A Deep Dive on Edge+Scalability

Edge computing is everywhere, and it is a trend that happened overnight. For enterprises and organizations that are eyeing or using the technology across their operations, it means learning how edge computing can make them more agile, more efficient and better prepared to grow their businesses in the global marketplace.

The rapid proliferation of edge computing in the last couple of years is amazing, fueled by the cascading growth of data that is used and located far from old world data centers. That includes data on mobile devices, on cash register networks in stores, and in a myriad of other places inside a company's operations. It also includes data growth created through the steady acceleration in IoT device deployments and by other business use cases where enterprises are realizing the value of having and using their data at the edge of their networks.

Edge computing is one of the fastest growing technology needs in the world today for enterprises of all types and sizes. But like the broad topic of cloud computing, edge computing can be complicated, difficult to deploy and manage and require specific expertise to make it work as a well-lubricated part of an organization's vast technology infrastructure.

One critical topic when discussing edge computing is how it can scale across organizations and their infrastructure, wherever and whenever more growth is required. Scalability of edge computing is something that requires consideration from the start of an edge computing deployment because once companies begin seeing the benefits of the technology, they will also likely want to increase its use and grow its possibilities.

Management Challenges in Edge Computing Deployments

The need for efficient scalability in edge computing for enterprises is one of its biggest management challenges as the volume of a company's data continues to expand. Scalability of edge architectures allows organizations to process and utilize data where it is needed and most of value. These edge computing management challenges can be found across a wide range of businesses, including telco, shipping, cruise lines, industrial facilities, factories, retail deployments and more, wherever there are satellite locations that are not connected to traditional data center compute power.

Dealing with increasing numbers of devices. The need to grow and scale edge computing capabilities will closely follow increases in the number of devices and endpoints that require oversight and management inside a company. These numbers can fluctuate, and this requires efficient and easy systems management to track and process them as needed by IT staff.

Accurate planning is a must. To make this linear growth in edge computing deployments happen without hitches, companies must plan carefully for their future system expansion needs. This means monitoring edge computing initiatives for satisfactory performance even after a deployment is fully rolled out so that additional capacity requirements can be spotted and added before system bottlenecks occur.

System complexity in edge computing. Edge computing can simplify a distributed IT environment, but edge infrastructure is not always simple to implement and manage. By scaling out edge servers to small remote sites, configuration and management can be more complicated than adding the equivalent capacity to a single core data center where IT workers are located. The increased overhead of remote physical locations can be difficult for smaller companies to manage, adding to complexity and worries. And if something fails in a remote edge computing site, there needs to be infrastructure in place that is repairable by non-technical local labor and managed centrally by IT workers who are located elsewhere.

Avoiding vendor lock-in hassles. Organizations each have unique IT infrastructures, made up of old and new hardware and software from diverse groups of vendors and their partners. When building out an edge computing strategy

and system, enterprises must have interoperable infrastructure with components and applications sourced from various vendors. This prevents vendor lock-in issues and allows organizations to grow, change, adapt, and scale without completely re-evaluating or re-engineering their edge solutions and vendors each time their business needs change. Avoiding vendor lock-in at edge sites — where organizations directly generate value and revenue — is critical for continuous innovation and competitive differentiation. Avoiding vendor lock-in allows enterprises to develop their own intellectual property to improve their critical business operations and offerings, giving them an advantage in highly competitive markets.

Easily reproducible remote site management is critical. A company's remote edge computing sites must be highly and easily reproducible to dramatically simplify their management. This will allow for easier troubleshooting when problems arise, while also keeping those issues common between locations because the systems are uniform across an enterprise's infrastructure. Vexing problems can occur more often when software applications are set up differently at each site.

Unexpected problems compound at the edge. Because edge computing is managed remotely, away from centralized or on-premises data centers, the undetected and unexpected issues that always seem to arise in a broad range of technology projects are compounded. This makes these complications even more threatening because they happen away from IT quick response team personnel. Edge computing initiatives must include methods to respond to these issues.

Dealing with multiple clusters. By its nature, edge computing adds complications for certain kinds of configurations and system needs. One example is enterprises that need to deploy and manage multiple compute clusters in locations where physical space is at a premium and connectivity can be difficult. This kind of challenge arises due to the unique profile of edge computing, where compute, connectivity and space resources diminish by default as they get farther from the data center where the work is addressed.

