

Brought to you by:



Confidential Computing

for
dummies
A Wiley Brand

Protect data and
applications in use

Secure data
sharing

Protect intellectual
property



Ambuj Kumar

Anand Kashyap

Fortanix Special Edition

About Fortanix

Fortanix is the data-first security company. As the industry's first and largest provider of Confidential Computing solutions, Fortanix decouples data security from infrastructure. Fortanix solutions empower organizations with centralized controls to secure data spread across clouds, applications, Software as a Service (SaaS), databases, and data centers. Enterprises worldwide, especially in privacy-sensitive industries like healthcare, fintech, financial services, government, and retail, trust Fortanix for data security and privacy. Fortanix investors include Intel Capital, Foundation Capital, Neotribe Ventures, and In-Q-Tel. For more information, visit <https://fortanix.com>.



Confidential Computing

Fortanix Special Edition

**by Ambuj Kumar
and Anand Kashyap**

**for
dummies®**
A Wiley Brand

Confidential Computing For Dummies®, Fortanix Special Edition

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

Copyright © 2022 by John Wiley & Sons, Inc.

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. Fortanix and the Fortanix logo are registered trademarks of Fortanix. 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: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY 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 ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

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: 978-1-119-89663-0 (pbk); ISBN: 978-1-119-89664-7 (ebk). Some blank pages in the print version may not be included in the ePDF version.

Publisher's Acknowledgments

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

Project Manager:

Carrie Burchfield-Leighton

Sr. Managing Editor: Rev Mengle

Acquisitions Editor: Ashley Coffey

Business Development

Representative: Cynthia Tweed

Special Help: Steve Kaelble

Content Refinement Specialist:

Saikarthick Kumarasamy

Table of Contents

INTRODUCTION 1

 About This Book 2

 Foolish Assumptions 3

 Icons Used in This Book..... 3

 Beyond the Book..... 4

CHAPTER 1: Introducing Confidential Computing 5

 Outlining the Data Life Cycle..... 5

 Moving to a Data-First Approach..... 6

 Recognizing the Need for Confidential Computing 12

 Enabling Innovation 12

CHAPTER 2: Complying with Privacy Regulations 15

 Protecting the World 15

 Reading the Rulebooks..... 16

 Getting into the GDPR 16

 Understanding Schrems II 17

 Staying healthy under HIPAA..... 18

 Dreaming of California data 18

 Putting on SOX 19

 Charging into PCIDSS compliance..... 19

 Encountering Roadblocks..... 20

CHAPTER 3: Putting Confidential Computing to Use..... 21

 Healthy Innovations 21

 Facilitating Fintech..... 23

 Banking on Data 24

 Revolutionizing Retail..... 26

 Serving Government 27

CHAPTER 4: Understanding the Confidential Computing Ecosystem..... 29

 Introducing the CCC 29

 Growing the Market 31

CHAPTER 5:	Choosing a Solution	33
	Boiling the Ocean	33
	Connecting With Experts	35
CHAPTER 6:	Sharing Confidential Computing with the Team	37
	Gaining Buy-In	37
	Moving Forward.....	40
CHAPTER 7:	Three Examples of Deploying Confidential Computing.....	41
	Advancing Genomics.....	41
	Connecting with the Best Data	42
	Overcoming Cloud Hurdles.....	43

Introduction

In the world of information technology (IT), the next big, new thing comes along just about every week, it seems. Innovations are frequently amazing, sometimes stunning, often quite disruptive. But also, just about every week, along comes another story of cybercrime and costly data breaches.

Fitting right into both contexts is a fast-growing technology known as *Confidential Computing*. It's far more than a "flavor of the week" IT innovation in that it tackles a key aspect of data security that hasn't been satisfactorily addressed in the past. But it's also more than simply a data security advancement because its use can accelerate innovations of all kinds — even literally lifesaving breakthroughs.

Until now, you could feel reasonably good about the security of your data while it was at rest, idle, and encrypted in a database on a server somewhere on-premises or in the cloud. You could feel okay about that data while it was in transit across public or private networks, too, because it can still be encrypted before traveling from one place to another.

But take it out of its secure lockbox to process it, and it can be a sitting duck, waiting for a malicious actor to pounce. And it *has* to come out of that box and be unencrypted because an application has to see the data in order to use or analyze it. No surprise, then, that many organizations are wary about allowing their most sensitive data to be processed on infrastructure over which they have little control.

Confidential Computing provides intrinsic security to applications and data while in this critical stage of existence — in use — wherever they might be deployed. The technology isolates sensitive data and code during data processing. Data in use is protected by RAM encryption and hardware-based technologies. If it's in the cloud, even the cloud provider can't see the data during processing, and neither can any potentially malicious, privileged users.

This really is a game-changer, for a number of reasons. On the surface, it's a powerful advancement in security because it protects data during that unavoidable point in its life when it has to

shed its protective armor. But that's only part of the value. The biggest causes for celebration are the ways Confidential Computing enables your organization to do things you may not have been able to do before.

The protections it provides — along with its capabilities for attestation, among others — open the door for doing business in some hyper-regulated places that had previously been less than welcoming. Regions of the world, such as Europe, have stringent data protections that cause some businesses to steer clear, and the same can be said about industries such as health care. Confidential Computing opens the door to new markets.

But it also enables new activities. For example, different and even competing organizations are much better able to collaborate with their data. Federated machine learning may allow them to achieve greater things together than they could separately — maybe even lifesaving healthcare breakthroughs — and they can use their data collaboratively without letting the other parties actually see it.

Unlike some disruptive new technologies, this one didn't emerge from someone's garage. Confidential Computing is supported and boosted by some of the biggest names in the world of IT, all working together.

About This Book

Confidential Computing For Dummies, Fortanix Special Edition, is your introduction to this exciting new world of advanced data security and the possibilities it enables. You get a high-level overview of what this technology is, why it's needed, what it can do for you, and how to get started. Flip through the pages for an outline of the data life cycle and the three states of data, and why organizations need security that focuses on the data itself, not just the infrastructure.

You find out more about the many different systems of privacy regulation that make the use of data complicated — and make mistakes very costly. You get a sense for the many innovative things you can do with Confidential Computing, whether your business is healthcare, finance, retail, or government. Find out more about the ecosystem bringing this all to life, and get pointers for getting started and bringing your teams onboard.

Foolish Assumptions

In preparing this book, we made a few assumptions about you, the reader:

- » You're part of the cybersecurity team, or perhaps the regulatory/compliance team, at a medium-size to large enterprise.
- » Your organization relies on the maintenance and use of sensitive data, perhaps related to healthcare, finance, retail, or government.
- » You may want to collaborate with another organization that's using your sensitive or regulated data, but you haven't yet found a viable solution that meets your security and compliance requirements while also delivering the performance that the application requires.

Icons Used in This Book

Throughout this book, we included some helpful icons. More than just decorations, they're there to highlight important pieces of information.



REMEMBER

You may be in a hurry while you're reading, but please don't skip over the key points spotlighted by this icon.



TIP

The whole point here is to create a useful tool for you, and this icon points out an actionable idea for you to consider.



TECHNICAL
STUFF

Confidential Computing aims to make complex security simpler, but this icon identifies some technical details.



WARNING

The whole topic of data security is worrisome, and this icon points out a specific issue to watch for.

Beyond the Book

There's a lot of information between the covers of this book, and we hope you come away with a greater understanding of Confidential Computing, why you need it, what it can do for you, and where to go from here. That said, there's plenty more to know than we could fit into these pages.

