


**THE
GORILLA
GUIDE TO...**®
EXPRESS EDITION



Enabling IT at the Edge

Alan R. Earls

Inside the Guide

- The Compelling Case for HCI in Edge Computing
- 5 Requirements for Edge Computing Success
- Where HCI and Edge Fit in a Cloud Computing World

THE GORILLA GUIDE TO...

Enabling IT at the Edge

Express Edition

By Alan R. Earls

Copyright © 2021 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review. Printed in the United States of America.

ACTUALTECH MEDIA

6650 Rivers Ave Ste 105 #22489
North Charleston, SC 29406-4829
www.actualtechmedia.com

PUBLISHER'S ACKNOWLEDGEMENTS

EDITORIAL DIRECTOR

Keith Ward

DIRECTOR OF CONTENT DELIVERY

Wendy Hernandez

CREATIVE DIRECTOR

Olivia Thomson

SENIOR DIRECTOR OF CONTENT

Katie Mohr

PARTNER AND VP OF CONTENT

James Green

WITH SPECIAL CONTRIBUTIONS FROM SCALE COMPUTING

David Demlow, Marlena Fernandez, Dusty Koekenberg,
Craig Theriac

TABLE OF CONTENTS

Introduction: Discovering Your Edge	7
Chapter 1: Computing at the Edge: Keep It Simple, Do More.....	9
Edge in Action: Edge Deployments in the Real World.....	11
Edge Delivers.....	17
Chapter 2: The Top 5 Requirements for a Successful Edge Deployment	20
A Modest Physical Footprint, with Simplified Cooling and Power.....	21
Affordable but Effective.....	22
Resilient and Survivable.....	24
Simplified Resource Additions (Scale Out) and Hardware Replacement.....	25
Repeatable, with Zero-Touch Provisioning.....	25
A Sound Edge Approach	28
Chapter 3: Edge Computing, the Cloud, and You.....	29
Likely Topologies.....	30
Edge “Fleet Management”	32
Rugged and Reliable.....	37
Jumping to the Edge.....	38

CALLOUTS USED IN THIS BOOK



The Gorilla is the professorial sort that enjoys helping people learn. In the School House callout, you'll gain insight into topics that may be outside the main subject but are still important.



This is a special place where you can learn a bit more about ancillary topics presented in the book.



When we have a great thought, we express them through a series of grunts in the Bright Idea section.



Takes you into the deep, dark depths of a particular topic.



Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



KNOWLEDGE CHECK

Tests your knowledge of what you've read.



PAY ATTENTION

We want to make sure you see this!



GPS

We'll help you navigate your knowledge to the right place.



WATCH OUT!

Make sure you read this so you don't make a critical error!



TIP

A helpful piece of advice based on what you've read.

INTRODUCTION

Discovering Your Edge

Welcome to The Gorilla Guide To...[®] (Express Edition)
Enabling IT at the Edge!

Edge computing is a new frontier for IT organizations. Demand for near-real-time response in retail applications, manufacturing, and beyond is driving a need to put computing power, especially efficient hyperconverged infrastructure (HCI) computing, where it can be best used. That means getting it as close as possible to the devices, users, and data. In other words, at the edge of the network.

But edge computing can be highly disruptive. It allows the organization to deploy applications in new ways that will drive growth and improve services. It does mean that IT teams will need to look at deployment differently. However, if you plan it properly and roll it out carefully and systematically, edge deployments can be just as secure, efficient, and robust as your local data center.

This book can serve as your starting point on the journey to that “edgy” future. In its pages you’ll get a solid

overview of the current state of edge computing, see how other companies have successfully set up and implemented their edge environments, and get plenty of tricks and tips to help you do edge computing right.

This guide is aimed at anyone with responsibility for choosing, implementing, and managing edge computing. That means IT architects, software integrators, CIOs and CTOs, and virtualization, network, and storage administrators will find value here. Ideally, each of those individuals will read this book, and work together as a cohesive team to bring a solid edge computing platform to fruition.

