



INDUSTRY 4.0: OPERATIONAL TECHNOLOGY AND EDGE COMPUTING

Manufacturing at the Edge

Technical White Paper
From the Experts at Scale Computing

Table of Contents

- Introduction**3
- Industry 4.0**3
- Manufacturing Industry Challenges**4
 - Latency.....4
 - Security.....5
 - Operational Technology.....6
 - Predictive Maintenance6
 - Unplanned Downtime6
 - Complexity7
 - Augmented and Virtual Reality (AR/VR).....7
 - Manageable Data Analytics7
 - Easy Interoperability7
 - Meeting Sustainability Goals.....8
- Edge Computing**.....8
- Green Edge Computing**.....8
- The Harrison Story**.....10
 - Challenges10
 - Research11
 - The Lightbulb Moment.....11
 - Savings11
- Storage Cost Reduction**12
- Central Monitoring and Management**.....12
- Modernizing with Low-Code and Containers**13
- Conclusion**13

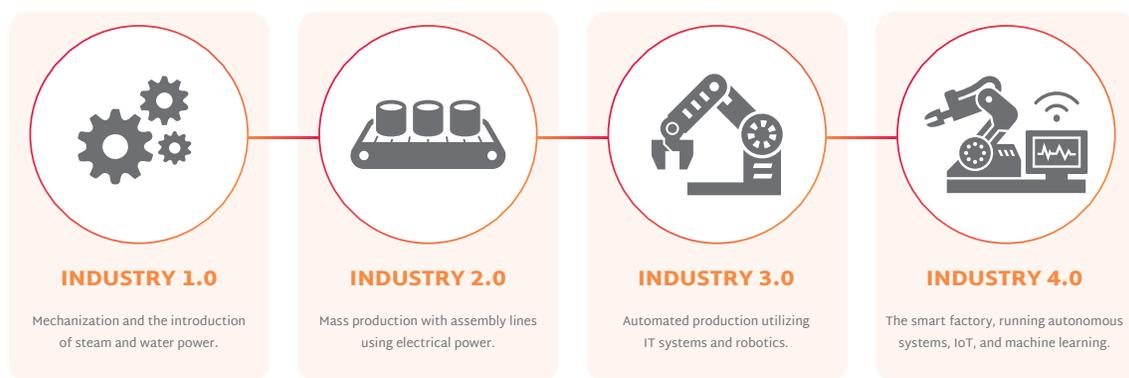
Introduction

The ever-expanding expectations of our digital world demand more applications run outside the cloud or data center, close to where people, devices and IoT technologies use them. But managing the sprawl of complex infrastructure to run applications locally in distributed locations is an escalating burden on IT departments. Attempts to handle these edge workloads with tools built for the cloud or traditional data centers result in excess cost, complexity, and fragility that fail to support manufacturing processes in the Industry 4.0 era.

Industry 4.0

Industry 4.0 originated in 2011 from a project in a high-tech strategy by the German government, which promotes the computerization of manufacturing.

Industry 4.0 uses technologies such as artificial intelligence (AI), machine learning (ML), augmented reality (AR), cloud and edge computing, robotics, the Internet of Things (IoT), 3D printing, and more, allowing manufacturers to integrate their end-to-end value chain. According to a [Gartner](#) study, almost 75% of data processing across diverse industries will occur at the network's edge by 2025, up from about 10% in 2018.



Industry 4.0 basic principles:

- **Interconnection.** The ability of machines, devices, sensors, and people to connect and communicate via the Internet of things.
- **Information transparency.** The transparency afforded by Industry 4.0 technology provides operators with comprehensive information to make decisions.
- **Technical assistance.** The technological facility of systems to assist humans in decision-making and problem-solving, and the ability to help humans with difficult or unsafe tasks.
- **Decentralized decisions.** The ability of systems and people to make decisions independently and perform their tasks as autonomously as possible.

Today's industry is shifting towards a green and digital transformation. Implementing Edge Industry 4.0 removes silos, brings IT and OT together, optimizes production, streamlines core functions, and fuels innovation. Implementing green technology solutions at the edge helps the overall sustainability performance of the manufacturer.