If you have an appetite for more information about Confidential Computing, check out these resources:

- » **An online overview with more details on Confidential Computing:** fortanix.com/products/confidential-computing-manager/what-is-confidential-computing
- » **An e-book exploring Confidential Computing, its business value, and industry-specific benefits:** resources.fortanix.com/demystifying-confidential-computing-ebook
- » **An overview of how Fortanix can help with your Confidential Computing journey:** fortanix.com/products/confidential-computing
- » **Industry players who are accelerating the adoption of Confidential Computing:** confidentialcomputing.io
- » **Research and in-depth detail from the experts at the Confidential Computing Consortium:** confidentialcomputing.io/white-papers-reports

- » Understanding the life cycle of data
- » Protecting data first
- » Outlining the concepts of Confidential Computing
- » Using Confidential Computing to enable innovation

Chapter 1

Introducing Confidential Computing

The technology world, and everything in our lives that it touches, revolves around data. There's unimaginable power in data but also incredible risk if that data falls into the wrong hands.

This chapter gives an overview of the three basic states of data and why it's increasingly important to focus security efforts on the data itself instead of infrastructure perimeters. It introduces the concept of *Confidential Computing*, which protects data even while it's in the midst of being processed. And this chapter touches on how Confidential Computing can accelerate the wonders of innovation.

Outlining the Data Life Cycle



REMEMBER

To get a handle on protecting data, it's helpful to take a step back and think about the life cycle of data and its various states of existence. Though there's little that's simple about the world of data, there's one basic truth that's surprisingly elementary. Data exists

in one of three basic states — it's either at rest, it's in use, or it's in transit.

Data is considered to be “at rest” when it's being stored. If it's being processed somehow, it's said to be “in use.” And if data is traveling across a network, that data is “in transit.”

Information technology security specialists have dedicated unfathomable time and effort figuring out how to protect data at rest, as well as data in transit. Powerful encryption technology can provide solid protection for data at rest, as well as data that is in transit from here to there.

Data at rest may, for example, be stored on a hard drive inside a server in a data center, whether on-premises or in the cloud. Someone could physically walk off with that hard drive and be unable to do anything harmful with the data on it, as long as there's proper encryption in place.

Data in motion is safe, too, if it's encrypted end-to-end. Even if a bad actor intercepts that data on its journey from one place to another, he won't be able to access the data.



WARNING

Of course, data doesn't just sit on a server or move from one server to another. Data is processed and used — and that's when data can be particularly vulnerable. Data in use is vulnerable to unauthorized access at runtime. That's prime time for bad actors to tamper with valuable data.

With more and more data crunching happening in cloud infrastructure, the task of protecting data is more complicated than ever. It's no longer sufficient to build protection solely around the places where data rests or the paths data takes while in transit. Protecting data across the whole life cycle requires protecting it while in use, and that means you need methods that focus on the data itself.

Moving to a Data-First Approach

It wasn't all that long ago that data security relied largely on securing the on-premises infrastructure. With this line of thinking, if the infrastructure was secure, the data could be considered secure, too. In this era, firewalls were seen as sufficient protection.

That way of thinking became outdated as users began employing cloud infrastructure. With data and operations in the cloud, and not within a data center, it was no longer possible to simply build a perimeter and keep the operation secure.

The world of the cloud brought along such security challenges as dynamic hosts, dynamic machines, upscaling and downscaling. The focus turned from securing on-premises operations to securing cloud infrastructure.

That is, of course, much more complicated. Data may be in public clouds, in private clouds, in software-as-a-service (SaaS) clouds. Data may live in SQL databases, NoSQL databases, or other storage systems. That's a lot of different infrastructure to protect, and you may be using multiple variations of this infrastructure at the same time.



REMEMBER

The shocking headlines you read about data breaches suggest that bad actors are pretty good at staying a step ahead of such security efforts. It's not for lack of trying on the part of data security professionals, and it's worth noting that all past security efforts have had the ultimate goal of keeping data safe. But they've focused on the infrastructure rather than the data itself — and data security and infrastructure security are really two separate things.

As the world continues to become more and more data-first, the focus of data security must become data-first, too. This way of thinking is sometimes called “Security 3.0.” It's no longer about the on-premises infrastructure or the cloud infrastructure. It's focused on securing data, wherever it is across the three states, and wherever it happens to be residing at any given moment.

Think of it as decoupling security from infrastructure. The concept of data-first multi-cloud security is an entirely different paradigm that secures data even under circumstances that may normally seem dire.

Like when a network is breached. Or when a cloud services provider is served government papers. Or when root users and their passwords are compromised. Decouple your data security from your infrastructure security, and you'll also decouple your data from problems that might arise with that infrastructure. Attack vectors that once seemed deadly become irrelevant.

WHY DO CUSTOMERS SPEND \$200 BILLION/YEAR ON CYBERSECURITY?

Take the following quiz to see if you can figure out why customers spend \$200 billion each year on cybersecurity. Is the answer

- A. To make security companies a ton of money?
- B. To secure their cloud, network, firewall, API, applications, logs, patches, VPC, hardware, software, or firmware?
- C. To secure their data?

If you answered “C,” you are correct because data is perhaps the most valuable enterprise asset. And there are serious consequences for enterprises in the event of data compromise or privacy violations.

The data-first sensibility isn’t a futuristic fantasy but is happening today within many organizations that are successfully securing their data regardless of infrastructure.



REMEMBER

How is this possible? Think for a moment about what data security really is. As Figure 1-1 illustrates, it’s a combination of identity and access control:

- » **Identity:** This refers to how the organization keeps track of different users, applications, and data elements.
- » **Access control:** After the organization is able to identify users, applications, and data elements, it can assign access controls. A user can be granted access to certain groups of data but not others. An application can be granted access to specific sets of data and denied access to others.

Data security, then, is a matter of enforcing the rules that are established through these identity-based access controls. Sounds simple enough, but the devil is in the details. How do you enforce the rules in a reliable manner?

In the past, the infrastructure provided well-defined perimeters that helped make that enforcement possible. In today’s multi-cloud world, that’s challenging, given that you are dealing with multiple providers with different policies and topologies and access control systems.



FIGURE 1-1: The basic elements of data security.



REMEMBER

The solution comes through the use of encryption. It's a simple enough concept to understand, in that encryption is mathematically derived and requires possession of a key before data can be accessed and processed. As long as the data is encrypted, it's safe from trouble even if it falls into the wrong hands through a breach.

Confidential Computing takes this simple concept and makes it work in a complicated multi-cloud world. It ensures that data protections remain solid not just when the data is at rest, not just when it's in transit, but also when it's in use.

That last state is where data is most vulnerable today. It has made it safely from the server, across cyberspace to the place where it will be accessed and used. After it has been decrypted there, it has a target on its back, in danger of foul play.



TIP

That's where Runtime Encryption (a concept developed by Fortanix) comes into play. It's pretty much what the name implies — encryption that keeps data safe while an application is running and processing the data. With this protection, the data can only be accessed and used by the application that has been specifically granted permission to access it.

This concept is the missing link, essentially the Holy Grail of data security that defines how Confidential Computing is different from other approaches. It's an essential, additional layer of protection that remains intact even if there's a breach in the infrastructure where the data is residing.

With Confidential Computing, processing happens within what's known as a *secure enclave*. An enclave is a place that's different from what surrounds it. In the case of Confidential Computing, what's different is that it's a safe place where the data within is protected even as it's being processed by applications running within the enclave.



REMEMBER

A secure enclave is a protected memory region — a *Trusted Execution Environment* (TEE). The secure enclave is a special section reserved within the CPU, with memory encrypted and protected by a key that's unique to the CPU and the application running in the enclave.

The enclave offers a place for an application to execute securely, enforced at the hardware level by the CPU. Data is encrypted in memory, and it's only decrypted when it's being used inside the CPU. Indeed, it can *only* be decrypted within that enclave on that CPU.