So let's get started! We begin with a quick overview of edge computing and how it's being used in the real world.

CHAPTER 1

Computing at the Edge: Keep It Simple, Do More

While edge computing isn't a new concept, it has become increasingly important for almost every kind of organization. The central idea is to locate computing capabilities away from the data center and closer to where data is created and used. Examples of this kind of distributed computing include the heavy processing associated with a retail site, factory floor operations, Internet of Things (IoT) devices, and content delivery or even transportation or agricultural systems.

A common thread is a growing preference for doing some or all of the data processing in place. This helps eliminate latency and reliability concerns of remote network delivery and empowers new kinds of applications that require processing data locally. It also unburdens networks by transferring *some* information from edge locations rather than all information, and it relieves centralized compute and storage facilities from handling all that data.

However, edge computing doesn't mean you can't centrally monitor and manage what's going on out there—it simply makes the monitoring process much more efficient.

Edge computing makes the case even more compelling by incorporating virtualization and, in some cases, convergence (for example, hyperconverged infrastructure, or HCI). This means companies can do even more with their edge investment by running multiple applications and workloads on the same hardware.

With the growth in IoT devices, sensors, cameras and more at the edge of distributed networks, there has been a sharp uptick in the generation of data, and pushing all of that data to and through a central data center or cloud no longer makes sense: It taxes resources and swallows bandwidth, and it can slow response time both to the edge and to other key systems, such as those for transaction processing.

In placing compute services at the edge, organizations can implement strategies for content caching, IoT management, improving response time, and faster data transfer rates. This also represents capabilities that can't really be matched by cloud technologies, while delivering those capabilities at a very competitive price point. Furthermore, there are no concerns about vendor lock-in or trouble moving or reclaiming future data.

Edge in Action: Edge Deployments in the Real World

Edge deployments are growing in scale, ambition, and impact across a variety of industries, and they showcase some of the edge's most convincing advantages, as shown in **Figure 1**.

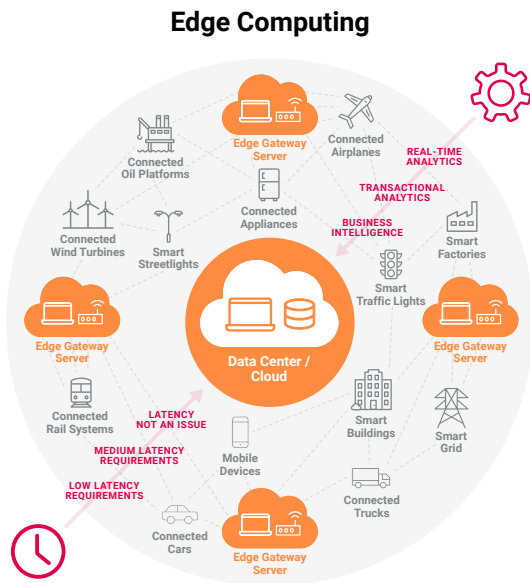


Figure 1: Edge computing opens the door to implementing powerful applications to use local data locally

Using Data Where It's Needed and Reducing 'Trips' to the Data Center

Moving data away from the edge and back to a centralized data center is generally economically unsound. And as data volumes continue to grow, it's becoming technically infeasible—there's simply too much data. Yet, competitive pressures mean accessing and using that data is ever more critical. The answer is data reduction and analysis at the edge, combined with edge use where possible and/or sharing with central IT for retention, reporting, and further analysis.

It's really a matter of having the right data in the right place at the right time. Unlike simply putting data or processing in the cloud from a local data source, which often leads to trouble in getting the data back for rapid use, edge opens the door to implementing powerful applications to use local data locally.

A good example of this strategy in action is Telford Offshore, which provides safe access and support for offshore work sites for the oil and gas industry. The company took a future-proof approach to edge that could be expanded to its entire fleet of vessels at sea.

