience, and competitive differentiation.

IT operations teams responsible for the deployment of remote edge systems and their ongoing lifecycle management should look to leverage modern tools purpose-built for the management of widely distributed enterprise edge systems. These tools can efficiently—and securely—deploy bare-metal systems, provision their operating systems and hypervisors, and then install and configure applications and services.

These systems include automation features that often feature a template-based approach that enables rapid and consistent deployments—and ongoing management—of the entire stack.

Edge computing is now positioned at the forefront of IT infrastructure management. According to Gartner, 75% of enterprise data will be processed outside traditional data centers/clouds by 2025—a sizable increase from less than 10% in 2019.[1]

This eBook serves as a resource guide for organizations looking to accelerate the deployment of distributed edge infrastructures at scale through automation, orchestration, and secure lifecycle management.

The Need for Scalable Edge Deployment

Deploying edge infrastructure across diverse environments can be complex. However, with a modern “DevEdgeOps” approach—combining DevOps principles with edge-specific operational needs—and standards-based systems management platforms, organizations can achieve rapid, consistent, and scalable deployments while minimizing manual intervention.

A wave of innovations in IoT compute, AIOps-based systems management applications, microservices architectures, and AI application development platforms have all converged to enable a common framework for distributed edge deployment and management. Previously, IT teams often lacked the systems and expertise necessary to efficiently provision and manage highly distributed edge infrastructures. Attempts to leverage legacy systems and centralized management tools frequently led to inefficiency, security risks, and inflated costs.

Taking this further, DevEdgeOps is a unified integration of DevOps practices and principles applied to the unique requirements of a distributed fleet of edge devices. DevEdgeOps unites two apparently competing forces into a holistic approach. While the operational best practices of DevOps function consistently in cloud environments which tend to focus more on applications, services, and virtualized systems, the wide variety and physicality of edge deployments require the flexibility to manage highly diverse fleets and infrastructures.

In its ideal form, the DevEdgeOps paradigm empowers IT and OT operations teams to deploy modern high-value solutions, whose infrastructure management is based on flexibility, consistency, scalability, extensibility, and security.

Look for the following features and functionality included in the leading-edge management systems:

- Highly integrated and purpose-built for management of widely distributed edge locations
- Flexibility to run virtual and container applications and a multitude of operating systems
- Consistent interoperability across edge fleets, centralized data centers, and cloud providers
- Security features at the application, device, network, and data layers
- Platform extensibility and flexibility

Key Focus Areas for Modern Edge Management

1. Zero-touch Provisioning

Zero-touch provisioning (ZTP) enables edge platforms to configure devices automatically when connected. This minimizes manual intervention and configuration drift.

Key technologies include:

- Near zero-touch provisioning (nZTP) for device discovery and secure registration
- FIDO device onboard (FDO), which uses ownership vouchers and secure key exchanges to onboard devices

Benefits include automated device activation, secure provisioning via key vaults, and encrypted management tunnels. ZTP reduces the need for onsite personnel and supports sustainability goals by minimizing transport-related emissions.

2. Standardized Deployment Models

Consistency means security and stability, and edge management systems must leverage template-based deployment and configuration management that also allows for the diversity at the edge. On the other hand, flexibility and extensibility are critically important, as well. Edge compute systems take many forms and are based on a variety of platforms and reference architectures. The ability to provide integrations and "plug-ins" to manage the intricacies of each unique platform is critical. This extends to unique industry applications and AI solutions based on multiple vendors and products in a single solution. Embracing vendor specificity while avoiding vendor lock-in is an achievable art form.

The use of standardized deployment models applies to all layers of the stack. The leading modern edge compute platforms apply this approach to each of the following layers of the full stack, in which automated provisioning and deployment, software patches and updates, and uniform configuration management are required:

- Bare-metal deployment
- Virtual machine and OS provisioning
- Container orchestration
- Application installation and updates

An industry-leading example of standardized deployment models is using orchestration and configuration templates, or "blueprints," as in the Dell NativeEdge solution. Blueprints are the foundation of task automation within Dell NativeEdge. A blueprint is a comprehensive automation plan that includes the entire configuration required for deploying an end-to-end edge solution. As such, a blueprint combines components such as infrastructure resources, network configurations, application settings, custom workflows, and scripts. IT operations teams can deploy an application by applying the blueprint to many edge devices throughout multiple provisioning and deployment stages, from test-dev to production. Blueprints also streamline Day-2 operations by automating post-deployment tasks such as software updates and configuration changes at all layers of the solution stack.