Think of secure enclaves as a safe private place where nosy neighbors cannot see what you are doing. Various CPU manufacturers and cloud service providers offer their own implementations of enclaves, including Intel SGX and TDX, AMD SEV, AWS Nitro, and Arm TrustZone.

After an enclave is created, the data being processed by the application running inside is secure. It's protected in that third state of data: in use. It's secure even if the hypervisor is compromised or if a bad actor has gained the kind of root access that would doom data security under other paradigms.



REMEMBER

Because the TEE runs in an isolated environment, it provides full data confidentiality and integrity, as well as code integrity. Unauthorized actors can't view data while it's in use inside the TEE, and they can't add to it, remove it, or alter it. Same goes for code — there's no way that unauthorized people can alter it, add to it, or remove code that's executing within the TEE.

Not only is your data protected, but also you can get proof of that protection through Confidential Computing. You don't have to take your cloud provider's word that security is under control. Confidential Computing technology has an attestation feature confirming to third parties that the data is safe inside the enclave.

Also good for your peace of mind is the centralized control you can gain through the use of such Confidential Computing solutions. This is a multi-cloud world, and your data may be in any number of different clouds, in databases, in SaaS applications, storage systems, data lakes, and data warehouses. Wherever it is, it's secure and under control from one centralized point of control.

As shown in Figure 1-2, security has been about infrastructure up until now, focused on identity, endpoint, and network. What has now emerged is the fourth pillar of security, centered on data.




FIGURE 1-2: The fourth pillar of security, focused on data.



You may have familiarity with some other technologies that aim to provide this kind of data security. There are, indeed, some options, though they have some comparative drawbacks. For example:

- » **Fully homomorphic encryption (FHE):** While Confidential Computing supports all databases and applications, FHE offers quite limited support in that regard. It has high levels of privacy and no constraints on sample size, but its longitudinal privacy varies. Its performance is poor compared with Confidential Computing, and deployment is difficult. And it doesn't provide for the integrity of the ciphertext nor the computation, which means it needs to be combined with other cryptographic or Confidential Computing techniques.
- » **Multi-party computation (MPC):** MPC cryptography, too, is extremely limited in support for databases and applications, with varying levels of privacy depending on the

implementation. Performance is poor and ease of deployment is not guaranteed. Security benefits can also be limited especially if the multiple parties use the same technology to generate their keys, creating a single point of failure.

Recognizing the Need for Confidential Computing

How much data is out there across the world? An almost unimaginable quantity. By some estimates, somewhere over 30 zettabytes of data were created, captured, copied, and consumed around the world in 2018. That's about 30 trillion gigabytes. By 2020, that number tripled, and by 2025 it could potentially be three times what it was in 2020. Suffice it to say that more data than ever is in need of protection today.

And it's also a fair statement that there are more threats than ever. The FBI gets a couple of thousand internet crime complaints a day, more than half of all consumers have been cybercrime victims, and the global cost of cybercrime is expected to grow at least 15 percent annually.



WARNING

The stakes, thus, are incredibly high, given the damage that cyber incidents can do to business operations and reputations. Regulatory interest (check out Chapter 2 for more info) is incredibly high, too — raising the stakes even further.

Throw in the fact that data storage and processing are happening in more and more disparate places every day, and you realize the need for a whole new approach to data security. Yesterday, you could guard the perimeters. Today, you need to guard the data itself, wherever it is, including when it's in use.

Enabling Innovation

Confidential Computing is certainly a way to rest easier at night, knowing that data has much greater protection from the ill intentions of bad cyber actors. But that peace of mind is really only part of the story.

Consider this. If your data itself is protected, and its security doesn't rely on the infrastructure, your sensitive data or code can more safely venture out into the otherwise dangerous world and find new ways to be powerful and useful.



TIP

That means you can share your most sensitive data outside the safe confines of your own organization. It's still safe, even when it's being used in infrastructure over which you have no control. Your sensitive data can be sent to any server, any cloud, and you can rest assured it is still safe from harm.

With that protection in place, your organization is free to collaborate and innovate with others using that data. Without Confidential Computing, your innovators likely feel confined to their safe spaces, hamstrung by the need for more infrastructure-based data protection. With Confidential Computing, that harness is removed, the doors are opened, and the potential for innovation is unbridled.

Check out Chapters 3 and 7 for some specific insights of how these doors can open. For example, financial institutions that are competitors in most respects can collaborate on efforts that help uncover potential money laundering. That's a whole lot easier to achieve with access to the data of multiple institutions, and the technology lets them upload encrypted data into an enclave where such work can take place — without anyone gaining even a glimpse of a competitor's sensitive data. Similar collaborative innovations are enlightening the world of healthcare.

Ultimately, Confidential Computing moves the security conversation from what users *can't* do to what they *can*. In the past, data security efforts have sometimes been seen as necessary but inconvenient. Confidential Computing creates a secure environment that doesn't get in the way of users and application teams.

And that reality is a fantastic “what's in it for me?” element (see Chapter 6 for more info). It's one more selling point to help you gain buy-in for implementing Confidential Computing in your organization.

- » Keeping data private globally
- » Getting familiar with privacy regulations
- » Encountering regulatory hurdles

Chapter 2

Complying with Privacy Regulations

The more the world relies on data to function, the more people realize how valuable data is and how much it needs to be protected. Privacy regulations are proliferating all around the world. This chapter discusses why those privacy regulations are there and how they can impede innovation and goes into some specifics about the various regulatory systems.

Protecting the World

There are few countries that don't have at least limited privacy regulations. And the number of countries covered by robust or even heavy regulation is growing all the time. Regulations have been the norm in specific industries for a while now. Healthcare is one key example, and it's easy to see why it's important to keep health information safe, given the intimate details they include, as well as financial data. No wonder hackers want to access them so much.

Financial industries have long had stringent data protection regulations, too, because bad actors can cause terrible trouble if they can access financial data. That said, the latest regulations are much more sweeping. They cover not just specific industries but whole populations. Before long, about three-quarters of the world's population will be covered by some sort of data privacy regulation.



WARNING

For companies that deal with data, the challenges are even more complicated than they sound. Your organization may be covered by more rules than you think because when it comes to privacy regulations, they apply not just where you're based, but where you're doing business. For example, if you're in America and read about the stiff regulations in Europe, don't breathe a sigh of relief and think it's not your problem. If your organization does business with people in Europe, then you're on the hook, like it or not.

Reading the Rulebooks

No matter what you're playing, it's hard to win the game if you don't know the rules. And when it comes to regulations related to data and privacy, the rulebooks vary depending on where the game is being played.

Europe is governed by strict data privacy rules, for example, and you'll find a set of regulations that is specific to California. But it's more than just geography. If you're operating anywhere near the world of United States healthcare, there's a set of rules that you definitely don't want to ignore. If the technology work you're doing may impact investors, you need to know different rules. If you're dealing in the business of payment cards, there's a whole set of standards to observe.

This section gives you an overview of some of the various regulatory systems that can impact what you're up to from a technology perspective.

Getting into the GDPR



REMEMBER

The European Union's General Data Protection Regulation (GDPR), which took effect in 2018, is one of the world's most stringent privacy and security laws. It's also one of the most feared because there are significant consequences for anyone who runs afoul of its requirements.

The GDPR is large and far-reaching, and it governs what organizations must do to protect the privacy and security of the data they have about individuals. There's a lot to it, and not all of it is as specific as you may imagine, so you need to err on the side of caution. One part you hear a lot about is the “right to be forgotten,” which allows people to ask that organizations delete their personal data.

