

Transformation of the Edge: Architecting your infrastructure for performance and reliability



Executive summary

Many companies are looking to their edge computing systems for improvements in the way they collect and process data resulting from the explosive growth of Internet-connected devices. Yet for most operational technology (OT) professionals who have long been responsible for edge-based computing systems in their organizations, the evolution of the edge may seem like a moving target.

As business and technology trends continue to reshape how computing happens at the edge, many companies are searching for the best way to select, implement, and capitalize on the opportunities that future-proof edge computing strategies offer today.

As business and technology trends continue to reshape how computing happens at the edge, many companies are searching for the best way to select, implement, and capitalize on the opportunities that future proof edge computing strategies offer today.

This white paper offers a closer look at edge computing and highlights various examples, industries, and use cases where it delivers the most benefit. It also describes the evolution that is taking place in edge computing today while demonstrating how next-generation solutions can help organizations execute on the new opportunities presented by the transformation of the edge, today and tomorrow.

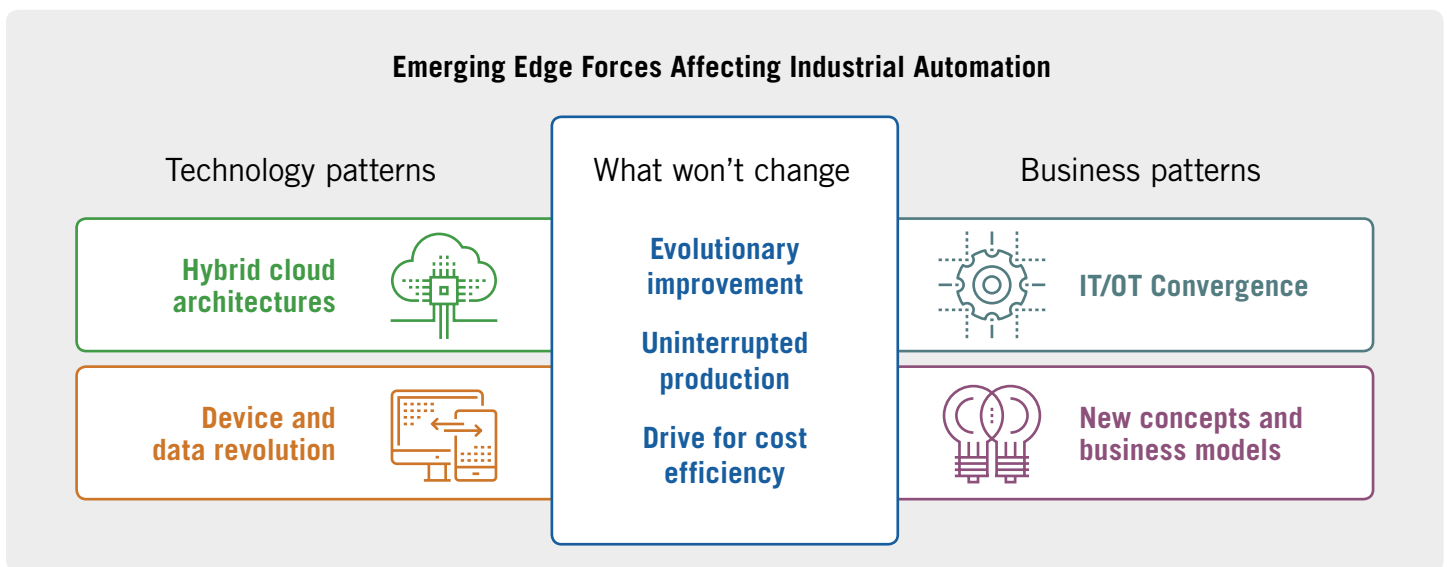
Welcome to the new edge

Today, many critical operations and processes take place at the edge of an organization’s network, such as unmanned machinery, public safety systems, transportation systems, and power and energy production.

