

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



Taking Data to the Edge

Tim Parker and Jason Carolan

INSIDE THE GUIDE:

- Customer Stories, Including:
 - Smart Cities
 - Sports and Video
 - Healthcare

**TAKE A QUICK WALK
THROUGH THE IT JUNGLE!**

Compliments of



THE GORILLA GUIDE TO...

Taking Data to the Edge

Express Edition

AUTHORS

Tim Parker, VP, Network Strategy

Jason Carolan, Chief Innovation Officer

Copyright © 2019 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

Okatie Village Ste 103-157

Bluffton, SC 29909

www.actualtechmedia.com

TABLE OF CONTENTS

| | |
|---|-----------|
| Introduction: The Edge of the Revolution..... | 5 |
| Chapter 1: An Overview of Edge Computing..... | 7 |
| Understanding Edge Computing..... | 9 |
| What's Driving Edge Computing?..... | 11 |
| Bridging to the Cloud..... | 12 |
| Challenges to Edge Computing..... | 13 |
| Edge Computing Adds Value..... | 14 |
| Chapter 2: Smart Cities are Smart Infrastructures... | 17 |
| Smart Cities Depend on Sensors and Data..... | 18 |
| The Journey of 1,000 Miles Begins with the First Step..... | 20 |
| How Hybrid Edge Services Help..... | 21 |
| Chapter 3: MLB Takes Baseball to the Edge..... | 24 |
| MLB's Many Data Platforms..... | 24 |
| Room to Grow with the Game..... | 27 |
| Chapter 4: Taking Health Monitoring into the Home..... | 29 |
| Beyond Electronic Health Records (EHR)..... | 31 |
| When IoT Sensors Monitor People and Health... | 32 |
| Diabetes Applications..... | 34 |

| | |
|--|-----------|
| Frontiers in Managing Chronic Conditions..... | 34 |
| Expediting and Automating Healthcare at the Edge..... | 35 |
| Technology Underpinnings Make Things Happen..... | 37 |
| Chapter 5: Put Hybrid Edge Computing to Work..... | 38 |
| Now Is the Time to Develop Edge and Distributed Strategies..... | 38 |
| The Tether That Binds for a Successful Future.. | 40 |
| About Flexential..... | 44 |


INTRODUCTION

The Edge of the Revolution

Welcome to this Gorilla Guide Express to Taking Data to the Edge. It covers one of the most important emerging areas of IT infrastructure, but one that still remains a mystery to many. This short book aims to take the mystery out of edge computing; by the time you're finished reading, you'll have a firm grasp of the ins and outs of the technology, as well as see ways that it's being used to solve real-world problems.

The real-world problem here is that people are using more devices than ever to stay connected. Those devices include phones, watches, tablets, laptops, desktops, fitness trackers, streaming sticks, wearables, and game consoles.

And it's only going to continue: The number of active Internet-connected devices is expected to grow to 20 billion by 2020, which means the soaring bandwidth and network demands will also explode—after all, we expect instant service, low latency, and zero bandwidth issues from these devices. That's why companies as well as their customers are demanding the data be at “the edge” for the fastest service and productivity. See **Figure 1**.



**By 2020, more than
65% of enterprises
(up from 30% today)
will adopt IoT products.**

Source: Gartner

Figure 1: Internet of Things (IoT) device adoption is soaring, and shows no signs of slowing.

In this book, we'll explore how several large companies in different industries are applying edge technology and computing. We'll consider:

- **Smart cities:** Some forward-thinking U.S. cities are implementing “smart” infrastructures, with strategically-placed devices, to provide better service for its residents and visitors.
- **Sports and entertainment:** We'll show you how MLB and Flexential take fans to the next level of interactivity, creating a better experience for fans.
- **Smart health:** You'll see how edge technology is being leveraged to report real-time monitoring of patient health, resulting in more effective treatments that even save lives.

CHAPTER 1

An Overview of Edge Computing

The infrastructure of the Internet is sometimes visualized in terms of “speed zones,” also called “latency arcs.” Think about it like a highway from a densely-populated area to a rural area. At the outermost periphery of the Internet, on the country roads, are end-users or consumers. This zone of the Internet is usually served by Internet Service Providers (ISPs) such as cable and telecommunications carriers. It’s referred to as “the last mile” because it represents the final—and slowest—link between clients (and the apps or applications they run) and servers (and the data and services that they provide). Internet speeds on the last mile range from under 10 Mbps for wireless and older Ethernet connections, to as high as 1 Gbps for 5g wireless and Gigabit Ethernet (available mostly in major metro areas).