Because of the extreme expense of Internet connections away from land (via satellite), it was vital that an edge system be able to function independently, without the

need for maintenance or continual connectivity to a data center.

Selecting a hyperconverged appliance to fulfill the required edge role delivered Telford Offshore, in effect, a complete IT infrastructure in one easy-to-manage solution, all while being cost-effective and efficient.

Similarly, Jerry's Foods, a retail grocery, liquor, and hardware store operator with some 50 U.S. locations, sought a modern approach to IT that would minimize the need for staff interventions, scale to meet future demand, and take advantage of modern technology at each retail location.

Jerry's Foods had a traditional virtualization system that couldn't meet those growing needs and required regular support to remain reliable. When it wasn't reliable, tasks like EBT food stamp transactions, critical for many customers and vital for the bottom line, sometimes couldn't be processed at all.

With a very small IT staff, rarely available at individual stores, Jerry's Foods needed an alternative that could be clustered for high availability, managed remotely, and deliver sufficient computing power for applications at each site. It also wanted to better support new and emerging retail and point-of-sale technologies that could eventually customize store visits for each

individual and enhance the marketing of specific products. Edge turned out to be the answer

Other retailers are using IoT and sensor tech to provide rapid response when staff needs to be redeployed—for example, to open up more checkout lines when multiple customers complete their shopping at the same time.

But not all edge options are the same. Some are focused more on IoT, others on POS and retail check-out. Ultimately, Jerry's Foods selected converged edge computing technology that allowed them to support the tried-and-true legacy applications it needed for transaction processing, inventory, and payroll tasks—accomplished locally but operating in concert with corporate headquarters—as well as new, leading-edge retail applications like data-intensive video surveillance.

For Jerry's Foods and many other organizations, edge computing has also shown its ability to help meet data location requirements (often driven by compliance mandates) while avoiding the many complexities of the cloud. Edge computing can also provide for autonomous operation that's not dependent on remote connectivity.

Getting More Value from Data

Placing the computing power right where the data is being generated and processing it locally allows

businesses to create more value from a single unit of data. With edge computing, processing of data will invariably be timelier and more actionable.



There are several old and many new protocols and transmission modes that are likely to extend the reach of each edge computing node as IoT matures. Using “near-field” communications, familiar technology like Bluetooth® and Wi-Fi, various 5G technologies, and more to connect and “wire” just about everything, it should be possible to ensure that investments in edge computing will be relevant through future technology evolutions.

Moreover, edge computing gave Jerry’s Foods an opportunity to make immediate use of data from refrigerated display units and storage areas. In the past, this information would have gone to the data center for storage and analysis. Now, edge computing makes quicker trend spotting possible, opening up opportunities to better coordinate activities to save energy (for example, by reducing staff trips in and out of food storage areas to reduce loss of cold air).

It also permits quicker identification of an improperly closed door on a storage unit and may allow coordination with point-of-sale data to ensure quicker restocking of displays than would have previously been possible.

Manufacturers have also found edge computing to be a great leap forward. For example, a recent article in [Machine Design](#) highlighted the improvements for in-process inspection made possible by edge computing harnessed to machine vision. The pairing of recognition technology with powerful local processing means that problems can be detected and acted on in near-real time. Rather than just generating trend data for later analysis and corrective action, sensors can empower immediate remediation.

For its part, Telford Offshore found an immediate improvement in database performance from its move to edge computing. That choice allowed further nodes to be added on an as-needed basis, eliminating the usual tendency to over-provision storage (which, of course, increases costs). And edge still supports key capabilities such as replication, snapshot, and cloning capabilities for both continuity and redundancy.

Enhancing the Customer and Employee Experience

Edge can enable many capabilities. On the near horizon are numerous applications, some now being implemented, under the rubric of Smart Cities. Examples range from plotting the whereabouts of municipal property to prevent misuse, to monitoring social distancing, managing traffic more efficiently, and providing better services to individuals with disabilities.

In warehouse operations, edge can speed up receiving activities by making sure the right people and equipment are at the correct loading dock the moment a shipment arrives.