These operations, and the equipment they rely on, are connected to a rapidly growing number of Internet of Things (IoT) devices, which is forecasted to exceed 20.4 billion endpoints by 2020.¹ These sensors and devices collect real-time data that can be used to streamline production cycles, improve product or service quality, or shorten response times. As a result, better management of edge data can help the entire business improve critical processes, reduce costs and gain a significant competitive advantage.

¹Gartner, *The Need for Real-Time Insights Offers Market Opportunities in Edge Computing*, “Forecast: Internet of Things – Endpoints and associated Services Worldwide, 2016,” February 24, 2017.

Emerging Edge Forces Affecting Industrial Automation



Yet reaping these benefits is more difficult than it might appear, especially when attempting to process and manage edge data through centralized cloud-based applications or on-premise data centers. While these computing models have their advantages, they are not particularly well suited to supporting environments where IT staff is limited, connectivity is poor or expensive, and operations are especially time-sensitive. All of these conditions are usually found at the edge.

Edge computing today is going through significant changes, bringing new pressures as well as new opportunities that together are driving the rapid transformation of what is happening at the edge.

The potential of the edge

To overcome such inherent challenges, many organizations have begun to rethink their edge computing infrastructures, which generally have limited or no connectivity to a remote data center and/or the cloud. Edge data may be processed directly at the point of origin, or it can be sent to a gateway networking device or intermediary server. But either way, the usefulness of the data is often constrained by the scale and capacity limitations of hardware – sometimes 30 years old – that make up most edge infrastructure.

Edge computing today is going through significant changes, starting with increased usage. This offers new pressures as well as new opportunities that together are driving the rapid transformation of what is happening at the edge.

For organizations that capitalize on these opportunities, the payoff can be tremendous. For example, applications related to analytics, smart traffic lights, vehicle-to-vehicle communications, power production, and manufacturing equipment benefit from improved edge computing infrastructure that more effectively processes and uses streaming data from IoT devices and sensors. In these cases, these applications can reroute traffic, improve predictive maintenance to prevent defects, and optimize production.

Benefits to OT Professionals

OT professionals have a lot on their plates: In addition to keeping the assembly line running, the water supply flowing, or the trains running on time, they are now being tasked with updating their IA systems to take advantage of the evolving edge.

A technology-plus-service approach to optimizing edge infrastructure can greatly simplify this challenge and deliver significant benefits to OT professionals and their organizations. These benefits include:

- **Risk mitigation:** No unplanned downtime in their IA systems helps organizations avoid the disruptions that result from even brief system failures.
- **Operational efficiencies:** Proactive monitoring, automation, and better visibility into every component of their IA systems enable OT professionals to quickly make operational improvements, saving time and money.
- **Reduced TCO:** Integration of all the hardware, software, and services required to manage IA applications into a single, unified solution translates into the lowest total cost of ownership for a high-availability edge solution.
- **OT empowerment:** Simplifying the management of edge architectures helps reduce OT professionals' dependence on IT departments and empowers them to take the lead in deploying technologies in their sphere of influence – to IIoT and beyond.

Changing requirements of edge computing

Edge systems are usually deployed in remote locations beyond the typical data center, or by definition, up to the very edge of computing networks. Simplicity is key at the edge, but so too are elements such as serviceability and usability. As talented people become increasingly scarce, edge environments can

often only afford to rely on OT professionals – not IT – so if anything goes wrong, the idea of flying someone in to fix it is not feasible from a time or cost perspective. OT staff may not be IT-specialized but are still able to get the product serviced and up and running quickly.

Many new edge applications are either holding a significant amount of data or streaming data to other locations, which makes data integrity extremely important. Any failures could cause a break in the data stream, presenting several issues to existing processes or compliance efforts. For example, in a dairy farm, if a system fails as milk is moving through a production line, farmers may not be able to collect the required data for compliance, creating increased risk in the event of a potential recall.

Applications at the edge generally do not get updated very often. This means companies need technology with a service model and tools that enable them to deploy the system and essentially leave it alone for a very long time.