You want to work with compliance specialists to ensure your organization is doing what it's supposed to be doing, but there are a couple of key points that are fairly easy to grasp (and a bit scary, too). First of all, the GDPR applies not just to organizations within the EU but those based elsewhere that offer their products or services within the EU. It also applies to any organization processing personal data of an EU citizen or resident — again, no matter where in the world that organization may be. Secondly, the costs of messing up are high. Penalties may be as high as 4 percent of your organization's global revenue, or 20 million Euro, whichever is higher. Plus, data subjects are allowed to seek damages.

In the context of this book, think of the GDPR as a legal framework enforcing effective data protection. If you have the data, you're responsible for it, so you need to employ the best protections you can find.



TIP

The Confidential Computing solution from Fortanix, for example, has a range of built-in features that help with GDPR compliance — everything from encryption and tokenization, to automatic policy synchronization, to a centralized tamper-proof audit trail.

Understanding Schrems II

Organizations in the United States and Europe have done business across the Atlantic for centuries. Goods went back and forth via ships, of course, but information has also been flowing back and forth since the 1850s, when the first transatlantic cables were run across the ocean floor.

Today, about half of the data flowing to and from the United States is transatlantic. More than half of Europe's global data flows also cross the Atlantic. These data flows are a pillar of digital commerce, and for a few years they were regulated by a framework known as the EU-US Privacy Shield. Among other things, the framework made it easier for American companies to accept personal data from entities in the EU.

Note that we said data flows “were” regulated by this framework. Not anymore. That’s because of a court decision known as the Schrems II ruling. The Court of Justice of the European Union in 2020 invalidated the EU–US Privacy Shield, deciding that it didn’t provide adequate GDPR protection for the data of EU citizens.



WARNING

As a result, many companies that had enjoyed liability protection were suddenly at risk. They were enjoying a bit of a pass from the GDPR but now had to get onboard with compliance, including with data security.

Staying healthy under HIPAA

Before the GDPR came along, the most feared acronym related to data privacy was probably HIPAA. It’s short for Health Insurance Portability and Accountability Act, and it’s a U.S. law enacted in 1996 that sets privacy standards relating to patient data.



REMEMBER

HIPAA focuses on what’s defined as protected health information (PHI). Organizations that handle PHI must establish specific processes and security measures to maintain privacy. A lot of HIPAA compliance revolves around who can access PHI. It basically has to stay private from anyone and everyone except those who specifically need access for providing care to the patient or handling the billing for that care. Guidelines exist for how PHI is shared and with whom.



WARNING

If protected data gets into the wrong hands through a breach, it’s a violation of HIPAA. The penalties for violating HIPAA are steep and costly, and if you think you’re safe because your organization isn’t actually a healthcare provider, think again. The Health Information Technology for Economic and Clinical Health (HITECH) Act came along in 2009 to broaden the scope of HIPAA. Business associates, partners, and vendors are all liable for compliance, too. Be sure patient data is well protected by technologies such as Confidential Computing.

Dreaming of California data

While the United States doesn’t at the federal level have the same kind of sweeping data privacy regulations found in Europe through the GDPR, you can find something along the same lines in California. That would be the California Consumer Privacy Act of 2018, or CCPA.



The CCPA aims to give consumers more control over the personal information that businesses collect about them and creates a variety of privacy rights for people in California. For example, they have the right to know what info a business collects about them and how it's used and shared. They have the right to have much of that info deleted. They can opt out of having their personal information sold, and they must not be discriminated against if they exercise these rights.

As with the other data privacy regulations, organizations will have a hard time complying if they don't have solid controls over their data.

Putting on SOX

The Sarbanes-Oxley Act of 2002 (SOX) established a series of processes applying to publicly traded companies in the United States, along with some foreign companies, plus accounting firms that work with these organizations. It didn't grow out of data security concerns specifically because it was enacted in the aftermath of corporate financial scandals.



Still, it sets some high bars for how data must be treated in order to be in compliance. Beyond addressing financial records and reporting, SOX has regulations related to information technology and data security. Therefore, best practices for SOX include a strong focus on data security.

Charging into PCIDSS compliance

Electronic payments offer what may be the most obvious place where data security is critical. Every business that accepts card payments, every bank that handles those payments, every individual pulling a card out of a wallet — everyone can agree on the importance of payment security.



Enter the Payment Card Industry Data Security Standards (PCIDSS). It's a set of technical and operational requirements for those that accept or process payment transactions. The PCIDSS also apply to software developers behind the applications used in those transactions and the companies that make any devices involved.

The standards go into detail about secure networks, access control, encryption, monitoring and testing, and information security policies. Confidential Computing can be an essential part of compliance, as the transactional environment moves more and more into the multi-cloud world.

Encountering Roadblocks

Privacy regulations, such as the ones listed in the previous section “Reading the Rulebooks,” provide valuable protection for people, but they have the ability to get in the way of effective uses of data as well as valuable potential collaborations for advancements in services and solutions that positively impact people’s lives. Companies must balance their compliance obligations with their need to allow data usage across a multi-cloud environment. They need to be able to share insightful data with partners and collaborators for analysis and such applications as machine learning, without putting themselves in danger of costly breaches and the penalties that would result.



REMEMBER

Confidential Computing offers a comforting assurance that companies can keep moving forward while maintaining safety. They gain the ability to operate in secure enclaves with attestation, confident that their data is completely protected while it’s being processed. Sensitive applications engaged in analytics and artificial intelligence can go about their work, running safely in the cloud while in compliance with regulations. The ability to use attestation to identify specific geographic locations of secure enclaves also enables isolation of workloads to ensure compliance with data sovereignty requirements.

IN THIS CHAPTER

- » Innovating better health through data
- » Facilitating the financial technology sector
- » Helping banking protect its data
- » Creating new retail experiences through data
- » Giving governments new ways to protect data

Chapter 3

Putting Confidential Computing to Use

The potential benefits of Confidential Computing are far-reaching and cross industry boundaries. It's about protecting data, of course, but the most important message is how keeping data safe while in use frees up a whole lot of different kinds of innovations that were stifled before.

This chapter explores how Confidential Computing can literally save lives through better healthcare and how it can protect financial accounts and battle money laundering. We also cover the details on how Confidential Computing can make retailing more robust and customer-focused and how it can lead to more effective government.

Healthy Innovations

We're only just beginning to understand the potential power of artificial intelligence (AI) and machine learning (ML), and it's likely that some of humankind's most dramatic innovations will

happen in healthcare. Just consider some of the many possibilities for how AI can lead to better outcomes, faster diagnoses, and potentially lower medical costs:

- » Cancer is treated most effectively when it's identified early, and thanks to the miracle of ML, computers are getting better and better at spotting tiny tumors in medical imaging.
- » Rare maladies may be caused by certain genetic conditions that computers can track down much faster than humans can.
- » Computers can sift through millions of medical records and help doctors discern what treatment will be most effective for a patient's specific situation, based on many other patients' experiences.

Advances in this area require access to high-quality, diverse data sets that are representative of global patient populations. It often requires access to data from individual electronic medical records, some of the most sensitive and highly protected data there is. Healthcare organizations have both legal and ethical requirements to treat this data carefully, prevent inappropriate access, and avoid any breaches of privacy.



REMEMBER

Confidential Computing can step in and help facilitate what promises to be remarkable medical progress. The technology can secure this highly sensitive data across its entire life cycle. Researchers can access real-world health data to develop and deploy algorithms and analytics without any risk of compromising the data or causing privacy violations.

The COVID-19 pandemic provided some timely examples of how Confidential Computing can make a difference in healthcare. Contact tracing of people exposed to the coronavirus, for example, tested the limits of accessing and using personal data without divulging it in unintended places.