The next zone of the Internet is called “the edge.” It straddles the boundary between the last mile and the rest of the global Internet. Here, local distribution networks connect via head end servers from cable providers, and where fiber or twisted-pair concentrators

connect network segments onto higher-speed links. In wireless terms, it's where cell towers connect into back-haul networks in carrier networks. This is also where Metropolitan Area Networks (MAN) or Carrier Ethernet fits into the overall Internet.



Interior zones on the Internet represent the highest-speed, highest-volume data highways available anywhere.

The middle zone that connects the edge to the core and hyperscalers is called the “near edge.” This zone consists of data centers with quality space and power, high-density carrier and fiber connectivity, and an aggregation point for edge deployments that requires more compute and connectivity than is typically found at the edge.

All these things represent the term “the edge,” and reside at small- to mid-scale aggregation points or data centers where servers may also be found. Thus, points on the edge get as close to users who connect to the Internet as computing can come. From a user's perspective, round trip times to and from their primary edge connection is as fast as the Internet ever gets. Any connections deeper into the Internet add to overall

access and round-trip times involved, no matter how fast those links may be.

Interior zones on the Internet represent the highest-speed, highest-volume data pathways available anywhere. Picture the Autobahn in Germany—a fast speed highway. This part of the Internet contains “the core” at peering points where high-speed optical switch fabrics handle terabit or even petabits of data. It also includes high-speed, high-volume optical links (usually IP/MPLS, MPLS-TP, or Carrier Ethernet) between peering points, or for long-haul terrestrial or undersea optical cable runs. This is the bulk data transport to the max.

Understanding Edge Computing

The idea that drives edge computing is bringing latency-sensitive applications as close to the consumer as physically possible (see **Figure 2**). Ideally, that means running such applications where users bridge from the last mile to the Internet edge. Practically speaking, it means finding a location close enough to the users near the edge where latencies are minimal (at or under 10 msec). Such proximity to users ensures minimal response times and the best possible user experience.

Edge computing may also be understood as a method for replacing traditional on-premises computing. With

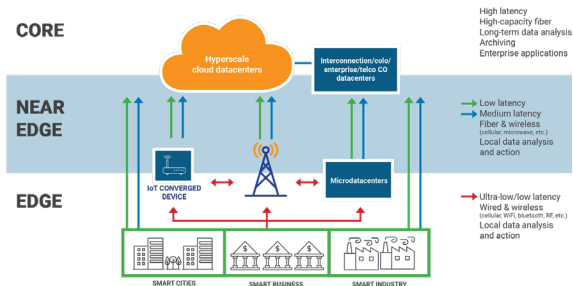


Figure 2: The “near edge,” which connects the edge to the core and to hyperscalers.

workloads moving into the cloud, companies and organizations must consider the costs and trade-offs involved in staying local versus making that move. Recognizing that public clouds can’t solve all IT problems—issues with regulatory compliance and latency will persist—implementations may take on hybrid characteristics. See **Figure 3**. It can also become very expensive to transport the billions of packets to the public clouds or core compute. This means a decision about what’s managed at the edge versus what is sent to the core is critical to a successful architecture. This usually means some local computing, plus a mix of private and public cloud use. That in turn is shaped by considerations like latency, applications used, privacy and control concerns, and other factors. This is called “hybrid IT.”

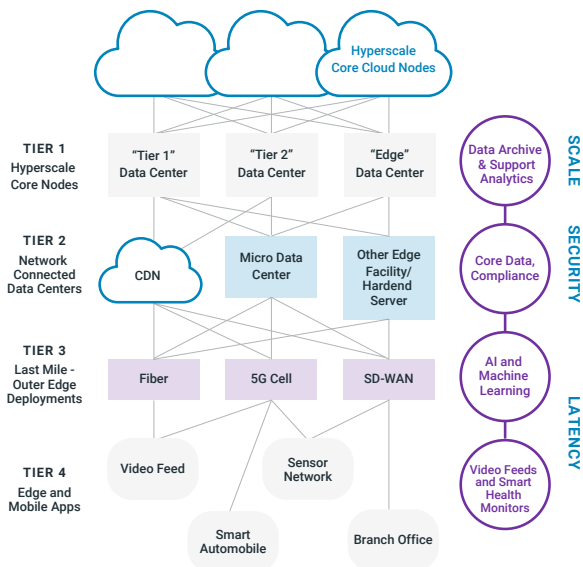


Figure 3: An overview of what the new “3 tier architecture” for applications looks like.