Edge also is growing in importance in medical situations, ensuring that patient data stays secure while improving the speed and quality of applications available to practitioners trying to deliver the best treatment options to patients.

Edge Delivers

What can edge do for your organization? Unlike some IT buzzwords or trends, edge computing is a natural development of robust, long-term trends, not a flash in the pan. It can confer advantages on large enterprises or small.

Edge computing offers a doorway to groundbreaking and exciting opportunities that harness artificial intelligence and machine learning, yet also provides a compelling delivery option for everyday office applications, database services, and enterprise applications. At the same time, it's a bridge to the future and a foundation for effective IT practices today. Converged edge, in particular, delivers the best of the old and the promise of the new.

Ubiquitous Computing

Not to change the subject, but one way of thinking about edge versus other types of computing is that it's part of the trend toward *ubiquitous* computing (computing involved in just about everything). We are getting there. Our devices, IoT sensors, smart machines, smart appliances, smart cars and autonomous vehicles are all part of this growing trend toward computerizing everything. Edge is important for tying all those pieces together and then uniting them with central data systems and/or the cloud.



That makes edge a real win-win for almost any use case. But there are still things to consider regarding your specific needs and what it takes to succeed with edge computing.

CHAPTER 2

The Top 5 Requirements for a Successful Edge Deployment

Success with edge computing is based on knowledge. In particular, organizations on that path should be aware of characteristics of edge computing implementations, qualitative and quantitative, that need to be considered.

Edge computing deployments have unique constraints that are significantly different from the ones with which typical data center deployments contend. After all, by definition, edge deployments are away from normal support services, far from the sanitized data center, and deeply enmeshed in the real work of the organization where they must deliver high value without disrupting other business activities.

So, organizations pursuing edge deployments should step back and carefully consider all the different factors that must be weighed for edge computing and then look for a solution that addresses the most crucial. These include: a physical footprint with minimal needs for cooling and power, a footprint that's affordable, easy to expand and maintain, simplified resource node additions

(scale out) and hardware replacement, a combination of hardware and software that's failure tolerant, and a configuration that's deployable in multiple locations, consistently and simply.

A Modest Physical Footprint, with Simplified Cooling and Power

It might seem an obvious point, but many products sold as “edge” equipment are not actually designed for that purpose. They're simply marketed for that role or, at best, minimally adapted for edge requirements. Because some vendors just sell standard data center equipment for edge use without accounting for the less-than-perfect environment that may be encountered there, edge installations are acquiring a reputation for being unexpectedly troublesome. For example, data center gear designed to work when provided with the highest quality cooling can suddenly develop reliability issues when located in a poorly ventilated storage space at an edge installation.

That's why it's vital that edge equipment should be conceived from the start for that purpose, with sufficient ruggedness built in to handle the wider range of issues that are typical at the edge, whether that edge is the shop floor, a warehouse, or somewhere in a retail establishment.

That pertains also to the need to have equipment that's largely self-contained and can be installed simply, ideally by existing staff. And, of course, maintenance needs to be minimal.

All of this can also be described as flexibility. In other words, edge components and systems need to be thought of as “universal” products that can be deployed when and where needed, with few limitations, and made appropriately secure in any given environment.

Affordable but Effective

As noted, the physical footprint matters. For the many industries that operate remote sites, ranging from finance and retail to manufacturing and so-called remote office branch office (ROBO), there's a need for reliable computing to support security, point-of-sale (POS), inventory management, and more. But none of these scenarios can afford large, dedicated spaces or complexity.

Edge adopters must consider the size of the actual equipment and its requirements for access space, air flow, cabling, and so forth. So, smaller and more compact equipment generally helps drive flexibility because it allows more freedom to choose deployment locations and makes it less likely a deployment will disrupt other activities. It typically also implies less onerous cooling and power needs.

