Stratus Special Edition

Edge Computing



Build edge computing expertise

Compete in an "always on" world

Build your business case

Brought to you by





Lawrence Miller, CISSP

About Stratus

For leaders digitally transforming their operations in order to drive predictable, peak performance with minimal risk, Stratus ensures the continuous availability of business-critical applications by delivering zero-touch Edge Computing platforms that are simple to deploy and maintain, protected from interruptions and threats, and autonomous. For 40 years, Stratus has provided reliable and redundant zero-touch computing, enabling global Fortune 500 companies and small-to-medium sized businesses to securely and remotely turn data into actionable intelligence at the Edge, cloud and data center – driving uptime and effciency.

Stratus platforms are:

Simple – Our Edge Computing platforms are easy to install, deploy and manage across applications and existing infrastructure as well as scale to new areas of operation.

Protected – We mitigate operational, fnancial and reputational risk by ensuring "always on" availability and securing data from cyber threats or data loss.

Autonomous – Our reliable, rugged, proven zero-touch technology runs in any environment without human monitoring, maintenance, repairs or support.

For more information, please visit **www.stratus.com** or follow on **Twitter** @StratusAlwaysOn and LinkedIn @StratusTechnologies.



Edge Computing

Stratus Special Edition

by Lawrence Miller, CISSP



These materials are @ 2020 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited

Edge Computing For Dummies[®], Stratus Special Edition

Published by John Wiley & Sons, Inc. 111 River St. Hoboken, NJ 07030-5774 www.wiley.com

Copyright © 2020 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. States and other countries, and may not be used without written permission. Stratus and ftServer are registered trademarks of Stratus Technologies. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom For Dummies book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/ custompub. For information about licensing the For Dummies brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-70803-2 (pbk); ISBN 978-1-119-70805-6 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor: Jennifer Bingham Acquisitions Editor: Katie Mohr Editorial Manager: Rev Mengle Business Development Representative: Karen Hattan Production Editor: Tamilmani Varadharaj

Introduction

o matter the industry, companies today are under constant pressure to optimize business performance and productivity. Whether it's an assembly line, an oil and gas pipeline, a water treatment plant, a retailer with multiple stores, or a pharmaceutical lab, almost every business needs to turn data into actionable results.

In addition to monitoring and control, edge computing provides the capability to collect and analyze data at or near its source, rather than sending it to a central server or cloud for processing and analysis. This capability ensures that potential problems can be quickly detected and corrected before they cause lost productivity and costly downtime. Edge computing doesn't replace onpremises data centers or cloud environments — it complements them. The edge delivers the right information to the right place to enable real-time decisions and actions.

About This Book

This book consists of four chapters that explore:

- The modern trends that drive the need for edge computing (Chapter 1)
- Edge computing challenges and how an edge computing platform can solve them (Chapter 2)
- >> Industry use cases for edge computing (Chapter 3)
- How to assess your organization's edge computing maturity level (Chapter 4)

Each chapter is written to stand on its own, so if you see a topic that piques your interest, feel free to jump ahead to that chapter. You can read this book in any order that suits you (though we don't recommend upside down or backwards).

Foolish Assumptions

It's been said that most assumptions have outlived their uselessness, but this book assumes a few things nonetheless!

Introduction 1

Mainly, it assumes that you are a director, manager, operations specialist, or IT specialist in an industry such as food and beverage, manufacturing, oil and gas, pharmaceuticals, power and utilities, retail, transportation, water and wastewater, buildings management, or the equipment manufacturers who support these industries. With such a broad audience across diverse industries, we've written this book primarily for readers who perhaps have only a cursory or general understanding of edge computing challenges and solutions.

If any of these assumptions describe you, then this is the book for you! If none of these assumptions describe you, keep reading anyway! It's a great book and after reading it, your knowledge of edge computing will be on the leading edge!

Icons Used in This Book

Throughout this book, you will find special icons to call attention to important information. Here's what to expect.



This icon points out important information you should commit to your noggin — along with anniversaries and birthdays!



REMEMBER

This icon explains the jargon beneath the jargon and is the stuff legends — well, legendary nerds — are made of.



Tips are appreciated, but never expected. This icon flags useful nuggets of information.

Beyond the Book

There's only so much we can cover in this short book, so if you want to learn more, check out www.stratus.com/edge-computing.

- » Defining edge computing and recognizing why businesses need it
- » Competing in an always-on business environment.
- » Looking at virtualization, cloud, and mobile computing

Chapter **1** The Role of Edge Computing in Digital Transformation

n this chapter, you'll learn about key digital transformation trends and technologies — such as virtualization, cloud, and mobile computing — and how these trends have enabled the always-on enterprise. You'll also learn what edge computing is and why it's important to businesses in every industry.

Understanding What Edge Computing Is

The era of edge computing is bridging the gaps between information technology (IT), business managers, and operational technology (OT) by collecting, managing, and analyzing data used across the organization to increase efficiencies. With the availability of sensors and greater connectivity giving rise to the Internet of Things (IoT), digital transformation — the process of integrating digital technology into areas of a business, resulting in a transformation of both business operations and business culture norms — is the next imperative. In this next wave, data is

CHAPTER 1 The Role of Edge Computing in Digital Transformation 3

the new currency. Analyzed data is knowledge and this knowledge drives successful business transformation. It empowers decision makers and helps organizations increase productivity and create better products.