Edge computing can help to adapt to the local needs across many sites and grow as demands change. The addition of the Internet of Things (IoT) and other endpoint devices generates vast data sets, causing high latency, bandwidth congestion, and data bottlenecks in centralized networks. These challenges become more significant as IoT in [manufacturing is expected to grow](#) from \$62.1 billion in 2021 to \$300.3 billion in 2030.



Adopting edge computing requires transformative thinking. Implementing the tools and processes necessary at potentially thousands of sites with little to no IT staff is challenging at the best of times. Manufacturers who already implement edge solutions see a variety of hardware from various technology providers. Interoperability depends on obtaining resources from multiple OT vendors to create consistency across the edge architecture. This causes dramatic complexity and difficulty to manage and maintain. Standardizing on a low code platform, providing integrated virtualization and other key tools simultaneously help overcome the complexity and improve edge TCO.

Automation engineers face many challenges. Retrofitting an existing system for industrial IoT applications can be a challenge. It requires access to machine data and running advanced software solutions without disrupting the infrastructure. Edge computing provides a platform to enable retrofitting and reduce costs.

Modernize your production from an IT infrastructure point of view. Building smart factories means being smart about what technologies to deploy in the cloud, the data center and at the edge, on the manufacturing floor. Especially in manufacturing, any downtime can become extremely costly.

Manufacturing Industry Challenges

Manufacturers face massive challenges in deploying technology at the edge, such as latency, security, implementing highly available and manageable operational technology as well as predictive maintenance solutions, while minimizing downtime at the same time. But there is more: implementing augmented and virtual reality, limiting complexity, analyzing vast amounts of data generated by thousands of sensors, sustainability and the interoperability between sensors and systems complete the list.

Latency

Manufacturers must use data from multiple machines, processes, and systems to adapt the manufacturing process in real-time. These processes generate huge amounts of data and need machine learning to determine the best action based on the data analysis.

Collecting, aggregating, and filtering the data to send outcomes to a central server requires powerful computing at the edge, critical for AI and ML, rather than in the cloud. For example, latency over 3-5 milliseconds makes it impossible to accurately control manufacturing robotics.

Edge computing and networks also provide bidirectional information - data sent and collected at the same time – to deliver uninterrupted inventory updates 50 times faster than a cloud environment and avoid overstocking, stock outages, and shrinkage, particularly during periods of peak demand or when extreme situations occur.

Sending a request to a data center across the world and waiting for a response to return takes time. As a result, cloud computing isn't ideal for many mission-critical manufacturing applications. Manufacturing without this low latency can't experience the full benefits of the OT/IoT. If a connected machine on an assembly line recognizes a malfunction, any delay in transmitting that signal could be costly. It could lead to mechanical damages or even employee injury.



Security

Security should concern all smart manufacturers. Digital manufacturing requires agility and security. The continuously increasing number of IoT devices in the network means more potential vulnerabilities open to cyberattacks. The paths used by bad actors to gain unlawful access to networks may easily penetrate endpoint devices or nodes because they have limited encryption or firewall protections.

The centralized nature of cloud computing makes it vulnerable to many forms of cyberattacks. If a factory distributes processing and storage functions throughout the edge, however, no single attack could bring down the whole network. Since computing happens closer to the data source, less data is at risk at any given moment.

As more companies install edge networks, the addition of devices capable of performing increasingly complex functions further away from centralized servers will also increase. Addressing security risks, achieving economies of scale, and protecting sensitive data from bad actors are some barriers remaining before edge computing gains widespread acceptance is a must.

Manufacturers also want to know how to sync IT and OT at scale securely in a workplace where staff skilled in both OT and cybersecurity disciplines are scarce.

The OT applications from one or more sensors, IoT devices, robots or other equipment run as a small virtual machine (VM) on a Scale Computing HyperCore cluster. Securing the Scale Computing cluster means all VMs running applications from various OT solutions on the cluster can be protected with security solutions of choice. You no longer have to worry about the protection of individual devices. Scale Computing Platform provides built-in disaster recovery options to protect your data and applications. Even in those worst-case scenarios, we can help make sure you experience minimal downtime to keep you in business.

