



IBM Cloud™ HIPAA Guidance

IBM Cloud Compliance Guide

Disclosure

This document is provided for informational purposes only. It represents IBM Cloud's product offerings and practices as of March 2021.

IBM is committed to helping our clients and prospects with the knowledge to enable them to make decisions regarding their own client base needs.

This document does not constitute any explicit, or implicit warranties, contractual commitments, conditions or assurances.

The intended audience of this guide is technical experts (e.g., solution architects and Chief Technology Officers) seeking to understand safeguards in IBM's cloud environment and services that are designed to help customers assess and meet the HIPAA requirements. It should be noted that IBM is providing a "HIPAA-Ready" solution as that term is used in this Guide. HIPAA-ready simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

Table of Contents

1	Introduction	5
2	HIPAA, HITECH and IBM Cloud Compliance	6
3	Business Associate Relationships and Requirements	6
4	Introduction to IBM Cloud	7
4.1	IBM's public cloud services and HIPAA	7
4.2	Information Security Policy & HIPAA Administrative Safeguards	8
4.3	Physical Security of IBM Cloud & HIPAA Physical Safeguards.....	8
4.4	IBM Cloud & HIPAA Technical Safeguards	9
5	The Shared Responsibility Model	10
5.1	Technical Shared Responsibilities	10
5.2	Administrative Shared Responsibilities.....	11
5.3	Physical Shared Responsibilities	11
5.4	Responsibility Summary	11
6	Architecture Examples	12
6.1	Account Setup and Managing User Identity, Access, and Authentication	12
6.2	IBM Cloud IaaS Architecture Examples as HIPAA-Ready	13
6.2.1	IBM Cloud VSIs and IBM Cloud Bare Metal as HIPAA-Ready.....	13
6.2.2	Using Firewalls and Load Balancers with VSIs or Bare Metal Server.....	15
6.2.3	Using Storage with VSIs or Bare Metal Servers	15
6.2.4	Using VMware Virtualization Architecture	15
6.2.5	IBM Cloud Infrastructure Services HIPAA Controls	19
6.3	IBM Cloud PaaS Architecture Examples as HIPAA-Ready	21
6.3.1	IBM Cloud PaaS Architecture Components	22
6.3.2	Setting up a PaaS Environment	25
6.3.3	IBM Kubernetes Service (IKS) Configuration.....	25
6.3.4	Summary: Secure Flow for an IKS Application	26
6.3.5	IBM Cloud PaaS HIPAA Controls	28
7	IBM Cloud IaaS and PaaS Services and HIPAA.....	30
7.1	IBM HIPAA-Ready Services	30
7.2	IBM Cloud Infrastructure Services Offerings and HIPAA	30
7.2.1	IBM Cloud Bare Metal.....	30
7.2.2	IBM Cloud Virtual Servers	31
7.2.3	Storage (Block Storage, Cloud Object Storage, and File Storage).....	32
7.2.4	IBM Cloud for VMware Solutions (Dedicated options).....	34
7.2.5	IBM Cloud Direct Link ("1.0 on Classic")	35
7.2.6	Hardware Security Module	35
7.3	IBM Cloud PaaS Offerings and HIPAA.....	36
7.3.1	IBM Cloud App ID.....	37
7.3.2	IBM Cloud Certificate Manager.....	37
7.3.3	IBM Cloud Databases	39
7.3.4	IBM Event Streams for IBM Cloud Enterprise.....	41
7.3.5	IBM Cloud Functions	41
7.3.6	IBM Cloud Internet Services	41
7.3.7	IBM Key Protect for IBM Cloud.....	43
7.3.8	IBM Cloud Kubernetes Service.....	43
7.3.9	LogDNA services: IBM Log Analysis and IBM Cloud Activity Tracker	47
7.3.10	IBM Cloud SQL Query	47

8 IBM HIPAA-Neutral Offerings	49
8.1 IBM Cloud Container Registry	49
8.2 IBM Cloud Identity and Access Management (IAM)	50
8.1.1 How IAM access policies provide access	50
8.1.2 Assigning access to access groups	50
8.3 IBM Cloud Load Balancer Options	51
8.4 IBM Cloud Security Advisor	51
9 Conclusion	52
10 Appendix A: HIPAA Definitions	53
11 Appendix B: HITECH	55
12 References and Resources	56
13 Disclaimers	57

1 Introduction

As organizations transition to the cloud, security is a top priority. With the rapidly expanding volume of personal information in the cloud, including Protected Health Information (PHI), it is critical to describe how the cloud is secured via critical services such as authentication, authorization, auditing, and end-client access.

This guide outlines how an IBM Cloud client can build environments and applications which are ready for HIPAA (Health Insurance Portability and Accountability Act).



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

This guide is for information purposes only and does not constitute legal or regulatory advice. IBM Cloud clients must independently analyze their environments and use cases to verify that their own control environment meets the HIPAA requirements.

This document is intended to discuss and assist clients with assembling a solution, considering such elements as: architecture design examples and descriptions of shared responsibilities across various stakeholders — healthcare entity, development firm, and cloud solution provider. This document also provides detailed features of IBM's public cloud portfolio and identifies those offerings which are HIPAA-ready, recognizing the client will ultimately decide on the components it wishes to utilize for its solution.

The intended audience of this guide is technical experts (e.g., solution architects and Chief Technology Officers) seeking to understand safeguards in IBM's cloud environment and services that are designed to meet the HIPAA requirements. Technical information is provided for Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) solution / product evaluators and decision makers. Readers should be familiar with HIPAA, and technical readers should have knowledge of IaaS and PaaS architecture.

2 HIPAA, HITECH and IBM Cloud Compliance

IBM designs certain offerings and services with the intention of making them HIPAA-ready. HIPAA is a 1996 law that was updated in 2009 via the Health Information Technology for Economic and Clinical Health (HITECH) amendment. For this paper, references to HIPAA are used as the aggregation of laws covering both HIPAA and HITECH.

IBM compliance processes assess, test, and confirm each IaaS and PaaS offering in the IBM Cloud Catalog that is identified as HIPAA-ready.

In order to be HIPAA-ready, offerings must be deployed in those IBM Cloud data centers which are themselves HIPAA-ready.

More details regarding HIPAA basics and definitions are available in [Appendix A](#).

3 Business Associate Relationships and Requirements

IBM is often the business associate of the client covered entity in cases where it receives PHI. IBM Cloud has policies and procedures to help IBM comply with its HIPAA obligations as a business associate, including cases where IBM stores and transmits PHI. **IBM's responsibility to the covered entity client is specified in the applicable Business Associate Agreement (BAA).**

A digital client can [configure an IBM Cloud account to utilize HIPAA-ready services](#). During that process, a client must accept an IBM BAA. The IBM Cloud BAA can be located on the [IBM SLA terms BAA page](#). Once configured, HIPAA-ready offerings are identified in the IBM Cloud Catalog to help clients know whether or not they have selected a HIPAA-ready offering.

IBM Cloud also executes BAAs with its vendors who qualify as business associates, requiring of them the same safeguards for HIPAA regulated data.

4 Introduction to IBM Cloud

IBM's public cloud is a suite of cloud computing services that offers an extensive array of IaaS and PaaS capabilities to help enhance the security, accessibility, and usability of clients' business-critical needs. IBM Cloud leverages strategic services from third-party IBM Business Partners. IBM terms do not govern or warrant the compliance claims by any third-party offering.

With IBM Cloud IaaS, organizations can deploy and access virtualized IT resources—such as compute, storage and networking resources—remotely using the internet. For compute, organizations can choose bare metal or virtual server instances.

With IBM Cloud PaaS, developers can use IBM services to create, deploy, run and manage various types of applications, including those used for HIPAA-compliant workloads. Developers can leverage various programming languages supported by IBM Cloud, including Java, Node.js, PHP and Python.

4.1 IBM's public cloud services and HIPAA

IBM Cloud is continuously deploying new and innovative services into the IBM Cloud Catalog. Based upon market demand, the capabilities of these services continue to be enhanced. Please note that not all offerings included in the catalog are necessarily HIPAA-ready.



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

A list of HIPAA-ready IBM Cloud services can be found at the IBM Cloud Compliance site located at <https://www.ibm.com/cloud/compliance/industry>. Other IBM Cloud services not listed may also be HIPAA-ready, have readiness in-progress, or have been deemed HIPAA-neutral. HIPAA-neutral means a capability which operates without implicating HIPAA. For instance, IBM Cloud has several PaaS services that are HIPAA-ready or may be HIPAA-neutral based on the inherent nature of the service.

IBM Cloud follows the HIPAA Security Rules, which are divided into 3 categories: administrative, physical & technical safeguards.

HIPAA Security Rules: HHS.GOV Safeguard Summaries

Administrative safeguards:

“administrative actions, and policies and procedures, to manage the selection, development, implementation, and maintenance of security measures to protect ePHI and to manage the conduct of the covered entity’s workforce in relation to the protection of that information”

Physical safeguards:

“physical measures, policies, and procedures to protect a covered entity’s electronic information systems and related buildings and equipment, from natural and environmental hazards, and unauthorized intrusion”

Technical safeguards:

“the technology and the policy and procedures for its use that protect ePHI and control access to it.”

Visit [HHS.GOV](https://www.hhs.gov) for these summaries and further details on the HIPAA Security Rule.

4.2 Information Security Policy & HIPAA Administrative Safeguards

When clients are creating a secure cloud solution that addresses HIPAA requirements, IBM recommends that clients adopt strong security policy and governance processes to mitigate risk and meet accepted standards for security and HIPAA readiness.

For successful cloud adoption, both clients and IBM should follow their respective cloud security policies. These security policies are often aligned to the cloud consumption and delivery models for IaaS, PaaS, and Software as a Service (SaaS).

IBM builds security into its cloud solutions. IBM maintains and follows IT security policies and practices that are integral to IBM's business and mandatory for all IBM employees. IBM reviews its IT security policies at least annually and amends such policies as IBM deems reasonable to maintain protection of cloud services and content processed therein. IBM employees complete security and privacy education annually and certify each year that they will comply with IBM's ethical business conduct, confidentiality, and security policies, as set out in IBM's Business Conduct Guidelines.

IBM Cloud meets strict industry security guidelines and policies as detailed at [IBM Cloud Architecture Center > Security policy and compliance > Policy, governance, risk, and compliance](#).

4.3 Physical Security of IBM Cloud & HIPAA Physical Safeguards

Adding to the HIPAA Security Rule’s administrative and technical safeguards, the physical security safeguards are another line of defense for protecting PHI.

A key component to secure and restrict physical access to PHI in a HIPAA-regulated environment is the security of the physical infrastructure and facilities that house the system. Physical security controls are in place for IBM Cloud as a cloud service provider. The IBM Cloud data center—from location and accessibility to power density and redundancy—is designed for security, resiliency, and efficiency. This section assists clients in understanding how IBM implements physical and environmental security at our global cloud data centers.

Because physical security in IBM Cloud is dependent on the underlying infrastructure, clients will want to understand how IBM implements physical and environmental security at all IBM Cloud data centers. Physical security for data centers is the protection of personnel, hardware, software, networks, and data from physical actions and events that could cause serious loss or damage to an enterprise, agency, or institution. Actions could arise due to human intervention or natural disasters.

IBM Cloud adopts several measures for increased physical security, adding enhanced protections for regulated environments:

- Physical security of the data center perimeter
- Entry and exit access controls and logging
- Secure offices, rooms, and facilities
- Protection against external and environmental threats
- Redundancy of power and network equipment
- Secure disposal of equipment during de-provisioning
- Corporate HR business policy and security for onboarding, training, and offboarding

Please refer to [IBM Cloud Architecture Center > Physical security architecture](#) for an overview of the physical security that is built into IBM Cloud data centers.

4.4 IBM Cloud & HIPAA Technical Safeguards

Healthcare applications that host HIPAA workloads must comply with HIPAA. [Section 6](#) of this guide includes use cases and architecture examples, as well as technical configuration information to help meet policies and procedures to protect PHI.

5 The Shared Responsibility Model

IBM proactively builds security into the cloud architecture. As the cloud provider, IBM Cloud secures the hardware and software of the cloud itself, as well as the physical security of the data centers. The Covered Entity is responsible for the security of their assets within the cloud and ensuring encryption of data in-transit and at rest, and, as applicable, selecting appropriate security.

As detailed earlier in [Section 3, “Business Associate Relationships and Requirements”](#), IBM puts a BAA in place with the Covered Entity and also with downstream suppliers who qualify as business associates.

Each party may have varying degrees of responsibility for addressing and/or mitigating specific controls; a party will be responsible (**R**), contribute or is consulted (**C**), or does not have any direct input and is simply informed (**I**).

The [HIPAA Security Rule](#) separates actions primarily into the three areas focused on in the following Responsibility Matrix: Technical Safeguards, Administrative Safeguards and Physical Safeguards. The tables in sections 5.1, 5.2, and 5.3 below are examples of shared responsibilities when using IBM Cloud. Specific responsibilities will vary based on each client’s HIPAA use case, and each service selected for each client’s use case.