Edge computing is a distributed computing model in which computing takes place near the physical location where data is being collected and analyzed, rather than on a centralized server or in the cloud. This new infrastructure involves sensors to collect data and edge servers to securely process data in real time onsite, while also connecting other devices, like laptops and smartphones, to the network.

Edge computing is important because it creates new and improved ways for industrial and enterprise-level businesses to maximize operational efficiency, improve performance and safety, automate all core business processes, and ensure always-on availability. It is a leading method to achieve full digital transformation of your business.

Here are a few of the things edge computing can do:

- Enhance solutions by addressing cloud computing issues such as performance, latency, bandwidth, security, and proximity.
- Power the next industrial revolution, transforming manufacturing and services.
- Optimize data capture and analysis at the edge to create actionable business intelligence.
- Create a flexible, scalable, secure, and more automated technology, systems, and core business process environment.
- Promote an agile business ecosystem that is more efficient, performs faster, saves costs, and is easier to manage and maintain.

Increasing computing power at the edge is the foundation needed to establish autonomous systems, enabling companies to increase efficiency and productivity while enabling personnel to focus on higher-value activities within the operation.



A leading global research firm forecasts that by 2022, approximately 75 percent of all data will need analysis and action at the edge.



By 2025, 175 zettabytes (or 175 trillion gigabytes) of data will be generated around the globe and edge devices will create more than 90 zettabytes of that data, according to IDC (*Data Age 2025: The Digitization of the World, From Edge to Core*).

Edge computing works hand-in-hand with the cloud to provide a flexible solution based on the data collection and analysis needs of each organization. For real-time collection and analysis, the edge is ideal for certain workloads. At the same time, the cloud can provide a centralized location for large-scale analytics. Together they provide real-time and longer term insights into performance and power initiatives like machine learning and asset performance management.

Delivering Performance in an Always-On World

As technologies become more intelligent, digital transformation is becoming a practical reality. The path to complete digital transformation is complex — applications and technologies are needed to analyze data at the edge of company networks where there are limited or no IT resources. We live in an always-on world and whether your applications run in your data center or on the plant floor, in a control room or any remote location, you can't afford for them to be down, or to be too complicated to operate and manage. With these escalating demands comes greater pressure to prevent even the smallest amount of application downtime. When it comes to your key business applications, even a little downtime at the wrong time can spell disaster. Traditionally, when hardware fails, applications go down.

Application and business operations downtime is an issue that plagues the vast majority of organizations. When these applications go down, many — if not all — of an organization's most business-critical operations become unavailable, and the cost of being unable to do business increases by the second.

Broadly, downtime costs consist of:

Business costs: These are the first costs that come to mind for most people. Overtime and remedial labor costs all add up during an outage. Other business costs include lost sales (and future repeat business), lost inventory and lost work in progress, potential legal penalties for not delivering on service-level agreements (SLAs), and litigation costs due to third parties seeking compensation for losses incurred during a system outage.

- Productivity costs: During an outage, employees can't perform their regular duties. The impact of this idle time varies by industry. For example, in an office environment, an employee may not be able to access the Internet but can work on a desktop spreadsheet program, so perhaps his or her productivity would be cut in half. But in a manufacturing environment, if the line stops, employees may be 100 percent idle.
- Recovery costs: These costs include the price paid to repair the system, IT staff overtime, and third-party consultants or technicians needed to restore services. There's also an opportunity cost sacrificed when IT needs to focus on system recovery instead of working on other critical projects for the organization.
- Customer loss: The effects of indirect costs can be felt long after an outage is resolved. Previously loyal customers can lose faith and take their business to competitors. Once a company is seen by its customers as unreliable, it can be very difficult to undo the perception.
- Reputation damage: Bad publicity can cause major damage to an organization — and not just large ones. It's true that the traditional press loves a good headline about bad news at a big company. But what can a complaint on Twitter or a negative post on Facebook cost you? Convergys found that one bad tweet can cost a company 30 customers. And while industry websites and bloggers don't always have a large audience base, they do have the rapt attention of your target market. This means that one negative blog post about your company can make a huge impression on your customers and prospects.
- Shareholder value impact: Bad press can also devalue a company's stock and reduce its market capitalization. Especially in shaky economic times, the stock market reacts to negative press about a company, even more so if the news is about a significant sales loss an event that is entirely possible when servers go down.



In the world of the Internet of Things (IoT), data is the currency that drives decision making and results. The growth of data and the number of devices is at an unprecedented stage. It is this growth and need for data and analysis that is truly fueling technology trends such as edge computing and the digital transformation of business operations. Technology concepts, such as computing at the edge and virtualization, depend on the devices and applications creating or analyzing this data to always be available. Downtime is the sworn enemy of core business applications. What is the value of these applications and how much does it cost you if they are inaccessible? Without knowing the true cost of downtime, your organization can't properly and costeffectively protect itself.