Operational Technology

One of the challenges faced by manufacturers is the ability to access data from their machines, processes, and systems.

The factory floor has many proprietary systems that do not talk to one another. Operational technology is still very traditional and does not use the same standards as IT, driving IT/OT convergence projects. Extracting data - for example, monitoring real-time scanning of sheet metal to detect fatigue, oven filling levels, flow through pipes, or keeping track of automated machine cycles - can overload a central server.

IoT connects sensors, equipment, controls, and industrial computing platforms to provide visibility and control of the equipment and processes that warehouses and factories depend on for efficient operation. These sensors and applications generate large volumes of data and need to deliver analytics in real-time for plant visibility and decision-making.

Edge computing allows manufacturers to filter data to reduce the amount sent to a central server, either on-premises or in a cloud.

Predictive Maintenance

Predictive maintenance is not an easy process to implement. OT and IT systems must be connected to provide the data to make this analysis possible. Detecting when a machine might fail through data analytics and mitigating this by conducting maintenance before a potential breakdown limits unplanned downtime risk. Predicting the outcome of the analytics is hindered by a limitation of measured variables and immature ML.

Running edge computing systems closer to the source means large amounts of data can be processed without the cost of data transportation to the cloud and latency issues.

Unplanned downtime

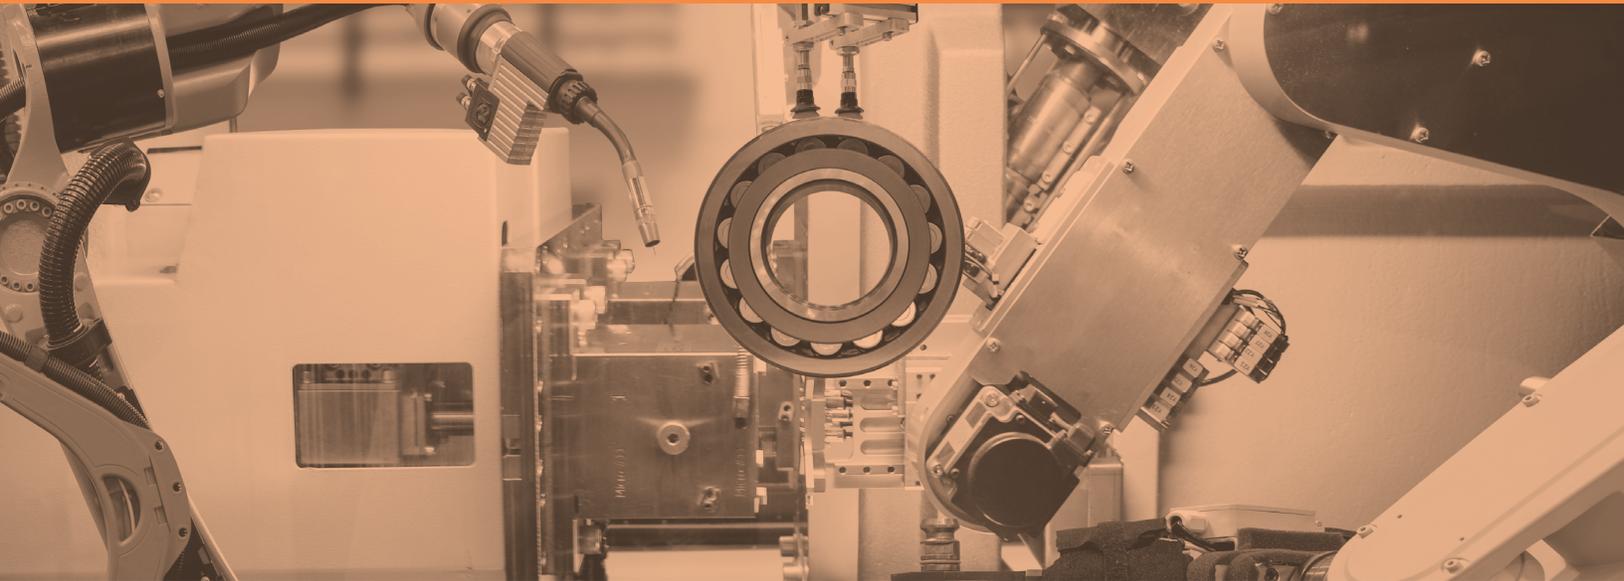
The cost of downtime increases in every industry. As business moves to a 24-hour economy, there are no more maintenance windows available. This is one of the more difficult areas to calculate. Based on the reports of analysts and research organizations, the impact is massive. According to [Aberdeen Group](#), the average cost of downtime across all businesses now adds up to \$260,000 per hour.

