Redbooks

ibm.com/redbooks

Maximizing Security with LinuxONE

Lydia Parziale

Gayathri Gopalakrishnan

Divya Konoor

Abhiram Kulkarni

Youssef Largou

Raydo Matthee

Security

IBM LinuxONE

IBM

Redpaper

IBM Redbooks

Maximizing Security with LinuxONE

October 2024

REDP-5535-02

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

Third Edition (October 2024)

This edition applies to LinuxONE, 2024.

This document was created or updated on January 3, 2025.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Cognos®
Db2®
DS8000®
FICON®
IBM®
IBM Blockchain®
IBM Cloud®
IBM Cloud Pak®
IBM Security®
IBM Spectrum®
IBM Z®
RACF®
Redbooks®
Redbooks (logo)  ®
Think®
WebSphere®
z/VM®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

LinuxONE is a hardware system that is designed to support and use the Linux operating system based on its distinctive architecture. LinuxONE can be used within a private and multi-cloud environment and caters to diverse workloads and requirements.

Security is deeply integrated into both the hardware and software of LinuxONE.

This IBM® Redpaper publication provides a broad understanding of how to use the various security features that make the most of and complement the LinuxONE hardware security features, including the following examples:

- Hardware accelerated encryption of data, which is delivered with near-zero overhead by the on-chip Central Processor Assist for Cryptographic Function (CPACF) and a dedicated Crypto Express adapter.
- Virtualization and industry-leading isolation capabilities with PR/SM, EAL 5+ LPARs, DPM, KVM, and IBM z/VM®.
- Other technologies that use LinuxONE security capabilities and practical use cases for these technologies.

Additionally, this IBM Redpaper publication gives a high level overview of Quantum-safe computing and confidential AI, which includes a discussion on the IBM z16 Telum chip.

This publication was written for IT executives, architects, specialists, security administrators, and others who consider security for LinuxONE.

Authors

This paper was produced by a team of specialists from around the world working at the IBM Redbooks Center, Poughkeepsie.

Lydia Parziale is a Project Leader for the IBM Redbooks® team in Poughkeepsie, New York, US, with domestic and international experience in technology management, including software development, project leadership, and strategic planning. Her areas of expertise include business development and database management technologies. Lydia is a certified PMP and an IBM Certified IT Specialist with an MBA in Technology Management and has been employed by IBM for over 25 years in various technology areas.

Gayathri Gopalakrishnan is an IT Architect for IBM India and has over 23 years of experience as a technical solution specialist, working primarily in consulting. She is a results-driven IT Architect with extensive experience in spearheading the management, design, development, implementation, and testing of solutions. A recognized leader, applying high-impact technical solutions to major business objectives with capabilities transcending boundaries. Adept at working with management to prioritize activities and achieve defined project objectives, ability to translate business requirements into technical solutions.

Divya Konoor is a Senior Technical Staff Member in IBM Infrastructure division. She is currently the Chief Architect for zSystems Hyper Protect Services offerings and has over 17 years in Security, Cloud / Virtualization Technologies and Systems Management domain. Currently she drives multiple offerings across Hyper Protect Services, On-Prem and in IBM Cloud. Previously, she was an architect for different public and private cloud offerings

across IBM zSystems and Power Systems. She has also been a Project Team Lead for one of the Opensource (OpenStack) community projects. She is a passionate speaker and product evangelist with multiple blogs, technical articles, patents and IP publications, amongst which include WeQuity Spotlight, IBM AoT, OpenStack Summit, ACM K, CDAC Cloud Auditing, IEEE CCEM speaker and technical papers, IEEE Wintechon technical committee member, GHC speaker and reviewer, OpenStack India Days Speaker, and PowerVM OpenStack Community Project Lead.

Abhiran Kulkarni is the Software Architect for IBM Hyper Protect at IBM India Systems Development Lab, Bangalore. He received his Bachelor of Engineering degree at PESIT Engineering College, Bangalore. He has over 15 years of experience in Software development and joined IBM in 2013. In his recent roles, he has worked in various development projects across IBM Z. He has worked on projects that are related to Secure Service Container, Secure Execution catering toward the clients that need zero trust architecture and Confidential Computing in hybrid cloud environment. He is author or co-author of 3 patents.

Victor Recio is the Software Architect for IBM Hyper Protect at IBM India Systems Development Lab, Bangalore. He received his Bachelor of Engineering degree at PESIT Engineering College, Bangalore. He has over 15 years of experience in Software development and joined IBM in 2013. In his recent roles, he has worked in various development projects across IBM Z. He has worked on projects that are related to Secure Service Container, Secure Execution catering toward the clients that need zero trust architecture and Confidential Computing in hybrid cloud environment. He is author or co-author of 3 patents.

Youssef Largouis the founding director of PowerM, a platinum IBM Business Partner in Morocco. He has 22 years of experience in systems, HPC, middleware, and hybrid cloud, including IBM Power, IBM Storage, IBM Spectrum®, IBM WebSphere®, IBM Db2®, IBM Cognos®, IBM WebSphere Portal, IBM MQ, ESB, IBM Cloud® Pak, SAP HANA and Red Hat OpenShift. He has worked within numerous industries with many technologies. Youssef is an IBM Champion 2020, 2021,2022, 2023 and 2024, an IBM Redbooks Platinum Author and has designed many reference architectures. His company has been recognized as an IBM Beacon Award Finalist in Storage, Software-Defined Storage, and LinuxONE five times. He is a regular speaker at IBM Think®, IBM TechXchange and Common Europe Congress. He holds an engineer degree in Computer Science from the Ecole Nationale Supérieure des Mines de Rabat and Executive MBA from EMLyon.

Raydo Matthee is a South African entrepreneur and technology leader with over 15 years of experience in IT. He is currently a Solutions Architect and course developer at Skunkworks (Pty) Ltd, specializing in security, integration, and applications. Raydo is an expert in cloud computing, API development, and IT security, with a focus on IBM products. He has authored numerous courses on topics such as OWASP, Okta, HashiCorp, and Asterisk, demonstrating his commitment to advancing IT education and consultancy.

Thanks to the following people for their contributions to this project:

Robert Haimowitz
IBM Redbooks, Poughkeepsie Center

Reinhard Bündgen, STSM
IBM Boeblingen

Oliver Bodemer
IBM Amsterdam

Thanks to the authors of:

- *Maximizing Security with LinuxONE*, REDP-5535, published in August, 2020:

  Lydia Parziale, Leticia Alexander, Yongkook Kim, Rushir Patel and Narjisse Zaki

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience with leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies last from 2 to 6 weeks, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Introduction

At IBM, we believe that your data is yours alone. The insights and advantages that come from your data are yours to use in the pursuit of your business objectives. IBM is dedicated to this mission, and our LinuxONE platform was designed around this core belief.

The world is experiencing a time of exponential growth in the sheer volume of data, fueled by the digital transformation of systems, services, and interconnected devices that all require strong data serving capabilities. Businesses must manage, store, and most importantly, protect this information while they use it to gain competitive advantage.

IBM LinuxONE is an all-Linux enterprise platform for open innovation that combines the best of Linux and open technology with the best of enterprise computing in *one* system.

In this IBM Redbooks publication, we explore the features and capabilities of the LinuxONE platform that extend enterprise-grade security to the Open Source world.

This chapter includes the following topics:

1.1 Introduction to LinuxONE

Linux adoption grew dramatically over recent years, expanding from initial use by startups for web servers into its use today for a vast range of enterprise computing workloads. It grew up alongside the open-source community, which is a unique resource that uses a network of passionate and dedicated developers who are willing to contribute to various projects.

IBM LinuxONE is an all-Linux enterprise platform for open innovation that combines the best of Linux and open technology with the best of enterprise computing in ONE system. It is designed to support customers who want an efficient and cost-effective solution for protecting their data and hosting various enterprise-grade Linux workloads. The LinuxONE platform is defined by security, speed, scalability, reliability, and openness.

The hardened Linux-based software stack of LinuxONE can run most open-source software packages, such as databases and data management (that is, MariaDB, PostgreSQL, MongoDB, Apache Kafka, Websphere Liberty, and Apache Spark), virtualization platforms, and containers (IBM z/VM, KVM, Docker and Podman), automation and orchestration software (Red Hat OpenShift, Kubernetes, OpenStack, Puppet, Node.js, Juju, and Chef), and compute-intensive workloads, such as blockchain.

1.2 Enterprise Security Challenges

The IBM LinuxONE platform provides unique capabilities to help with overcoming security challenges and differentiating your business offerings. These challenges can range from maintaining regulatory compliance so that you are not subject to fines and penalties, to ensuring that your critical systems are not compromised or taken over by malicious entities.

When you are considering an infrastructure platform, you must understand the security features that are inherent on the platform in the cloud and on-premises. IBM LinuxONE is engineering from the ground up to protect your business from all manner of cyber threats. By providing a highly securable, massively scalable, data serving platform, LinuxONE can help any business that wants to thrive in a data-centric economy.

1.2.1 Data protection and privacy

According to the IBM Cost of a Data Breach Report for 2023, several key findings emerge:

- The average total cost of a data breach hit a record high in 2023, reaching USD 4.45 million. This reflects a 2.3% increase from 2022, which stood at USD 4.35 million. Over the long term, there has been a notable 15.3% increase from the 2020 report, where the average cost was USD 3.86 million.
- Around 51% of organizations are considering increasing their security investments following a breach. Despite the escalating costs associated with data breaches, there's a near even split among respondents regarding plans to boost security spending. The primary areas earmarked for additional investment include incident response (IR) planning and testing, employee training, and deploying threat detection and response technologies.
- The impact of robust security AI and automation on mitigating the financial fallout of a breach is significant. Investments in these technologies have proven effective in reducing costs and minimizing the time required to detect and contain breaches. Organizations that extensively leverage these capabilities in their strategies experienced an average 108-day reduction in the time taken to identify and contain breaches.

Additionally, they reported USD 1.76 million lower data breach costs compared to counterparts that did not utilize security AI and automation capabilities.

The report suggests several recommendations to mitigate the cost of a data breach:

- Integrate security measures throughout all phases of software development and deployment processes, conducting regular testing.
- Upgrade data protection strategies to align with the demands of hybrid cloud environments.
- Employ security AI and automation to enhance efficiency and precision in threat detection and response.
- Enhance resilience by thoroughly understanding your organization's attack surface and regularly practicing incident response protocols.

In Figure 1-1, notable shifts were observed in the top countries or regions with the highest average cost of a data breach from 2022

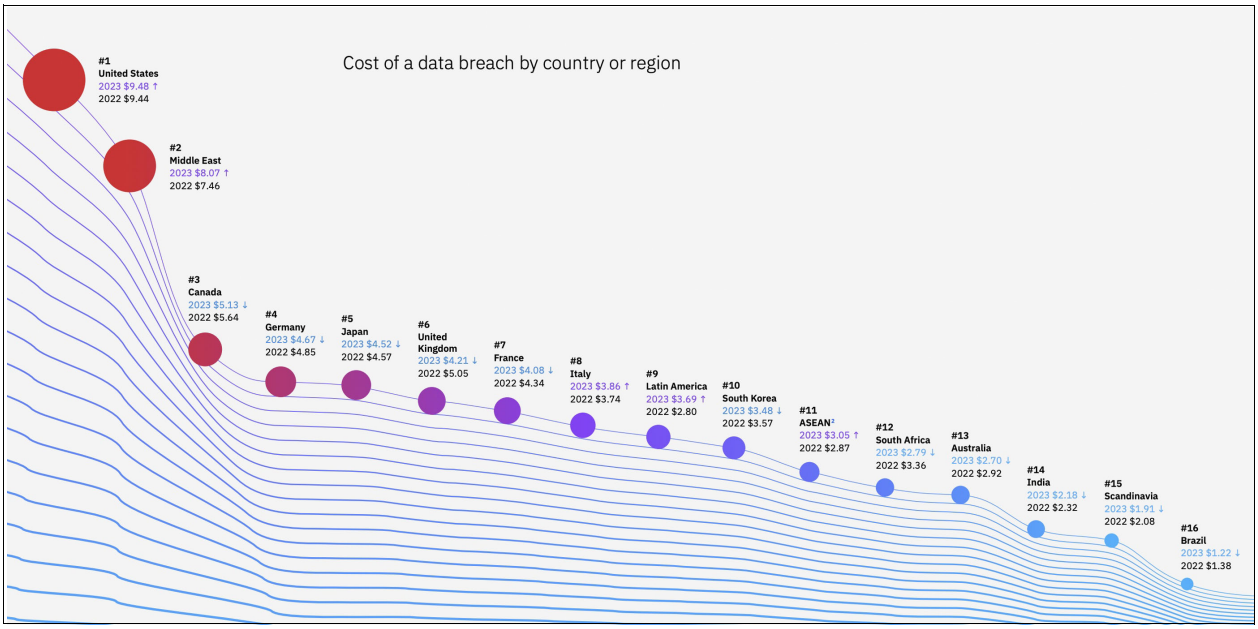

Figure 1-1 Top countries with the highest average cost of a date breach

Protecting data involves measures aimed at preventing unauthorized access or use of information. Encryption stands out as the most efficient method to shield sensitive data from unauthorized access or misuse, consequently lowering breach costs. Surprisingly, despite the prevalence of data breaches over the years, only a small fraction of breached records was found to be encrypted.

This lack of encryption is because until recently, encrypting data was time-consuming, expensive, and it severely degraded system performance. Businesses that chose to encrypt did it selectively, leaving the rest of their data exposed to threats.

Four levels of data encryption deployment are available. Figure 1-2 on page 4 shows that, as you move higher up in the pyramid, you gain more security control of your data at the cost of a more complex and intrusive encryption implementation. Conversely, as you go lower, complexity and costs are reduced, but with a less granular approach to encryption. Deploying data encryption on one layer or another is a tradeoff that depends on the context and regulatory environment of each client.

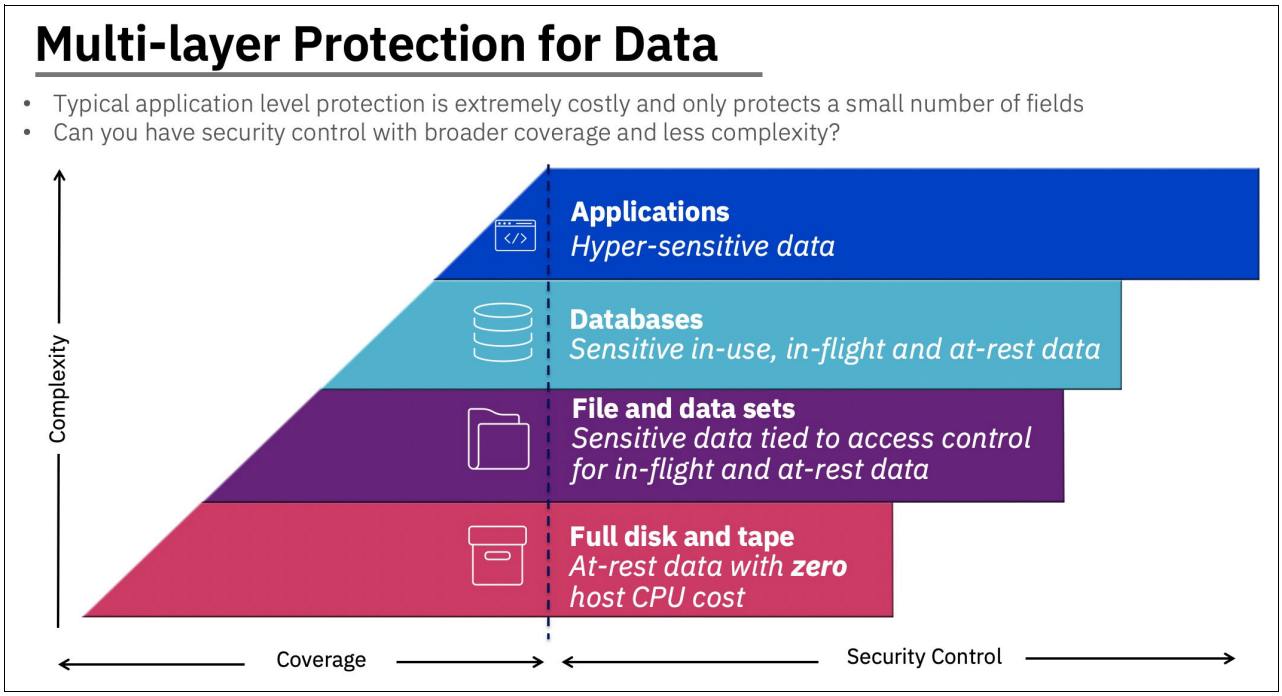

Figure 1-2 Various levels of data protection

With the latest generations of LinuxONE hardware, IBM embedded encryption logic and processing onto each processor chip in the system. This configuration allows you to encrypt massive amounts of data with little effect on your system performance.

The on-chip encryption capability is further enhanced by the IBM Crypto-Express adapter, which enables industry-leading key protection technology. This combination of integrated cryptographic hardware makes the LinuxONE platform an efficient and cost-effective solution for securely hosting various enterprise-grade Linux workloads.

For more information about these cryptographic hardware capabilities, see 2.1, “Secure cryptographic hardware” on page 18.

1.2.2 Cyber resiliency and availability

In today’s digital economy, being continuously open for business is a competitive advantage. Your customers expect to transact business with you 24 x 7 with no excuses for interruptions or outages. To this end, cyber resiliency is an often overlooked method for providing business value through security. You might lose access to your core systems because of a ransom ware attack or face unplanned downtime as a result of a distributed denial of service (DDoS) attack. These types of outages can cause costly business disruptions and productivity losses.

To be competitive, enterprises must provide trusted services with high uptime to their clients, while consistently delivering new value and features. This demand from clients requires a computing platform that accommodates your developers’ creative genius and a highly secure infrastructure that provides instantaneous data delivery at any time, whether you have thousands or millions of simultaneous users.

IBM LinuxONE machines have the industry’s highest reliability levels for over a decade, with up to 99.999999% or greater availability. In fact, the experts who track downtime say that the underlying hardware infrastructure is “in a class of its own.”

Engineered to help protect against insider and outsider threats in multi-tenanted cloud environments, the LinuxONE III generation introduced a new capability called the Secure Execution for Linux. It is a hardware-based security technology that is designed to protect workloads from internal and external threats to help our clients prevent security breaches. IBM Secure Execution can help protect and isolate workloads on-premises, or on IBM LinuxONE and IBM Z® hybrid cloud environments. Users, and even system administrators, cannot access sensitive data in Linux-based virtual environments.

For more information, see, “IBM RACF® for z/VM provides security systems that include access control and auditing functionality as the backbone for Linux security.” on page 33.

The LinuxONE platform also features technology to address the next evolution in cyber attacks. The potential of quantum computing is quickly advancing, and soon will explode. This shift will force the entire industry to evolve as quantum computing might break currently secure cryptographic algorithms.

The new IBM Crypto-Express 8S introduces support for quantum-safe signing algorithms to ensure that data can be secured today and well into the age of pragmatic quantum computing. The current generation of quantum-safe cryptographic algorithms were developed internally by IBM to help prevent the eventual quantum computing attacks of the future.

IBM LinuxONE 4 is the industry’s first quantum-safe enterprise Linux system, which integrates new hardware encryption capabilities, that allows users to apply quantum-safe encryption to protect data workloads and infrastructure. Each core includes a dedicated coprocessor for cryptographic functions, which is known as the Central Processor Assist for Cryptographic Functions (CPACF). CPACF supports pervasive encryption and is providing hardware acceleration for encryption operations.

The compression capabilities have been improved, delivering greater compression throughput than previous generation systems. This on-chip compression co-processor uses industry standard compression algorithms and can reduce data storage requirements and costs. This compression can also increase data transfer rates to boost throughput above comparable x86 CPUs; all without adversely impacting response times.

The new Telum chip is the first server-class chip with a dedicated on-chip AI accelerator that provides IBM LinuxONE with the capacity to execute real-time inferences at speed and scale by co-locating data and AI.

For more information, see IBM LinuxONE.

1.2.3 Industry and regulatory compliance

Many businesses operate on the assumption that adhering to regulatory compliance standards is enough to mitigate business risk and protect a company’s data. Although this assumption might have been true in the past, it is no longer an option to implement a static compliance policy only.

Cyber threats in the modern era are constantly evolving and move too quickly for an organization to sit back with passive data protection policies. Protecting only enough data to achieve compliance should be viewed as the bare minimum, not a best practice.

Corporate risk management and compliance is an ongoing cost and effort that only increases over time. The scope of these industry and government regulations is constantly in flux. It also is expanding with more compliance requirements, and the introduction of new mandates, such as General Data Protection Regulation (GDPR).

At a high level, global security regulations feature varying requirements. However, a common thread is that most include specific requirements regarding encryption of data and access to that data. LinuxONE addresses both of these needs as a core feature of the platform.

A comparison of the how regulatory compliance is handled by LinuxONE and x86 is shown in Table 1-1.

Table 1-1 Data compliance on LinuxONE versus x86

IBM LinuxONE capabilities	x86 capabilities
► Encrypt everything quickly and economically with hardware acceleration. ► Protect against side channel the top insider threats with tamper resistant encrypted keys and confidential computing with IBM Secure execution. ► Isolate LPARs at the architectural level with EAL5+ designed and HSM crypto isolation. ► Protect keys in memory with tamper resistant design (FIPS 140-2 Level 4 HSM) and encryption. ► Hardware-accelerated SSL/TLS encryption. ► Limit access to data by encrypting everything ► Remove entire groups of users from audit scope.	► Slow/costly encryption due to lack of hardware acceleration. ► Vulnerable attacks due to clear keys. ► Weak isolation. ► Selectively encrypt sensitive data at best due to slow/costly encryption. ► Little protection against insider threats. ► Large audit scope.

LinuxONE provides secure hosting environments where an organization that runs on-premises or cloud-based services can ensure that their user's data is protected always. Even IT administrators with physical access to hardware cannot access data, including sensitive and personally identifiable information.

1.3 IBM LinuxONE servers

IBM LinuxONE 4 is a family of enterprise servers with advanced features designed with on-chip AI inferencing and industry-first, quantum-safe technologies.
The latest generation, LinuxONE 4 LA1 and LinuxONE LA2 and AGL4 were introduced in 2022 and 2023. The most recent addition, LinuxONE 4 Express, was unveiled in February 2024:

- IBM LinuxONE 4 LA1

  The newest member of the LinuxONE family, the IBM LinuxONE 4 was generally available in late 2022 and maintains the new form factor introduced in LinuxONE III, featuring a 19-inch frame that flexibly scales 1 - 4 frames. It is designed around the new Telum processor with 8 core per chip with Dual Chip Module packaging, 5.2 Ghz processor and is configurable with up to 200 processor cores, up to 40 TB of RAM, and 10 TB of Redundant Array of Independent Memory (RAIM) per central processing drawer.

- IBM LinuxONE 4 LA2

  Released in May 2023, the IBM LinuxONE 4 LT2 is the newest entry model into the IBM LinuxONE family of servers. It delivers a 19-inch single-frame (versus the option of up to four frames for the LT1) with an efficient design with a low entry cost that can easily coexist with other platforms in a cloud data center. This model is designed around the new Telum processor with 8 core per chip with Dual Chip Module packaging.

It is configurable with up to 68 cores running at 4.6 GHz, up to 16 TB of RAM, and 8 TB of RAIM per central processing drawer.

- IBM LinuxONE 4 AGL

  Released in May 2023, the IBM LinuxONE 4 AGL is a new entry model option which consists of a rack mount model. It delivers a rack mount option from 10U to 39U that can be collocated with other technologies with a client-supplied 19 inch rack. This model is designed around the new Telum processor with 8 core per chip with Dual Chip Module packaging. It is configurable with up to 68 cores running at 4.6 GHz, up to 16 TB of RAM, and 8 TB of RAM per central processing drawer.

- IBM LinuxONE 4 Express

  It is a special offering addition that is based entirely on IBM LinuxONE 4 AGL hardware but preconfigured and bundled specifically to three sizes: small, medium and large. The small size consists of 4 Cores and 384GB of RAM. The medium size consists of 6 cores and 512GB of RAM. The large size consists of 12 Cores and 736GB of RAM.

For more information on the LinuxONE 4 family, see IBM LinuxONE 4.

IBM LinuxONE 4 LA1 data sheet

The data sheet for LinuxONE 4 LA1 model is shown in Table 1-2. The complete data sheet can be found at IBM LinuxONE Emperor 4.

Table 1-2 LinuxONE 4 LA1 at a glance

IBM LinuxONE 4 LA1 features		
LinuxONE 4 Models	**Cores**	**Memory: Min - Max**
LA1	Up to 200	512 GB - 40 TB
Cryptography		
Crypto-Express8S (2-port adapters)	Maximum: 30 adapters, 60HSMs	
Crypto-Express8S (1-port adapter)	Maximum: 16 adapters: 16HSMs	
Connectivity		
IBM FICON® Express32S SX/LX	Maximum: 384 ports	
IBM Adapter for NVMe1.1	1 Slot	
OSA Express 7S 1.2 (1000BT/1G)	Maximum: 96 ports	
OSA Express 7S 1.2 (10G/25G)	Maximum: 48 ports	
Inter-LPAR Connectivity		
HiperSockets	Up to 32 high-speed virtual local area networks.	
Shared Memory Communications - Direct Memory Access (SMC-D)	Up to 32 ISM virtual CHIPDs.	
Supported distributors		

IBM LinuxONE 4 LA1 features	
Linux	Canonical, Red Hat and SUSE with their latest supported releases and versions; for the certified levels see IBM tested platforms.
Supported hypervisors	
IBM z/VM	z/VM V7.2, z/VM 7.3, z/VM 7.4.
KVM	KVM hypervisor, which is offered with the Linux distribution.
IBM partitioning technology	Up to 85 LPARs for secure workload isolation.

LinuxONE 4 LA2 data sheet

The data sheet for LinuxONE 4 LA2 model is shown in Table 1-3.