Environments are often already committed to productive, revenue-generating activities, so an edge solution shouldn't impact that. Edge should minimize the need to take up or dedicate space or create new spaces. It is worth noting that compact form factors can also be helpful for enhancing physical security. For example, a smaller form factor means equipment can be secured and ceiling mounted for example, where it becomes harder to tamper with. The same attributes allow equipment to be placed in an already-secure location without impacting other activities there. All of this is very important when equipment is away from a dedicated, secured data room or data closet.

Lifecycle planning also represents a best practice. Hardware may eventually need to be replaced or upgraded. With a planned and simplified replacement process and with regular monitoring and maintenance, it should be possible to anticipate most replacements and all upgrades. This is an important point, because unanticipated or unplanned hardware replacement can require expensive "emergency" truck rolls, as downtime is rarely acceptable in modern organizations.

It's far better to have a plan for equipment types that follow a clear upgrade path, use standard configurations and connections, and, therefore, can be replaced or upgraded with minimal effort and little or no downtime.

Resilient and Survivable

Edge computing is no place for daintiness: It's where real work gets done, some of it dirty, messy, hot, and noisy. So, in addition to needing only a basic physical environment and simple power and connectivity, prudent adopters make sure the edge setup (hardware and software) is designed to be failure resistant, able to recover from many problems autonomously, to protect data, and to maintain operations in almost any circumstance.

What does that look like in practice? As much as possible, physical interventions onsite need to be avoided—no buttons that say “press here” to restart something. Instead, everything should be callable and able to respond when power is made available. Furthermore, autonomy should be part of the basics—delivering no-nonsense reboots and allowing most other maintenance tasks to be initiated remotely. On a physical level, equipment shouldn't be overly sensitive to voltage dips or spikes. Some data center equipment is intended for use with sanitized power systems, designed to protect equipment from any variations in power quality. In contrast, the real world where edge systems live may include older, less adequate circuitry, multiple competing uses on the same system (such as welding and production equipment, large HVAC systems, and so on).

Edge equipment needs to be ready to handle that kind of stress without generating performance issues.

Simplified Resource Additions (Scale Out) and Hardware Replacement

Edge environments are very dynamic with new applications being deployed regularly and data volumes growing exponentially creating new demands on edge infrastructure. It's critical that infrastructure is designed to accommodate that growth and expand and upgrade the edge micro-datacenter with new resources and applications as easily as the initial edge deployment. Failure to plan for expansion of the edge environment can lead to expensive forklift upgrades or multiple independent islands of infrastructure to manage, with all the complexity and cost associated with that kind of choice.

Repeatable, with Zero-Touch Provisioning

For all but the smallest of organizations, this is perhaps the most important consideration because edge may involve multiplying sites and types of equipment on the network. If approached haphazardly, without a plan, edge can quickly spawn hard-to-manage complexity that can strain IT staff and have company-wide

implications. To keep from becoming a nightmare, edge systems should take a standardized approach requiring little or no customization and minimal skills in installation.

When possible, edge should offer or embrace infrastructure as code (IaC), which simplifies change control.

Explaining Zero-Touch Provisioning

When people not familiar with zero-touch provisioning (ZTP) first hear the term, it can sound a little strange—but for most people it's actually a familiar experience. Mobile devices are now normally updated and adjusted through zero touch. The same is true for operating system updates on personal computers. This matters tremendously. Imagine what life would be like if every phone update required a trip to a retailer, visit from a tech, or interfacing with a human tech online. These days, ZTP is even showing up as a term of art in business and medicine because it speaks to efficiency. ZTP matters.



Repeatability means that service and support is standardized so staff doesn't need to research each installation before responding to a problem but, instead, can count on using a consistent approach and methodology. It's a model for efficiency used across every other domain, from manufacturing to medicine, but too often ignored in edge deployments.

Similarly, management must not require specialized IT staff on site; upgrades and infrastructure scaling must be non-disruptive; the foundation must be self-healing; and IT specialists must be able to manage the entire edge fleet seamlessly at scale. This is a logical corollary of deploying similar systems at every node. Those systems will have identical software and applications, identical mechanisms for deployment, and every opportunity to enhance and improve standardization through repetition. With this approach, even inexperienced staff can quickly become experts.