Digital Transformation and Other Trends

Digital transformation is a key component of every modern business strategy today, regardless of industry. Digital transformation initiatives enable key capabilities and deliver benefits including business agility, operational efficiency, advanced analytics, and process automation, among others. Several technology trends driving digital transformation include:

- Cloud: The cloud enables on-demand self-service access to computing, storage, and other virtual resources that can rapidly scale based on demand in a pay-as-you-go consumption model. The three main cloud service models are:
 - **Software as a service (SaaS):** The customer is provided access to an application that is hosted and managed in the cloud by the service provider. The provider is only responsible for managing access to the application and the data in the application.
 - Platform as a service (PaaS): The customer can deploy applications in the cloud environment and manage certain configuration settings for the platform (for example, a SQL database instance), but the service provider is responsible for managing the compute, storage, networking, and operating system components, as well as the underlying physical infrastructure (servers, storage arrays, and network switches) and the data center.

- Infrastructure as a service (laaS): The customer manages the compute, storage, networking, operating system, and application components, but the service provider is responsible for managing the underlying physical infrastructure and the data center.
- Remote monitoring and control: Safely maintaining or scaling operations requires capabilities that enable greater and more reliable remote and secure monitoring and control of critical systems. This capability enables manufacturers at a local facility or plant-level — to remotely drive operational efficiency and performance safely, while freeing limited personnel to focus on higher-value activities.
- Virtualization: The concept of virtualization was a breakthrough in computer technology when it was developed more than 40 years ago to enable shared use of relatively limited computing resources. Virtualization allows the capabilities of a physical machine to be distributed across multiple environments and takes several forms, which reduces the need for CapEx spending on server hardware, improves utilization of available resources, and reduces OpEx costs such as maintenance, power, cooling, and equipment rack space. Some common types of virtualization include:
 - **Desktop virtualization** is used to create one consistent environment that is simultaneously shared with multiple physical machines.
 - Server virtualization allows a server to be partitioned so that multiple functions can be run simultaneously on a single physical machine.
 - Operating system virtualization can be used so that one physical machine can run multiple operating systems (such as Windows and Linux) on the same physical machine in separate operating system environments (OSEs).



There are many other types of virtualization such as application, storage, and network virtualization, among others. These virtualization types are beyond the scope of this book.

These trends have created a highly competitive business culture in which companies in all industries are expected to always be up and running, ready to serve their customers and be productive it's a 24/7/365, always-on world.

- » Recognizing business challenges at the edge
- » Addressing challenges with an edge computing platform

Chapter **2** Addressing Pain Points at the Edge

n this chapter, you'll learn about the challenges that businesses must deal with at the edge and how an edge computing platform helps solve these challenges.

Understanding Real-World Challenges at the Edge

Unlike a corporate headquarters location that may have an onpremises data center, high-speed Internet connectivity for access to the cloud, and onsite IT staff, remote edge locations have relatively limited resources. Compared to the corporate headquarters, remote edge locations may experience challenges such as:

Less IT equipment and space: Few organizations can afford to install redundant high-end servers, storage, and networking equipment in all their remote locations. Instead, they may install critical applications on commodity servers (and in some cases, desktop PCs) that don't have redundant components (such as power supplies and hard drives). Worse still, edge locations may not have a server room or

CHAPTER 2 Addressing Pain Points at the Edge 9

other environmentally controlled space for IT equipment. Oftentimes, critical IT equipment may be installed in an industrial or otherwise unsecure area that may not have adequate power, cooling, or ventilation, leading to a higher risk of equipment failure.

- Slower and more expensive Internet connectivity: Many remote locations today are connected to the corporate data center and the cloud via a broadband Internet connection from a local Internet service provider (ISP). These connections are typically "best effort" connections that offer asynchronous connectivity (different upload and download speeds) with minimal service-level agreements (SLAs). If broadband Internet isn't available in certain remote locations, more expensive private circuits may need to be installed, or the site may need to use 3G or 4G cellular connections with relatively limited — and costly — data plans. For organizations that don't deploy any on-premises infrastructure at their remote locations, an Internet outage can mean a complete work stoppage at a remote location.
- Fewer IT staff: Specialized IT staff are generally located at the corporate headquarters. Remote locations may have little or no onsite IT staff, which means a full stop when systems are down. For sites that have no onsite IT staff, remote access from the corporate headquarters to commodity IT equipment in the edge location can be challenging, with limited remote management tools and no hands-on capabilities.
- Different IT skillsets and priorities: Professionals don't have the time or skillsets to manage computing infrastructure. Thus, IT must deliver platforms at the edge that are simple to use and maintain by staff.

Ultimately, these challenges cause lost productivity due to application performance issues and equipment downtime (see Chapter 1 to learn more about the cost of downtime). Businesses need an effective solution to successfully address these and other challenges at the edge.

Solving Challenges with an Edge Computing Platform

Edge computing enables monitoring and control at the edge and supports innovation and new ways to do business. The ability to monitor and control critical systems at the edge is imperative. At the same time, the capability to collect and analyze data in real time enables data-driven decision making for better outcomes. Potential problems can be detected and corrected more quickly than if the data were sent to a central data center or cloud for processing and analysis. As edge computing grows in capabilities and adoption, more organizations are looking at how to integrate the edge with their existing architectures.