Table 1-3 LinuxONE 4 LA2 at a glance

IBM LinuxONE 4 LA2 features		
LinuxONE 4 Models	**Cores**	**Memory: Min - Max**
LA2	Up to 68	64GB - 16TB
Cryptography		
Crypto-Express8S (2-port adapters)	Maximum: 20 adapters, 40HSMs	
Crypto-Express8S (1-port adapter)	Maximum: 16 adapters: 16HSMs	
Connectivity		
IBM FICON Express32S SX/LX	Maximum: 96 ports	
IBM Adapter for NVMe1.1	1 Slot	
OSA Express 7S 1.2 (1000BT/1G)	Maximum: 96 ports	
OSA Express 7S 1.2 (10G/25G)	Maximum: 48 ports	
Inter-LPAR Connectivity		
HiperSockets	Up to 32 high-speed virtual local area networks.	
Shared Memory Communications - Direct Memory Access (SMC-D)	Up to 32 ISM virtual CHIPDs.	
Supported distributors		
Linux	Canonical, Red Hat and SUSE with their latest supported releases and versions; for the certified levels, see IBM tested platforms.	
Supported hypervisors		
IBM z/VM	z/VM V7.2, z/VM 7.3, z/VM 7.4.	
KVM	KVM hypervisor, which is offered with the Linux distribution.	
IBM partitioning technology	Up to 40 LPARs for secure workload isolation.	

IBM LinuxONE 4 AGL data sheet

The data sheet for IBM LinuxONE 4 AGL model is shown in Table 1-4.

Table 1-4 LinuxONE 4 AGL at a glance

IBM LinuxONE 4 AGL features		
LinuxONE 4 Models	**Cores**	**Memory: Min - Max**
AGL	Up to 68	64GB - 16TB
Cryptography		
Crypto-Express8S (2-port adapters)	Maximum: 20 adapters, 40HSMs	
Crypto-Express8S (1-port adapter)	Maximum: 16 adapters: 16HSMs	
Connectivity		
IBM FICON Express32S SX/LX	Maximum: 96 ports	
IBM Adapter for NVMe1.1	1 Slot	
OSA Express 7S 1.2 (1000BT/1G)	Maximum: 96 ports	
OSA Express 7S 1.2 (10G/25G)	Maximum: 48 ports	
Inter-LPAR Connectivity		
HiperSockets	Up to 32 high-speed virtual local area networks.	
Shared Memory Communications - Direct Memory Access (SMC-D)	Up to 32 ISM virtual CHIPDs.	
Supported distributors		
Linux	Canonical, Red Hat and SUSE with their latest supported releases and versions; for the certified levels, see IBM tested platforms.	
Supported hypervisors		
IBM z/VM	z/VM V7.2, z/VM 7.3, z/VM 7.4	
KVM	KVM hypervisor, which is offered with the Linux distribution.	
IBM partitioning technology	Up to 40 LPARs for secure workload isolation.	

IBM LinuxONE 4 Express data sheet

The data sheet for IBM LinuxONE 4 Express model is shown in Table 1-5.

Table 1-5 LinuxONE 4 Express at a glance

IBM LinuxONE 4 Express features		
LinuxONE 4 Models	**Cores**	**Memory: Min - Max**
Express	Up to 16	384GB - 864GB
Connectivity		

IBM LinuxONE 4 Express features		
IBM FICON Express32S SX	Maximum: 8 ports	
OSA Express 7S 10GbE SR 1.2	Maximum: 4 ports	
Inter-LPAR Connectivity		
HiperSockets	Up to 32 high-speed virtual local area networks.	
Shared Memory Communications - Direct Memory Access (SMC-D)	Up to 32 ISM virtual CHIPDs.	
Supported distributors		
Linux	Canonical, Red Hat and SUSE with their latest supported releases and versions; for the certified levels, see IBM tested platforms.	
Supported hypervisors		
IBM z/VM	z/VM V7.2, z/VM 7.3, z/VM 7.4	
KVM	KVM hypervisor, which is offered with the Linux distribution.	

1.4 Reasons to choose IBM LinuxONE

IBM LinuxONE delivers the best of enterprise Linux on the industry's most reliable and highly scalable hardware. These systems are specialized scale-up enterprise servers that are designed exclusively to run Linux applications.

IBM LinuxONE provides the highest levels of availability (near 100 percent uptime with no single point of failure), performance, throughput, and security. End-to-end security is built in with isolation at each level in the stack, and provides the highest level of certified security in the industry.

IBM Secure Execution for Linux is a continuation and expansion of well-known security features of IBM Z and LinuxONE. It supplements pervasive encryption, which protects data at-rest and data in-flight, to also protect data in-use. IBM Secure Execution for Linux is a hardware-based security technology that provides a trusted execution environment (TEE) for "Confidential Computing". It provides scalable isolation for individual workloads to protect them from not only external attacks, but also insider threats. For more information, see IBM Hyper Protect Platform: Applying Data Protection and Confidentiality in a Hybrid Cloud Environment.

The CryptoExpress cards (CEX) available on IBM LinuxONE 4 provide access to a FIPS 140-2 Level 4 Hardware Security Module (HSM) which is exploited by Hyper Protect Services to provide encryption key management services and quantum safe features for signing and encapsulation with Dilithium and Kyber.

Additionally, LinuxONE Systems facilitate transparent use of redundant processor execution steps and integrity checking, which is necessary in financial services industries. LinuxONE servers typically enable hot-swapping of hardware, such as processors and memory. This swapping is typically transparent to the operating system, enabling routine repairs to be performed without shutting down the system.

IBM LinuxONE delivers on the promise of a flexible, secure, and smart IT architecture that can be managed seamlessly to meet the requirements of today's fast-changing business climate.

1.4.1 Best of Enterprise Linux and open-source

LinuxONE provides the following benefits:

- Premium Linux experience with subsecond user response times and virtually unlimited scale.
- Broad portfolio of open-source and other vendor products and tools delivered on the platform.
- Choice of Linux (RHEL, SUSE Enterprise Linux, and Ubuntu) and tools that best fit your environment.
- Eliminates risks by running Linux on industry's most secure and resilient hardware platform.
- Easy integration of data and applications with existing IBM z Systems® based solutions.
- Overall increases the operational IT efficiency.

1.4.2 Hardware strengths

IBM LinuxONE provides these hardware strengths:

- Reliability:
  - Redundant processors, I/O, and memory.
  - Error correction and detection.
  - Remote Support Facility.
- Availability:
  - Fault tolerance.
  - Automated failure detection.
  - Non-disruptive hardware and software changes.
  - IBM LinuxONE machines have the industry's highest reliability levels for over a decade, with up to 99.99999% or greater availability.
- Virtualization:
  - High-performance logical partitioning by using IBM Processor Resource/Systems Manager (IBM PR/SM)[1].
  - Up to 85 (LinuxONE 4 LA1) or 40 (LinuxONE 4 LA2, AGL) logical partitions (LPAR) with independent virtual resources.

- PR/SM is one of the most secure systems available, having achieved Common Criteria Evaluation Assurance Level 5+ (EAL5+) for partition isolation. This is one of the highest levels of certification that can be achieved by commercially available hardware.

  Note: For more information about Common Criteria, Evaluation Assurance Levels, Protection Profiles, and a list of certified products, see The Common Criteria.

- IBM Dynamic Partition Manager provides facilities to define and run virtualized computing systems by using a firmware managed environment that coordinates the physical system resources shared by the partitions. The partitions resources include processors, memory, network, storage, crypto, and accelerators.
- Both the industry-leading virtualization hypervisor z/VM and the open-source hypervisor kernel-based virtual machine (KVM) are supported on all IBM LinuxONE models.
- PR/SM, z/VM, and KVM employ hardware and firmware innovations that make virtualization part of the basic fabric of the IBM LinuxONE platform.
- IBM HiperSockets allows up to 32 virtual LANs, thus allowing memory-to-memory TCP/IP communication between partitions.

► Scalability:

- IBM LinuxONE 4 Model LA1 scales to 200 physical processors and up to 40 TB of memory.
- IBM LinuxONE III Models LA2 and AGL scale to 68 physical processors and up to 16 TB of memory.
- LinuxONE 4 Model LA1, can process 1 trillion web transactions per day and can support thousands of virtual servers or up to two million containers on a single system.

► Security:

- The pervasive encryption capabilities of LinuxONE allow you to encrypt massive amounts of data with little effect on your system performance. The LinuxONE hardware benefits from encryption logic and processing on each processor chip in the system.
- The Central Processor Assist for Cryptographic Function (CPACF) is well-suited for encrypting large amounts of data in real time because of its proximity to the processor unit. CPACF supports:

  - DES
  - TDES
  - AES-128
  - AES-256
  - SHA-1
  - SHA-2
  - SHA-3
  - SHAKE
  - DRNG
  - TRNG
  - PRNG

  With the LinuxONE 4, CPACF supports Elliptic Curve Cryptography clear key, improving the performance of Elliptic Curve algorithms.

[1] PR/SM is standard component of all IBM LinuxONE models, which enables LPARs to share system resources. PR/SM divides physical system resources, both dedicated and shared, into isolated LPARs. Each partition is like an independent system running its own operating environment. It is possible to add and delete resources like processors, I/O, and memory across partitions while they are actively in use.

The following algorithms are supported:

- EdDSA (Ed448 and Ed25519)
- ECDSA (P-256, P-384, and P-521)
- ECDH (P-256, P-384, P521, X25519, and X448)

Protected key signature creation is also supported.

– Optional cryptography accelerators provide improved performance for specialized functions:

- Can be configured as a secure key coprocessor or for Secure Sockets Layer (SSL) acceleration.
- Certified at FIPS 140-3 and Common Criteria EAL 4+.

– IBM's Hyper Protect Virtual Server offering is exclusive to IBM LinuxONE because it delivers more security capabilities to protect Linux workloads from internal and external threats throughout their lifecycle, build, management, and deployment phases. Some of the security benefits include:

- Building images with integrity, which secures continuous integration and delivery.
- Managing infrastructure with least privilege access to applications and data.
- Deploying images with trusted provenance.

– The IBM LinuxONE 4 maintains the Secure Execution for Linux. It is a hardware-based security technology that is designed to protect and isolate workloads on-premises, or on IBM LinuxONE and IBM Z hybrid cloud environments. Users, and even system administrators, cannot access sensitive data in Linux-based virtual environments.

► Regulatory compliance

Most security regulations feature include specific requirements regarding encryption of data and access to that data. LinuxONE addresses those two aspects at the core of its platform through capabilities, such as EAL 5+ isolation between LPARs, pervasive encryption, and protection against side-channel attacks and insider threats with tamper resistant encrypted keys.

► Just-in-time deployment of resources:

– On/Off Capacity on Demand provides temporary processing capacity to meet short-term requirements or for testing new applications.

– Flex Capacity which allows to dynamically shift production capacity between different LinuxONE servers in multiple sites.

– Capacity Backup (CBU) allows you to replace model capacity or specialty engines to a backup server in the event of an unforeseen loss of server capacity because of an emergency. CBU ensures that customers can access additional capacity during a disaster recovery situation without having to purchase more capacity. Typically, this system allows customers to sign up for CBU on an IBM LinuxONE at another site and use this capacity for a number of contracted disaster recovery tests or for a contracted time during a declared disaster at the customer site.

► Power and cooling savings

With its low power and cooling requirements, IBM LinuxONE is an ideal platform for the consolidation of distributed servers.

1.5 Main Infrastructure Security Features on IBM LinuxONE

LinuxONE offers enhanced security features, including robust guest isolation mechanisms, ensuring that each virtualized environment remains securely separated.

The platform also provides end-to-end virtualization, enabling comprehensive control and management of virtual resources throughout the system. Additionally, cryptographic hardware support is integrated, offering hardware-accelerated encryption capabilities for enhanced data protection and integrity.

1.5.1 Hardware isolation

Hardware enabled protections move clients closer to realizing a Zero Trust security environment through workload isolation and hardened access restrictions over their data. LinuxONE isolates individual workloads to protect them from both external attacks and internal threats and utilizes a hardware-based trusted execution environment (TEE) without the need to create separate LPARs.

A trusted execution environment (TEE) is a secure area of a main processor. It guarantees code and data loaded inside to be protected with respect to confidentiality and integrity. A TEE as an isolated execution environment provides security features such as isolated execution, integrity of applications executing with the TEE, along with confidentiality of their assets. In general terms, the TEE offers an execution space that provides a higher level of security for trusted applications running on the device than a rich operating system (OS) and more functionality than a 'secure element' (SE).

1.5.2 Chain of trust

A hardware root of trust is the foundation on which all secure operations of a computing system depend. It contains the keys used for cryptographic functions and enables a secure boot process. It is inherently trusted, and therefore must be secure by design. The most secure implementation of a root of trust is in hardware making it immune from malware attacks.

IBM LinuxONE III and LinuxONE 4 have secure boot, which provides the possibility for an operating system to boot in a secure fashion. The chain of trust is verified from the hardware throughout the whole firmware up into the operating system.

1.5.3 End-to-End virtualization

z/VM has been around for over 40 years and represents world-class security and reliability. It hosts Linux guests and can also host more traditional workloads. A fundamental strength of z/VM is the ability for virtual machines to share system resources and is designed with very high levels of resource utilization. z/VM is designed for extreme scalability, security, and efficiency to create opportunities for cost savings, while providing a robust foundation for cognitive computing on the LinuxONE platforms.

1.5.4 Hardware cryptography

The Crypto Express (CEX8S) card is a coprocessor and can support a wider range of callable services that include secure key and clear key support for PKA decrypt, digital signature verification, digital signature generation, including RSA and ECC variants. Alternatively, the card can be configured as an accelerator (CEXCA). In this mode, the card supports only three clear key cryptographic APIs, associated with RSA public key encryption, decryption, and verification. When the cryptographic coprocessor is configured as an accelerator it provides better throughput at the expense of supporting fewer services.
Further information on the CEX8S card can be found in 2.1.3, “IBM Crypto Express8” on page 25.

1.5.5 SELinux

Security-Enhanced Linux (SELinux) defines access controls for the applications, processes, and files on a system. It uses security policies, which are a set of rules that tell SELinux what can or can't be accessed, to enforce the access allowed by a policy.

When an application or process, known as a subject, makes a request to access an object, like a file, SELinux checks with an access vector cache (AVC), where permissions are cached for subjects and objects.

If SELinux is unable to make a decision about access based on the cached permissions, it sends the request to the security server. The security server checks for the security context of the app or process and the file. Security context is applied from the SELinux policy database. Permission is then granted or denied.

Security-Enhanced Linux is a Linux kernel security module that provides a mechanism for supporting access control security policies, including mandatory access controls. SELinux is a set of kernel modifications and user-space tools that have been added to various Linux distributions.

SELinux provides a flexible Mandatory Access Control (MAC) system built into the Linux kernel. Under standard Linux Discretionary Access Control (DAC), an application or process running as a user (UID or SUID) has the user's permissions to objects such as files, sockets, and other processes. Running a MAC kernel protects the system from malicious or flawed applications that can damage or destroy the system.

SELinux defines the access and transition rights of every user, application, process, and file on the system. SELinux then governs the interactions of these entities using a security policy that specifies how strict or lenient a given Red Hat Enterprise Linux (RHEL) installation should be.

1.5.6 IBM Multi-Factor Authentication

Mainframe systems are the foundation of trusted digital experiences for most of the world's largest companies and organizations. However, passwords protecting critical users, data and applications are a relatively simple point of attack for hackers to exploit because passwords rely on user education and compliance for both implementation and control. Using a variety of methods such as social engineering and phishing, criminals have exploited employees, partners, and general users to hack into even the most secure platforms.

IBM Multi-Factor Authentication (IBM MFA) raises the level of assurance of your mission-critical systems with expanded authentication capabilities and options for a comprehensive, user-centered strategy to help mitigate the risk of compromised passwords and system hacks.

MFA enhances the security of z/VM sign-ons by incorporating a diverse range of factors. These encompass support for various authentication methods such as RSA SecurID®, Gemalto SafeNet, IBM Security Access Manager (ISAM), IBM Cloud Identity Verify (CIV) via RADIUS, Generic RADP, LDAP, Native timed one-time password (TOTP), and SmartCard usage. Many of these factors also include biometric support, adding an extra layer of identity verification. MFA operates on the principle of three key components: knowledge (what you know), possession (what you have), and inherence (what you are). Utilizing "what you are" or "what you have" alongside traditional passwords offers increased confidence in verifying the authenticity of individuals accessing critical systems.

2

Core security technologies on LinuxONE

A secure digital business starts at the hardware level. LinuxONE helps you to establish foundational system integrity across physical and virtual infrastructure to address rapidly evolving security threats to your enterprise.

Data security is typically achieved through complicated encryption keys that are difficult for hackers to crack. Unfortunately, along with all the positive potential that quantum computing brings, there is also the unintended use of quantum computing for cracking encryption keys. Estimates vary on when it will be possible for cybercriminals to use quantum computers to crack encryption keys, but the hackers are not waiting to steal data they can't decrypt with existing technology. They are harvesting data now to decrypt later when they have access to quantum computing.

Defending against this can be an even larger concern for organizations that have long lifecycles for data. The data that organizations must keep for several years could become easily decryptable by quantum computing before the data reaches the end of its lifecycle. Organizations that are going to replace their encryption infrastructure should consider the risk of quantum cracking before their next encryption upgrade.

LinuxONE is a secure platform. The LinuxONE Infrastructure Platform includes security features that are embedded in the platform for workloads in the cloud and on-premises.

In this chapter, we cover the core technologies that are embedded in the LinuxONE hardware for ensuring a pervasive level of data secure protection.

This chapter includes the following topics:

2.1 Secure cryptographic hardware

Traditionally, encryption of all your data requires a large amount of time and computation overhead. This need is the result of the limitations of software-based encryption, which shares computing resources with the rest of your system, and can slow down other shared applications.

Although it can be cheap to get started with software-based encryption, it quickly becomes prohibitively expensive as you scale your business. Especially with more secure forms of encryption (more complex algorithms, longer bit values), it becomes resource-intensive. Therefore, you typically must pay more for extra processing power for the same level of performance.

The LinuxONE solution to this problem is through dedicated hardware that is tuned for encryption and can encrypt 100% of data that is at-rest and in-flight (by default) with minimal compute overhead. The security is embedded in every feature of the hardware and software stack. This level of protection is achieved by using hardware accelerated encryption capabilities and does not require any code or application changes.

The LinuxONE solution provides improved trust and reduced risk through hardware encryption that is fast and strong enough to encrypt all data it manages, with designed-in redundancy to deliver a highly available single point of truth.

LinuxONE allows for secure platform simplification by providing the following benefits:

- Protecting all applications and fully encrypting their data without any changes to the business applications, including built-in hardware acceleration that allows for faster encryption.
- Enabling bulk encryption that is simple, transparent, and features optimized performance to secure your cloud infrastructure.

Types of cryptographic keys

An important starting point for understanding the cryptographic hardware on LinuxONE is the types of encryption keys that are used. LinuxONE cryptography uses the types of encryption keys that are listed in Table 2-1.

Table 2-1 Encryption key types for LinuxONE

Key type	Description
Data	A data-encrypting key that is used to encrypt and decrypt data.
Key-encrypting	A key that encrypts or wraps other keys.
Effective	A type of data-encrypting key; also called a *data key* that is wrapped by a key-encrypting key (KEK).
Master	A special KEK that is in a tamper-responding, Crypto Express adapter only and sits at the top of a KEK hierarchy. Loading and managing the master key can be done by using the Trusted Key Entry (TKE) workstation.
CPACF wrapping	A special key-encrypting key that is generated at LPAR activation and is in the Hardware System Area, which is inaccessible to applications and the operating systems. It is used to create protected keys.
Secure	Key values are encrypted under a Master Key and no key ever appears in decrypted form outside of the Crypto Express HSM. Crypto operations are performed only within the Crypto Express HSM.

Key type	Description
Clear or plain	A data-encrypting key that is not encrypted by any other key. The key material is in plain text.
Protected	Key values are encrypted under a CPACF wrapping key. Crypto operations are performed by using only CPACF. When the key is not in use, it is protected by the Crypto Express HSM. In the case of Linux in a native LPAR, the wrapping key is specific to the LPAR. However, for guests of z/VM or KVM, the wrapping key is specific to the guest.
Operational	A key that is not a master key or KEK, such as a data-encrypting key (which can be clear, secure, or protected).

The tradeoffs between clear key, protected key, and secure key encryption implementations are shown in Figure 2-1.

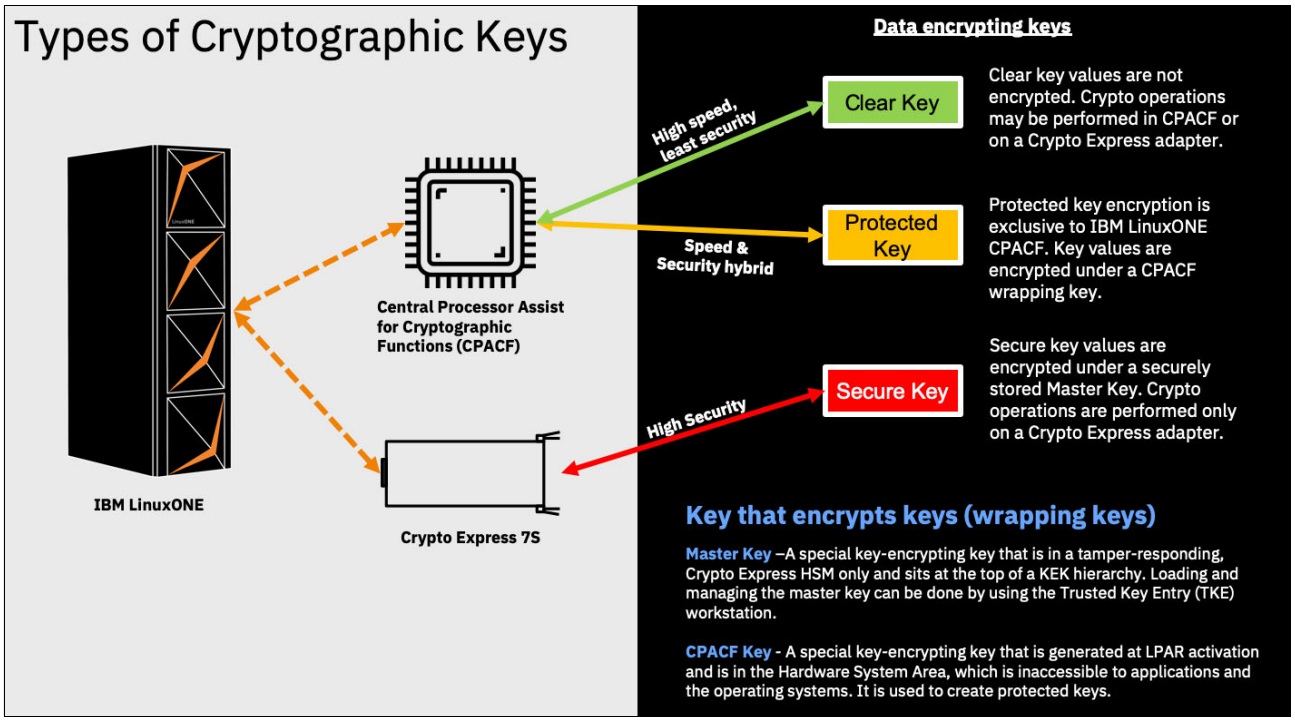

Figure 2-1 Encryption key implementations and tradeoff

Cryptographic key management

Cryptographic key management is a complex task that must be managed according to strict policies. You must account for various legal, regulatory, and compliance requirements. Your key management system should allow authorized persons a method for key identification, exchange, separation, update, backup, and management.

For more information about key management, see 3.4, “Cryptographic Key Management for LinuxONE” on page 42.

2.1.1 Central Processor Assist for Cryptographic Functions

The hardware accelerated encryption capabilities of LinuxONE are enabled by Central Processor Assist for Cryptographic Functions technology (CPACF).

The CPACF on-chip encryption co-processor is on every compute chip that is next to the main processor and can encrypt up to 13 GB of data per second per core. This configuration results in performance improvements of up to 6x and is best suited for symmetric, high-speed bulk encryption.

The Central Processor Assist for Cryptographic Function (CPACF) is well-suited for encrypting large amounts of data in real time because of its proximity to the processor unit. CPACF supports:

- DES
- TDES
- AES-128
- AES-256
- SHA-1
- SHA-2
- SHA-3
- SHAKE
- DRNG
- TRNG
- PRNG

With the LinuxONE 4, CPACF supports Elliptic Curve Cryptography clear key, improving the performance of Elliptic Curve algorithms.

The following algorithms are supported:

- EdDSA (Ed448 and Ed25519)
- ECDSA (P-256, P-384, and P-521)
- ECDH (P-256, P-384, P521, X25519, and X448)

CPACF is a no-charge feature of LinuxONE. If you have the hardware, you can enable the technology and start to get the benefits right away.

True Random Number Generator

Another feature of LinuxONE cryptographic hardware is the ability to generate unreproducible, unique data with the on-chip true random number generator (TRNG). This capability is the basis for generating high-quality cryptographic keys. TRNG is an improvement of Deterministic RNG because the numbers that are generated are more random.

Protected key encryption

The ability to use protected key encryption with CPACF is a differentiating feature on LinuxONE. Many encryption services use plaintext “clear keys” that are stored unsecured in main system memory. These clear keys are visible and vulnerable during the encryption and decryption process. Clear keys can be stolen from memory or system dumps, meaning that your encrypted data is now at risk.

Protected key technology uses CPACF for encrypting data at high speeds without exposing keys to main system memory. The data-encrypting (protected) keys are encrypted or “wrapped” by a special key that is stored in a secure environment. In that environment, it is inaccessible to applications, hypervisor, and operating system (known as the HSA, or Hardware System Area). Among the many use cases, this technology enables fast and highly secure encryption and decryption of complete disks (volumes) or selected partitions.

With the introduction of LinuxONE III, CPACF supports the creation of protected key signatures.

2.1.2 IBM Crypto Express adapter

LinuxONE can also include IBM Crypto Express adapters, a Hardware Security Module (HSM), and a cryptographic co-processor that supports high-speed, asymmetric encryption. This co-processor also supports a symmetric cryptographic function to assist key encryptions, encrypt data, compute message authentication codes, protect financial PINs, secure EMV card transactions, and many other functions. This specialized hardware performs AES, DES, TDES, RSA, ECC, SHA-1, SHA-2, SHA-3, and other cryptographic operations.