Finally, look for zero-touch provisioning. This is a device-configuration process that can be operated automatically and eliminates most of the burden on IT administrators when setting up, maintaining, or upgrading an edge system (see **Figure 2**).



Figure 2: Zero-touch provisioning allows devices to be configured automatically

A Sound Edge Approach

The point of edge deployments should be to help the whole business succeed. That means avoiding excessively complex edge technologies and requirements and instead making sure edge delivers on its promise of cost-effectiveness and empowerment beyond the data center. By keeping those points in mind and watching out for the five key concepts noted here, organizations will have the basis of a sound edge approach, one that's highly affordable, highly functional, and highly reliable.

Test-drive innovative, right-sized, and reliable application infrastructure for 30 days at no cost. Take advantage of the industry's only Edge Computing trial program, including hardware (you can sign up [here](#)). But don't forget to think about the broad approach to deployment that you may wish to employ. That's where thinking about edge computing topologies comes in.

CHAPTER 3

Edge Computing, the Cloud, and You

So, you've come this far on your edge computing journey. It's time to familiarize yourself with just a few more concepts. Then, you'll be ready to rule the edge! This time it's the overall structure or topology of edge computing that you want to deploy.

The infrastructure topology of the future will be strongly focused on hybrid environments. Most organizations are likely to employ a combination of data center and cloud-based resources. Meanwhile, edge computing applications will play an important role in handling local processing needs. This “new” topology is in contrast to both the traditional all-on-premises approach, as well as the “born in the cloud, all cloud” perspective. It recognizes that latency and bandwidth issues matter, as do regulatory and autonomy issues, which are other common drivers for on-premises infrastructure, especially as data generated locally becomes ever more important.

Edge computing is increasingly important for many industries that embrace full-scale digitalization or are simply modernizing and growing to meet real business needs. Edge computing has taken on an expanding role in these developments over the past decade and is poised to extend its role further in the years ahead. Recent events have only heightened the importance of this growing role.

The increased focus on edge computing introduces complexities and presents opportunities. Among the prerequisites for adopting a hybrid cloud + edge infrastructure is not only understanding and selecting a topology, but also bearing in mind the “fleet management” strategies needed for the edge computing systems, which demand thorough consideration of a range of factors.

Likely Topologies

While the never-ending evolution of information technologies makes it hard to predict exactly what topologies will emerge or seem most appealing in the future, there are currently several strong contenders for edge computing. Generally speaking, the most common topologies represented in edge computing and recognized by leading analysts differ in the functions placed




Figure 3: A simplified comparison of edge topologies

at the edge and the relationship of those functions to cloud and/or on-premises data center computing (see **Figure 3**).

- **Regional Data Center Edge—CDN, Telecom DC, Colocation.** This might be a service provider configuration with multiple tenants and, in comparison with other edge computing scenarios, is typically a very large-scale operation that differs from a traditional data center only in its relationship to an even larger central data center and in usually having a narrower focus.
- **Local Data Center Edge—Small Data Center, Micro Data Center.** This type of edge computing is likely general-service-oriented, perhaps for a remote office or branch office and is characterized by low or no staffing.

- **Gateway Edge—Intelligent Local/Field Gateway.** This typically comprises a small cell or access point, such as a video management software (VMS) surveillance system and offers zero touch provisioning (ZTP) and configuration management.
- **Device Edge—Embedded Computing Devices and Traditional PLCs.** This type of edge, typically a single machine or work cell, has only enough intelligence to assist with a specific operation and provide some degree of reporting.
- **Compute Edge—Edge Server/Storage Outside of a Data Center.** Examples of this type of true edge computing potentially have the ability to include specialized services, such as video analytics-based applications and likely include ZTP.

It's important to note that these topologies aren't rigid. There are gray areas and some edge implementations can actually include multiple topologies. However, in most instances, one is clearly predominant.