However, edge computing requires a thoughtful architecture and implementation, which can be a challenge without the right expertise. Having multiple systems collecting and analyzing data can mean more sites that need to be configured and monitored, which adds complexity to the environment. But having too few systems can result in critical data being missed.



An edge computing platform solves the following challenges:

- >> Monitoring and control at the edge
- Internet of Things (IoT) and Industrial Internet of Things (IIoT) data collection and analysis
- Cloud/enterprise IT architecture issues such as performance, latency, bandwidth, security, and proximity

EXPLORING STRATUS EDGE COMPUTING PLATFORMS

Whether your applications run in your data center or on the plant floor, you can't afford downtime or operational/management complexity.

Stratus offers a simple, protected, and autonomous cost-effective way to deploy and maintain a continuously available processing

(continued)

CHAPTER 2 Addressing Pain Points at the Edge 11

(continued)

environment for your virtualized applications — both in the data center and at the network's edge.

By removing the complexity from availability, Stratus makes it easy for you to protect your critical workloads, efficiently manage your systems, and generate insights from business and operational data.

ztC Edge: Zero-touch computing for your critical edge applications

ztC Edge is a zero-touch, reliable, secure, highly automated computing platform, specifically designed for edge environments. Purpose-built for non-IT specialists, ztC Edge is easy to set up, configure, and manage. With a built-in virtualization and availability layer, automated data protection, and application recovery, ztC Edge significantly reduces the dependence on IT for virtualized computing at the edge. Its selfprotecting and self-monitoring features help reduce unplanned downtime and ensure the continuous availability of business-critical applications.

ztC Edge is a versatile, secure, highly automated computing platform that helps understaffed organizations to increase productivity and reduce risk. A redundant pair of rugged hot-swappable nodes, with built-in virtualization, automated recovery, and cloud-based systems health management services, ztC Edge helps business managers and engineers deliver IoT and IIoT applications quickly, easily, and reliably.

ftServer: A fully integrated continuously available platform that's easy to deploy, manage, and service

Stratus ftServer is a fully integrated, continuously available hardware and software solution for running tier 1 business-critical workloads, manufacturing operations, and control applications. Offering better than five-nines availability, ftServer eliminates unplanned downtime and data loss, and ensures all transactions are processed and applications are available, helping companies maximize revenue, production quality, and productivity. Unlike other high availability solutions, ftServer is easy to deploy, manage, and service, making it ideal for remote locations or environments with IT resource constraints. A single, turnkey solution leverages industry-standard components to deliver industry leading reliability, availability, and serviceability (RAS), accelerating time to value, simplifying management, and lowering total cost of ownership.

IN THIS CHAPTER

- » Drilling into oil and gas at the edge
- » Curing bio manufacturing challenges at the edge
- » Building the case for smart machines
- » Getting smart about Smart Building management
- » Looking at discrete manufacturing
- » Making the case for an edge computing platform

Chapter **3** Exploring Edge Computing Use Cases

n this chapter, you learn about different industry use cases for edge computing and real-world customer success stories, as well as how to build a business case for deploying edge computing platforms in your organization's remote edge locations.

Oil and Gas

The global oil and gas industry is always in a cycle of constrained economics driven by price pressure and increased regulation. The challenge is to deliver efficiencies and to control costs in the face of these increasing challenges. This all needs to be done in a sustainable way. Automation systems are a key area where investment can generate rapid returns through a variety of optimizations. However, whether you're a gathering company, a pipeline transmission company, a refinery, or a distribution company with retail outlets — when your operational systems fail, you are exposed to safety threats, compliance violations, inaccurate balancing, and a host of other issues. These disruptions can have major economic, safety, and societal impacts.

The supervisory and operational layers of an oil and gas process control system are the most critical areas for maintaining efficient and cost-effective operation. Failure of these systems, for even short periods of unplanned downtime, creates a whole host of problems with resulting economic, regulatory, environmental, and even life-threatening impacts. Many oil and gas companies have tried but failed to solve these problems with conventional servers running critical applications such as supervisory control and data acquisition (SCADA) systems in multiple locations. Downtime problems continually plague this approach and it doesn't scale. Additionally, not all locations have onsite technicians to maintain servers and other IT equipment. The solution is a reliable edge computing platform.

An edge computing platform allows you to virtualize multiple applications and run them on one platform. Instead of a one-toone server-to-application ratio, virtualization supports a oneto-many server-to-application ratio, saving you infrastructure costs. But the more applications you have running on any one platform, the more critical that platform becomes. An edge computing platform must be designed for high availability, resiliency, and centralized management.

CASE STUDY

An oil and gas pipeline company that transports natural gas to a variety of industries including hospitals and manufacturing, as well as households, needed to ensure the availability of their infrastructure. If the system were to go down, the operator would be running blind with no way to control their units or to see what's going on with the units at that point.

The Stratus solution provides hardware redundancy and configuration capabilities with a fault-tolerant computing platform at each compressor station (every 50mto 100 miles) to provide the human-machine interface (HMI) capability for the operators. This allows them to do their job and keep the facility running due to overall system uptime.