In the automotive industry, downtime can cost up to \$50,000 per minute (\$3 million per hour). High available IT infrastructure becomes essential.

A [Vanson Bourne](#) study reports that unplanned downtime results in loss of customer trust and productivity, in which 46% couldn't deliver services to customers, 37% lost production time on a critical asset, and 29% were unable to service or support specific equipment or assets.

Downtime in manufacturing can have a big impact on the bottom line. Almost every factory loses at least 5% of productivity, with experiences as much as a 20% loss due to downtime. Minimizing downtime in manufacturing is just as pivotal as maximizing quality and output to maintain contribution margins. This is at least what most think of when talking about downtime in manufacturing. It doesn't end there.

Due to software updates and upgrades, downtime often takes careful planning as systems won't be available for a certain period. Continuous operations, whether during a software update or even when a disk or a complete server fails, is critical. Systems need to keep running.



Complexity

Businesses departing from less expensive, established system architectures will require a substantial investment coupled with increased maintenance and IT support. Building an edge infrastructure is costly, involves significant planning and logistical preparations, and requires considerable software and hardware resources. For example, adding more endpoint devices or a new inventory storage area requires more servers and peripheral devices. This also increases the overhead expenses necessary to maintain them. For small businesses without remote locations or those producing low-complexity items, an edge network may yield some efficiency gains but may not be cost-effective.

In contrast, a centralized server utilizes the same infrastructure regardless of changes or additions to the factory layout.

Augmented and Virtual Reality (AR/VR)

Adopting AR and VR is another challenge which manufacturers face. These technologies mean you can introduce new, more efficient ways to train staff on health and safety, new production processes, fault detection inspection and more. The cloud, however, would prevent a good experience for the factory worker. Running the AR/VR software on a local edge device, processing data and rendering the stream, overcomes the latency challenge, providing a more seamless experience.

Manageable Data Analytics

Data is the foundation of the Industry 4.0 revolution. IoT devices and sensors collect massive amounts of data, providing a valuable base for improved analytics. Analyzing all this data requires considerable storage, bandwidth, and computing power.

Edge computing overcomes these concerns in two ways. First, processing data at or near the source, so the overall process is much faster. Second, each data point within the smart factory processes its own information, so no one system must handle everything. Processing data at the edge reduces the strain on individual systems and refines the process. Since it's segmented, it's easier to sift through and find relevant information.

Easy Interoperability

Seamless interoperability between devices and systems is key for every smart manufacturer. At the same time, this is one of the bigger challenges since there is no standard protocol. If these machines could translate protocols, though, they would make themselves more compatible.

Moving compute functions to edge computing systems, located close to the machinery – and their data providers like sensor and plc back boxes - on the factory floor eliminates some of the need for a universal standard. Located on an edge computing system, devices can convert signals themselves and work with a greater variety of systems.



Meeting Sustainability Goals

In a fast-changing, highly competitive landscape, manufacturers collaborate and innovate with customers and suppliers to accelerate progress on their sustainability goals. Manufacturers lacking innovation will face challenges in meeting these goals.

To overcome these challenges, manufacturers must strengthen their monitoring and evaluation systems — and these systems should encompass operational processes, including sustainable manufacturing and strategic initiatives and projects. A continuous improvement process must be put in place with the goal of increasing efficiency and effectiveness.