5.1 Technical Shared Responsibilities

Specification	Description	IBM	Client
Access Control	Implement policies, procedures and technical controls for assigning a centrally controlled unique set of credentials for each user. Establish governance procedures for validation and authentication of an entity. Processes should include updating and terminating access to services that access PHI when appropriate and emergency access to PHI when needed.	R	R
PHI Authentication	Confirm whether PHI has been altered or destroyed in an unauthorized manner.	I	R
Encryption & Decryption	Implement capabilities for encrypting and decrypting PHI such that it cannot be viewed and/or modified by an unauthorized entity.	I	R
Audit Controls	Implement logging of attempted access to PHI and record what is done with the data once it has been accessed.	R	R
Automatic Session Timeouts	Implement controls to automatically log off active sessions after a specific period of time.	R	R
Backup & Restore	Implement adequate backup and restore of PHI so that in the case of a failure/disaster, PHI can be recovered and accessed in case of an emergency.	I	R
Integrity Controls	Implement security measures to ensure that electronically transmitted electronic protected health information is not improperly modified without detection until disposed of. Include period scans of the systems responsible for handling PHI for vulnerabilities.	I	R

(**R**) responsible (**C**) contributes/is consulted (**I**) no direct input / informed

5.2 Administrative Shared Responsibilities

Specification	Description	IBM	Client
Risk Management	Establish a program responsible for risk management/reporting and sanction policies associated with security breaches. Conduct periodic risk assessments of all areas which PHI is used and determine all the ways in which breaches to PHI could occur.	R	R
Workforce Security & Training	Establish a program to train and keep employees up to date regarding their responsibilities for safeguarding PHI. Establish onboarding and termination procedures to manage employees' access to systems that store, process, or transit PHI.	R	R
Contingency Planning	Develop and test on a periodic basis a contingency plan for accessing PHI in the case of an emergency, disaster, or outage.	C	R
Business Associates Agreement	Ensure a BAA is in place with every covered entity that will be handling PHI.	R	R

(R) responsible (C) contributes/is consulted (I) no direct input / informed

5.3 Physical Shared Responsibilities

Specification	Description	IBM	Client
Facilities Access	Control who has physical access to the location where PHI can be created, received, maintained, or transmitted, including software engineers, facility personnel, etc. Procedures must include safeguards to prevent unauthorized physical access, tampering, and theft.	R	I
Device Management	Define the policies and procedures for secure use of all devices (workstations, laptops, and mobile devices) that have access to PHI.	R	R
Inventory Management	Maintain a full inventory of the hardware and infrastructure that will handle PHI including maintenance records and records of the movements of each item. Copies of PHI and appropriate disposal must be done prior to movement of any equipment.	R	R
Disposal	Implement policies and procedures to address the final disposition of electronic protected health information, and/or the hardware or electronic media on which it is stored.	R	R

(R) responsible (C) contributes/is consulted (I) no direct input / informed

5.4 Responsibility Summary

IBM is responsible for the readiness and delivery of the IBM Cloud offerings designated as HIPAA-ready. Processes created on or above the IBM Cloud standard offering by the client or any third party are outside the scope of IBM Cloud's responsibility. For IaaS offerings (Bare Metal servers, Virtual Server Instances, VMware, etc.), the client is responsible for the implementation of HIPAA controls on the stack above the cloud services provisioned directly from IBM Cloud. It is the responsibility of the client to document and operate the service offering in a HIPAA compliant manner. Section 6.2.5 provides the details of IBM Cloud Infrastructure Services HIPAA Controls and Section 6.3.5 provides the details of IBM Cloud PaaS HIPAA Controls. Additional information is available in [IBM Cloud Docs > Shared responsibilities for using IBM Cloud products](#).

6 Architecture Examples

In this section, examples of IBM Cloud IaaS and PaaS architecture for HIPAA readiness are described. Sample use cases are provided to frame architecture components, as well as configuration recommendations to support HIPAA-regulated environments.

6.1 Account Setup and Managing User Identity, Access, and Authentication

Identity and Access Management (IAM) is crucial to managing access to secured IaaS and PaaS environments by granting access to authorized users while keeping intruders away. IBM Cloud IAM allows clients to easily create a new IBM Cloud account with resource access restriction and compliance auditing. To accomplish this, clients need a production-quality environment that is fully isolated from any development/QA environments and that includes IaaS and PaaS architecture.

Establishing proper user roles and permissions limits who can access resources. For example, clients can set policies such that only a particular user has administrative access to create virtual server instances (VSIs) or Kubernetes clusters. Different IAM roles can be defined for developers, operators, administrators, cluster administrators, and users.

- **Developer:** a programmer, for instance, who can develop the custom components of a HIPAA application.
- **Operator:** a multidisciplinary role that manages a client's cloud capabilities.
- **Administrator:** a role responsible for creating teams and assigning resources to other teams. Administrators have access to the resources that are assigned to a team by the cluster administrator.
- **Cluster Administrator:** this role can connect to an LDAP directory and add users and assign them to IAM roles. Cluster Administrators also create namespaces and manage workloads, infrastructure, and applications across all namespaces.
- **User:** A healthcare practitioner, for example, who accesses a patient's Electronic Medical Record (EMR).

Learn more about IAM via the [IBM Cloud Docs > Managing your account, resources, and access > Assigning access to resources by using access groups](#), > [Leveraging context-based restrictions to secure your resources](#), and > [Controlling access to resources by using tags](#).

For application user authentication, clients can secure resources and add authentication with IBM Cloud App ID. App ID helps protect the application by redirecting users to the authentication page. The client can use IAM service access roles to enable developers to perform tasks in App ID instances, such as configuring identity providers, managing users, customizing authentication UI and more.

For more details about App ID setup and configuration for HIPAA-ready usage, please refer to [Section 7.3.1, "IBM Cloud App ID"](#) in this guide.

6.2 IBM Cloud IaaS Architecture Examples as HIPAA-Ready

In this section, architecture examples of HIPAA-ready IBM Cloud IaaS offerings such as IBM Cloud Bare Metal servers, IBM Cloud Virtual Servers (also known as Virtual Server Instances (VSIs)) and IBM Cloud for VMware Solutions for clients are described. With bare metal servers and VSIs clients receive an operating system, configure network and storage options, and deploy their workload application onto IBM Cloud. IBM Cloud for VMware Solutions deliver a fully automated deployment of a software-defined platform based on a VMware validated design, providing virtualized compute, storage, and networking on single-tenant bare metal infrastructure.



The client is responsible for HIPAA readiness for their application / workload such as a healthcare portal application / database workload and for ongoing compliance. As discussed in [Section 4.2, “Information Security Policy & HIPAA Administrative Safeguards”](#), IBM manages the physical security of all data centers and infrastructure for IBM Cloud in a manner designed to restrict physical access to data.

6.2.1 IBM Cloud VSIs and IBM Cloud Bare Metal as HIPAA-Ready

The following use case leverages an example of a health clinic using virtual or bare metal servers to build a highly available and scalable web application to enable patients to view their laboratory results. The client’s application is a simple PHP frontend with a database.

If clients have lower cost and shorter-term usage needs for the workload, they create VSIs, install PHP and a database, and use cloud storage for persistent application files and database backups. When clients create a virtual server, they can choose between a public (multi-tenancy) environment or a dedicated (single tenancy) environment. Due to their single tenancy nature, dedicated hosts are a good choice as all resources are dedicated to and accessible by one user/application while remaining isolated from other IBM clients.

If the same workload requires high performance and long-term usage, bare metal servers are preferred as they provide the raw horsepower needed for processor-intensive and disk I/O-intensive workloads, including high performance, flexibility, on-demand provisioning, and control. Clients can compare IBM Cloud Bare Metal Servers vs. IBM Cloud Virtual Servers using [these IBM demos](#).

The architecture example below reflects both VSI and bare metal solutions for clients to select based on their needs. The services used in this example include:

- [Edge Services: IBM Cloud Internet Services](#)
- [IBM Cloud Load Balancers](#)
- [IBM Cloud Virtual Servers](#)
- [IBM Cloud Bare Metal](#)
- Storage including:
 - [IBM Cloud Block Storage](#)
 - [IBM Cloud File Storage](#)
- [IBM Cloud Direct Link "1.0 on Classic"](#)

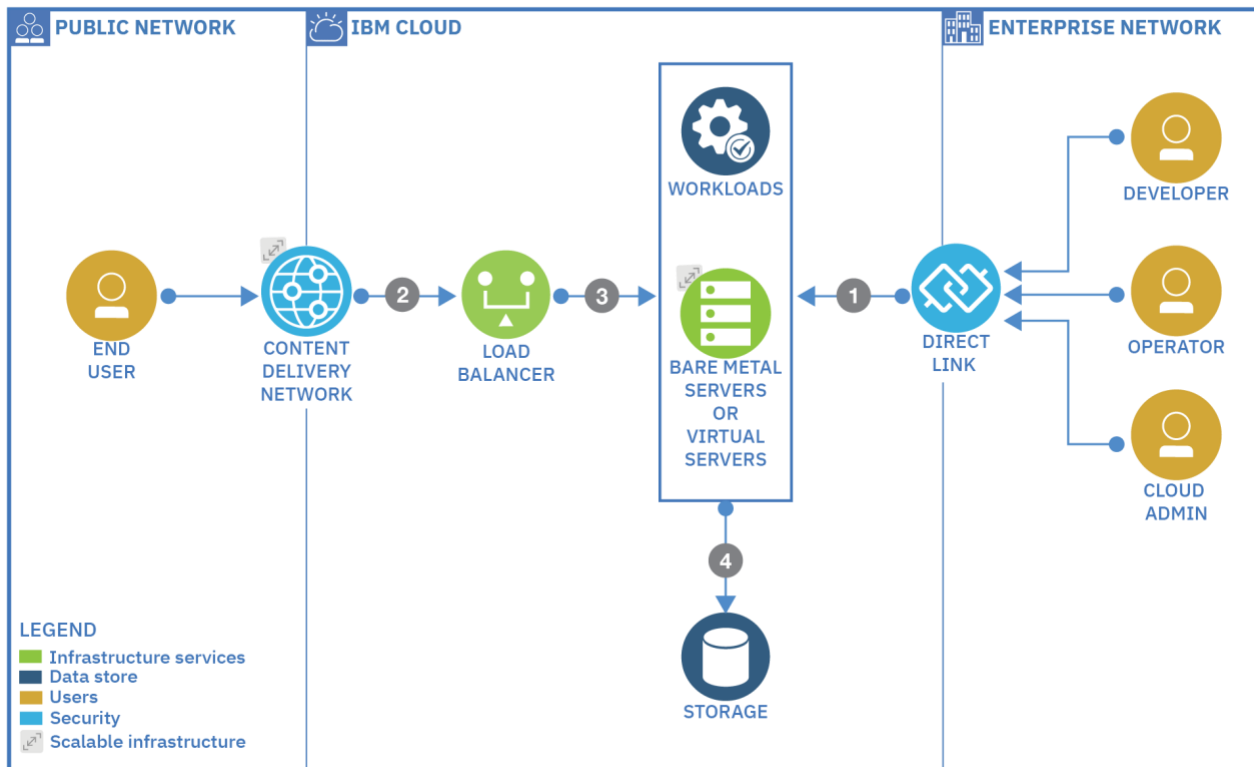


Figure 1: IBM Cloud Virtual Servers and IBM Cloud Bare Metal Server HIPAA-ready architecture example

The capabilities shown in the above diagram represent the following process example:

- 1 Roles including developers, operators, and administrators can connect through the Enterprise Network via IBM Cloud Direct Link over a multiple protocol layer switching (MPLS) routing technique. They can design their networking infrastructure using VSI or bare metal servers and then deploy the workload.
- 2 Users connect to the application to access lab result.
- 3 IBM Cloud Load Balancers are provisioned to distribute requests across application servers to load balance traffic within a location.
- 4 The client's selected server (VSI or bare metal) uses storage such as IBM Cloud Block Storage or IBM Cloud File Storage to persist application files and database backups.

The above solution may be extended to another location via Edge Services (IBM Cloud Internet Services) for increased resiliency and higher availability.

For more information, see [Section 7.2.1, "IBM Cloud Bare Metal"](#) and [Section 7.2.2, "IBM Cloud Virtual Servers"](#).

6.2.2 Using Firewalls and Load Balancers with VSIs or Bare Metal Server

Several firewall and load balancer options are available with VSIs and bare metal servers:

- IBM Cloud Internet Services (CIS): A client could use a CIS firewall (FW) and load balancer (LB). CIS also has content delivery network (CDN), distributed denial-of-service (DDoS), Internet Protocol Fire Wall (IP FW), global load balancing (GLB), and more. For more information, see [IBM Cloud Docs > CIS > Getting Started with IBM Cloud Internet Services](#).
- Clients can also explore additional firewall and load balancers options for their specific requirements.

For provisioning a load balancer server to distribute workloads across application servers, see [IBM Cloud Docs > Solution Tutorials > Use Virtual Servers to build highly available and scalable web app](#).

6.2.3 Using Storage with VSIs or Bare Metal Servers

When clients require additional, persistent storage, they can order IBM Cloud Block Storage and IBM Cloud File Storage (20 - 12,000 GB) when they provision a VSI or bare metal server. Clients need to connect the add-on storage after the server provisioning process is completed. Learn more about [Block Storage](#) and [File Storage](#) via IBM Cloud Docs.

6.2.4 Using VMware Virtualization Architecture

[IBM Cloud for VMware Solutions \(Dedicated options\)](#) delivers a fully automated deployment of a software-defined platform based on a VMware validated design, providing virtualized compute, storage, and networking. This platform is ordered, provisioned, and managed from both the IBM Cloud console as well as within the VMware environment. Using advanced automation and single-tenant bare metal infrastructure, the entire VMware environment is rapidly deployed to the IBM Cloud and made available to the client in a matter of hours. Additional capacity can be added (or contracted) to meet changes in workload demands. This virtualization architecture use case shown in the diagram below uses an example of a health clinic where a healthcare practitioner user accesses a patient's Electronic Medical Record (EMR) from workloads deployed in a portal application and database. The client is responsible for HIPAA readiness for their portal application / database workload.



Note, the newer IBM Cloud for VMware Solutions Shared offering has not yet been assessed for HIPAA readiness; only the IBM Cloud for VMware Solutions (Dedicated options) has been assessed and is asserted as HIPAA-ready.

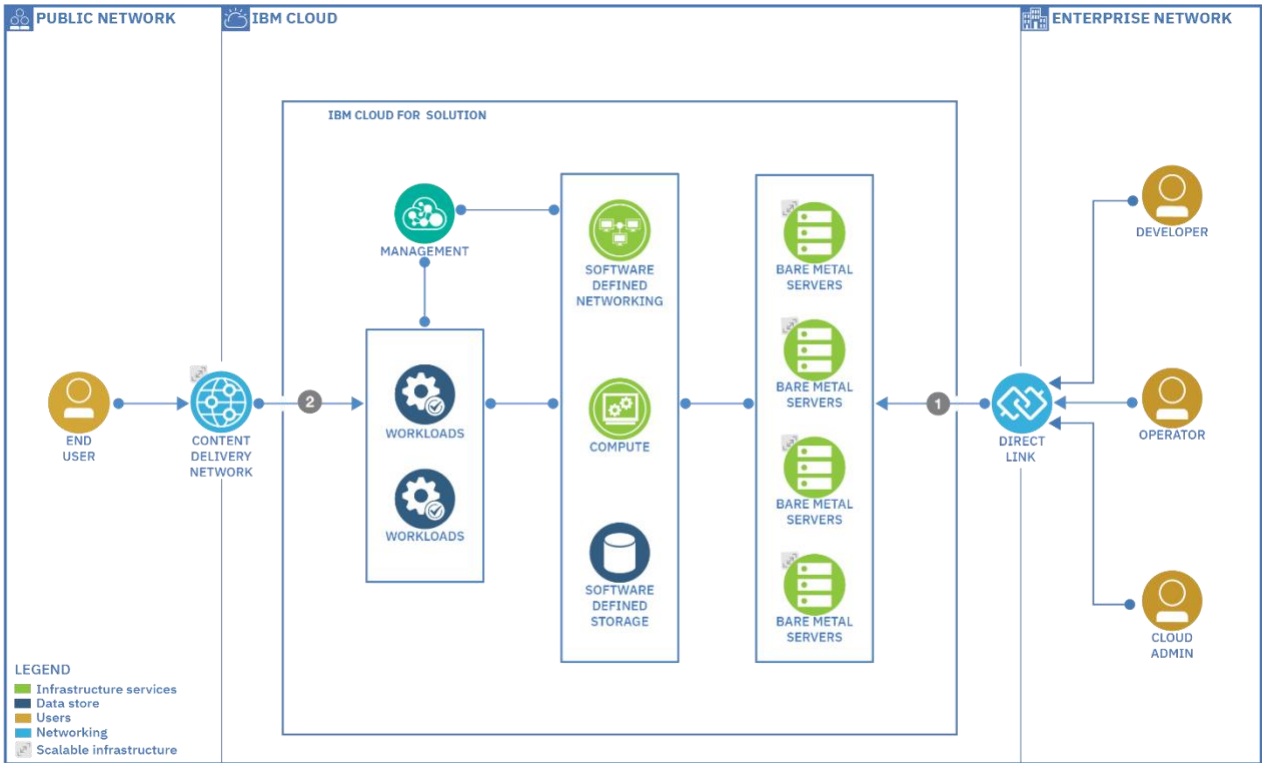








Figure 2: IBM Cloud for VMware Solutions (Dedicated options) virtualization architecture example for HIPAA Readiness

The capabilities depicted in the diagram above exemplify the following process:

- 1 Roles including developers, operators, and administrators can connect through the Enterprise Network via IBM Cloud Direct Link over a multiple protocol layer switching (MPLS) routing technique. They can design/automate/manage their IBM Cloud for VMware Solutions (Dedicated Options) and then deploy the Healthcare portal application and database workloads.
- 2 A user (healthcare practitioner) accesses a patient's Electronic Medical Record (EMR) via Edge Service (IBM Cloud Internet Services) from workloads deployed in a portal.