Another pandemic collaboration had multiple entities sharing a common AI model for medical imaging diagnosis of COVID-19. Multiple hospitals trained AI models for diagnosis, then sent their models to the cloud where they could be aggregated and improved — all without sharing personal health information.

The healthcare industry is super-regulated in everything from its finances to its infection controls to the way it handles its data.

With the use of Confidential Computing, it's possible to fully audit data confidentiality and deliver solid proof of data and code integrity.



REMEMBER

Indeed, there's a lot that can be done through the implementation of Confidential Computing in healthcare. There are wide-ranging possibilities, from drug discovery to genomics to supply chain operations to electronic medical records to federated learning. It's not a stretch to suggest that lives can be improved, or even saved, through the application of Confidential Computing in healthcare.

Facilitating Fintech

Was it all that long ago that financial services were largely within the realm of traditional banking and lending institutions? It's hard to believe how much financial technology, or fintech, has changed all of that.

Now, thanks to ubiquitous digital technologies and the widespread flow of data, loans and insurance can be arranged on a smartphone app — even mortgages. Investments are far easier to make and access on the go.

Cryptocurrencies are changing a lot of people's way of financial thinking.



REMEMBER

Fintech is highly disruptive and heavily reliant on the secure exchange and processing of data. Its on-the-go nature is facilitated by the cloud. It's living within highly regulated business environments. And it's an attractive target for bad cyber actors, which means it's in dire need of cutting-edge data security, such as Confidential Computing.

The fintech world is particularly vulnerable because of the way it's bringing about modularization of financial services. Today, products come and go quickly, and even traditional banks need to get new digital products up and running quickly.

That often means bringing in a third-party fintech provider which has algorithms and credit scoring models already figured out. Banks can plug in the outsourced solution and quickly be in a new business. But it requires giving fintech providers access to sensitive customer data.

It's not just digital lending solutions, either. Banks are often turning to third-party providers for such tasks as origination, analytics, compliance, and risk management.



WARNING

Fintech solutions rely on data sharing, which creates lots of points of potential data interception. Fintech solutions typically revolve around apps, which can be prime targets for attack because they may be easier to bypass than a company's networks. Fintech succeeds because of data ubiquity, but everyone involved in that data sharing is on the hook for protecting the data.

Consider what must happen to make the world of decentralized finance a reality. It's based on blockchain records of data, digital assets, and currencies involving a network of participants. Blockchain infrastructure allows transactions to be recorded and validated without any need for a centralized third party.

Again, this all relies on interactions with sensitive, personally identifiable information. Application developers must ensure that data isn't stored on the immutable blockchain.



TIP

A smart answer to these challenges is to conduct fintech and decentralized finance activity within the secure enclaves of a Confidential Computing implementation. Confidential Computing can combine with blockchain technologies to allow users to provide attestation and verification services for privacy and security.

Users can execute contracts in Trusted Execution Environments (TEEs) instead of independently accessing and validating data and contracts for themselves. TEEs prove the reliability of transactions, so there's no need for participants to perform verification again.

Confidential Computing has already proven quite helpful in securing work involving cryptocurrency. Confidential Computing can combine with hardware security modules to ensure that access is limited only to the right people, and only for specific authorized tasks.

Banking on Data

Financial institutions increasingly may choose to tap into third-party fintech to maintain a competitive position in an increasingly digital and disruptive world. Meanwhile, they're also struggling

to overcome such perils as digital theft, fraud, and money laundering activities.

These kinds of threats are why banks are usually subject to what are referred to as anti-money-laundering (AML) regulations. One reason is that financial operators can detect when a transaction may be used to fund terrorism or other nefarious activities.

But a successful AML solution requires constant observation and analysis of sensitive data, often in conjunction with other financial institutions. This has tended to be a highly resource-intensive process, complicated by cross-institute and cross-border data flow.



REMEMBER

Confidential Computing provides an attractive solution to some of the complexities of financial crime detection. It can allow financial institutions to collaborate without exposing private data to each other. Use it in combination with federated ML, and it's possible to make predictions about who is engaged in money laundering with enhanced accuracy.

This concept allows financial institutions to work together, tapping into such information as bank account details while ensuring sensitive data remains safe within a TEE. The federated ML algorithm can make inferences about illicit activities based on transactions at any of the collaborating institutions, but partners don't have to share private data with one another.

It works because encrypted data is loaded into the TEE. The ML algorithm crunches data from all of the collaborating institutions, and returns inferences based on that aggregated data. By confidentially pooling their data, institutions end up with a more accurate assessment about the client, but the actual data on which that assessment is based remains confidential within the TEE.

Confidential Computing ensures that banks remain in full control of their data, even as it is in use, being processed in the cloud. Code that's running in the cloud is fully verifiable and well protected. And both data and code are completely opaque to the cloud platform outside of the TEE.

Revolutionizing Retail

Smart retailers have been trying to maximize customer data for years to individualize the customer experience and drive greater sales. Tracking of individual shopping habits has become increasingly commonplace ever since. This often leads to concerns around the privacy of consumers.

Confidential Computing offers a secure and private computing architecture to retail. What makes it increasingly powerful is the ability to pull in individualized data from many more places — not just in-store purchases but internet searches and purchases from other retailers.

Hyper-personalized content can find us when we're visiting retailers, whether online or in a brick-and-mortar location. Likewise, decentralized customer data platforms can help retailers easily build audiences for marketing campaigns, based on shared customer data while respecting customers' privacy.



TIP

Needless to say, with all of that data flitting about, security is of tremendous importance. And because data must be “in use” to make all this magic happen, it's vital that data be protected while it's being processed, as with Confidential Computing. A secure, hardware-based TEE is the place for performing computation in order to protect in-use data.

Retail would be nowhere without payment processing, and TEEs are already commonly used in the payment processing industry.

Payment data must be protected throughout the process. Cash registers, tablets and chip-reading hardware should keep payment processing code protected. It's important to protect the means of data entry and ensure data can't be read or tampered while it's entered. With that in mind, number pads can be isolated so that input is only readable within the TEE, thus keeping it out of reach of any malware or unauthorized access by third parties.

Serving Government

Governments make a lot of rules regarding the protection of data and the imperative of privacy (see Chapter 2 for more information). Those regulations dictate how entire industry sectors operate, mandating a wide range of technologies and sparking creation of innovations such as Confidential Computing.

Governments also hold massive amounts of data themselves. Data about individual citizens, for example, including financial information, demographic information, citizenship information, and countless other subjects. All such data requires significant protection and privacy.

National security data obviously needs to be well-protected, as well as wide-ranging regulatory data.

The U.S. federal government pours billions of dollars into cybersecurity every year. It's placing increased emphasis on zero-trust technologies. And it's more and more interested in figuring out how to protect not just data at rest, not just data in transit, but data in use.

Confidential Computing is an ideal approach for dealing with government security and compliance requirements for the protection of highly sensitive data. Ensuring that data can be decrypted only inside a TEE is an excellent solution for such activities as performing analytics on a database.



TIP

Governmental agencies are all about documentation and attestation, and Confidential Computing has an answer for that need, of course. The technology includes attestation that assures the data is safe in the enclave.

And while earlier generations of Confidential Computing technology had some limitations in enclave size, today's processors offer capacity for significant amounts of enclave memory. That's a good thing for the federal government in particular, for which most tasks must happen on a massive scale.

Ultimately, Confidential Computing transforms the way governmental agencies can approach security in the cloud. It's no longer just a matter of placing trust in the cloud service provider and hoping data will remain secure while in use. Confidential Computing provides that assurance.

It also offers options for some of the most sensitive needs of government, such as the clouds that federal agencies manage in classified environments. These are air-gapped clouds that are not directly linked to the internet, which means any Confidential Computing operation must take place without that connection.