Edge Computing

Edge computing extends the capabilities of computation, network connection, and storage from the (private) cloud or data center to the edge of the network. It enables the application of business logic between the downstream data of the (private) cloud service and the upstream data of the Internet of Things (IoT). In Industrial IoT, edge computing provides added benefits of agility, real-time processing, and autonomy to create value for intelligent manufacturing.

Market Categories and Deployment Types

To better understand market and vendor positioning, assess how well edge computing solutions position themselves to serve specific market segments.

- **Near edge:** In this category, solutions meet the needs of organizations deploying a large number of data center locations, such as telco locations, cellular network sites, or large manufacturers with many plants or retail stores.
- **Far edge:** Here, offerings can be deployed in non-traditional locations for servers. These solutions will support smaller form-factor machines and will not require data center services such as server racks, air conditioning, and protected power.

Edge computing has existed in the manufacturing sector for a few years. Manufacturing plants have significant processing power on-premises, whether it be in programmable logic controllers (PLCs), the machines themselves, or an on-premises data center. Edge computing fits into this broader context by allowing manufacturers to use more flexible, standard hardware and software to access and share data relevant to their manufacturing processes.

Green Edge Computing

Scale Computing solutions are very efficient compared to traditional IT infrastructure environments. Innovative servers, specially designed for the edge, are on the leading edge regarding power savings.

New IT infrastructure solutions like the CO₂ emission friendly [Intel NUC Enterprise Edge Compute](#) (EEC) edition built with Scale Computing (introduced in May 2022) are designed for the edge. Lightweight software from Scale Computing is packaged on the Intel NUC EEC product for uncontrolled, non-IT environments and managed centrally through the Scale Computing user interface, making it easy to run applications anywhere they are needed while also reducing workloads for IT teams.

With the award-winning Scale Computing software built in, the Intel NUC EEC is an out-of-the-box solution ideal for running any critical applications on-premises with immediate high availability. The Scale Computing software includes centralized fleet management, rolling upgrades, ease of deployment and VM creation, and automated intelligence for maximum uptime.

The Intel NUC EEC edge computing servers are green in “all” areas:

- **Power consumption.** Internal testing shows the power consumption of the Intel NUC EEC servers is just 10-30% compared to traditional 19” rack servers, enormous savings when you implement hundreds of clusters across the factory floor at various sites. Just on power alone, savings can add up to \$500 (depending on country and electricity contracts in place) per server per year.
- **Production.** The production of the Intel NUC EEC is a fraction compared to a rack server. Typical CO2 emission reduction due to the manufacturing cost and material usage is 75% or more.
- **Transportation.** You can ship 8-10 Intel NUC’s EEC servers for the same price as one rack server. This directly results in 75% or more savings in fuel consumption needed for transportation.
- **Zero-Touch Provisioning (ZTP).** Scale Computing Fleet Manager with ZTP is the fastest time to application deployment. ZTP allows administrators to centrally configure clusters in SC//Fleet Manager before nodes arrive on-premises, decreasing the installation time of nodes and clusters by 75%. ZTP minimizes the disruption of installing new infrastructure at the edge.

With SC//Fleet Manager’s ZTP, you avoid:

- » Manual on-site configuration
- » Configuration inconsistencies with many devices
- » Potential for human error
- » Sending expensive IT resources
- » Wasting productivity better spent on other tasks
- » Needing IT tools onsite (keyboard/monitor)



- **Eliminating on-site repairs.** With the Intel NUC EEC, you send the server by mail, something you would never consider with a traditional server. Engineers no longer need to travel to a location for on-premises repairs. The cost of time and petrol consumed compared to onsite repairs is eliminated.



The Harrison Story

More than 100 years ago, J.W. Harrison established the National Car Coupler Company in Attica, Indiana. Today, that company is Harrison Steel Castings Company, a world leader in the production of highly engineered carbon and low/medium alloy steel castings with customers like Caterpillar and many others.