The various components of IBM Cloud for VMware Solutions (Dedicated options) are explained in detail as follows:

Architecture Item	Description
 MANAGEMENT	Management is a centralized platform for managing the entire software-defined data center. VMware vCenter Server provides a centralized platform for managing VMware vSphere environments, allowing clients to automate and deliver a virtual infrastructure across the hybrid cloud with confidence.
 WORKLOADS	Workloads refers to all client's portal application and database workloads.
 SOFTWARE DEFINED NETWORKING	Software Defined Networking provides a network overlay virtualizing the physical network to provide a large number of client-defined networks, intelligent network routing, and micro segmentation for enhanced firewall capabilities. Network virtualization platform (VMware NSX) for the Software-Defined Data Center (SDDC) delivers the operational model of a virtual machine for entire networks. With NSX, network functions including switching, routing, and firewalling are embedded in the hypervisor and distributed across the environment. This effectively creates a "network hypervisor" that acts as a platform for virtual networks and services.
 COMPUTE	<p>The Compute component enables many virtualized Linux and Windows servers to run concurrently on the same physical bare metal server, providing high levels of server utilization and capacity. VMware vSphere provides a powerful, flexible, and secure foundation for business agility that accelerates the client's digital transformation to hybrid cloud.</p> <p>With vSphere, clients can support new workloads and use cases while keeping pace with the growing needs and complexity of their infrastructure. vSphere is the heart of a secure SDDC, securing applications, data, infrastructure, and access.</p>
 SOFTWARE DEFINED STORAGE	Software Defined Storage as a local storage of the physical host is aggregated into a high performance, highly available software-defined SAN. VMware Virtual SAN powers industry-leading Hyper-Converged Infrastructure solutions with a vSphere-native, high performance architecture.
 BARE METAL SERVERS	IBM Cloud Bare Metal Servers provide a dedicated, single tenant basis for deploying the client's private infrastructure. Clients can locate their deployment in any of the dozens of IBM Cloud locations around the globe.

6.2.4.1 VMware Solutions Provisioning

IBM Cloud's VMware automation deploys and configures the entire standardized VMware software-defined stack, so it is ready for clients' use. For more detail see [IBM Cloud Docs > IBM Cloud for VMware Solutions > Getting started with IBM Cloud for VMware Solutions](#).

For IBM Cloud for VMware Solutions overview and provisioning, see [IBM Demos for VMware Solutions on IBM Cloud](#).

IBM Cloud hosts the network and bare metal servers in the underlying VMware instance and provides the ability to expand capacity by ordering additional bare metal servers or additional network attached storage (NAS). IBM provides basic monitoring of these resources, but the VMware instance is not offered as a fully managed service. Clients have the responsibility for ongoing configuration, security, management, and monitoring of all components of the instance.

For more information, see [IBM Cloud Docs > VMware Solutions > Post-deployment considerations for your VMware instance](#), which enumerates many of the responsibilities and activities client should plan for in order to fully operationalize their VMware instance.

For more information, see [Section 7.2.4 "IBM Cloud for VMware Solutions \(Dedicated options\)"](#).

6.2.5 IBM Cloud Infrastructure Services HIPAA Controls

IBM Cloud has instrumented controls in its platform to host healthcare application workloads that are designed to assist with HIPAA compliance. For HIPAA-ready IBM Cloud Infrastructure Services, examples of these controls are described below:



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

Control Focus Area	Examples
Secure, manage and fully document all access controls	<ul style="list-style-type: none">– For logical access, IBM Cloud Infrastructure Services has in place defined access management policies and procedures for administering logical access to the Infrastructure Services environment. Access is defined by role-based access control groups and two factor authentications.– Access requests to groups require documented approval by the employee's manager or a designated member of senior management.– Access to the Infrastructure Services environment requires employees to authenticate through bastion hosts prior to gaining access. Access is logged and monitored through a Security Information and Event Management (SIEM) tool. Logs are retained for one year.– IBM Cloud Infrastructure Services has implemented an Identity and Authentication Management Procedure which outlines requirements for passwords including complexity and when to change the password. To enforce password rotation and complexity requirements, a Lightweight Directory Access Protocol (LDAP) group policy is implemented to ensure passwords meet or exceed IBM Cloud Infrastructure Services' defined policies.– For physical access, IBM Cloud Infrastructure Services has defined physical access management policies and procedures for administering access to data center facilities.– Access is controlled through physical badge and only granted after documented management approval based on the employee's job role. Entry into and from the data centers is monitored 24/7 by video surveillance. Surveillance footage is retained for a minimum of 90 days. Visitors to the data center are required to sign in, provide government issued ID, and be escorted at all times by IBM Cloud personnel. Upon termination of an employee, their badge is physically collected, where possible, and the local site manager or employee's manager will contact the 3rd party facilities provider to disable the badge.

Control Focus Area	Examples
Implement a mechanism to encrypt electronic protected health information whenever deemed appropriate	<ul style="list-style-type: none"> – By default, Cloud Object Storage (COS) uses AONT (All or Nothing Transformation) AES 256-bit encryption when data is uploaded to a COS storage pool. – By default, client communication to and from COS public endpoints are encrypted using the TLS 1.2 protocol with a minimum of AES 128-bit encryption. – All other encryption, at rest, or in transit is the sole responsibility of the client
Audit Logging	<ul style="list-style-type: none"> – IBM requires that activity logs must be maintained for at least 9 months (270 days) from the date of activity. – IBM Cloud Infrastructure Services has defined audit logging and monitoring policies and procedures in place for the environment. – System components within the Infrastructure Services environment forward their logs to a Security Information and Event Management (SIEM) tool for long term storage analysis. Logs are maintained for one year. – IaaS system component logs do not contain PHI. – Clients are solely responsible for establishing audit logging and monitoring policies and procedures for their HIPAA related workloads and data. Log aggregation through a SIEM tool should be considered for log monitoring and long-term storage. Audit logs should be retained in accordance with applicable laws and regulations.

Additional information is available in [IBM Cloud Docs > Shared responsibilities for using IBM Cloud products](#).

6.3 IBM Cloud PaaS Architecture Examples as HIPAA-Ready

HIPAA may apply to Platform as a Service (PaaS) microservices handling PHI. This use case identifies and defines guidance to develop and deploy HIPAA-ready microservices.

As an example, let's define a patient/pharmacy monitoring use case where there are three microservices managed by a single IBM Cloud Kubernetes Service container cluster. These microservices include:

- Patient microservice in Patient Pod: Represents the patient record (Electronic Health or Medical Record (EHR or EMR)) and will consolidate and store all vital signs as part of the process.
- Monitor microservice in Monitor Pod: Collects real-time vital signs from the patient, e.g., blood pressure, heart rate, temperature.
- Pharmacy microservice in Pharmacy Pod: Consolidates the patient data such as medication information/history, allergies and receives and processes new and refill requests based on diagnosis.

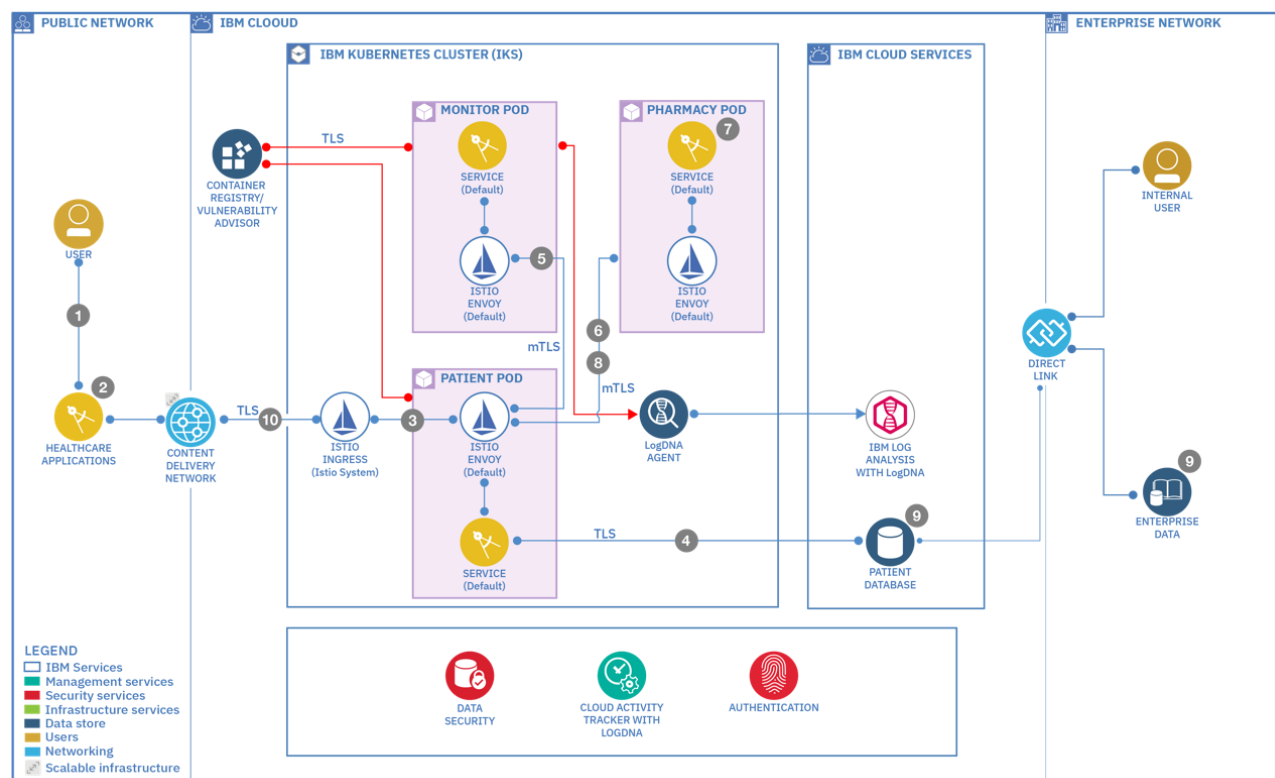


Figure 3: IBM PaaS Architecture example for HIPAA Readiness

The capabilities are described in the following step-by-step process:





- 1 A healthcare professional logs into the healthcare application based on the remote patient check-in schedule (e.g., twice a day).
- 2 A healthcare professional picks the first patient on the list – Ted Williams – and enters his name into the healthcare application.
- 3 The healthcare application places a request to the Patient microservice to get a current view of the patient.
- 4 The Patient microservice queries the Patient database (e.g., IBM Cloudant for IBM Cloud) to obtain the latest Electronic Medical Record (EMR) with history and medication schedule.
- 5 The Patient microservice makes a request to the Monitor microservice at the patient's home to obtain the latest vitals.
- 6 The Patient microservice combines all patient data into a single entity, and then submits to the Pharmacy microservice.
- 7 The Pharmacy microservice evaluates the patient history, medication schedule and current vitals, and suggests any medication changes, if needed.
- 8 The Pharmacy microservice generates the recommendation and returns the result to the Patient microservice.
- 9 The Patient microservice consolidates all data into an updated patient entity and updates the Patient database and transmits through the Enterprise Network to the Enterprise database for archiving purposes.
- 10 The Patient microservice returns the consolidated patient record with recommendations from the Pharmacy microservice to the healthcare applications alerting, if necessary, any interventions.
- 11 The healthcare professional takes the appropriate action.









6.3.1 IBM Cloud PaaS Architecture Components

Cloud-native microservice development requires a multi-party responsibility model. When considering microservices, IBM is the container service provider, and the client is responsible for the development, security and operational lifecycle of the application. IBM manages the security for the control plane, which includes all the functions and processes that determine which path to use. The client has the responsibility to secure the data plane, which includes all the functions and processes that forward packets/frames from one interface to another.

For this use case example, IBM, as the container service provider, recommends an architecture design using IBM Kubernetes Service (IKS) in IBM Cloud for running all 3 microservices.

The table below explains PaaS architecture components to build an IBM cloud environment for the use case example and identifies the IBM cloud service recommended to perform the function/task/service.