Tools are available for making that happen. Government agencies are able to achieve Confidential Computing provisioning, updating, and attestation without internet connections.



REMEMBER

Confidential Computing also is opening new doors for effective government policymaking. Just like businesses, governments work best when they're basing decisions on the most current information possible.

As policymakers work to help their jurisdictions solve problems and achieve goals for the future, they have often been hamstrung by the quality of their data. Real-time, detailed information is increasingly necessary to keep up with change.

One problem with real-time information is that the fresher it is, the more sensitive it is. If it's information about individuals, it requires careful handling to maintain privacy. If it's industry information, government officials may have a hard time accessing the best data if they can't ensure it will remain confidential.



REMEMBER

Confidential Computing allows policymakers to collect data from multiple sources as they analyze trends. They can do so by assuring that data is unencrypted and analyzed only in a secure enclave, keeping actual data points out of the hands of competitors or malicious actors.

The ability to use current, otherwise-sensitive data to help drive policymaking would be a tremendous step forward for effective government.

- » Bringing Confidential Computing experts together
- » Expanding the market for Confidential Computing

Chapter 4

Understanding the Confidential Computing Ecosystem

Not everyone has heard of Confidential Computing yet, but the word is spreading fast. An ecosystem populated by major industry players is collaborating to make that happen and keep the technology innovations moving forward.

This chapter introduces the consortium of industry expertise that is helping Confidential Computing improve and continually expand, and shares forecasts of how the market for this technology will grow.

Introducing the CCC

Innovation can start as a single great idea, but it can take a village to bring that fledgling idea fully to life. Confidential Computing benefits from the support of lots of incredibly influential players across the world of technology, including various hardware vendors, cloud providers, and software developers.



REMEMBER

The many backers of Confidential Computing work together through a collaboration known as the *Confidential Computing Consortium* (CCC). It's a project community of the Linux Foundation, and its primary purpose is to accelerate adoption of Trusted Execution Environment (TEE) technologies in a world where greater data security is an imperative.



REMEMBER

Fortanix is a founding member of the Confidential Computing Consortium. The list of premier members is a veritable who's who of technology players, all of whom are committed to bringing Confidential Computing into broader use and acceptance. Examples, in alphabetical order, include

- » **Accenture:** A multinational professional services company with a specialty in IT and consulting.
- » **Ant Group:** An affiliate of the Chinese Alibaba Group, it owns China's largest digital payment platform.
- » **Arm:** A United Kingdom-based technology company that specializes in automotive processors.
- » **Google:** This California-based technology giant needs little introduction, other than to say it's so much more than its humble search engine beginnings, and its cloud services touch just about everyone in one way or another.
- » **Huawei:** A Chinese provider of information and communications technology with customers across the globe.
- » **Intel:** The world's biggest maker of semiconductor chips, Intel is based in California.
- » **Meta:** The tech giant formerly known as Facebook is much more than just a social media company.
- » **Microsoft:** Another big name that needs no introduction, its software and hardware products serve billions of people, and its cloud services enable all kinds of innovation.
- » **Red Hat:** This North Carolina company is now part of IBM and is known for innovative open-source software products.

The CCC also includes a host of other prominent technology players. They all work together to define and accelerate the adoption of Confidential Computing by establishing open-source software and offering developers tools for innovating the work of securing data in use.

The fact that so many major technology players are in on this effort is a clear sign of just how important it is to bring Confidential Computing into widespread use. As computing continues to spread across multiple environments, companies need better control over their workload data and sensitive IP — the CCC is working to make that happen.

The CCC is helping developers add secure enclave technology to their applications that could perhaps be used in the following ways:

- » Building machine learning models geared toward secure multi-party data sets
- » Enabling confidential query processing in database engines within secure enclaves
- » Helping to protect sensitive data in Internet of Things edge devices
- » This is involved in such areas as patient information, billing and warranty activity, and machine learning model execution

The organization is also educating the industry and developers about Confidential Computing and best practices for its implementation. Ongoing research is looking into attestation and interoperability of TEEs, and experts are working to ensure these technologies aren't co-opted for nefarious purposes.

Growing the Market

Confidential Computing is stepping into the data security spotlight in a big way, according to experts. It's a solution to a serious problem that has business and IT leaders alike up at night fretting — the security of their sensitive data in a world with ever-increasing threats.

According to research that Everest Group conducted on behalf of the CCC, the total addressable market for Confidential Computing was roughly \$2 billion in 2021, and it's likely to build at a compound annual growth as high as 95 percent through 2026, if the most optimistic growth scenarios prevail. That's what Everest calls "hyper growth," spurred by cyber risks, regulations, and avenues for incremental revenue position.

Some two-thirds of that market is likely to involve the Confidential Computing software segment. But the hardware segment will see a lot of action, too, and is likely to double every year through 2026. Confidential Computing, says Everest Group, has the potential to become a standard for end-to-end security.

Regulated industries made up more than three-quarters of the demand in 2021, Everest Group reports. That includes such sectors as banking, finance, insurance, healthcare, life sciences, the public sector, and defense. And awareness and the willingness to invest are only going up from what the data from 2021 show.

The biggest drivers of growth are expected to be multi-party computing and security-specific use cases. Most of the market can be found in North America, the Asia-Pacific countries outside of China, and Europe. Specific hot spots for growth are likely to correlate with places where stricter privacy laws are enacted or being considered.

- » Taking that first step
- » Finding resources for your journey

Chapter 5

Choosing a Solution

Progress is a great thing. But progress can be difficult and daunting, sometimes so much so that those first steps cause great anxiety. When it comes to Confidential Computing, try to relax! As this chapter reveals, there's no need to tackle everything at once. Your progress can be made in easy-to-digest bites. What's more, there's no need to proceed alone. This chapter helps you discern how to choose the resources for getting started and be sure you're choosing a path that's not unnecessarily complicated.

Boiling the Ocean

The wisdom has been shared for generations. It may have had its roots in a Chinese proverb, and elders of many cultures have passed it along ever since. It's been printed on motivational posters and calendars, and more recently, it's shown up as social media memes in various incarnations.

"The journey of a thousand miles begins with one step," the ancient wisdom goes. "There is only one way to eat an elephant: a bite at a time," according to another variation. And some in the business world of process improvement will tell you, "Don't boil the ocean."



REMEMBER

Turns out that's not bad advice when it comes to the adoption of Confidential Computing. Organizations that have settled on the need for Confidential Computing must then move toward choosing a solution and beginning implementation. But it's easy to get bogged down by thinking too big at first.



TIP

An ideal approach is to first focus on the most immediate pain points and the highest priorities. Or perhaps get a start with some newer, relatively minor projects that are not already protected, then evolve into broader needs.

It works well in this case because of the nature of Confidential Computing. For contrast, think for a moment about the imperative of protecting your organization from phishing schemes, and how important it has been to educate every single employee about being wary while opening every single email, because it takes only one wrong step to subject an entire system to ransomware. That's a pretty big ocean, and you pretty much have to boil the whole thing to protect against that tiny but dangerous drop.

By comparison, when adopting Confidential Computing you can step down the path with a single application if you so choose. You can get a start with a specific analysis involving one data set in one secure enclave, then add to the work as you gain some experience.

One reason that works is the fact that the right Confidential Computing implementation can succeed with the apps you have, with no need for substantial rewrites or adoption of new ones. That makes incremental implementation so much easier than, say, going live with a new layer of security that must launch across the organization to be effective.