An HSM is designed to withstand physical and logical attacks and features special hardware to perform cryptographic operations and protect keys. The HSM is accessed from a host computer that uses a set of generic or specialized API functions. The IBM Crypto Express HSM can support the following primary secure key cryptographic modes, APIs, and one clear key accelerator mode. You also can reload your HSM firmware at any time to switch from one to the other:

- IBM Common Cryptographic Architecture (CCA)

  CCA provides a set of general-purpose cryptographic functions. Its primary strength is support of finance industry payments applications. This mode is also often called *cryptographic co-processor mode*, and many IBM software products support CCA modes to enhance security:

  - Provides comprehensive support for cryptographic operations in compliance with CCA 7.0 standards.
  - Suitable for environments requiring high security and compliance with FIPS 140-2 Level 4 and PCI-SSC HSM requirements.
  - Includes support for enhanced cryptographic functions, such as ASC X9 TR-34, X.509, SHA3, EdDSA, EdDH, and Dilithium.

- IBM Enterprise PKCS #11 (EP11) mode.

  This mode supports the industry standard PKCS #11 API. EP11 is designed for customers who seek support for open standards and enhanced security. It offers a various general purpose, secure-key-only cryptography functions:

  - Designed to support the PKCS #11 v2.4 standard, offering robust key management and cryptographic operations.
  - Ideal for applications needing secure and flexible cryptographic processing.
  - Provides Protected Key support and includes enhancements for improved performance and security.

- Accelerator

  This mode makes Crypto Express hardware become an asymmetric cryptographic function accelerator. This mode might be useful for enterprise web services:

  - Optimized for maximum cryptographic performance, offering more performance improvement when configured as an accelerator.
  - Suitable for environments with high-performance cryptographic needs where speed is critical.

Figure 2-2 on page 22 shows the Crypto Express 8S configuration options.

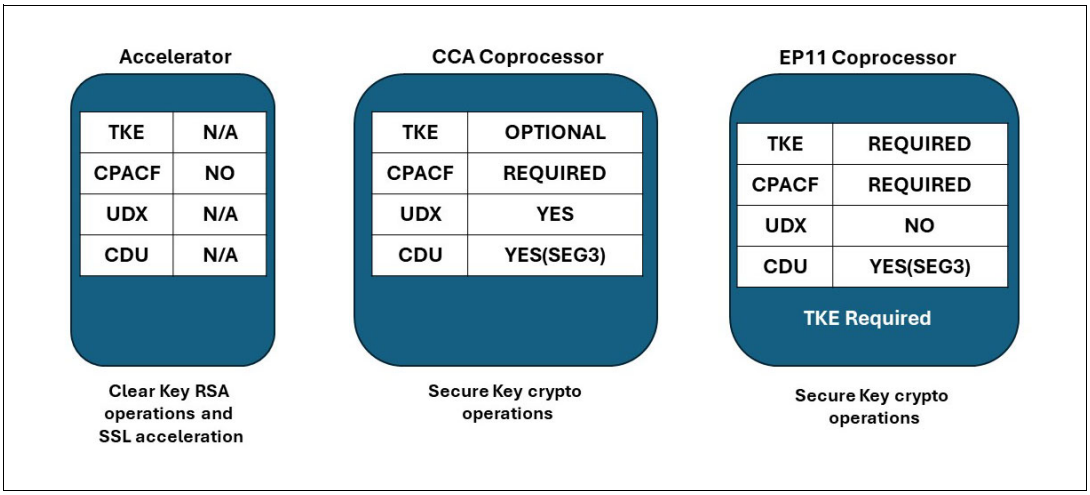

Figure 2-2 Crypto Express 8S configuration options

Note:

- Only one configuration mode can be active at a time for each HSM.
- Changing the configuration from one mode to another will erase all card secrets to ensure security and integrity.
- The only exception to this rule is when switching between CCA and Accelerator modes or vice versa. This switch does not result in the erasure of card secrets, allowing for more flexible operational adjustments without compromising stored cryptographic data.

HSM Certifications

The IBM Crypto Express adapter is a Hardware Security Module that is certified for FIPS 140-2 Level 4. This certification applies to all adapter modes. This means the cryptographic co-processors are protected within a tamper-resistant and tamper-responsive environment that erases encryption keys if it senses an attack, eliminating the risk of exposing cryptographic keys during a breach.

The IBM Crypto Express adapter is also certified for PCI-HSM, which applies to the Common Cryptographic Architecture (CCA) firmware / cryptographic coprocessor mode. Common criteria applies to the EP11 firmware coprocessor mode.

The following section describes how the Crypto Express adapter HSM can provide the maximum protection for your encryption keys.

FIPS certification levels

The Federal Information Processing Standard (FIPS) Publication 140-2 (FIPS PUB 140-2), is a US government computer security standard that is used to approve cryptographic modules. It defines what areas of security requirement to meet the standard in terms of design and behaviors of the module. These requirements include specification, ports and interfaces, roles and services, and authentication; finite state model; physical security; operational environment; cryptographic key management; electromagnetic interference/electromagnetic compatibility (EMI/EMC); self-tests; design assurance; and mitigation of other attacks.

For more information about the FIPS 140-2 standards, see the US Department of Commerce's federal information processing standards publication *Security Requirements for Cryptographic Modules*.

Figure 2-3 shows a summary of the following levels of FIPS 140-2 that LinuxONE is certified to protect against.

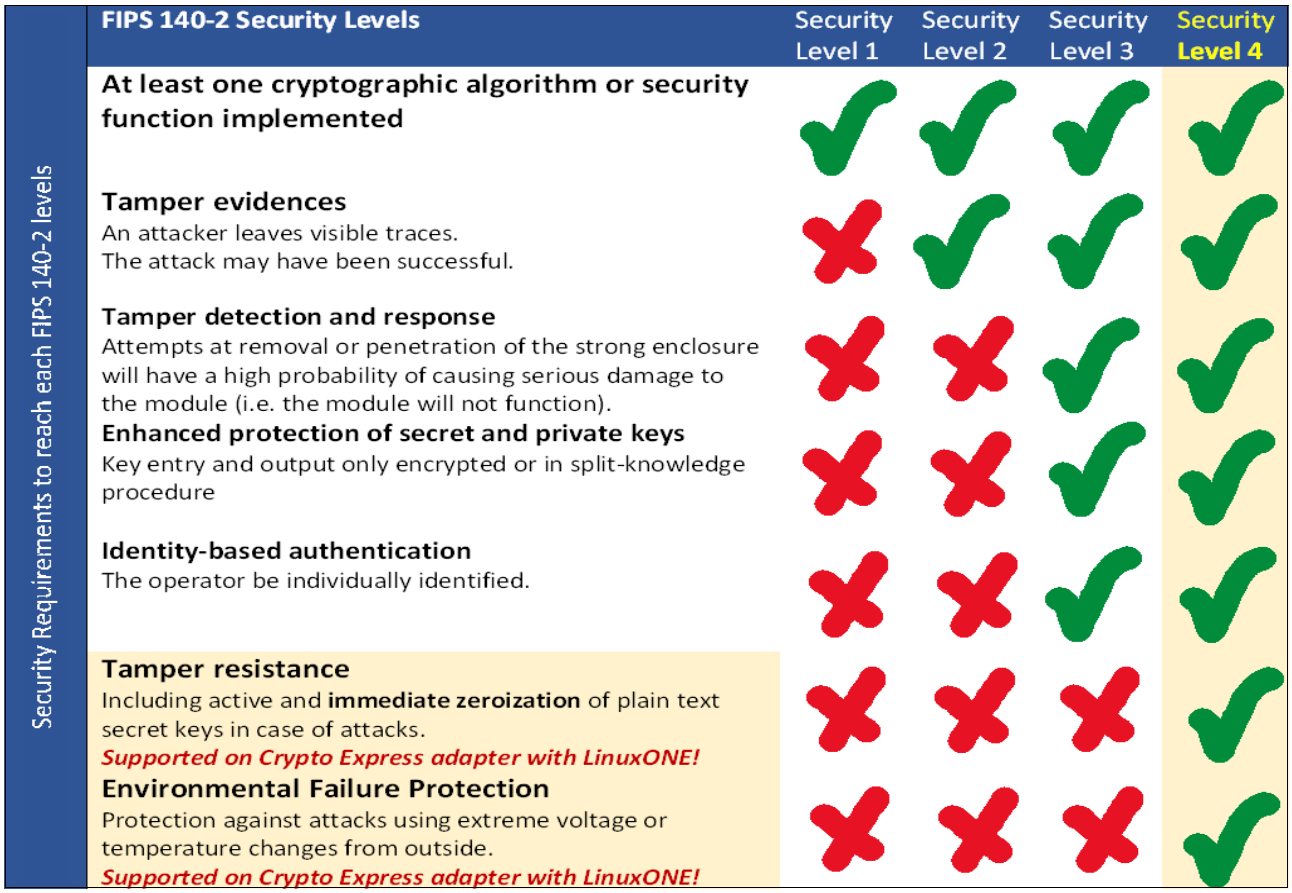

Figure 2-3 Different certification levels of the FIPS 140-2 Standard

Secure key encryption

The HSM feature of the IBM Crypto Express adapter enables secure encryption with the capability to protect cryptographic keys that use a special key called Master Key. (We use the terms *HSM* and *Crypto Express adapter* interchangeably in this document.)

The master key is used to encrypt the keys that are used by your applications, which are stored outside of the Crypto Express HSM. Those keys are fully protected because the Master Key is protected by the security features of the HSM.

The Master Key is stored in Crypto Express adapter hardware. Loading and managing the Master Key into the Crypto Express adapter can be done in software, but it is highly recommended to perform it with the Trusted Key Entry (TKE) workstation. For more information about TKE and key management, see “Cryptographic key management” on page 19 and 3.4, “Cryptographic Key Management for LinuxONE” on page 42.

Secure key encryption uses a wrapped key for encrypting data that never appears in clear text outside of a secure environment. The Cypto Express adapter ensures that these keys are never exposed in the clear. Any unauthorized attempts to access the Crypto Express adapter enclosure result in the deletion of the stored keys.

Figure 2-4 shows a high-level workflow of how master keys and secure keys are used in LinuxONE. Four different types of Master Keys for Crypto Express Adapter are available: AES, DES, RSA, and ECC. For each Master Key type, the adapter can set up to 85 domains, which is designed to support up to 85 LPARs in LinuxONE II and III system.

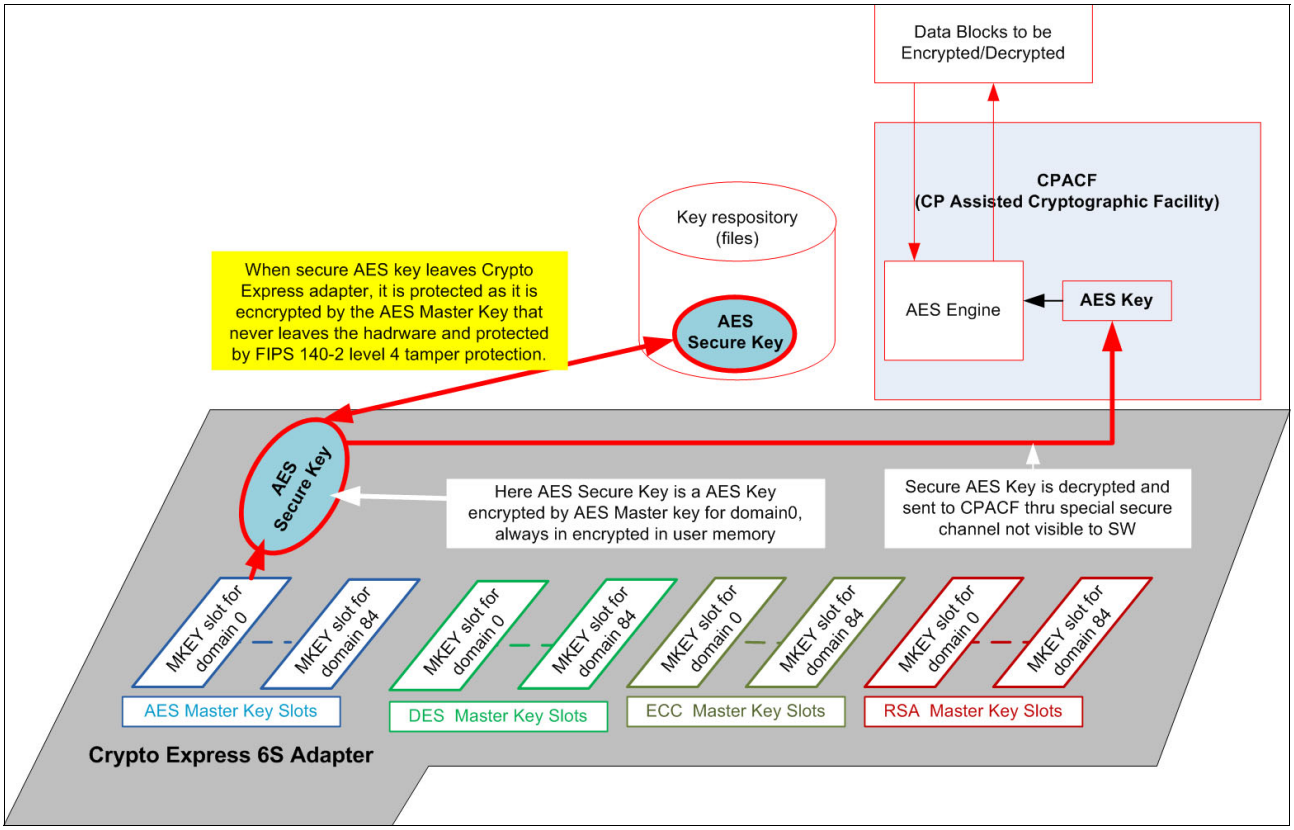

Figure 2-4 Master Keys and Secure Keys with Crypto Express adapter for LinuxONE

Creating a secure key

At the time of this writing, the Crypto Express for LinuxONE supports AES encryption for the user-defined secure key. The hardware is designed to support RSA, DES, AES, and ECC to be the secure keys. For more information about updates, see the IBM Knowledge Center.

Figure 2-5 on page 25 shows an example of a secure key use case, which is a high-level process that is used to create a secure key for an LUKS2 format volume to encrypt a disk volume by using a secure key and protected key.

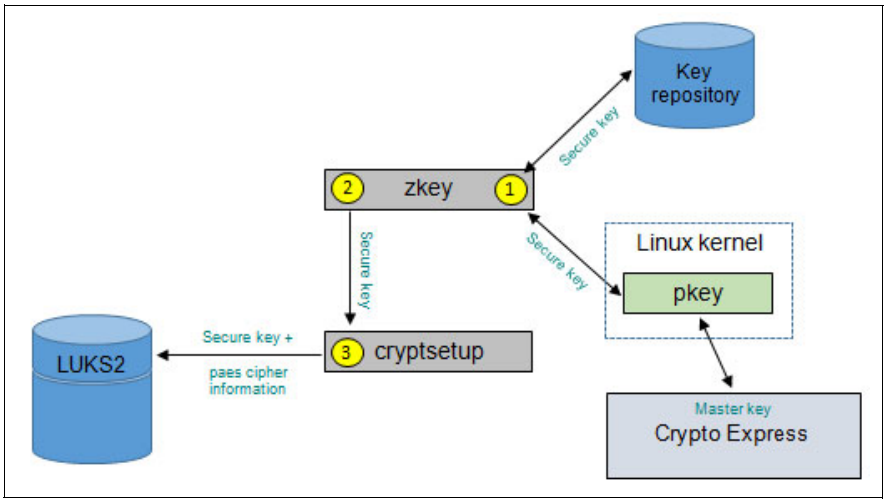

Figure 2-5 Creating a secure key

This process includes the following steps:

1. A secure key is created by using a **zkey** command. The **zkey** utility generates the secure key with the help of the **pkey** utility and an assigned Crypto Express adapter (with master key). The secure key is also stored in the key repository.
2. The use of the **zkey cryptsetup** command generates output strings that are copied and pasted to the **cryptsetup** command to create the encrypted volume with the appropriate secure key.
3. The **cryptsetup** utility formats the physical volume and writes the encrypted secure key.

An example of this process is shown in Figure 6-1 on page 65. For more information, see *Getting Started with Linux on Z Encryption for Data At-Rest*, SG24-8436.

2.1.3 IBM Crypto Express8

To help protect client's sensitive data, IBM introduced the Crypto Express8S card. It is a hardware security module (HSM) that allows applications to get access to the quantum- safe algorithms. IBM Crypto Express8 will be useful for application modernization and for building new applications. Crypto Express8 allows clients to leverage both quantum- safe and classical cryptography.

Cryptography hardware exists inside the machine. Every central processor (CP) has a small cryptographic engine that can be used with CP Assist for Crypto Functions (CPACF). There are a limited number of functions CPACF performs. However, they also work in conjunction with Crypto Express; in this case, in this generation of the 8S, those go in the IO drawer. They use a single or dual hardware security module.

Note: All models of LinuxONE 4 use the industry's first quantum-safe hardware security module (HSM), the Crypto Express8S, which can protect sensitive data with quantum-safe cryptography along with creating an inventory to help with migration and modernization.

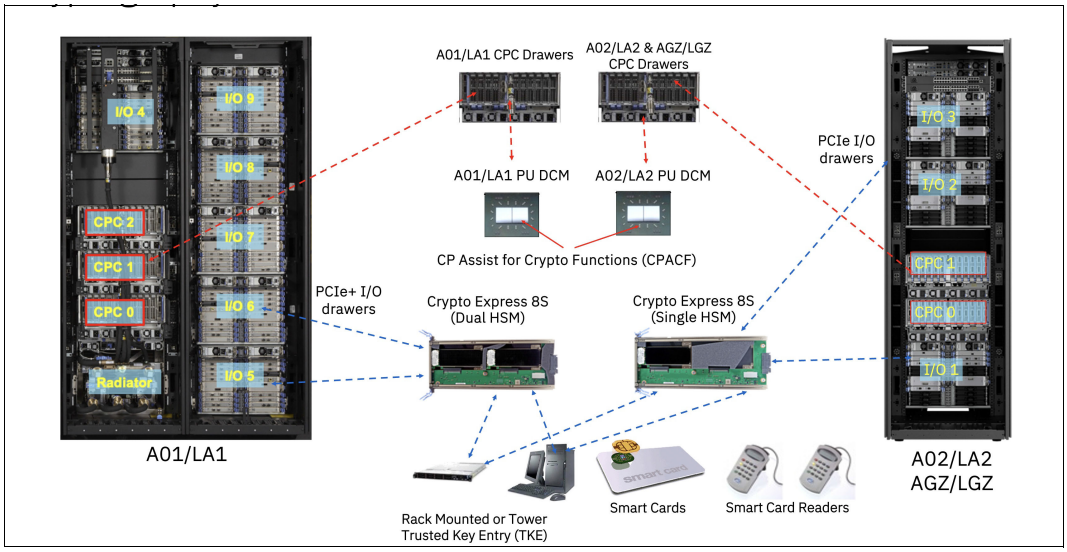

Figure 2-6 Cryptography hardware in IBM LinuxONE

With IBM LinuxONE 4, clients have the capabilities necessary to leverage the quantum-safe algorithms in their applications. IBM provides the Crypto Express8S and the Integrated Cryptographic Service Facility (ICSF), which allow users to gain access to both quantum-safe and classical cryptography. There are two supported application programming interfaces (APIs) - Common Cryptographic Architecture (CCA) and enterprise PKCS #11. There are APIs for managing quantum-safe keys, as well as APIs for generating and verifying quantum-safe digital signatures, and for encapsulating keys and protecting data.

There is a unique API call ciphertext translate. For encrypting data at the application level, clients can use the cipher text translation API to decrypt data that was encrypted with a quantum-vulnerable algorithm and re-encrypt it with a quantum-safe algorithm inside the HSM. IBM supports a hybrid key exchange scheme that allows clients to exchange keys in certain use cases with IBM partners. z/VM guest operating systems can also take advantage of the Crypto Express8 features. The IBM LinuxONE 4 quantum-safe APIs enable clients to begin using quantum-safe cryptography along with classical cryptography as they begin modernizing existing applications and building new applications.

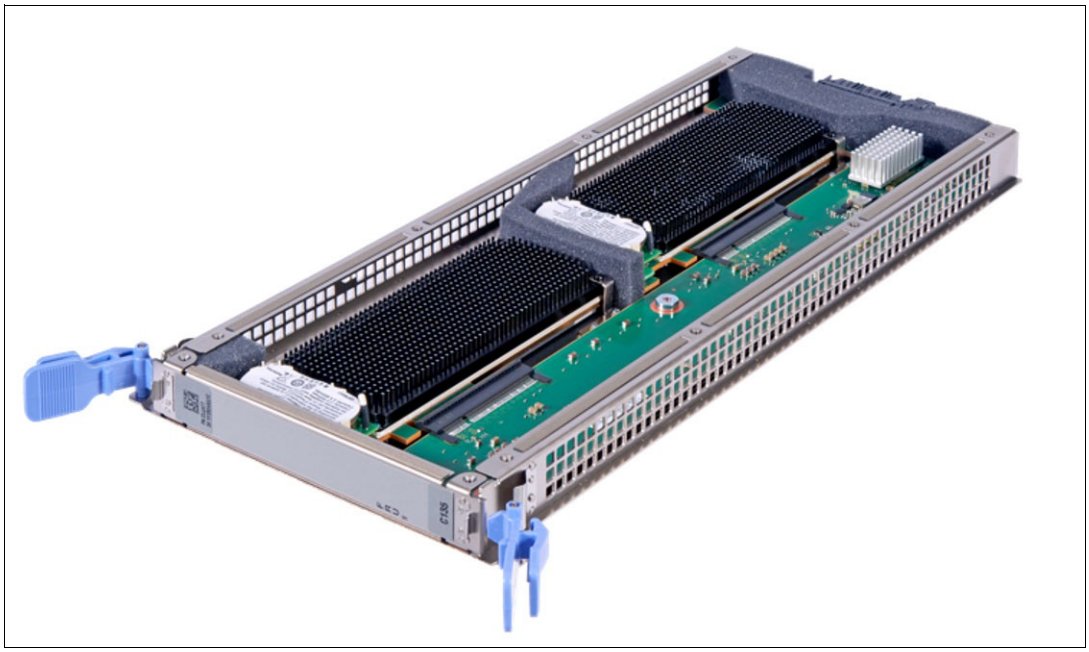

Figure 2-7 Crypto Express8S dual angled without cover

For more information on Crypto Express cards, see:

https://www.ibm.com/docs/en/cryptocards?topic=4770-overview

2.2 IBM Secure Execution for Linux

IBM Secure Execution for Linux is a hardware-based security technology that is built into LinuxONE III generation systems. It is designed to protect workloads from internal and external threats to help our clients prevent security breaches. IBM Secure Execution can help protect and isolate workloads on-premises, or on IBM LinuxONE hybrid cloud environments.

Current approaches to security address data-at-rest and data-in-transit. Not many users secure data when it is in use, which creates a window of vulnerability that insiders or criminals can use. Confidential computing is the industry movement around the use of technology to address this vulnerability.

Secure Execution is designed to further this agenda by protecting data that is in-use through the implementation of a hardware-based Trusted Execution Environment (TEE). Hardware-enabled protections can move clients closer to realizing a Zero Trust environment through workload isolation and hardened access restrictions over their data (see Figure 2-8).

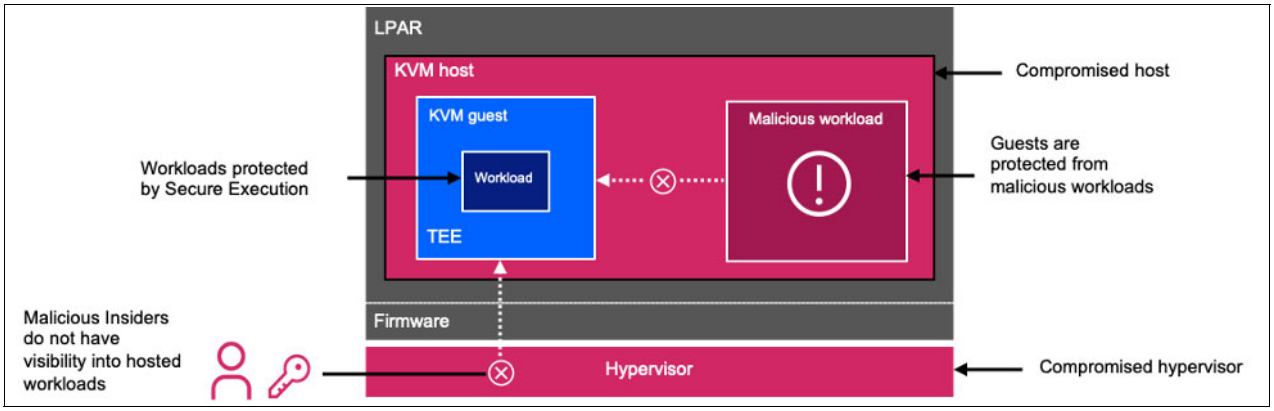

Figure 2-8 Solution overview

IBM Secure Execution for Linux protects your applications and their data from insider threats and external attacks by using the hardware-based environment that is provided rather than using current software-based approaches. By isolating the hardware that is hosting the workloads from the rest of the device, Secure Execution can enable sensitive workloads to run securely on untrusted or compromised infrastructure.

By providing a high level of isolation for workloads, Secure Execution is designed to help prevent security breaches that can result in large financial penalties, regulatory scrutiny, and company discharges.

At its core, Secure Execution provides a KVM-based VM that is fully isolated and protected from the hypervisor with encryption keys to which only LinuxONE hardware and firmware can access. This TEE is designed to protect and isolate workloads better than a standard software environment from internal and external threats.

As more companies move their on-premises workloads to public cloud, the need for a highly secure and trustworthy multi-tenant hosting solution becomes necessary to help support the confidentiality and integrity of each application and its data.

Secure Execution gives you the ability to use hardware-based security technology (TEE) to provide a mechanism by which a hosted workload can run without its memory or execution state being visible to the host or any other workload that is hosted in the same environment.

Enterprises can now protect data and code in-use in their hosted workloads by using protection mechanisms that are offered by Secure Execution. It also provides effective access controls so that only authorized users can access sensitive workloads.

Secure Execution is designed to eliminate the window of opportunity for hosts and guests that are infected with malicious code to use security lapses and gain full privileges to your hosted core business systems. Workload owners can use Secure Execution to help protect sensitive data from corruption and help support data confidentiality and integrity.

For more information about Secure Execution for Linux, see the IBM Knowledge Center.