Also, applications at the edge generally do not get updated very often. This means companies need technology with a service model and tools that enable them to deploy the system and essentially leave it alone for a very long time. In addition, in many manufacturing environments, edge systems are now being deployed with the intent of keeping them in place for ten years or more. This contrasts with a traditional data center where hardware infrastructure tends to be refreshed every three to five years, and its applications and middleware are routinely updated.

New trends, complications, and opportunities

Organizations computing at the edge today may believe that their existing approach is sufficient – for now. However, many trends and shifts are occurring in technology that will require them to change how they think and act about edge computing:

- **Device evolution:** Increased volumes will drive down unit costs, making IoT devices cheaper, more capable, and standardized. For example, programmable logic controllers (PLCs) are already becoming more intelligent, open standards-based, and much more cost-effective.
- **OT will evolve:** The OT role will continue to evolve because IT departments are already stretched too thin, and the edge is simply not a top priority for them. OT will assume more responsibility for those technologies that are critical to them, such as edge computing.

New Considerations for Edge Solutions

Reliable infrastructure

Can we deliver redundant compute, network and storage with minimal downtime?



Application monitoring

Can the applications work together and respond to my data needs?



Effective
Edge
Solutions

Predictive serviceability

Can we minimize IT involvement by remotely managing and predictively servicing the Edge stack?



Secure connectivity

Can we ensure the data is not tampered with and has not been stolen?



- **Increasing machine-based automation:** Emerging technologies such as real-time analytics and artificial intelligence will increase the criticality of edge data. Automation will become increasingly faster and much more capable, even surpassing the limits of human ability. For example, railways can now process edge data to support 300 km per hour travel safely.
- **New vendors and new business models:** New technological solutions, such as AI, machine learning, and advanced analytics will continue to drive innovation, agility, and capital efficiency. Organizations must consider partnering with vendors, technology, and solutions that are best suited to their long-term strategies.
- **Hyper-converged platforms,** such as Hewlett Packard Enterprise's Converged Edge Systems, integrate data capture, control, compute, and storage capabilities in a single hardware device. They emphasize operational simplicity and can perform real-time analytics at the edge, but at the sacrifice of continuous availability and affordability.
- **High-availability solutions,** such as VMware's vSphere, rely on virtualization to deliver higher levels of availability at acceptable cost, but sacrifice operational simplicity by requiring IT-level skill sets. Many also lack the computational strength and performance required for real-time data synthesis and analysis.

Clearly, operating at the edge will continue to require new ways of thinking and acting. These trends must be carefully evaluated by any organization deploying an edge computing strategy.

Today's solutions fall short

There is widespread recognition that today's edge-based systems need new enabling technologies to meet changing demands. However, the supply of such technologies lags behind the need.

There are computing/communications technologies available that address the industrial automation (IA) sector. However, they tend to focus on the needs that IA end user customers have had for decades, rather than on edge-specific pressures (such as the need to perform predictive maintenance or integrate with Big Data analytics) that are driving the IA field today. As such, they represent a suboptimal combination of capabilities and deficiencies.

Examples of these technologies include:

- **IA network computers** from traditional data-center vendors, such as Dell's Edge Gateways for IoT, can collect data from sensors or devices on the edge and send it to a data center or public cloud. They are less expensive than other classes of edge solutions, but do not scale up well (as opposed to scaling outward), and do not provide continuous availability or operational simplicity – both of which are essential on the edge.

Clearly, the 30-year-old industrial control systems that are still the norm in many manufacturing and IA settings need a substantial upgrade to move toward the Industrial Internet of Things, pervasive clouds, and/or fog computing. However, a rip-and-replace approach will not work. The challenges facing these companies during their transition to the edge includes continuing to collect data from their legacy systems, extending the lifecycle of those systems where possible, and incrementally adding capabilities.

Moreover, not all companies need to move all the way along the continuum of the edge. Their business needs will dictate how much of an edge transformation they require. Yet in order for any organization to reap the maximum benefits that the evolving edge can deliver to its business, a new technology approach is needed.