The need for common management tools. Having an edge computing architecture that requires different management tools than the rest of the critical IT systems inside a company should be a major concern for enterprises. By ensuring that the same tools and processes used in other centralized infrastructure, such as a company's core data center or chose cloud provider, can also work in the edge computing environment, companies can simplify and streamline their already complex system management responsibilities. This would include things like seeking centralized tools for automated provisioning, management, and orchestration of potentially tens of thousands of sites that have minimal or no IT staffing across edge, core and cloud?

Fail-safe measures are mandatory. Not only is remote management of edge computing sites a constant concern, but edge sites also must continue to operate in the event of network failures. In response, the inclusion of fail-safe provisions is critical to ensure that high availability is maintained 24-7.

Custom edge computing requirements for each remote site. Every edge computing site is not going to be the same. Each will have different technical requirements on different tiers, including the size of the hardware footprint, challenging environments, and costs. Enterprises must have access and flexibility to use hybrid workloads that consist of virtual machines, containers, and bare-metal nodes running network functions, video streaming, gaming, AI/ML, and business-critical applications.

Edge computing and its demands can be challenging, but they certainly can be successful, particularly when also adopting appropriate and advanced system management layers to help smooth the path to success for the projects.

What is ACM and How it Addresses Edge Computing Challenges

To help battle these edge computing challenges, Red Hat's Advanced Cluster Management (ACM) targets these difficult issues that cause great pain for organizations. Red Hat ACM provides edge cluster management and applies security policies, making

it easier for enterprises to configure and manage their Red Hat OpenShift edge clusters while reducing complexity, increasing scalability and removing vendor lock-in.

ACM allows enterprises to deploy and manage the lifecycles of their OpenShift clusters as well, giving them the tools to create, update, scale and even remove them reliably and consistently in an open source programming model. Application lifecycle tools are built-in so that needed applications can be added and configured using integrated open standards and CI/CD pipelines. Governance, risk, and compliance tools are in ACM as well, allowing the creation of policies that automatically configure and maintain the consistency of controls and desired states. This provides powerful observability capabilities that visually see system alerts, critical application metrics and overall system health through a single pane where issues can be identified and resolved.

CUSTOMER SUCCESS STORY

In the process of evaluating Red Hat and its Edge offerings, we were able to hear from a customer on how Red Hat's OpenShift platform has helped them achieve their goals.

5G network expansion offers new data opportunities for telco companies. **Verizon and its former division, Verizon Media**, wanted a way to use a large-scale AI platform to help process data at the edge to improve decision making and build innovative new services. The company built Leo, an AI platform with a modular, container-based foundation on Red Hat OpenShift. Its platform was able to now process more than 1 million AI inferences per second. As a result, Verizon improved prediction accuracy for thousands of edge appliances and endpoints.

How Red Hat Solves Edge Computing Scalability Issues

The growth and importance of edge computing is continuing to inspire a broad ecosystem of supported products and services to help make the technology more manageable for enterprises. And powering that growth is open source, which is becoming the prevalent approach to deliver on edge computing interoperability, scale and ease of operation across public clouds, private data centers and edge locations. Red Hat is well-positioned for this market due to its ability to leverage a diverse ecosystem of providers and software vendors, its in-house innovation – and most crucially – its device-agnostic platforms that focus on delivering scale, automation and security of operation.

Red Hat's broad product lines use a common platform and tool sets from edge to core to cloud, which help enterprises reduce the required skills needed to keep them running and bringing in business value. Red Hat OpenShift, Red Hat Enterprise Linux (RHEL) and its other platforms provide operational consistency and portability of applications while ensuring consistent application lifecycles and development processes.

And when building and operating edge computing deployments, comprehensive systems management and scaling capabilities are essential for success, making Red Hat an excellent partner candidate for edge environments. The need to drive robust operations is never more relevant when you factor in the sheer breadth and scale of edge computing deployments.

Red Hat has long led the market with its powerful open source platforms, expertise, services and commitment to its customers around the world. The edge computing marketplace is just the latest enterprise growth area where Red Hat's work, reputation and quality will help businesses take their technology infrastructures to the next level.