Another example is VMware's Edge Compute Stack (ECS), or the latest generation VMware Cloud Foundation Edge, which is a streamlined configuration of VMware Cloud Foundation (VCF) customized for edge use cases. Since VCF Edge is built upon VCF, it provides a private cloud infrastructure platform that delivers a highly integrated enterprise-class collection of compute, storage, networking, management and security capabilities. And being based on VCF, VCF Edge provides not only a scalable and flexible edge infrastructure, but the automated deployment and lifecycle management features intrinsic to VCF—thus simplifying the management and orchestration of edge deployments across multiple sites.

3. Containerization and Virtualization

VMware's Cloud Foundation Edge is a clear example of a platform that provides unified management of both virtual machines and containers running simultaneously, with container orchestration provided by the underlying Tanzu platform.

Dell's NativeEdge provides multiple deployment models, from VMware Edge Compute Stack running on multi-node clusters, to Kubernetes clusters based on the lightweight K3s. As lightweight Kubernetes distribution created by Rancher Labs, K3s is a perfect fit for Edge deployments based on its lightweight binaries, low resource consumption, and high availability features. And K3s maintains compatibility with standard Kubernetes, allowing the use of tools and workflows that are typical of full-blown enterprise Kubernetes deployments. It even runs all Kubernetes components in a container.

This brings us again to Dell's NativeEdge blueprints to manage such container-based environments. NativeEdge blueprints automate the deployment and configuration of both VMs and containers, and even the applications themselves. Blueprints leverage orchestration tools like Ansible and Helm and can deploy multi-node clusters on virtual machines (typically 1, 3 or 5 modes). To deploy directly to bare-metal or native containers, NativeEdge can even leverage Docker Compose files. And of course, NativeEdge easily configures parameters required by each application, such as network and CPU settings, data volumes and repositories, and more.

4. Application Deployment and Updates

Edge management systems employ several emerging standards to consistently deploy applications, software updates, code libraries, and industry-standard code frameworks, such as those used for edge-based AI applications. Fundamental to these standards is the use of preconfigured templates and workflows for both application deployment and ongoing configuration or "state management."

Many edge management systems utilize a stacked combination of technologies, commercial and open-source products, applications, and code libraries to achieve a unified approach to deployment. For example, many edge management systems utilize their own software within their own systems management portals as well as how they structure deployment workflows and runbooks. Most also leverage industry-standard projects, applications, and utilities such as Ansible, Fabric, Utilities, Helm, GitHub, Docker Compose, FreeMarker, YAML, Flux CD, and more.

For context, let's first look at a simplified workflow for an automated application deployment sequence:

- An application bundle and template are uploaded to a central code repository (such as GitHub).
- An edge deployment site is identified, and an edge cluster or target device is selected.
- The edge management system connects to and prepares the edge device(s) for deployment.
- The application (and supporting code libraries) is deployed to the cluster or edge device(s).
- The application is tested and validated.

One exciting development in the world of application deployment and management is in the realm of independent software vendors (ISV) integration. Several edge management system vendors are beginning to partner with ISVs to create and maintain readily accessible catalogs of runbooks or blueprints that optimize deployment, configuration, and lifecycle management of their solutions. Examples of these runbooks, templates, or blueprints include:

- Solutions for targeted use cases in security, observability, computer vision, and machine learning, across multiple industries and verticals and deployment theaters
- NVIDIA AI software for inferencing at the edge
- Data collection and aggregation systems that convert the data gathered from various IoT devices and sensors into stream or file sets that can be efficiently transferred central data centers, cloud, or other edge locations

5. Security Features for Full-Stack Management

We round out this eBook with a summary of edge management system security features and functionality not previously discussed, which enable safe and trusted processes—all the way from bare-metal device provisioning to application deployment and updates. Dell research showed that 81% of organizations recognize emerging technologies such as AI, IoT, and edge computing present significant data protection challenges.[2]

Maintaining security at each edge location presents unique challenges, since edge nodes are often not housed in a well-protected environment as they would be at a dedicated data center.