Architecture Item	Description
 HEALTHCARE APPLICATION	<p>Healthcare Application is a user interface (web browser or mobile based UI) where a healthcare professional can log in to access remote patient details for the use case.</p>
 EDGE SERVICES	<p>Edge Services provide the network capability to pass information through the internet (DNS, CDN, firewall, load balancer) to IBM Cloud.</p> <p>IBM Cloud Internet Services (CIS) provides a combined service for firewall, load balancer, CDN, DNS.</p>
 KUBERNETES CLUSTER	<p>IBM Cloud Kubernetes Service (IKS) provides a Kubernetes environment to create a cluster of compute hosts and deploy highly available containers. A Kubernetes cluster lets a client securely manage the resources that they need to quickly deploy, update, and scale applications.</p> <ul style="list-style-type: none"> Worker nodes carry the deployments and services that make up an application. When hosting workloads in the public cloud, clients want to ensure that an application is protected from being accessed, changed, or monitored by an unauthorized user or software. The client owns the worker node and is responsible for securing it. Clients should not run production workloads on free clusters. Clients can explore various firewalls to protect clusters. Learn more in IBM Cloud Docs > Classic: Opening required ports and IP addresses in your firewall. When data from IKS moves to a cloud service, security services for encryption such as TLS are used to protect data in motion. Using Transport Layer Security (TLS), clients can establish Mutual TLS (mTLS) between nodes, which is designed to protect PHI and safeguard theft and/or unauthorized modification of PHI outside the node.
 CONTAINER REGISTRY/ VULNERABILITY ADVISOR	<p>IBM Cloud Container Registry provides a multi-tenant, highly available, scalable, and encrypted private image registry that is hosted and managed by IBM. The client can use IBM Cloud Container Registry by setting up their own image namespace and pushing Docker images to their namespace. PHI must never be stored in the actual container image.</p> <p>The Vulnerability Advisor for IBM Cloud Container Registry checks the security status of container images that are provided by IBM third parties or added to a client's organization registry namespace.</p>

Architecture Item	Description
	<p>Ingress exposes HTTP and HTTPS routes from outside the cluster to services within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.</p>
	<p>Setting up proper user roles and permissions is key to limit who can access resources and the damage that a user can do when legitimate permissions are misused.</p> <p>Ingress exposes HTTP and HTTPS routes from outside the cluster to services within the cluster. Access to Ingress nodes is controlled via IBM Cloud Identity and Access Management (IAM).</p>
	<p>Clients access logs to troubleshoot problems and pre-empt issues. Clients configure the LogDNA agent on every worker (node) in a cluster and provision an instance of the IBM Log Analysis with LogDNA service.</p>
	<p>IBM Log Analysis with LogDNA service helps managing IKS logs (such as worker logs, pod logs, app logs, or network logs). These logs are then available for clients to analyze.</p>
	<p>The client can use HIPAA-ready database services available in the IBM Cloud Catalog, but need to encrypt PHI data in their application before storing in a database. Patient microservice consolidates all data into an updated patient entity and updates the Patient database.</p>
	<p>The connection to the enterprise network is established through the transformation and connectivity component. IBM Cloud Direct Link on Classic helps ensure the security of sensitive data to and from the IBM Cloud.</p>
	<p>Enterprise data is typically used by applications and users in an enterprise. Enterprise data is accessed over the direct link to the enterprise network. Clients can use their enterprise-specific on premises database in their enterprise network for archiving cloud patient database data.</p>
	<p>Data security discovers, categorizes, and protects cloud data and information assets with a strong focus on protection of data at rest or in transit. Data security services such as IBM Key Protect for IBM Cloud help key management. IBM Cloud Certificate Manager helps manage and deploy SSL/TLS certificates for client apps and services.</p>

6.3.2 Setting up a PaaS Environment

To set up and configure an IKS environment for managing microservices:

- Clients should establish a new IBM Cloud account to isolate development/QA environments from the production environment and manage user identity and access as suggested in [Section 6.1, "Account Setup and Managing User Identity, Access, and Authentication"](#).
- To create IKS clusters, follow the steps in [Section 7.3.8, "IBM Cloud Kubernetes Service"](#).
- To restrict access to the environment and secure the network, refer to [Section 7.3.8.6, "Restricting access to the environment and securing the network"](#).

6.3.3 IBM Kubernetes Service (IKS) Configuration

IKS provides many features for cluster components to allow clients to deploy containerized applications in a security-rich environment. Clients extend the level of trust in their cluster to ensure what happens within the cluster. For example, Figure 4 below demonstrates how clients can implement trust in their cluster in various ways.

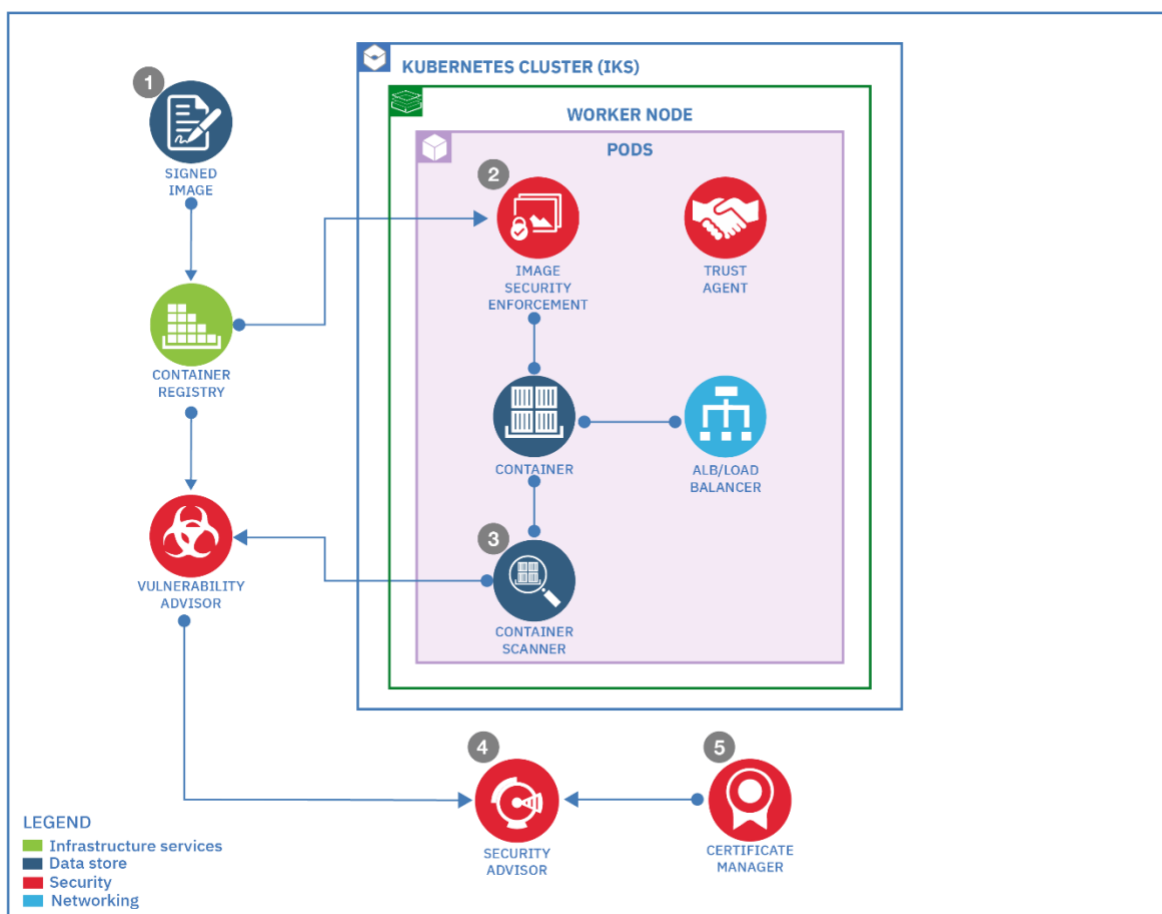


Figure 4: IKS configuration detail to enable trust in cluster

For more information about enabling trust in IKS clusters, see [IBM Cloud Docs > Security for IBM Cloud Kubernetes Service](#).

The steps for a runtime of images and containers are as follows:

- 1 Container Registry – Content Trust for client images:**
Help to protect the integrity of the images by enabling content trust in the [IBM Cloud Container Registry](#). With trusted content, the client can control who can sign images as trusted. After trusted signers push an image to the registry, users can pull the signed content so that they can verify the source of the image and this verification is required for HIPAA-ready applications. For more information, see [IBM Cloud Docs > IBM Cloud Container Registry > Signing images for trusted content](#).

- 2 Container Image Security Enforcement:**
With Portieris the client can enforce image security policies. Users can create image security policies for each Kubernetes namespace or at the cluster level, and enforce different rules for different images to meet [Vulnerability Advisor](#) policies or [content trust](#) requirements. If a deployment does not meet the policies that the client sets, security enforcement prevents modifications to their cluster. For more information, see <https://github.com/IBM/portieris>.

- 3 Vulnerability Advisor:**
By default, Vulnerability Advisor scans images that are stored in Container Registry to find potential security vulnerabilities. To check the status of live containers that are running in the cluster. For more information, see [IBM Cloud Docs > Managing image security with Vulnerability Advisor](#).

- 4 Security Advisor:**
With [IBM Cloud Security Advisor](#), the client can easily cohesively manage security on IBM Cloud workloads. The client will need to configure the appropriate security [alerts](#) for the environment. In IKS this may be related to high-risk behavior, logging changes, worker node updates, etc. For details on the available rule packages see [IBM Cloud Docs > Security and Compliance Center > Available Activity Insights rule packages](#).

- 5 IBM Cloud Certificate Manager:**
If the client wants to expose their application by using a custom domain with TLS, they can store their TLS certificate in Certificate Manager. Expired or about-to-expire certificates can also be reported in their [Security Advisor](#) dashboard. For more information, see [IBM Cloud Docs > Getting started with Certificate Manager](#).

6.3.4 Summary: Secure Flow for an IKS Application

No application architecture is complete without a clear understanding of potential security risks and how to protect against such threats. The following architecture and steps explain a security-oriented, end-to-end view that achieves these objectives:

- Encrypts content in storage buckets with their own encryption keys.
- Requires users to authenticate before accessing an application.
- Monitors and audits security-related API calls and other actions across cloud services.

The services used in the following architecture diagram include:

- [IBM Cloud Kubernetes Service](#)
- [IBM Cloud Container Registry](#)
- [IBM Cloudant for IBM Cloud](#)
- [IBM Cloud App ID](#)
- [IBM Cloud Object Storage](#)
- [IBM Cloud Activity Tracker](#)
- [IBM Key Protect for IBM Cloud](#)
- [IBM Cloud Certificate Manager](#)

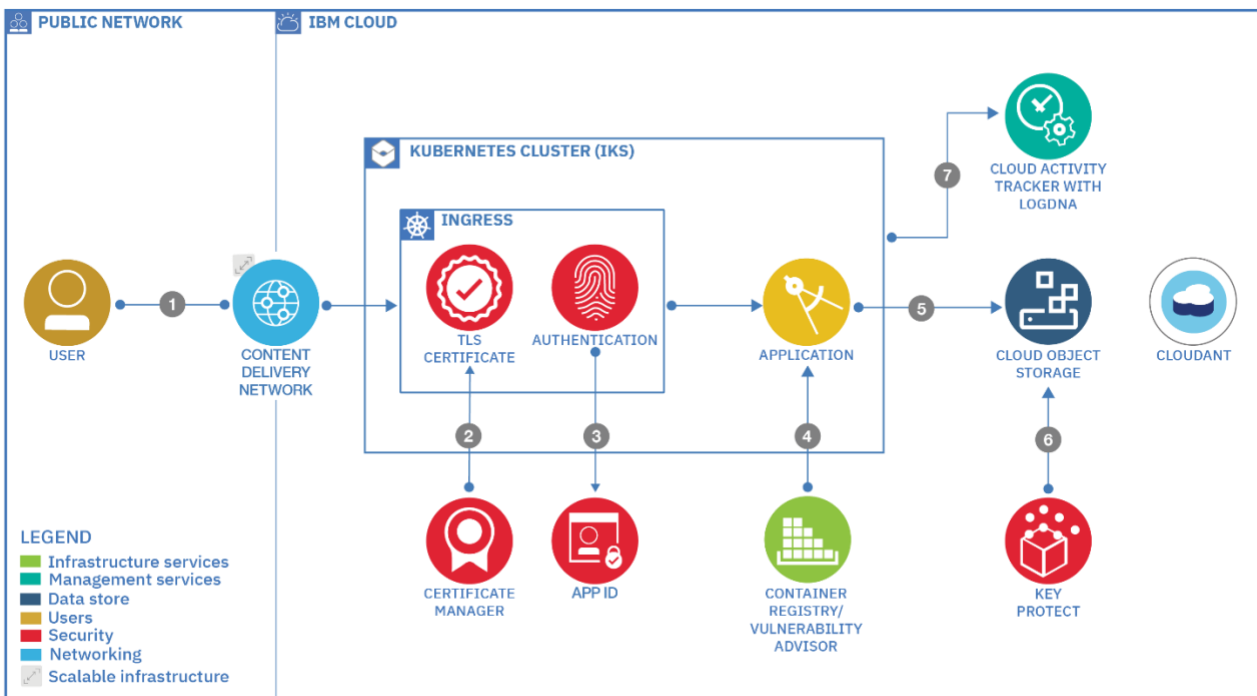


Figure 5: End-to-end secure flow for an IKS application

The steps to achieve this end-to-end view include:

- 1 User connects to the application.
- 2 If using a custom domain and a TLS certificate, the certificate is managed by and deployed from Certificate Manager.
- 3 App ID secures the application and redirects the user to the authentication page. Users can also sign up.
- 4 The application runs in a Kubernetes cluster from an image stored in the Container Registry. This image is automatically scanned for vulnerabilities.
- 5 Uploaded files are stored in Cloud Object Storage (COS) with accompanying metadata stored in IBM Cloudant.
- 6 File Storage buckets leverage a user-provided key to encrypt data.
- 7 Application management activities are logged by IBM Cloud Activity Tracker with LogDNA.

For more details, see the [IBM Cloud Docs > Solution Tutorials > Apply end to end security to a cloud application](#) tutorial.

6.3.5 IBM Cloud PaaS HIPAA Controls

IBM Cloud has instrumented controls in its platform to host healthcare application workloads which utilize microservices and must comply with HIPAA. For IBM Cloud PaaS HIPAA-ready offerings, examples of these controls are described below.



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

Control Focus Area	Examples
Secure, manage and fully document all access controls into microservices and image registries	<ul style="list-style-type: none">– IBM Cloud provides tools that can be used to implement access:<ul style="list-style-type: none">• based on job functions.• to roles that can be defined and assigned to individual members.• to enforce separation of duties aligned with best practices.• to only those who have a defined business need access.– IBM's public cloud clients are responsible for:<ul style="list-style-type: none">• defining roles granularly to support separation of duties and least privileged access.• following least privilege best practices to grant only the access required to perform the job role.• using minimum authorization to reflect data owner's approval of the access.• establishing password policies and enforcing for all containers (not just those running).
Define all communication in transit, that is, into, out of and in between worker nodes, including those running (healthcare) microservices	<ul style="list-style-type: none">– IBM Cloud clients are responsible for using the features of the platform to:<ul style="list-style-type: none">• secure all communications between microservices, regardless of data type (PHI) or microservice type (healthcare focused).• use highly secure protocols such as TLS and mTLS.• encrypt PHI itself within the encrypted TLS channels.