2.2.1 IBM Secure Execution for Linux support for Crypto Express adapters

With IBM LinuxONE 4 servers, IBM Secure Execution for Linux supports Secure Execution guests with secure passthrough access to up to 12 Crypto Express 8S adapter domains in accelerator or EP11 coprocessor mode. This enables customers to achieve the highest level of cryptographic key protection (FIPS 140-2 level 4 certified) for their sensitive data by deploying workloads as Secure Execution KVM guests with access to Hardware Security Modules (HSMs) on LinuxONE 4 servers equipped with Crypto Express 8S adapters. This configuration offers substantial value for applications involving key and certificate management, multi-party computation, and digital assets. As confidential computing becomes increasingly widespread, the need to leverage highly certified HSMs grows, extending use cases to protect AI models and ensure data sovereignty across different organizational and infrastructure boundaries.

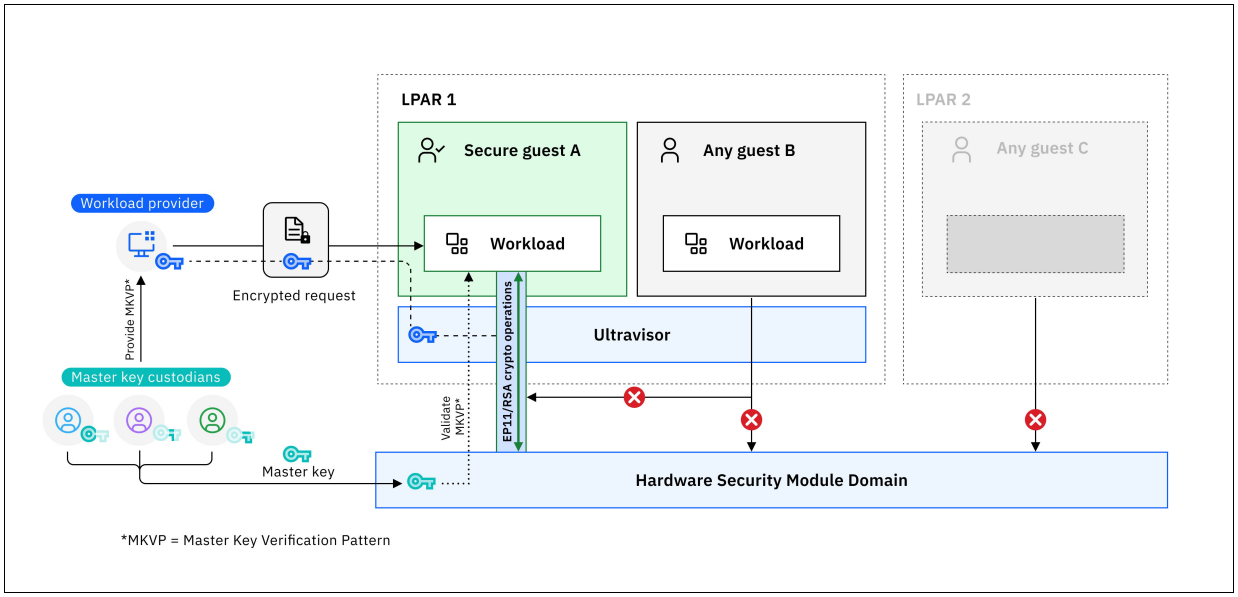

Figure 2-9 Solution overview

To proceed, ensure you have a secure-execution boot image capable of supporting the insertion of secrets into the ultravisor. You will need Crypto Express8S adapters that can be configured in the following modes:

- Accelerator Mode: The adapters can be set up to function as accelerators, providing enhanced cryptographic performance.
- Enterprise PKCS #11 Coprocessor Mode: For this mode, ensure you are using Enterprise PKCS #11 version 5.8.30. This mode allows the adapter to act as a PKCS #11 compliant cryptographic coprocessor.

Additionally, the adapter domains for the Crypto Express8S adapters must be configured in passthrough mode (dedicated). This configuration supports up to a maximum of 12 adapter domains per secure guest. To exploit this new function, IBM LinuxONE 4 severs with firmware bundles S30 and S31b are needed. To use a Crypto Express 8S adapter in EP11 mode the minimal EP11 firmware version loaded must be 5.8.30.

2.3 IBM Secure Boot for Linux

IBM LinuxONE 4 supports quantum-safe secure boot. Secure Boot is the mechanism that verifies the integrity of code being loaded, before it's allowed to execute. The secure boot process includes checking code for the proper signature by an approved signer. IBM LinuxONE uses a quantum-safe mechanism that leverages a dual-signing scheme which includes both quantum-safe and classical crypto algorithms. IBM LinuxONE protection is anchored in a hardware root of trust. The dual-signing scheme helps ensure the authenticity of the firmware that is launched on the system. This helps defend against ransomware attacks that attempt to inject malware into the bootup process.

IBM Secure Boot for Linux brings boot integrity to the LinuxONE platform, which is a complete chain of trust from trusted source to boot loader. Secure Boot is part of the Unified Extensible Firmware Interface (UEFI), which is a central interface between the firmware, operating system, and individual components of a computer.

This capability protects your system from root level attacks and viruses that target vulnerabilities during the boot process. The system checks images at boot time for a vendor-signed cryptographic key to verify that the image is from an official provider, and that the image was not tampered with or replaced by malicious third parties.

This feature can be enabled through a simple interface option on the Hardware Management Console. The system firmware first confirms that the system boot loader is signed with a verified cryptographic key. The system then confirms that the key was authorized by a database that is contained in the firmware and only recognized keys allow the system to boot.

- Helps protect system integrity by using quantum-safe and classical digital signatures to perform a hardware-protected verification of the IBM LinuxONE firmware loaded while booting the machine.
- This protection is anchored in the IBM LinuxONE hardware root of trust for its firmware chain of trust.
- Both quantum-safe and classical cryptography provide a double layer of protection for the IBM LinuxONE firmware chain of trust that spans the classical and quantum computing paradigms.
- Helps protect IBM LinuxONE firmware from quantum attacks through a built-in dual signature scheme with no changes required.
- Dual signing scheme is designed to ensure only authentic firmware is launched, thereby helping to protect the system from the threat of fraudulent firmware attacks.
- Helps defend against ransomware attacks that attempt to inject malware into the boot up process.

2.3.1 PR/SM and LPARs

The IBM LinuxONE system has a unique capability to implement a hypervisor at the hardware and firmware level. The hardware hypervisor is IBM Processor Resource/Systems Manager (PR/SM), which is informally referred to as PRISM. PR/SM is implemented in firmware as a part of the base system that fully virtualizes the system resources and runs without any extra software.

PR/SM is a Type-1 hypervisor that runs directly on bare metal. With it, you create multiple isolated partitioned environments on the same physical server. These isolated environments are known as *logical partitions* (LPARs).

EAL 5+ isolation and cryptographic key protection

LinuxONE systems feature EAL 5+ isolation and cryptographic key protection. EAL5+ is a regulatory certification for LPARs that verifies the separation of partitions to improve security. Therefore, you can run many virtual servers concurrently, and use the ability of LinuxONE to isolate and protect each virtual server as though they were running on physically separated servers.

LinuxONE LPARs provide excellent isolation between each other, but not between the VMs or containers within the same LPAR. Secure Execution for Linux is a LinuxONE III hardware capability that hypervisors can use to isolate virtual machines and containers from each other within an LPAR

The PR/SM-based LPARs of LinuxONE are the only technology that is commercially available that can provide this highly certified level of isolation between workloads.

Isolation and cryptographic key protection are achieved by using a dedicated cryptographic coprocessor. The CP Assist for Cryptographic Function (CPACF) delivers cryptographic and hashing capabilities in support of clear-key operations. The Crypto Express adapter is used to create the fortified data perimeter by using the IBM LinuxONE protected key in which the keys that are used in the encryption process are not visible to the applications and operating system.

Each LPAR on a LinuxONE system has its own uniquely generated and assigned cryptographic keys that are held in a secure hardware area. This configuration provides a level of cryptographic isolation between secure environments that is called for under many regulatory compliance frameworks (for example, PCI-DSS).

2.3.2 Kernel-based virtual machine

Kernel-based virtual machine (KVM) is a popular open-source Linux hypervisor and a key technology for the LinuxONE platform. It is a Type-2 hypervisor that provides simple, cost-effective virtualization technology for Linux workloads. It also allows sharing of CPU, memory, and I/O resources and can coexist with other types of virtualization technologies simultaneously running on LinuxONE. KVM on LinuxONE frees operators to adopt and switch from various hardware platforms, and more familiar interfaces.

For more information about use cases and examples, see KVM on IBM Z and LinuxONE.

One of the advantages for KVM virtualization is the familiar standard Linux user interfaces for open-source developers, offering a low barrier to adoption and easy integration with hybrid environments.

KVM on LinuxONE is supported through the following Linux distribution partners:

- Red Hat Enterprise Linux (RHEL)
- SuSE Linux Enterprise Server
- Canonical Ubuntu

2.3.3 z/VM

IBM z/VM provides high levels of extreme security, scalability, and efficiency; therefore, it also provides a robust foundation for on-premises cloud computing. The z/VM virtualization technology is designed to run hundreds to thousands of Linux servers on a single IBM Z or IBM LinuxONE server with the highest degrees of efficiency, elasticity, and security. Its ability

to support numerous machine images and solution's architectures provides a highly flexible production and test environment for IBM Z and LinuxONE operating systems

Do more with less. Virtualization helps to deploy more servers, networks, applications, and data on less physical hardware, and the capacity of Z/LinuxONE is outstanding. With LinuxONE, you can achieve nearly 100% utilization of system resources nearly 100% of the time. Virtualization on LinuxONE provides high levels of resource sharing, I/O bandwidth, and system availability.

Consider the following points:

- You can reduce costs on a bigger scale by running software that is paid by core on fewer cores, which saves on software license fees.
- High server density helps to use less power and floor space.
- The high resiliency of z/LinuxONE minimizes hardware that is needed for business continuance and disaster recovery (DR).

Manage growth and complexity in the following ways:

- In support of the high density and mass of virtual servers, facilities for life cycle management are provided (for example, workload management, monitoring, security, charge back, patching, backup, and recovery).
- Hardware resources can be added to a running system without disruption, which helps on a continuous business operation and availability of the services.
- LinuxONE can scale horizontally and vertically, scaling workload deployment on such a “scale up” machine means fewer components to manage.

Provide more flexibility and minimize lead time for new projects by using:

- Workload deployment to a single IBM Z or IBM LinuxONE server offers significant advantages in terms of flexibility.
- Rapid provisioning, which reduces lead time for new IT projects and helps to increase business agility.
- Resources that can be assigned dynamically and efficiently between workloads, whenever and wherever they are needed.
- Virtual machines for Cloud deployment. IBM Cloud Paks are based on Red Hat OpenShift Container Platform; OCP 4.2 needs z/VM-based VMs for the implementation.

z/VM is IBM’s internally developed Type-2 hypervisor that manages the sharing of a LinuxONE system’s physical resources between virtual guests. The z/VM hypervisor typically runs on an LPAR and manages Linux VMs. It also can manage other types of operating systems, including z/VM, on top of the hypervisor.

z/VM is a proven and established virtualization platform with industry-leading capabilities for efficient vertical and horizontal scaling that was proven over many decades. For more information, see IBM z/VM.

The following operating systems are supported:

- Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 7, Red Hat Enterprise Linux 6, and Red Hat Enterprise Linux 5
- SUSE Linux Enterprise Server 15, SUSE Linux Enterprise Server 12, SUSE Linux Enterprise Server 11, and SUSE Linux Enterprise Server 10
- Ubuntu 18.04 and Ubuntu 16.04

The smart economics of the use of z/VM are shown in Figure 2-10.

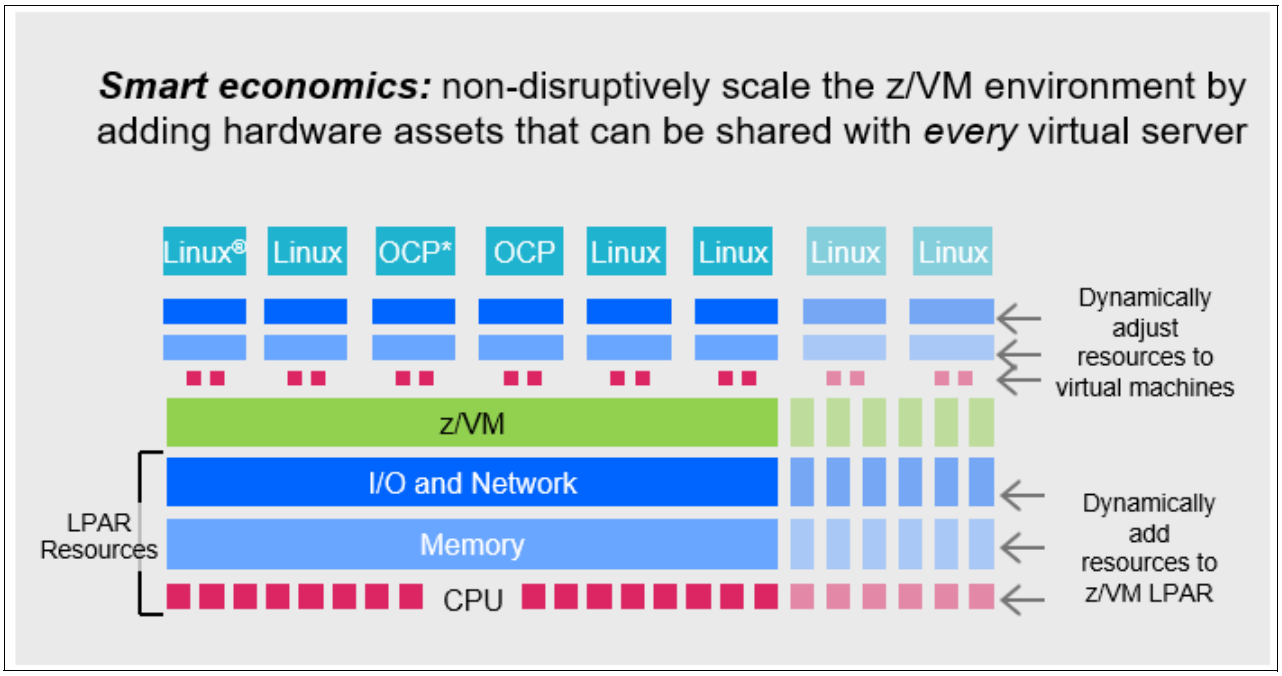

Figure 2-10 Smart economics of the use of z/VM

Other security capabilities on z/VM

z/VM supports encrypted paging in support of the philosophy of encrypting all data that is in-flight and at-rest (available with z/VM V7.1 and z/VM V6.4). Ciphering occurs as data moves between active memory and a paging volume.

IBM RACF® for z/VM provides security systems that include access control and auditing functionality as the backbone for Linux security.

3

Users of security on LinuxONE

Building on the secure foundation of the LinuxONE platform, IBM is continuously developing new technology further up the stack that can use the system's industry-leading security capabilities.

Organizations can take advantage of these unique offerings to differentiate their services from competitors with the power and speed that users demand, the security-rich environment that businesses and regulators require, and efficiencies that lower operational expenditures.

This chapter includes the following topics:

3.1 IBM Hyper Protect Services

Built on IBM LinuxONE technology, IBM Hyper Protect Services provide built-in data-at-rest and data-in-use protection to help developers easily build applications with highly sensitive data.

These services are available on-prem and in IBM Cloud with IBM LinuxONE, and are infused with enterprise-grade data protection.

Now, developers and clients can build, deploy, and host applications with industry-leading data protection that encrypts information that is at rest and in use. This technology is designed to help protect against threats inside and outside of an organization.

The IBM Hyper Protect family provides the following services and intends to expand to include others that are crucial for providing protected cloud capabilities:

- IBM Cloud offerings
  - IBM Cloud Hyper Protect Crypto-Services
  - IBM Cloud Hyper Protect Virtual Servers for Virtual Private Cloud (VPC)
- IBM On-Premises Solution
  - IBM Hyper Protect Virtual Servers
  - IBM Crypto Express Network API for Secure Execution Enclaves
  - IBM Hyper Protect Offline Signing Orchestrator
- IBM Cloud and On-Premises
  - IBM Hyper Protect Secure Build

For more information on Hyper Protect Services, see section 7.4, “Hyper Protect Services” on page 83.

3.1.1 IBM Cloud Hyper Protect Crypto Services

IBM Cloud Hyper Protect Crypto Services is an as-a-service (aaS) key management, Hardware Security Module (HSM) and encryption solution, which gives you full control over your encryption keys for data protection.

The integrated Unified Key Orchestrator acts as a secure key repository for distributing and orchestrating keys across multiple clouds, enabling quick recovery from key loss or disasters. With Hyper Protect Crypto Services, you can:

- Build on the highest level of security with FIPS 140-2 level 4 certified hardware.
- Experience a worry-free approach to multi cloud key management through the all-in-one as-a-service solution. Benefit from automatic key backups built-in high availability secure business continuity and disaster recovery.
- Manage your keys seamlessly across multiple cloud environments. Create keys securely and bring your own key seamlessly to hyperscalers such as Microsoft Azure AWS and Google Cloud Platform to enhance the data security posture and gain key control.
- Protect data by pervasively encrypting data at rest and in transit with Keep Your Own Key (KYOK).

Built on IBM LinuxONE technology, the service helps ensure that only you can access your keys. A single-tenant key-management service with key vaulting that is provided by dedicated customer-controlled HSMs helps you to create encryption keys with ease.

Alternatively, you can bring your own encryption keys to manage. The managed cloud HSM supports industry standards, such as PKCS #11, so that your applications can integrate cryptographic operations like digital signing and validation. Keep your own keys for cloud data encryption that is protected in a dedicated cloud HSM. Maintain control of the key hierarchy, including the HSM master key.

For more information, see IBM Cloud Hyper Protect Crypto Services.

3.1.2 IBM Cloud Hyper Protect Virtual Servers for VPC

IBM Cloud Hyper Protect Virtual Servers for Virtual Private Cloud (VPC) is a fully managed confidential compute container runtime that enables the deployment of sensitive containerized workloads in a highly isolated environment with technical assurance. It protects instances in all states within the data lifecycle: at-rest, in-transit, and now in use, with confidential computing. Unique to the market, it utilizes IBM Secure Execution for Linux to enhance data protection to achieve data privacy and protection over containerized workloads with sensitive data or business intellectual property (IP).

Through encrypted contract, different personas can provide their contribution while ensuring no access to the data, integrity of workloads and environments. The deployment can be human audited through a signed and encrypted attestation record to certify integrity. Workloads are locked down by individual, instance level secure boundaries. Technical assurance that unauthorized users, including IBM Cloud admin, will not be able to access the environment and the data.

For more information, see Auditable deployment of trustworthy container images in a tamper-proof environment.

3.1.3 IBM Hypertext Virtual Servers (On-Premises)

IBM Hyper Protect Virtual Servers (HPVS) are available as an on-premises solution as well, based on IBM LinuxONE, with all the capabilities as described in section 3.1.2, “IBM Cloud Hyper Protect Virtual Servers for VPC” on page 37. For more information, see IBM Hyper Protect Virtual Servers 2.1.x.

3.1.4 IBM Hyper Protect Offline Signing Orchestrator

IBM Hyper Protect Offline Signing Orchestrator (OSO) is a technology to deploy cold storage solutions for digital assets. OSO is designed to address limitations of current cold storage offerings for digital assets, including the need for people to perform manual procedures for the execution of a cold storage transaction. It helps protect high-value transactions by offering additional security layers including disconnected network operations, time-based security and electronic transaction approval by multiple stakeholders.

Hyper Protect OSO provides a policy engine that brokers communication between two different applications that have been designed not to communicate directly with each other for security purposes, providing an efficient and securable solution to facilitate digital asset transactions.

For more information, see Getting started with Hyper Protect Offline Signing Orchestrator.

3.1.5 IBM Hyper Protect Secure Build

IBM Hyper Protect Secure Build is a process to build and deploy images by using a trusted source and method. It creates a repeatable and definable process to create the images that are approved by the organization. This includes security checks throughout to verify that nothing abnormal or malicious is included in the containerized image. If any of these checks fail, the image will be rejected. In addition, the build automation process itself is encrypted within its own enclaves, protected from all tampering. In summary, secure build allows you to inject audit capabilities in the build process, which ensures all images that are published and deployed are approved and signed by the security audit team.

The secure containers built through the Secure Build process can be run as container workloads on IBM HPVS (On-Premises) and HPVS for VPC.

For more information, see the on-premises documentation, Building your applications with Hyper Protect Secure Build and the cloud documentation, Hyper Protect Secure Build.

3.2 IBM Cloud Infrastructure Center (ICIC) with Secure Execution

IBM Cloud Infrastructure Center is an advanced infrastructure management (OpenStack compatible) software that includes on-premises cloud deployments of KVM-based Linux® virtual machines on the LinuxONE platforms. It provides lifecycle management for the virtual infrastructure that is based on Red Hat KVM and enables the automation of infrastructure services. The IBM Cloud Infrastructure Center now supports the life cycle management of Secure Execution (SE) guest for KVM. KVM guest images created using IBM Secure Execution technology, can be uploaded to ICIC to perform life cycle operations of a Secure Execution enabled KVM guest.

For more information, see Planning for secure execution.

3.3 IBM Fibre Channel Endpoint security

Fibre Channel is a high-speed data transfer protocol for storage area networks (SAN). Yet, as with any other component of a data center, the need to implement security measures in the SAN exists to reduce and eliminate insider threats of unauthorized access to data. Managing a SAN involves not only providing highly available data access and optimal performance, but it is also essential that all data on the SAN be secure at all times.

IBM Fibre Channel Endpoint Security (IFCES) offers an end-to-end solution that ensures all data flowing on FICON and Fibre Channel Protocol (FCP) links from IBM LinuxONE to IBM storage or supported SAN switches and directors, or between IBM LinuxONE platforms over FICON Channel-to-Channel connections, is encrypted and protected. This offering provides in-flight protection for all data, independent of the operating system, file system, or access method in use. IFCES integrates seamlessly within the IBM LinuxONE ecosystem, safeguarding all data transmitted over FICON and FCP links. This ensures secure data exchanges regardless of the operating system, file system, or access method, thereby providing comprehensive protection within the IBM LinuxONE, storage systems, and supported SAN switches and directors.

In this chapter, we delve into the security features and mechanisms associated with IBM Fibre Channel (FC) Endpoint Security.

This section provides a comprehensive overview of how IBM enhances data protection in fibre channel environments, ensuring secure and reliable data transmission.
End-to-end protection features the following benefits:

- Enabled automatically between host and storage endpoints that are security-capable.
- Each established link must “prove” its identity as a trusted component.
- Trusted connections are identified and visible to both the Operating System and HMC.
- Policy can be established to enforce that only trusted connections can be made.
- Each time a link goes down or up, reauthentication or negotiation of device encryption keys occurs.
- Integrated key management by using IBM Security™ Key Lifecycle Manager (ISKLM).
- Can be used immediately after Power-on-Reset or enabled later by running IBM LinuxONE with minimal or no disruption.

Because of these benefits, IBM Fibre Channel Endpoint security can help you realize the following:

- Meet regulatory and compliance mandates.
- Minimize the enterprise risk and effect of receiving and storing sensitive data.
- IBM Fibre Channel Endpoint security is another data security technology that contributes to the IBM LinuxONE approach of encryption everywhere. It extends the value of pervasive encryption, which further minimizes the risk of security breach, potential noncompliance, and financial liability.

For more information, see *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455.

3.3.1 Overview of Fibre Channel Security

Given Fibre Channel’s critical role in data infrastructure, ensuring the security of FC endpoints is paramount to protect data integrity, confidentiality, and availability.

Key Concepts:

- **Data Integrity:** Ensuring that data remains accurate and unaltered during transmission.
- **Confidentiality:** Protecting data from unauthorized access and breaches.
- **Availability:** Ensuring that data and resources are available to authorized users when needed.

3.3.2 Security features and mechanisms

IBM’s FC Endpoint Security incorporates a range of features designed to safeguard data transmissions and endpoint devices from potential threats. Security features include the following:

- **Authentication**:
  - **Mutual Authentication:** Ensures that both the initiator and target authenticate each other before data transmission begins. This prevents unauthorized devices from accessing the network.
  - **Digital Certificates:** Utilized for authenticating endpoints, enhancing security by verifying identities through trusted certification authorities.

- **Encryption**:
  - **Data-in-Transit Encryption:** Protects data as it moves across the network, preventing interception and unauthorized access. IBM uses AES-256 encryption to ensure robust data protection.
  - **Hardware-Based Encryption:** Offloads encryption processes to dedicated hardware, reducing the impact on system performance and ensuring high-speed encrypted data transfers.
- **Access Control:**
  - **Zoning:** Logical segmentation of the SAN to control access to devices. Only devices within the same zone can communicate, limiting the potential for unauthorized access.
  - **LUN Masking:** Controls access to specific storage volumes (LUNs), ensuring that only authorized hosts can access certain data sets.
- **Monitoring and Logging:**
  - **Activity Logs:** Comprehensive logging of access and transmission activities, enabling detailed audits and tracking of data interactions.
  - **Intrusion Detection Systems (IDS):** Monitors network traffic for suspicious activities and potential security reaches, providing real-time alerts and automated responses.

Tip: Regularly reviewing and analyzing logs can help identify and mitigate potential security threats before they escalate.

3.3.3 Best practices for implementing Fibre Channel Endpoint Security

To maximize the effectiveness of IBM FC Endpoint Security, it is crucial to follow best practices during implementation and ongoing operations. The following are some of the best practices:

- **Network Design and Segmentation:**
  - **Implement Zoning:** Use zoning to segregate devices based on function, security requirements, or operational needs.
  - **Isolate Sensitive Data:** Ensure that sensitive data is only accessible within secure zones, minimizing the risk of unauthorized access.
- **Regular Updates and Patch Management:**
  - **Firmware and Software Updates:** Keep all FC hardware and software components up to date with the latest security patches and firmware updates.
  - **Automated Patch Management:** Utilize automated tools to manage and deploy patches, reducing the risk of human error and ensuring timely updates.
- **Robust Authentication Mechanisms:**
  - **Use Strong Authentication Methods:** Implement mutual authentication and use strong digital certificates to verify endpoint identities.
  - **Periodic Credential Reviews:** Regularly review and update credentials and access permissions to ensure only authorized users have access.
- **Comprehensive Monitoring and Auditing:**
  - **Enable Detailed Logging:** Ensure that all access and activity logs are enabled and stored securely for audit purposes.