Bio Manufacturing

Bio (life sciences) manufacturing companies rely on consistency and automation to meet strict regulatory requirements and rigid operational standards. Any small amount of unplanned downtime can result in a decline in product quality and brand reliability.

In applications such as manufacturing execution systems (MES), for example, unplanned downtime results in lost production time. In regulated industries like pharmaceuticals, loss of data and/or control can compromise the integrity of a batch record and require in-process product to be destroyed. Even minor system interruptions can call into question the value of an IT solution.

CASE STUDY

Ensuring uptime and preventing data loss is a key challenge in biotech because batches are typically manufactured in thousands of vials. Because it's impossible to take each vial back to the lab and test the quality attributes of each and every vial, the U.S. Food and Drug Administration (FDA) requires quality to be built into the design of biotech manufacturing systems. Every process that's involved in the manufacture of a drug has to be validated. If you don't have the systems and data to support the continuous monitoring of the product, then according to regulatory guidelines you can't release that product to the market, which can represent millions of dollars.

A leader in data acquisition systems that is heavily involved in the biotech manufacturing industry provides very high accuracy sensors and measurement systems that are used to validate each of these processes. To provide the required levels of reliability, accuracy, and uptime, the company has created an integrated package validated with their own sensors and input/output (I/O) hardware, in conjunction with Stratus fault-tolerant platforms.

CHAPTER 3 Exploring Edge Computing Use Cases 15

Equipment (Machine) Building

In today's data-driven operations where Industry 4.0, digital transformation, and the Industrial Internet of Things (IIoT) are fast becoming common initiatives, the need for smart machines and equipment is critical. Manufacturers are quickly embarking on digital transformation journeys. They are automating and modernizing their operations using innovative machines to create smarter factories.

In addition to monitoring and controlling applications, machine builders' customers require machines or equipment that enable them to achieve key business results using information or data that is visible and actionable from the edge to the data center and in the cloud.



Smart factories are connected, digitalized plants that depend on integrated supply chain solutions, optimized asset performance management, intelligent automation, and smart machines.

Smart machine and equipment builders who want to separate themselves from the competition must evolve to achieve digital transformation. Digitalization is not a choice. Machine builders need to adopt this direction in order to create smart machines. Smart equipment is the foundation of Smart Manufacturing. These machines — equipped with Internet of Things (IoT) capable sensors, intelligent controllers, and HMI/SCADA systems take advantage of the innovations in digital technology. Realizing the full benefit of these technologies, while ensuring their safe and reliable operation, requires edge computing.



The edge is the physical location where operations take place, where products are produced, and where operational data exists and can be collected and analyzed, allowing for improvements in production. It's where machines are installed — on the plant floor or production line.

There are a variety of computing platforms available for machine and equipment builders. The best choice reliably meets the requirements of your customers' operational environment and is highly scalable, able to handle the complexity and size of your machine or piece of equipment.

Purpose-built edge computing platforms combine powerful computing solutions with inherent redundancy and built-in virtualization that you can embed in your machines. They improve machine monitoring and control and can be used in a myriad of applications, including historian servers and analytics for collecting information from IIoT devices and sensors passing data to the cloud and the enterprise.

Building Management

As the environmental and security demands on building infrastructures get increasingly complex, there is an opportunity to rethink the approach to how buildings are managed. Today's buildings and campuses support a wide range of control systems — access, environmental, video monitoring, energy efficiency, and more. These solutions are typically deployed in a stovepipe fashion on disparate platforms. This lack of integration has led to complexity, high costs, and inefficiencies.



In this next era of the Smart Building, the situation is changing fast. Three key technologies are enabling this transformation:

- IoT devices: In the past, the endpoint devices such as cameras, badge readers, and thermostats in a building were proprietary and expensive. The introduction of low-cost devices coupled with the adoption of consistent communications standards is resulting in more devices and increased integration between systems.
- Virtualization: One of the things that has held back virtualization and the integration of building technologies has been the requirement for each solution to have its own unique infrastructure. Now that the devices are IoT-based, there is an opportunity to simplify the building infrastructure and lower costs through virtualization.
- Analytics: The analysis of data from the increased number of end devices can lead to changes to business processes to drive efficiencies and manage cost.

CHAPTER 3 Exploring Edge Computing Use Cases 17

CASE STUDY

Over the past 30 years, a large international airport has undertaken numerous improvement and modernization projects in various stages in one part of the airport or the other. As a result, various solutions were installed using more than nine different and separate controls interfaces to manage and modernize the different parts of the airport. Over time, these silos became extremely difficult to operate, maintain, and support.

They were able to integrate the different islands of automation into a complete and centrally connected facilities and building management system. They decided to implement a single system that could view the entire facility by unifying the operating silos.

In effect, they would build a Unified Operations Center (UOC) that would have the capability to integrate all the currently installed and different technologies into one system. They also wanted to make it scalable so it would be easier to connect to new solutions, devices, or systems in the future. They were able to integrate the different islands of automation into a complete and centrally connected facilities and building management system.