Control Focus Area	Examples
Ensure the integrity of containers running (healthcare) microservices	<ul style="list-style-type: none"> – To satisfy these rules, clients are responsible for: <ul style="list-style-type: none"> • not running containers as root. • enabling container images as (pre-) configured with minimum privileges and malware instrumentation for running containers. • ensuring that all running containers and microservices are subject to frequent vulnerability scans using available features, and findings are remediated as soon as possible according to local best practices.
Ensure the integrity of stored containers and images	<ul style="list-style-type: none"> – To satisfy these rules, clients are responsible for using the features and functions of IBM's public cloud to: <ul style="list-style-type: none"> • not run containers as root. • enabling container images as (pre-) configured with minimum privileges and malware instrumentation for running containers. • not storing container images containing PHI or end user data. • instrumenting containers on deployment with required security tools and verify it is not part of stored image. • ensuring that stored container images are kept up to date with patches, version levels, etc.

For further details, please refer to the U.S. Department of Health & Human Services HIPAA documentation available on <https://www.hhs.gov/hipaa>.

Additional information is available in [IBM Cloud Docs > Shared responsibilities for using IBM Cloud products](#).

7 IBM Cloud IaaS and PaaS Services and HIPAA

Previous sections of this guide laid the groundwork for sample use case architectures leveraging IBM services. This section dives deeper, with details and configuration guidance, for HIPAA workloads that may leverage IaaS and PaaS services available with IBM Cloud.

This section (Section 7) is for reference only. Use of any the services and features described below does not guarantee compliance with HIPAA requirements, but they can assist with HIPAA readiness. IBM Cloud clients are responsible for implementing the controls necessary to comply with HIPAA on an ongoing basis.



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

This guide is for information purposes only and does not constitute legal or regulatory advice. IBM Cloud clients must independently analyze their environments and use cases to verify that their own control environment meets the HIPAA requirements.

7.1 IBM HIPAA-Ready Services

The Pharmacy use cases discussed earlier leverage HIPAA-ready services. HIPAA-ready IBM Cloud services are listed on IBM's compliance webpage at <https://www.ibm.com/cloud/compliance/industry>. Additional IBM services may be HIPAA-ready but may not be listed on the IBM.com page.

In addition to considering secure engineering best practices, clients deploying IBM Cloud services for HIPAA workloads should keep in mind the following recommendations in this section.

7.2 IBM Cloud Infrastructure Services Offerings and HIPAA

Earlier in this guide, [Section 6.2, "IBM Cloud IaaS Architecture Examples as HIPAA-Ready"](#) provided use case and architecture examples for building an environment for HIPAA workloads using IBM Cloud Infrastructure Services. Additional details for HIPAA use cases for IaaS offerings follow below in this section.

7.2.1 IBM Cloud Bare Metal

[IBM Cloud Bare Metal](#) servers provide users with sole access to the entire server. IBM Cloud Bare Metal servers can be acquired in a preconfigured form or custom-configured to exact specifications.

Data transfers which include PHI must be encrypted in transit. For traffic between external sources (such as the internet or a traditional IT environment) and IBM Cloud, clients should use industry-standard transport encryption mechanisms such as TLS or IPsec virtual private networks (VPNs), consistent with the [HHS.gov "Guidance to Render Unsecured Protected Health Information Unusable, Unreadable, or Indecipherable to Unauthorized Individuals"](#).

The client retains all responsibility for configuring any resource provisioned onto IBM Cloud Bare Metal servers in order to meet HIPAA requirements.

A use case example including IBM Cloud Bare Metal servers is featured earlier in this guide in [Section 6.2.1, "IBM Cloud VSIs and IBM Cloud Bare Metal as HIPAA-Ready"](#).

For more information on provisioning bare metal servers, see the [IBM Cloud Docs > Bare Metal Servers > Getting started tutorial](#).

7.2.1.1 Encryption

There are several different options available on IBM Cloud to encrypt client data. It is recommended that clients use one or more of the following methods for HIPAA workloads:

- Local Hard Drive options that support SED or “Self-Encrypting Drives.”
- Block storage volumes that have encryption capabilities.
- Cloud Object Storage that supports BYOK for encryption.

7.2.1.2 Access and Authentication

Clients are advised to ensure login credentials on all servers are complex and changed periodically. Additionally, it is suggested that remote access (Windows® Remote Desktop Protocol (RDP) and Secure Shell (SSH) be disabled over the public network.

7.2.1.3 Storage Recommendations

It is recommended that data is stored on IBM Cloud is always encrypted via one of the encryption methods listed above.

7.2.2 IBM Cloud Virtual Servers

[IBM Cloud Virtual Servers](#) are scalable and come with dedicated core and memory allocations. The hypervisor is fully managed by the IBM Cloud. A client can perform configuration and management tasks by using both the IBM Cloud client portal and the API. Virtual servers are deployed to the same VLANs as bare metal servers, enabling workloads to be spread across virtual servers and bare metal servers while maintaining interoperability.

The client retains all responsibility for configuring any resource provisioned onto IBM Cloud VSIs in order to maintain HIPAA compliance. IBM clients use the IBM Cloud Console to provision and configure the public virtual server instance.

A use case example including IBM Cloud Bare Metal servers is featured earlier in this guide in [Section 6.2.1, "IBM Cloud VSIs and IBM Cloud Bare Metal as HIPAA-Ready"](#).

To learn more about how to provision and configure VSI servers, see [IBM Cloud Docs > Virtual Servers for Classic > Provisioning Public Instances](#).

7.2.2.1 Encryption

There are several different options available on IBM Cloud to encrypt client data. It's recommended that clients use one or more of these methods to ensure compliance, provided that ultimately the client will determine what it feels best on needs and compliance assessment.

- Encrypted VDI – Encrypts client data at the hypervisor.
- Block storage volumes with encryption capabilities.
- Cloud Object Storage that supports BYOK for encryption.

7.2.2.2 Access and Authentication

A use case example including IBM Cloud Bare Metal servers is featured earlier in this guide in section [6.1 Account Setup and Managing User Identity, Access, and Authentication](#).

Clients are advised to ensure login credentials on all servers are complex and changed periodically. Additionally, it is suggested that remote access (Windows® Remote Desktop Protocol (RDP) and Secure Shell (SSH) be disabled over the public network.

7.2.2.3 Storage Recommendations

It is recommended that data is stored on IBM Cloud is always encrypted via one of the encryption methods listed above.

7.2.3 Storage (Block Storage, Cloud Object Storage, and File Storage)

IBM Cloud includes options for Block Storage, Cloud Object Storage (COS), and File Storage. The client retains all responsibility for configuring any resource provisioned onto IBM Cloud storage options in order to maintain HIPAA compliance.

Storage options were mentioned earlier in this guide in [Section 6.2.3, “Using Storage with VSIs or Bare Metal Servers”](#).

7.2.3.1 IBM Cloud Block Storage

[IBM Cloud Block Storage](#), sometimes referred to as block-level storage, is a technology that is used to store data files on Storage Area Networks (SANs) or cloud-based storage environments. Developers favor block storage for computing situations where they require fast, efficient, and reliable data transportation.

Additional Block Storage details for HIPAA use cases:

- By default, there is encryption of data using AES-256 standard, including snapshots and replicas of encrypted volumes, and encryption keys are managed in-house using industry standard Key Management Interoperability Protocol (KMIP). For more information, see [IBM Cloud Docs > Block Storage - Classic > Securing your data in Block Storage](#).
- Clients using RedHat Enterprise Linux (RHEL) can enable full disk encryption of block devices with Linux Unified Key Setup-on-disk-format (LUKS). This aids in protecting the contents on mobile devices and removeable media. For more information, see [IBM Cloud Docs > Block Storage - Classic > Achieving full disk encryption with LUKS in RHEL6](#).

For more information on configuring Block Storage, see [IBM Cloud Docs > Block Storage - Classic > Getting started with Block Storage](#).

7.2.3.2 IBM Cloud Object Storage

[IBM Cloud Object Storage](#) makes it possible to store practically limitless amounts of data, simply and cost effectively. It is commonly used for data archiving and backup, for web and mobile applications, and as scalable, persistent storage for analytics. Flexible storage class tiers with a policy-based archive help clients effectively manage costs while meeting data access needs.

Additional IBM Cloud Object Storage details for HIPAA use cases:

- **Encryption of data:** By default, all objects stored in IBM Cloud Object Storage are protected at-rest by using randomly generated keys and an all-or-nothing-transform (AONT). If it is necessary to control encryption keys, root keys can be provided on a per object basis as detailed in [IBM Cloud Docs > Cloud Object Storage > Server-Side Encryption with Customer-Provided Keys \(SSE-C\)](#). Additionally:
 - For more granular control over the encryption keys, you can leverage COS integration with IBM Key Protect for IBM Cloud. By using Key Protect, you can enable the security benefits of Bring Your Own Key (BYOK) by importing your own root of trust encryption keys, called Customer Root Keys (CRKs), into the service. With Key Protect, you can use a CRK to wrap (encrypt) and unwrap (decrypt) the Data Encryption Keys (DEKs) that are associated with your data resources, so you control the security of your encrypted data in the cloud. COS buckets should be encrypted and protected using IBM Key Protect for IBM Cloud.
 - Key Protect does not process Personal Health Information (PHI). However, Key Protect keys can be used to encrypt/decrypt data encryption keys used by data services to protect PHI data. For more information about Key Protect, see [Section 7.3.7, “IBM Key Protect for IBM Cloud”](#).
 - **Setting bucket level permissions:** IBM Cloud Identity and Access Management (IAM) access policies and credentials management can be used to control access to the individual COS buckets which are used to create logical segregation of objects stored. Bucket-level permissions can be set via UI or API to grant specific access roles to certain users.
 - **Setting COS bucket firewall:** COS provides the ability to restrict access to buckets by using a bucket-level firewall that will only allow access if the request originates from a trusted network. Access can be restricted to a specific IP address within your network. Read more about this feature at [IBM Cloud Docs > Cloud Object Storage > Setting a firewall](#).
 - **Backup:** Periodic backup and restore functionality needs to be in place to ensure appropriate data availability. To achieve a zero RPO, for example, it requires a mirrored data store with real time replication.

For more information on configuring Cloud Object Storage, see:

- [IBM Cloud Docs > Cloud Object Storage > Getting started with IBM Cloud Object Storage](#)
- [IBM Cloud Docs > Cloud Object Storage > Your responsibilities when using IBM Cloud Object Storage](#)

7.2.3.3 IBM Cloud File Storage

With [IBM Cloud File Storage](#), clients can deploy and customize flash-backed NFS-based file storage from 25 GB to 12,000 GB capacity with up to 48,000 IOPS, and increase storage capacity or adjust performance on the fly to quickly adjust to changes in workload demands.

By default, there is encryption of data using AES-256 standard, including snapshots and replicas of encrypted volumes, and encryption keys are managed in-house using industry standard Key Management Interoperability Protocol (KMIP). For more information, see [IBM Cloud Docs > File Storage - Classic > Securing your data in File Storage](#).

For more information on configuring File Storage, see [IBM Cloud Docs > Getting started with File Storage](#).

7.2.4 IBM Cloud for VMware Solutions (Dedicated options)

With [IBM Cloud for VMware Solutions \(Dedicated options\)](#), a client can take advantage of virtualization to transform data centers into simplified cloud computing infrastructures and enable their organizations to deliver flexible and reliable services. VMware virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the data center. Clients can deploy their applications that include PHI data on IBM Cloud for VMware Solutions, but clients are responsible for HIPAA readiness for their application / workload.

IBM Cloud for VMware Solutions (Dedicated options) includes VMware's vCenter Server, which offers VMware's vSphere hypervisor, VMware vCenter Server, VMware NSX, and optionally VMware vSAN.

There are two client installations options — vCenter and vSphere:

1. **vCenter:** VMware vCenter Server on IBM Cloud offers automated installation and configuration of these VMware virtualization components on IBM Cloud bare metal servers. IBM and VMware have worked closely to standardize and validate the architecture of this offering according to best practices.
2. **vSphere:** For enterprises seeking full flexibility or configuration Control, vSphere may be suitable. A client can take advantage of a customizable virtualization service that combines VMware-compatible bare metal servers, hardware and licenses to build your own IBM-hosted VMware environment.

A use case example including IBM Cloud for VMware Solutions is featured earlier in this guide in [Section 6.2.4, "Using VMware Virtualization Architecture"](#). Additional information for HIPAA use cases:

- For vCenter and vSphere:
 - Firewalls: Clients are accountable for the firewalls created and any communications between IBM Cloud for VMware components.
 - Encryption: Clients need to ensure appropriate encryption of the data in their system.
- Additional information for vCenter:
 - Automation: Upon initial deployment, the offering's automation and client's account are separate. The client is allowed to and responsible for changing all passwords that IBM provides in the portal.
 - Access and credentials are created during initial deployments and provided to the client. As part of the requirement of the offering, IBM Support must retain full access to the management layer to provide lifecycle management as well as support to clients.

Further details related to architecture and responsibilities can be obtained at [IBM Cloud Docs > Customer versus IBM responsibility for vCenter Server](#).

7.2.5 IBM Cloud Direct Link (“1.0 on Classic”)

[IBM Cloud Direct Link on Classic](#) helps ensure the security of sensitive information such as PHI between the client’s enterprise network and the IBM Cloud. It is implemented to support private, hybrid, and cross-provider workloads, as well as large or frequent data transfers.

A use case example including IBM Cloud Direct Link is featured earlier in this guide in [Section 6.2.1, "IBM Cloud VSIs and IBM Cloud Bare Metal as HIPAA-Ready"](#).

IBM Cloud Direct Link is configured with an isolated virtual routing and forwarding (VRF) table that effectively removes it from the Internet. Information and restrictions regarding the use of VRF are available at [IBM Cloud Docs > Direct Link on Classic \(1.0\) > Configuring Direct Link on Classic](#).

For more information on configuring Direct Link, see [IBM Cloud Docs > Direct Link on Classic \(1.0\) > Getting started with IBM Direct Link on Classic](#).



Note, the newer IBM Cloud Direct Link “2.0” offerings have not yet been assessed for HIPAA readiness. Only the IBM Cloud Direct Link “on Classic” offerings have been assessed and are identified as HIPAA-ready.

7.2.6 Hardware Security Module

[IBM Cloud Hardware Security Module](#) is a centralized, high-assurance capability for cryptographic processing, key generation, and key storage. As client data will likely contain PHI, clients can use HSM which is designed to protect the cryptographic infrastructure by securely managing, processing, and storing cryptographic keys inside a hardened, tamper-resistant device. The IBM Cloud HSM offering uses the network-attached general purpose SafeNet Luna Network HSM, manufactured by Thales (formerly Gemalto). A FIPS 140-2 Level 3 validated, single-tenant, password-authenticated device, this HSM is made available by IBM Cloud for client use and configuration.

