Data Sovereignty with IBM Hyper Protect Services

with a special focus on Key Management

Why IBM for confidential computing?

IBM’s portfolio of cloud service offerings can help your enterprise:

- **Protect sensitive data at rest, in transit and in use.** With IBM’s security-first approach and framework you can attain your data protection and privacy requirements and mitigate risks by meeting any regulatory requirements. Protecting sensitive data requires a holistic approach—spanning compute, containers, databases and encryption.
- **Deliver the highest level of commercial technical assurance.** IBM Cloud Hyper Protect Cloud Services enables end-to-end protection for business processes in the cloud and is built on secured enclave technology. It delivers the highest level of key management protection through FIPS 140-2 Level 4 HSM.
- **Gain complete authority over your data.** Single-tenant key management services, with integrated HSMs, provide complete control of cloud data encryption keys for data encryption at rest and private keys related to data in transit. The portfolio enables the span of confidential databases, confidential servers and confidential containers, which enable you to have complete authority over your data with technical assurance.

Data protection and privacy challenges in the cloud

As we enter a new normal period of accelerated digital transformation post-COVID, the vast number of organizations are now relying heavily on public and hybrid cloud services. And companies in highly regulated industries, now more than ever, find themselves needing cloud services that offer a greater level of protection and privacy.

As a result, data privacy and protection outside of the traditional perimeter and in the cloud have become a chief information security officer’s (CISO’s) imperative. The global average cost of a data breach in 2020 was USD 3.86 million and 52% of those breaches were caused by malicious attacks.[1] With these increases in data breaches, an enterprise’s data protection and privacy in the cloud is at stake as it needs one single point of control that provides a holistic view of threats and mitigates complexity.

The data protection needs of organizations are driven by the concerns about protecting sensitive information, intellectual property, and meeting compliance and regulatory requirements. In today’s digital global economy, data is one of the most valuable assets so data must be protected end to end – when it’s at rest, in motion and in use.

Data is often encrypted at rest in storage and in transit across the network, but applications and the sensitive data they process — data in use — are vulnerable to unauthorized access and tampering while they are running. Even when encrypted at rest, depending on where it’s encrypted, either the data or the encryption keys could be vulnerable to unauthorized access. According to Gartner, by 2025, 50% of large organizations will adopt privacy-enhancing computation for processing data in untrusted environments to protect data in use.[2]

The dilemma for organizations is how do they independently retain ownership and control of their data while still driving innovation? Protecting sensitive data is vital to an enterprise’s cloud data security, privacy and digital trust.

As enterprises contemplate moving sensitive data and workloads to the public cloud, they’re looking for ways to address the following concerns:

1. Is my data and my customers’ data safe in the cloud?
2. How do I meet regulatory and privacy requirements?
3. How can I ensure that my cloud provider has no access to my data?
4. How do I protect personal identifiable information (PII)?
5. How do I preserve privacy of user and business data?
6. How do I preserve privacy of data while performing analytics and AI modeling or sharing data with other third parties?

When hosting their data with cloud providers, companies want to have complete authority over their valuable data and associated workloads, including no access to sensitive data for even their cloud providers.

So how can you protect your sensitive data in the public cloud?

Encryption is a key technical measure to safeguard data in the cloud. The loss of data often leads to loss of customer trust with serious financial consequences. Regulatory compliance often mandates encryption of data at rest and in transit or strongly encourages it as a technical measure to protect data. And regulatory compliance requirements can be complex and the penalties significant. Extensive use of encryption, data loss prevention, threat intelligence sharing, and integrating security into the development, security and operations process (DevSecOps) were all associated with lower-than-average data breach costs.

Among these safeguards, encryption had the greatest impact. Deploying extensive encryption can be a substantial cost-mitigating factor in the event of a data breach — as the average total reduction in the cost of a breach due to extensive encryption was USD 237 thousand in 2020.[1]

Yet, data protection through encryption is only as strong as your ability to protect the keys used to encrypt the data. With constant threats of external cyberattacks and insider threats, now, more than ever, there’s a need for workload isolation, data encryption, trusted execution environments, and other security practices and tools to protect your most sensitive workloads.

How can you mitigate these concerns and risks?

The current approaches to securing data is through data at rest and data in transit encryption. However, the challenging problem resides in gaining technical assurance that only you have access to your data or keys and protecting sensitive data in use to provide protection at all stages of data usage. Due to the growing understanding of the need for data in use protection, the adoption of confidential computing is increasing.