Thus, edge management systems must be intrinsically secure, from the operations console to the repositories from which applications and updates are deployed, to the connectivity by which systems are managed. And the edge devices themselves must be validated and trusted by the management system itself. Encryption techniques are used in different ways up and down the stack, to secure data in transit (including command-and-control) and data at rest (which can include code repositories, configuration files and templates, and a company's own proprietary data being processed at their edge locations).

Edge management systems integrate these security features at multiple levels (in a multitude of ways) within the management stack:

Device Level

- End-to-end (E2E) edge device supply chain and authenticity validation hardware and software
- Secure onboarding (via FIDO FDO and SZTP, for example, as previously outlined)
- Hardened BIOS/UEFI that secures the boot sequence up to the OS boot level
- Secure Boot/TPM with unique digital identity

Secure Boot ensures that only a digitally signed and trusted operating system image can be used at the time of system boot. This is often enabled using an integrated **trusted platform module (TPM)**, a server hardware feature that enables accelerated advanced cryptographic methods, such as digital signatures and remote attestation, based on an intrinsically secure chain of trust.

User Level

- Privileged access management (PAM) and role-based authentication (RBAC)
- Continuous authentication and authorization of users to systems and applications
- Multifactor authentication (MFA)

Network Level

- Micro-segmentation
- Application isolation
- End-to-end (E2E) transport encryption (data in transit and data at rest)
- Edge gateways: Outbound session initialization enforcement; edge gateways do not allow inbound connections by design

Data Level

- Data in-flight encryption
- Data access control
- Data-at-rest encryption

Data-at-rest encryption prevents access to data repositories without encryption keys by which to unlock the data. In the event of media loss or theft, data remains secure without the presence of the unlocking key.

Connected-device (Peripheral) Level

- Peripheral isolation

Isolation of the edge system OS from peripherals helps eliminate attack vectors that could originate from externally connected devices. This is a critical feature since edge devices can be installed in locations that may be accessible to internal attackers.

Edge Management Traffic Level

- Command-and-control in-flight encryption

Many edge management systems are designed to implement mutual TLS (mTLS) on port 443 for communication between edge gateways and the edge management system itself. Critically, the edge management architecture typically enforces outbound communication only from edge gateways to the edge management portal while prohibiting inbound communication to the gateways.

- Edge site firewall protection with network address translation (NAT)
- Site-to-site VPN between edge sites and data center
- Secure connections between edge clusters and other supporting cloud services

DevEdgeOps: Bridging DevOps and Edge Infrastructure

DevEdgeOps adapts agile DevOps workflows to edge-specific constraints. It merges development automation with real-time operational oversight, helping teams deploy infrastructure with:

- Flexibility
- Scalability
- Consistency
- Security

DevEdgeOps empowers IT/OT teams to manage heterogeneous edge fleets just as effectively as centralized or cloud environments.

Key attributes of DevEdgeOps-ready platforms include:

- Multiplatform interoperability (VMs, containers, OSs)
- Extensible plugin support
- Security-first architecture
- Template-driven automation for all provisioning and lifecycle phases

Business Benefits of Edge at Scale

Edge computing is no longer optional—it's a strategic priority. It is no surprise that the growth forecast for edge systems looks strong—IDC forecast a 15.4% increase in worldwide spending on edge computing, totaling an expected $232 billion in 2024, and $350 billion in 2027.[3]

By embracing automation, standardization, and security best practices, businesses can:

- Reduce time-to-deploy for edge systems
- Minimize operational overhead
- Ensure fleet-wide consistency
- Enable real-time insights from AI and IoT workloads

Ultimately, organizations that adopt scalable, automated edge rollouts are positioned to capitalize on the decentralization of compute. With a DevEdgeOps model and modern edge management platforms, enterprises can build secure, scalable edge environments that support next-gen innovation while driving agility, resilience, and competitive advantage.

How Connection Can Help

Connection is your partner for modern infrastructure solutions and services. From hardware and software to consulting and customized solutions, we're leading the way in infrastructure modernization.

Explore our Solutions and Services
Modern Infrastructure

Reach out to one of our Connection experts today:
1.800.998.0067

[1] Compunnel, 2025, The Convergence of Edge and Cloud: How Edge Computing Enhances Capabilities

[2] Dell Technologies, 2025, Dell NativeEdge Security Solution Brief

[3] FutureCIO, 2024, Edge Computing Investments Will Reach 232 Billion in 2024