7.2.6.1 HSM Password and Authentication

IBM employs and makes available extensive password and authentication capabilities:

- **Enablement of TLS Ciphers:** The SafeNet Luna Network HSM uses a default set of cipher suites for TLS communications, such as client connections; if the default list is not suitable, it can be modified.
- **Appropriate setting of the System Date and Time:** Functionality is available to set the date and time manually using the appliance's internal clock, or by synchronizing the appliance with a network time protocol (NTP) server. NTP provides a reliable, consistent, and accurate timing mechanism using Coordinated Universal Time (UTC) and is the recommended option for providing an accurate date and time. Accurate time is paramount for security auditing and troubleshooting of logs.
- **Generation of the HSM Server Certificate:** IBM recommends generating a new HSM server certificate before placing the HSM in service. Do not use the default certificate generated at the factory but regenerate the server certificate once the HSM is in service. If there is a need to generate a new certificate, update client NTLS links to use the new certificate.

- **Binding the NTLS or SSH Traffic to a Device:** It is possible to configure the service to restrict NTLS or SSH traffic to a specific network device (or IP address for SSH traffic):
 - NTLS is used to securely transport the cryptographic messages exchanged between a client and the HSM across the network. Bind the NTLS traffic to a specific network device, a bonded network device, or all network devices.
 - SSH is used to securely transport the administrative messages exchanged between LunaSH and the appliance or HSM across the network. By default, SSH traffic is unrestricted. SSH binding is optional.

7.2.6.2 HSM Initialization

Initialization prepares a new or existing HSM for use. There is a need to initialize the HSM before objects can be generated or stored. The following are key steps:

- **Creation of a Network Trust Link Between the Client and the Appliance:** Leverage the cryptographic resources to create a secure Network Trust Link (NTL). After that, it is possible to configure links to individual partitions on the appliance using NTL or Secure Trusted Channel (STC).
- **Creation of a Secure Trusted Channel (STC) Link Between a Client and a Partition:** If there is a need for higher level security for the network links than is offered by NTLS, utilize the STC to provide secure client-partition links. STC offers the following features to ensure the security and integrity of client-partition communications:
 - All data is transmitted using symmetric encryption; only the endpoints can decrypt message
 - Message authentication codes prevent an attacker from intercepting and modifying any command or response
 - Mutual authentication of the HSM and the endpoint ensure that only authorized entities can establish an STC connection
 - Configuring the SafeNet Luna Network HSM appliance to use a Network Time Protocol (NTP) server

For more information, please see [IBM Cloud Docs > Getting started with IBM Cloud HSM](#).

7.3 IBM Cloud PaaS Offerings and HIPAA

Earlier in this guide, [Section 6.3, “IBM Cloud PaaS Architecture Examples as HIPAA-Ready”](#) provided use case and architecture examples which may help a client build a HIPAA-ready environment using IBM Cloud PaaS offerings. Additional details for HIPAA use cases for PaaS offerings follow in this section.



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

This guide is for information purposes only and does not constitute legal or regulatory advice. IBM Cloud clients must independently analyze their environments and use cases to verify that their own control environment meets the HIPAA requirements.

7.3.1 IBM Cloud App ID

[IBM Cloud App ID](#) allows clients to easily add authentication controls to web and mobile applications with zero code changes and no redeploy required. It enhances client applications with advanced security capabilities such as multi-factor authentication (MFA), single sign-on (SSO) and user-defined password policies. Clients can also use App ID's scalable user registry to let users manage their own accounts.

App ID was included in [Section 6.3.4, “Summary: Secure Flow for an IKS Application”](#). App ID secures applications—including those with PHI data—and redirects the user to the authentication page. There are extensive advanced security features to strengthen the security of an application. These include:

- **Multi-factor authentication:** A capability to confirm a user's identity by requiring the user to enter a one-time passcode that is sent to their email or via SMS in addition to entering their email and password.
- **Password policy management:** Capability to enforce more secure passwords for Cloud Directory by configuring a set of rules that user passwords must conform to. Examples include the number of attempted sign-ins before lockout, expiration times, minimum time span between password updates, or the number of times that a password can't be repeated.
- **Transport level encryption:** Any connections which are not HTTPS and TLS 1.2 and above will be rejected. This ensures that the communication channel between an application and the App ID service is using the highest level of encryption
- **Security Assertion Markup Language (SAML) signature and encryption:** By signing and encrypting SAML payloads for both requests and responses, the application can enable an additional level of validating component authenticity and preventing man-in-the-middle attacks.
- **Logging and Auditing:** IBM Cloud Activity Tracker with LogDNA can be used to review and build charts for management and runtime activity. Audited events include configuration changes, user management, user logins, password resets and more.
- **User profile updates from applications:** The developer or operator can decide whether end-users should be able to up make updates to their profiles using client applications. Users should not update their profiles; clients should use the App ID dashboard to disable this capability.
- **Service Access Roles:** The client can use IBM Cloud IAM service access roles to enable developers to perform tasks in App ID instances, such as configuring identity providers, managing users, customizing authentication UI and more.

For more information, please see [IBM Cloud Docs > IBM Cloud App ID > Getting started with App ID](#).

7.3.2 IBM Cloud Certificate Manager

[IBM Cloud Certificate Manager](#) helps clients manage and deploy SSL/TLS certificates for client applications and services, providing a secure repository for certificates and their associated private keys. It helps prevent outages by sending notifications when certificates are about to expire.

Secure Sockets Layer / Transport Layer Security (SSL/TLS) is a standard security protocol for establishing an encrypted link between a server and a client. SSL certificates are an essential component of SSL/TLS to make internet transactions secure. The SSL certificate's job is to initiate secure (SSL) sessions and act like digital passports to provide authentication to protect the confidentiality, integrity of the communications.

Certificate Manager was included earlier in this guide in [Section 6.3.4, “Summary: Secure Flow for an IKS Application”](#). Additional information for HIPAA use cases:

- Sensitive healthcare data is vulnerable if it is transmitted in open networks. HIPAA readiness requires PHI, whether in transit or at rest, to be secured. SSL allows sensitive information such as health information to be transmitted securely.
- A client may use the following features of Certificate Manager to strengthen the security for the HIPAA-enabled applications.

Feature	Benefit
Order SSL/TLS Certificates	Easily order and renew public SSL certificates from external certificate authorities.
Manage SSL Certificates	Import third-party issued SSL certificates into a central repository, store them securely, track their expiration and usage, set permissions on who can use them, and download them to deploy to the IBM Cloud resources.
Get Notified about Expiring Certificates	The client should configure Certificate Manager to send notifications well in advance of certificate expiration to avoid outages. The expiration notification can be configured to trigger an automated process to renew and deploy certificates.
Service Access roles	Use service access roles in Certificate Manager to enable approved clinical healthcare users to complete tasks in Certificate Manager instances, such as importing, downloading, editing, or deleting certificates.
Activity Tracking	Certificate Manager is integrated with the IBM Cloud Activity Tracker with LogDNA service to track how users and applications interact with the Certificate Manager service in IBM Cloud. Authentication and authorization are components to ensure appropriate logging of the access of PHI data. A covered entity client can monitor the activity of the IBM Cloud account. A client can investigate inappropriate activity and actions to comply with regulatory audit requirements. The events that are collected comply with the Cloud Auditing Data Federation (CADF) standard .
Availability / Data Residency	The Certificate Manager service is highly available in multiple regions. In each supported location, the service exists in multiple availability zones with no single point of failure. This is critical for disparate and de-centralized healthcare facilities in multiple regions that need to ensure uptime and resiliency. Data that is stored in the Certificate Manager database is backed-up daily. If recovery of a location is required, the data is available to be restored. A client can decide on data residency considerations and create a Certificate Manager instance in the required region.
Key Management	Any keys stored in the services should be managed according to key management best practices. For instance, access policies for certificates and keys should be configured. All access to and activities regarding keys and certificates should be audited, with the logs retained for at least a year.

For more information Certificate Manager:

- See [IBM Cloud Docs > Certificate Manager > Managing security and compliance with Certificate Manager](#) to learn more on the security features that Certificate Manager provides.
- See [IBM Cloud Docs > Certificate Manager > Getting started tutorial](#) details on configuring IBM Cloud Certificate Manager.

7.3.3 IBM Cloud Databases

[IBM Cloud Databases \(ICD\)](#) offer a complete database portfolio for data and analytics managed by clients. This is the core service that houses sensitive healthcare data such as patient lab results, medication, treatment, etc.

ICD is a collection of databases that address the data-intensive needs of application developers, data scientists and IT architects to deliver immediate and longer-term benefits:

- Build and access through a public cloud
- Allow enterprise users to host databases without buying dedicated hardware
- Managed by the user or offered as a service and managed by a provider
- Support SQL and NoSQL databases
- Access through a web interface or vendor-provided API

7.3.3.1 IBM Cloudant for IBM Cloud

[IBM Cloudant for IBM Cloud](#) provides a fully managed, distributed JSON document database optimized for handling heavy workloads that are typical of large, fast-growing web and mobile apps. Cloudant elastically scales throughput and storage independently to meet application requirements.

IBM Cloudant was included earlier in this guide in [Section 6.3.4, “Summary: Secure Flow for an IKS Application.”](#) The following are designed to store PHI data securely in an IBM Cloudant database, which the client should consider:

- Encryption of sensitive data is critical and should be invisible to IBM Cloudant. While IBM does not access any client data as a matter of policy and all data is encrypted at rest, IBM operations teams may occasionally be exposed to client data during the provision of services. For data which must remain confidential, encrypting data before sending it to any cloud database service is a best practice.
- For logging purposes, IBM Cloud Activity Tracker with LogDNA and IBM Log Analysis with LogDNA are integrated with IBM Cloudant to provide auditing and logging access for the service. Consider archiving logs into IBM Cloud Object Storage to retain long term access to the information with ad-hoc analytics through IBM Cloud SQL Query.
- PHI should not be logged and/or otherwise stored in the files in a container based microservice. The container should be provisioned with an encrypted file system, and the microservice (database and/or data storage service) needs to ensure that PHI is adequately encrypted and protected from theft and unauthorized modification at the application level.
- While Cloudant is highly available across 3 data centers in a multi-zone region, further application availability can be created by deploying two or more Cloudant instances in other IBM Cloud regions and setting up real-time replication between them.

More information on IBM Cloud security is available via [IBM Cloud Docs > Cloudant > Security](#).

7.3.3.2 IBM Cloud Databases (DataStax, Elasticsearch, EnterpriseDB, etcd, MongoDB, PostgreSQL, Redis, and Messages for RabbitMQ)

[IBM Cloud Databases](#) is a fully managed collection of databases available through a consistent consumption, pricing, and interaction model. For HIPAA use cases, the following should be considered:

- **Data Access:** While IBM does not access any client data as a matter of policy and all data is encrypted at rest, IBM operations teams may very occasionally be exposed to client data during the operations of the service. For data which **must** remain confidential to a client company, encrypting data before sending it to any Cloud Database service is best practice. Where sensitive client personal information is stored, rather than encrypting the data, a client may wish to instead pseudonymize this data to protect the identity of the persons involved while allowing the cloud service to process and analyze the data.
- **Database Access:** The IBM Cloud Databases detailed in this document support the ability to connect over the IBM Cloud private network through Cloud Service Endpoints. It is recommended to solely use private database endpoints and set appropriate IP whitelisting on client databases for each database instance. IBM Cloud currently allows a client to select “Private network only” or “Public network only”. Therefore, IBM recommends selecting “Private network only” which automatically shuts off the public endpoint.
- **Database Upgrades:** Clients should ensure regular upgrading and maintenance of their database to the latest supported version. IBM recommends pro-active planning for major version upgrades, as they may include changes that impact APIs. IBM Cloud Databases seek to retire major versions after they become end of life. Clients will then stop receiving security and performance patches and remediation.
- **Encryption:** By default, all database instances and backups are encrypted at rest and enforce connecting to the database over TLS. IBM Cloud Databases allow clients to bring their own encryption key through IBM Key Protect for IBM Cloud for data at rest and backups, at the time of provisioning.
- **Credential Rotation:** IBM Cloud Databases employ two layers of access control. The first is access to the resource instance through IBM IAM, which includes operations like scaling, taking backups, or IP whitelisting. The second is through the native authentication mechanism. A client should ensure understanding the distinction and assign roles appropriately to the team. Please make sure to rotate or offboard users when they no longer need access to the database.
- **Backup:** Database backups reside in IBM Cloud Object Storage and are also encrypted at rest. Backups are stored across three data centers for highest resiliency and availability.
- **Availability:** IBM Cloud Databases provides replication, fail-over, and high-availability features to protect your databases and data from infrastructure maintenance, upgrades, and failures. If a client deploys to an IBM Cloud Multi-Zone Region (MZR), the database instance is spread over the region's availability zone locations. If a client deploys to an IBM Cloud Single-Zone-Region (SZR), the database instance is spread over multiple hosts within the data center. IBM Cloud Databases automatically detects disruptions to the service and routes connections away from unavailable hosts.

Additional shared responsibility information for ICD is available at [IBM Cloud Docs > IBM Cloud Databases > Responsibilities for Cloud Databases](#).

7.3.4 IBM Event Streams for IBM Cloud Enterprise

Built on open-source Apache Kafka, [IBM Event Streams](#) is an event-streaming platform that helps clients build smart applications that can react to events as they happen. This can include real-time monitoring of patients who have heart failures or those admitted to the Intensive Care Units (ICUs) in hospitals.

The following are critical recommendations in using the service:

- **Do not place sensitive information in topic names:** Topic names and consumer groups are encrypted for transmission between Event Streams and clients as a result of TLS. However, Event Streams does not encrypt these values at rest. Therefore, names should not have any PHI sensitive data such as name, address or medical condition. Learn more in [IBM Cloud Docs for Event Streams > Data security and privacy](#).



Note, IBM Event Streams for IBM Cloud Enterprise is HIPAA-ready. The Lite and Standard plans for IBM Event Streams for IBM Cloud are not HIPAA-ready.

For more details, please refer to [IBM Cloud Docs > IBM Event Streams for IBM Cloud Enterprise > Managing encryption](#).

7.3.5 IBM Cloud Functions

Based on Apache OpenWhisk, [IBM Cloud Functions](#) is a polyglot functions-as-a-service (FaaS) programming platform for developing scalable lightweight code.

Cloud Functions leverages other IBM services such as IBM Cloud Object Storage (COS) or IBM Event Streams for IBM Cloud Enterprise. To ensure that PHI remains encrypted while using Cloud Functions, connections to these external resources need to be encrypted via protocols such as HTTPS or SSL/TLS. For example, when COS is accessed, data should be obtained by a secure HTTP service endpoint request such as: `https://s3.us-south.cloud-object-storage.appdomain.cloud`. Furthermore, the Cloud Object Store object name should not contain any PHI even though the object itself can contain PHI data.