– **Regular Audits:** Conduct regular security audits and reviews to identify and address potential vulnerabilities.

3.3.4 Compliance and regulatory considerations

Implementing Fibre Channel Endpoint Security must also align with industry standards and regulatory requirements to ensure compliance and mitigate legal risks.

Key regulations

The following are some of the key regulations that must be followed.

- **General Data Protection Regulation (GDPR):** Ensures data protection and privacy for individuals within the European Union.
- **Health Insurance Portability and Accountability Act (HIPAA):** Mandates the protection of health information in the United States.
- **Payment Card Industry Data Security Standard (PCI-DSS):** Sets requirements for protecting payment card information.

Compliance strategies

The following are some of the compliance strategies that should be considered.

- **Documentation and Reporting:** Maintain comprehensive records of security measures, configurations, and access logs to demonstrate compliance during audits.
- **Regular Compliance Audits:** Conduct internal and external audits to ensure ongoing compliance with relevant regulations and standards.

Tip: Leveraging automated compliance tools can streamline the process of maintaining and demonstrating regulatory compliance, reducing the administrative burden on IT teams

3.3.5 Summary

IBM Fibre Channel Endpoint Security provides a robust framework for protecting data in SAN environments. By integrating advanced authentication, encryption, access control, and monitoring features, IBM ensures the confidentiality, integrity, and availability of critical data transmissions. Adhering to best practices and compliance requirements further enhances the security posture, enabling organizations to securely manage their data infrastructure.
Table 3-1 outlines the key features of IBM Fibre Channel Endpoint Security.

Table 3-1 Key features of IBM Fibre Channel Endpoint Security

Feature	Description
Mutual Authentication	Ensures both initiator and target authenticate each other
Data-in-Transit Encryption	Protects data moving across the network with AES-256 encryption
Zoning	Segregates devices within the SAN for controlled access
LUN Masking	Restricts access to specific storage volumes

Feature	Description
Activity Logs	Comprehensive logging for audit and tracking
Intrusion Detection	Monitors for suspicious activities and potential breaches

For more information on IBM Fibre Channel Endpoint Security, see IBM Fibre Channel Endpoint Security.

3.4 Cryptographic Key Management for LinuxONE

As described in Chapter 2, “Core security technologies on LinuxONE” on page 17, different types of keys are used for different encryption algorithms. The “Secrecy of Data” is maintained by the encryption keys, not by the encryption algorithms. Although in the early days of cryptography, the encryption algorithm was regarded as the protector of the information, this belief proved to not always be true.

The method of storing keys is important. It is also important how often you change the keys before they get too old. It is often referred to as the ‘”lifecycle” of a key or a pair of keys that are generated (born), used (lives), and then destroyed or changed (dies).

When keys are created and stored, it also is important to decide who can create and distribute the key and decide on owners for a key that is split into different segments. The National Institute of Standards and Technology (NIST) defines this role as a *key custodian*, and every enterprise must assign that role to a party that is responsible for key lifecycle management.

In this section, we discuss this key lifecycle management in two parts:

- “Operational Key Lifecycle Management” on page 42 - Operational Encryption Key Management with IBM Security Key Lifecycle Manager (ISKLM)
- “Master Key Lifecycle Management” on page 44 - Master Key Management for Crypto Express Adapter with Trusted Key Entry (TKE)

3.4.1 Operational Key Lifecycle Management

Key lifecycle management is a critical aspect in any encryption strategy. Cryptographic keys feature a lifecycle that includes tasks, such as key creation, key activation, key deactivation, key archival, and key deletion. Some regulations, such as European Union (EU) General Data Protection Regulation (GDPR), Payment Card Industry Data Security Standard (PCI-DSS), and Health Insurance Portability and Accountability Act (HIPAA), require key management processes to be created and well-documented.

Encryption Algorithms and Sharing Keys

When an application must encrypt and decrypt data, it chooses standard encryption algorithms to communicate with other applications and the other entity of the application. Examples of these encryption standards are RSA, DES, AES, ECC, and SHA. Each standard has a specific purpose for its use. When an application encrypts data, it typically uses multiple sets of encryption algorithms to make it more difficult for unauthorized parties to decrypt the data.

For example, if a database uses Advanced Encryption Standard (AES) from the US National Institute of Standards and Technology (NIST) to encrypt its table data, the database also uses the RSA cryptosystem to encrypt the AES encryption key. As a result, the key can be shared with other applications in a secure way.

Many variations and combinations of encryption algorithms exist that can be used and shared. Managing the keys for those methods also varies, but some popular standards, such as the Key Management Interoperability Protocol (KMIP)[1] and Public-Key Cryptographic Standards (PKCS)[2], are available.

IBM Security Key Lifecycle Manager - ISKLM

IBM Security Key Lifecycle Manager (ISKLM) offers a simple integration to your enterprise applications and infrastructures to manage operational cryptographic keys. ISKLM centralizes, simplifies, and automates the encryption key management process to help minimize risk and reduce operational costs of encryption key management. It offers secure, robust key storage, key serving, and key lifecycle management for IBM and non-IBM storage solutions that use the OASIS Key Management Interoperability Protocol (KMIP).

IBM Security Key Lifecycle Manager helps customers meet regulations, such as the Payment Card Industry Data Security Standard (PCI DSS), Sarbanes-Oxley, and the Health Insurance Portability and Accountability Act (HIPAA), by providing centralized management of encryption keys.

ISKLM runs on LinuxONE and supports Red Hat Enterprise Linux and SuSE Linux Enterprise Server. It is also integrated with KMIP and PKCS#11. Therefore, you can use it to manage keys for storage devices, such as IBM DS8000®, Spectrum Scale, Virtual Tape libraries, and middleware, such as IBM Db2® for LinuxONE.

For example, if Db2 for LinuxONE is set up to use Db2 Native Encryption, it uses the IBM GSkit interface to use CPACF hardware in CPU to accelerate AES encryptions and decryptions. Db2 Native Encryption also supports KMIP, where ISKLM can be used to manage operational keys for the encryption key lifecycle management.

If ISKLM is hosted on a LinuxONE server with the proper JAVA SDK level (that is, JAVA SDK v7/v8), it uses IBMPKCS11Impl to use the Crypto Express Adapter to use hardware acceleration for PKCS#11 functions. You can also use the **`pkcsconf -m`** command to display the supported mechanisms for each slot on a LinuxONE system.

[1] KMIP is an open standard method to standardize the key management within the companies.
[2] PKCS is a set of specifications that were developed for public key cryptography and initially developed by RSA Data Security Inc.

An overview of ISKLM is shown in Figure 3-1.

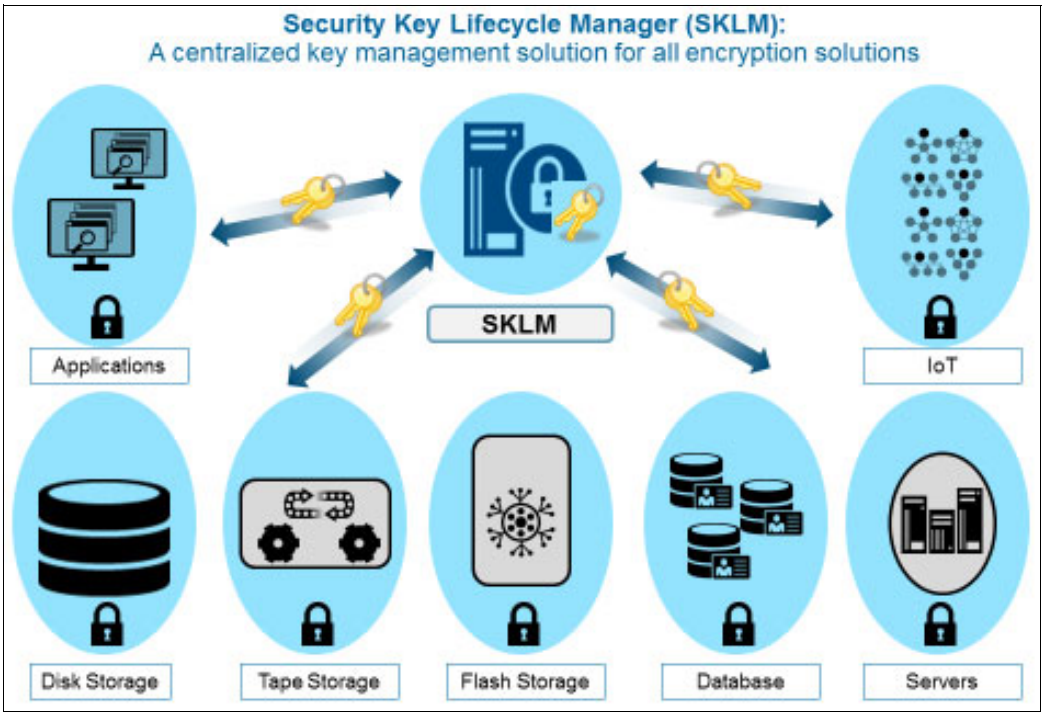

Figure 3-1 ISKLM overview

3.4.2 Master Key Lifecycle Management

The Trusted Key Entry (TKE) workstation is an optional feature of LinuxONE that manages cryptographic keys in a secure environment. The TKE workstation is an integrated solution that contains a combination of hardware, firmware, and software to provide a basic key management tool for the cryptographic coprocessors. TKE securely manages multiple cryptographic modules that run in Common Cryptographic Architecture (CCA) or IBM Enterprise PKCS#11 (EP11) and uses compliant-level hardware-based key management techniques from a single point of control.

Trusted Key Entry (TKE) workstation is a platform that provides convenient way to manage Master Keys in the HSM (in this case, Crypto Express Adapter for LinuxONE).

Master keys on LinuxONE

When a Crypto Express Adapter is configured to wrap an encryption key (that is, a key encrypting key, also known as *secure key*), the root key that is used to encrypt secure keys (also known as the *master key*) is stored in the hardware. This master key that is stored in the hardware never leaves its entity, and is processed only within the hardware to decrypt a key when called by a hardware function. Therefore, a secure key (encrypted key) that is transported into a Crypto Express adapter is used inside that adapter to encrypt or decrypt data with a specific encryption algorithm that is tied to that key.

Many organizations have security policies or specific requirements about how to set the master keys, and how often they should be changed/rotated (see section 7.2.3, “Key rotation best practices” on page 79). Most of the time, the master keys are divided into multiple pieces and distributed to multiple key custodians. At the time of rotating the master keys or generating new master keys, all the key custodians get together along with witnesses to set the new master keys.

This process often is called *Key Ceremony*, in which all the required persons get together to securely change the master key. This process adds another layer of security into the process, which makes it more difficult for the keys to be exposed.

TKE: Why do you need it?

You can use a `panel.exe` software tool to change master keys in a Crypto Express Adapter for LinuxONE. It is a free command-line tool that provides a basic interactive session to set up the master keys and change them when needed. For more information about the use of this command-line interface, see the following resources:

- *Getting Started with Linux on Z Encryption for Data At-Rest*, SG24-8436.
- The panel.exe utility topic in IBM Knowledge Center.

To make a new master key or change the current master key in the HSM, you must enter the current master key into the terminal by using `panel.exe` to load the keys into the Crypto Express Adapter. This method has many potential security exposures because it is connected through the user's terminal to LinuxONE. The keys that are entered on terminal might be exposed in the following ways:

- In host memory on the system where they are entered.
- On the network channel for communication to the LinuxONE host.
- In host memory for the Linux partition.

It also gets tricky to enter all the keys in hexadecimal number form in the terminal without having any errors because one mistake can lead to the loss of an encryption key, which can in turn lead to losing access to the data.

Therefore, IBM is offering the TKS workstation as a solution. It makes the master key management tasks easy and secure compared to the software-only method.

TKE: How it works

Trusted Key Entry is a specialized appliance that is built with custom hardware and software. The hardware is composed of a x86 core-based workstation, which is equipped with a PCI-express cryptographic co-processor. This adapter is the same Crypto Express Adapter that is used in LinuxONE. It also features a pair of smart card readers that are attached so that Master Keys are stored in the smart cards.

The TKE smart panel is shown in Figure 3-2 on page 46.

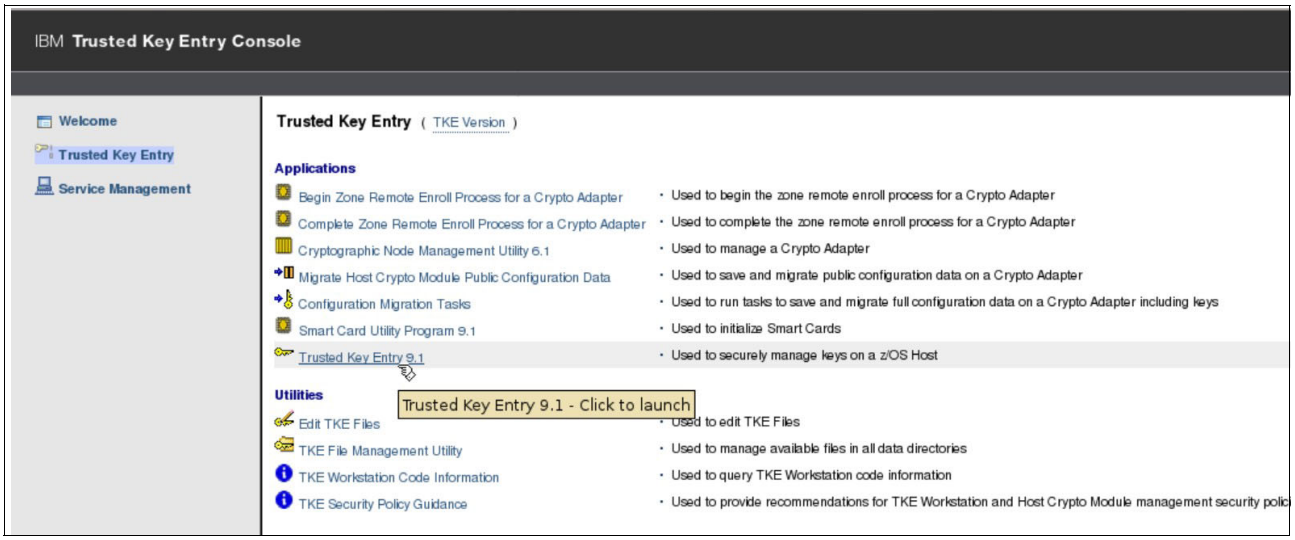

Figure 3-2 Trusted Key Entry Menu panel

The TKE application is designed to give operators and key custodians the most secure way to manage master keys. By using six smart cards (two for certificate authority [CA] cards, two for module [domain] admin cards and two for split master keys), TKE gives maximum security for creating or changing master keys and maintaining them in the most secure format.

A Linux server that runs on LinuxONE has a TKE daemon running that is called `catcher.exe`. This daemon allows the TKE workstation to be connected and communicate downward to the Crypto Express Adapters. By using the GUI on the TKE workstation, each key custodian performs the tasks that each person is assigned to. Figure 3-3 shows the TKE panel with the statuses of various master keys.

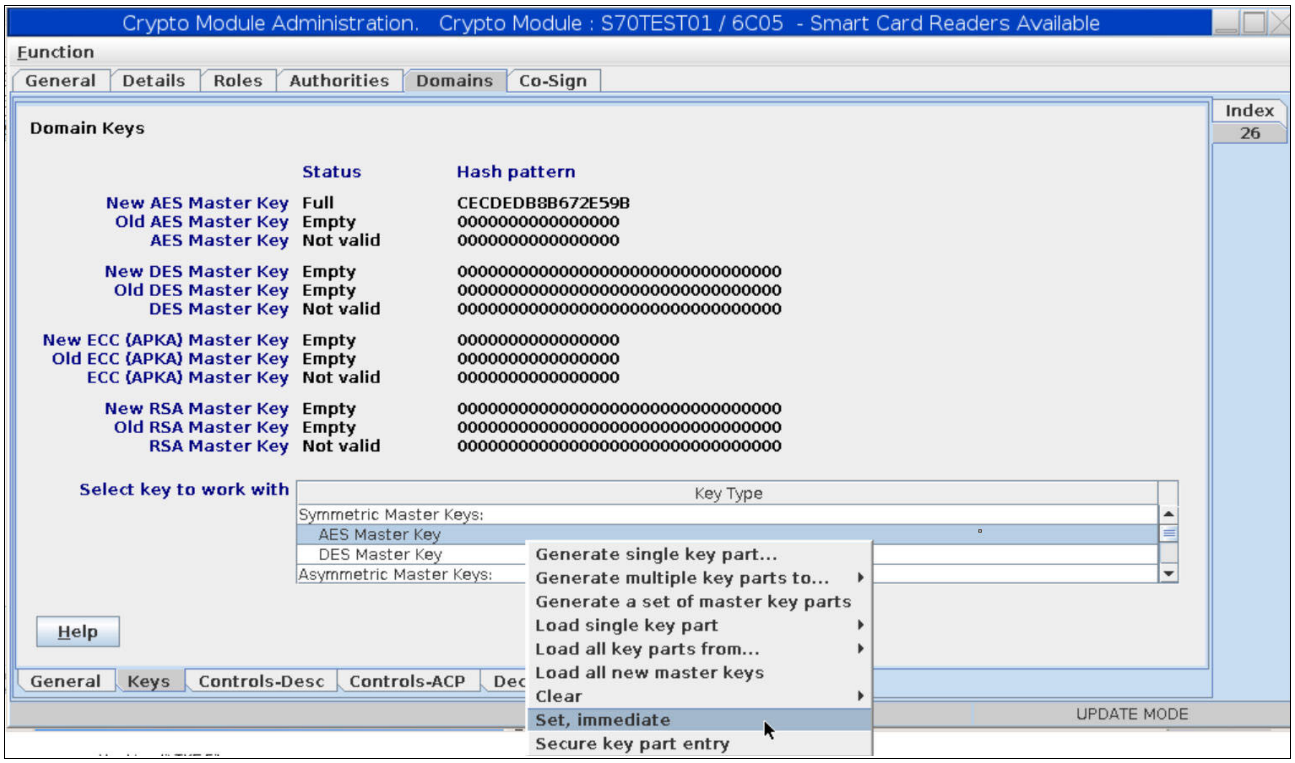

Figure 3-3 Setting the AES Master Key on TKE

Because TKE does not allow a single user to change the master key or allow actions without proper procedures by key owners, this appliance requires the steps that are necessary to provide the most secure way to manage the keys.

The enterprise always must prepare service availability. With data encryption, the production servers and disaster recovery servers must maintain the same master keys, even though they might not be in the same data center. Therefore, after the master keys are created or regenerated in the production servers, they must be loaded into other servers that share and use the encrypted data, including cold-state disaster-recovery servers.

For more information about how to configure and operate TKE, see the IBM Knowledge Center.

4

Quantum-safe computing

Quantum cryptography (also known as quantum encryption) refers to various cybersecurity methods for encrypting and transmitting secure data based on the naturally occurring and immutable laws of quantum mechanics.

Quantum-safe computing, also known as post-quantum cryptography or quantum-resistant cryptography, refers to the development of cryptographic algorithms and protocols that remain secure even in the presence of powerful quantum computers. Quantum computers, once fully realized, have the potential to break many of the cryptographic systems that currently underpin the security of digital communications and transactions. Quantum-safe computing aims to mitigate this threat by designing cryptographic techniques that are resistant to attacks by quantum computers.

LinuxONE provides a robust and secure platform for deploying and managing the infrastructure and applications required for quantum-safe computing and artificial intelligence (AI). Its combination of security, scalability, reliability, and compatibility makes it a strategic choice for organizations looking to harness the power of these emerging technologies while ensuring the integrity and security of their data and workloads.

This chapter includes the following topics:

4.1 Quantum-safe algorithms

Organizations use cryptographic algorithms that are used to protect their data at rest and data in transit. Cryptographic algorithms are used for protecting data with respect to Confidentiality, Integrity, Authentication and Non Repudiation. Cryptographic algorithms such as symmetric (AES, 3DES, and others) and public-key schemes (such as ECC, RSA, ECDSA, ECDH, and others) are used for encryption of Data. As Quantum Computing advances, these algorithms can potentially be threatened by attackers, exposing the data. Use of Quantum computer with Shor's and Grover's algorithms can break or weaken some of these Cryptographic algorithms.

Although attackers cannot easily break most encrypted data today, they might be able to decrypt that data in the future by using a large quantum computer, also known as a cryptographically relevant quantum computer (CRQC). This calls for quantum-safe algorithms.

New algorithms are being developed to safeguard against attacks from conventional or quantum computers.

The algorithms are separated into the following categories:

1. Digital signature algorithms
2. Key encapsulation mechanisms and key-establishment algorithms

The primary algorithms NIST recommends to be implemented for most use cases are CRYSTALS-Kyber (key-establishment) and CRYSTALS-Dilithium (digital signatures).

IBM has created and adopted Quantum Safe Cryptographic algorithms. IBM is leveraging quantum-safe technologies and data from quantum computer attacks.

4.2 Quantum-safe capabilities with IBM Z

IBM LinuxONE 4 supports the quantum-safe secure boot feature (see section 2.3, “IBM Secure Boot for Linux” on page 30) which is protected through different firmware layers during the boot process. Only trusted IBM approved firmware is loaded.

This hardware-protected verification uses dual-signature scheme which use quantum safe and classical digital signatures.

The Crypto Express Hardware Security Module (HSM) also now uses a quantum-safe dual-signature scheme. IBM Crypto Express8S (CEX8S) includes implementations of the CRYSTALS-Dilithium and CRYSTALS-Kyber algorithms; the IBM Crypto Express7S (CEX7S) for IBM z16 and IBM z15™ includes CRYSTALS-Dilithium support.

Section 6.3.2, “Hyper Protect proposed solution architecture” on page 73, “Pervasive Encryption Pyramid for LinuxONE” on page 74 discusses the pervasive encryption functions that have been updated to use quantum-safe mechanisms for key management.

The Common Cryptographic Architecture (CCA) provides a variety of services for cryptography and data security. See 2.1.2, “IBM Crypto Express adapter” on page 21 as well as the PCIe Cryptographic Coprocessor product page (release CCA 8.1) for more information.

CCA provides API functions to support quantum-safe algorithms by using the CRYSTALS-Kyber Key Encapsulation Method (Kyber-KEM) and CRYSTALS-Dilithium Digital Signature Algorithm (CRDL-DSA). For more information, see PKA key algorithms.

IBM LinuxONE applications that use a PKCS #11 API can take advantage of the Enterprise PKCS #11 (EP11) coprocessor mode of IBM® cryptographic adapters. Section 2.1.2, “IBM Crypto Express adapter” on page 21 provides an overview of IBM Enterprise PKCS #11 (EP11) mode. You can also find more detail in Exploiting Enterprise PKCS #11 using openCryptoki - IBM Documentation and Quantum safe cryptography with the EP11 token.

PKCS#11 provide API functions to support quantum-safe cryptography signatures by using CRYSTALS-Dilithium. For more information on how to transition to quantum-safe cryptography, see the IBM Redbooks *Transitioning to Quantum-Safe Cryptography on IBM Z*, SG24-8525.

5

IBM Support for Hyperledger Fabric with IBM LinuxONE

The IBM Blockchain Platform (IBP) provides an easy way to create and deploy Hyperledger Fabric servers in the public cloud or in an on-premises data center.

In this chapter, we describe how an enterprise client can deploy IBM Blockchain Platform. We also and explain the security benefits of building a blockchain on LinuxONE.

This chapter includes the following topics:

5.1 Blockchain, Hyperledger, IBM Blockchain Platform, and IBM Support for Hyperledger Fabric

Well-known from its use for crypto currencies, such as Bitcoin, blockchain technology is a hot topic in the technology industry. The core components of blockchain technology are not new. However, as computing power grew alongside the expansion of IoT and mobile devices, adopting blockchain technology into applications became a feasible reality. Blockchain can provide a way for enterprises to share data among one another securely, but also apply permissions on who can create and view that data.

Blockchain technology provides a way to securely store data, most of the time as a metadata to point to more detailed database tables. The data that is stored in blockchain is distributed among participants' servers, which are also called *decentralized data stores*.

Data that is stored in blockchain is also inherently encrypted with strong cryptographic algorithms and key management procedures. When a blockchain stores data with these cryptographic procedures, it stores them with time stamps, which is sequence information that indicates the order of the data blocks, and hash information that is derived from a previous data block. As a result, it can prove its position in the blockchain that is based on its parent block.

This strong cryptographic process makes blockchain "immutable", which means that after data is stored, it stays in its original form and cannot be modified. Unlike other commonly used databases, blockchain normally does not allow inserts or modification of its data. For this reason, blockchain is a strong candidate for verified ledger applications between many distinct parties.

Large enterprises in industries, such as banking, insurance, supply chain and logistics, and healthcare are thinking about how they can adopt blockchain technology. IBM worked with these customers on requirements for blockchain technology. During that process, IBM discovered that public network blockchain protocols, such as Bitcoin and Ethereum, were not a good fit for enterprise use cases.

Hyperledger blockchain was created from the needs of the enterprise. It is a permission-based private blockchain network, so only allowed participants can see and access the data from the shared, distributed ledger (database).

Some of the more recent enhancements include:

Some of the more recent enhancements include:

- Hyperledger Fabric 2.x:
  - Fabric Gateway, a novel service operating on peer nodes, oversees transaction submission and processing for client applications.
  - Updated Gateway SDKs (v1.0.0) are now accessible for Node, Java, and Go.
  - A capability has been added to capture a snapshot of a peer's channel information, including its state database, and integrate new peers (within the same organization or across different organizations) to the channel using the snapshot.
  - Decentralized Smart Contract Lifecycle: Enhances governance and management of smart contracts.
  - External Chaincode Launcher: Eliminates dependency on Docker daemons.
  - Raft Consensus: Provides a new consensus mechanism improving fault tolerance.