REMEMBER

In short, you don't have to worry right now about how you're going to get everyone onboard this plane right now. To help discern whether it's time to begin this journey, you can focus your thoughts on answering a few questions:

- » Are you dealing with securing data that's dispersed across various locations, such as in the cloud, on-premises, maybe within a Software-as-a-Service (SaaS) environment?
- » In your use case, do you run into regulatory challenges related to keeping that data secure, meeting privacy compliance requirements, generating proof of compliance?

- » Are your data security and compliance challenges getting in the way of using that data the way you need to, in order to achieve your objectives? Are you encountering roadblocks while trying to share or pool data in order to really unleash its potential?

If these are familiar frustrations, now's the time to explore Confidential Computing. No need to be intimidated! You'll be glad you took that first step. Or that first bite of the elephant.

Connecting With Experts

As long as we're quoting sources of wisdom, it's worth considering one more thought: "Don't complicate what's simple." There may be a source for this idea among Chinese proverbs, too, but Stephen King wrote these particular words. Fitting perhaps, given that he is a master of horrifying fiction, and you're trying to prevent a horrifying reality through IT security.

Here's the point: This book has put forth the idea that Confidential Computing is an elegant solution to a daunting problem, and it can be implemented without many of the hassles one might expect. But in order to really cash in on the simplest path, you need the right solution and expertise.



WARNING

With Confidential Computing, it actually *is* possible to complicate what should be a relatively simple, scalable, and easy-to-manage solution. That's why you must choose your path carefully. As you weigh one Confidential Computing solution against another, here are a few questions to ask:

- » **Managing your secure enclaves:** How much manual work will be on your team's plate as you adopt Confidential Computing? Will you need to integrate multiple infrastructure components and engage in a lot of manual configuration? Or is the product you're considering able to manage the entire Confidential Computing environment and enclave life cycle through a turnkey solution?
- » **Using your applications:** To get your applications up and running in a secure enclave, will your developers need to rewrite them using open-source software development kits (SDKs)? Or will the solution you're considering enable your

existing applications to run within the enclave, alongside whatever enclave-native applications or prepackaged applications suit your fancy? Will you incur extra development and integration costs, or will there be a simpler path?

- » **Engaging the security team:** Will your teams need to establish the various critical services required to securely implement Confidential Computing, for authentication, encryption, attestation, and audit? Or will your solution deliver cryptographically enforced policy and auditing, managing security policies including identity verification, data access control, and attestation? Will the solution you're looking at let you implement geo-fencing? Will it provide audit logs to verify compliance requirements?
- » **Gauging performance:** For the solution you're considering, will it scale up nicely or are there memory or sample size constraints? Can you be certain of excellent performance?
- » **Ensuring support:** How proactive and available is the support for the solution you're considering? Will you be able to maintain and continually update the platform as needed to ensure you're keeping up with the latest insights? Are subject-matter experts available to provide guidance as needed? Is 24/7 support available?



REMEMBER

What's great about Confidential Computing is that there are many powerful players working collaboratively to make it succeed. But as with any other purchase or implementation, your organization must evaluate your options carefully to find the solution that works best for you, achieves your goals, makes your life easier, and lets you sleep at night.

- » Bringing your teams onboard
- » Looking into the future

Chapter 6

Sharing Confidential Computing with the Team

If you've been reading these chapters in order and have made it to this point, you may well be sold on Confidential Computing as an idea whose time has come for your organization. If so, you still may need to sell others on your team. This chapter sets forth some thoughts for the conversations you can expect to have as you try to gain buy-in for adopting Confidential Computing technologies. It looks ahead to where things may be moving — in case you want to add a few more enticing thoughts to your meetings.

Gaining Buy-In

IT security experts have an extraordinarily important job, but they may not always be the most popular people in the organization. All the warnings of threats and terrible things that could and probably will happen if everyone isn't super-careful — not exactly happy conversation starters.

The data and applications teams in particular don't always like to hear from the security team. The warnings of threats may be accompanied by protective measures that are, at the least, inconvenient. Who wants to hear that their innovative idea, their exciting light-bulb moment, may run into hurdles from data security measures?



REMEMBER

Confidential Computing, on the other hand, brings welcome news. The data security professionals don't have to say "no" or "hold on a minute." They're now walking into planning meetings with "yes" answers.

"Access any data you need to access," they may say. "Bring in the data sets you already have and take advantage of standards already in use. You're not going to run into clunky interfaces, either." The bottom line is getting buy-in for Confidential Computing is aided by the simple fact that you're providing new ways for your data and applications teams to do the things they want to do.

For example, consider that they may have a situation where the most appealing option is to deploy a particular application in the public cloud. But the task involves sensitive data, and in the past, it has seemed rather risky to run a workload on someone else's system, where the level of trust is less than ideal.

With Confidential Computing, you're able to make that happen safely. You can run a private workload with sensitive customer data or corporate data and run it on a system for which you're not the administrator.

So now, imagine you've delivered that piece of positive news, and you're still getting some nervous glances from the people who need to buy into this concept. Here's the next key point to share: Confidential Computing makes this magic happen because you're performing computations in a hardware-based Trusted Execution Environment (TEE).

You're not just protecting data with software, which some may worry about compromise involving malicious actors. You're protecting data using hardware, so that data is decrypted within a designated CPU boundary. The memory controllers involved use embedded hardware keys that the cloud provider doesn't have. Your applications are running in an isolated environment that's both hardware-protected and attestable. Secure, easy to use, and not requiring changes to any existing data pipeline, Confidential Computing, has something for everyone.



That's an extraordinary safeguard. That means a bad actor with kernel access can't touch your data. Supervisor access, admin access, hypervisor access — it matters not because those people can't look into the application that's running within the Confidential Computing enclave. And you can attest to the fact that your data is safely in use where it's supposed to be.

So, the data is incredibly well protected, not just at rest, not just in transit, but while it's actually in use, being processed by the application. Confidential Computing is secure, easy to use, and doesn't require changes to any existing data pipeline — that message of safety and simplicity is a great start for buy-in.

But that's only part of the story. Imagine now that you get some questions from skeptics in the room.

"How well does Confidential Computing perform?"

"How difficult is it to deploy?"

"Is it going to be a hassle to use?"

"Will we run into issues with large workloads?"

"Are we going to find out that our datatypes aren't supported, or our applications?"

These questions are all good and understandable. We're all accustomed to things that sound too good to be true, and then turn out to have a lot of strings and complications attached.



But there's a positive answer to each of those questions. The performance levels for Confidential Computing are downright excellent. There are no sample size constraints to worry about, and all datatypes and applications are supported. Existing applications are fine, including prepackaged applications.

Confidential Computing's deployment and usability are just like any cloud-native app and aren't going to cause frustrations. Fortanix, for example, offers a turnkey solution that manages the whole Confidential Computing environment and enclave life cycle. It oversees the infrastructure across multiple public clouds as well as on-premises environments. What could be more straightforward?

Suffice it to say, Confidential Computing is enough of a game-changer that gaining buy-in won't be as challenging as it sometimes can be. As a cutting-edge concept that solves one huge

problem without creating a bunch of headaches, the “what’s in it for me?” question has some positive answers.