IBM Cloud Functions may utilize at-rest PHI data. In those situations, the data should be encrypted by a service such as IBM Key Protect for IBM Cloud, which utilizes keys to ensure privacy and security in idle data.

7.3.6 IBM Cloud Internet Services

[IBM Cloud Internet Services \(CIS\)](#), leveraging Cloudflare®, includes Domain Name Service (DNS), global load balancer (GLB), distributed denial-of-service (DDoS) protection, web application firewall (WAF), Transport Layer Security (TLS), and caching to bring market-leading security and performance to client internet applications. CIS is critical, for example, in deploying externally facing web applications that utilize and display healthcare data or PHI information.

IBM Cloud Internet Services was included, along with example usage and process flow, in this guide in [Section 6.2.1, "IBM Cloud VSIs and IBM Cloud Bare Metal as HIPAA-Ready"](#) as well as [Section 6.3, "IBM Cloud PaaS Architecture Examples as HIPAA-Ready."](#) Additional information for HIPAA use cases:

- **IP Firewall:** IBM Cloud Internet Services offers several tools for controlling the client's traffic to protect domains, URLs, and directories against volumes of traffic, certain groups of requesters, and specific requesting IPs.
- **IP Rules:** The IP Rules allow clients to control access for specific IP addresses, IP ranges, specific countries, specific ASNs, and certain CIDR blocks. Available actions on incoming requests are:
 - Whitelist
 - Block
 - Challenge (Captcha)
 - JavaScript Challenge (IUAM challenge)

For example, if clients notice that a specific IP is causing malicious requests, they can block that user by IP address.

- **Domain Lockdown:** Domain Lockdown allows the client to whitelist specific IP addresses and IP ranges such that all other IPs are blacklisted. Domain Lockdown supports:
 - Specific sub-domains. For example, clients can allow IP 1.2.3.4 access to the domain bee.example.com and allow IP 5.6.7.8 access to domain kind.example.com, without necessarily allowing the reverse.
 - Specific URLs. For example, clients can allow IP 1.2.3.4 access to directory example.com/bee/* and allow IP 5.6.7.8 access to directory example.com/kind/*, but not necessarily allow the reverse. This capability is useful when the client needs more granularity in their access rules. With the IP Rules, they can either apply the block to all sub-domains of the current domain, or all domains on their account, and they cannot specify URIs.
- **User-Agent Blocking Rules:** User-Agent Blocking rules allow the client to take action on any User-Agent string they select. This capability is similar to a Domain Lockdown, except the block examines the incoming User-Agent string rather than the IP. Clients can choose how to handle a matching request with the same list of actions as they have established in the IP Rules (Block, Challenge, and JS Challenge). Note that User-Agent blocking applies to their entire zone. The client cannot specify sub-domains in the same manner they can Domain Lockdowns. This tool is useful for blocking any User-Agent strings that are deemed suspicious.
- **Challenge Passage:** Located in the Advanced security settings, Challenge Passage allows the client to control how long a visitor that passed a challenge or JavaScript challenge will gain access to their site before being challenged again. This is based on the visitor's IP and, therefore, does not apply to challenges presented by WAF rules, because they are based on an action the user performs on their site.
- **Browser Integrity Check:** Located in the Advanced security settings, the Browser Integrity Check looks for HTTP headers that are commonly abused by spammers. It denies traffic with those headers access to their page. It also blocks or challenges visitors that do not have a user agent, or who add a non-standard user agent (this tactic is commonly used by abuse bots, crawlers, or APIs).
- **Disable Content Caching:**
 - As per best practices the origin webserver should set no-cache in the Cache-Control header for the regulated content.
 - The CIS Page Rules can be used to disable caching for any content on a specified path, even if the origin does not send a no-cache Cache-Control header.

More information on IBM Cloud Internet Services is available at [IBM Cloud Docs > Getting started with IBM Cloud Internet Services](#).

7.3.7 IBM Key Protect for IBM Cloud

[IBM Key Protect for IBM Cloud](#) helps secure sensitive PHI data from unauthorized access or inadvertent employee release while meeting compliance auditing standards. Key Protect is a regional highly available service, deployed in multiple data centers located within a multi-zone geographical region.

Key Protect was included earlier in this guide in [Section 6.3.4, “Summary: Secure Flow for an IKS Application.”](#) To ensure security considerations, the following are needed as part of the design and implementation:

- Users and key access policies for creating, deleting, or rotating encryption keys is managed by the IBM Cloud IAM (identity & access) rules established when setting up an IBM Cloud account. Ensuring appropriate access rights to healthcare data is necessary to ensure compliance for covered entities.
- Clients can enable the security benefits of Bring Your Own Key (BYOK) by importing their own root of trust encryption keys, called Customer Root Keys (CRKs), into the service. With the Key Protect API, clients can use a CRK to wrap (encrypt) and unwrap (decrypt) the keys that are associated with data resources, for clients to control the security of their encrypted data in the cloud.
- If using Bring Your Own Key (BYOK), it is important to provide for secure transport of keys from on-premises systems into Key Protect. Key Protect offers two features to increase key transport security:
 - Creating a transport encryption key for the Key Protect service instance. Transport keys are used to encrypt and securely import root key material into Key Protect based on the policies that the client specifies.
 - Connect to Key Protect by using a private endpoint. This is accomplished by enabling virtual routing and forwarding (VRF) and service endpoints for the infrastructure account. When a client enables VRF for the account, please connect Key Protect by using a private IP that is accessible only through the IBM Cloud private network.
- Key Protect does not process, store, transmit, or otherwise interface with PHI data directly. However, Key Protect keys can be used to encrypt/decrypt data encryption keys used by data services to protect data which may include PHI that is stored within the cloud.
- It is recommended to not to use any personal, financial, or health related information for the encryption key names. (“Key Name”: A unique, human-readable alias for easy identification of the key.) Key Protect has no mechanism to validate that PHI is used in the key name. If it is necessary to import keys into Key Protect, verify that no part of the key contains PHI data.
- If importing keys into Key Protect be sure to not upload PHI information as part of the key (i.e. usernames, passwords, IP addresses, etc.)
- For security auditing, all Key Protect API call logs are sent to the IBM Cloud Activity Tracker with LogDNA service where users and applications interaction may be monitored for abnormal activity.

More information is available at [IBM Cloud Docs > Key Protect > Getting started tutorial](#).

7.3.8 IBM Cloud Kubernetes Service

[IBM Cloud Kubernetes Service](#) (IKS) is a managed container service for the rapid delivery of applications that can bind to advanced services like IBM Watson® and blockchain. IKS is a critical service for deploying healthcare solutions within various workflows such as critical point-of-care treatment, claims processing or imaging analysis.

As a certified K8s provider, IBM Cloud Kubernetes Service provides intelligent scheduling, self-healing, horizontal scaling, service discovery and load balancing, automated rollouts and rollbacks, and secret and configuration management. The IKS service also has advanced capabilities around simplified cluster

management, container security and isolation policies, the ability to design your own cluster, and integrated operational tools for consistency in deployment.

IBM Cloud Kubernetes Service was featured earlier in this guide in [Section 6.3, “IBM Cloud PaaS Architecture Examples as HIPAA-Ready”](#), with examples and configuration guidance throughout the subsections.

To create an IBM Cloud Kubernetes Service cluster for HIPAA use cases, follow these steps:

- Create an IKS cluster (standard subscription) on a dedicated virtual machine using the account.
- Select Classic infrastructure. Select hardware VM size based on workload. Select master service end point as Private endpoint only and check the Encrypt local disk.
- Select OS (Ubuntu 16 or 18).
- For more details of how to create IKS clusters, please visit [IBM Cloud Docs > Creating Kubernetes clusters](#).

For the worker node and the PODs within them supplied by IBM, all fix packs and patches will be published by IBM. The client must apply these patches to ensure security compliance for their worker nodes.

The client is responsible for meeting all requirements of this section on any software components / applications that they deploy into the worker nodes.

Additional general information about IKS is available at [IBM Cloud Docs > Getting Started with IBM Cloud Kubernetes Service](#). Specific additional information and links are contained in the IKS sections below in this guide.

Additional information for HIPAA use cases:

7.3.8.1 Package a single process per container

A container is not a virtual machine, and, as such, the client should not package applications in the same way. IKS Containers are designed to have the same lifecycle as the application that it hosts, so each container should only contain one application. When the application stops, so should the container. Failure to follow this guideline can result in containers ending up in an unknown state where core components have crashed or become unresponsive. This in turn means that Kubernetes cannot tell whether the container needs to be restarted without additional health checks.

7.3.8.2 Do not run container processes as Root

Namespaces isolate processes that are running in one container from processes that are running in another on the same host system. By default, the user namespace for containers is the same as that of the host. Specifically, the root user inside the container is the root user of the host system, which means that if a process breaks out of the container sandbox it has the potential to compromise the entire host.

Therefore, where possible, applications in containers should not be run as root. This can be achieved by using the USER statement inside a Dockerfile and setting it to a number that does not correspond to an existing user on their worker node. Picking a value like ‘1001’ would achieve this, with the end result being that if a process escapes the container; it would not have any permissions on anything on the host that it doesn’t already own.

7.3.8.3 Do not store secrets in Containers

Many applications might need secret information, such as private keys (including those for SSL certificates), encryption keys etc. Such secrets should never be stored inside a container image, its Dockerfile, environment variables, etc., because they are then available to anyone who pulls the image from a Registry (or can view the config of the running container). They should instead be stored in a secret management solution with the values mounted into the container at runtime, if needed.

7.3.8.4 Build the smallest image possible

To protect their applications from attackers, it is best practice to reduce the possible attack surface by including as little unnecessary software as possible. The best way to do this is to package single binaries in a scratch container, but where this is not possible the client should use a Linux distribution that is optimized to contain as little additional software as possible, such as Alpine, rather than a large distribution such as Ubuntu.

7.3.8.5 Monitor image vulnerabilities using IKS Vulnerability Advisor (VA)

IBM Cloud provides CLI and API access to VA for determining the number of known vulnerabilities in a given image stored in the IBM Cloud Container Registry service, as documented at [IBM Cloud Docs > IBM Cloud Container Registry CLI](#) and [IBM Cloud Docs > Vulnerability Advisor for IBM Cloud Container Registry](#). Application owners should periodically check the status of their images and remediate vulnerabilities as required.

7.3.8.6 Restricting access to the environment and securing the network

All IKS containers are protected by [predefined Calico network policy settings](#) that are configured on every worker node during cluster creation. By default, all outbound network traffic is allowed for all worker nodes. Inbound network traffic is blocked, except a few ports that are opened so that network traffic can be monitored by IBM and for IBM to automatically install security updates for the Kubernetes master. Access from the Kubernetes master to the worker node's kubelet is secured by an OpenVPN tunnel. For more information, see the [IBM Cloud Docs > Kubernetes service > Architecture and dependencies of the service](#).

If the client wants to allow incoming network traffic from the internet, they must expose the apps with a NodePort service, a network load balancer (NLB), or an Ingress application load balancer (ALB). Learn more at [IBM Cloud Docs > Kubernetes Service > Choosing an app exposure service](#).

Additionally, for PHI, authentication and authorization are critical in deploying client applications. As such, the following mechanisms should be enabled.

7.3.8.7 Authentication and Authorization

- Identity management for microservices is provided by Service Accounts. It is accepted practice to establish one account per application. When this is not unique to the end-user, the end-user and consumer account should be cross-referenceable in logs produced by the microservice. Please note that PHI should never be logged. Instead, logging the account and end-user should be such that logs can be cross referenced to stored PHI.
- Password policy breaches are scanned for as part of Vulnerability Advisor (i.e. password older than 90 days, password length less than 8, etc.). A client can set password readiness by adding the following to the Docker file:

```
RUN \
sed -i 's/^PASS_MAX_DAYS.*/PASS_MAX_DAYS 90/' /etc/login.defs && \
sed -i 's/^PASS_MIN_DAYS.*/PASS_MIN_DAYS 1/' /etc/login.defs && \
sed -i 's/sha512/sha512 minlen=8/' /etc/pam.d/common-password
```

- Container network accessibility should be limited by leveraging ingress/egress traffic rules in Kubernetes natively.
- Only private registries should be used to minimize PHI tampering.
- Container images stored in a private registry must be scanned periodically and frequently for known vulnerabilities. If an image in a registry or currently running container is identified as a security risk, the appropriate processes and procedures must be followed to remove the risk and correct. For IBM Cloud, the Vulnerability Advisor can be used to address this requirement for Kubernetes. It must be configured per the documentation to scan actively running containers.
- An administrator needs to provision the IKS under their IBM Cloud account, and as such can define access policies for the Kubernetes cluster to create different levels of access for different users. For example, an administrator can authorize certain users to work with cluster resources while others can deploy containers only.
- Credentials needed by the cluster and pods should always be stored in Kubernetes Secrets (Secrets).
- Access to the control plane should be minimized to only those roles responsible for managing the Kubernetes cluster. These roles would not have privileges to access the worker node containers or microservices and, therefore, would not have direct access to PHI.
- IBM Log Analysis with LogDNA service can be used to store and analyze container logs and Kubernetes cluster logs that are collected automatically by the IKS in Public and in Dedicated deployments. For more information on LogDNA, see [Section 7.3.9, “LogDNA services: IBM Log Analysis and IBM Cloud Activity Tracker”](#).
- IKS will retain logs for at least one year, and they will be protected against unauthorized access. Clients can choose what events they want to log for their cluster and if and where to forward the logs.
- SSH should be disabled into containers running healthcare microservices. Instead, use Kubectl exec to gain authorized access into a running container. This will provide the appropriate logging needed.
- Each provider tenant must be isolated from one another. As a result, it is required that each cluster is dedicated to one and only one healthcare provider. The underlying Container/Cloud Provider must ensure that adequate controls are in place to isolate tenants on all infrastructure including compute, storage and network.

7.3.9 LogDNA services: IBM Log Analysis and IBM Cloud Activity Tracker

[IBM Log Analysis](#) and [IBM Activity Tracker](#) services offer easy to use and powerful log and event collection, search, and archive abilities. Both services are built on common LogDNA services hosted on IBM Cloud.

- IBM Log Analysis is offered to aggregate and analyze application logs either originating from within IBM Cloud or elsewhere.
- IBM Cloud Activity Tracker aggregates the Cloud activity events applications made the IBM Cloud do, including actions users of the IBM Cloud UI performed.

LogDNA services do not have mechanisms to prohibit PHI data from collection, and PHI should not be included in application logs.