  - Private Data Enhancements: Ensures data privacy within blockchain networks.
  - Performance Improvements: Optimizes the efficiency of blockchain operations.
- IBM Blockchain Platform and IBM Support for Hyperledger Fabric:
  - Support for Multicloud Environments: Expanded to support LinuxONE, IBM Z, and various cloud providers like AWS and Red Hat OpenShift. This flexibility helps meet data residency requirements and prevents vendor lock-in.
  - Enhanced Security: Integration with Hardware Security Modules (HSMs) and compliance with FIPS encryption standards. Regular security scans and vulnerability assessments are conducted to protect against threats.
  - Developer Tools: Introduction of Red Hat CodeReady Workspaces and Ansible Collections for streamlined development and deployment.
  - Fabric Operations Console: Simplifies the management of blockchain networks, providing a unified interface for monitoring and configuration.
- IBM Contributions to Open Source:
  - IBM has donated significant portions of its blockchain platform code to the Hyperledger community, enhancing standardization and interoperability across different blockchain networks.
- Enterprise Adoption:
  - The platform is widely adopted across industries such as banking, healthcare, and supply chain, leveraging blockchain to improve transparency, efficiency, and security in operations.

Figure 5-1 shows how each block of data on a blockchain network is stored in each peer node's database. A peer node is served as data processing host for the blockchain network. It validates and endorses a transaction that is requested from an application. Then, it ensures that the rules (business rules) that are described in chain code are checked when the transaction occurs.

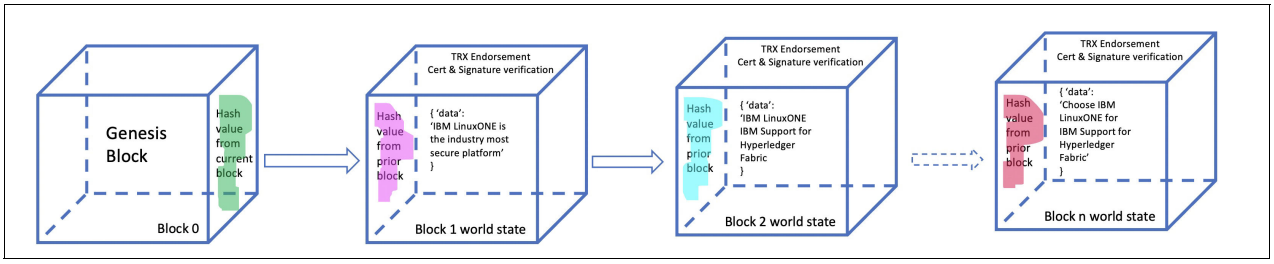

Figure 5-1 High-level concept of Hyperledger blockchain data creation

The peer also communicates with the orderer unit, where it keeps the sequence of transactions and distributes to proper peer nodes when a transaction is requested.

A certificate authority also is available that manages certificates for each member's entity to verify that it is genuine and to validate authenticity.

On a peer node, the blockchain always starts with a genesis block, which is an initial data block that includes a specially encoded hash value. Its hash value is also transferred to the next block in sequence as an input. The next block generates a new hash value from the previous block's hash and the payload of the new block's data; then, it repeats for the next blocks.

Keeping the previous block's hash and making a new hash from it makes blockchain data unable to be reversed or modified. When the contents of the data payload (world state) changes, blockchain changes the hash for the next block that was created and affects all of the world states behind the block that was modified. This idea is the key idea of blockchain data being immutable and it requires a high-speed hash process to generate new blocks.

Hyperledger also uses an asymmetric cryptographic function to create certificates and validate the signature of the transactions. The workflow of those asymmetric keys can be complicated in Hyperledger Fabric because it involves multiple members and different channels for various application uses.

If a person or an organization plans to develop Hyperledger applications, they must plan carefully where to host the service because it can use a large number of resources, depending on transaction sizes.

IBM LinuxONE is a great choice for a large enterprise Hyperledger Fabric service platform because it features the best cryptographic hardware in the industry. Also, it includes a feature called *On/Off Capacity on-demand*. You use it to enable extra CPU capacity to boost processing power in a few seconds for any need that is caused by spiking workloads.

5.1.1 IBM Blockchain Platform

IBM Blockchain Platform allows a consortium of organizations to effortlessly create and join blockchain networks, whether on-premises or across private, public, or hybrid multicloud environments using Kubernetes. Clients have the flexibility to deploy their nodes on their preferred cloud platform and seamlessly connect to any IBM Blockchain Platform network. This connectivity extends to networks hosted on either their own Kubernetes clusters or through IBM Blockchain Platform on IBM Cloud. Powered by Hyperledger Fabric 2.2.10, IBM Blockchain Platform 2.5.4 supports deployment on various Kubernetes distributions. A notable advantage of this platform is IBM's rigorous daily testing of the open-source code to identify and mitigate security vulnerabilities. Additionally, it offers around-the-clock, 365-days-a-year support with service level agreements (SLAs) tailored for production environments.

Note: The IBM Blockchain Platform Software Edition has been replaced with IBM Support for Hyperledger Fabric. As of April 30, 2023, IBM Blockchain Platform Software Edition is no longer under support. To migrate from IBP to IBM Support for Hyperledger Fabric, see section 5.1.3, "Migrating from IBM Blockchain Platform (IBP) to IBM Support for Hyperledger Fabric" on page 58.

5.1.2 IBM Support for Hyperledger Fabric

IBM Support for Hyperledger Fabric offers certified images of Hyperledger Fabric, the Fabric Operations Console, and a Kubernetes Operator, facilitating blockchain component deployment across diverse platforms, including popular Kubernetes distributions like Rancher and OpenShift Container Platform. This solution caters to customers seeking to deploy, store, and manage blockchain assets across varied environments, whether on-premises or across public and private clouds, driven by security, risk management, or compliance needs. Key features include:

- 24x7 Break-fix Support: Ensures continuous operation by providing around-the-clock assistance for critical Hyperledger Fabric issues.
- Security Enhancement: Proactively strengthens security and compliance through regular vulnerability scans.

- Expertise and Updates: Access to the latest fixes and automated regression testing via community updates.

IBM Hyperledger Fabric support also includes a flexible management platform running on Kubernetes, encompassing essential components such as operators, management consoles, peers, certification authorities (CAs), ordering nodes, and smart contract containers. This solution supports Fabric versions 2.2.15 or v2.5.8 and is compatible with Kubernetes distributions like Red Hat OpenShift Container Platform on IBM LinuxONE 4.12, 4.13, 4.14 and 4.15.

The IBM Support for Hyperledger Fabric is a Hyperledger Fabric blockchain protocol as-a-service that has following properties:

- Modularity

  Blockchain networks must incorporate a wide range of new and existing “pluggable” features, depending on the enterprise and industry. As a result, Hyperledger Fabric was developed to be modular to support networks as new features emerge. Modularity in Hyperledger Fabric allows the IIBM Support for Hyperledger Fabric to use industry-leading security practices to serve production-ready networks, including GDPR and HIPAA best practices.

- Scalability

  Organizations across industry sectors demand solutions that scale as they move past initial explorations and proof-of-concepts.

  Hyperledger Fabric was built to support growing business networks, which must dynamically add participants and support. These additions increase the amounts of transaction processing. Many aspects of scalability depend on the network configuration of consensus, membership, and security. IBM Support for Hyperledger Fabric uses Hyperledger Fabric to provide a modular platform that supports the ability to configure a network to support the throughput numbers and network growth that are required.

- Consensus

  An important feature in the security, scalability, and maturity of any blockchain framework is a clearly defined and implemented consensus protocol. Consensus in Hyperledger Fabric is pluggable and fit specific enterprise use cases.

  Therefore, Hyperledger Fabric allows you to choose the best consensus protocol to fit your specific business networks’ needs. Hyperledger Fabric’s success to date is driven by the massive amount of community support it received through Hyperledger. Open governance of the code base with a clear purpose allowed it to emerge as the industry-leading protocol and framework for enterprise production networks.

Running the IBM Support for Hyperledger Fabric components outside of IBM Cloud provides you with more flexibility to grow or join a blockchain network. It also helps network initiators grow their networks by allowing new members to join while they use the platform of their choice. It allows organizations that are interested in joining blockchain networks to collocate their peers with their existing applications or to integrate with their systems of record.

Figure 5-1 on page 55 shows an overview of IBM Support for Hyperledger Fabric.

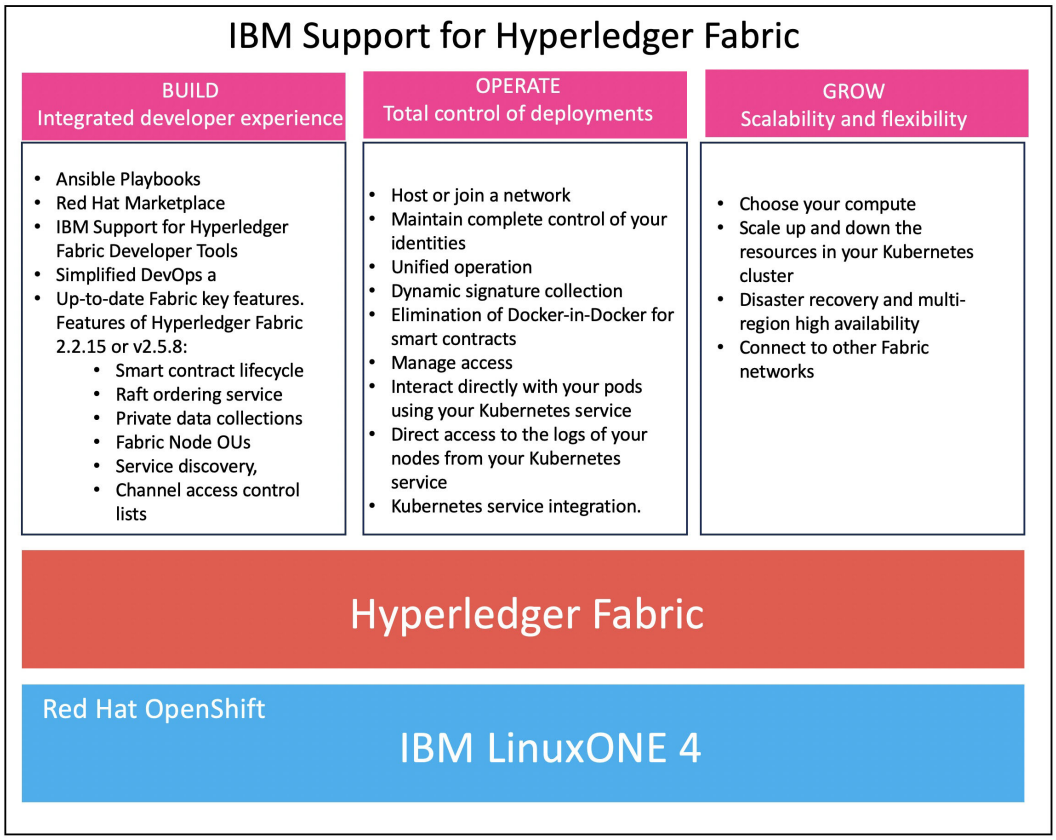

Figure 5-2 IBM Support for Hyperledger Fabric Overview

IBM Support for Hyperledger Fabric offers a robust and scalable platform designed to facilitate rapid and secure deployment of applications and data. This guide focuses on securing your IBM Support for Hyperledger Fabric service instance, where the blockchain console operates, and outlines best practices for securing the Kubernetes cluster hosting the blockchain nodes. To enhance security, IBM leverages built-in features of IBM LinuxONE, including:

- Hardware Security Module (HSM) Configuration: Users can configure a Hardware Security Module (HSM) to generate and securely store private keys for nodes (using LinuxONE CPACF and Crypto Express security hardware features)
- Encryption for **`etcd`** Data: Implementing encryption for **`etcd`** data stored on the local disk of the Kubernetes master enhances security at the datastore layer of your cluster. This measure safeguards sensitive information against unauthorized access or tampering.

5.1.3 Migrating from IBM Blockchain Platform (IBP) to IBM Support for Hyperledger Fabric

While IBM internally tests the migration tool for production use, there are essential steps you should take to prepare for migrating from IBP to IBM Support for Hyperledger Fabric:

- Upgrade Hyperledger Fabric Components:
  - Ensure any blockchain components using Hyperledger Fabric 1.4 are upgraded.
  - Upgrade to at least Hyperledger Fabric v2.2 before March 31, 2023.

- Check Kubernetes Version Compatibility:
  - Verify your existing IBM Kubernetes Service is at 1.23 or 1.24.
  - For networks on OpenShift Container Platform (OCP), ensure 4.9, 4.10, or 4.11 are used.
- Verify Component Prerequisites:
  - Ensure Hyperledger Fabric v2.2 components meet the following prerequisites:
    - Peers: 2.2.10 or higher
    - Ordering nodes: 2.2.10 or higher
    - Certificate authorities: 1.5.5 or higher
  - For Hyperledger Fabric 2.4 components, ensure:
    - Peers: 2.4.8 or higher
    - Ordering nodes: 2.4.8 or higher
    - Certificate authorities: 1.5.5 or higher
- Update IBM Blockchain Platform Component Image:
  - Change the IBM Blockchain Platform component image to the equivalent IBM Support for Hyperledger Fabric image. This change will briefly interrupt normal traffic flow when IBM Blockchain Platform components restart with the new images.
  - Note that no components or data are moved or migrated from their current location during this process.

For more information, see Migrating to IBM Support for Hyperledger Fabric.

5.2 IBM Support for HyperLedger Fabric for LinuxONE

To run a Hyperledger application for your enterprise, or even for personal projects, you need to have at least one Hyperledger Fabric Network running somewhere. It is an architecture form that consists of multiple servers and services. They require many different access controls, permissions, and business rules. In the Hyplerledger world, these business rules are called *Smart Contracts*, and it is represented in the form of *Chain Code* (“chain” as in “blockchain”).

You can choose to download the open-source version of Hyperledger codes to your personal computer or your choice of servers. If you are an enterprise blockchain application developer or architect, you need at least have one Hyperledger Fabric Network running across your team to develop and test more extensively.

When the business applications interact with other business entities, the Hyperledger Fabric must serve those other entities over the internet and communicate to handle the chaincode’s requests. This task often is not a simple. Although you are running multiple servers, they must always be up and redundant to keep the data replicated and consistent among participating members of the blockchain network.

5.2.1 Red Hat OpenShift Container Platform

Red Hat OpenShift Container Platform (OCP) is a reliable and scalable cloud platform that can run on your on-premises infrastructure. It is built on open-source frameworks, such as containers and Kubernetes. In addition, it offers common services for self-service deployment, monitoring, logging, and security, and a portfolio of middleware, data, and analytics with IBM Cloud Pak® solutions.

OCP on LinuxONE brings all of the benefits of reliability and scalability of Red Hat OpenShift and provides excellent security without any changes in the application or containers. It also can scale vertically and horizontally, which helps process surges in workloads that occur because of any business environment changes (for example, a stock market surge that is the result of a catastrophic event).

For more information, see IBM LinuxONE.

5.2.2 How IBM Support for Hyperledger Fabric for LinuxONE can use CPACF and Crypto Express security hardware features

IBM Support for Hyperledger Fabric can use the cryptographic hardware that is built into LinuxONE. Starting with CP Assisted Cryptographic Functions (CPACF), when new blockchain data is generated, it uses a HASH function that is called (SHA-256) in CPACF. It then encrypts the data block in the file system with an AES block cipher, which also is in CPACF. IBM Support for Hyperledger Fabric automatically uses these hardware features when used with supported Kubernetes platforms for LinuxONE.

LinuxONE can also be equipped with the Crypto Express co-processor feature, which is an HSM that is plugged into the PCI slot on the server. It meets the highest level of FIPS 140-2 standard, a level-4.

When equipped with the Crypto Express co-processor, IBM Support for Hyperledger Fabric can use protected key encryption for file systems. In this case, the encryption key for the block ciphers is always protected with another “key-encrypting-key” in the HSM. This function dramatically reduces the chance of exposing the encryption keys because it is protected by an HSM in the IBM Cloud and on IBM LinuxONE. In contrast, most of the other public cloud solutions store keys in memory as clear text.

The Crypto Express co-processor also helps accelerate blockchain processing times. It contains custom-built ASICs that accelerate asymmetric cryptographic functions, such as RSA or Elliptic Curve Cryptography (ECC). When encryption and authentication are required for blockchain processes, Crypto Express hardware can handle them much faster, compared to the use of only CPU cycles for the computations.

Hyperledger Fabric uses Elliptic Curve Digital Signature Authentication (ECDSA) for certificate authority enrollment and for encrypting blocks of data. Therefore, when you use ECDSA acceleration through an ECC unit in CPACF (LinuxONE III and LinuxONE 4), you can process more data in a shorter time. Its TCP/IP communications are also SSL/TLS enabled. Therefore, having RSA digital signature generation and verification helps to serve SSL/TLS traffic better with a Crypto Express coprocessor.

5.2.3 Secure Service Container

IBM developed its Secure Service Container (SSC) technology, which is exclusive to IBM LinuxONE, to provide an easy-to-deploy secure hosting appliance for container-based applications that run in hybrid cloud environments. SSC is a secure computing environment for microservices-based applications that can be deployed without any application code changes, which makes it an easily consumable solution for cloud-native development. It provides several unmatched security benefits, such as automatic pervasive encryption of data in-flight and at-rest, protection from privileged administrators, and tamper protection during installation and start time to protect against malware.

The Secure Service Container (SSC) on IBM LinuxONE enhance the security and integrity of blockchain networks in several ways:

- SSC is designed to provide a highly secure environment by isolating blockchain components from the rest of the system. It ensures that sensitive cryptographic operations and key management are protected against unauthorized access or tampering. This isolation helps in maintaining the confidentiality and integrity of blockchain transactions and data.
- Blockchain networks often deal with valuable digital assets and sensitive information. SSC on LinuxONE is designed to enhance protection by leveraging hardware-based security features such as cryptographic coprocessors (CPACF), Hardware Security Modules (HSMs), and Trusted Platform Modules (TPMs). These features provide strong safeguards against various forms of cyber threats, including attacks targeting cryptographic keys and data integrity.
- Enterprises operating blockchain networks must adhere to regulatory compliance requirements and demonstrate trustworthiness to stakeholders. The SSC supports these efforts by offering a secure execution environment.

6

Use cases

In this chapter, we explore examples of encryption use cases that use the security features of IBM LinuxONE. The chapter includes more technical details to help you understand of what occurs “under the hood” of the LinuxONE system.

This chapter includes the following topics:

6.1 Containers and data encryption use case

This chapter describes how you might start to use LinuxONE hardware cryptographic features within a Docker environment and running on LinuxONE. The following sections describe how to benefit from CPACF hardware acceleration to secure client/server communications through an OpenSSL example:

- "Context and challenges"
- "Solution"
- "Implementation" on page 65
- "Summary" on page 69

6.1.1 Context and challenges

We are in an era where container technologies are gaining more popularity. The companies are increasingly adopting this virtualization technology to build and scale cloud-native applications. Undoubtedly, containers provide the flexibility and rapidity to pursue business agility for speed-to-market. This fact is understood and adopted by companies to improve their efficiency and competitiveness.

Docker is one of the most popular container platforms, as it is adopted by companies across all industries to start the journey of cloud transformation. However, this new world of container infrastructure brings many questions about security. Again, IT decision-makers, IT architects, and developers must address all the security concerns that might arise during each phase of containers lifecycle. That way, they ensure the success of the containerization strategy within the company. Some of these security concerns were addressed by different solutions.

If we look at image authenticity, it is possible to use a cryptographic signature mechanism: the Docker Content Trust and Notary functions. This mechanism ensures that a container image is not processed or modified by an unauthorized third party.

If we look at container isolation level, some Linux features, such as namespaces, control groups, AppArmor, or SELinux, can help to keep the containers isolated from each other.

But what about protecting container data? This security concern is a key requirement for companies to comply with regulatory standards when they use containers. Protecting sensitive data in Docker containers can be addressed with encryption. However, this issue again brings us back to the challenges related to data encryption (performance overhead, application changes, and so on). Thus, how can LinuxONE help clients to overcome these challenges that come with the use of containers?

6.1.2 Solution

The first step toward a secure container system is to have a secure host environment with advanced security capabilities to run the entire set of containerized applications.

To help our clients with their journey toward secure hybrid cloud, LinuxONE extends its security features to containers. Thus, containers can use LinuxONE cryptographic hardware features. Encrypting data applications on containers is simplified and the encryption overhead is significantly reduced after the containers are configured to use CPACF or cryptographic coprocessor cards.

So that containers can use CPACF for hardware accelerated encryption on every core, you must enable access to CPACF in the Docker host where Docker images are deployed. If the Docker host can access CPACF, the Docker containers that run on this host automatically can access CPACF without extra configuration.

The same approach is used for the cryptographic coprocessor cards (Crypto Express adapters) approach. The Docker host must access the cryptographic cards in addition to loading the `zcrypt` device driver. Then, the container images can use the cryptographic cards through the `/dev/z90crypt` device node.

In this document, we use the encryption of data in-flight with OpenSSL as an example to show how to enable and use CPACF function with containers. Figure 6-1 shows the components that are part of the solution for encrypting data in flight with OpenSSL based on CPACF.

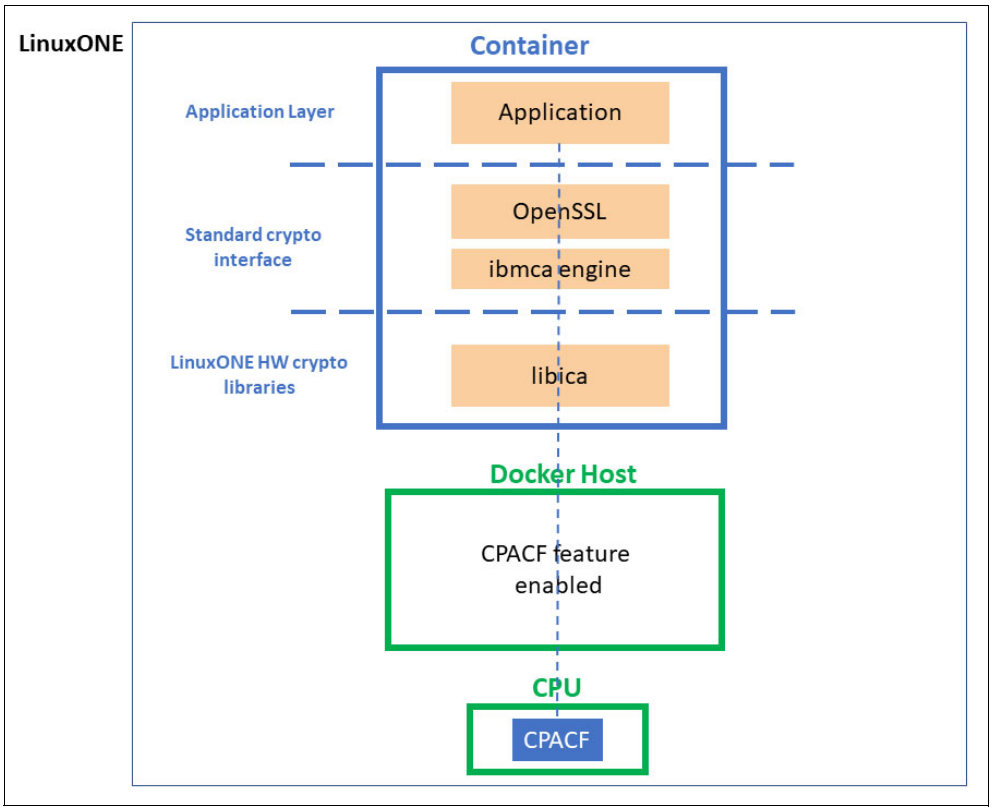

Figure 6-1 Components used for data in flight encryption with OpenSSL

Figure 6-1 also shows an example of containers that are backed by LinuxONE hardware-cryptographic acceleration with CPACF, and the components that are used for data in flight encryption with OpenSSL.

6.1.3 Implementation

In this section, we complete the following steps to enable a Docker container to use CPACF hardware encryption acceleration for OpenSSL.

1. “Checking the CPACF ennoblement in Docker host” on page 66.
2. “Installing required packages in the container” on page 66.
3. “Configuring OpenSSL” on page 67.

For this demonstration, we use a CentOS Linux server release 7.6 for Docker host and container, running on a z/VM LPAR. The Docker image that we use is an image that supports s390x architecture. The `s390x` is used as a suffix by the image providers to specify that the Docker images are compiled for LinuxONE architecture.

Note: For more information about the different options to acquire Docker images for LinuxONE, see IBM Knowledge Center.

Checking the CPACF ennoblement in Docker host

CPACF can be used in your environment if the Licensed Internal Code (LIC) feature 3863 is installed in your LinuxONE. This feature is available at no extra charge. You can confirm that this cryptographic feature is enabled by going directly to the HMC console in one of the following ways:

- Follow the steps that are described in the “Verification of installed LIC 3863 using the SE” section of the IBM Redbooks publication, *Security and Linux on z Systems*, REDP-5464.
- Run the following command in the Docker host:

```
[root@openshiftmaster ~]# cat /proc/cpuinfo | grep features
```

If the `features` list contains **`msa`**, as shown in Example 6-1, the CPACF feature is enabled in the LinuxONE central processors.

Example 6-1 Resultant features list

```
features        : esan3 zarch stfle msa ldisp eimm dfp edat etf3eh highgprs te vx
vxd vxe gs sie
```

Installing required packages in the container

Now that it is confirmed that CPACF is enabled in the Docker host, we know that it is also available automatically to the container. We switch to the Docker container to allow OpenSSL to use CPACF hardware encryption acceleration. For this purpose, we install the following packages, as described next:

- `libica`
- `openssl`
- `openssl-ibmca`

Installing libica library

As shown in Figure 6-1 on page 65, the `libica` library ensures communication between CPACF and OpensSSL. It contains CPACF interfaces that allow applications to use CPACF. Complete the following steps to install the `libica` library:

1. Connect to the Docker container with root user by using the following command, where `2371cdc247ee` is the ID of the container image:

   ```
   [root@openshiftmaster ~]# docker exec -it --user root 2371cdc247ee bash
   bash-4.2#
   ```

2. Install the `libica` library by using the following command:

   ```
   bash-4.2# yum install libica-utils
   ```

The installed `libica` library provides the **`icainfo`** command. You can use this command to display the list of encryption algorithms that are supported by the hardware and therefore available to the container.

Installing OpenSSL library

OpenSSL is an open-source cryptographic library that provides an implementation of the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols. These protocols secure communication between two parties: client and server. Different applications rely on OpenSSL to perform the encryption requests.

To install the OpenSSL library in your environment, use the following command:

```
bash-4.2# yum install openssl
```

Installing the openssl-ibmca library

As shown in Figure 6-1 on page 65, OpenSSL needs the `ibmca` engine to communicate with the `ibmca` library. In this case, the encryption requests are transferred to the `ibmca` engine and not processed directly by OpenSSL. Then, the `ibmca` library communicates with CPACF to allow applications to start using hardware-based cryptographic acceleration.