The term confidential computing refers to cloud computing technology that protects data while in use. The technology helps reduce security concerns as companies adopt more cloud services. The primary goal of confidential computing is to provide greater privacy assurance to companies that their data in the cloud is protected and confidential and instill confidence in moving more of their sensitive data and computing workloads to any location, including public cloud services.

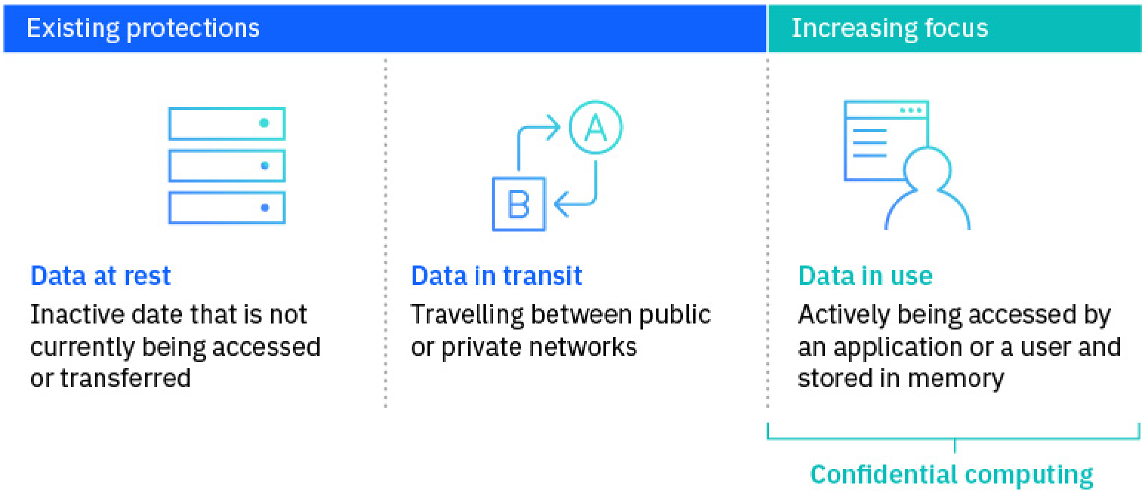

Existing protections

Increasing focus

Data at rest
Inactive date that is not currently being accessed or transferred

Data in transit
Travelling between public or private networks

Data in use
Actively being accessed by an application or a user and stored in memory

Confidential computing

Protect data across the compute lifecycle. To achieve the highest level of commercial privacy assurance, IBM goes beyond confidential computing to help protect your sensitive data across the entirety of the compute lifecycle — providing you with complete authority over your data at rest, in transit and in use.

What should you know about protecting your data across the lifecycle? Explore the following chapters to learn more about confidential computing and how it can help with data privacy and protection in your hybrid cloud environments.

[1] Cost of a Security Breach Report, Ponemon Institute, 2020
[2] Top Strategic Technology Trends for 2021, Gartner, January 2021

How confidential computing addresses total privacy assurance

Why adopt a confidential computing approach?

1. **To protect sensitive data, even while in use and to have full authority over your data and keys.**
   Confidential computing protects sensitive data when used together with data encryption at rest and in transit, along with exclusive control of keys. It thereby eliminates the single largest barrier to moving sensitive or highly regulated data sets and application workloads from an inflexible, expensive on-premises IT infrastructure to a more flexible and modern public cloud platform.

2. **To protect intellectual property.**
   Confidential computing isn't just for data protection. Data owners and data scientists can use this approach to protect proprietary business logic, analytics functions, machine learning (ML) algorithms or entire applications.

3. **To collaborate securely with partners on new cloud solutions.**
   For example, one company can combine its sensitive data with another company's proprietary calculations to create new solutions — without either company sharing any data or intellectual property it doesn't want to share.

4. To eliminate concerns when choosing cloud providers.
   Confidential computing lets a company choose the cloud computing services that best meet its technical and business requirements without worrying about storing and processing customer data, proprietary technology, and other sensitive assets. This flexibility helps alleviate any additional competitive concerns if the cloud provider also provides services to competing businesses.

5. To protect data processed at the edge.
   Edge computing is a distributed computing framework that brings enterprise applications closer to data sources, such as Internet of Things (IoT) devices or local edge servers. When it's used as part of distributed cloud patterns, the data and application at edge nodes can be protected with confidential computing.