Harrison Steel operates a state-of-the-art, technology-driven international engineering and manufacturing organization. The physical USA plant encompasses 700,000 square feet under one roof and is capable of castings ranging from 400 to 18,000 pounds net casting weight. Sophisticated engineering software, tightly controlled manufacturing processes, superior quality programs, a motivated workforce, an emphasis on safety, and a continuous improvement culture all inform and fuel Harrison Steel's business.

Challenges

Shane Rogers, IT Director of Harrison Steel had an open mind when looking for technologies that would help to optimize the manufacturing operations and business system processes, improving its competitive abilities.

For years, Harrison Steel collected data. Data points collected from PLC devices have been there for decades. It used a sneakernet to grab these data points, but it never looked at data for other areas like PMs or the quality of the parts. As more engineers and technology got involved, Harrison Steel realized it could use that data to make better quality parts. All of this was based on human data collection, starting 15-20 years ago. The IOPS of this data collection, however, was extremely low, making automated data collection and real-time analytics impossible. This was a major concern for Harrison Steel.

Working with machines that have “tons” of sensors, Harrison Steel started to look at what more it could do with the data it could collect and what type of networking it would need. Rogers began to realize, after the company coupled the first machines on its network to automate data collection, that it was losing data and overloading the network. It had to stop because there was simply not enough bandwidth.

Harrison Steel runs very modern sensor technology but also still has robots running, for example, Windows CE. The problem with these older systems is that they cannot connect to anything other than just to dump data, which is still a manual process.

Research

Harrison Steel started to research the market for a separate manufacturing production network such as Rockwell and others just to collect OT data from the machines and ovens. One of the objections to adopting this was the high cost of expanding or upgrading the network when it would need more bandwidth. As more and more machines, and such vast amounts of sensory data come in every year, the cost and flexibility became a concern.

Halfway through the research process, Rogers learned that Scale Computing, a vendor he was already familiar with, had introduced the HE100 Series, a new category of edge computing devices. Harrison Steel was running its data center operations on various Scale Computing HyperCore clusters for years and with great satisfaction.

Light bulb moment

For Rogers, this announcement was a light bulb turning on. He realized he could place little boxes with three-node clusters in a full, highly available setup close to the machines. With machine data stored on the cluster, no network congestion would occur. It made real-time analytics possible. At regular intervals, the cluster would be offloaded over the network back to the data center in a controlled way.

Harrison Steel has implemented the HE150 clusters running SC//HyperCore and Rockwell Automation in the factory. With these clusters, repeated across the factory floor and connected to all types of machinery and robots, extensive data collection and analysis of mold fill time, flow rates, vibration, and temperature data (among others) takes place. This way, Harrison Steel can make better-informed business decisions at greater speed.

Another great advantage of the setup is the seamless integration with the clusters in their data center running core applications like the Electronic Production System (PMS), Reverse Scheduler, ERP, and InforEAM. Although Harrison Steel uses Scale Computing in the data center as well, the Scale Computing edge solutions can also work with other data center systems and infrastructure.

Savings

From an ROI perspective, building out a dedicated network for the data collection would have cost Harrison \$500,000 to \$1M. This is a great deal more than the dozen HE150 edge clusters implemented (ranging from \$7-12,000 per 3-node cluster, including the multi-year SC//HyperCore subscriptions). Besides these noticeable savings, the self-healing technology embedded in the SC//HyperCore system minimizes ongoing management hours.



Storage Cost Reduction

Implementing Industry 4.0 manufacturing processes means large amounts of data require substantial storage. Traditional local storage options are inconvenient, and cloud storage can be expensive. Storing data at the edge is more efficient. The edge becomes the gateway, allowing local data analytics, sending results or summarized data to the cloud, and saving central storage costs. Storing data on edge computing devices enables manufacturers to select what gets sent to the cloud.

The result is a segmented, organized, and affordable data storage solution.

Scale Computing appliances efficiently use integrated hybrid, all-flash or NVMe storage. Starting as small as 250 GB, a 3-node micro edge cluster for small, highly available OT applications used on the factory floor to edge systems can run 10s of virtual machines.