Install the `openssl-ibmca` library by using the following command:

```
bash-4.2# yum install openssl-ibmca
```

Configuring OpenSSL

Now that all required packages are installed, OpenSSL must be prepared to use the `ibmca` engine. This step is the last step to enable CPACF hardware acceleration of cryptographic functions in OpenSSL.

The `Openssl-ibmca` package contains an `openssl_cnf` configuration file that we use to configure OpenSSL to use the `ibmca` engine. Complete the following steps:

1. Locate the `openssl_cnf` file:

   ```
   bash-4.2# find / -name openssl.cnf /etc/pki/tls/openssl.cnf
   ```

2. Make a copy of the `openssl_cnf` file:

   ```
   bash-4.2# cp -p /etc/pki/tls/openssl.cnf /etc/pki/tls/openssl.cnf.v0
   ```

3. Locate the `openssl.cnf.sample.s390x` file. This file is in the `openssl-ibmca` package. It allows the loading of the `ibmca` engine for all the applications that feature OpenSSL support:

   ```
   bash-4.2# find / -name openssl.cnf.sample.s390x
   ```

4. Add the content of the `openssl.cnf.sample.s390x` file to the `openssl_cnf` file:

   ```
   bash-4.2# tee -a /etc/pki/tls/openssl.cnf <
   /usr/share/doc/openssl-ibmca-1.4.1/openssl.cnf.sample.s390x
   ```

5. Insert a reference to the `ibmca` engine in the `openssl_cnf` file by adding the line `openssl_conf = openssl_def` under `RANDFILE = $ENV::HOME/.rnd` line, as shown in the following command:

   ```
   bash-4.2# vi /etc/pki/tls/openssl.cnf
   # Using the following parameters prevents the configuration file from hanging
   if HOME isn't defined.
   HOME = .
   RANDFILE = $ENV::HOME/.rnd
   openssl_conf = openssl_def
   ```

6. Confirm that the support of `ibmca` engine is enabled for OpenSSL. The `ibmca` engine must be listed in the result of the following command, as shown in the right column:

```
bash-4.2# openssl engine
(dynamic) Dynamic engine loading support
(ibmca) Ibmca hardware engine support
```

Testing the CPACF hardware encryption acceleration in the container

To verify that your container was properly configured to use the CPACF hardware encryption acceleration with OpenSSL, run the following command:

```
bash-4.2# openssl speed sha1 -elapsed | tail -n 3
```

The results of that command are shown in Example 6-2.

Example 6-2 Verify that container is configured to use OpenSSL with sha1 encryption algorithm

```
You have chosen to measure elapsed time instead of user CPU time.
Doing sha1 for 3s on 16 size blocks: 6411336 sha1's in 3.00s
Doing sha1 for 3s on 64 size blocks: 5656337 sha1's in 3.00s
Doing sha1 for 3s on 256 size blocks: 4747020 sha1's in 3.00s
Doing sha1 for 3s on 1024 size blocks: 2938084 sha1's in 3.00s
Doing sha1 for 3s on 8192 size blocks: 612107 sha1's in 3.00s
The 'numbers' are in 1000s of bytes per second processed.
type             16 bytes     64 bytes    256 bytes   1024 bytes   8192 bytes
sha1             34193.79k   120668.52k   405079.04k  1002866.01k  1671460.18k
```

To verify that the OpenSSL cryptographic calls are using CPACF, run the following command:

```
bash-4.2# icastats
```

Typical results of this command are shown in Example 6-3.

Example 6-3 Results of icastats command

```
 function     |           hardware       |            software
--------------+--------------------------+-------------------------
              |      ENC    CRYPT   DEC  |      ENC    CRYPT   DEC
--------------+--------------------------+-------------------------
        SHA-1 |        34525344          |                0
      SHA-224 |               8          |                0
      SHA-256 |               8          |                0
      SHA-384 |               8          |                0
      SHA-512 |               8          |                0
    SHA3-224  |               0          |                0
    SHA3-256  |               0          |                0
    SHA3-384  |               0          |                0
    SHA3-512  |               0          |                0
   SHAKE-128  |               0          |                0
   SHAKE-256  |               0          |                0
```

Notice that the hardware column in Example 6-3 shows that the OpenSSL cryptographic calls are using CPACF.

6.1.4 Summary

In addition to the advantages related to scalability and performance, LinuxONE provides your containerized applications with a secure host environment. By running your containers on LinuxONE, you can benefit from its hardware security features to secure and accelerate the encryption of your data. One of the main advantages of the use of these hardware security features is encrypting your data without modifying your applications and with minimal performance overhead.

The use of OpenSSL for data in-flight encryption is one of the various examples of what you can implement on LinuxONE. You also can use other Linux libraries, such as `dm-crypt`, to encrypt your data at-rest while you continue to benefit from the hardware cryptographic acceleration with CPACF or cryptographic cards.

With its differentiating security capabilities, LinuxONE is a great starting point in your transition to cloud-native applications. LinuxONE also can help you meet compliance requirements for protecting sensitive data in a container system.

6.2 Database and volume encryption use case

This section describes how you can use LinuxONE hardware cryptographic acceleration capabilities features to address the challenges that are related to protection of sensitive data that is stored in databases. This approach is illustrated through a database volume encryption use case with `dm-crypt` LUKS- and CPACF-protected keys to protect data at-rest.

6.2.1 Context and challenges

Databases contain different types of sensitive information, including the following examples:

- Personally Identifiable Information (PII), which is related to data that can help to identify a specific individual.
- Business Information that includes data that represents a risk to the company if disclosed to competitors or the general public.
- Classified Information that refers to top-secret government data.

All these types of sensitive data must be protected, especially in the light of data breaches that are affecting more businesses. Different regulations and laws exist to ensure that this sensitive data remains in good hands. Otherwise, companies can incur heavy financial repercussions.

All these factors make the protection of data that is stored in databases more urgent and important in our century. It becomes one of the highest priorities for companies around the world.

This context brings back a significant focus on data-at-rest encryption. This focus addresses some of the data security issues and complies with various regulations, such as Payment Card Industry Data Security Standard (PCI DSS) or the Health Insurance Portability and Accountability Act (HIPAA).

Consider the following scenario. Your company uses a database that stores sensitive information. To meet the data protection requirements, you must encrypt all data in-flight and at-rest. How you can meet this requirement without altering your applications' response times, increasing your costs, or suffering from lack of skills in your organization?

For data in-flight, these issues can be easily addressed by encrypting your data based on SSL encryption by using the OpenSSL Linux library, for example. On LinuxONE, OpenSSL can benefit from the hardware cryptographic acceleration with CPACF or cryptographic cards. Thus, you can transparently encrypt connection by the database server, in addition to securing the Linux sessions through SSH, for example.

Regarding the protection of data at rest, one of the most common use cases is the use of Filesystem-level encryption, which encrypts data before it is written out to physical volumes. However, this option can expose you to the same issues (costs, performance impact, and so on). Thus, how can LinuxONE help you to minimize the effect of encrypting your data at-rest?

6.2.2 Solution

Encryption at the file system level can be implemented by using `dm-crypt` LUKS, which is an encryption subsystem that is included in the Linux Kernel. With this transparent encryption mechanism, you can encrypt disks, software RAID volumes, partitions, and logical volumes to protect your data at rest. It allows the encryption of all data that is written to disk and decrypts all data that is read from disk. The data appears only in the clear in the application.

Note: The `dm-crypt` provides plain format volumes or Linux Unified Key Setup (LUKS) volumes. In plain mode, `dm-crypt` does not add any headers or metadata to the volumes. The `dm-crypt` LUKS adds a metadata header to the encrypted volume data, and therefore offers more features than plain mode.

LUKS2 format is the preferred option for LinuxONE data at-rest encryption. For more information, see section 3.9.3, “Volume format considerations” in *Getting Started with Linux on Z Encryption for Data At-Rest*, SG24-8436.

By implementing this encryption option on other platforms, you might see increased performance overhead, which is always involved when you use software-encryption mode. On LinuxONE, `dm-crypt` LUKS can use the CPACF on-chip encryption co-processor to remove this overhead.

LinuxONE also provides a unique security enhancement for `dm-crypt`, which consists of the use of the CPACF protected keys support to resolve the security issues that are related to storing keys. If `dm-crypt` is used with clear key, keys are readable if the memory system is dumped. With a CPACF protected key, no encryption key is stored in clear in the operating system.

The `dm-crypt` cryptographic operations are performed by using CPACF. CPACF code generates the wrapping keys. These keys are unique to each LPAR, and they are stored in the Hardware System Area (HSA). HSA is a reserved memory area that is separated from client-purchased memory that is used for internal system functions. HSA also is accessible through only the firmware.

As shown in Figure 6-2 on page 71, at the cryptographic card level, the secure key is encrypted by the master key. This secure key is used as the source key of the protected key. It is decrypted and sent to CPACF in clear text. Then, at CPACF level, the key is wrapped by the CPACF wrapping key, and the protected value is stored in memory. As a result, protected keys can never be seen in plain text by the operating system or by applications. With each encryption or decryption operation, the protected keys are sent to CPACF to be unwrapped by the CPACF wrapping keys.

Figure 6-2 shows an overview of data-at-rest encryption that uses `dm-crypt` on LinuxONE. It also shows that `dm-crypt` is located between the file system and the physical disks, which allows data encryption to happen when they are written to physical disks.

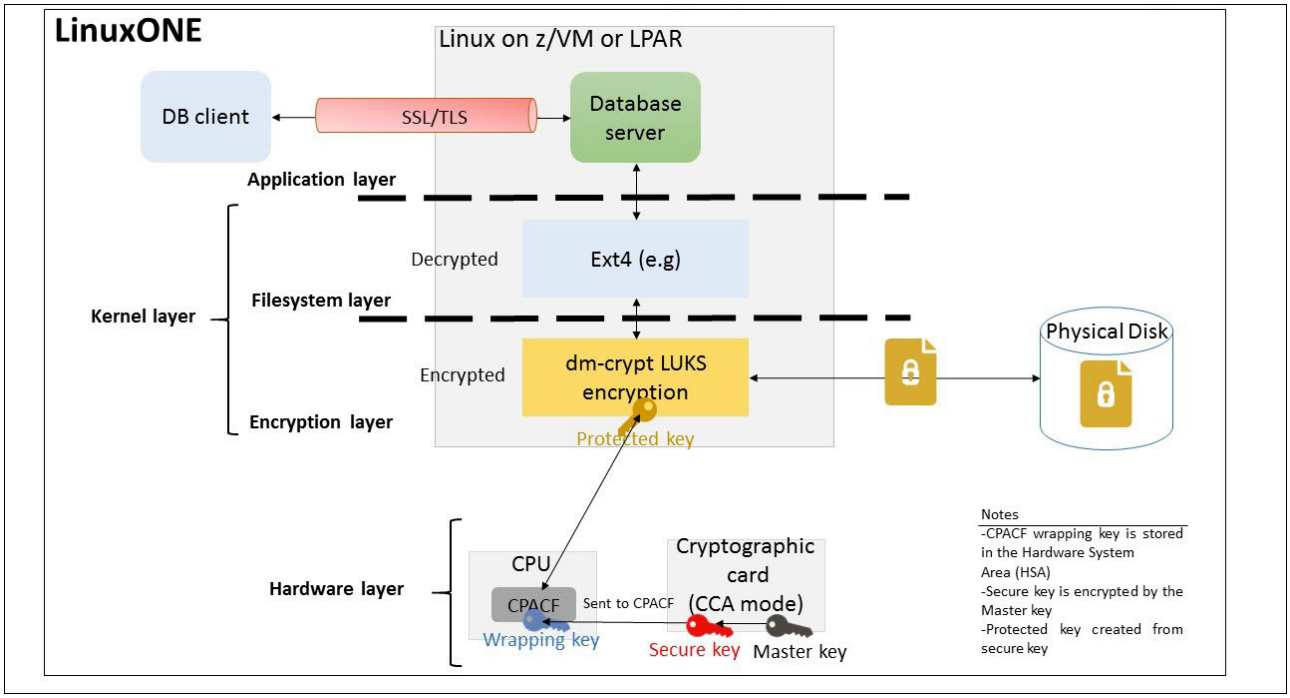

Figure 6-2 Solution overview: LinuxONE hardware-cryptographic acceleration with dm-crypt

Figure 6-2 also shows the key wrapping process that is used to provide `dm-crypt` with CPACF protected keys.

6.2.3 Getting started

The following steps refer to the main steps that compose the data at-rest encryption process that uses `dm-crypt` LUKS with CPACF protected keys.

1. Check CPACF enablement in your system. Licensed Internal Code (LIC) feature 3863 must be installed.
2. Check the availability of Crypto Express adapters (use CCA mode), and add the domains that are needed to hold the master key and secure key. (This process can be done by using the HMC console.)
3. Load the master key into the Crypto Express adapter domain.
4. Install the `pkey` kernel module for protected key management.
5. Use the `zkey` utility to generate and manage secure key.
6. Set up `dm-crypt` as usual (install the `cryptsetup` package).
7. Create the encrypted volumes and set up the `dm-crypt` LUKS2 header in the volumes.

For more information about these steps, see *Getting Started with Linux on Z Encryption for Data At-Rest*, SG24-8436.

6.2.4 Summary

As an enterprise data-serving platform, LinuxONE provides a secure environment to protect the sensitive data that is stored in your databases. Protecting data at rest is one of the important components in the Data Security journey. To help clients with the different challenges that are related to the implementation of data-at-rest encryption, LinuxONE offers hardware cryptographic acceleration capabilities with CPACF.

Encryption of database volumes with `dm-crypt` LUKS is one of the solutions that can be implemented on LinuxONE to protect your data at rest. That type of protection is one of the most common use cases. Because this mechanism works at the kernel level, it can improve security in the following fundamental areas:

- Performance: You avoid the overhead that is associated with traditional encryption.
- Encryption: You better protect the encryption keys against eventual attacks by using the protected keys that CPACF provides.

6.3 Hyper Protect Digital Asset Platform

In the early days of Bitcoin's introduction, crypto currency and blockchain technology were recognized as somewhat interchangeable. As the industry and community adopted crypto-currency and understood its uses better, it is clear that the two technologies serve different purposes.

Blockchain on the enterprise side was adopted as a distributed ledger; for example, Linux Foundation’s Hyperledger project.

As individuals, private sectors and governments acknowledged this type of alternative currency likely is to remain. Also, many architectures and practical implementations were proposed to protect digital assets, such as crypto currencies.

IBM recently announced Hyper Protect Digital Assets Platform, which uses the strong security features of IBM LinuxONE to protect those digital assets against any intrusions.

6.3.1 The Importance of Digital Assets

Bitcoin was developed by Satoshi Nakamoto as a form of peer-to-peer electronic cash transactions. He saw the problems of current (at the time) electronic payment systems and wanted to solve the issue with his invention of a peer-to-peer distributed transaction system, without any central authority involved.

Since Bitcoin’s genesis block was created over 10 years ago, many other crypto currencies were introduced, and disappeared. Many crypto currency-related businesses were founded, and one popular area was crypto currency exchanges in which central bank-managed currency can be converts into a crypto currency and vice versa.

Some instances exist of crypto wallets being stolen, or even an owner of the crypto currency exchange passing away while no one else knew the master encryption key for the crypto-currency asset repository.

Now, imagine what similar situations can happen to you even if you do not own any crypto currencies. If you use any coffee franchise application, they might feature their own rechargeable wallets in your local currency value, or their own points value translated from the local currency.

If something happens to their system that stores your money or points and loses them, there is almost nothing your central or local bank can do because it is out of the governed and regulated system. After you purchase points, it is a completed transaction from the store's perspective.

As another example, we review application-based money transfers. More banks allow people to transfer money to other banks by using their own franchised money transfer applications. You need only the recipient's email or phone number but no bank routing and account number to make the wire transfer. Some countries even use their wireless phone company to manage money transfers and make payments, even bypassing banks.

Many online payment companies enable the use of their application to charge central currency into their central system and make payments or transfer to anyone that uses the same application. You can think of all of these applications as *digital assets*, where a valued asset is stored in digital format and can be moved from one owner to another easily and quickly without a traditional central agency being involved.

Digital assets are also not only currency but anything that represents ownership of physical assets someone or organization can own legally, such as property titles or deeds that are transformed into digital form (that might contain digital signatures), or even intellectual properties that do not feature any associated physical properties.

If an incident (including cyber attacks) occurs to a system that stores these digital assets, such as deeds, it can be catastrophic if no other way method exists to restore proof of ownerships or the property.

Many countries, including the US, are making efforts to understand how Central Bank Digital Currency (CBDC) can work for them and how to properly prepare for the next generation of currency exchanges and safeguarded infrastructures. If or when a central bank announces a plan for digital currency, it might greatly affect private banks and accompanied payment systems. Because it has a potential to adopt Satoshi Natakmoto's idea of peer-to-peer payment system, it is critical to understand how a digital asset exchange enterprise can ensure that their infrastructure is safe and secure.

IBM's Hyper Protect Digital Asset Platform can provide the highest level of security and safeguards for the services handling those digital assets. Next, we describe how the proposed architecture with Hyper Protect Digital Asset Platform works.

6.3.2 Hyper Protect proposed solution architecture

In this section, we examine a proposed solution architecture.

Root of Trust: Keeping the master key secure

As described Chapter 2, "Core security technologies on LinuxONE" on page 17 and Chapter 3, "Users of security on LinuxONE" on page 35, IBM LinuxONE offers a strong hardware security module (HSM) by a Crypto Express adapter (CEX).

By using the CEX, Hyper Protect Crypto Services enables you to keep and manage operational encryption keys secure. Those keys are encrypted with your own master key that is visible only to you and never leaves the hardware (HSM, CEX in LinuxONE, and so on).

The master keys can be backed up using smart cards through another secured service that is called Trusted Key Entry (TKS) as described in Chapter 3, "Users of security on LinuxONE" on page 35. This master key management with encrypted operational keys provides the "root of trust" to make Hyper Protect Digital Asset Platform most secure.

Providing secure application hosting server: Vertical security

Now, the digital asset management application (we call it a *digital wallet application* or *hot wallet* here) runs on the LinuxONE Hyper Protect Virtual Server. This server starts with tamper resistant secure start process, and the storage that is accessed by the virtual server is encrypted with a protected key. This protected key is decrypted only in a special memory area of LinuxONE. When access to the encrypted disk is needed, encryption or decryption of data that is transferring in or out of the disk is run.

After the server is up, it deploys a container image that is also encrypted and signed by the Secure Build process. During the build process of the secure Docker (or equivalent OCI-compliant container) images, many security checks and endorsement are performed to ensure no malicious code is being implanted, and the application code that is saved in the container is a legitimate application from the software vendor (in this case, a digital wallet application).

If any reason exists to capture the memory dump from Hyper Protect Virtual Server for debugging purposes, the memory dump also is encrypted. Therefore, you must use a private key to access the encryption key for the memory dump. This feature is a key differentiator for LinuxONE compared to other platforms.

Communicating between container images by using secure channel: Horizontal security

Because Hyper Protect Virtual Server can limit the login access to an operating system shell, the server communicates externally only with APIs. Secure communications channels also can exist between application containers that use secure protocols, such as mTLS.

The asymmetric key pairs for applications and APIs are managed by EP11 over gRPC (GREP11) where the master key for EP11 service also comes from the HSM (Crypto Express card), which provides maximum protection of asset transfers when a digital asset is transferred by using the user's private/public key pair.

If the digital wallet must be a cold wallet (that is, a wallet that is offline and used for storing crypto-currencies), the master key that is stored in HSM is zeroed out to make all the access to the assets inaccessible. The cold wallet can be accessed again after the master key is restored with key restore ceremony.

Pervasive Encryption Pyramid for LinuxONE

Either from external or internal threats, the security of applications must be protected from all possible accessible threats by bad actors. Hyper Protect Digital Asset Platform architecture provides maximum security protection of every layer that any security professional would want to adopt.

When IBM LinuxONE was designed, it was proposed to achieve maximum security with an architecture called Pervasive Encryption. Pervasive Encryption addresses protection from many threats by encrypting data in-transit and at-rest based on FIPS 140-2 Level 4 HSM module. Hyper Protect Digital Asset Platform architecture is a good example of how enterprise applications can achieve Pervasive Encryption with LinuxONE.

Figure 6-3 shows how the Pervasive Encryption Pyramid of IBM LinuxONE is matched with Hyper Protect Digital Asset Platform. You can see it provides end-to-end full data protection over various layers of infrastructure and application layers.

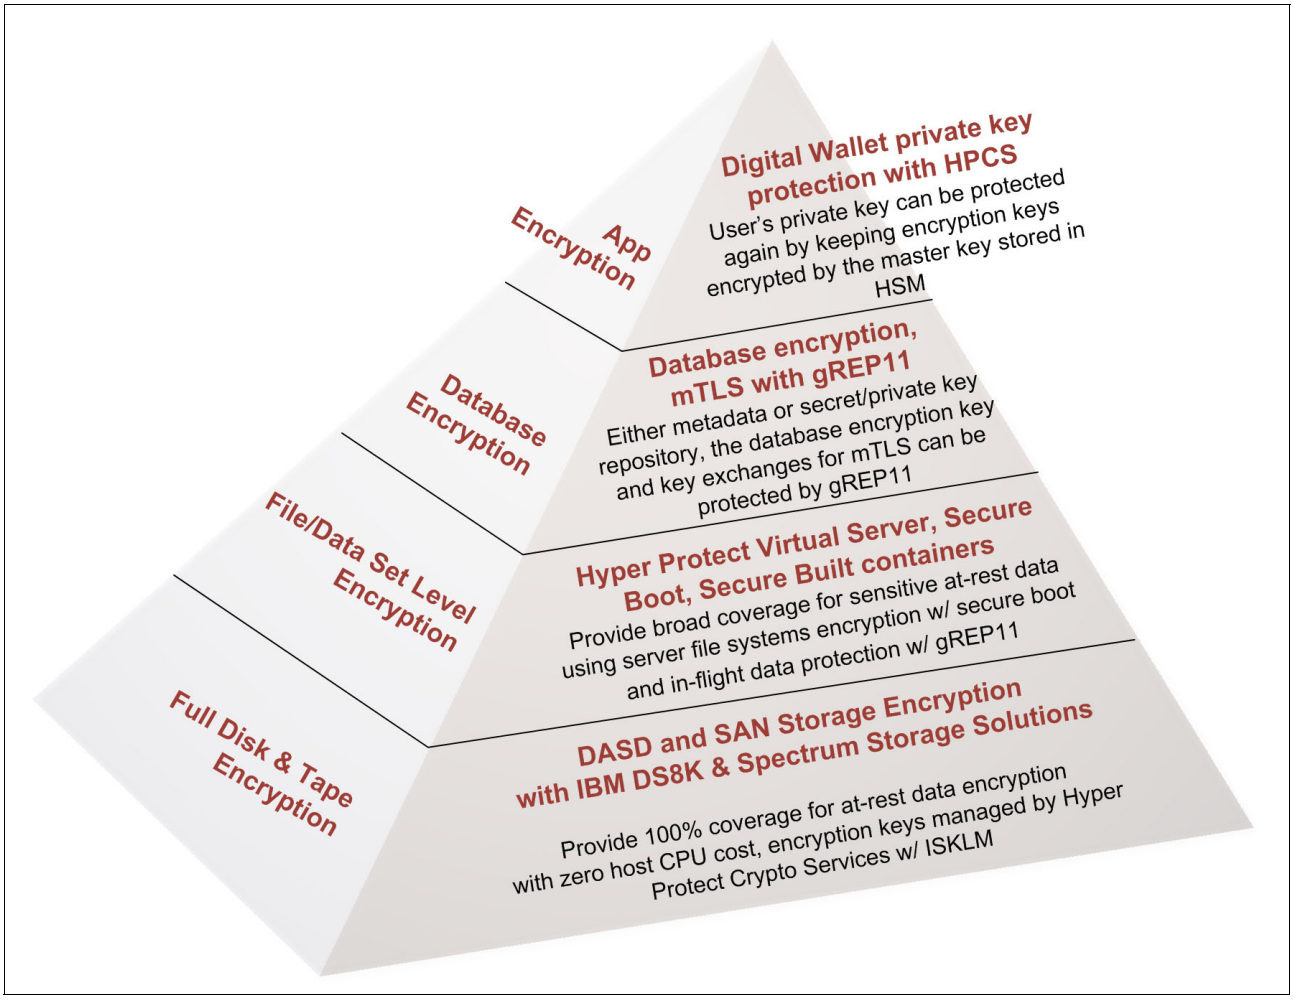

Figure 6-3 LinuxONE Pervasive Encryption with Hyper Protect Digital Asset Platform

7

Confidential AI

Confidential artificial intelligence (AI) is a paradigm that emphasizes the protection of sensitive data throughout the entire AI lifecycle, from data collection and processing to model training and deployment. This chapter delves into the principles, benefits, and technical implementations of confidential AI, highlighting how it can be integrated into modern AI systems to ensure data privacy, integrity, and security.

In 2022, IBM introduced the groundbreaking IBM® z16™ featuring an on-chip AI inference accelerator that is focused on accelerating AI model execution to deliver real-time insights that can meet the needs of the most demanding business workloads. This IBM Telum® processor with the Integrated Accelerator for AI has been designed to run inferencing for high volume workloads at scale and with security in mind.

With a carefully selected suite of AI software (AI Toolkit for IBM Z and LinuxONE and IBM Cloud Pak® for Data on IBM Z and LinuxONE), clients can manage AI model lifecycles in one place allowing for a quick deployment of wide range of use cases.

The AI Bundle for IBM Z® and LinuxONE builds on the success of the IBM z16 or LinuxONE 4 by providing new dedicated capacity for AI workloads with a highly optimized software stack.

In this chapter, we discuss how IBM LinuxONE can help you to achieve confidential AI. Included in this chapter are the following sections:

7.1 Telum Chip

In this section, we explore how the Telum chip leverages advanced security features to enhance AI processing and data protection. This section includes technical details to help you understand the inner workings of the Telum chip.

The IBM z16 (and future generations) is equipped with an on-chip accelerated AI processor named Telum. Each Telum chip consists of two 7nm, eight-core / sixteen-thread processors designed to run at a base clock speed above 5GHz. A typical IBM z16 will have sixteen of these chips in total, arranged in four-socket “drawers“.