What is confidential computing?

Confidential computing is a cloud computing technology that isolates sensitive data and code in a protected CPU enclave during processing. The contents of the enclave — the data being processed, and the techniques used to process it — are accessible only to authorized programming code, and invisible and unknowable to anything or anyone else, including the cloud provider.

It protects data during processing and, when combined with storage and network encryption with exclusive control of encryption keys, provides end-to-end data security in the cloud.

For years cloud providers have offered encryption services for protecting data at rest in storage and databases, and data in transit, moving over a network connection. Confidential computing eliminates the remaining data security vulnerability by protecting data in use — that is, during processing in a runtime.

How confidential computing works

IBM's approach is to help provide total privacy assurance with confidential computing. Protecting sensitive data requires a holistic approach — spanning compute, containers, databases and encryption.

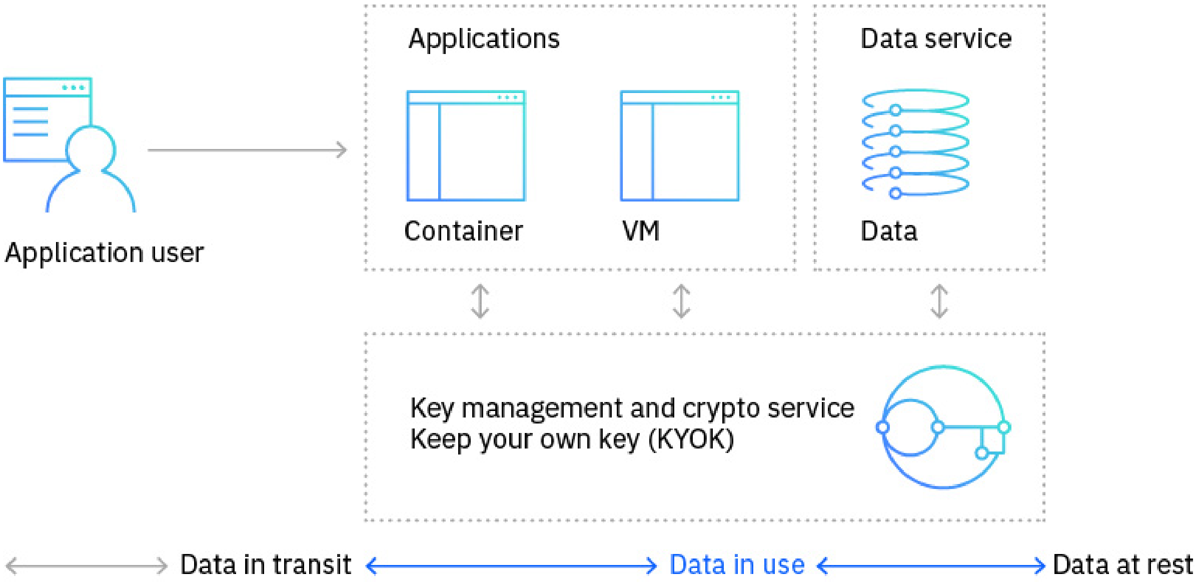

Before data can be processed by an application, it's unencrypted in memory. This step leaves the data vulnerable just before, during and just after processing to memory dumps, root-user compromises and other malicious exploits.

Confidential computing solves this problem by using a hardware-based trusted execution environment (TEE), which is a secure enclave within a CPU. The TEE is secured using embedded encryption keys, and embedded attestation mechanisms that help ensure the keys are accessible to authorized application code only. If malware or other unauthorized code attempts to access the keys, or if the authorized code is hacked or altered in any way, the TEE denies access to the keys and cancels the computation.

In this way, sensitive data can remain protected in memory while it's decrypted within the TEE to processing. While decrypted and throughout the entire computation process, the data is invisible to the operating system, other compute stack resources, and to the cloud provider and its employees.

Technical assurance delivers the highest level of privacy and protection

Operational assurance means your cloud provider will not access your data based on trust, visibility and control. Technical assurance makes certain your cloud provider cannot access your data based on technical proof, data encryption and runtime isolation — and can protect your CI/CD pipeline from bad actors.