What’s Driving Edge Computing?

Though the impetus to use the cloud for conventional computing is undeniable and growing strongly, there’s more to IT than supporting line-of-business and other applications. Machine-to-Machine (M2M) communication and related workloads most likely occur outside of major corporate offices and campuses.



M2M workloads can also be highly latency sensitive, because they demand real-time responses to instrumentation and change orders and must integrate with logistics and supply chain management needs. M2M communication occurs anytime a sensor is read, a website handles a purchase (e.g., manages inventory, schedules shipping, tracks order status), or quality controls get used.

In some situations, life-and-death situations—for instance, an ambulance reworking traffic signals to reduce delays, or the ability for an EMT to update the hospital with the vital signs of a patient en route. Even that few seconds saved can make a huge difference.

Bridging to the Cloud

In modern organizations, access to big data and public clouds is mandatory. Public cloud providers deliver effective, easy-to-use versions of critical IT services. These run the gamut from facial recognition to big data analytics. But it's essential to choose cloud components that make things better, not worse. The costs of Internet

connectivity and bandwidth pose key considerations for (and occasional impediments to) cloud use. Hybrid edge service providers offer unique services, because they focus heavily on carrier interconnectivity.

Peering agreements and network congestion between carriers and backhaul providers can pose problems, too. A 100 msec round trip for a sensor interrogated through a public cloud might be better served through a 10 msec connection to a nearby hybrid edge data center, itself only 10 msec away from that same cloud provider through a private and highly available transport network.

Ultimately, it's about the latency, and understanding that the arrangements that minimize latency are always straightforward or single-hop transitions. Picture a traffic jam on the 405 in Los Angeles, California, taking longer than the smooth-sailing drive on a county road in Portland, Oregon.

Challenges to Edge Computing

Using the Internet efficiently and effectively can pose challenges: it's a complex place. Hybrid edge computing offers a rational approach to coping with such complexity. Organizations must learn how to make use of the

capabilities available to provide the best user experience at the lowest cost.

More importantly, organizations must also learn how to manage M2M communications so that things happen quickly enough to allow business to proceed and succeed. This requires a deep and profound understanding of the workloads in need of handling, and how various hybrid edge computing solutions can deliver the right access to the right data in the right timeframe.

Edge Computing Adds Value

Hybrid edge computing solutions offer a quick transition from the Internet edge to high-speed connection between their data centers. Hybrid edge service providers also offer SDN and SD-WAN capabilities, which can be combined to solve tricky real-world problems.

For example, hybrid edge service providers offer geographically dispersed, readily accessible data centers called near edge data centers. These provide excellent ingress for far-flung collections of sensors (oilfields, utilities, and water/flood gauges, for example) which gain ready and immediate access to data acquisition and reduction/filtering services, with the ability to route important or mission-critical information quickly wherever it needs to go. A comparison between FedEx

Data services are exploding, and traffic patterns are changing.




Figure 4: The changing nature of data delivery.

Ground and FedEx Next Flight service is apt here: the unmanaged public cloud maps to the former, and a hybrid edge service provider to the latter.

Easy Addition and Subtraction



Flexential offers access speeds from 100 Mbps to 100 Gbps to its customers, depending on their latency requirements (and many customers mix and match speeds to fit specific scenarios). Flexential's 100 Gbps network backbone minimizes the impact of long-haul communications on latency.

Flexential's services, based on SDN and SD-WAN, are easy to provision and deploy, and help to manage costs by adding capacity as demand increases, and lowering capacity as it wanes.

All in all, a hybrid edge service provider offers the kind of glue that can tie private IT together with public and private cloud assets and services, seamlessly and affordably (see **Figure 4**).

CHAPTER 2

Smart Cities are Smart Infrastructures

Smart cities are an emerging phenomenon, both in the United States and globally. Such locales are best understood as taking the elements of any city's infrastructure—buildings, roads, utilities (water, power, gas, waste handling, and so forth)—and adding sensors for reporting, communications to send data along, and backend computing and analytics to handle and make sense of that data.

Smart infrastructures (much like cities, spread over larger, more dispersed areas like counties or metro areas, or far-flung systems, like railroads, telecommunications systems, watershed management, or the power grid) work in similar fashion. Usually, they're more spread out geographically, and involve more parties in ownership, control, governance, and regulatory aspects.