The 30-year-old industrial control systems that are still the norm in many manufacturing and IA settings need a substantial upgrade to move toward the Industrial Internet of Things, pervasive clouds, and/or fog computing. However, a rip-and-replace approach will not work.

The ideal solution

To overcome these shortcomings, an ideal solution would satisfy enterprise edge needs for reliable, scalable, high performance computing. Furthermore, it would deliver these near-term benefits, and meet current requirements, as well as accommodate an evolving edge infrastructure.

More specifically, the essential attributes for a move to the edge include:

- **Operational simplicity** that provides low-touch to zero-touch operation, serviceability, and usability – all critical since IT resources are scarce at the edge or on the plant floor.
- **Uninterrupted production** for no unplanned downtime, plus the assurance of data integrity, since newer edge applications either hold large volumes of data or stream data to remote locations (such as the cloud).
- **Virtualization and reliability capabilities** that enable manufacturers and IA organizations to deploy highly available – yet highly efficient – business-critical systems and databases.
- **Interoperability designs** that leverage current standards (e.g., operating systems) as well as emerging standards that will enable devices and systems to operate with each other.
- **Cyber security**, including built-in protection of all components of the edge computing system, the data they handle, and their communications with each other and externally.

An example of edge-optimized architecture and delivery

Today, industry-leading technology vendors are delivering an innovative new edge infrastructure approach that incorporates all the attributes described above. More, this solution is comprehensive – a complete, turnkey compute platform that enables IA companies to update aging systems, virtualize industrial control applications, and increase operational efficiency.

This new edge approach consists of three essential elements:

- **A zero-touch computing platform:** An automated, fully integrated, and self-protecting industrial computing platform that reduces the need for IT support while increasing operator efficiency and lowering downtime risk.

- **Software-defined edge infrastructure:** A single interface to a full stack of essential applications, including virtualized compute, storage, and networking; cyber security; IoT gateways; VPNs; routers; analytics; and artificial intelligence.
- **A proactive managed service:** A revolutionary approach to deploying, monitoring, and maintaining edge solutions and infrastructure. Unlike yesterday's reactive service, such a proactively managed service can be easily overseen by OT professionals, executed by the system itself, and supported by a single, locally based service provider.

The combination of all three elements – a zero-touch computing platform, the software-defined edge infrastructure, and proactive managed service – creates a new infrastructure solution that is optimized for organizations at the edge.

Key advantages of the optimized solution are:

- It is **ideal** for running business-critical industrial applications quickly and reliably in remote, locations with limited or no IT resources.
- It **increases** operational efficiency with pre-installed virtualization software and intuitive, user-friendly configuration and management tools.
- It is **easy** to deploy and can be installed quickly by users, reducing the time it takes to get critical applications up and running.
- It **reduces** IT's burden with self-monitoring and self-protecting capabilities that make it ideal for unmanned facilities.
- It predictively **prevents** unplanned downtime, via redundant on-premise systems backed by a managed service cloud.
- It **supports** multiple ecosystems – a wide range of architectures and applications, without modification.

Such a comprehensive edge platform satisfies the essential requirements of the today's edge environment (see section "Today's Solutions Fall Short"), including operational simplicity, virtualization, uninterrupted production, and interoperability.



About Stratus Technologies

Stratus takes the complexity out of keeping business critical applications running 24/7. Stratus' technologies proactively prevent instances of unplanned downtime both in the data center and at the edge, and our services ensure any issues are addressed before customers need to. Global Fortune 500 companies and small to medium-sized businesses in a wide range of industries across the globe have been relying on Stratus for operationally simple, continuous availability for more than 35 years.

Learn more at stratus.com today.



www.stratus.com

Specifications and descriptions are summary in nature and subject to change without notice.

Stratus and the Stratus Technologies logo are trademarks or registered trademarks of Stratus Technologies Bermuda Ltd. All other marks are the property of their respective owners. ©2018 Stratus Technologies Bermuda Ltd. All rights reserved. 170000