Edge “Fleet Management”

For all edge topologies, but particularly for “true” edge, organizations need to be cognizant of lifecycle management; that is, of “fleet management.” They must plan for longevity to get the maximum value from

expenditures and to avoid dead-end investments. Their thinking should include concepts such as the following:

- **ZTP and configuration management.** This is relevant not just initially, but as hardware fails or as capacity needs change. It's a rationale for designing many edge systems around scale-out clustered architectures like hyperconverged infrastructure rather than a single monolithic "box."

ZTP is an important "Day 1" feature, as well as a long-term configuration attribute, because it greatly simplifies all ongoing activities related to edge computing. It means the edge system itself, and often the devices attached to it, can be configured automatically, without the need for hands-on intervention or with only minimal remote involvement by staff. This has obvious implications for the lifetime costs of edge computing and is achieved through approaches that provide automatic provisioning and configuration as well as ongoing system monitoring. Essentially, once an edge device with ZTP is powered and connected, it reaches out to acquire the information it needs to become fully functional.

- **Centralized management for "Day 2."** This includes automation of ongoing management where and when possible. The Day 2 concept is a way of talking about

all the operational realities that become important after initial set up. Day 2 can sometimes be a surprise, as “simple” implementations become more complex when day-to-day management challenges turn out to be larger than expected. So, an edge computing approach should supplement ZTP, as noted, and also offer ongoing visibility that’s simple yet comprehensive, building on the advantages of edge computing.

- **Cloud integrations.** Edge computing offers simplicity and cost advantages, and those advantages should extend to integration. One element that should always be considered is how cloud fits in. Cloud can, in some cases, replace a traditional, centralized on-premises data center or can supplement that data center. It can also provide bridging functions to other resources (for example, storing data that’s later accessed by a data center). And cloud can offer advantages including relatively simple integration through the use of APIs. An edge computing solution should have a readiness to meld with cloud computing, whether Infrastructure as a Service (IaaS) or Software as a Service (SaaS), including cloud storage integration.

It will often be found that cloud capabilities are more than sufficient to address edge computing storage needs, allow greater geographic resilience, and

centralize access for analysis. Supplemental capabilities (typically SaaS) can also enrich a given edge computing installation.

To ensure the safety and protection of edge computing data and the ability to get up and running after disruptions, whether natural or human caused, organizations should consider cloud storage options in particular. The cloud storage integration aspect, namely getting edge computing data to the cloud, can help provide off-site resiliency, centralized data sharing, and so forth.

- **Monitoring and support.** It's crucial that edge computing be highly visible, easy to support, and free from surprises. It should “just work,” and problems should be anticipatable as much as possible. An edge computing installation should start with designed-in reliability. Edge computing systems should be built for field use, not just for operations in “sanitized” data centers or air-conditioned offices. But when problems do occur, they should be easily and quickly identifiable with a clear path to problem rectification.

Some additional issues to consider as you select a topology and plan your edge computing implementation include a “start at the end” suggestion: Define what success is and how it will be measured. This could include the view of the customer and other stakeholders.

Consider, too, how your edge computing choice can support the evolution of your other IT goals and initiatives. Then, look at important nitty-gritty issues. For example, in most cases edge computing will help with latency and response times, but to ensure optimal results, you'll want to identify and examine the factors that will contribute to improved performance.

Measuring Latency

Latency is never good and what users notice, but in the real world, almost everything has some degree of latency. The key is understanding how much is too much and how much different users and applications can tolerate. And, of course, latency is usually in part a product of how many activities are competing for the same resources. If you want to get deeper into the topic there are tools and techniques to help and new terms that you'll encounter like Round Trip Time (RTT) and Time to First Byte (TTFB). But don't worry. Picking the right topology should get you close to where you want to be!



Finally, it's prudent to look at the conceptual components of an edge computing implementation, including storage, computing capacity, and electrical power, as well as analytics capabilities, network and communication needs, and both physical and logical security. Vendor solutions will address most of these issues, but having clarity about your needs will help ensure that those needs are met.