Smart Cities Depend on Sensors and Data

In concept, a smart city seems simple. Start by identifying infrastructure elements that need to be monitored and managed (and often also, charged to customers, or charged back to organizational units). Equip them with sensors to provide useful information about how they're being used, errors or problems, service quality and delivery, and so forth. House that data properly, so that real-time information, events and alerts can be communicated to interested (or obligated) parties and gather that data to be sliced and diced for analytics to help improve service, reduce waste, maximize returns and more.

Sounds simple from a 10,000 foot view, doesn't it? But on (or under) the ground, it's a huge, complicated collection of disconnected systems, each with its own existing legacy of data, reporting, billing, and management tools (**Figure 5**). Then there's the work involved in putting (or adding) sensors in the field, establishing communications to acquire the data they gather, filtering and managing that data, and sharing that information. This is where things get interesting. Think of the parties that will be involved:

- **Service organizations** (trouble tickets, work orders, field staff assignments)



Figure 5: Smart cities require fast connections, lots of sensors and mountains of data that need to be processed fast.

- **Billing organizations** (metered or usage fees, fixed and special fees, plus applicable taxes that span the range across multiple jurisdictions: city, county, state, and federal)
- **Regulatory bodies** (rate or utility commissions)
- **Reporting bodies and requirements** (budget management, resource management, people management,

environmental impacts and assessments, planning commissions, state and federal grant applications and management, consumer/voter input, and more)

In some cities, city agencies or departments handle most or all of these functions. In other cities, public-private partnerships prevail. This adds the complexity of acquiring and sharing data across organizational boundaries to all the issues already inherent in acquiring, filtering, managing, archiving, communicating, and otherwise responding to the immense volumes of data that smart city initiative inevitably create.

The Journey of 1,000 Miles Begins with the First Step

The important step in taking the smart city plunge is to craft a flexible and capable M2M data architecture. This will involve many sources of incoming data, which must be properly filtered and handled. It also requires a set of robust and open-ended data acquisition, analysis and reporting tools, because data volumes (and uses for such data) only keep increasing. And finally, it means being able to provide controlled access to at least some of that data for all kinds of uses when and where it's needed.

Cities—and other infrastructure providers facing the challenges of “smartening up”—will want to survey their

current services and resources. They'll need to inventory what kinds of monitoring, reporting, data acquisition and handling, and so forth, are already in place for each one.

The architecture work comes from deciding how to bring all that information together, and then how to store, archive, analyze and expose that data to meet all the various communication requirements around that data. This is an area of endeavor where a hybrid edge service provider can be particularly and directly helpful in making sure that all the various requirements for access, reporting, and communication get handled.

To make smart cities and 5G successful, collaboration between private and public organizations is critical. Identifying partnerships between mobile carriers, WiFi providers and local municipalities is extremely important. Coupled with the need for edge data center support, it can be successful, but only through joint cooperation and support.

How Hybrid Edge Services Help

To begin with, smart cities (and others of that ilk) should recognize that an edge-centric architecture is ideal for their needs. It makes sense to handle the huge volumes of sensor and M2M data that smart city initiatives generate right at the edge of the Internet.

Smart Initiatives

Flexential's edge data centers offer cities the kinds of connections, compute and storage facilities that smart initiatives need. They offer colocation and peering presences to bring existing players (third-party service providers, such as power companies, utility operators, waste management organizations, and so forth) into the mix for incoming data streams.



Flexential also offers predictive scaling so that capacity can adapt to accommodate end-of-cycle billing and reporting needs, while reducing capacity during more idle periods.

This takes the well-known database principle of “selection before projection” and morphs it into “reduce the data before communicating about that data.” Given that edge data centers can offer raw data storage at low costs, cities can elect to archive “the whole thing” at the edge before reducing data volumes to only those items that actually need report inclusion or analysis. Then, they need communicate only a much-reduced set of data

items to other consumers of that data, thereby limiting bandwidth consumption and overall traffic volumes.

For smart cities (and other infrastructure players), near edge solutions offer security and compliance, to keep data secure, reliable, and available. Near edge hosting services can deliver a broad range of compute, storage and networking capabilities, so that data filtering and reduction can be applied as needed, along with storage for big data sets and archival/regulatory retention policy compliance.