- For data that must remain confidential to a client company, clients must design the application to not log—or at least mask—the sensitive data.
- Implement de-identification methods on PHI and sensitive data. An example is using a record number as it appears in the database in lieu of the data itself. Note that de-identification means that the data cannot be reconstructed.
- Regularly review log and event data for anomalies and PHI leakage.

When using the offerings, it is recommended to take the appropriate configuration steps to ensure the service helps fulfill business needs.

LogDNA was included earlier in this guide throughout [Section 6.3, “IBM Cloud PaaS Architecture Examples as HIPAA-Ready”](#). Additional information for HIPAA use cases:

- **Setup Absence Alerting:** Absence alerting tests for the absence of data flowing into a client’s service instance. If data is not flowing into the system, it may indicate an issue in the application or environment. Absence duration is unique to workload and can be customized within the UI.
- **Prepare a proper Object Storage location:** LogDNA can archive to client-configured IBM Cloud Object Storage (COS). There are many COS configurations helping clients meet a variety of needs. Data may need to be replicated across Regions to meet business and regulated requirements. Alternatively, data may need to be restricted to certain locations to meet data locality requirements.
- **Setup Archival Logs:** Archive, when sent to the properly configured COS account, may provide the application or environment the necessary backup of data.

Note, LogDNA as a service does not store an independent backup copy of client data.

More information about LogDNA is available at:

- [Cloud Docs > Activity Tracker > Getting started with IBM Cloud Activity Tracker](#)
- [Cloud Docs > Log Analysis > Getting started tutorial](#)

7.3.10 IBM Cloud SQL Query

[IBM Cloud SQL Query](#) is a serverless data lake service that reads data in IBM Cloud Object Storage (COS), processes it with full standard SQL syntax, and writes the results back to Cloud Object Storage. With SQL Query, a client can analyze and process the data where it is stored - there is no ETL, no databases, and no infrastructure to manage.

Microservices-based healthcare applications should follow recommendations for persisting PHI in COS

in a secure manner. IBM Cloud SQL Query has no permanent access to data stored on COS but only persists the SQL statement and any error messages. For HIPAA use cases, clients should not use catalog management (Hive Metastore). Healthcare applications must connect IBM Cloud SQL Query to Key Protect, such that the SQL statement and error messages are encrypted with keys that are managed by clients.

For more details, please refer to [IBM Cloud Docs > SQL Query > Overview](#) and [IBM Cloud Docs > SQL Query > Security and compliance](#).

8 IBM HIPAA-Neutral Offerings



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

A list of HIPAA-ready IBM Cloud services can be found at the IBM Cloud Compliance site located at <https://www.ibm.com/cloud/compliance/industry>.

Other IBM Cloud services not listed may also be HIPAA-ready, have readiness in-progress, or been deemed HIPAA-neutral. HIPAA-neutral means a capability which operates without implicating HIPAA. For instance, IBM Cloud has several PaaS services that are HIPAA-ready or may be HIPAA-neutral based on the inherent nature of the service.

The following are services referenced in this document that are HIPAA-neutral based on service scope and functionality – no transmission, storing or processing of PHI data. However, they can be used in a HIPAA environment.

- IBM Cloud Container Registry
- IBM Cloud Identity and Access Management (IAM)
- IBM Cloud Load Balancer
- IBM Cloud Security Advisor

8.1 IBM Cloud Container Registry

[IBM Cloud Container Registry](#) allows storage and distribution of container images in a fully managed private registry. Private images can be pushed to conveniently run them in the IBM Cloud Kubernetes Service and other runtime environments. Images are checked for security issues, so informed decisions can be made about your deployments.

Container Registry was included earlier in this guide, with details and configuration information in [Section 6.3.3, “IBM Kubernetes Service \(IKS\) Configuration”](#) and [Section 6.3.4, “Summary: Secure Flow for an IKS Application”](#). Additional information for HIPAA use cases:

- **Establish a Private Docker registry for images to be stored:** This can be accomplished via the IBM Cloud Container Registry Setup. IBM Cloud enforces an account/namespace hierarchy for granting access to resources like IBM Cloud Container Registry.

More information on configuring IBM Cloud Container Registry can be found at [IBM Cloud Docs > Container Registry > Getting Started with IBM Cloud Container Registry](#).

8.2 IBM Cloud Identity and Access Management (IAM)

[IBM Cloud Identity and Access Management](#) is designed to enable clients to securely authenticate users for platform services and control access to resources consistently across IBM Cloud. A set of IBM Cloud services are enabled to use IBM Cloud IAM for access control and are organized into resource groups within an IBM Cloud user account to give users quick and easy access to more than one resource at a time. IBM Cloud IAM access policies can be used to assign users and service IDs access to the resources within an account. Users and service IDs can be added into an access group to easily give all entities within the group the same level of access.

IAM was previously referenced in this guide in [Section 6.1, “Account Setup and Managing User Identity, Access and Authentication”](#), [Section 6.3, “IBM Cloud PaaS Architecture Examples as HIPAA-Ready”](#), and [Section 6.3.1, “IBM Cloud PaaS Architecture Components”](#). Additional information for HIPAA use cases:

- To control access to their Kubernetes clusters and environments, clients can utilize the IAM tool to set authentication and authorization rules.
- IAM gives clients the ability to tie in authentication into an existing LDAP or SAML service and manage authorization rules based on pre-defined roles in IAM.

8.1.1 How IAM access policies provide access

A policy consists of a subject, target, and role. The subject in this case is the access group. The target is the subject to access, such as a set of resources, a service instance, all services in the account, or all instances of a service. The role defines the level of access that is granted to a user.

The most commonly used roles are viewer, editor, and administrator. The viewer role provides the least amount of access for viewing instances and resource groups in an account. The editor role has more access for creating, editing, deleting, and binding service instances. The administrator role includes everything for working with a service instance and can assign access to others. However, two different categories of roles are available to consider: platform and service.

For more information about the roles that can be assigned, see the [IBM Cloud Docs > Managing your account, resources, and access > IAM roles and actions](#).

8.1.2 Assigning access to access groups

Clients can organize resources in a resource group and users and service IDs into an access group to make assigning access as simple as possible. After setting up each one, create access policies for the access groups to give users in the account access to the resources created.

1. Click Manage > Access (IAM) and select Access Groups.
2. Select the name of the access group to assign access.
3. Select the **Access policies** tab, and then click **Assign access**. The following options are available for assigning access:

- **Assign access to resources within a resource group:** Use this option to give the two-part policy that is needed for users who create resources from the catalog and assign the resources to a resource group. When this option is used, access can be given to the resource group itself, all resources in a particular resource group or just one service or instance in the resource group.
- **Assign access to resources:** Use this option to assign access to all IAM-enabled services across the account or to a single service in the account, but not to an instance level.
- **Assign access to Account Management Services:** Use this option to provide a user access to account management services as a way to delegate some of the account owner capabilities. For example, delegate the ability to view billing and usage, invite and remove users, manage access groups, manage catalog services, or manage service IDs. Access can be provided to all account management services or just one.

More information about configuring IBM Cloud Identity and Access Management can be found at [IBM Cloud Docs > Managing your account, resources, and access > Setting up your IBM Cloud account](#).

8.3 IBM Cloud Load Balancer Options

[IBM Cloud Load Balancers](#) distribute traffic among client application servers residing locally within IBM data centers. Clients can explore various [load balancers](#) options to meet their specific requirements.

Load Balancers were mentioned earlier in this guide in [Section 6.2.1, "IBM Cloud VSIs and IBM Cloud Bare Metal Server as HIPAA-Ready"](#).

For more information, please see [IBM Cloud Docs > Load Balancer > Configuring load-balancing options and placing your order](#).

8.4 IBM Cloud Security Advisor

[IBM Cloud Security Advisor](#) enables centralized security management through a unified dashboard that alerts security admins to issues and guides them to understand, prioritize, manage, and resolve security issues that are related to their cloud applications and workloads.

- **Security risk and posture:** Application security remains important with constant news articles that announce a new data breach or hack. Security risks will always be a part of development and although attacks can be difficult to predict, one way to prevent them is by closely monitoring cloud deployments. For example, the risks can be related to vulnerabilities in container images that are in use, expiring certificates that can cause outage of a cloud service or application or suspicious clients or servers with a known bad reputation interacting with clusters.
- **Centralized security management:** A consolidated view of all of IBM Cloud security services and integrated partner services is available as is selection of and subscription to different services from the IBM Cloud catalog.

Security Advisor was included in [Section 6.3.3, "IBM Cloud Kubernetes Service \(IKS\) Configuration"](#).

Additional configuration information and detailed tutorials are available in [IBM Cloud Docs > Security Advisor > Getting started](#).

9 Conclusion

This guide is designed to help practitioners and clients understand IBM's public cloud capabilities, with reference to HIPAA, and includes examples of use cases and architecture designs. With a high-level overview, shared responsibility analysis, and product-level guidance, covered entities are enabled to address readiness guidelines.



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

This guide is for information purposes only and does not constitute legal or regulatory advice. IBM Cloud clients must independently analyze their environments and use cases to verify that their own control environment meets the HIPAA requirements.

As security remains top of mind for healthcare organizations moving to the cloud and security strategies evolve and mature to accommodate the transition from on site to cloud data services, IBM Cloud continues to release new capabilities to enhance how enterprises modernize and build new business applications without jeopardizing security and control.

Offering both services and enterprise-grade cloud capabilities, IBM has a unique position in the healthcare industry providing:

- Business and technology services that help healthcare clients leverage cloud to develop executable strategies and transform their businesses, operations and organizations.
- HIPAA-ready cloud solutions with unprecedented service level control.
- A plethora of evolving cloud services for computing, storage, backup, SAP, security, and unified communications.
- An infrastructure to allow healthcare covered entities to mitigate operational risk and leverage a scalable public cloud.

10 Appendix A: HIPAA Definitions

Business Associate: A person or entity that performs functions or activities on behalf of covered entity, other than in the capacity of a member of the workforce of such covered entity. Activities include creating, receiving, maintaining, or transmitting protected health information. For more information, see the U.S. Department of Health and Human Services (HHS) web page on [HSS.gov: Business associates – 45 CFR 164.502€, 164.504\(3\), 164.532\(d\) and \(e\)](#).

Business Associate Agreement/Addendum (BAA): An agreement between a covered entity and a business associate, or between business associates, providing for, among other things, safeguarding of Personal Health Information (PHI).

Covered Entity: Any health plan, healthcare clearinghouse, or healthcare provider that transmits any health information electronically in connection with a covered transaction, such as submitting healthcare claims to a health plan.

HIPAA Privacy Rule: An established set of standards for handling Personal Health Information (PHI). More details can be found on the HHS.gov site regarding the HIPAA Privacy Rule. A summary of the HIPAA Privacy Rule is available at <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html>.

HIPAA Security Rule: Standards for protecting Personal Health Information (PHI). More details can be found on the HHS.gov site regarding the HIPAA Security Rule. A summary of the HIPAA Security Rule is available at <https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html>

PHI: Personal Health Information, including demographic information collected from an individual, which:

1. is created or received by a health care provider, health plan, employer, or health care clearinghouse; and
2. relates to the past, present, or future physical or mental health or condition of an individual; the provision of health care to an individual; or the past, present, or future payment for the provision of health care to an individual; and
 - a. identifies the individual or there is a reasonable basis to believe the information can be used to identify the individual
 - b. is transmitted by an electronic media; maintained in electronic media; or transmitted or maintained in any other form or medium.

PHI data elements:

1. Names full or last name and initial
2. All geographical identifiers smaller than a state
3. Dates (other than year) directly related to an individual, including birth date, admission date, discharge date, date of death
4. Telephone numbers
5. Fax numbers
6. Electronic mail addresses
7. Social security numbers
8. Medical record numbers

9. Health plan beneficiary numbers
10. Account numbers
11. Certificate/License numbers
12. Vehicle identifiers and serial numbers, including license plate numbers
13. Device identifiers and serial numbers
14. Web Universal Resource Locators (URLs)
15. Internet Protocol (IP) address numbers
16. Biometric identifiers, including finger and voice prints
17. Full face photographic images and any comparable images; and
18. Any other unique identifying number, characteristic, code, or combination that allows identification of an individual.

De-identification: A process or mechanism for patient information to not be identified. One example is no use of any of the PHI data elements. Other methods include the Safe Harbor and Expert Determination. For further details, please reference the [HHS.gov page regarding de-identification](#).

11 Appendix B: HITECH

HITECH Summary: <http://www.hipaasurvivalguide.com/hitech-act-text.php>

HITECH Enforcement Details: <https://www.hhs.gov/hipaa/for-professionals/special-topics/hitech-act-enforcement-interim-final-rule/index.html>

12 References and Resources

HIPAA Compliance Checklist: <https://www.hipaajournal.com/hipaa-compliance-checklist/>

IBM Business Conduct Guidelines:

https://www.ibm.com/investor/att/pdf/IBM_Business_Conduct_Guidelines.pdf

IBM Cloud Catalog: Sign up for a free IBM Cloud account at <https://cloud.ibm.com>

IBM Cloud HIPAA-ready Services: The latest list of IBM Cloud HIPAA-ready services can be found at <https://www.ibm.com/cloud/compliance/industry>. IBM clients who are subject to HIPAA and who wish to use IBM Cloud products to manage or process PHI must enter into a Business Associate Agreement (BAA) with IBM.



HIPAA-ready, as used in this document, simply means the offering is ready to accept HIPAA data. HIPAA compliance, as distinguished from HIPAA-ready, involves actually meeting the HIPAA requirements on an ongoing basis. The client is responsible for its own compliance to the extent it has control over elements of compliance, and it is the client's responsibility to understand, assess and comply with its applicable requirements.

IBM Cloud Terms and Conditions: Go to <https://www.ibm.com/support/customer/csol/terms/> for documents related to data security & privacy for IBM Cloud Services offerings, including:

- Cloud Services terms, such as:
 - Service descriptions (IBM Cloud SDs, Standard SDs, Trial/beta SDs, US Federal SDs)
 - Cloud Services Agreement (CSA)
 - Business Associate Addendum (BAA)
- Data Protection terms, such as:
 - IBM Data Processing Addendum
 - IBM Data Security and Privacy (DSP)
 - DPA Exhibits

13 Disclaimers

© Copyright IBM Corporation 2021

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
December 2021

IBM, the IBM logo, ibm.com, and IBM Cloud are trademarks of International Business Machines Corp., 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 www.ibm.com/legal/copytrade.shtml.

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

Clients are responsible for ensuring their own compliance with various applicable laws and regulations. Clients are solely responsible for obtaining professional legal advice as to the identification and interpretation of any relevant laws and regulations that may affect the clients’ business and any actions the clients may need to take to comply with such laws and regulations. IBM does not provide legal, accounting or auditing advice. IBM also does not represent or warrant that its services or products will ensure that clients are compliant with any applicable laws or regulations.