The System of Systems Unified Operations Center approach unifies all the silos into one system, with remote clients to operate and maintain any part of the airport from anywhere. The UOC provides overall visibility into the facility infrastructure, using the open and scalable architecture of the system platform and the reliability and redundancy of always-on Stratus fault-tolerant platforms to prevent the downtime of critical business applications, ensuring continuous availability, even if underlying hardware components fail. Benefits of the approach include:

- Vendor agnostic: The airport doesn't depend on any single vendor, original equipment manufacturer (OEM), or manufacturer.
- **Instantaneous alarms:** All events and alarms are integrated; the immediate notification enables airport personnel to immediately identify problems in any location.
- Scalable system based on standards: Users can build, test, deploy, maintain, and scale industrial applications by sharing a standards-based application development environment.

These three advances are enabling the Smart Building, but the true core of Smart Building technology is the control applications that collect the IoT device data and transform that data into analytics. These control applications may also have a role in managing the IoT endpoint devices and enabling audits and compliance in critical locations. An edge computing platform consolidates and virtualizes these applications on a single, highly available and resilient platform that eliminates the numerous disparate systems of the past.

Discrete Manufacturing

Manufacturers are adopting new IT solutions on an unprecedented scale to meet productivity, efficiency, quality, and regulatory compliance goals.

In the manufacturing industry, the bottom line depends on your ability to keep your plant infrastructure up and running at all times. Downtime in SCADA, HMI, historians and other critical industrial automation systems can mean lower yields, lost revenue, regulatory fines, compromised quality, or a damaged reputation. In critical industrial automation applications, any disruption can mean millions of dollars of equipment and product are at risk as the site is flying blind. Automation is empowering industrial enterprises to dramatically increase productivity. Yet modernizing industrial automation requires underlying technologies designed to improve efficiency and quality, while eliminating unplanned downtime.



To maximize productivity and quality you need operationally simple continuous availability and continuous visibility.

One of the top benefits of implementing edge computing is the capability to collect and analyze data where it is collected, catching and correcting problems that might not be identified as quickly if the data were to be sent to a central server or cloud for processing and analysis. Keeping data onsite also reduces the security risk associated with porting data, which can be important in financial organizations, for example. It also reduces bandwidth costs by processing some data onsite, rather than sending all data to a cloud or central server.

CASE STUDY

When the IT system for a large aluminum rolling and casting plant was discontinued a few years ago, the organization set out to find a replacement system. The management and control of both its pit and pusher furnaces were dependent upon it. After considering a cluster system, the manufacturer eventually chose a system built on fault-tolerant servers from Stratus.

Because of its consistently positive experiences, the manufacturer has now expanded its use of Stratus fault-tolerant platforms. Today, there are four fault-tolerant servers in the hot-rolling mill for the pit and push furnaces, which run around the cloud 365 days a year. These systems have run for six years without a single second of downtime.

Successful edge computing requires a thoughtful architecture and implementation, which can be a challenge without the right expertise. Having multitudes of sites collecting and analyzing data can mean more sites that need to be configured and monitored, adding complexity. Having too few can mean critical data is missed. Decentralized locations can also mean fewer technical personnel onsite, meaning nontechnical operations staff may be called in to troubleshoot. These challenges can be addressed by working with knowledgeable system integrators and using the right edge technology.

Building Your Business Case for an Edge Computing Platform

Increasing computing power at the edge is the foundation needed to establish simple, protected, autonomous systems, enabling companies to increase efficiency and productivity while enabling personnel to focus on higher-value activities within the operation.



Edge computing is important because it creates new and improved ways for industrial and enterprise-level businesses to maximize operational efficiency, improve performance and safety, automate all core business processes, and ensure "always-on" availability. It is a leading method to achieve the digital transformation of how you do business. A turnkey edge computing platform uses virtualization and a resilient architecture to deliver real business benefits that can help you make the business case for an edge computing platform. These benefits include:

- Data availability, integrity, and protection: Critical business operations and applications must be always available in order to collect and analyze the data needed to make decisions that impact the business. There can be no downtime. But this data also must be protected, secure, and free from integrity issues. Data security and redundancy ensures that no data is lost. Edge computing platforms, such as the Stratus portfolio, provide the highest levels of availability as part of a simple, protected, and autonomous set of solutions for customers.
- Reduced engineering hours and greatly improved productivity: Instead of performing a single task multiple times on multiple physical machines, the task is only performed once.
- Simplicity: Edge computing platforms are easy to install, deploy, and manage across applications and existing infrastructure, and scale to new areas of operation.
- Stronger competitive advantage: Moving from physical machines to virtual machines provides a competitive advantage for end users. Virtualization protects data analytics and systems in a simple and secure environment that's easy to deploy and helps reduce the number of PCs and software licenses needed, in order to save costs while driving strong performance. This approach also allows for protected data by offering high availability and software fault tolerance.
- Reduced ongoing support burden: In the same way that reducing repetitive tasks saves time, fewer physical machines will reduce the time your IT staff spends troubleshooting hardware problems, managing upgrades and patches, and performing backups.
- Greater flexibility: As customer demands change, edge computing enables the scaling or shifting of operations to meet their needs.