Finally, Flexential's access to the far edge means it can convey data and communications wherever and however it might be needed: a core deployment in Portland, Oregon, edge deployments in Florida or Virginia, or at partner/service provider premises around the globe.

CHAPTER 3

MLB Takes Baseball to the Edge

Few, if any, games are as data-driven as baseball. Some might argue that statistics drive the game, rather than the other way around. Beyond the huge volumes of data that the game itself generates, fans not only consume those statistics, but they also combine and evolve them with statistics-driven pseudo-games like rotisserie baseball or fantasy leagues. And, as immortalized in the movie *Moneyball*, data miners arguably reinvented the way the game is played, players get acquired and traded, and baseball teams are built.

MLB's Many Data Platforms

Nobody does baseball—and baseball data—better than Major League Baseball (MLB). With America's game going big-time into data, MLB has built a modern infrastructure designed to support and scale the league's digital product platforms. Not only does MLB bring data to ball clubs and fans, it does so in real-time, minute-by-minute as the games are played, and on-demand

through a variety of apps designed to support a variety of data consumption scenarios.

The MLB Ballpark app is designed to support fans as they visit major league ballparks around North America. The MLB at Bat app provides access to live gameday audio, in-game highlights, breaking news, and expert analysis from teams across the entire league. The MLB First Pitch app lets fans (and teams) follow rookies and young players from all 30 team organizations (including AA and AAA) with news, pitch-by-pitch game information, scores, schedules, video clips, live game broadcasts, and audio streams.

In the same vein, MLB offers additional apps for prizes based on breaking records (MLB Beat the Streak), in-game animations and recreations (RBI Baseball 19), plus Home Run Derby and Line Drive games. All of these apps are intensely data-driven, with audio and video content as well as reams and reams of stats, images, models, baseball physics simulations, and more.

Suffice it to say that MLB understands that data is life and breath for baseball. The more ways they can slice, dice and deliver that data, the greater the returns they can get from the game—emotionally, physically, and financially. This makes timely, efficient access to its platforms and data a critical ingredient for current and future success.

MLB Beefs Up Its Infrastructure



In mid-2019, MLB partnered with Flexential to add a new high-end data center presence in the Denver-Englewood area of Colorado. The league sought to implement a data management system built around highly reliable, scalable and proximity-based data centers. Flexential's near edge data centers provide a key component in permitting the league to provide the kinds of infrastructure it needs to support league ball clubs and their fan bases.

The new Flexential data center in Denver-Englewood complements the league's East Coast presence (of the 30 teams in MLB, nearly half are in the northeastern U.S., plus Toronto). This new central location is not only reliable and secure, it also reduces average data retrieval latency by half.

At the same time, Flexential's near edge hybrid cloud architecture provides state-of-the-art connectivity and resiliency capabilities for MLB, its teams and fans. To lower environment stresses and costs, the new location sports an industry-leading Power Usage Effectiveness (PUE) rating. Xcel Energy provides it with electricity,

of which 28 percent comes from renewable resources (solar and wind).

Overall, the move closer to the customers in MLB's data network provides improved real-time performance, outstanding resilience and availability, and a positive, enjoyable user experience for fans across the world.

Room to Grow with the Game

Given that MLB continues to develop and deliver new data platforms and apps for baseball fans and professionals alike, a data center architecture that can handle increased demands for reliability, connectivity, and comprehensive service packages is critical to success. The league's growing online media platforms and their presence and use can be accommodated, and capacity and capability added as demand (and consumption) continue to increase.

The number of active Internet-connected devices is expected to increase to 125 billion by just 2030, which would put about 15 connected devices into the hands of every global consumer. At the same time, bandwidth demand and networking requirements will also expand



Figure 6: Major League Baseball is one example of how data is being moved to the edge to enhance user experience.

dramatically. By taking connectivity for media and data delivery to the edge, MLB remains well-positioned to provide a pulse-pounding user experience for each and every baseball fan.

Near edge data centers minimize latency for user activities, and deliver the best and fastest access to media, games and data. Far-edge connections lead to a high-speed backbone, which supports nationwide access and services for MLB, its platforms, partners and services (see **Figure 6**).

Even bandwidth-intensive streaming media gets the service and handling it needs, so fans can watch games, highlights, play-by-play analyses, and more. Together, MLB and Flexential take fans to the edge of their seats—and beyond.