Operational assurance

Cloud provider will not access your data

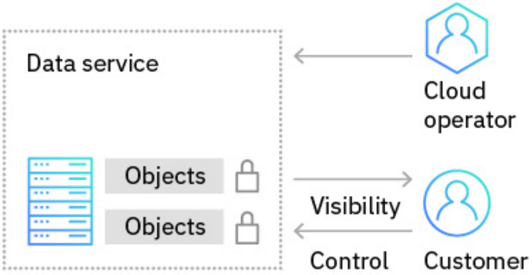

Technical assurance

Cloud provider cannot access your data

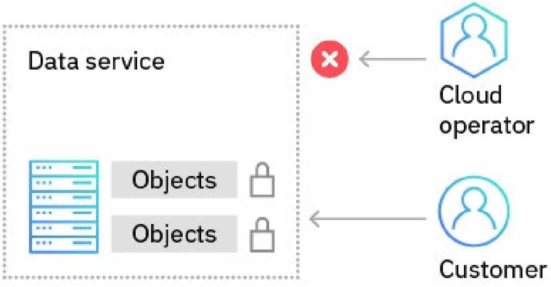

So, who do you have to protect against? You want to have the highest technical assurance that cloud administrators, vendors, software providers and site reliability engineers (SREs) can't access your data while in use.

Explore the next chapter to learn more about the services to protect your sensitive data.

A holistic approach to total privacy assurance

How to protect data with full authority — at rest, in transit and in use

Protecting sensitive data requires a holistic approach — spanning compute, containers, databases and encryption. The key is controlling access to the data as tightly as possible and provide a way to securely process unencrypted data. It's important to have technical assurance that only you have access and control over your data and to ensure your cloud service operators can't access the data or keys. The protection of these data states is complementary and doesn't supersede or replace the other existing protections.

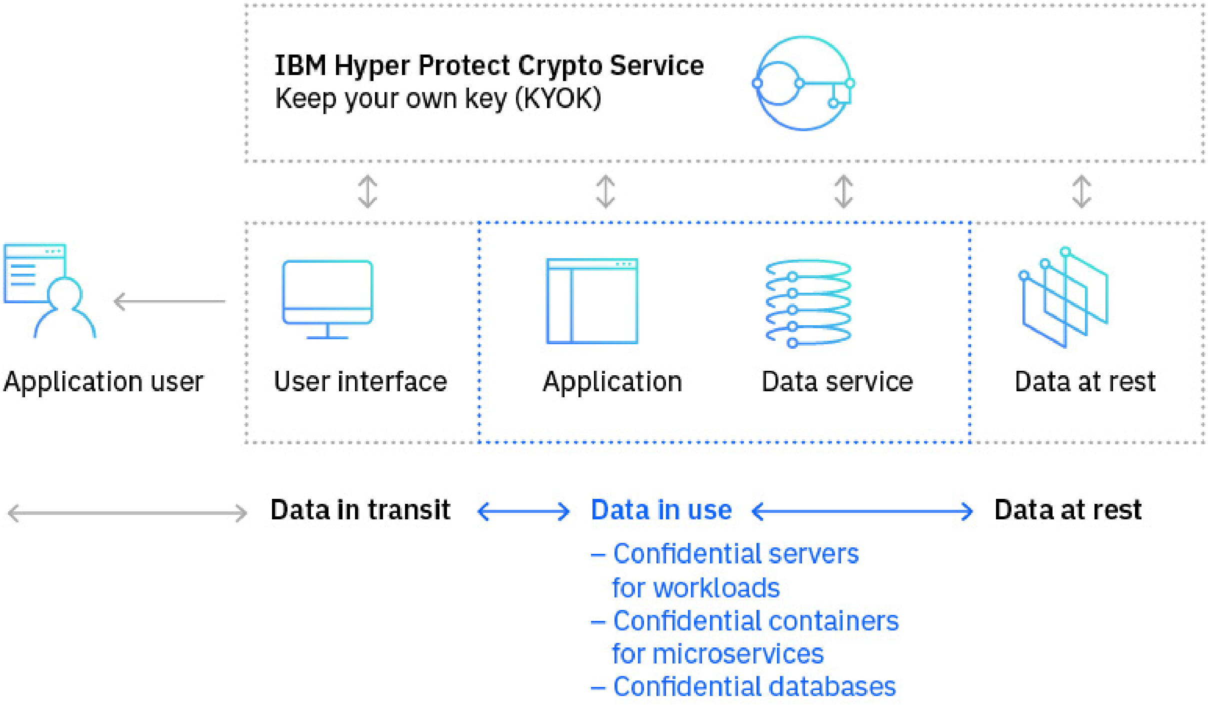

Hyper protect your sensitive data and workloads in the cloud.