Telum has been designed to bring real-time AI for transaction processing at scale and make AI inferencing capabilities available. The following features of Telum help businesses maximize their security:

- **Quantum-Safe Protection:** Protect today’s data against current and future threats with quantum-safe cryptography APIs and crypto-discovery tools.
- **Flexible Capacity:** Enhance resiliency by dynamically shifting system resources across locations to proactively avoid disruptions.
- **Cost Reduction and Compliance:** Simplify and streamline compliance tasks, reducing costs and keeping up with changing regulations.

7.1.1 Integrated security features

In this section, we discuss the integrated security features of the Telum chip.

- **Hardware-based roots of trust**

  The Telum chip uses hardware-based roots of trust to provide a secure foundation for all operations. This includes:

  - **Secure Boot:** Ensures that the system boots only with trusted software, preventing unauthorized firmware and software from running on the device. This establishes a foundation of trust for subsequent software layers.
  - **Hardware Validation:** Continuously validates the integrity of hardware components, safeguarding against tampering and ensuring that all hardware operates as intended.
  - **Unauthorized Modification Prevention:** Protects against any unauthorized modifications to the firmware or software, maintaining the integrity and security of the system throughout its operation.
  - **Firmware and Microcode Updates:** Securely delivers updates to firmware and microcode to address vulnerabilities and enhance security features without compromising system integrity.
  - **Physical Security Measures:** Incorporates physical security measures such as tamper-evident packaging and secure hardware destruction procedures to prevent physical attacks on the hardware components.

- **Real-time threat detection and response**

  The Telum chip is equipped with capabilities to detect and respond to threats in real-time:

  - **Integrated Security Systems:** Monitors and reacts to anomalies at the hardware level, significantly reducing the latency usually involved in software-based threat detection. This proactive monitoring allows for immediate identification of potential security breaches.

- **Anomaly Detection:** Utilizes AI-driven analytics to identify potential threats and anomalies in real-time. This includes behavioral analysis to detect deviations from normal operation patterns that may indicate a security threat.
- **Immediate Response:** Automatically initiates countermeasures to mitigate identified threats, ensuring continuous protection of data and operations. This real-time response capability is crucial for maintaining the security and availability of critical systems.

Note: The combination of these features provides a robust security framework, enhancing the overall security posture of systems leveraging the Telum chip

7.2 Data protection mechanisms

In this section, we discuss data protection mechanisms.

7.2.1 On-chip encryption capabilities

The Telum processor includes the following advanced encryption technologies to ensure data security:

- **AES (Advanced Encryption Standard):** Provides hardware-accelerated encryption for data at rest and in transit, reducing reliance on external software encryption solutions.
- **Real-Time Data Encryption:** Automatically secures data processed by the chip, ensuring that sensitive information is protected throughout its lifecycle.

7.2.2 Secure key management

In section 3.4, “Cryptographic Key Management for LinuxONE” on page 42, we discussed key lifecycle management. In this section, we expand upon key lifecycle management by including how the new Telum chip is designed to help secure your keys.

Effective key management is critical for any encryption system. The following are a few of the secure key management features:

- **Key generation and storage:** The Telum chip securely generates and stores cryptographic keys, protecting them against unauthorized access.
- **Key Lifecycle Management:** Manages the entire lifecycle of cryptographic keys, from generation to destruction, in compliance with industry security standards.
- **HSMs:** Utilizes Hardware Security Modules (HSMs) to securely store and manage cryptographic keys, adding an extra layer of security.

Tip: Utilizing hardware-based key management enhances security by ensuring keys are never exposed in plaintext.

7.2.3 Key rotation best practices

The following are some of the best practices for key rotation.

- **Regular Rotation:** Rotate keys regularly to reduce the risk of compromise.
- **Automate Key Rotation:** Automate the process to ensure consistency and reduce human error.

- **Secure Old Keys:** Archive or destroy old keys securely.
- **Synchronize Key Rotation:** Ensure all instances of the key are rotated simultaneously.
- **Audit and Monitor:** Regularly audit and monitor the key rotation process.
- **Plan for Rotation Failures:** Have a plan for handling key rotation failures.

7.2.4 Mainframe advantages for key management

The following are some of the advantages for key management on the mainframe.

- **Reliability:** High reliability ensures key management system availability.
- **Security:** Provides unmatched security for storing and managing encryption keys.
- **Efficiency:** Handles large-scale workloads efficiently.
- **Scalability:** Can handle increasing key management demands.
- **Compliance:** Meets strict compliance standards.
- **Centralized Management:** Enables easier control, rotation, and auditing of keys.

7.2.5 Secure boot and firmware integrity

Secure boot means that bad actors cannot inject malware into the boot process to take over the system during startup. The Telum chip has secure boot to protect firmware with no changes required. In addition to the features discussed in 2.3, “IBM Secure Boot for Linux” on page 30, the following are additional boot integrity features:

- **UEFI Secure boot:** Verifies each component before use to prevent malware.
- **Maintaining firmware integrity:** Uses policy objects to verify the next entity before execution.
- **Root-of-Trust (RoT):** Establishes a trusted core to verify software authenticity.
- **Signed and verified firmware:** Ensures integrity by signing and verifying firmware images.
- **Robust firmware update processes:** Keeps firmware up-to-date with security patches.

7.2.6 Firmware attacks and mitigation strategies

The Telum chip can help fight against firmware attacks. The following are some mitigation strategies to assist:

- **Update Software Immediately:** Apply patches rapidly to reduce exploitation risks.
- **Defend Privileges and Accounts:** Use Privileged Access Management (PAM) for fine-grained access control.
- **Signed Software Execution Policies:** Enforce policies for scripts, executables, and firmware.
- **System Recovery Plan:** Ensure data restoration as part of disaster recovery (DR).
- Inventory and Regular Firmware Updates: Maintain an updated inventory and regularly update firmware.
- **Hardware with Firmware Protection:** Upgrade to hardware offering built-in firmware protection.
- **Backup Firmware:** Regularly back up firmware and configurations for restoration.

7.2.7 Cryptographic hardware and AI acceleration

Dedicated hardware support for cryptographic operations is provided, enhancing performance and ensuring that security functions do not introduce significant latency.

The IBM® Telum™ processor has industry-first on-chip integrated accelerators to predict and automate with AI and is designed for unprecedented speed, scale and extremely low latency.

7.2.8 Secure enclave technology

The IBM® Telum™ processor employs secure enclaves to isolate sensitive computations and data, ensuring that even if other parts of the system are compromised, the enclave remains secure.

By implementing these strategies and best practices, IBM's LinuxONE enhances security and ensures robust key management and firmware integrity. Learn more about the key management solutions and firmware integrity practices available with IBM LinuxONE.

7.3 Use cases for confidential AI

In this section, we discuss some use cases where confidential AI is of utmost importance and is enhanced by using the Telum processor.

Table 7-1 summarizes the use cases we will discuss.

Table 7-1 Use cases for confidential AI

Industry	Use Case
Healthcare	Protecting patient data during diagnostics
Financial Services	Securing transaction data in fraud detection
Government	Safeguarding sensitive data in intelligence

7.3.1 Healthcare

Confidential AI is instrumental in protecting patient data during diagnostics and treatment planning. By ensuring compliance with healthcare regulations such as Health Insurance Portability and Accountability Act (HIPAA), confidential AI enables healthcare providers to utilize advanced AI models for patient care without compromising data privacy. For example, AI can assist in identifying patterns in medical imaging while keeping patient data secure through encryption and secure processing environments.

The health care industry would use confidential AI to encrypt patient data during AI model training, ensuring privacy and compliance with healthcare regulations such as HIPAA. This approach not only protects sensitive information but also allows for the secure deployment of AI-driven diagnostic tools.

One example use case for the health care industry and protecting patient data during AI-driven diagnostics is shown in Table 7-2 on page 82, where we have outlined the key considerations for confidential AI.

Table 7-2 Example use case - Healthcare

Objective	Ensure the security and privacy of patient data while leveraging AI for diagnostics.
Data-at-Rest Encryption	Encrypt patient records and medical data using robust encryption techniques (for example, AES-256).
Data-in-Transit Encryption	Secure data as it moves between healthcare systems using TLS encryption.
Data-in-Use Protection	Utilize secure enclaves during AI model training and inference to prevent unauthorized access.
Key Management	Manage cryptographic keys securely using hardware-based solutions (for example, HSMs).
Compliance	Adhere to healthcare regulations (such as HIPAA) to maintain patient privacy.

7.3.2 Financial services

In the financial services sector, Confidential AI secures transaction data and customer information, particularly in applications like fraud detection and risk management. By leveraging secure enclaves for data processing, financial institutions can enhance the accuracy of their AI models while ensuring the confidentiality and integrity of sensitive financial data. This approach helps in detecting fraudulent activities in real-time, providing an additional layer of security to customers' financial information.

The financial sector would implement confidential AI to protect customer data in fraud detection systems. By leveraging secure enclaves and hardware-based encryption, financial institutions can process data securely, enhancing the reliability and security of their fraud detection algorithms.

A sample use case for the financial services industry for securing transaction data in fraud detection systems is shown in Table 7-3.

Table 7-3 Example use case for financial services

Objective	Protect sensitive financial information and prevent fraudulent activities.
Data-in-Transit Encryption	Encrypt transaction data during communication between banking systems.
Data-in-Use Protection	Process data within secure enclaves to prevent exposure during fraud detection algorithms.
Key Management	Safeguard cryptographic keys throughout their lifecycle.
Industry Standards	Comply with financial industry standards (for example, PCI DSS) to ensure data security.

7.3.3 Government

Confidential AI plays a crucial role in safeguarding sensitive government data used in intelligence analysis and public service optimization. Ensuring compliance with stringent regulatory requirements, Confidential AI allows government agencies to process and analyze data securely. This includes using hardware-based encryption and isolated key management to protect data during AI model training and inference, thereby enhancing the security and efficacy of public sector services.

Table 7-4 provides a sample use case for government to safeguard sensitive data in intelligence analysis.

Table 7-4 Example use case for government

Objective	Ensure national security by protecting classified information.
Data-at-Rest-Encryption	Encrypt sensitive intelligence data stored in databases or repositories
Isolated key management	Manage cryptographic keys in isolated environments (for example, HSMs) to prevent unauthorized access.
High availability and disaster recovery	Implement redundant architectures for continuous operation.
Regulatory compliance	Adhere to stringent regulatory requirements specific to government agencies.

7.3.4 Summary

Confidential AI provides a robust framework for protecting sensitive data throughout the AI lifecycle, ensuring data privacy, integrity, and security. By leveraging advanced encryption techniques, secure execution environments, and effective key management, organizations can confidently integrate AI technologies while maintaining compliance with regulatory requirements. The principles, benefits, and technical implementations of confidential AI highlighted in this chapter demonstrate its critical role in modern AI systems.

Table 7-5 summarizes the key takeaways from this section.

Table 7-5 Table 7-2: Key takeaways from confidential AI

Principle	Benefit to confidential AI
Data Privacy	Ensures sensitive information is protected
Data Integrity	Maintains tamper-proof systems
Regulatory Compliance	Helps meet data protection regulations

7.4 Hyper Protect Services

In this section, we explore examples of how IBM Hyper Protect Services (HPS) leverages advanced security features to enhance AI processing and data protection. This section includes more technical details to help you understand what occurs “under the hood” of Hyper Protect Services.

IBM Hyper Protect Services (HPS) is designed to provide unparalleled security for data and workloads, ensuring that organizations can confidently leverage AI while maintaining strict data confidentiality and integrity. HPS combines advanced encryption, exclusive control over cryptographic keys, automated compliance features, and robust hardware-based security measures to offer a comprehensive solution for protecting sensitive data throughout its lifecycle (see Figure 7-1).

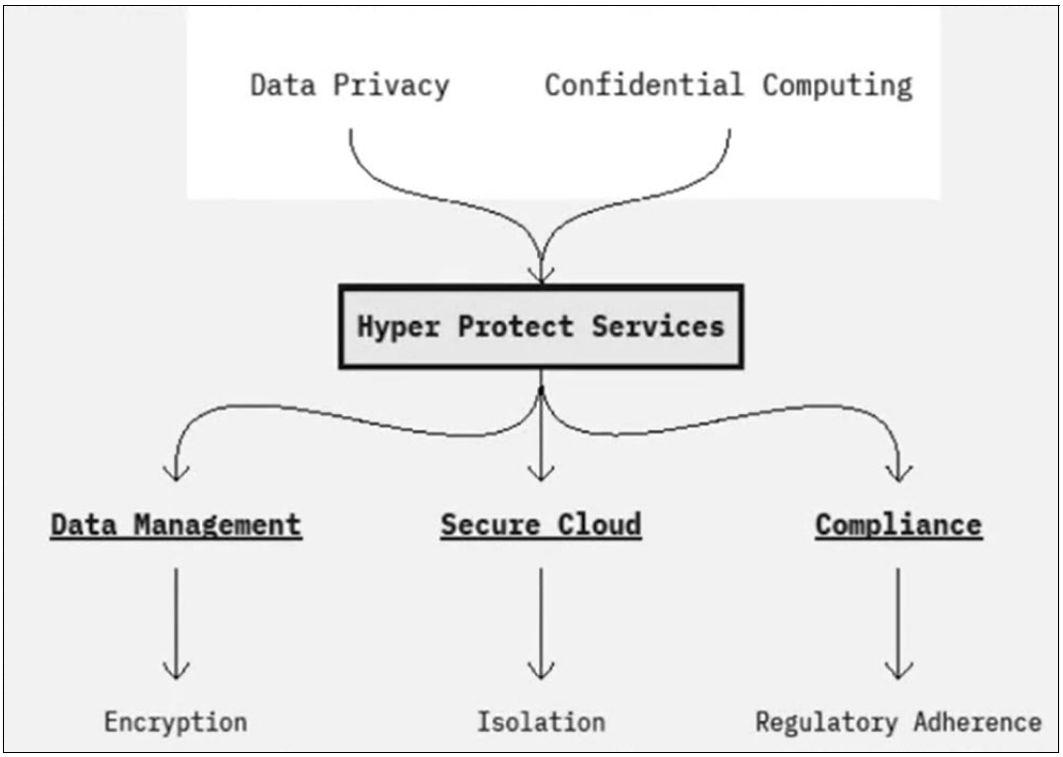

Figure 7-1 Hyper Protect Services and confidential computing

7.4.1 Hyper Protect Services Key features

Some of the key features of Hyper Protect Services include the following:

Comprehensive security

Hyper Protect Services encompass a wide array of security measures designed to protect data at every stage and includes virtual servers. Virtual servers securely isolate workloads by using encrypted virtual machines, providing a secure environment for running critical applications.

- **Encrypted Databases:** Utilize AES-256 encryption to secure data stored in databases, ensuring that sensitive information is protected from unauthorized access.
- **Crypto Services:** Manage cryptographic operations in secure hardware modules, leveraging hardware security modules (HSMs) to ensure the integrity and confidentiality of cryptographic keys. IBM Cloud Hyper Protect Crypto Services is an as-a-service (aaS) key management and encryption solution, which gives you full control over your encryption keys for data protection.
- **Virtual Servers:** Securely isolate workloads using encrypted virtual machines, providing a secure environment for running critical applications.

Exclusive control

You retain full control over your data and cryptographic keys.

- **Key Ownership:** Clients manage their own encryption keys with Keep Your Own Key (KYOK) capabilities, ensuring that no third party can access or control their keys.
- **No Third-Party Access:** Service providers have no access to client data or keys, maintaining the highest level of data privacy and security.

Automated compliance

Features that streamline adherence to regulatory requirements include the following:

- **GDPR:** Ensure data protection and privacy in compliance with European regulations, safeguarding personal data and maintaining regulatory compliance.
- **HIPAA:** Secure healthcare data in accordance with U.S. standards, protecting patient information and ensuring compliance with health regulations.
- **PCI-DSS:** Protect payment data as per industry standards, securing financial transactions.

Tip: Hyper Protect Services offer automated compliance features that can significantly reduce the burden of regulatory adherence. By leveraging these capabilities, organizations can ensure that their data handling practices are always up-to-date with the latest regulations.

7.5 LinuxONE and confidential AI

In this section, we discuss confidential AI and LinuxONE.

7.5.1 Data privacy

Data privacy focuses on the individual rights of data subjects—that is, the users who own the data. For organizations, the practice of data privacy is a matter of implementing policies and processes that allow users to control their data in accordance with relevant data privacy regulations.

Telum has been designed with significant innovation in security, with transparent encryption of main memory. Telum's Secure Execution improvements are designed to provide increased performance and usability for Hyper Protected Virtual Servers and trusted execution environments, making Telum an optimal choice for processing sensitive data in Hybrid Cloud architectures.

LinuxONE 4 with the Telum processor complements data privacy in the following ways:

- **Secure Data Handling:** Ensures that data remains private throughout the AI lifecycle, protecting sensitive information from unauthorized access.
- **Confidential Computing:** Uses secure enclaves to protect data during computation, ensuring that data is isolated and processed securely.

7.5.2 Integrity and auditability

Compliance audits are often a major time drain, taking staff away from more pressing efforts to satisfying auditor's demands and stitching together data from multiple sources. Fighting mainframe compliance "drift" (maintaining, updating and adding new processes for compliance) can be costly.

The IBM LinuxONE Security and Compliance Center is an integrated set of microservices that collects evidence data from participating IBM software components and products and provides the following:

Model Integrity: Protects AI models from tampering, ensuring that they remain reliable and trustworthy throughout their lifecycle.

End-to-End Security: Secures AI models from development to deployment, ensuring that data and models are protected at every stage.

7.5.3 Secure AI model lifecycle management

The AI model lifecycle consists of the following four main steps:

1. Collect: Make data simple and accessible.
2. Organize: Create a business-ready analytics foundation.
3. Analyze: Build and scale AI with trust and transparency.
4. Infuse: Operationalize AI throughout a business.

The following are key points to remember for secure AI model lifecycle management.

- **End-to-End Security:** Secures AI models from development to deployment, ensuring that data and models are protected at every stage.
- **Policy Enforcement:** Ensures consistent security and compliance, implementing policies that govern data handling and access throughout the AI lifecycle.

While response time, quality, fairness, explainability, and other elements must be managed as part of the whole AI lifecycle, security during every facet of this lifecycle must also be considered and managed.

One example of a tool that will help you to manage this lifecycle is IBM Cloud Pak for Data, a multicloud data and AI platform with end-to-end tools for enterprise-grade AI model lifecycle management.

The solution areas that IBM Cloud Pak for Data addresses are:

- Data fabric and AI Lifecycles – Automating integration, metadata management and AI governance
- Data Governance and Security
- Data Integration
- Data Observability
- Master Data Management

7.5.4 Use cases

In Table 7-6, we have identified three industries and how they would use Hyper Protect services.

Table 7-6 Hyper Protect Services Use Cases

Industry	Use case
Healthcare	AI-driven diagnostics

Industry	Use case
Financial Services	Fraud detection systems
Government	Intelligence analysis

1. **Healthcare provider:** A healthcare provider can use Hyper Protect Services to encrypt patient data during AI model training. This ensures complete confidentiality and regulatory compliance. The encrypted data can only be accessed through secure enclaves, ensuring that sensitive patient information is not exposed during the model training process.
2. **Financial services:** Financial institutions can secure customer data in fraud detection systems by leveraging secure enclaves for processing. This ensures that customer data remains confidential and secure, even while it is being processed for fraud detection.
3. **Government agency:** Government agencies can protect data in AI-based intelligence platforms using hardware-based encryption and isolated key management. This ensures that sensitive data remains secure and confidential, even when it is being processed and analyzed.

7.5.5 Technical implementations

In this section, we summarize with a high-level overview of a variety of techniques to enable confidential AI.

Encryption techniques

Encryption can protect data at rest, in transit and while being processed (in-use), regardless of whether the data is in a computer system on-premises or in the cloud. For this reason, encryption has become critical to cloud security efforts and cybersecurity strategies more broadly. The following are some of the ways you can protect your data with encryption techniques.

1. **Data-at-Rest Encryption:**
   - **AES-256 Encryption:** IBM Cloud Hyper Protect Crypto Services (HPCS) employs AES-256 encryption to protect stored data, ensuring sensitive information is secure from unauthorized access.
2. **Data-in-Transit Encryption:**
   - **TLS Encryption:** HPCS uses TLS to secure data in transit, protecting data as it moves between systems.
3. **Data-in-Use Encryption:**
   - **Secure Enclaves:** HPCS leverages secure enclaves to protect data during computation, ensuring sensitive information remains protected even while being processed.

Tip: Isolated key management in HSMs ensures that cryptographic keys are never exposed in plaintext, providing an additional layer of security for sensitive data.

Key Management

Careful management of keys is vital to the effective use of cryptography in your cybersecurity strategy. The following are some of the implementations you can employ. A dedicated key management service and Hardware Security Module (HSM) provides you with the Keep Your

Own Key capability for cloud data encryption. IBM Cloud Hyper Protect Crypto Services provides you with exclusive control of your encryption keys in the highest security level.

1. **Isolated Key Management:**
   - **Hardware Security Modules (HSMs):** Keys are managed in FIPS 140-2 Level 4 certified HSMs, ensuring cryptographic keys are never exposed in plaintext. For more information, see IBM Cloud Hyper Protect Crypto Services.
2. **Key Lifecycle Management:**
   - **Comprehensive Management:** HPCS provides robust key lifecycle management, from generation to secure destruction, ensuring secure handling of cryptographic keys throughout their lifecycle. For more information, see NIST Guidelines on Key Management.

High availability and scalability

The IBM LinuxONE 4 family extends the latest performance, security and AI capabilities of LinuxONE. Additionally, the following enhances high availability (HA) and scalability.

1. **Redundant Architectures:**
   - **High Availability and Disaster Recovery:** IBM Cloud Hyper Protect Crypto Services includes redundant architectures to ensure high availability and disaster recovery (DR), maintaining continuous operation even in the face of hardware failures. For more information, see IBM Cloud HA and DR.
2. **Elastic Scalability:**
   - **Dynamic Resource Allocation:** HPCS supports elastic scalability, dynamically allocating resources for optimal performance, ensuring the system can scale to meet the needs of demanding workloads. For more information, see IBM Cloud Pak for Integration.

7.5.6 Summary

IBM Hyper Protect Services provides a robust framework for securing AI environments, offering advanced encryption, exclusive key control, and comprehensive compliance features. By integrating HPS into their operations, organizations can protect sensitive data, maintain regulatory compliance, and leverage AI technologies with confidence. The detailed technical implementations and real-world use cases highlighted in this chapter demonstrate the practical benefits and applications of Hyper Protect Services in enhancing data security and integrity across various industries.

Tip: By leveraging IBM Hyper Protect Services, organizations can achieve a higher level of data security and compliance, enabling them to fully utilize the power of AI while protecting sensitive information.

A

Reference guide

Table A-1 lists some of the common cryptographic hardware libraries, tools, and drivers that can be used on LinuxONE.

Table A-1 Cryptographic hardware libraries, tools, and drivers

Name	Function
libcsulcca	The libcsulcca library provides APIs for CCA and secure key cryptography functions that are provided by cryptographic express coprocessor.
libica	The libica library provides hardware support for cryptographic functions. It is part of the openCryptoki project in GitHub. It is primarily used by OpenSSL through the IBM OpenSSL CA engine or by openCryptoki through the ICA token. A higher level of security can be achieved by using it through the PKCS #11 API implemented by openCryptoki.
libzpc	The IBM Z Protected-key Crypto library libzpc is an open-source library targeting the 64-bit Linux on IBM Z (s390x) platform. It provides interfaces for cryptographic primitives. The underlying implementations make use of the extensive performance-boosting hardware support of z/Architecture and its protected-key feature which ensures that key material is never present in main memory at any time.
openCryptoki	openCryptoki is an open-source implementation of the Cryptoki API that is defined by the PKCS #11 Cryptographic Token Interface Standard. Therefore, openCryptoki supports several cryptographic algorithms according to the industry-wide PKCS #11 standards. The openCryptoki library loads the so-called tokens that provide hardware- or software-specific support for cryptographic functions.
z90crypt	z90crypt is a cryptographic device driver. It acts as the interface to the PCI cryptographic card coprocessor. This driver must be loaded to use CEX features.
zKVM	
virsh	You can create, delete, run, stop, and manage your virtual machines from the command line by using a tool that is called virsh. Virsh is useful for advanced Linux administrators who are interested in script or automating some aspects of managing their virtual machines.
qemu	Qemu is a machine emulator that can run operating systems and programs for one machine on a different machine. Mostly it is not used as emulator but as virtualizer in collaboration with KVM kernel components. In that case, it uses the virtualization technology of the hardware to virtualize guests.

libvirt	Although qemu has a command-line interface (CLI) and a monitor to interact with running guests, it is rarely used that way for other means than development purposes. Libvirt provides an abstraction from specific versions and hypervisors and encapsulates some workarounds and best practices. libvirt is an open-source API, daemon, and management tool for managing platform virtualization. It can be used to manage KVM, QEMU, and other virtualization technologies. These APIs are widely used in the orchestration layer of hypervisors in the development of a cloud-based solution.
vfio	The VFIO driver is an IOMMU/device-agnostic framework for making available direct device access to userspace in a secure, IOMMU-protected environment. That is, it allows safe, non-privileged, userspace drivers.
Linux tools	
`chzcrypt`	The **`chzcrypt`** command is used to configure cryptographic adapters that are managed by the cryptographic device driver and modify the AP bus attributes.
`lszcrypt`	The **`lszcrypt`** command is used to display information about cryptographic adapters that are managed by the cryptographic device driver and its AP bus attributes.
`icainfo`	The **`icainfo`** command is used to determine which libica functions are available on your Linux system.
`icastats`	The **`icastats`** command is used to determine whether libica uses hardware acceleration features or works with software fallbacks. icastats collects the statistical data per user and not per system.
`zcryptctl(KVM)`	The **`zcryptctl`** command is used to control access to AP queues and functions.
z/VM	

`QUERY CRYPTO`
 	
`DOMAINS USERS`
	These command queries can be used to show the status of the cryptographic hardware.
	
	When the DOMAINS operand is specified, the status of the installed AP domains is displayed. When the USERS operand is specified after the DOMAINS operand, the users who are authorized for CRYPTO APVIRT in the directory are listed.

REDP-5535-02

ISBN 0738461792

Printed in U.S.A.