Central Monitoring and Management

With Scale Computing Platform, you can run applications and process data outside centralized data centers, at the edge of your network closest to where that data is created and utilized. Leverage Scale Computing Fleet Manager to centrally monitor and manage hundreds or thousands of distributed edge infrastructure deployments with few or no onsite IT personnel.

SC//Platform consists of two key applications: **SC//HyperCore**, the hyperconverged IT infrastructure server environment and **SC//Fleet Manager**, no matter if you need to manage just a few or 50,000 clusters or locations.

The award-winning Scale Computing HyperCore delivers on-premises edge computing with high availability and disaster recovery to remote locations at an affordable entry-level cost. All servers can be deployed quickly, managed locally or remotely, and can almost instantly self-heal. Benefit from affordable edge computing infrastructure without sacrificing reliability, ease of deployment, and ease of use.

SC//Fleet Manager is the first cloud-hosted monitoring and management tool built for hyperconverged edge computing infrastructure at scale. It's now easier than ever to securely monitor and manage your entire fleet of clusters running SC//HyperCore. For IT Managers with multiple clusters, Scale Computing's industry-leading low maintenance time will be reduced by 50% or more.

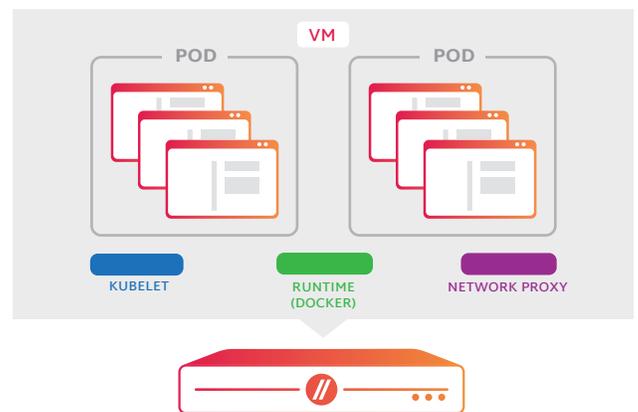


Modernizing with Low-Code and Containers

To modernize operational technologies, manufacturers adopt low-code and container technologies to boost their application efficiency more frequently to modernize operational technologies.

We often hear customers ask if it is possible to run **containers and Kubernetes** on SC//Platform. The answer is, yes! We fully expect the adoption of containers to increase over the coming years as more and more application vendors take advantage of the provided portability. The more complex these container-based applications become, the more likely they will be deployed with an orchestration tool such as Kubernetes.

- **Containers.** A container bundles together the entire runtime environment, including the application, its dependencies, libraries/binaries, and the configuration files needed to run it. Containers are outstanding in their ability to provide consistency from one environment to another (e.g., from the developer's test environment to a VM running on SC//Platform).
- **Kubernetes.** Kubernetes is a widely adopted, open-source orchestration tool for managing and running containers in production environments. It provides a framework for containers' resiliency, networking, load-balancing, and scaling. It is common to see these tools as a part of a larger CI/CD (continuous integration and continuous delivery) pipeline.
- **Low-code applications.** Low-code platforms give you the tools to assemble the right business experience across domains, keeping your digital thread unbroken from strategy to planning to design to manufacturing. Across the value chain, manufacturers see the possibilities for new efficiencies, new connections, and new ways of engagement.



Conclusion

Manufacturers who upgrade their existing factory IT infrastructure or those designing a new factory should consider Scale Computing Platform.

Scale Computing combines highly available and extremely easy to manage clusters with a low TCO, from extremely small to large edge implementations, and easily adapts to continuously changing requirements.

For more information, contact your local Scale Computing representative, email Scale Computing at info@scalecomputing.com or visit: scalecomputing.com/manufacturing-edge-hci-solutions

Corporate Headquarters
525 S. Meridian Street - 3E
Indianapolis, IN 46225
P. +1 317-856-9959
scalecomputing.com

EMEA B.V.
Europalaan 28-D
5232BC Den Bosch
Pays-Bas
emea@scalecomputing.com