Rugged and Reliable

Consider two industry sectors that are frequently adopting edge computing: retail and manufacturing. And further consider how these deployment styles can work with the specific needs of those sectors. For example, manufacturers often have use cases that focus on asset tracking, remote operations, and logistics. They may also extend to warehousing, operational automation, security, maintenance, and diagnostics.

Similarly, retail edge computing use cases often include supply chain control and optimization, digital signage, in-store experience, and (recently) proximity marketing.

Comparable experiences will likely be found in all industries. Keep in mind that choices regarding both topologies and vendors can be important in avoiding

proverbial “truck rolls” (emergency service calls). Clearly, edge computing loses some of its appeal if it becomes a source of vulnerability and unreliability. Rugged and reliable is definitely a choice and selecting the right vendor can help ensure that your organization won’t have to incur unexpected or unscheduled services.

Understanding edge computing topologies is the first step toward building an effective edge computing deployment to support these and other use cases. It takes a clear grasp of business needs and edge computing capabilities to complete the job successfully. That’s where Scale Computing comes in. Its HC3 edge solution is a self-healing, right-sized infrastructure that runs critical apps, right where they’re needed.

Jumping to the Edge

We hope you’ve enjoyed the Gorilla Guide, which explained why edge computing is trending to be one of the biggest IT growth areas. Of course, as noted, there are wide variations in what people mean when they discuss edge and the different topologies under which it is implemented.

But the common thread is the need to collect and handle data faster and do more with it at a point that’s closer to

its source, the activity that needs the data, or both. And the reward is a reduction in the “time to value” for any given data.

Now that you have a more complete view of the concepts, terminology, and tradeoffs involved in edge computing, it’s time to step back and consider your situation and the needs of your organization. Where is it stressed? Where should it grow? What’s the best way to make the most of existing investments in IT and beyond? Take your time. And get a feel for how edge computing can figure into your infrastructure today and your vision for the future.

Then, go looking for information. Where you’re unsure or need more details—track them down. Effort put into this can get you to the point where you understand the choices available to you and which ones make the most sense. Don’t be afraid. Make the jump and discover how the right edge choices can provide real competitive advantage.

When that time comes, consider Scale Computing solutions. Its [HC3 Edge](#) platform brings the efficiency and scalability of HCI to edge computing. This is especially important at the network edge, where things like small space requirements and lack of staffing become critical concerns that must be addressed.

If you'd like to try out Scale Computing for free, you can [fill out a form](#) and get the process started. What have you got to lose? Get started today!

ABOUT SCALE COMPUTING



Scale Computing is a leader in edge computing, virtualization, and hyperconverged solutions. Using patented HyperCore™ technology, the company's HC3 self-healing platform automatically identifies, mitigates, and corrects infrastructure problems in real-time, enabling applications to achieve maximum uptime, even when local IT resources and staff are scarce. Edge Computing is the fastest growing area of IT infrastructure, and industry analysts have named Scale Computing an outperformer and leader in the space, including being named the #1 edge computing vendor by CRN. Scale Computing's products are sold by thousands of value-added resellers, integrators, and service providers worldwide. When ease-of-use, high availability, and TCO matter, Scale Computing HC3 is the ideal infrastructure platform. Read what our customers have to say on Gartner Peer Insights, Spiceworks, TechValidate and TrustRadius.

www.scalecomputing.com

ABOUT ACTUALTECH MEDIA



ActualTech Media is a B2B tech marketing company that connects enterprise IT vendors with IT buyers through innovative lead generation programs and compelling custom content services.

ActualTech Media's team speaks to the enterprise IT audience because we've been the enterprise IT audience.

Our leadership team is stacked with former CIOs, IT managers, architects, subject matter experts and marketing professionals that help our clients spend less time explaining what their technology does and more time creating strategies that drive results.

For more information, visit

www.actualtechmedia.com