CHAPTER 4


Taking Health Monitoring into the Home

Healthcare IT makes up an increasingly large component of the overall IT industry, as healthcare providers from individual practitioners to global hospital and clinic systems continue to increase their investments in technology and security. In fact, there has probably never been a bigger government-driven digital transformation than that brought about by Electrical Medical Records (EMR) legislation, in effect in the United States since 2012. Although small, individual practitioners can apply for waivers, most other medical practices and businesses are well into this transformation, if not completely done with the process.

By themselves, EMR represents huge volumes of data. According to the Centers for Medicare & Medicaid Services,¹ healthcare expenditures in the U.S. accounted for 17.9 percent of GDP in 2017. In that same year, healthcare spending hit \$3.5 trillion, which translates to

¹ <https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical.html>

ANNUAL GROWTH IN HEALTH EXPENDITURES



Source: Centers for Medicare and Medicaid Services,
Office of the Actuary: National Health Expenditure Projections 2018-27

Figure 7: Healthcare expenditures in the IT field are exploding.

\$10,739 for each person in the U.S. population at large. In 2017, healthcare spending in the U.S. grew at a rate of 3.9 percent, 1.8 percent ahead of that year's inflation rate. In fact, growth of healthcare spending in the U.S. has outpaced inflation for at least two decades, and is poised to continue into the future, as **Figure 7** shows.

Today, IT stands at the forefront of healthcare spending because it supports and even drives so much of its capacity and capability. That explains why growth in IT healthcare spending,² at 8.8% for 2018, is more than double the growth rate for overall spending on healthcare in general. Physicians, public health professionals, hospitals and clinics, and insurance companies, are just starting to extract value from the huge masses of data routinely acquired and managed during the healthcare delivery process.

Over time, legislators and medical professionals alike have come to understand that data analytics can help lasso healthcare costs and wasted resources that sometimes appear ready to burst beyond human control.

Beyond Electronic Health Records (EHR)

But there's a new data-driven frontier opening up in healthcare that goes beyond information captured during office or hospital visits, medical consultations, routine checkups and the like.

Consider an analogy: Acquisition and use of medical data is advancing from the snapshot era to the streaming

² <https://www.healthdatamanagement.com/news/2018-tech-budgets-to-rise-about-88-for-healthcare-organizations>

era. That is, rather than relying on occasional visits or checkups when health information is acquired, assessed and actions taken (these are essentially snapshots of patient condition at a single point in time), certain medical conditions may instead be monitored constantly so that medical data can track patients throughout daily, weekly, monthly and even yearly health cycles (this is the streaming model).

When IoT Sensors Monitor People and Health

Such an approach to healthcare monitoring involves the use of sensors. Some of these qualify truly as “Internet of Things” devices in the sense that they are dedicated bits of equipment, usually specialized to monitor specific types and sets of medical data. In the past decade we’ve seen the introduction of numerous kinds of implants or plug-in medical devices for patients with specific, chronic conditions.

While older versions of such devices have been more or less standalone, programmed to acquire data and manage certain medical parameters (heart rate, blood sugar, brain waves, and so on), newer versions of these devices include wireless communications capabilities that permit them to upload data either periodically or

constantly. This is starting to define a whole new frontier for medical data.

Then, too, the near-universal proliferation of smartphones, along with smartwatches and strap-on fitness monitors, has opened up opportunities to acquire—and wirelessly share—medical data of all kinds. Because the smartphone offers the ability to acquire short-range communications from implants or plug-ins (usually via Bluetooth or NFC), it can also act as a relay to convey data from such short-range, low-powered devices to medical software on some server for acquisition, analysis, and filtering.



The near-universal proliferation of smartphones, along with smartwatches and strap-on fitness monitors, has opened up opportunities to acquire—and wirelessly share—medical data of all kinds

Perhaps more importantly, ongoing monitoring of health conditions through such means can also prevent or preempt serious medical emergencies. The Apple Watch's ECG feature, for example, is FDA-cleared³ to

³ <https://www.self.com/story/cardiologists-apple-watch-heart-monitoring-features>

detect atrial fibrillation, and has been credited with saving some users from possible heart attacks.

Diabetes Applications

Normally, medical implant or plug-in devices make most sense for monitoring and managing chronic health conditions. According to the CDC,⁴ as of 2018 more than 100 million Americans have been diagnosed with diabetes or prediabetes. Taking 100 million as a hard and fast number, that represents a staggering 30.5 percent of the U.S. population. This makes diabetes not just a big healthcare concern, but also makes diabetes devices a huge global market of approximately \$22 billion for 2018,⁵ with the U.S. representing 25 percent – 30 percent of that overall outlay.