CHAPTER 3 Exploring Edge Computing Use Cases 21

Better security: The addition of sensors, programmable logic controllers (PLCs), video cameras, access control systems, and so on into the environment can create exposures that edge computing can help manage through deployment of security protocols closer to the source. Secure operations can be maintained with modern edge computing platforms, closing the gap created by sensors, interprocess communication (IPCs), standard servers, and other smart edge devices.



Because edge computing is distributed, the security risk is different than in a centralized environment. The security controls found in private data centers or public clouds, like firewalls or antivirus tools, don't automatically transfer. Experts recommend a few simple steps, including hardening each host, real-time network monitoring, encrypting data, and adding physical security measures.

IDC identified edge computing as one of the top ten key drivers for IT over the next five years.



Virtualization is an essential component of edge computing, allowing IT administrators to quickly and easily manage workloads and shift between servers. Virtualization plays a critical role in many edge scenarios, including gateways or microdata centers that process data produced by sensors at the edge, or apps running in containers that are hosted on virtual machines.

When building your business case for an edge computing platform, consider these tips:

Gain agreement and focus: Recognize the various priorities of your operations staff and IT managers, and obtain their buy-in on where to start. Doing so will ensure goals and outcomes are agreed upon, increasing the likelihood of success. Alignment will help the team determine and install the most appropriate and valuable edge computing technology that complements that complements enterprise & industrial systems.



Start with a scalable, agreed upon project in an area that has immediate tangible benefits, rather than an overly complex, long-term project with aspirational benefits.

Assign ownership and expect collaboration: IT, business managers, and staff are experts at their craft. Don't allow either department to consume the responsibilities of the other. Instead, you should assign clear ownership that encourages close collaboration.



Break down siloes by bringing together the right skills and expertise from within and from outside the company to form a team with clearly defined ownership of the project's outcomes.

- Strive for 100 percent overall equipment efficiency (OEE): All businesses, notably in manufacturing, must consider OEE for optimal production productivity. Owning easy-to-use, self-managed, and protected platforms can ensure processes run as fast as possible, have zero unplanned downtime, and the products being produced lack defects and require no rework.
- Establish clear requirements for systems to encompass reliability and resiliency to ensure always-on availability with no downtime or unintended consequences.
- Consider the total cost of ownership (TCO): It is easy to be lured into obtaining short-term low costs without fully considering the TCO. Machine or support failure can lead to significant unexpected expenses and will more than justify the cost of a more reliable solution. Research and select autonomous machines that allow equipment to permanently run without constant onsite management, thereby reducing the risk of errors and more expenditures.



In addition to the cost of downtime, consider licensing and support costs in your TCO calculations.

>> Thoroughly understand the project and your environment: Have a clear scope of your project and recognize the implementation requirements, especially in remote or rugged locations of the facility or plant. Select platforms that provide increased efficiency, simplified security, self-monitoring, dashboard-type performance monitors, and less downtime to function within those environments.

CHAPTER 3 Exploring Edge Computing Use Cases 23



- » Moving from Isolated to Informed
- » Taking the next steps to Insightful and Intelligent
- » Achieving the Invisible enterprise

Chapter **4** Five Stages on the Edge Computing Journey

Stratus has developed the following maturity model to help organizations plan their edge computing journey and deploy a simple, protected, and autonomous edge solution that securely delivers data to applications at the edge, cloud, and data center.

Isolated

Many organizations begin the edge computing journey at this level. Capabilities are limited and include digitally-enabled assets at a single site with static operation and time-based maintenance schedules. The focus is on maximizing the productivity of individual assets that are managed manually (such as data collection via clipboard). These assets are labor-intensive, requiring onsite staff due to the need for direct monitoring of physical assets.



On the technology side, there is no connectivity, and closed or proprietary applications run on interprocess communication (IPC) systems with basic physical security safeguards. Often, these systems have been pieced together over time to meet individual

CHAPTER 4 Five Stages on the Edge Computing Journey 25

project requirements, but there is no consolidated overview or dashboard for monitoring status. This makes management difficult, produces limited actionable data, and impedes real-time decision making. An edge computing platform solves these challenges by using a "System of Systems" approach that unifies all the disparate silos into one system, with remote clients able to operate and maintain any part of the edge computing solution from anywhere.

Informed

At this level, connected assets may be grouped together at a single location but maintenance is still based on a timed schedule and requires manual intervention. There is some flexibility in the operational setup, but there are still challenges, and a choice must often be made between optimizing productivity or optimizing quality — not both, due to the fact that labor needs are still quite high with a lot of hands-on management and decision making required.



On the positive side, there is more interoperability with virtualized servers providing some real-time capabilities. In addition, there are some security measures, mainly air-gap, which keeps data safe by not connecting machines to the Internet while allowing them to connect to other machines at the same physical location.

Insightful

For organizations that have made the leap from Informed to Insightful, connected assets are now in a predictive environment. This means that assets are providing data about performance and maintenance, allowing operators to move away from scheduled updates and maintenance to as-required updates, saving unnecessary work hours. Connected sites have multiple connected assets supporting flexible operations that are driven by analytics and possibly even artificial intelligence (AI).