Moving Forward

Confidential Computing is poised to see explosive growth in the coming years, across a variety of industry sectors. We discuss that in more detail in Chapter 3. Different sectors have their own reasons and advancements they hope to improve, according to Everest Group research. For example:

- » **Banking, financial services, insurance:** Players in this sector are driven by the need for collaborative analytics to help detect fraud and money laundering. They’ll be searching for new ways to achieve secure decentralized finance access, and privacy in blockchain transactions.
- » **Healthcare and life sciences:** There’s a world of discovery to be made through research and analytics involving patient data, which must be secure. Drug discovery and treatment modeling are driving the need for confidential collaboration, and the increasingly ubiquitous IoT devices in health care must remain secure.
- » **The public sector:** There are new ideas for cross-agency collaboration that will drive Confidential Computing, new concepts in citizen data analytics, innovations in predictive maintenance for military equipment, and ongoing need for security for publicly exposed assets.
- » **Retail and consumer packaged goods:** Tracking and provenance are increasingly employing distributed ledger technology, and Confidential Computing can allow more powerful buyer and supplier analytics. New business models involving multi-party computing will drive increased adoption, too.
- » **High-technology sectors:** Confidential Computing will be needed for key management, federated learning, multiparty analytics, intellectual property and client data protection, and security for IoT edge devices and data.
- » **Manufacturing:** Watch for more distributed ledger technology tracking and provenance to drive adoption of Confidential Computing, more sophisticated analytics, edge computing, and the need to crunch encrypted data for autonomous vehicles in the factory.

- » Tapping into the genome
- » Easily accessing vital data sets
- » Breaking down cloud barriers

Chapter 7

Three Examples of Deploying Confidential Computing

Confidential Computing isn't just a cool-sounding idea for the future. It's happening today in more and more places. This chapter shows you how three organizations made a Fortanix Confidential Computing solution work to boost innovation.

Advancing Genomics

Incredibly powerful applications for data exist in the world of healthcare (we cover this more in Chapter 3). A great example is the field of genomics, which focuses on studying and mapping genomes to gain new understandings about health conditions and their potential treatment.

The human genome, of course, is a huge data set of its own, with genetic discoveries waiting to be tapped into and processed. That's the work of the Translational Genomics Research Institute

(TGen), an Arizona-based institute seeking to develop diagnostics and targeted therapeutics to achieve better outcomes for such conditions as cancer and Alzheimer's disease.

TGen needs high-quality data to make its work happen. Accessing the quality data sets it requires is a significant challenge, in large part because of the stringent data privacy regulations in the healthcare sector. The U.S. Health Insurance Portability and Accountability Act (HIPAA) sets a high bar for data privacy, along with significant penalties for violating its requirements related to individually identifiable, personal health information (PHI). TGen also works with European data covered by the super-strict General Data Protection Regulation (GDPR).

Data encryption is an obvious need, and that need has for some time had decent solutions in place with regard to data at rest and in transit. But TGen needed encryption capabilities while processing or sequencing individual data.



REMEMBER

The Trusted Execution Environment (TEE) that's at the heart of Confidential Computing was what TGen needed. And Fortanix Confidential Computing protects complex artificial intelligence workloads. This protection is available on any enabled platform, whether it's in the cloud or on-premises.

The capabilities of Confidential Computing definitely have come in handy for the rigors of regulatory compliance. To keep European regulators happy, the solution used geo-fencing to enable policies based on location of data. It includes support for multiple data processes in various clouds and locations and centralized logs to provide a global audit trail of all data access.

TGen reports that the solution has been easy to use, without adding personnel or new technical skills. Intellectual property is well protected, regulatory requirements are met, and costs are reduced.

Connecting with the Best Data

The BeeKeeperAI platform also aims to accelerate artificial intelligence (AI) innovations in healthcare. It was created at the University of California San Francisco's Center for Digital Health Innovation, and like the work of TGen (see the preceding section), it deals with stringent requirements for data security that are part and parcel in the world of healthcare.

Experts believe machine learning (ML) and AI algorithms can improve patient outcomes by 30 or 40 percent, and potentially cut treatment costs in half. Success requires access to diverse sets of high-quality data about global patient populations. Gaining access to those data sets can take months to clear the data stewardship hurdles, and the costs are high, too.



REMEMBER

Confidential Computing, in this case provided with the help of Fortanix and partners Microsoft and Intel, has provided a groundbreaking solution. Simply put, it allows data to remain in the control of the data steward, removing the need for costly and time-consuming secure acquisition of data.

Rather than bringing the data to the algorithm, this concept sends the algorithm to the data via secure cloud, where it can compute against the data. The data steward is able to curate a specialized data set for the purpose, and the BeeKeeperAI node runs within the data steward's regulatory-compliant cloud environment. The data and algorithm are encrypted in that location, so there's no risk of exposure or data breach.

Ultimately, the Confidential Computing solution offers greater protection of and access to the real-world data sets needed to succeed in the work of nonbiased, generalizable healthcare AI. It provides full protection of the AI model's intellectual property (IP), because it's published in a secure enclave. And the time and cost of accessing data sets is reduced significantly, which means healthcare innovation happens faster.

Overcoming Cloud Hurdles

NEC is a multinational information technology company that provides IT systems and network solutions to clients in a wide range of industries. Governments rely on its expertise, as do medical institutions, power companies, manufacturers, retailers, financial services companies, and many other kinds of businesses.

Clients are eager to manage and analyze data in the most accessible and effective way possible but have not always been eager to upload private data to cloud computing environments. There has been a sense that it's not the easiest place for analyzing private data, and clients sometimes also feared the potential for breaches while an application is running in the cloud using that

data. Breaches, of course, create negative publicity and risk costly regulatory penalties.

As powerful as cloud computing is for such tasks, these concerns have in the past created a psychological barrier. NEC turned to Fortanix for a solution, and that answer turned out to be Confidential Computing, using Intel Software Guard Extensions and Fortanix-created Confidential Computing enclaves. The hope was to break through that psychological barrier and get customers using powerful cloud applications to analyze their data.



REMEMBER

NEC was looking for three key features:

- » **Management:** It was crucial to find a system designed specifically to manage Confidential Computing and the life cycles of enclaves.
- » **Application support:** The platform needed to support running existing applications in the enclave to avoid any requirement for rewriting applications or creating new ones specifically for the enclaves.
- » **Security:** All data needed to be encrypted while in memory, and the key needed to be secured in a hardware security module integrated in the system.

Fortanix Confidential Computing solution checked off these three important boxes, and clients are now achieving powerful things in highly secure ways.

The new approach has opened the door to collaborations with partners on new tech initiatives. A wide range of analytic use cases have cropped up in such sectors as retail, hospitality, and government, and more opportunities can be found across financial services, healthcare, manufacturing, and others.

DISCOVER MORE SUCCESS STORIES

The examples in this chapter have been sourced from fortanix.com. If you want to read more and discover other successful case studies, too, visit <https://fortanix.com/customers>.

A note from Anil Rao at Intel

I would encourage every security and compliance stakeholder to read this book and discover how Confidential Computing can digitally transform your organization. Confidential Computing has experienced significant growth as every industry works to ensure privacy and security for their data. This book really hits the target in educating a broader audience about the game changing nature of Confidential Computing.

Anil Rao, VP and GM, Systems Architecture and Engineering

Office of CTO, Intel

Secure your data across its entire life cycle

This book introduces you to the exciting new world of Confidential Computing. Enterprises need to ensure security across the entire data life cycle — at rest, in motion, and also in use. Confidential Computing provides intrinsic security to applications and data in use when, perhaps, it's the most vulnerable both from external and internal threats. Besides, you can now safely collaborate with your data and not have to worry about compromising security or privacy. Confidential Computing is a game changer.

Inside...

- Defining Confidential Computing
- Complying with privacy and regulations
- Putting Confidential Computing to use
- Understanding the ecosystem
- Choosing a solution
- Gaining internal buy-in
- Seeing deployment examples



Ambuj Kumar is the co-founder and CEO of Fortanix. He is a prolific inventor and cryptography expert with a master's degree from Stanford University. **Anand Kashyap** is the co-founder and CTO of Fortanix. He is an expert in networking and security with a PhD from Stony Brook University.

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

ISBN: 978-1-119-89663-0

Not For Resale

for
dummies
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

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