Frontiers in Managing Chronic Conditions

Indeed, diabetes poses economic challenges as well as health concerns for those coping with the disease. A 2000 study⁶ estimated that individuals with the diabetes

⁴ <https://www.cdc.gov/media/releases/2017/p0718-diabetes-report.html>

⁵ <https://www.globenewswire.com/news-release/2019/06/10/1866428/0/en/Diabetes-Devices-Market-To-Reach-USD-38-53-Billion-By-2026-Reports-And-Data.html>

⁶ <https://care.diabetesjournals.org/content/24/2/257>

can expect an annual loss of about one-third of their normal earnings, owing to lost productivity and medical absences. Do the math, and that amounts to 10 percent of the entire workforce's earning capacity.

Furthermore, serious complications of diabetes include heart disease, nerve damage, kidney damage, blindness, and circulatory issues sometimes severe enough to require amputation of one or more extremities. All this makes a strong and compelling case to use diabetes devices to help patients monitor blood sugar and other factors and keep them as healthy as their conditions permit.

Expediting and Automating Healthcare at the Edge

The key to successful treatment for chronic conditions is careful, proactive and—where appropriate—timely management and intervention. This makes communicating diabetes health data to physicians and monitoring applications not just important, but also time sensitive.

Also, healthcare governance requirements (for example, HIPAA) mandate that health information be kept private and secure, both in motion and while at rest in storage. The cloud is an obvious kind of glue to tie together widely-dispersed diabetes monitors in patients' homes (or

on/in their persons) and monitoring applications and communications.

But the public cloud may introduce latencies that could, in some circumstances, adversely affect patient outcomes. A near edge hybrid cloud solution, with its ability to minimize latency and ensure that responsible and involved parties (caregivers, physicians, clinics, insurance companies, and so on) receive timely notifications, is a much safer and better fit for this kind of data-handling situation.

Making Healthcare Better

Flexential's powerful far-edge infrastructure and high-speed links between data centers means that important alarms and alerts can get to physicians and caregivers quickly enough to make sure proper medical responses are timely and apt.



Also consider that Flexential's built-in security and record protections comply fully with HIPAA and other privacy regimes. This creates an ideal fit for medical devices and information of all kinds, with diabetes devices as just one particularly key example.

Technology Underpinnings Make Things Happen

Near edge hybrid cloud solutions make it easy for diabetes devices to do their jobs with maximum dispatch and effectiveness. Near edge data centers make sure that incoming diabetes data is received as quickly as modern technology (and last mile Internet infrastructures) permit. At-hand data filtering and analysis capabilities can quickly zero in on diagnostic indicators that suggest or demand immediate action or intervention.

CHAPTER 5

Put Hybrid Edge Computing to Work

As explored and explained in the previous chapters, hybrid edge computing is uniquely well-suited when it comes to making use of the nearly limitless capabilities and capacities, and incredible services, found in today's myriad of private and public cloud offerings.

What makes hybrid edge computing so powerful and effective, in fact, is its ability to position data handling and compute capability at the near edge of the Internet. That means it's close enough to users and endpoint devices to minimize latency. In turn, this provides positive and efficient experiences for end-users, since M2M interactions and transactions are fast and reliable enough to keep the wheels of business turning and patients healthy.

Now Is the Time to Develop Edge and Distributed Strategies

In Chapter 3, you learned about Major League Baseball's partnership with Flexential to build upon its already-startling successes in generating interest,

excitement and revenue from the media- and data-rich experience that America's pastime provides. Looking forward, MLB will continue to expand its platforms and services, and look for more and better apps with which to entertain and inform its ball clubs and fan base. Savvy organizations, both public and private, must do likewise to help them meet and exceed their goals and objectives.



What makes hybrid edge computing so powerful and effective, in fact, is its ability to position data handling and compute capability at the near edge of the Internet. That means it's close enough to users and endpoint devices to minimize latency.

Thus, it's essential to design and develop hybrid edge and distributed strategies for data and services right now. Your competitors are surely doing likewise, and you can't afford to trail behind and risk losing not only a competitive advantage but fail to take advantage of key opportunities to grow and expand your solutions and offerings.