There is sitewide control (like at the Informed level) and the supply chain is fully integrated. Data collection is semi-automated and real-time analytics allow a focus on both plant productivity

and quality, rather than having to choose one or the other. Labor intensity is also decreasing at this level, with fewer people managing more assets. Operations are supported by IT, and systems are more open with multiple vendors and systems centrally managed. In addition, cloud computing comes into the mix, along with virtualization and containers.

Intelligent

At the Intelligent level, the organization advances its capabilities across the board with enhanced human-to-machine semiautonomous systems. There are multiple connected sites (or plants), building on the connected assets and systems at each location that were optimized at the Insightful stage. This allows operations to be more flexible and dynamic, able to change in response to real-time conditions or changes in requirements. Enterprises are now able to maximize supply chain efficiency, site productivity, and overall quality.

In terms of labor resources, there is increased collaboration between IT and operations, with each side understanding and supporting the needs of the other. The labor force is highly skilled, and each individual can manage many technical resources.

Technology at the Intelligent level features enterprise-level analytics and controls with shared services. There is enterprisewide connectivity with embedded cybersecurity. Machine learning is ramping up, and cloud, edge, and enterprise systems are integrated.

Invisible

The Invisible enterprise is highly automated with fully autonomous, zero-touch, secure machine-to-machine (M2M) systems extending across multiple organizations within the enterprise. Operations are agile and self-adapting, with AI-driven operations at all levels. Areas of focus include enterprise-wide advanced automation across business operations and applications, including all verticals such as Smart Buildings, oil & gas midstream, OEM machine builder, and many others. These processes incorporate

CHAPTER 4 Five Stages on the Edge Computing Journey 27

AI-driven decision making, automated data collection, continuous analytics across all functions, and self-optimizing edge-based applications.



Because the level of automation is so high and AI has been integrated into so many decision-making processes, the level of labor intensity at the Invisible level has decreased significantly. Labor is highly skilled, and while there is a limited amount of supervisory work needed, for the most part, workers can focus on strategy and business-critical applications. Labor resources include supervisory personnel and limited, very high-skilled labor. The workforce is self-managing and focused on high value operations and strategic initiatives.

On the technology side, there is interoperability between the edge at the plant, fog, and the cloud. All the assets can talk to each other and share information, taking action or analyzing at the right stage in the data pipeline. Any security issues are detected and resolved in real time as needed. Technology includes fog computing and everything-to-everything interoperability, real-time adaptive cybersecurity, and hyper-connectivity enabling selfoptimization between cloud, enterprise, and edge environments.



Both edge and fog computing involve moving processing closer to where data is collected. However, in fog computing, data is transmitted from the point of collection to a gateway for processing, then sent back to the edge for action. Fog computing uses edge devices and gateways with a local area network for the processing. Fog computing acts as a bridge that connects the cloud and the edge by combining the ability to run applications at the edge while leveraging the capacity of the cloud.

Edge computing enables the Invisible enterprise to collect and analyze data where it is collected, catching and correcting problems that might not be identified as quickly if the data were to be sent to a central server or cloud for processing and analysis. Keeping data onsite also reduces the security risk associated with transporting data, which can be important in financial organizations, for example. It also reduces bandwidth costs by processing some data onsite, rather than sending all data to a cloud or central server.



Edge computing creates new and improved ways for industrial and enterprise-level businesses to maximize operational efficiency, improve performance and safety, automate all core business processes, and ensure always-on availability. Key capabilities and benefits of edge computing in the Intelligent enterprise include:

- Powers the next revolution, transforming enterprises, manufacturing, and services.
- Optimizes data capture and analysis at the edge to create actionable business intelligence.
- Creates a flexible, scalable, secure, and more automated technology, systems, and core business process environment.
- Promotes an agile business ecosystem that is more efficient, performs faster, saves costs, and is easier to manage and maintain.

CHAPTER 4 Five Stages on the Edge Computing Journey 29



ACCELERATE YOUR DIGITAL TRANSFORMATION JOURNEY



Learn the 10 Best Practices for Edge Computing at stratus.com/edgebestpractices

terials are © 2020 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited

Companies today are facing mounting pressure to be more productive and efficient. At the same time, new technologies are becoming more powerful and more connected. The influx of data from new devices, and the need to analyze this data where little to no skilled IT resources exist, is driving the need for—and growth of—new edge computing infrastructure. In this book, you'll learn how edge computing helps you deliver the right information to the right place in real time.

Inside...

- Deliver performance and reliability
- Discover what edge computing is all about
- Recognize real-world challenges at the edge
- Build your business case for edge computing
- Explore industry use cases for edge computing
- Map your edge computing journey



Lawrence Miller served as a Chief Petty Officer in the U.S. Navy and has worked in information technology in various industries for more than 25 years. He is the co-author of CISSP For Dummies and has written more than 150 For Dummies books on numerous technology and security topics.

Go to Dummies.com[™] for videos, step-by-step photos, how-to articles, or to shop!



ISBN: 978-1-119-70803-2 Not For Resale

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.