IBM's capabilities include industry-leading security services for cloud data, digital assets and workloads. They're built on IBM® LinuxONE security-rich enclaves, which offer built-in protection for data at rest and in flight, plus protection of data in use. The services are designed to make it easy for application developers to build applications that deal with highly sensitive data while helping companies meet regulatory compliance requirements. Explore how IBM Cloud Data Shield and IBM Cloud Hyper Protect Services can help protect your sensitive data across the compute lifecycle.

IBM Cloud Hyper Protect Services. These services enable enterprises to have complete authority over their sensitive data, workloads and encryption keys. Not even IBM Cloud® administrators have access.

IBM Cloud Hyper Protect Crypto Services

Keep your own key for data encryption

Hyper Protect Crypto Services is a single-tenant, hybrid cloud key management service. Unified Key Orchestrator, a part of Hyper Protect Crypto Services, enables key orchestration across multicloud environments. Hyper Protect Crypto Services is built on FIPS 140-2 Level 4 certified hardware (link resides outside ibm.com), the highest level in the industry.

- Cloud data encryption that's protected in a dedicated cloud hardware security module.
- KYOK, single-tenant key management service with key-vaulting provided by dedicated, customer-controlled HSMs and that supports industry standards, such as PKCS #11.
- IBM Cloud Hyper Protect Crypto Services is also the only service in the cloud industry that is built on FIPS 140-2 Level 4-certified hardware.
- Manage security policies and orchestrate across multicloud environments from a single point of control (UKO)

IBM Cloud Hyper Protect Virtual Servers

Control workloads with sensitive data or business IP

Hyper Protect Virtual Servers are fully developer-friendly and able to use industry-standard, Open Container Initiative (OCI) images with a standard user interface to provision, manage, maintain and monitor within the Virtual Private Cloud (VPC) infrastructure of IBM Cloud. By leveraging VPC, this next generation of Hyper Protect Virtual Servers is able to offer additional network security.

- Complete authority over your IBM® LinuxONE Virtual Servers for workloads with sensitive data or business IP.
- Keeps out unauthorized users, designed to address your top security concerns, and provides a confidential computing environment even IBM Cloud administrators can't access.
- Use your own public Secure Socket Shell (SSH) key to maintain exclusive access. Support for the industry's bring your own key (BYOK).
- For more info, read the Technical Whitepaper

Hybrid-Multi Cloud Environments: Key Management

Tackling Challenges in a complex environment through Key Orchestration

- How do you establish a robust key management solution where you don't accidently lose keys? Especially as multiple systems need their own key combinations and lifecycles, if keys are lost, you may not be able to decrypt the datasets using these keys.

- How do you establish an efficient key management solution where you quickly can restore a key to one or more key stores, where you can easily add new key stores to a group and have all relevant keys (controlled by the policies) quickly installed?

- How do you deal with a complex environment with multiple environments, such as many IBM Z® or multiple clouds each having many LPARs and targets, where datasets are shared by applications across the systems in many ways?

- How do you ensure to use and distribute high quality keys for encryption while being you are the owner of your masterkeys?

The IBM Approach for a more efficient and effective Encryption

The highest level of security for your keys involves generating them with FIPS 140-2 level 4 certified Hardware Security Modules (HSMs), where only you are in possession of the master key. In the IBM Cloud, these HSMs are provided by Hyper Protect Crypto Services (HPCS). On z/OS®, you will be utilizing the IBM Crypto Express adapters. From there keys can be installed into z/OS keystores or into cloud keystores using bring your own key.

If a key is lost, you will not be able to decrypt your data anymore. In addition to secure key generation, it is therefore essential that these keys can be easily backed up and restored. Manual processes involving different tools are time consuming and error prone leading to long recovery times once applications start failing due to missing keys.

Unified Key Orchestrator for IBM z/OS (UKO)will let you monitor that all keys are installed and present in their target keystore location. Lost keys can be restored within seconds.

Furthermore, additional keystores can be added quickly to ensure a fast distribution of all keys across cloud instances or z/OS Sysplexes.

All in on Hybrid Cloud – can choose deployment option, that fits to your needs!

This solution is available as UKO for z/OS or as a SaaS offering in the IBM Cloud based on HPCS.

For more information, please contact your IBM representative.

IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.
12/23

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

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

THE INFORMATION IN THE DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION ON NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.