Looking forward, it's essential to factor edge computing into the design and development of future products and services. This enables companies and organizations to

count on the kinds of positive user experiences that keep them coming back for more and looking for reasons to do more with what's made available to them.

On the M2M side of the business, this means finding more synergies and opportunities to grow through further deployment of IoT, better use of big data and analytics, and more effective combinations of—and communications among—systems, all the way from research and development, to manufacturing or implementation, supply chain management, the entire sales cycle, and ultimately to customer service and support.

With positive end-user and customer experiences at each step along this path through the lifecycle, more and better opportunities must present themselves, along with improved options for productivity, revenue, and earnings.

The Tether That Binds for a Successful Future

Edge data center providers with strong managed services and support can help companies manage data, provide improved interactivity and response times, and get the most out their network traffic and data (in terms of economy, efficiency, and effectiveness).

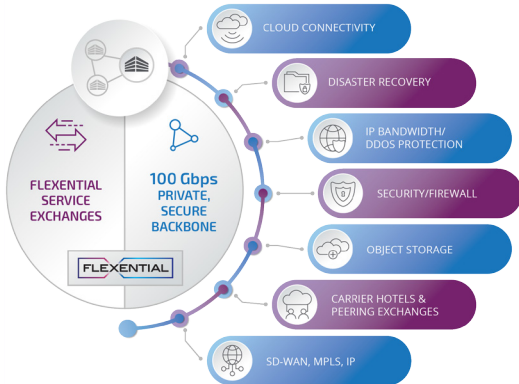
At the same time, organizations can improve their connectivity and capacity from the near edge, into the far edge, and onto the core. Proper situation and selection of compute, storage and networking facilities at each step along this path will help organizations take care of business, users and customers, while making best (and most cost-effective) use of resources to tie everything together.

The key to online success is to tether data, services and interactivity as close to the near edge as possible, and to elaborate the supporting connections and infrastructure behind the scenes to support growth and new developments.

Organizations that understand how to keep time-sensitive services close to the edge will gain user/customer allegiance and enthusiasm. Learning how to move data between the edge and the core quickly and efficiently keeps the overall environment humming and offers the best returns on technology investment. Manage transitory or time-sensitive data at the edge and deliver on what's actionable or needs to be kept to the core. The technology of taking data to the edge will continue to evolve, and so should your strategy and processes.

About Flexential

FLEX ANYWHERE™



FlexAnywhere™ Infrastructure

- 100 Gbps network backbone
- IP BW with automated DDoS Scrubbing
- >13,000 cross connects
- 25 carrier exchanges
- 40 Hybrid edge data centers
- 14 Private cloud nodes
- 12 PoP connections to carrier hotels
- Direct connectivity to multiple subsea cables to the Asia-Pacific region
- 80+ on-net premier carriers to choose from
- 300+ network, data center, managed service and cloud providers

Figure 8: Flexential provides the necessary infrastructure to move your data to the edge.

Get More Information

As they move into the cloud, companies and organizations must remember the FOG strategy (“Feet On the Ground, head in the clouds”). Flexential is ready to help, with consulting and architectural input to make best use of its edge computing capabilities (see **Figure 8**). Visit the website,¹ where you’ll find additional information and reading about their data centers,² solutions (colocation,³ connectivity,⁴ data protection,⁵ cloud,⁶ security and compliance,⁷ managed solutions,⁸ and FlexAnywhere⁹), and professional services.¹⁰



¹ <https://www.flexential.com/>

² <https://www.flexential.com/data-centers>

³ <https://www.flexential.com/colocation>

⁴ <https://www.flexential.com/connectivity>

⁵ <https://www.flexential.com/data-protection>

⁶ <https://www.flexential.com/cloud>

⁷ <https://www.flexential.com/security-and-compliance>

⁸ <https://www.flexential.com/managed-solutions>

⁹ <https://www.flexential.com/flexanywhere>

¹⁰ <https://www.flexential.com/flexential-professional-services>



Flexential offers flexible and essential services that help organizations optimize their journey of IT transformation while simultaneously balancing cost, scalability, compliance and security. The company is committed to building trusted relationships and delivering tailored solutions that suit the individual needs of its customers. Flexential is deeply invested in the success of its customers, who trust it to deliver core data center solutions of colocation and connectivity, as well as cloud, managed solutions and professional services. Flexential's robust suite of assets spans 21 markets and comprises 40 highly redundant and connectivity-rich data centers. For more information on Flexential, please visit www.flexential.com or email info@flexential.com

